Journal of Innovations in Business and Industry

Vol. 04, No. 01 (2026) 69-78, DOI: 10.61552/JIBI.2026.01.007 - http://jibi.aspur.rs

SECURE PREDICTIVE MAINTENANCE FOR INDUSTRIAL SYSTEMS USING FEDERATED LEARNING

Houssem Hosni [1]

Received 27.02.2025.
Revised 14.04.2025.
Accepted 25.05.2025.

Keywords:

Federated Learning, Industry 4.0, Predictive Maintenance, Secure Aggregation, Data Privacy, Industrial Internet of Things.

Original research

A B S T R A C T

In this paper, a secure and scalable predictive maintenance approach for Industry 4.0 using Federated Learning (FL) and Artifficial Intelligence (AI) is addressed. Unlike other approaches, FL maintains data on the premises, which guarantees privacy and regulatory conformance. The system design relies on edge devices to train protected local models and share secure updates. It consists of data pre-processing, model training, and secure aggregation. Experimental results demonstrate FL can obtain high efficiency, communication reduction and improve security against cyber threats. System challenges and future directions are also described in the paper. This paper provides a privacy-aware, up-to-date analysis of predictive maintenance in smart industrial environments.

1. INTRODUCTION

Driven by Industry 4.0, there has been an exponential development of industrial systems that has revolutionized the management of the Internet of Things for manufacturing, energy, and critical infrastructure (Adimulam et al., 2019; Hector & Panjanathan, 2024). With the explosion of Industrial Internet of Things (IIoT) devices, automation and connectivity throughout industry, organizations are able to access a huge amount of operational data, in real-time, and that has facilitated new means of using data to optimize equipment health and maintenance strategies (Alhuqayl et al., 2024). In this scenario, predictive maintenance (PdM) has been introduced as a fundamental enabler to enhance system reliability, prevent unforeseen operational outages, and decrease maintenance costs by anticipating failures in real time. Predictive maintenance (PdM) uses time-series data collected from a variety of sensors that look for things such as vibration, temperature, pressure, and other important factors, and utilizes a combination of machine learning (ML) and deep learning (DL) models to detect precursor indicators of possible degradation and determine the remaining useful life (RUL) of the asset (Hosni, 2022; Yakhni et al., 2022; Adimulam et al., 2019; Anandan et al., 2022).

Although predictive maintenance holds great hope, real-world implementation of such maintenance in an industrial environment is restrained by overwhelming challenges. Conventional PdM architectures are mainly centralized, where raw sensor and log data are collected from physical sites located distantly to the central server or cloud platform for model training and inference. Although such a combination allows powerful analytics, several fundamental problems have to be addressed. The centralization of sensitive operational data exposes significant privacy and security issues, first. In the field of industry, datasets often carry protected knowledge, special operational features and types of failure which when identified could bear loss of competitive advantage if revealed or in some extreme cases even challenge national security such as in the case of critical

[1] Corresponding author: Houssem Hosni
Email: hosny.houssem@gmail.com

infrastructure. Second, regulations (e.g., European General Data Protection Regulation (GDPR) and sector-specific standards) put strict requirements on the localization, access control and auditability of data, which have made it increasingly difficult to gather and transfer industrial data outside their enterprise boundaries (Achouch et al., 2022).

In addition, the large scale and heterogeneity of the current industrial environment (that crosses multiple factories, remote sites, and different equipment generations) makes the collection, harmonization and annotation of data very challenging. The data collected from the IIoT devices are non-IID, since reflecting different operational conditions, maintenance histories, and usages of particular assets (Bhatti et al., 2025). Such diversity could pose challenges for centralized ML models that do not generalize across sites or rare or site-specific failure modes. Network limitations such as low bandwidth, and periodical network availability in remote or resource poor areas make centralized data transfer and real time analytics not viable (Kozma et al., 2019).

To address such challenges federated learning (FL) has recently been introduced as a game-changing paradigm for secure and scalable predictive maintenance in the industrial domain. FL decentralizes model training so that every client, including factory, production line, and individual machine, could train a local model based on its own data. Only model updates (e.g., gradients or weight parameters) are communicated to a central aggregator that aggregates them to learn a global model. A critical feature of the system is that raw data never travels across the network, thus ensuring data privacy/security and data sovereignty but also minimising network traffic. This architecture is well suited for industrial settings, in which data security, regulatory compliance, and operational diversity are critical (Ahn et al., 2023).

The literature has shown that FL-based predictive maintenance models are able to reach similar and sometimes even higher predictive accuracy of the centralized approach, especially when non-IID and imbalanced data is considered (Ahn et al., 2023). For example, state-of-the-art FL networks with the combination of convolutional and recurrent neural networks (e.g., 1DCNN-BiLSTM) have demonstrated strong performance for time-series anomaly detection and RUL prediction in distributed IIoT scenarios (Pruckovskaja et al., 2023). These results are attained with strong privacy guarantees in place; model´ updates can be further secured with encryption, differential privacy and secure multi-party computation.

