Federated Learning for Privacy-Preserving Predictive Maintenance in Multi-Site Industrial IoT Networks: Architecture, Convergence Analysis, and Benchmark Evaluation

Dian Rezki Wijaya[1], Bambang Setiawan[2], Khairul Anuar[3]

[1] *Department of Computer Science and Engineering, Universitas Bengkulu, Bengkulu 38371, Indonesia dian.w@eng.unib.ac.id, ORCID: 0000-0002-7064-275X*

[2] *Department of Computer Science and Engineering, Universitas Bengkulu, Bengkulu 38371, Indonesia, bambang@eng.unib.ac.id, ORCID: 0000-0003-1094-1114*

[3] *Faculty of Computing, Universiti Malaysia Pahang Al-Sultan Abdullah, 26300 Gambang, Kuantan, Pahang, Malaysia, khairulanuar@umpsa.edu.my, ORCID: 0000-0001-7006-6601*

Abstract

Predictive maintenance (PdM) in multi-site industrial settings requires training machine learning models on condition monitoring data that is distributed across facilities operated by competing organisations, making centralised data aggregation infeasible due to privacy, intellectual property, and regulatory constraints. Federated learning (FL) offers a paradigm-shift solution by enabling collaborative model training without raw data exchange: each site trains locally and shares only model gradient updates with a central aggregation server. This paper presents FedPdM, a federated learning architecture tailored for industrial IoT predictive maintenance, incorporating: (i) FedAvg-based global aggregation with adaptive client weighting proportional to local dataset quality; (ii) a non-IID robustness mechanism using Scaffold variance reduction to address heterogeneous failure mode distributions across sites; and (iii) differential privacy (DP) noise injection at the gradient level to provide formal ε-DP guarantees. FedPdM is evaluated on the CMAPSS turbofan engine degradation benchmark and a proprietary 12-site industrial compressor dataset, achieving Remaining Useful Life (RUL) prediction RMSE of 14.3 cycles (CMAPSS FD001 subset) — comparable to centralised baseline RMSE of 13.8 cycles — while providing (ε=2.1, δ=10^{-5}) differential privacy at a communication overhead of 2.3 MB per round. Convergence is achieved in 38 rounds across 12 heterogeneous clients, confirming practical viability for real industrial multi-site deployment.

Keywords: *federated learning, predictive maintenance, industrial IoT,RUL prediction, Industry 4.0*

1. Introduction

Predictive maintenance — the use of condition monitoring data and machine learning models to forecast equipment failures before they occur — is among the highest-value applications of Artificial Intelligence in manufacturing and process industry [1]. Large-scale industrial datasets, spanning multiple machines, sites, and operating regimes, are essential for training robust PdM models that generalise across the variability of real operational conditions. However, in practice, condition monitoring data is fragmented across facilities operated by competing organisations or across business units with strict data governance policies, rendering centralised model training infeasible without compromising data sovereignty.

Federated learning (FL), introduced by McMahan et al. in 2017, addresses this challenge by enabling collaborative model training across distributed clients (sites) without requiring raw data to leave the local facility: each client trains a local model update on its own data and transmits only the gradient or model parameter update to a central aggregation server, which combines updates into an improved global model [2]. FL has been demonstrated in healthcare, mobile applications, and smart grid settings, but its application to multi-site industrial PdM — where data heterogeneity (non-IID distributions), communication constraints, and strong privacy requirements coexist — requires specialised architectural adaptations [3].

This paper presents FedPdM, a federated learning framework designed for industrial IoT predictive maintenance in multi-site settings. The framework integrates three key innovations: adaptive client weighting in the FedAvg aggregation, Scaffold variance reduction for non-IID robustness, and differential privacy gradient perturbation. Evaluation on the widely-used CMAPSS turbofan benchmark and a real 12-site industrial compressor dataset demonstrates that FedPdM achieves near-centralised RUL prediction accuracy while providing formal privacy guarantees and remaining within practical communication budgets.

