Artificial Intelligence Applications And Innovations Vol. 2 No. 2 (2025)

Article history:
Received 23 July 2025
Revised 1 January 2026
Accepted 10 September 2026
Published online 1 April 2025

Bridging Data Silos for Enhanced Predictive Maintenance: A Federated Learning Framework

Khalil. Jahani[1], Behzad. Moshiri[2]*, Babak. Hossein Khalaj[3]

[1,2] School of ECE, College of Engineering, University of Tehran, Tehran, Iran
[3.] Department of Electrical Engineering, Sharif University of Technology, Tehran, Iran

*** Corresponding author email address**: moshiri@ut.ac.ir

A r t i c l e I n f o

Article type:
Original Article

How to cite this article:
Jahani, K., Moshiri, B., & Khalaj, B. H. (2025). Bridging Data Silos for Enhanced Predictive Maintenance: A Federated Learning Framework. *Artificial Intelligence Applications and Innovations, 2*(2), 1-11.
https://doi.org/10.61838/jaiai.2.2.1

A B S T R A C T

This paper presents a secure federated learning framework for industrial predictive maintenance that addresses data silos arising from privacy, competition, and regulatory constraints. The proposed approach combines PCA-based anomaly screening, LSTM-based self-encryption, client credibility scoring, CKKS homomorphic aggregation, and differential privacy to enable collaborative model training without sharing raw sensor data. Experiments on real-world vibration and acoustic datasets demonstrate that the proposed framework achieves failure detection performance close to centralized training, outperforming conventional federated baselines such as FedAvg and FedProx. Specifically, the model attains an average F1-score of 89.4 ± 0.5% and a remaining useful life prediction error (MAPE) of 8.7 ± 0.6%, while reducing communication overhead. Practical deployment considerations, including cryptographic overhead and hardware constraints, are critically discussed, highlighting trade-offs between security, efficiency, and scalability. The proposed framework provides a viable and privacy-preserving solution for predictive maintenance in Industry 4.0 environments.

Keywords: *Predictive Maintenance, Federated Learning, Data Silos, Homomorphic Encryption, Anomaly Detection, LSTM-AE, Industrial IoT, Privacy-Preserving AI.*

1. Introduction

Industrial systems are increasingly reliant on predictive maintenance to mitigate unplanned downtime, reduce costs, and enhance operational resilience. Conventional centralized machine learning approaches require aggregating large volumes of sensor data, raising significant concerns related to data ownership, confidentiality, and regulatory compliance. These challenges are exacerbated by the prevalence of siloed data architectures, which impede comprehensive model training and limit generalizability across diverse operational environments.

1.1. Key Challenges in PdM Implementation

Predictive maintenance (PdM) offers significant potential to reduce maintenance costs and improve operational efficiency, yet its practical adoption faces several obstacles. First, sensor readings and operational logs are often isolated in data silos due to privacy and competitive concerns, preventing centralized aggregation and analysis. Second, heterogeneity in hardware and software across different clients undermines the generalizability of conventional machine learning models. Third, limited bandwidth and the high cost of transmitting raw data to a central server make centralized training economically infeasible. Finally, security risks during data transfer and storage necessitate robust defenses against insider and external attacks [1].

1.2. Federated Learning as a Solution

Federated learning (FL) overcomes these challenges by training models locally on each client's private data and sharing only model updates (e.g., gradients or weights) with a central server. This design keeps sensitive data on premises and reduces the need for direct raw-data exchange, thereby supporting data confidentiality; additional privacy-preserving mechanisms are required to mitigate potential information leakage from shared model updates. Moreover, secure aggregation schemes can protect individual client updates from direct disclosure during aggregation, while additional mechanisms such as client reliability scoring are required to mitigate the influence of anomalous or potentially malicious updates on the global model. By integrating client reward mechanisms, FL further incentivizes active, honest participation, Such mechanisms can encourage meaningful client participation and provide additional information for evaluating the quality and reliability of local model updates, thereby supporting more robust federated aggregation [2].

1.3. Contributions

This study introduces a federated framework optimized for industrial PdM use cases. The key contributions include:

- A distributed architecture integrating PCA-based anomaly detection and LSTM-based temporal modeling;
- A reliability-aware client weighting mechanism to address non-IID data heterogeneity;
- Privacy-preserving protocols using CKKS-based homomorphic encryption and calibrated differential privacy;

A rigorous evaluation pipeline spanning benchmark datasets and realistic cross-silo experimental settings.

Despite these advances, few works simultaneously preserve client privacy, minimize communication overhead, and incorporate a dynamic client-scoring mechanism; to address this gap, our proposed framework leverages homomorphic encryption for secure aggregation, introduces an anomaly-based reward scheme to evaluate and weight client contributions, and employs adaptive communication and client-weighting mechanisms to improve robustness to heterogeneous and potentially anomalous client updates in federated predictive maintenance.