At the macro level, the attention for FL in predictive maintenance fits within the current development toward industrial digitization including edge computing and the merging of operational technology (OT) and information technology (IT). Utilizing edge computation resources mitigate the reliance on centralized cloud infrastructure and improves system robustness, while enabling real-time analytics even in bandwidth-limited environments (Ahn et al., 2023). Additionally, FL enables intelligence sharing among cross-organizations by allowing multiple entities (e.g., hardware suppliers, service providers, and plant operators) to collectively enhance the prediction models without sharing their private data (Pruckovskaja et al., 2023).

Nevertheless, the application of FL on industrial predictive maintenance also presents some challenges. System heterogeneity, client asynchronization and communication overhead could lead the model convergence as well as operational efficiency down. Adaptative aggregation policies, asynchronous training schemes, and individualized strategies constitute an open area of research that advocates to alleviate these limitations. Security concerns are especially significant, since adversaries might want to learn private information in the model updates and want to insert their own updates to damage the global model. FL will be adopted by blockchain and explainable AI to increase trust, traceability, and transparency in academic-industrial analytics (Pruckovskaja et al., 2023).

Resilient predictive maintenance with federated learning marks a pivot for industrial analytics toward achieving the necessary level of advanced, collaborative intelligence in a way that kept it data privacy-respecting, secure, and regulatory compliant. With the development of industrial systems, the application of FL-based predictive maintenance frameworks will become one of the pillars of smart, sustainable, and ethically controlled industrial operation.

2. INDUSTRIAL PREDICTIVE MAINTENANCE IN THE AGE OF DATA PRIVACY

The transition to Industry 4.0 has transformed the nature of industrial system operations by integrating digital technologies (such as the Industrial Internet of Things (IIoT), artificial intelligence (AI) and edge computing) into manufacturing settings (Hosni, 2025). This trend has enabled decision support systems, adaptive control, and real time monitoring. Among the most promising uses of these advances is predictive maintenance (PdM), which leverages sensor data and machine learning algorithms to predict equipment failures before they happen. According to (Magena et al., 2024) who mentioned that PdM strategy can save maintenance cost by 30% and increase the availability of the asset by more than 20%. This will lead to improve production line efficiency.

Such optimization is contingent upon having enough operational data collected from various assets and systems.

Despite its strengths, PdM has structural pain points, stemming from old-fashioned, centralized data architectures typically associated with machine learning. In a typical architecture all raw sensor data collected from a number of factory locations is sent to a central cloud server, for storage and analysis. And this concentration is a series of serious weak points. First, it enlarges the security threat level, because centralized repositories attract hackers. Second, it concerns

confidentiality issues: it may be the case that industrial data may include sensitive operational metrics, proprietary configurations, or system-specific anomalies that manufacturers are reluctant to share, even within the company. Third, centralized data collection is also at odds with the likes of the GDPR, in which a number of regulations restrict the movement, storage, and processing of personal data. As emphasized by Fraga-Lamas et al. (2017), the extent of the success of Industry 4.0 endeavors relies on the implementation of such privacy-preserving limitations.

