

Secure and Scalable Federated Learning for Predictive Maintenance in Industry 4.0 Environments

Prerna Dusi

Assistant Professor, Department of Information Technology, Kalinga University, Raipur, India
 Email: ku.PrernaDusi@kalingauniversity.ac.in

Article Info	ABSTRACT
Article history: Received : 12.10.2024 Revised : 14.11.2024 Accepted : 16.12.2024 Keywords: Federated Learning, Predictive Maintenance, Industrial Internet of Things (IIoT), Data Privacy and Security, Industry 4.0, Edge Intelligence	The developing industry 4.0 and the spread of the Industrial Internet of Things (IIoT) devices transformed the concept of predictive maintenance (PdM) into consistent tracking and decision based on data. Regardless, classical centralized machine learning-based PdM solutions have brought along essential challenges such as data privacy breach, excessive communication overhead, and insufficient heterogeneity across industrial settings. To overcome these shortcomings, in this paper a new federated learning (FL) framework is provided that allows predictive maintenance to be done in an Industry 4.0 ecosystem in a decentralized, secure, and scalable manner. Our framework enables collective training of global model in collaboration with each other, in contrast to centralized methods that require transfer of raw data, thus protecting the secrecy of operational functions and aligned with the data governance requirements. In the proposed system potentially numerous privacy-preserving mechanisms are used, such as differential privacy (DP) to obfuscate local gradients and homomorphic encryption (HE) to allow secure aggregation of encrypted model updates. Moreover, we proposed an algorithm allowing the selection of a client in an adaptive way that emphasizes high-quality clients (and stable ones) according to statistical contribution and reliability to quicken convergence and lower training variance. Though a slight trade-off to model accuracy is created by privacy enhancements, experimental evaluation on the NASA C-MAPSS turbofan engine degradation dataset shows that the proposed secure FL-based PdM framework offers near-centralized prediction accuracy with dramatic decreases of communication cost and an effective defense against data inference attacks. In this paper, the author underlines the possibility of federated learning as a premise of secure, intelligent, and scalable maintenance approaches in future industrial automations.

1. INTRODUCTION

Industry 4.0 has transformed the manufacturing industry with its models of industrial process management using the concept of cyber-physical systems (CPS), the Industrial Internet of Things (IIoT), edge computing, and artificial intelligence (AI). Predictive maintenance (PdM) can be listed among the most effective instances of such transformation as it relies on real-time data point out the interrelationships of machinery to predict equipment failures and increase asset reliability, save unplanned downtime, and optimize the maintenance schedules of different assets to streamline the overall operations. Nonetheless, the implementation of PdM solutions in IIoT areas is extremely challenges. Conventional centralized machine learning systems require massive amounts of sensor measurements to be gathered in distributed industrial assets and relayed to a

centralized server to train. Such a solution has several disadvantages such as privacy/data security, risks (raw sensor data usually includes information about the operations that can be used as a target in the cyber-attack) and expensive communication overheads and its low scalability with the exponentially increasing number of connected devices. Federated Learning (FL) has been identified to address such concerns as becoming a decentralized paradigm that lets the IIoT gadgets collaboratively teach a shared model without exchanging raw information. In its place, only the local changes to a model like gradients or weights are sent to an aggregator. Although FL presents the benefits associated with having data privacy and minimized bandwidth, the use of this technology in industrial PdM continues to pose a serious challenge. These are, the security of model updates against adversarial attacks, treatment of

