CYBERSECURITY PEDAGOGY & PRACTICE JOURNAL

Volume 5, No. 2
October 2026
ISSN: 2832-1006

In this issue:

4. **Bridging the Cybersecurity Skills Gap by Leveraging Cybersecurity Clinics: Developing the Arizona Cybersecurity Clinic**
Paul Wagner, University of Arizona
Robert Honomichl, University of Arizona
Shengjie Xu, University of Arizona

34. **What is Cybersecurity (or Cyber Security)? Answer [Still] Loading…**
Maguire Parsley, University of North Carolina Wilmington
Luke Warpula, University of North Carolina Wilmington
Ulku Clark, University of North Carolina Wilmington
Geoff Stoker, University of North Carolina Wilmington

59. **Enhancing Cyber Defense: A Federated Multi-Modal Deep Learning Framework for Privacy-Preserving Zero-Day Attack Detection**
Samuel Sambasivam, Woodbury University

78. **Evaluating Eleven Online CTF Platforms for Undergraduate Cybersecurity Competition Preparation**
Brandon P. Grech, Anderson University (SC)

104. **Federated Learning and Edge Computing for Privacy-Preserving Real-Time Predictive Maintenance in Industrial IoT Systems**
Kiran Kumar Vejendla, City University of Seattle
Kyong (Jin) Chang, City University of Seattle
Brian Maeng, City University of Seattle
Sherin Aly, City University of Seattle

117. **Perceived Usefulness of Password Complexity: A Methodology Framework**
Kyle Herman, Georgia Southern University
Hayden Wimmer, Georgia Southern University

132. **Beyond Awareness: Building Digital Resilience Through Comprehensive Cybersecurity Foundations for K-5 Students**
Ellie Ebrahimi, University of North Carolina Wilmington
Yasin Emre Gokce, University of North Carolina Wilmington
Majorie Pare, University of North Carolina Wilmington
Berker Koyuncu, University of North Carolina Wilmington
Jude Vargas, University of North Carolina Wilmington
Jacob Blount, University of North Carolina Wilmington

iscap
Information Systems & Computing Academic Professionals

The **Cybersecurity Pedagogy and Practice Journal** (**CPPJ**) is a double-blind peer-reviewed academic journal published by **ISCAP** (Information Systems and Computing Academic Professionals). Publishing frequency is two times per year. The first year of publication was 2022.

CPPJ is published online (https://cppj.info). Our sister publication, the proceedings of the ISCAP Conference (https://proc.iscap.info) features all papers, panels, workshops, and presentations from the conference.

The journal acceptance review process involves a minimum of three double-blind peer reviews, where both the reviewer is not aware of the identities of the authors and the authors are not aware of the identities of the reviewers. The initial reviews happen before the ISCAP conference. At that point, papers are divided into award papers (top 15%), and other accepted proceedings papers. The other accepted proceedings papers are subjected to a second round of blind peer review to establish whether they will be accepted to the journal or not. Those papers that are deemed of sufficient quality are accepted for publication in the CPPJ journal.

While the primary path to journal publication is through the ISCAP conference, CPPJ does accept direct submissions at https://iscap.us/papers. Direct submissions are subjected to a double-blind peer review process, where reviewers do not know the names and affiliations of paper authors, and paper authors do not know the names and affiliations of reviewers. All submissions (articles, teaching tips, and teaching cases & notes) to the journal will be refereed by a rigorous evaluation process involving at least three blind reviews by qualified academic, industrial, or governmental computing professionals. Submissions will be judged not only on the suitability of the content but also on the readability and clarity of the prose.

Currently, the acceptance rate for the journal is under 35%.

Questions should be addressed to the editor at editorcppj@iscap.us or the publisher at publisher@iscap.us. Special thanks to members of ISCAP who perform the editorial and review processes for CPPJ.