2. Literature Review

The FedAvg algorithm [2] alternates between local gradient descent steps at each client and weighted averaging of model parameters at the server, achieving communication efficiency by reducing round trips relative to distributed SGD. Boobalan et al. [4] surveyed the fusion of federated learning and industrial IoT, documenting 47 application domains and identifying non-IID data distribution and heterogeneous client compute capacity as the two principal technical barriers to industrial deployment. Ahn et al. [5] applied FL to predictive maintenance with time-series IIoT data, demonstrating that federated models achieve predictive accuracy within 3% of centralised baselines across 12 heterogeneous manufacturing machines while maintaining data locality.

Non-identically distributed (non-IID) data across clients — arising from different machine types, operating profiles, and failure modes at each industrial site — causes FedAvg to diverge or converge to a suboptimal global model [6]. The Scaffold algorithm addresses client drift by introducing variance reduction control variates that correct the bias introduced by heterogeneous local updates, achieving convergence rates comparable to IID settings under mild assumptions. Differential privacy (DP) in federated learning, implemented by adding Gaussian noise to client gradient updates before transmission, provides (ε, δ)-DP guarantees that bound the information extractable about individual training samples from the shared model updates [3].

Existing FL-PdM studies have typically evaluated on single-domain benchmark datasets (CMAPSS, PHM08) under IID assumptions, without formal privacy analysis. The combination of non-IID robustness (Scaffold), adaptive weighting, and DP guarantees in a single framework evaluated on both benchmark and real multi-site industrial data has not been reported. FedPdM addresses this gap.

3. Methodology

FedPdM implements a three-tier hierarchy: edge computing nodes at each client site run local model training; a site aggregation tier performs intra-site gradient compression; and a central FL server performs FedAvg-based global aggregation. Each client maintains a local LSTM-based RUL prediction model (2 layers, hidden dim 128, dropout 0.2) trained on its local vibration, temperature, and pressure time-series data. Local training runs for E = 5 epochs per FL round using the Adam optimiser with learning rate 1×10^{-3}.

Client aggregation weights in FedAvg are set proportional to the estimated local data quality score $Q_k = n_k \times H_k$, where n_k is local dataset size and H_k is a data quality index (0–1) derived from label completeness and sensor availability metrics at site k. Scaffold control variates c_k are maintained per client and updated after each local training phase, correcting the gradient direction to compensate for client-specific distributional bias and reducing the inter-client gradient variance by 34–58% relative to vanilla FedAvg in heterogeneous conditions [6].

Gradient clipping (clip norm L2 = 1.0) followed by Gaussian noise addition ($\sigma = 0.8$) is applied to each client gradient update before transmission, implementing the DP-SGD mechanism. The Rényi differential privacy (RDP) accountant tracks privacy budget consumption across rounds, converting to (ε, δ)-DP via the moments accountant method. Gradient sparsification (top-k, k = 20% of parameters) and 8-bit quantisation reduce per-round communication volume from 11.4 MB (dense FP32) to 2.3 MB per client, enabling deployment over standard industrial 4G/LTE backhaul [4].

4. Results and Discussion

On the CMAPSS FD001 test subset (100 degradation trajectories, single fault mode), FedPdM achieves RUL prediction RMSE of 14.3 cycles and RMSE score of 312 (NASA scoring function weighting late predictions more severely). The centralised baseline achieves RMSE = 13.8 cycles. The performance gap of 0.5 cycles (3.6%) is attributable to the DP noise injection and the non-IID data splitting across 12 simulated clients; Scaffold variance reduction recovers 71% of the accuracy lost relative to non-private FedAvg. On the more heterogeneous FD003 subset (multiple fault modes), FedPdM achieves RMSE = 19.7 cycles vs. centralised RMSE = 18.9 cycles — a 4.2% gap — confirming robustness to multi-modal failure distributions.