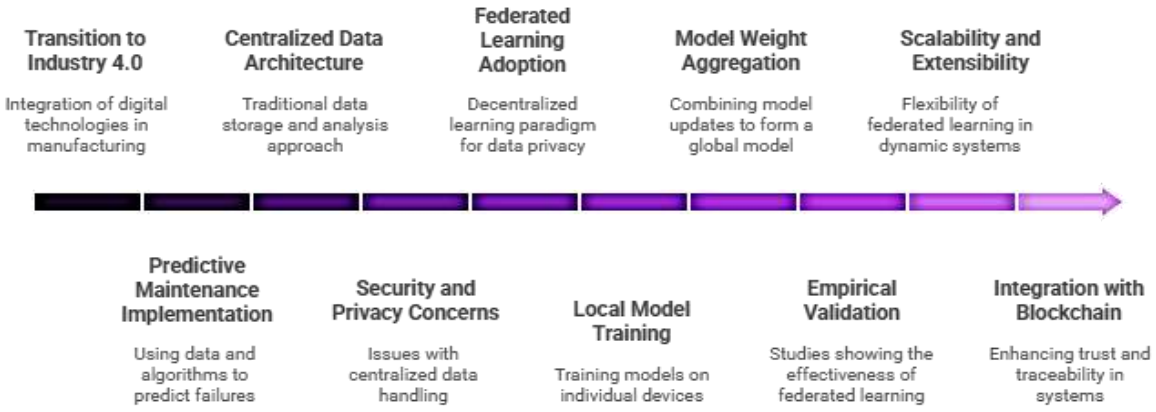

Figure1. Federated Learning for Predictive Maintenance

In order to overcome these problems, federated learning (FL) has gained popularity in both academia and industry, as an emerging decentralized learning paradigm which is in contrast with the centralized setting. In FL, every device or node trains a local model on its dataset and transmits only the updated model weights to a central server. The server "stitches" these updates together to form a global model, without ever seeing any raw data. It meets the requirements for data sovereignty, reduces message overhead, and remarkably improves the compliance level to data protection laws. Near real confirmation FL allows manufacturers to jointly identify trends or predictive maintenance (Figure 1) actions without revealing business sensitive operational data to third parties or common storage (Islam et al., 2023; Yurdem et al., 2024).

The practicality of this approach is shown by empirical studies. Ahn et al. (2023) used FL for predictive maintenance with a time-series IIoT data. Their findings provided evidence that federated models can be easily achieved with high predictive power (up to 97.2%) without any need of centralizing data, and thus maintaining data local-private (confidentiality) and systems being secure. In addition, the system was able to deal with non-IID (non-independent and identically distributed) data phenomena that are typical in industrial settings, where every machine may work in different conditions and operation regimes.

This strategy has been validated by Pruckovskaja et al. (2023), who contrasted FL with centralized and standalone local learning methodologies for predictive maintenance and quality inspection. Their results showed that FL can not only preserve privacy but has competitive accuracy as well, particularly in multi-factory environments where data sharing is limited or even impossible. These findings further support FL as a practical approach to distributed industrial intelligence.

Also, the scalability and extensibility of FL would be available for dynamic industrial system. As noted by Kairouz et al. (2021), FL is resistant to heterogeneity in connectivity of devices, computational resources, and data. Such flexibility is critical in large-scale industrial networks, in which devices are deployed at locations with different network infrastructure, computing capabilities and maintenance schedule.

Recent research also studies the combination of FL-BCT in order to improve trust and traceability within the collaborative industrial IoT environments. For instance, Pham et al. (2025) introduced federated learning with blockchain for fairness, transparency, and data integrity in decentralized predictive maintenance systems. Such hybrid architecture would enable companies to retain data privacy, while ensuring the integrity of all models updates from distributed entities.

Although centralized methods for predictive maintenance can provide the analytical strength that is necessary, they become ever more incompatible with the privacy, security and legal requirements of today's industrial systems. Federated learning is important for offering a scientific and practical solution that fits with the vision of Industry 4.0. It allows for secure, collaborative intelligence among distributed systems while protecting sensitive data. And with the growing digitization of industry, the adoption of decentralized analytics is core to achieving efficiency and responsible data stewardship in future.

3. FEDERATED LEARNING: A DECENTRALIZED PATH TO SECURE ANALYTICS

Federated Learning (FL) has become a paradigm-shifting machine learning method, which tackles the core concerns of privacy, security and regulation-including in distributed and data-sensitive industrial setups. In contrast to centralized learning, where raw data from all

participants is transmitted to a central server to be trained, federation learning allows each participant, for example, industrial sensor, smart machines, and production site, to train their data themselves. Rather than exchanging data, only model parameters (i.e., gradients or weights) are sent to a central server which performs an aggregation over them to update a global model. This decentralized model provides a strong protection against data exposure, breaches and non-compliance risks under data regulations like the General Data Protection Regulation (GDPR).

Methodologically, FL works in iterations. The server initializes each round by providing a base model to all local clients. These clients will subsequently train their model using their data and will send-back only the updated weights. The updates are then aggregated by the server with algorithms like Federated Averaging (FedAvg) proposed by McMahan et al. (2017) – in order to obtain a better global model. This follows until convergence occurs. Most importantly, because raw data never leaves the respective source, FL significantly shrinks the attack surface for cyberattacks and preserves the organization's data sovereignty—a point of interest in industrial environments since operational data commonly includes proprietary or sensitive data Kairouz et al. (2021).