2026 ISCAP Board of Directors

Amy Connolly James Madison University President	Michael Smith Georgia Institute of Technology Vice President	Jeff Cummings Univ of NC Wilmington Past President
David Firth University of Montana Director	Mark Frydenberg Bentley University Director/Secretary	Leigh Mutchler James Madison University Director
RJ Podeschi Millikin University Director/Treasurer	Bryan Reinicke Rochester Institute of Technology / Director	Jeffry Babb West Texas A&M University Director/Curricular Matters
Eric Breimer Siena University Director/2026 Conf Chair	Tom Janicki Univ of NC Wilmington Director/Meeting Planner	Xihui "Paul" Zhang University of North Alabama Director/JISE Editor

Copyright ©2026 by Information Systems and Computing Academic Professionals (ISCAP). Permission to make digital or hard copies of all or part of this journal for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial use. All copies must bear this notice and full citation. Permission from the Editor is required to post to servers, redistribute to lists, or utilize in a for-profit or commercial use. Permission requests should be sent to editorcppj@iscap.us.

CYBERSECURITY PEDAGOGY AND PRACTICE JOURNAL

Editors

Jeffrey Cummings
Co-Editor
University of North Carolina Wilmington

Anthony Serapiglia
Co-Editor
Saint Vincent College

Thomas Janicki
Publisher
University of North Carolina Wilmington

2026 Review Board

Brandon Brown
Coastline College

Shawn Clouse
University of Montana

Jeff Landry
Univ of South Alabama

Li-Jen Lester
Sam Houston State Univ

Jamie Pinchot
Robert Morris University

Samuel Sambasivam
Woodbury University

Shwadhin Sharma
California State University Monterey Bay

Sushma Mishra
Robert Morris University

Kevin Slonka
Saint Francis University

Geoff Stoker
Univ of NC Wilmington

Paul Wagner
University of Arizona

Paul Witman
California Lutheran University

Federated Learning and Edge Computing for Privacy-Preserving Real-Time Predictive Maintenance in Industrial IoT Systems

Kiran Kumar Vejendla
vejendlakiran@gmail.com

Kyong (Jin) Chang
changkyong@cityu.edu

Brian Maeng
maengjooyol@cityu.edu

Sherin Aly
alysherin@cityu.edu

City University of Seattle
Seattle,WA 98121

Abstract

This paper presents a framework involving Federated Learning (FL) and Edge computing (EC) (FL-EC) to achieve privacy-preserving and real-time predictive maintenance in the Industrial Internet of Things (IIoT) systems. FL reduces the privacy risk by training machine learning models on local edge-devices without the transmission of sensitive data, whereas EC executes data nearer to the source decreasing latency and increasing responsiveness. The FL-EC framework demonstrates high accuracy, F1-score, and recall compared to traditional centralized models, establishing its utility for real-time fault detection. Scalability is also provided by the framework, making the system effective to support different and heterogeneous edge devices. FL-EC helps to solve the most important problems of traditional IIoT systems by maintaining data confidentiality and limiting latency. The research also contrasts FL-EC with the available literature and shows that it can be applied practically in the automotive and oil and gas industry as well as in industries where real-time monitoring and maintenance are essential. The results indicate that FL-EC has the potential to achieve significant cost-saving and efficiency improvement in industry. The future research must aim at improving federated learning algorithms, solving the non-IID data problem, and enabling the optimization of edge devices to be better performing and customized, which will guarantee a wider application in IIoT settings.

Keywords: Federated Learning, Edge Computing, Predictive Maintenance, IIoT, Privacy-Preserving, Real-Time Fault Detection

Recommended Citation: Vejendla, K., Chang, K., Maeng, B., Aly, S. (2026). Federated Learning and Edge Computing for Privacy-Preserving Real-Time Predictive Maintenance in Industrial IoT Systems. *Cybersecurity Pedagogy and Practice Journal.* v5(n2) pp 104-116. DOI# https://doi.org/10.62273/GPYS1669

©2026 ISCAP (Information Systems and Computing Academic Professionals)

Federated Learning and Edge Computing for Privacy-Preserving Real-Time Predictive Maintenance in Industrial IoT Systems

Kiran Kumar Vejendla, Kyoung (Jin) Change, Brian Maeng and Sherlin Aly

1. INTRODUCTION

The Industrial Internet of Things (IIoT) has transformed how industries oversee and control equipment, making it possible to shift current maintenance approaches based on conventional maintenance tools to more dynamic ones, including predictive maintenance (PM) (Naidoo & Sibanda, 2024). Predictive maintenance involves real-time information collected through interconnected sensors and is used to predict equipment malfunction prior to its happening to enable an organization to maximize its operations, minimize downtimes and increase its output (Civerchia et al., 2017). Nevertheless, the increasing dependency on centralized machine learning models in IIoT systems causes a number of acute problems, in particular, with privacy, latency, and scalability.