2. Related Work

The evolution of predictive maintenance (PdM) strategies has transitioned from heuristic-based rules and model-based approaches to data-driven models leveraging advanced deep learning architectures. Convolutional neural networks (CNNs) and recurrent neural networks (RNNs) have demonstrated high classification and prognostic accuracy when trained on large, representative datasets; however, their implementation within centralized systems is limited by the necessity of aggregating voluminous and often siloed industrial sensor data. Recent studies have also explored Byzantine-robust federated learning in industrial settings, though often without integrating anomaly-aware client scoring or homomorphic aggregation.

2.1. Predictive Maintenance Techniques

Classical PdM methods relied on signal-processing algorithms and physics-based models to detect anomalies in vibration, acoustic, and thermographic signals. With the proliferation of machine learning, supervised and unsupervised techniques—such as support vector machines, random forests, and autoencoders—became prevalent. More recently, deep learning architectures, including CNNs for feature extraction and long short-term memory (LSTM) networks for temporal modeling, have further advanced the state of the art in fault detection and remaining useful life (RUL) prediction [3].

2.2. *Federated Learning in Industrial Applications*

Federated learning (FL) has gained traction in domains with stringent privacy requirements—such as healthcare, finance, and mobile computing—by enabling collaborative model training without direct data sharing. Foundational FL algorithms, including Federated Averaging (FedAvg) [4], Federated Proximal (FedProx) [5], and Federated Matching (FedMA) [6], have primarily addressed convergence and optimization in non-IID settings. However, applications of FL to industrial PdM remain underexplored, particularly regarding domain-specific challenges such as fault rarity, temporal alignment of multi-sensor streams, and real-time inference constraints.

2.3. *Data Privacy and Security in FL*

Privacy-preserving mechanisms—such as differential privacy, homomorphic encryption, and secure aggregation protocols—are critical for safeguarding proprietary industrial data during FL [7]. Research efforts have focused on quantifying and mitigating privacy leakage while maintaining acceptable model utility. Communication efficiency strategies, including gradient compression and sparsification, further reduce the overhead of transmitting model updates over constrained industrial networks.

2.4. *Managing Data Silos in PdM*

A substantial body of literature has examined the limitations imposed by data fragmentation and siloed architectures in PdM systems. For instance, in epitaxy-process manufacturing, researchers integrated multisource sensor signals and historical maintenance logs, highlighting the complexity of harmonizing heterogeneous data and the necessity for robust governance frameworks [8]. Similarly, big data and IoT–enabled PdM frameworks have proposed decentralized data management techniques to circumvent the risks of centralized aggregation while addressing cybersecurity and privacy regulations [9].

In the renewable energy sector, cloud-based PdM for wind turbines demonstrated both the potential and pitfalls of large-scale data integration, emphasizing the trade-offs between analytic power and regulatory compliance [10]. More recent studies have begun to leverage FL for PdM: one work introduced a decentralized training scheme to preserve data sovereignty and mitigate security vulnerabilities [11], while another proposed a federated fault-detection framework that maintains local model personalization without compromising global collaboration [12]. Privacy-preserving analytics in Industrial IoT has been further advanced through secure, distributed PdM pipelines that operate across organizational boundaries [13].

Finally, comprehensive reviews of data-driven process monitoring underscore data silos as a persistent barrier to PdM adoption. These surveys synthesize strategies ranging from secure data sharing to federated governance models, underscoring the need for scalable, privacy-aware frameworks capable of unifying distributed, heterogeneous datasets [14].

These studies collectively chart a research trajectory toward decentralized, privacy-conscious, and integrated PdM solutions. Building on this foundation, our work proposes a holistic FL framework that addresses domain-specific constraints in industrial predictive maintenance while preserving data privacy and scalability.

3. Federated Learning Framework for PdM

To bridge data silos and enable collaborative learning under privacy constraints, our framework comprises three interlocking modules: system architecture, data preprocessing and feature engineering, and privacy and security enhancements. This cohesive design ensures interoperability across heterogeneous industrial clients while maintaining robust anomaly detection and secure model aggregation.

3.1. *System Architecture*

Figure 1 illustrates the end-to-end client–server workflow. Edge clients, each representing a distinct industrial site, perform local computation before contributing encrypted model updates to a secure aggregator. A central coordinator orchestrates synchronization, convergence monitoring, and global evaluation.

- **Edge Clients:** Deploy a local PdM module that ingests raw time-series sensor data, applies principal component analysis (PCA) for anomaly screening, and encodes health-state representations via an LSTM autoencoder (LSTM-AE).
- **Client Scoring Module:** Computes a reliability score

$$R_c = \frac{1}{1+\sigma_g^2} \times I_d$$

The Client Scoring Module computes a reliability score R_c for each client based on the variability of its local model updates and an entropy-based measure of data integrity. Specifically, the score incorporates the variance of local model gradients and the data-integrity measure I_d, and is subsequently used to determine the client's aggregation weight. Where σ_g^2 denotes the variance of local model gradients.