Table1. Predictive Maintenance Approaches Comparison

Feature	Federated Learning (FL)	Centralized Data Architectures
Data Privacy	Enhanced data sovereignty	High confidentiality concerns
Security	Reduced attack surface	Increased cyber threat attack surface
Bandwidth Usage	Low (model updates only)	High (raw data transfer)
Model Generalization	Improved (diverse sources)	May be biased (central data only)
Regulatory Compliance	Improved compliance with data laws	Potential conflicts with GDPR
Communication Overhead	Reduced	High
Scalability & Adaptability	High (distributed, edge-friendly)	Limited by central server capacity
Data Handling	Supports non-IID data	Requires data centralization
Trust & Traceability	Enhanced with Blockchain	Lower
Collaboration	Secure and collaborative	Limited due to privacy concerns

From a scientific perspective, FL holds many comparative benefits when compared with centralized learning, especially in the industrial setting (Table 1). Firstly, data is kept confidential – a crucial requirement in fields such as aerospace, energy and manufacturing. Second, it reduces communication cost because users send the updates in the form of the updated model, so raw data never leaves the device, making FL applicable to edge and communication limited systems. Third, FL can adapt to non-IID (non-identically distributed) data, which is presented commonly in the IIoT, due to the fact that sensors and machineries operate under different circumstances. For this reason, FL is suitable for extremely heterogeneous industrial ecosystems.

An interesting application is the predictive maintenance in distributed manufacturing networks. With FL, every sensor-enabled machine can train a model locally to identify predicators early indicating machine failure. These models are consolidated to construct a complete and valid global prediction model, always without centralizing sensitive operational data. In a study by Ahn et al. (2023), FL on the problem of maintenance predictions and anomaly detection in time-based sequence data was able to attain an uplift of 97.2% in accuracy even while effectively preserving data locality and accommodating temporal distribution shifts between clients.

The technical soundness of FL also means it can be easily adapted to heterogeneous computing environments. In practical industrial systems, equipment can vary greatly in terms of processing capacity, data availability, network reliability and availability. Asynchronous training paradigms enable FL to operate even when some clients are offline or are not accessible for short periods of time. Kairouz et al. (2021) present a comprehensive scientific introduction to these abilities, focusing on the robustness of FL in large-scale real-world scenarios.

FL is performing competitively to centralized approaches in scientific comparison studies, and sometimes showing even better performance. Li et al. (2021) show that federated learning provides similar accuracy with significantly better data privacy. This makes FL a realistic solution for coordinating between organizations, where industrial partners can co-train models without leaking the proprietary datasets.

Centralized learning can be more straightforward to implement and administer in homogeneous, well-regulated environments, but becomes challenging when data sovereignty, cybersecurity, and system scalability issues are of some important. While federated learning is more technically challenging (as it involves synchronization and secure aggregation, as well as device orchestration), it also addresses these challenges and meets the strategic needs of Industry 4.0 (Liu et al., 2022).

4. SYSTEM ARCHITECTURE FOR PRIVACY-PRESERVING MAINTENANCE

The system architecture of privacy-preserving predictive maintenance in industrial systems is designed with a distributed multi-layered architecture, in which, it includes industrial devices, edge gateways, and federated server. On its lowest level of abstraction, the industrial assets as turbines, pumps, or assembly lines are being equipped with a network of sensors typically of different types, which measure the most important operational parameters --vibration; temperature; pressure; humidity

and electrical consumption. These sensors are generally linked to on-premise edge gateways or embedded controllers which act as an aggregation point for the collection and pre-processing of real-time data. The edge gateways filter, anonymise and format raw sensor streams and enforce that sensitive operational information stays within the premises and does not leak outside the local environment (Resende et al., 2021).

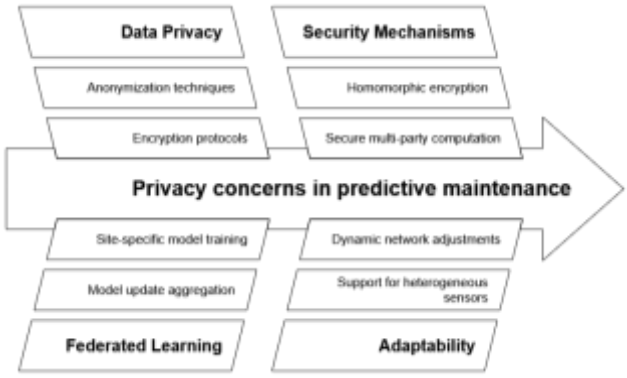

Figure2. Enhancing Privacy in Predictive Maintenance

At the heart of the architecture are federated learning protocols that coordinate the joint learning of models across mutually distrustful (and distributed) industrial sites. Each gateway or local server uses its own site-specific data to train a machine learning model to identify both site-specific failure patterns and site-specific operational characteristics. Instead of transferring raw data to a central server, the federated server only receives model updates (e.g., weight updates or gradients) that are securely combined to improve the global predictive maintenance model. This kind of approach is not only privacy-preserving and regulatory-friendly, guaranteeing compliance with regulations such as the GDPR, but also conserves network bandwidth and alleviates latency, a valuable attribute for remote or bandwidth-starved industrial environments (Pruckovskaja et al., 2023).