Conventional predictive maintenance systems face major challenges hindering their adoption in IIoT settings. To begin with, there is a privacy issue because vast amounts of sensitive operational information need to be sent between the edge devices to centralized servers to process information, which puts pressure on risks of data leakage and breaches of data protection laws, including the General Data Protection Regulation (GDPR) (Ali et al., 2021). Besides, centralized systems can create delays in real-time fault detection and decision-making since it requires time to send the data and make predictions at a central point. The delay can be incredibly problematic especially in those industries that demand instant responses, like automotive or oil and gas (Bhaskaran et al., 2020). Lastly, being scalable is also a major concern since centralized systems are not able to cope with large volumes of data that are produced by a wide range of diverse and geographically dispersed devices and equipment in industrial settings. Such constraints underscore the need to have a more decentralized, efficient and privacy preserving predictive maintenance in IIoT.

To address these challenges, this study proposes a combination of Federated Learning (FL) and Edge Computing (EC) that can facilitate privacy-protected real-time predictive maintenance in IIoT systems (Alshkeili et al., 2025). Federated Learning (FL) is a machine learning method that is decentralized, which enables the local training of models without requiring the transfer of raw data to central servers. FL can suppress transmission of data, as well as lower the risk of privacy by aggregating locally trained models (Beltrán et al., 2023). Conversely, EC processes data nearer to the source that is at the edge of the network, hence, greatly reducing latency and makes it possible to provide fault detection analysis in real-time. FL combined with EC is the promising solution to addressing the intrinsic issues of privacy, latency and scalability of traditional IIoT predictive maintenance systems (Liu et al., 2024).

This model provides the assurance that sensitive information is held on peripheral devices, and it complies with the stringent privacy rules and, nevertheless, allows the worldwide compilation of models to make correct predictions of faults. Furthermore, the capacity of edge computing to deal with real-time data processing at the edge prevents the reduction in the latency of the decision-making process, providing a more reactive method of predictive maintenance. Therefore, the integrated FL-EC framework can overcome the main shortcomings of centralized systems, namely the ability to scale to larger heights, ensuring security, and enabling maintenance of IIoT processes in real-time. This research addresses the following key questions:

1. How effectively does Federated Learning (FL) mitigate privacy risks while maintaining model accuracy?
2. How does the combination of FL and Edge Computing (EC) affect the process of anomaly detection and feature extraction in sensor data to predictive maintenance of IIoT?
3. What is the benefit of the proposed decentralized structure in terms of detecting faults in real time by learning temporal dependencies in industrial time-series data as well as the issue of privacy, latency, and scalability?
4. Is the edge computing/FL combination less latent than centralized predictive maintenance systems?

©2026 ISCAP (Information Systems and Computing Academic Professionals)

5. To what extent can the proposed framework be scaled when evaluated with a variety of industrial configurations with different types of machines and sensor configurations?

Using the advantages of Federated Learning and Edge Computing, this study will discuss a new form of predictive maintenance that will effectively solve the weaknesses of centralized systems. The suggested framework is an excellent chance to address the challenge of privacy, latency, and scalability of the traditional IIoT systems, which have become a part of the current process of digital transformation of industries. With the ongoing development of IIoT and the proliferation of this device type, this combined FL-EC model can potentially become a key enabler of safer, more efficient and reliable predictive maintenance systems and create operational gains and cost reductions in a variety of industries.