heterogeneity in client devices and data distributions, and scalable learning with minimal communication overhead and latency. Although federated learning (FL) has been thoroughly discussed in the context of mobile devices and healthcare, it is not widely applied in the context of predictive maintenance in industry. The available literature regularly does not acknowledge real industrial problems of pre-existent non-independence and identically distributed (non-IID) data-sets, unreliable edge devices, latency-sensitive practices, and multi-layered security requirements. Moreover, a very small number of works have incorporated strong privacy-preserving techniques (e.g., differential privacy (DP) and homomorphic encryption (HE)) into federated learning pipelines aimed at industrial applications. The bulk of the implementations relies on the use of static client participation techniques, which may interfere with the learning performance and convergence with the use of low-quality or lagging devices. The necessity to come up with safe, scalable, and flexible architectures of machine learning using predictive maintenance in the Industry 4.0 environment is an urgent need that fuels this study. With the rise of complex and inter-connected industrial systems, there is an increased pressure on finding a learning framework that can protect critically sensitive operational data besides adapting to changes in the environment. Using federated intelligence, the solution proposed will allow distributed learning without compromising data sovereignty, operational confidentiality and compliance with GDPR, NIST regulations, etc. -- hence playing a pivotal role in the future of intelligent and secured industrial automation. Towards that direction, a new federated learning infrastructure is proposed in the study that is tailored to the predictive maintenance problem in Industrial Internet of Things (IIoT) systems and has overcome the core shortcomings of current methods. The offered framework contains the scalable architecture with the account of heterogeneous and variable IIoT nodes in terms of computational capabilities and network characteristics, which allows its wide application to solve problems in various industrial contexts. It integrates the strong security features such as differential privacy and homomorphic encryption to give end to end security to model update or parameter to adversarial attacks. Besides, an adaptive client participation algorithm is derived that allows dynamically selecting participating nodes in regard to the quality of the data used, availability of the devices, and network reliability thus enhancing the speed at which a model converges and the overall training efficiency. The framework gets empirical validation based on the

NASA C-MAPSS dataset, which proves the effectiveness of the framework to improve the accuracy of the prediction, result in lowered communication overhead, and maintain privacy. In summary, this paper provides a safe, expandable, and futuristic option of federated predictive maintenance, which is in line with the objectives of resilient, smart, and robotized manufacturing of the Industry 4.0 period.

2. RELATED WORK

Predictive maintenance (PdM) has been developed substantially through the implementation of a deep learning approach to problems, especially the Long Short-Term Memory (LSTM) networks, Convolutional Neural Networks (CNNs) and autoencoders, which have been largely used in the industry in sensor anomaly detection and Remaining Useful Life (RUL) predictions [1] [2]. Although centralized deep learning algorithms have demonstrated valuable levels of accuracy, there is an intrinsic challenge associated with this strategy regarding data privacy, communication overheads, and scalability of the system, particularly when the IIoT system is supposed to be large-scale [3].

In the effort to circumvent these constraints, Federated Learning (FL) has nowadays appeared, being a decentralized learning framework allowing to train a model on a distributed set of devices, without importing raw information [4]. In privacy-intensive applications like mobile keyboard anticipation (e.g., Google Gboard [5]) and medical diagnosis [6], where centralized information cannot be collected due to governance and confidentiality needs, FL has proved to be successful. In the medical sphere, the potential of FL was also proven when Sheller et al. [7] jointly trained deep learning models in several hospitals without violating patient confidentiality. Equally, Brisimi et al. [8] examined FL-based learning using wearable sensors and maintained the sensitivity of individuals.

Nevertheless, the FL use in industrial predictive maintenance is taking shape. Li et al. [9] introduced FL to bear fault diagnosis based on vibration signals, but their research did not involve a full set of privacy protection and could not solve network unreliability. Zhao et al. [10] proposed a federated transfer learning framework to generalize PdM on non-homogeneous machines; nevertheless, they have not implemented methods to preserve privacy, such as Differential Privacy (DP) or Homomorphic Encryption (HE). Moreover, most of the current FL schemes in IIoT utilize fixed client selection mechanisms, and these methods incorporate every device regardless of quality and availability of data as well as the state of

communication, and this causes inefficient training and poor convergence [11].

In a recent survey, such key players as non-IID data, system heterogeneity, straggler mitigation, and secure aggregation in FL were discovered to be crucial issues that require solutions so that FL can be truly adopted in large-scale dynamic industrial settings [12], [13]. However, the literature does not recommend an end-to-end FL solution, which at the same time pursues scalability, communication efficiency, adaptive device participation, and the protection of multiple layers of privacy in the context of predictive maintenance requirements in Industry 4.0.

Our work, in this sense, closes the gap because we suggest an all-encompassing FL framework that satisfies the requirements of privacy-preserving, communication-frugal, and scalability requirements in predictive maintenance in industries. We have contributed to DP and HE integration, dynamic client selection strategy that fits the IIoT networks, and a large amount of validation on the NASA C-MAPSS dataset to define and represent a real-world operation scenario.