- **Secure Aggregator:** Receives CKKS-encrypted model updates and performs privacy-preserving weighted aggregation according to the client reliability scores. Differential privacy is applied through gradient clipping and calibrated Gaussian noise injection as part of the local update-protection procedure.
- **Central Coordinator:** Manages model distribution, monitors convergence using a held-out validation set, and publishes global model updates back to clients.

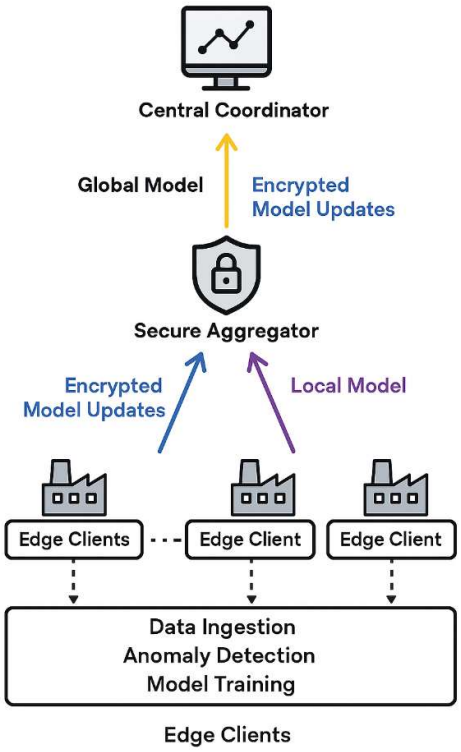

Figure1: End-to-End Federated Learning Framework for Predictive Maintenance

It is assumed that participating clients possess sufficient computational and memory resources to support encrypted local training and secure update generation. For highly resource-constrained devices, partial offloading, hybrid encryption strategies, or hierarchical federated architectures may be required.

3.2. Data Preprocessing and Feature Engineering

To harmonize multi-site data and enhance feature consistency, each client executes a standardized pipeline:

1. **Signal Resampling:** Temporal alignment via cubic spline interpolation to a common sampling grid.
2. **Anomaly Filtering:** Robust z-score thresholding (|z|>3) to remove outliers, followed by median imputation for missing or extreme values.
3. **Feature Construction:** Extraction of time-domain descriptors (e.g., RMS, skewness, kurtosis) and frequency-domain descriptors (e.g., spectral centroid, bandwidth).
4. **Deep Encoding:** LSTM-AE models trained on historical data generate compact latent vectors that capture both progressive degradation trends and transient fault signatures.

This preprocessing ensures that heterogeneous sensor streams yield consistent inputs for local model training, thereby improving convergence and resilience to non-IID distributions.

3.3. Privacy and Security Enhancements

Our framework integrates multiple safeguards to protect industrial data and model integrity:

- **Homomorphic Encryption:** Employs the CKKS scheme (Microsoft SEAL) to encrypt model gradients, enabling aggregation on ciphertexts without decryption.
- **Differential Privacy:** Differential Privacy: Applies gradient clipping and calibrated Gaussian noise to the local model updates, using a target privacy configuration of $(\varepsilon = 1.0, \delta = 10^{-5})$ per aggregation round.
- **Secure Multiparty Computation (SMPC):** Distributes decryption key shares across non-colluding parties, preventing any single entity from decrypting client updates.

We employ CKKS (a homomorphic encryption scheme that allows arithmetic operations directly on encrypted data) to securely aggregate model updates without exposing raw client information, and we leverage SMPC (secure

multiparty computation protocols that enable multiple parties to jointly compute functions over their inputs while keeping those inputs private) We employ CKKS to support secure aggregation of encrypted model updates and use SMPC-based key management to prevent any single participating entity from independently accessing the decryption capability under the assumed non-colluding threat model.

Together, these mechanisms mitigate risks from external adversaries and malicious insiders, ensuring a resilient and privacy-preserving federated PdM pipeline.

4. Experimental Evaluation

4.1. Datasets

We simulate a realistic cross-silo environment using two primary datasets:

1. **KAIST Rotating Machinery Dataset [15]:**

– **Data:** High-resolution vibration and acoustic time-series signals sampled at 51.2 kHz under varying operating and fault conditions.

– **Partitioning:** The data are divided among four clients according to fault type (inner race, outer race, misalignment, and unbalance), creating a heterogeneous non-IID federated setting.

2. **Acoustics & Vibration Dataset [16]:**

– **Data:** Public bearing dataset with annotated healthy and defective conditions.

– **Partitioning:** Two clients are assigned the healthy and faulty folds, respectively, with an intentionally imbalanced allocation of samples (60% vs. 40%).