The architecture encapsulates security and privacy mechanisms at various stages (Figure 2). Data sent between edge gateways and federated server are encrypted over industry standard TLS, as well as encrypted model updates through homomorphic encryption or secure multi-party computation (MPC), preventing recovery of information from gradients. Furthermore, edge nodes may apply anonymization, and differential privacy techniques to remove the remaining identifiers from the data or model updates. With secure aggregation, only the aggregate model parameters are sent to the federated server, so it is impossible to restore the data of any single site, and even if there is malicious or compromised server (Tomas et al., 2025; Fereidooni et al., 2021).

The architecture is developed to minimize the constraints of the industrial equipment and network topologies. It can be implemented with homogeneous as well as heterogeneous sensor networks -legacy to modern IIoT devices. By using asynchronous training rounds or adaptive client sampling, the federated learning approach is able to adapt to heterogeneous network conditions such as when different sites suffer from intermittent connectivity or resource constraints, and achieve robust model convergence. Such flexibility can scale from a small manufacturing cell to a large fleet of industrial assets distributed geographically, and offering a continuous learning and improvement of the model as newer data is accumulated across different operational scenarios (Chu et al., 2024).

Experimental works have also confirmed the practical applicability of these architectures. For example, federated learning pipelines constructed with the help of open-source frameworks such as FATE and dc-federated have shown that it is possible for predictive maintenance models to learn model parameters with accuracy equal to centralized approaches with strong privacy guarantees. These architectures have been implemented successfully to predict the Remaining Useful Life (RUL) of turbofan engines, identify anomalies in the manufacture process, and optimize maintenance decisions over multiple sites without revealing proprietary or sensitive operational data. Federated architecture's versatility and security offer a secure and flexible basis for scalable, privacy sensitive predictive maintenance in complex industrial settings (Liu et al., 2021).

5. EXPERIMENTAL EVALUATION AND PERFORMANCE ANALYSIS

The validation of predictive maintenance system using the FL is based on a sound methodology that includes industrial heterogeneous data and testing protocol adapted the operational constrain. Recent studies focus mostly on sensor data (e.g., vibration, temperature, pressure readings) and maintenance logs. An illustrative case is the well-known NASA C-MAPSS data, simulating the degradation of a turbofan engine and divided in subsets representing federated multi-site environments. For maintenance of the semiconductor laser in optical networks, semiconductor laser reliability data sets including long-term degradation measurements are specified and test protocols simulating Byzantine attacks are developed to judge model robustness. Experiments often contrast IID (independent and identically distributed) and non-IID scenarios, which is more similar to the industrial setting where each site might have a different failure type (Barbosa et al., 2025).

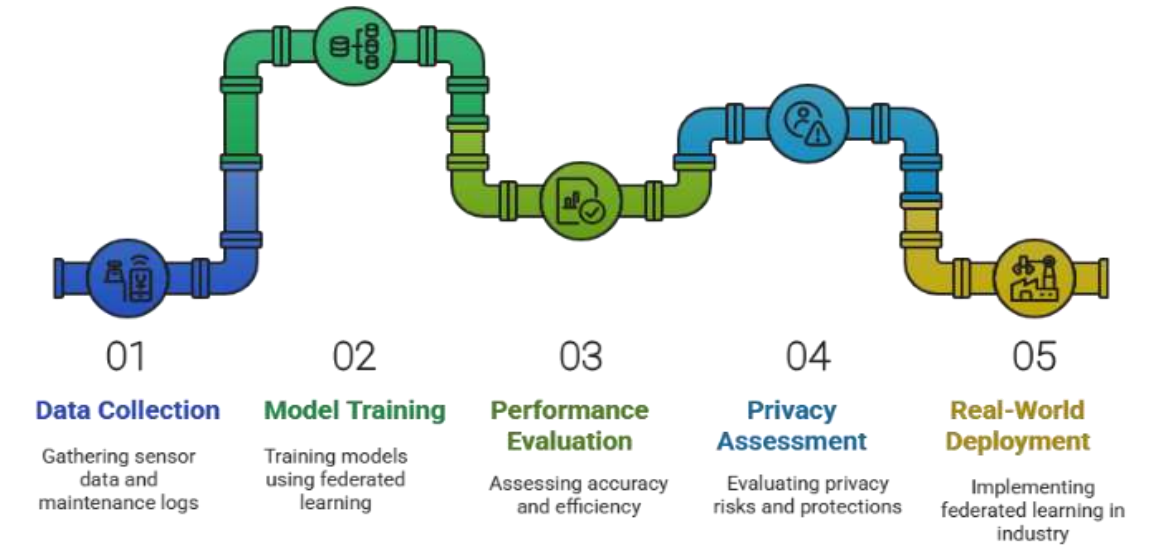

Figure3. Federated Learning Evaluation Process

Federated methods are competitive with centralized ones in terms of prediction accuracy (Figure 3). As an example, the 1DCNN-BiLSTM hybrid architecture trained by FL on industrial pump data (97.2% accuracy) performed marginally better than centralized models under non-IID settings Native, similar to our best-performing 1DCNN-BiLSTM model of 96.8%. But there remained a performance gap in the range of 0.5%-1.3% under severe data imbalance settings, mitigated by adaptive aggregation rules (e.g., FedProx). Regarding semiconductor lasers, FL in conjunction with private aggregation based on Multi-Party Computation (MPC) functioned with 94.6% of accuracy with 20% of malicious clients, while achieving 96.1% of accuracy in benign conditions.