3. System Architecture

3.1 Overview

The given system architecture is aimed at a secure, scalable, and privacy-preserving federated learning (FL) in predictive maintenance in the Industry 4.0 context. It consists of three main elements: Edge Clients, a Federated Server and a Secure Aggregator. The Edge Clients are the IIoT-based sensors or the embedded edge controllers installed on the machinery in the industry, which continuously gathers the operational data including - temperature, pressure, vibration, and RPM. Machine learning models are trained on local data and so the sensitive raw data does not need to be pushed over the network. The Federated Server is a centralized coordination entity that accepts encrypted version updates of the model by the involved edge clients and aggregates such versions of the model and finally distributes or releases a copy of the updated global version of the model to the clients. Secure Aggregator is used to reduce data confidentiality risks in data transmission and aggregation, and its privacy-preserving techniques use homomorphic encryption and differential privacy which achieve the goal of anonymity and the integrity and secretness of learning.




Figure 1a. High-Level Architecture of Secure and Scalable Federated Learning System

3.2 Workflow

The structure of the federated learning process is secured and efficient with the system. At first, all edge clients initialize an edge model based on a common baseline or a pre trained model of the global one. These local models are subsequently

trained with time-series sensor data belonging to a specific machine or subsystem. After a local training cycle is done, both clients incur privacy-preservation requirements, resulting in adding noise to the model parameters with respect to differential privacy (DP), and HE encryption of

updates. This two-fold security is to ensure that even the intercepted or enlarged updates can not present sensitive patterns. The privacy-preserving Federated Averaging (FedAvg) is executed by the Secure Aggregator in the Federated Server by training with encrypted updates without decryption of the single updates. Once the

aggregation is completed, the new global model will be broadcast back to all the edge participants and where the model will change the older version of model and iteration starts. The procedure is then repeated until some convergence criteria are satisfied after a number of communication rounds, e.g. the model accuracy or stability.




Figure 1b. Federated Learning Workflow with Differential Privacy and Secure Aggregation

3.3 Data Flow

Time-series sensor values of industrial assets present the focus of data flow in the system. Such inputs comprise the most important operation measures of temperature, levels of vibration, pressure, rotational speed, and acoustic signals, which are essential in modeling the health and the degradation patterns of equipment. Every edge client cleanse and normalizes the local data before sending it to the training pipeline. System output has 2 components: (1) Remaining Useful Life (RUL) forecasts, estimating how much time it is reasonable to expect before a component has a high probability of failure, or how much time is left

to schedule a repair; and (2) anomaly detecting scores, indicating the variations in the working parameters and that something is not acting according to the design or specification. These outputs can be calculated at a local level using the global model and its subsequent changes so that information could be obtained without affecting the data privacy. The proposed system architecture enables safe and efficient predictive maintenance of complex industries, since its implementation sustains a decentralized flow of data and does not transmit the model updates, leaving only the encrypted data accessible.




Figure 1c. End-to-End Data Flow for RUL Prediction and Anomaly Detection at the Edge

4. Security and Scalability Strategies

4.1 Differential Privacy (DP)

When applying the federated learning mechanism to a proprietary industrial data scenario, it is important to be sure that local model updates do not unintentionally leak some valuable secrets. To overcome this, Differential Privacy (DP) is on the client side of the proposed framework. Participating edge devices transmit Gaussian-noise-tinted versions of their updates (weight update or local model gradients) to the server before transmissions. The optimisation of this noise is done against privacy budget parameters (ϵ , δ) which stipulates the degree of the privacy

protection provided. The noise addition makes sure that there is minimal impact of presence or absence of any single data point in the local data of a client on the output of the model, making sure data privacy of an individual. Significantly, DP includes a mathematical guarantee which prevents attacks on reconstruction of data that happen to be of special concern in the IIoTs where sensor data can represent proprietary machine operational behavior. DP can therefore be used as a form of first defense whereby the summation world model is no longer sensitive toward the information of a particular participant.




Figure 2a. Client-Centric Differential Privacy Ring Model in Federated Learning

4.2 Homomorphic Encryption (HE)

Whereas Differential Privacy protects against data leakages due to individual update, Homomorphic Encryption (HE) protects data during data transmission and the aggregation of the model updates in the course of the training of the federated learning system. The HE enables the process of running computations on data that have not been decrypted such that, the federated server is allowed to perform aggregation on the encrypted gradients of many clients without accessing the raw or non-encrypted results of parameters. This ability means that, even in case the server or the

channel of communication is breached, the adversary will not have any access to valuable model updates. We use a lightweight additive homomorphic encryption scheme in our framework with a good balance between its computational overheads and cryptographic security guarantees and hence fits IIoT settings with limited edge devices. Combining HE and DP would provide a two-layer shield of privacy; DP encrypts the data at the client-side and HE encrypts the model update on transport and aggregation - providing confidentiality end-to-end across the federated network.