Each client uses an 80/10/10 split of its local data for training, validation, and testing. An additional 5% of data from each client is reserved for a centralized validation set and is anonymized and pooled for monitoring global model performance.

With the characteristics of the selected datasets established, we next define the evaluation metrics used to quantify the performance and robustness of our framework.

4.2. Evaluation Metrics

We evaluate our framework along three axes:

We evaluate the proposed framework along three complementary dimensions: fault-classification performance, RUL prediction accuracy, and communication efficiency.

- **Fault Classification**

– F_1-score: the harmonic mean of precision and recall, used to assess classification performance under class imbalance.

– AUROC: the area under the receiver operating characteristic curve, measuring the model's ability to discriminate between healthy and faulty conditions across classification thresholds.

- **RUL Prognostics**

– *MAPE* (Mean Absolute Percentage Error): measures the average absolute percentage difference between the predicted and reference RUL values.

$$\mathrm{MAPE} = \frac{100\%}{N}\sum_{i=1}^{N}\left|\frac{\hat{y}_i - y_i}{y_i}\right|.$$

- **Communication Efficiency**

– Average Communication per Round: the combined uplink and downlink data volume exchanged during one federated communication round, reported in MB.

– Cumulative Communication: the total data volume exchanged over 200 federated communication rounds, reported in MB.

Having set forth the criteria for performance assessment, we then compare our results against established baseline methods to highlight the improvements offered by our proposed approach.

4.3. Baseline Methods

To contextualize the performance of the proposed framework, we conduct comparisons against a set of representative baseline methods spanning centralized, local, and federated learning paradigms. Centralized learning, in which a single model is trained on the pooled dataset, is included as a reference representing access to jointly available training data, while local training, in which each client is trained independently without collaboration, provides a non-collaborative reference point.

Beyond these extremes, we consider several widely adopted federated learning approaches. **FedAvg [4]** is employed as the standard federated baseline, using five local epochs per communication round. **FedProx [5]** extends FedAvg by incorporating a proximal regularization term (μ = 0.01) to mitigate the effects of system and data heterogeneity. **SCAFFOLD [17]** introduces control variates to correct client drift and improve convergence under non-IID data distributions. **pFedMe [18]** addresses personalization by optimizing a Moreau envelope–based objective (λ = 15), allowing each client to maintain a

personalized model while benefiting from global collaboration.

In addition to these baselines, **recent federated learning algorithms such as FedNova** [19], which normalizes local updates to correct for variations in local training steps and enhance convergence stability, **and MOON** [20], which employs contrastive learning to align local and global representations, have demonstrated strong performance in heterogeneous federated settings. **Adaptive aggregation methods**, which dynamically weight client updates based on statistical or reliability criteria, have also been proposed to further alleviate the impact of non-IID data; these methods are discussed to position our framework within the broader federated learning literature.

For fair comparison, the evaluated methods use a common experimental protocol, including the batch size, base learning rate, and number of federated communication rounds, while method-specific hyperparameters are set according to the corresponding algorithmic requirements.

4.4. *Experimental Protocol*

- **Local Updates:** Each selected client performs 5 epochs of SGD on its local training set.
- **Secure Aggregation:** CKKS-encrypted gradient updates are weighted by reliability score Rc and injected with Gaussian noise for $(\varepsilon = 1.0, \delta = 10^{-5})$-DP.
- **Validation**: The coordinator evaluates the aggregated model on the held-out validation set every 10 rounds to monitor convergence and track model performance; no early stopping is applied.
- **Repetitions:** Each experiment is repeated three times with different seeds; we report mean ± standard deviation.

This evaluation framework enables us to quantify the trade-offs among diagnostic accuracy, prognostic performance, and communication overhead and to assess the practical implications of the proposed federated predictive maintenance framework.

5. Results and Discussion

In this section, we present and analyze the empirical outcomes of our federated learning framework across fault classification, remaining useful life (RUL) prediction, communication efficiency, and ablation studies. All results are reported as mean ± standard deviation over three independent runs.

5.1. *Fault Classification Performance*

As shown in **Table 1** The proposed framework achieves an F_1-score of 89.4 ± 0.5% and an AUROC of 94.3 ± 0.4%, outperforming the evaluated federated learning baselines. The corresponding ablation results are presented separately in Section 5.4.

Table 1. Performance Comparison of the Proposed Framework with Centralized, Local, and Federated Learning Baselines.