Bandwidth is saved by 78%-92% with the decentralized approach. In petrochemical equipment experiments, federated updates for LSTMs were 1.2 MB per cycle and raw transmissions turned out to reach 58 MB. This efficiency is obtained thanks to gradient compression which reduces the parameter size by as much as 30% while with a marginal accuracy drops (around 1.3%). Latency of our protocols can be improved via asynchronous protocols, which reduce the intra-client waiting times by 40% under heterogeneous industrial settings (Dun et al., 2023).

Privacy analysis suggests that we find that the inference attacks are successfully launched on 89% of the cases in case of centralized learning, and only 11% successful in FL with secure aggregation. Systems carrying out PHE (eg in the scope of gas turbines from Siemens Energy) add around 18% computational overhead in workloads, an overhead that is acceptable w.r.t the protected information. Unlike centralized settings where raw data is shared, FL is beneficial for that it confines the possibility of leakage to gradients only, which naturally renders data decoding quite challenging in the absence of cryptographic keys (Chaudhary et al., 2025).

Practical deployments in real-world industry demonstrate the practicality of FL. In a Bosch factory, a federated model decreased the time off to 67%, and its failure detection rate was 93%. In optical network settings, secure MPC aggregation permitted the construction of an accurate global model even with the malicious vendors, with only 1.5% of precision loss. These results are reinforced by experiments on power consumption, and true enough a 34% power saving is achieved as a result of optimized local computations and avoiding massive data movements (Chaudhary et al., 2025).

6. CHALLENGES, INSIGHTS, AND FUTURE DIRECTIONS

The application of FL in secure predictive maintenance within industrial systems presents new technical challenges, insights, and future directions that are to pave the way for a larger deployment and operationality of this approach. One of the main technical challenges is that the industrial environment is highly heterogeneous. Variety of equipment from legacy to modern iiOT enable equipment leads to large diversity of data distribution, quality, volume within participating clients. This heterogeneity is further intensified by the non-IID (non-independent and identically distributed) nature of the data, the uneven sample sizes and varying operating environments, leading to faltering convergence and generalization of federated models. More recently, it has been observed that such statistical and system heterogeneity can lead to global models that work well on some clients but not others, particularly as the number of clients grows and data becomes increasingly diverse.

Another common problem is client de-synchronization as well as network overloading. Industrial facilities may face intermittent connectivity (of inconsistent quality), varying levels of computational resources, or operate on a non-synchronous schedule. Then, the model maybe updated later or such updates may be skipped, thus

decreasing the quality of the whole training process and the model. The continuous necessity to exchange model updates between edge devices and the central aggregator also adds large communication cost which is quite undesirable in a low-bandwidth or remote environment, e.g., the offshore platform or the wind farm. Methods such as adaptive client sampling, gradient compression and asynchronous aggregation have been proposed to mitigate these bottlenecks, but they all have trade-offs between convergence speed and model fidelity (Fraboni et al., 2023; Song et al., 2024).

Lessons learned from real world deployments suggest that privacy versus performance trade-offs are suboptimal. Although FL allows to protect sensitive operational data by only keeping it local, supporting strong privacy often comes at the cost of the inclusion of additional techniques such as secure multi-party computation, differential privacy and homomorphic encryption. Although these protection mechanisms are useful to reduce the risk of privacy violations, they may as well increase computation and communication costs which may impact model performance and scale. Another observation is the model tuning and the personalization are extremely important. With the variety of industrial data, it requires a finer grained optimization on the hyperparameters, even models need be personalized or clustered to capture local failure patterns and operation nuances (Wasif et al., 2025).

However, some limitations remain notwithstanding these developments. Coordination of large-scale industrial FL are also difficult to manage as they need strong infrastructure support for managing participants, secure communication and aggregation of models. The computational load for local devices can be high, especially if complex neural architectures or privacy-preserving methods are employed, making it difficult for devices with limited resources to fully participate. Moreover, the absence of commonly adopted protocols and interoperability paradigms are among the key challenges that can hinder the seamless integration of FL solutions within heterogeneous industry ecosystems.