Figure 2b. Secure Aggregation of Encrypted Model Updates Using Homomorphic Encryption

4.3 Adaptive Client Participation

In federated learning the scalability also relies on effective client selection and its participation strategies besides secure communication. It is also possible that, in a non-uniform industrial setting, edge devices have different computational power, varying connectivity stability, and the quality of data. To take care of such variability, our framework develops an adaptive client participation mechanism relying upon a reputation scoring system. The clients are scored dynamically based on the historic performance metrics of the update quality, the training completion rate, the communication latency and the model divergence. The higher scorer clients are the top priority in each training session, and only reliable and good quality clients participate and generate updates on the model. The selective participation has the effect of lessening the burden of stragglers, providing little unused communication overhead, and hastening convergence of the model. Also, the strategy will be fair in the long-term as those clients who failed to perform well in the past will have a chance to re-enter the training process when they will be more reliable. Incorporating this adaptive mechanism with privacy-preserving techniques, our framework is thereby able to achieve better security and scalability demanded of industrial deployments into Industry 4.0 systems.

5. Experimental Setup

5.1 Dataset

To understand the effectiveness and feasibility of the proposed federated learning framework towards predictive maintenance, we considered the well known NASA C-MAPSS (Commercial Modular Aero-Propulsion System Simulation) dataset. This baseline data set represents a simulation of degradation of turbofan engines functioning during different operating and faulty conditions. It has several subsets (FD001-FD004) corresponding to various operating conditions and faults modes. Every example in the dataset relates time-series sensor readings that comprise temperature, pressure fan speed, and engine performance metrics, over several working cycles until death. Such features contribute to the realisation of the C-MAPSS as a good collection of data to study the Remaining Useful Life (RUL) prediction and the fault diagnosis in practical industrial conditions. The FD001 set in particular, the part of the data that has consistent operating conditions and only one fault mode, seems to be the most suitable subset in our cases (controlled assessment without losing industrial applicability).

5.2 Models

In order to determine the effectiveness of our suggested system, three different models were created and compared to one another. The first model is a baseline one and involves a centralized Long Short-Term Memory (LSTM) network, which is trained using all the data in the conventional way. The model provides the best learning capacity, but should have privileged access to all the data of the clients, which contravenes the principles of privacy. The second model is a standard Federated LSTM (FL-LSTM) in which each client trains one local LSTM model but has to send periodic encrypted updates to a central server that aggregates updates. This setup involves homomorphic encryption (HE) to arrange secure communication, but applies fixed client choice. Model three and the proposed model is on the basis of growing the FL-LSTM to an adaptive client selection and differential privacy (DP) to contribute to the robustness and privacy. We simulated 20 edge clients in our experiment, emulating a virtual instance of an IIoT node and a different number of computational resources, local dataset volumes, and network bandwidth limitations to account in a realistic way for heterogeneity across industrial settings. To test the fault tolerance of the system, a non-IID client data partitioning was used to emulate a device specific failure mode and introduced dropout behavior randomly. This model dynamically selects most reliable clients in each iteration of training and adds DP noise to protect local update. All the models have equal LSTM architectures to compare on a fair basis and trained through numerous rounds of communication until convergence.

5.3 Metrics

To compare and analyze the effectiveness of the suggested models, we applied a large list of quantitative evaluation measures covering the accuracy of predictions, system efficiency, safeguarding privacy, and realistic deployment settings. To ensure accuracy, the Mean Absolute Error (MAE) and the Root Mean Square Error (RMSE) were used to compute the predicted values on the RUL values to arrive at information on the precision and reliability of the models. Communication cost was measured by the number of bytes communicated each round of training, and the bandwidth used in the entire cycle of training. The convergence rate was determined by the amount of iterations in the communication pattern it took the global model to converge at a definite level of accuracy. To assess privacy leakage resistance, we employed both the model inversion and membership inference attacks on the aggregate global models and determined the success rate of the attack.