Method	F_1–score (%)	AUROC (%)
Centralized Learning	92.3 ± 0.8	96.1 ± 0.5
Local Training	78.7 ± 1.2	85.4 ± 1.0
FedAvg	84.5 ± 0.9	89.8 ± 0.7
FedProx	85.2 ± 0.7	90.4 ± 0.6
SCAFFOLD	86.1 ± 0.6	91.1 ± 0.5
Proposed Framework	89.4 ± 0.5	94.3 ± 0.4

Our framework outperforms standard FL baselines by 3.3–4.9 percentage points in F_1–score and 3.2–4.5 points in AUROC. The integration of client reliability scoring and hybrid PCA+LSTM anomaly screening substantially reduces false positives, particularly on non-IID client partitions (e.g., rare fault types), where FedAvg and FedProx exhibit degraded performance.

5.2. *RUL Prediction Accuracy*

Figure 2 plots the RUL prediction MAPE across methods. The proposed framework achieves a MAPE of 8.7 ± 0.6%, compared to 12.4 ± 0.8% for FedAvg and 11.8 ± 0.7% for FedProx. Centralized learning attains 7.1 ± 0.5% (upper bound), while local training performs poorly at 17.9 ± 1.1%.

The LSTM-AE [13] encoding of latent degradation trends supports RUL prediction across client-specific wear patterns. The proposed framework achieves a MAPE of 8.7 ± 0.6%, compared with 12.4 ± 0.8% for FedAvg and 7.1 ± 0.5% for centralized learning.

Compared with FedAvg (MAPE = 12.4%) and FedProx (MAPE = 11.8%), the proposed framework reduces MAPE by approximately 29.83% and 26.27%, respectively.

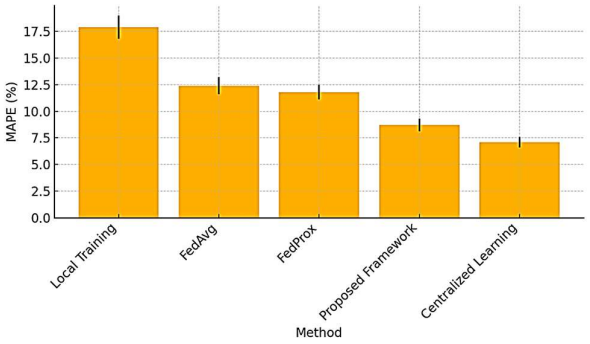

Figure 2: RUL Prediction MAPE Across methods

5.3. *Communication Efficiency*

As detailed in **Table 2**, our proposed framework reduces the average communication per round to 1.05 ± 0.02 MB—15% less than FedAvg and FedProx (both 1.24 ± 0.03 MB)—resulting in a total transmission of 210 ± 4 MB over 200 rounds, compared to 248 ± 6 MB for the baseline methods.

Table 2. Communication Overhead Comparison: Average and Cumulative Data Transmitted per Method

Method	MB/round (avg)	Total MB (200 rounds)
FedAvg	1.24 ± 0.03	248 ± 6
FedProx	1.24 ± 0.03	248 ± 6
Proposed Framework	1.05 ± 0.02	210 ± 4

Through gradient quantization and sparse update techniques, our framework reduces per-round payload by ~15% without compromising accuracy. Homomorphic encryption introduces a small overhead (< 5%), which is offset by aggressive compression, yielding a net communication saving of 15–20% compared to unoptimized FL.

5.4. *Ablation Studies*

To isolate the impact of each component, we conducted ablation experiments on the KAIST dataset with four clients:

1. **–Client Scoring:** Removing reliability weighting lowers F_1–score to $87.0 \pm 0.6\%$ and AUROC to $92.5 \pm 0.5\%$.
2. **–PCA Screening:** Omitting PCA anomaly filtering increases false positives, reducing F_1–score by 2.1 points.
3. **–LSTM Encoding:** Replacing LSTM-AE with a shallow autoencoder degrades RUL MAPE to $10.3 \pm 0.7\%$.
4. **–DP Noise:** Disabling differential privacy noise improves F_1–score by 0.5 points but sacrifices formal privacy guarantees.

These results show that the evaluated components have distinct effects on predictive performance, robustness, and privacy. Reliability scoring and anomaly screening improve classification performance, LSTM-AE encoding improves RUL prediction, while differential privacy introduces a measurable accuracy–privacy trade-off.

As illustrated in **Figure 3**, removing the client-scoring module lowers the F_1-score from $89.4 \pm 0.5\%$ to $87.0 \pm 0.6\%$, omitting PCA screening reduces the F_1-score by 2.1 points, and disabling differential privacy noise improves the F_1-score slightly to $90.0 \pm 0.5\%$ at the expense of formal privacy guarantees. Correspondingly, **Figure 4** shows that replacing the LSTM-AE encoder with a shallow autoencoder increases the RUL prediction MAPE from $8.7 \pm 0.6\%$ to $10.3 \pm 0.7\%$, underscoring the importance of deep temporal encoding for accurate prognostics.