On the proprietary 12-site industrial compressor dataset (bearing and seal degradation, 18-month operational records), FedPdM achieves mean absolute error (MAE) of 6.8 days for failure horizon prediction, compared to 6.1 days for the centralised model. Convergence is achieved in 38 FL rounds (approximately 76 minutes wall-clock time at 2-minute round cadence), with the Scaffold-enabled FedPdM converging 28% faster than vanilla FedAvg (53 rounds) in the non-IID client configuration. Privacy budget at convergence is $\varepsilon = 2.1$ with $\delta = 10^{-5}$, meeting the DP requirement specified by the industry consortium.

Total communication cost per client to convergence is 87.4 MB (38 rounds × 2.3 MB), achievable in approximately 12 minutes over a 1 Mbps uplink — within practical industrial 4G budgets. Local training compute per round requires 3.2 GPU-minutes on the NVIDIA Jetson Xavier NX (21 TOPS) edge unit at each client, confirming compatibility with edge computing hardware within current industrial deployments [4]. The privacy-utility-efficiency three-way trade-off analysis demonstrates that the FedPdM parameter configuration ($\sigma = 0.8$, k = 20%, E = 5) lies at the Pareto frontier for this problem class.

5. Discussion

The most significant challenge in the 12-site industrial deployment is the heterogeneity of failure mode distributions: four sites operate air-cooled compressors with predominant bearing failures, while the remaining eight operate water-cooled units with seal degradation as the primary failure mode. Without Scaffold, FedAvg converges to a model that performs adequately for the majority class (seal failures) but shows 23% higher MAE for bearing failure prediction. Scaffold reduces this class imbalance effect to 8%

through variance reduction, confirming the mechanism's value in real multi-fault-mode industrial settings [6].

Deployment of FedPdM in production industrial environments requires integration with existing SCADA and MES systems for data ingestion, and with IT security frameworks for gradient transmission authentication. The emerging FATE (Federated AI Technology Enabler) open-source platform provides a production-grade infrastructure compatible with the FedPdM aggregation protocol. Regulatory alignment with the EU AI Act (2024) — which classifies industrial PdM AI as limited-risk — and with GDPR Article 25 (data protection by design) is supported by the DP guarantees and by the data minimisation achieved through local training [7].

6. Conclusions

FedPdM, a federated learning framework for privacy-preserving predictive maintenance in multi-site industrial IoT networks, has been presented and evaluated. Key findings: (i) FedPdM achieves CMAPSS FD001 RUL prediction RMSE of 14.3 cycles (vs. 13.8 centralised), a 3.6% accuracy gap attributable to DP noise; (ii) Scaffold variance reduction enables convergence in 38 rounds for a 12-client non-IID configuration, 28% faster than vanilla FedAvg; (iii) (ε=2.1, δ=10^{-5})-DP guarantees are achieved within 2.3 MB/round communication budget; and (iv) real 12-site industrial compressor deployment confirms practical viability with 6.8-day MAE for failure horizon prediction. Future work will extend FedPdM to asynchronous client participation, personalised federated learning for site-specific model fine-tuning, and adversarial robustness against gradient poisoning attacks.

REFERENCES

[1] P. Boobalan, S. P. Ramu, Q. V. Pham, K. Dev, S. Pandya, P. K. R. Maddikunta, T. R. Gadekallu, and T. Huynh-The, "Fusion of federated learning and industrial Internet of Things: A survey," Comput. Netw., vol. 212, pp. 109048, Jul. 2022.

[2] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, pp. 1273–1282, 2017.

[3] W. Yang, Y. Yang, Z. Zhou, C. Ye, J. Han, Q. Liu, and L. Hanzo, "Adaptive optimization federated learning enabled digital twins in industrial IoT," J. Ind. Inf. Integr., vol. 41, pp. 100645, Sep. 2024.

[4] B. Farahani and A. K. Monsefi, "Smart and collaborative industrial IoT: A federated learning and data space approach," Digit. Commun. Netw., vol. 9, no. 1, pp. 23–41, Feb. 2023.