We have added two more convenient edge-oriented metrics to these conventional measures: energy efficiency and latency. Energy was measured in terms of floating-point operations (FLOPs), based on the estimate required to train on local models and communicate to the rest of the compute nodes per round which factored in device-specific power curves. End-to-end latency was counted as the start of a round of training on edge clients up to successful training without considering the calculation time on the client side and the delay on the network. These added-on measures will give a more down-to-earth estimation as to the applicability of the framework to the frames that apply in the industrial setting of energy-limited and latency-sensitive scenes of industrial frames.

6. RESULTS AND DISCUSSION

A set of comparative experiments was carried out to assess the effectiveness of the proposed secure and scalable federated learning (FL) framework of predictive maintenance on Industry 4.0 settings over four different configurations: (1) Centralized LSTM, (2) FL-LSTM (vanilla), (3) FL with Differential Privacy (DP) and Homomorphic Encryption (HE) and (4) the Proposed FL framework, which combines the effectiveness of DP, HE and adaptive client participation. The tested models were constructed on the basis of the NASA C-MAPSS dataset and evaluated on four dimensions that include prediction accuracy, communication cost and the risk of privacy breach. This is because the Centralized LSTM model had

the best prediction accuracy with a Mean Absolute Error (MAE) of 15.2 as it was not restricted on the amount of data to use given the fact that it could access the whole dataset. Nevertheless, the costs of such approach allow us to speak about considerable communication expenses and severe privacy risks because it operates with raw sensor data extracted on all edge clients. This type of model is not applicable in industrial contexts in which data confidentiality and adherence to data protection laws (e.g. GDPR) are of utmost concern. On the other hand, vanilla FL-LSTM, without privacy additions and adaptive solutions, realized the greatest MAE of 17.1, which demonstrated a slight drop in accuracy, caused by data locality and reduced global awareness. However, it has medium rates of communication efficiency as it did not transfer raw data and there is minimal privacy risk as well. With this setting, it is evident that FL has the fundamental benefit in privacy-preserving training with lack of protection over information leaking through model updates and the inefficient participation by low-quality clients.

The model that uses FL + DP + HE had an accuracy of (MAE = 17.5), but it reduced the communication cost and managed privacy protection significantly. Such a configuration proves the utility of implementing privacy-preserving mechanisms to the FL pipeline. It explains the relative small loss of accuracy due to the noise imposed by differential privacy and the computational cost of encrypted computing, but the model has potential in situations when security of data is of primary importance.




Figure 3. Comparative Radar Plot of Model Configurations Based on Performance Metrics

Table 1. Comparative Evaluation of Federated Learning Configurations on Accuracy, Communication Overhead, and Privacy Risk

Method	MAE ↓	Comm. Cost ↓	Privacy Risk ↓
Centralized LSTM	15.2	High	High
FL-LSTM (vanilla)	17.1	Medium	Medium
FL + DP + HE	17.5	Low	Low
Proposed FL (Full)	15.9	Low	Very Low

The improvement of all metrics was provided by the offered FL framework that incorporates DP, HE, and a new adaptive client selection mechanism. It produced a MAE of 15.9 which was close to the performance of the centralized model whilst guaranteeing low communication cost and minimal risk of privacy breach. Adaptive client selection algorithm was especially helpful to reduce model drift generated by unreliable/noisy client updates. The scheme ensures that solar PV installers who will continue with the training during the next round will be high quality and stable contributors. model consistency and fast model convergence that were maintained by the system. This finding proves that federated learning can achieve a comparable performance to a centralized solution with the help of an intelligent client orchestration and robust privacy protection measures.

Overall, the proposed system proves that it can meet high predictive accuracy, data protection and efficiency factors in communication, and hence falls in the range of deployment in large scale industrial communications where data security and varied nature of the application set-ups are highly sensitive. The findings confirm the architectural decisions and support the feasibility of federated learning as one of the fundamental premises of secure and intelligent predictive maintenance within the Industry 4.0.

7. CONCLUSION

This paper introduces a feasible end-to-end federated learning (FL) framework in securing, scalable, and protecting privacy in predictive maintenance in the Industry 4.0 setting. The architecture also poses little to no risks of privacy breaches, communication overhead and regulatory non-compliance as raw data is not transmitted using the framework and, instead, the training of the models occurs on edge devices within an Industrial Internet of Things (IIoT) system. Utilizing Differential Privacy (DP) and Homomorphic Encryption (HE) offer end-to-end model security in model update, and the adaptive mechanism of client participation generation dynamically picks reliable and high-quality edge clients to improve model robustness and increase the model convergence speed.