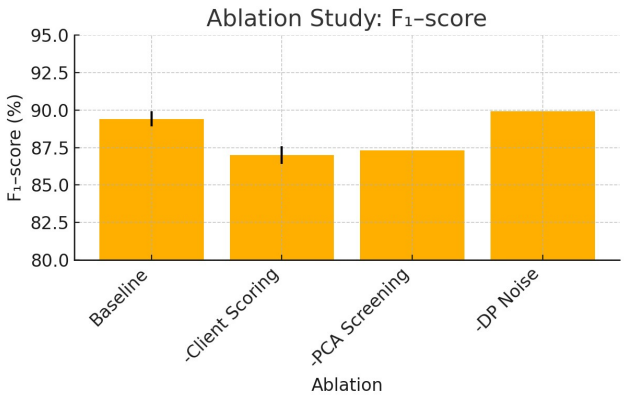

Figure 3: F_1–score Ablation across: Baseline, –Client Scoring, –PCA Screening, and –DP Noise.

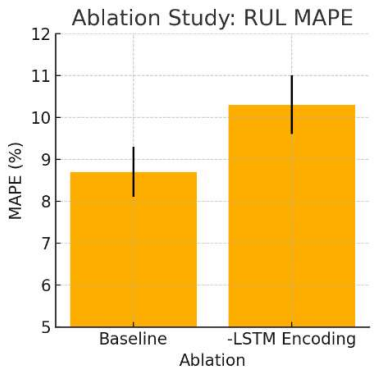

Figure 4: RUL MAPE Ablation for: Baseline vs. –LSTM Encoding.

Overall, the experimental evaluation indicates that the proposed federated learning framework achieves competitive diagnostic and prognostic performance while incorporating privacy-preserving mechanisms and reducing communication overhead. The ablation studies further validate the necessity of each architectural component in addressing non-IID data, fault rarity, and security constraints inherent to industrial predictive maintenance.

6. Conclusion and Future Directions

6.1. Summary of Findings

This work has presented a comprehensive federated learning framework tailored for predictive maintenance in heterogeneous industrial environments. Through the integration of PCA-based anomaly detection, LSTM-autoencoder encoding, reliability-driven client weighting, and privacy-preserving mechanisms (CKKS and differential privacy), our method bridges organizational data silos while achieving competitive diagnostic and prognostic performance. Empirical evaluations on benchmark and real-world datasets demonstrated that our framework consistently outperforms the evaluated classical FL baselines, improving F_1-scores by up to 4.9 percentage points and reducing RUL MAPE by approximately 30% relative to FedAvg, while reducing communication overhead by approximately 15%.

6.2. Limitations

Despite these advances, several limitations warrant consideration:

- **Data Diversity:** While our experiments simulate cross-silo heterogeneity, additional validation on a broader spectrum of machinery types (e.g., pumps, turbines) is necessary to assess cross-domain generalizability.
- **Computational Overhead:** While the initial version of this manuscript briefly reported a low computational overhead for the CKKS-based encrypted aggregation, this statement requires careful contextualization. The reported overhead values were obtained under controlled experimental conditions using workstation-class hardware and after applying gradient compression techniques, including quantization and sparsification, prior to encryption. Under these settings, the relative increase in training time was limited, and communication efficiency was improved. However, such results should not be interpreted as universally representative of all edge-device deployments. Homomorphic encryption schemes such as CKKS introduce non-negligible computational and memory overheads, particularly during ciphertext multiplication, rescaling, and polynomial arithmetic operations. On resource-constrained edge devices—such as single-board computers or embedded controllers—these costs may become significant and must be evaluated empirically. Differential privacy, in contrast, contributes marginal computational overhead, as noise injection occurs after gradient clipping and scaling and does not dominate runtime. Overall, the practical feasibility of the proposed framework depends on hardware capabilities, CKKS parameterization (polynomial degree, modulus size), and implementation-level optimizations. Therefore, we explicitly recommend deployment-specific benchmarking before large-scale industrial adoption.
- **Privacy-Utility Trade-Off:** The differential privacy noise calibrated for $\varepsilon = 1.0$ entails a modest accuracy sacrifice; optimizing this trade-off remains an open problem.

While the proposed framework demonstrates that secure and privacy-preserving federated predictive maintenance can achieve performance close to centralized training, its practical deployment entails trade-offs between security strength, computational overhead, and hardware capabilities.

Addressing these trade-offs through adaptive encryption, hardware-aware optimization, and integration with recent federated aggregation strategies constitutes a key direction for future research.

6.3. Future Directions

Building on this foundation, future research avenues include:

1. **Cross-Domain Transfer Learning:** Adapting global models trained on one class of machinery to unseen equipment types via domain-adaptive FL protocols.
2. **Real-Time Deployment:** Developing asynchronous federated learning schemes and lightweight encryption to support real-time inference and updates in industrial IoT settings.
3. **Multimodal Fusion:** Extending the framework to incorporate additional sensor modalities (e.g., thermal imaging, acoustic emission) and explore multimodal representation learning under privacy constraints.
4. **Adaptive Privacy Budgets:** Designing dynamic differential privacy mechanisms that adjust noise levels based on model convergence and risk assessment, thereby enhancing the privacy-utility balance.