Previous Research Overview

The Industrial Internet of Things (IIoT) is important in facilitating predictive maintenance (PdM) by constantly tracking machines and equipment to identify the initial indicators of possible failure. Recent innovations have focused on changing the conventional reactive maintenance methods to more effective predictive-based maintenance methods based on real-time data that are detected by different sensors in the industrial setting (Hafeez et al., 2021). An example of this, Civerchia et al. (2017) discussed the application of IIoT to advanced predictive maintenance describing a system applicable to a real-world electricity power plant. Their solution showed a notable decrease in downtimes and a better reliability of the equipment based on real-time sensor data. Nevertheless, one of the remaining issues is the sheer volume of information produced by IIoT devices that makes it difficult to process and transmit them to cloud-based systems.

Federated Learning (FL) has become a promising approach to privacy-aware machine learning in IIoT environments. FL supports collaborative model training across edge devices without transferring their raw operational data to a central repository. Fan et al. (2022) proposed a privacy-preserving aggregation scheme that combines secret sharing with the Paillier cryptosystem, an additive homomorphic public-key encryption method that permits selected arithmetic operations on encrypted values without first revealing the underlying plaintext. This combination is intended to protect model updates and reduce exposure to model-update reconstruction or inference attacks while controlling communication and computation costs. Odeh et al. (2024) subsequently introduced asynchronous FL for mobile-edge IIoT ecosystems, addressing device heterogeneity and network delay in decentralised training.

The edge computing (EC) has also become a part of the IIoT as it allows processing data nearer to the data generation point, minimizing the time and transmission expenses. Hafeez et al. (2021) emphasized the role of EC in addressing the high real-time demands of IIoT applications and how it is used in predictive maintenance in particular. They talked about the use of EC to handle the processing of large volumes of sensor data, decrease load on cloud infrastructure, and enhance the speed of decision-making. Besides, Liu et al. (2024) introduced a federated learning-based edge computing system that integrates the processing capabilities of edge nodes with the privacy-insurance characteristics of FL to be able to process data efficiently and provide security. This integration will be of particular importance in IIoT environments where the reduction of latency and the real-time performance are the primary considerations.

The combination of FL and EC are compatible and can help address the problems of IIoT-based predictive maintenance. By combining those two technologies, it is possible to train models in a decentralized manner, but at the same time, make sure that data privacy is not violated, and processing does not take too much time. Liu et al. (2024) examined the possibility of edge computing to minimize the latency by computing data on the device and introduced a frame of devices selection and computation offloading to improve the federated learning process further. Their architecture proves the cooperation of FL and EC to handle the huge volumes of data IIoT devices produce and provide more effective and more privacy-aware predictive maintenance solutions.

Privacy remains a concern in IIoT federated systems because adversaries may attempt to infer sensitive information from model updates. Fan et al. (2022) combined Practical Byzantine Fault Tolerance (PBFT), secret sharing, and Paillier-based encrypted aggregation to support fault-tolerant coordination when devices fail or behave maliciously. These mechanisms are designed to reduce exposure to reconstruction and tampering risks; however, their effectiveness depends on the assumed threat model, key management, and implementation details.

Although these are optimistic advances regarding federated learning and edge computing regarding IIoT-based predictive maintenance, a number of challenges still persist. Among the most prominent ones is that data in IIoT settings is not IID (Non-Independent and Identically Distributed) and, therefore, this aspect may adversely affect the functionality of federated learning models (Hongbin & Zhi, 2023). The other difficulty is that such systems are difficult to scale when used in large and heterogeneous industrial networks. More studies are required to come up with more efficient FL algorithms capable of addressing the challenges of non-IID data and still ensuring privacy and security. Moreover, it is necessary to have better strategies of device selection and computation offloading to make better use of edge computing resources and minimize the energy usage and latency of predictive maintenance systems (Liu et al., 2024).

As modern software architectures have shifted toward distributed, microservice-based, and cloud-native paradigms, Application Performance Monitoring (APM) has evolved from a reactive diagnostic capability into a fundamental operational requirement. Modern observability practices provide continuous runtime visibility into application behavior, enabling organizations to detect performance degradation, prevent service-level agreement (SLA) violations, optimize resource utilization, and maintain service reliability in increasingly complex distributed environments (Usman et al., 2022).