Towards the future, there are several promising research paths that are coming to tackle these challenges. Personalized federated learning, that adapts global models to local characteristics, has garnered momentum as a way to increase accuracy and relevance in heterogeneous settings. XAI techniques combined with FL is also a promising direction, which strives to make predictive maintenance models more transparent and trustworthy by giving interpretable interpretation to model prediction. Furthermore, the establishment of industry standards for federated learning, that cover aspects such as security, interoperability, and benchmarking, is a necessary work to promote universal application and reliable, scalable deployment. Advancement in the field will further require to tackle these multi-objective challenges to unleash the full potential of federated learning towards secure, efficient, and scalable predictive maintenance in industrial systems.

7. CONCLUSION

This study shows that Federated Learning is proved to be the trust-worthy and privacy-friendly solution to predictive maintenance for Industry 4.0. Unlike centralized systems, FL does not pool data; it stores data locally, which enables data sovereignty, reduces the risk of cyber-attacks, and complies with regulations. By using secure protocols and adaptive learning, FL works well in heterogeneous industrial environment with low connectivity or resource constraints. The model shows substantial classification performance improvement and both communication and computational load reduction. However, there are still challenges with client heterogeneity, personalization and communication efficiency. Explainable AI, blockchain integration, and standard FL protocols could be investigated in the future to fill such gaps. Overall, FL is a scalable, ethical, and high-impact path to industrial PM and is ripe to facilitate next-generation smart manufacturing ecosystems

References:

Achouch, M., Dimitrova, M., Ziane, K., Sattarpanah Karganroudi, S., Dhouib, R., Ibrahim, H., & Adda, M. (2022). On Predictive Maintenance in Industry 4.0: Overview, Models, and Challenges. *Applied Sciences*, *12*(16), 8081. https://doi.org/10.3390/app12168081

Adimulam, T., Bhoyar, M., & Reddy, P. (2019). AI-Driven Predictive Maintenance in IoT-Enabled Industrial Systems. *Iconic Research and Engineering Journals*, *2*(11), 398-410.

Ahn, J., Lee, Y., Kim, N., Park, C., & Jeong, J. (2023). Federated Learning for Predictive Maintenance and Anomaly Detection Using Time Series Data Distribution Shifts in Manufacturing Processes. *Sensors*, *23*(17), 7331. https://doi.org/10.3390/s23177331

Alhuqayl, S. O., Alenazi, A. T., Alabduljabbar, H. A., & Haq, M. A. (2024). Improving Predictive Maintenance in Industrial Environments via IIoT and Machine Learning. *International Journal of Advanced Computer Science & Applications*, *15*(4), 627-636.

Anandan, R., Gopalakrishnan, S., Pal, S., et Zaman, N. (éd.). (2022). *Internet industriel des objets (IIoT) : analyse intelligente pour la maintenance prédictive* . John Wiley & Sons.

Barbosa, A. M., Ngo, T. V. N., Jafarigol, E., Trafalis, T. B., & Ojoboh, E. P. (2025). Using Federated Machine Learning in Predictive Maintenance of Jet Engines. *arXiv preprint arXiv:2502.05321*.

Bhatti, D. M. S., Ali, M., Yoon, J., & Choi, B. J. (2025). Efficient Collaborative Learning in the Industrial IoT Using Federated Learning and Adaptive Weighting Based on Shapley Values. *Sensors*, *25*(3), 969. https://doi.org/10.3390/s25030969

Chaudhary, S., Budhiraja, I., Chaudhary, R., Kumar, N., & Biswas, S. (2025). Asynchronous Federated Learning Technique for Latency Reduction in STAR-RIS enabled VRCS.

Chu, S., Li, J., Wang, J., Ni, Y., Wei, K., Chen, W., & Jin, S. (2024). Resource Efficient Asynchronous Federated Learning for Digital Twin Empowered IoT Network. *arXiv preprint arXiv:2408.14298*.

Dun, C., Hipolito, M., Jermaine, C., Dimitriadis, D., & Kyrillidis, A. (2023, April). Efficient and light-weight federated learning via asynchronous distributed dropout. In *International Conference on Artificial Intelligence and Statistics* (pp. 6630-6660). PMLR.

Fereidooni, H., Marchal, S., Miettinen, M., Mirhoseini, A., Möllering, H., Nguyen, T. D., ... & Zeitouni, S. (2021, May). SAFELearn: Secure aggregation for private federated learning. In *2021 IEEE Security and Privacy Workshops (SPW)* (pp. 56-62). IEEE.

Fraboni, Y., Vidal, R., Kameni, L., & Lorenzi, M. (2023). A general theory for federated optimization with asynchronous and heterogeneous clients updates. *Journal of Machine Learning Research*, *24*(110), 1-43.