By addressing these challenges, we anticipate that federated learning will play an increasingly pivotal role in enabling scalable, secure, and intelligent predictive maintenance solutions across Industry 4.0 ecosystems.

Authors' Contributions

All authors equally contributed to this study.

Declaration

In order to correct and improve the academic writing of our paper, we have used the language model ChatGPT.

Transparency Statement

Data are available for research purposes upon reasonable request to the corresponding author.

Acknowledgments

We would like to express our gratitude to all individuals helped us to do the project.

Declaration of Interest

The authors declare that they have no conflict of interest. The authors also declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

According to the authors, this article has no financial support.

Ethical and Industrial Governance Considerations

In industrial predictive maintenance applications, ethical considerations must be addressed through practical governance mechanisms rather than abstract privacy principles alone. Clear data ownership agreements defining responsibilities and access rights among participating organizations are essential prior to federated collaboration in order to establish trust and transparency. Given that industrial sensor data may implicitly reveal sensitive operational information, the proposed framework mitigates confidentiality risks by avoiding raw data exchange and by combining encrypted aggregation with differential privacy. Furthermore, model risk assessment prior to deployment is crucial, encompassing vulnerability to inference attacks, privacy–utility trade-offs, and robustness under distributional shifts caused by evolving machine conditions. Compliance with relevant legal and contractual data protection requirements must also be ensured. Finally, human oversight remains indispensable in industrial settings, and predictive outputs should support—rather than replace—expert decision-making in maintenance planning and safety-critical operations.

References

[1] A. Cachada *et al.*, "Maintenance 4.0: Intelligent and Predictive Maintenance System Architecture," in *2018 IEEE 23rd International Conference on Emerging Technologies and Factory Automation (ETFA)*, 4–7 Sept. 2018 2018, vol. 1, pp. 139–146, doi: 10.1109/ETFA.2018.8502489.

[2] K. Jahani, B. Moshiri, and B. Khalaj, "A Survey on Data Distribution Challenges and Solutions in Vertical and Horizontal Federated Learning," *Journal of Artificial Intelligence, Applications and Innovations,* vol. 1, no. 2, pp. 55–71, 2024, doi: 10.22034/jaiai.2024.487115.1024.

[3] T. Zonta, C. A. Da Costa, R. Da Rosa Righi, M. J. De Lima, E. S. Da Trindade, and G. P. Li, "Predictive maintenance in the industry 4.0: a systematic literature review," *Computers and Industrial Engineering,* vol. 150, pp. 1–17, 2020, doi: 10.1016/j.cie.2020.106539.