Research gap and incremental contribution. Existing studies demonstrate that FL and EC can be combined, but they frequently concentrate on one component of the problem. Fan et al. (2022) emphasise privacy-preserving aggregation, whereas Liu et al. (2024) emphasise edge resource selection and computation offloading. However, existing frameworks have largely overlooked the critical component of monitoring. This paper does not claim to introduce the first FL-EC architecture, a new cryptosystem, or a new aggregation optimiser. Its contribution is an application-specific design-science artifact and a unified empirical matrix that evaluates centralised and federated CNN/LSTM models, multiple aggregation strategies, a hybrid CNN-LSTM, secure versus non-secure exchange, latency, scalability, and PySyft-based distributed simulation for the same predictive-maintenance context. The novelty is therefore incremental and integrative: it shows how architectural, learning, and security choices interact rather than evaluating them in isolation in an enterprise context.

2. RESEARCH METHODOLOGY

This study adopted a mixed-methods methodology with an exploratory sequential design. It began with a qualitative phase, utilizing a literature review of existing frameworks, followed by a quantitative, experiment-based evaluation to construct and assess a privacy-aware, low-latency predictive-maintenance artifact for an IIoT setting. The procedure comprised four linked phases: (i) requirements and problem formalisation, (ii) framework design and development, (iii) experimental simulation and comparative evaluation, and (iv) performance, latency, scalability, and security-impact assessment.

Phase 1: Requirements and problem formalisation. A focused review of IIoT predictive-maintenance challenges was used to identify three operational requirements: local data confidentiality, real-time responsiveness, and scalability across heterogeneous edge clients. These requirements were translated into measurable criteria comprising accuracy, precision, recall, F1-score, latency, and communication overhead.

Phase 2: Framework design and development (artifact creation). The artifact was derived by mapping each operational requirement to an architectural mechanism. The privacy requirement led to local data retention, parameter-only exchange, and support for protected aggregation; the latency requirement led to edge-level preprocessing and inference; and the scalability requirement led to modular client, communication, orchestration, and aggregation functions. The designs reported by Fan et al. (2022) and Liu et al. (2024) were used as reference points, but the components were reorganised into a three-layer predictive-maintenance artifact and linked to a common comparative evaluation plan. This requirement-to-component mapping explains how the proposed framework was created and how each element relates to a stated research objective.

Phase 3: Experimental simulation and comparative evaluation. The framework was evaluated with the publicly available AI4I 2020 Predictive Maintenance Dataset from the UCI Machine Learning Repository. The observations were partitioned to represent multiple edge clients under decentralised conditions. Centralised baselines were compared with FL-EC

variants using Federated Averaging (FedAvg), Federated Proximal (FedProx), Federated Normalised Averaging (FedNova), and Federated Dynamic Regularisation (FedDyn). Fault-prediction architectures comprised a Convolutional Neural Network (CNN), a Long Short-Term Memory network (LSTM), and a hybrid CNN-LSTM model.

Phase 4: Performance, latency, scalability, and security-impact assessment. Classification performance was measured using accuracy, precision, recall, and F1-score. End-to-end latency incorporated local inference and model-update time in the centralised and edge-enabled configurations. Security overhead was examined by comparing secure and non-secure federated variants, while scalability was assessed by increasing the number of simulated clients and observing model stability and communication overhead.

The four-phase procedure links each design decision to an operational requirement and each research claim to a corresponding experimental measure. The framework is therefore presented as a systematically derived and empirically assessed integration artifact rather than as a purely conceptual architecture.

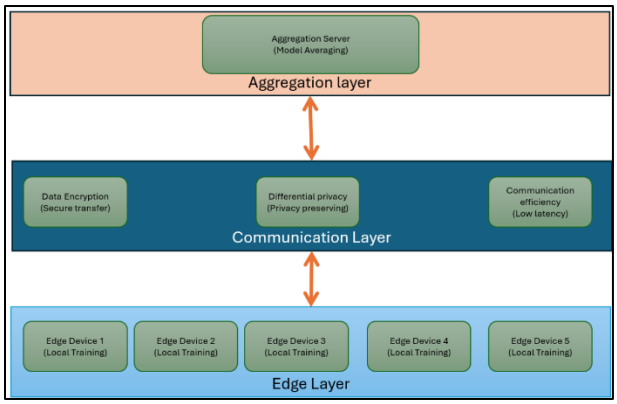


Figure 1: Proposed System Architecture

Figure 1 illustrates a three-layer FL–EC architecture. Edge devices perform local model training on IIoT data. The communication layer ensures secure, privacy-preserving, and low-latency parameter exchange. The aggregation layer combines local updates into a global model, enabling scalable, real-time, and privacy-aware predictive maintenance.