A NASA C-MAPSS turbofan engine dataset demonstrates using experimental validation that the proposed framework is able to strike a good

balance between the predictive accuracy and the system-level efficiency. The outcome reveals that the suggested FL system is quite comparable in terms of performance with centralized counterparts but cuts down communication expenses significantly and displays high tolerance to privacy inference assaults. The results validate the feasibility of federated intelligence as an underlying technology to support real-time, data-sensitive predictive maintenance of the industrial assets in the heterogeneous environments.

In a bigger picture, the paper establishes the framework of using smart, robust, and safe maintenance systems in smart manufacturing infrastructures. It facilitates the principal industry issues of the heterogeneity of devices, poor stability of communications and cybersecurity threats with a flexible, futureproof architecture that can handle a wide variety of real-life applications.

Nevertheless, one should also consider the computational cost that is incurred by privacy-preserving methods, specifically the differential privacy and homomorphic encryption. On the one hand, they massively enhance the safety of information; on the other, they also impose an extra burden on resources, which can be a burden on low-power IIoT devices. As future work, we will consider light weight cryptographic graduate, optimized encryption schemes and adaptive DP tuning as a method of reducing computational drag and preserving strong privacy guarantees.

Possible next steps in this work involve applying the system to real-time embedded and microcontroller-based IIoT systems, combining the framework with multi-modal sensor fusion (e.g. acoustic, thermal, and visual data) and testing integration with blockchain-based trust management and federated transfer learning to accomplish cross-domain generalisation at scale. Moreover, compressing the models with energy efficiency and the edge-aware optimization of training will become important in scaling the framework to resource-constrained settings in order to enable wider industrial usage of the framework.

REFERENCES

- [1] Y. Zhang, L. Peng, Y. Wang, and B. Wang, "A deep learning-based model for fault diagnosis of rotating machinery using time-

- frequency image," *Mechanical Systems and Signal Processing*, vol. 108, pp. 1–15, Aug. 2018.
- [2] R. Zhao, R. Yan, Z. Chen, K. Mao, P. Wang, and R. X. Gao, "Deep learning and its applications to machine health monitoring," *Mechanical Systems and Signal Processing*, vol. 115, pp. 213–237, Jan. 2019.
- [3] A. S. Razavi, M. M. Dehkordi, and B. Li, "A survey of machine learning techniques for condition monitoring and predictive maintenance of bearings in rotating machinery," *Journal of Manufacturing Systems*, vol. 56, pp. 659–680, Jan. 2020.
- [4] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated Learning: Strategies for Improving Communication Efficiency," in *Proc. NIPS Workshop on Private Multi-Party Machine Learning*, Barcelona, Spain, Dec. 2016.
- [5] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS)*, vol. 54, pp. 1273–1282, Apr. 2017.
- [6] J. Xu, B. Glicksberg, Y. Su, P. Walker, and J. T. Dudley, "Federated learning for healthcare informatics," *Journal of Healthcare Informatics Research*, vol. 5, no. 1, pp. 1–19, Jan. 2021.
- [7] M. J. Sheller, G. A. Reina, B. Edwards, J. Martin, and S. Bakas, "Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation," in *Proc. Int. Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2018, pp. 92–104.
- [8] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi, "Federated learning of predictive models from federated Electronic Health Records," *International Journal of Medical Informatics*, vol. 112, pp. 59–67, Sept. 2018.
- [9] X. Li, Q. Wang, and J. Zhang, "A federated learning-based intelligent fault diagnosis method for industrial rotating machinery," *IEEE Access*, vol. 11, pp. 6723–6734, Jan. 2023.
- [10] L. Zhao, M. Zhang, and J. Chen, "Federated transfer learning for predictive maintenance in cross-device industrial settings," *Journal of Intelligent Manufacturing*, vol. 35, pp. 941–956, Feb. 2024.
- [11] H. Haddadpour, M. Kamani, and M. Mahdavi, "Local SGD with periodic averaging: Tighter analysis and adaptive synchronization," in *Proc. 33rd Conference on Neural Information Processing Systems (NeurIPS)*, 2019, pp. 1–12.
- [12] P. Kairouz et al., "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [13] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 10, no. 2, pp. 1–19, Jan. 2019.