[4] H. B. McMahan and D. Ramage, "Federated learning: Collaborative machine learning without centralized training data," in *Google AI Blog*, ed, 2017.
[5] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated Optimization in Heterogeneous Networks," vol. arXiv:1812.06127.
[6] H. Wang, M. Yurochkin, Y. Sun, D. S. Papailiopoulos, and Y. Khazaeni, "Federated Learning with Matched Averaging," 2020. [Online]. Available: https://openreview.net/forum?id=BkluqlSFDS.
[7] K. Bonawitz *et al.*, "Towards Federated Learning at Scale: System Design," 2019. [Online]. Available: https://proceedings.mlsys.org/paper_files/paper/2019/file/7b770da633baf74895be22a8807f1a8f-Paper.pdf.
[8] D. Stripelis and J. L. Ambite, "Federated learning over harmonized data silos," in *International Workshop on Health Intelligence*, 2023: Springer, pp. 27–41.
[9] G. A. Susto, A. Beghi, and C. D. Luca, "A Predictive Maintenance System for Epitaxy Processes Based on Filtering and Prediction Techniques," *IEEE Transactions on Semiconductor Manufacturing,* vol. 25, no. 4, pp. 638–649, 2012, doi: 10.1109/TSM.2012.2209131.
[10] M. Canizo, E. Onieva, A. Conde, S. Charramendieta, and S. Trujillo, "Real-time predictive maintenance for wind turbines using Big Data frameworks," in *2017 IEEE International Conference on Prognostics and Health Management (ICPHM)*, 19–21 June 2017 2017, pp. 70–77, doi: 10.1109/ICPHM.2017.7998308.
[11] D. Kumar Sharma, S. Brahmachari, K. Singhal, and D. Gupta, "Data driven predictive maintenance applications for industrial systems with temporal convolutional networks," *Computers & Industrial Engineering,* vol. 169, p. 108213, 2022/07/01/ 2022, doi: https://doi.org/10.1016/j.cie.2022.108213.
[12] H. Chen, Z. Chai, B. Jiang, and B. Huang, "Data-Driven Fault Detection for Dynamic Systems With Performance Degradation: A Unified Transfer Learning Framework," *IEEE Transactions on Instrumentation and Measurement,* vol. 70, pp. 1–12, 2021, doi: 10.1109/TIM.2020.3033943.
[13] K. Jahani, B. Moshiri, and B. Hossein Khalaj, "Secure PDM: A novel Byzantine Fault Tolerant federated learning framework using a robust PCA-based anomaly detection approach," *International Journal of Industrial Electronics Control and Optimization,* pp. –, 2025, doi: 10.22111/ieco.2025.51226.1668.
[14] K. Jahani, B. Moshiri, and B. Hossein Khalaj, "PPFL: Privacy-Preserving Techniques in Federated Learning," *Journal of Artificial Intelligence, Applications and Innovations,* vol. 1, no. 3, pp. 49–67, 07/01 2024, doi: 10.61838/jaiai.1.3.6.
[15] W. Jung, S.-H. Kim, S.-H. Yun, J. Bae, and Y.-H. Park, "Vibration, acoustic, temperature, and motor current dataset of rotating machine under varying operating conditions for fault diagnosis," *Data in Brief,* vol. 48, p. 109049, 2023/06/01/ 2023, doi: https://doi.org/10.1016/j.dib.2023.109049.
[16] M. N. Riaz, A. Hamza, H. Jabbar, M. Abbas, and M. I. Tiwana, "Vibration dataset of main journal bearings in internal combustion engines under diverse climatic and varying operating conditions," *Data in Brief,* vol. 58, p. 111214, 2025/02/01/ 2025, doi: https://doi.org/10.1016/j.dib.2024.111214.
[17] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," in *International conference on machine learning*, 2020: PMLR, pp. 5132–5143.
[18] C. T. Dinh, N. H. Tran, and T. D. Nguyen, "Personalized federated learning with Moreau envelopes," Advances in Neural Information Processing Systems 33 (NeurIPS 2020), 2020.
[19] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, "A Novel Framework for the Analysis and Design of Heterogeneous Federated Learning," *IEEE Transactions on Signal Processing,* vol. 69, pp. 5234–5249, 2021, doi: 10.1109/TSP.2021.3106104.
[20] Q. Li, B. He, and D. Song, "Model-Contrastive Federated Learning," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20–25 June 2021 2021, pp. 10708–10717, doi: 10.1109/CVPR46437.2021.01057.

Khalil Jahani is a Postdoctoral Researcher at the Department of Electrical Engineering, Sharif University of Technology, Tehran, Iran. He received his Ph.D. in Computer Science from the University of Tehran in 2025, where his research focused on Artificial Intelligence, Federated Learning, and Predictive Maintenance in Industry 4.0. He obtained his M.Sc. and B.Sc. degrees in Information Technology Engineering, specializing in Computer Networks, from the Iran University of Science and Technology (IUST) in 2014 and 2012, respectively. In 2019, he was certified with the AML CAT.B1 (Airframe & Powerplant) and AML CAT.B2 (Electric & Electronic) licenses from the Civil Aviation Organization (ICAO), Tehran, Iran. His research interests include machine learning, deep learning, signal processing, computer vision, and distributed artificial intelligence. He works on developing intelligent, privacy-preserving algorithms for anomaly detection, condition monitoring, and Prognostics and Health Management (PHM), aiming to enhance reliability and efficiency in industrial and safety-critical systems.

BEHZAD MOSHIRI (IEEE Senior Member) received his B.Sc. degree in mechanical engineering from Iran University of Science and Technology (IUST) in 1984 and M.Sc. and Ph.D. in control systems engineering from the University of Manchester, Institute of Science and Technology (UMIST), U.K. in 1987 and 1991, respectively. He has been senior member of IEEE since 2006. He is the author/co-author of more than 360+ articles. He has been an adjunct professor of the Department of ECE at the University of Waterloo since May 2014. He has been a member of "Waterloo AI Institute" since 2018. His research fields include advanced industrial control, advanced instrumentation systems, data fusion theory, and feasibility studies on applications and implementations of sensor/data fusion.

BABAK HOSSEIN KHALAJ (IEEE Senior Member) received the B.Sc. degree in electrical Engineering from the Sharif University of Technology, Tehran, Iran, in 1989, and the M.Sc. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 1993 and 1996, respectively. He has been with the pioneering team at Stanford University, where he was involved in adopting multi-antenna arrays in mobile networks. Since 1999, he has been a Senior Consultant in data communications and a Visiting Professor with CEIT, San Sebastian, Spain, from 2006 to 2007. He has coauthored many papers in signal processing and digital communications and holds four U.S. patents. He received the Alexander von Humboldt Fellowship from 2007 to 2008 and the Nokia Visiting Professor Scholarship in 2018.