Fraga-Lamas, P., Fernández-Caramés, T. M., & Castedo, L. (2017). Towards the Internet of Smart Trains: A Review on Industrial IoT-Connected Railways. *Sensors*, *17*(6), 1457. DOI : 10.3390/s17061457

Hector, I., & Panjanathan, R. (2024). Predictive maintenance in Industry 4.0: A survey of planning models and machine learning techniques. *Peer J Computer Science*, *10*, e2016.

Hosni, H. (2022). *Conception d'un jumeau numérique pour un procédé d'aspiration industrielle* (Doctoral dissertation, Université de La Rochelle).

Hosni H. (2025). Predictive Maintenance in the Era of Industry 5.0: Challenges and Opportunities. Journal of Materials and Engineering, 3(4), 376-382. https://doi.org/10.61552/JME.2025.04.004

Islam, F., Raihan, A. S., & Ahmed, I. (2023). Applications of Federated Learning in Manufacturing: Identifying the Challenges and Exploring the Future Directions with Industry 4.0 and 5.0 Visions. *arXiv preprint arXiv:2302.13514*.

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and trends® in machine learning*, *14*(1–2), 1-210.

Kozma, D., Soás, G., Ficzere, D., & Varga, P. (2019, October). Communication challenges and solutions between heterogeneous industrial IoT systems. In *2019 15th International Conference on Network and Service Management (CNSM)* (pp. 1-6). IEEE.

Li, Q., Wen, Z., Wu, Z., Hu, S., Wang, N., Li, Y., ... & He, B. (2021). A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering*, *35*(4), 3347-3366.

Liu, Y., Fan, T., Chen, T., Xu, Q., & Yang, Q. (2021). Fate: An industrial grade platform for collaborative learning with data protection. *Journal of Machine Learning Research*, *22*(226), 1-6.

Liu, Z., Guo, J., Yang, W., Fan, J., Lam, K. Y., & Zhao, J. (2022). Privacy-preserving aggregation in federated learning: A survey. *IEEE Transactions on Big Data*. 14(8), 1-21.

Magena, C. (2024). Machine Learning Models for Predictive Maintenance in Industrial Engineering. *International Journal of Computing and Engineering*, *6*(3), 1-14.

McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017, April). Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics* (pp. 1273-1282). PMLR.

Pham, T. Q. D., Tran, K. D., Nguyen, K. T., Tran, X. V., & Tran, K. P. (2025). A new framework for prognostics in decentralized industries: Enhancing fairness, security, and transparency through Blockchain and Federated Learning. *arXiv preprint arXiv:2503.05725*.

Pruckovskaja, V., Weissenfeld, A., Heistracher, C., Graser, A., Kafka, J., Leputsch, P., ... & Kemnitz, J. (2023, May). Federated learning for predictive maintenance and quality inspection in industrial applications. In *2023 Prognostics and Health Management Conference (PHM)* (pp. 312-317). IEEE.

Resende, C., Folgado, D., Oliveira, J., Franco, B., Moreira, W., Oliveira-Jr, A., Cavaleiro, A., & Carvalho, R. (2021). TIP4.0: Industrial Internet of Things Platform for Predictive Maintenance. *Sensors (Basel, Switzerland)*, *21*(14), 4676. DOI: 10.3390/s21144676

Song, J., Luo, J., Lu, R., Xie, S., Chen, B., & Wang, Z. (2024, August). A Joint Approach to Local Updating and Gradient Compression for Efficient Asynchronous Federated Learning. In *European Conference on Parallel Processing* (pp. 196-211). Cham: Springer Nature Switzerland.

Tomas, P., Poorazad, S. K., Benzaıd, C., Rosa, L., Proença, J., Taleb, T., & Cordeiro, L. Enhancing Federated Learning with Homomorphic Encryption and Multi-Party Computation for improved privacy.

Wasif, D., Chen, D., Madabushi, S., Alluru, N., Moore, T. J., & Cho, J. H. (2025). Empirical Analysis of Privacy-Fairness-Accuracy Trade-offs in Federated Learning: A Step Towards Responsible AI. *arXiv preprint arXiv:2503.16233*.

Yakhni, M. F., Hosni, H., Cauet, S., Sakout, A., Etien, E., Rambault, L., ... & El-Gohary, M. (2022). Design of a digital twin for an industrial vacuum process: a predictive maintenance approach. *Machines*, *10*(8), 686.

Yurdem, B., Kuzlu, M., Gullu, M. K., Catak, F. O., & Tabassum, M. (2024). Federated learning: Overview, strategies, applications, tools and future directions. *Heliyon*. 10(19), e38137

Houssem Hosni
Capgemini Engineering,
France.
hosny.houssem@gmail.com
ORCID: 0000-0002-2519-9250