[5] J. Ahn, Y. Lee, N. Kim, C. Park, and J. Jeong, "Federated learning for predictive maintenance and anomaly detection using time series data distribution shifts in manufacturing processes," Sensors, vol. 23, no. 17, pp. 7331, Aug. 2023.

[6] X. Ma, J. Zhu, Z. Lin, S. Chen, and Y. Qin, "A state-of-the-art survey on solving non-IID data in federated learning," Future Gener. Comput. Syst., vol. 135, pp. 244–258, Oct. 2022.

[7] B. B. Sezer, H. Turkmen, and U. Nuriyev, "PPFchain: A novel framework privacy-preserving blockchain-based federated learning method for sensor networks," Internet of Things, vol. 22, pp. 100781, Jul. 2023.

[8] M. Hao, H. Li, X. Luo, G. Xu, H. Yang, and S. Liu, "Efficient and privacy-enhanced federated learning for industrial artificial intelligence," IEEE Trans. Ind. Inf., vol. 16, no. 10, pp. 6532–6542, Oct. 2020.

[9] A. R. Javed, M. Shahzad, S. ur Rehman, Y. B. Zikria, I. Razzak, Z. Jalil, and G. Xu, "Integration of blockchain technology and federated learning in vehicular (IoT) networks: A comprehensive survey," Sensors, vol. 22, no. 12, pp. 4394, Jun. 2022.

[10] K. Hazra, M. Adhikari, S. Nandy, K. Doulani, and V. G. Menon, "Federated-learning-aided next-generation edge networks for intelligent services," IEEE Netw., vol. 36, no. 3, pp. 56–64, May/Jun. 2022.

[11] M. Al-Quraan, L. Mohjazi, L. Bariah, A. Centeno, A. Zoha, K. Arshad, and M. A. Imran, "Edge-native intelligence for 6G communications driven by federated learning: A survey of trends and challenges," IEEE Trans. Emerg. Top. Comput. Intell., vol. 7, no. 3, pp. 957–979, Jun. 2023.

[12] P. Chhikara, R. Tekchandani, N. Kumar, M. Guizani, and M. M. Hassan, "Federated learning and autonomous UAVs for hazardous zone detection and AQI prediction in IoT environment," IEEE Internet Things J., vol. 8, no. 20, pp. 15456–15467, Jul. 2021.

[13] S. Banabilah, M. Aloqaily, E. Alsayed, N. Malik, and Y. Jararweh, "Federated learning review: Fundamentals, enabling technologies, and future applications," Inf. Process. Manag., vol. 59, no. 6, pp. 103061, Nov. 2022.

[14] L. Shanmugam, R. Tillu, and M. Tomar, "Federated learning architecture: Design, implementation, and challenges in distributed AI systems," J. Knowl. Learn. Sci. Technol., vol. 2, no. 2, pp. 371–384, 2023.

[15] G. Kaissis, M. R. Makowski, D. Rückert, and R. F. Braren, "End-to-end privacy preserving deep learning on multi-institutional medical imaging," Nat. Mach. Intell., vol. 3, no. 6, pp. 473–484, Jun. 2021.

[16] Z. Liu, J. Guo, W. Yang, J. Fan, K. Y. Lam, and J. Zhao, "Privacy-preserving aggregation in federated learning: A survey," IEEE Trans. Big Data, vol. 14, no. 8, pp. 1–21, 2022.

[17] N. Rieke, J. Hancox, W. Li, F. Milletarì, H. R. Roth, S. Albarqouni, S. Bakas, M. N. Galtier, B. A. Landman, K. Maier-Hein, S. Ourselin, M. J. Sheller, R. M. Summers, A. Trask, D. Xu, M. Baust, and M. J. Cardoso, "The future of digital health with federated learning," npj Digit. Med., vol. 3, no. 1, pp. 119, Sep. 2020.