Dataset

In this research, the predictive maintenance Dataset (AI4I 2020) is taken as the main source of data. This is a dataset consisting of sensor data of industrial machines, and is applied to predictive maintenance. The UCI Machine learning Repository offers AI4I 2020 data, which covers the values of different operational parameters like temperature, vibration, and pressure that are crucial in identifying possible failures in the machine (Matzka., 2020).

There are 10,000 instances in the dataset, 22 features and a target variable as operational condition of the machine. The characteristics are the number of sensor readings and the target variable is the failure state of a machine (1) or the normal operating state of a machine (0). In predictive maintenance, the exercise is to model and predict whether the machine can fail or not basing on its past sensor reading. Such secondary data will be a source of valuable information that may be used to train machine learning models to predict equipment failures in industrial IoT (IIoT) systems, with no need to gather primary data.

Federated Learning Framework

The study uses Federated Learning (FL) to a dataset of AI4I 2020 to design a predictive maintenance model that is privacy-preserving. FL permits the distributed construction of machine learning models in many edge-computers, without communicating the original data. It is especially appropriate in IIoT, where privacy and data security are the primary priority.

The federation of learning model is a system where edge devices (connecting to various subsets of the dataset) use their data to train local models. This local model is combined by a central server to form a global model. FL framework guarantees that the sensitive information, including individual machine sensor readings is not stored on any remote server thereby protecting the privacy (Fan et al., 2022). In this study, we have applied a FedAvg (Federated Averaging) model-training algorithm. The most used federated learning algorithm is FedAvg in which every edge device trains a local model on their local data and transmits the model weights back to the central server. The server then sums up the model weights to generate the global model. This procedure minimizes the necessity of having physical access to the raw information but still making sure there is an ability to learn on the model using the aggregated insights across the devices.

Edge Computing Setup

The Edge Computing (EC) system applied in the study is aimed at reducing latency and enhancing the efficiency of the predictive maintenance architecture. EC provides edge processing of

data, which is nearer to the data source (i.e. the machines), thereby limiting the volume of data sent to the cloud and decision-making is faster. Here, the edge devices have the responsibility of doing preprocessing sensor data prior to transmitting aggregated models to the server where they learn further (Hafeez et al., 2021).

To conduct the secondary data analysis, EC is modeled by local data processing through the use of AI4I dataset. We do not transmit the raw sensor data to a central server, but instead simulate the operation of machine learning model training on local bits of data and combining them through the FL framework. This minimizes the overhead in communication and faster model updates are guaranteed.

Model Development
The main aim of the research project is to come up with a predictive maintenance model that is based on machine learning algorithms, capable of forecasting the possible failures of the industrial machines. The models that will be applied in this study are:

Convolutional Neural Network (CNN): CNNs are applied mainly to spatial feature extraction, thus it is more applicable to sensor data that could have spatial patterns. CNNs are able to identify anomalies in the machine sensors data by learning the features of the temperature, pressure, and vibration patterns that are usually predictive of a machine failure (Scarpa et al., 2018).
Long Sho
rt-Term Memory (LSTM): LSTMs are a type of Recurrent Neural Networks (RNN) which is very effective in the learning of temporal dependencies (Zhang et al., 2021). In this study, LSTMs are employed to forecast machine failures based on time-series information, which is necessary to model time-specked machine health data in IIoT settings.

Hybrid CNN-LSTM Model: This model is a combination of the two CNN and LSTM. The sensor data is extracted into features with CNN and modeled with LSTM to capture the time dependencies of features (Cheng et al., 2019). This hybrid methodology is especially applicable to predictive maintenance where it is possible to combine both spatial (e.g. unusual vibrations) and temporal (e.g. gradual wear and tear over time) patterns in machine behavior.

Evaluation Metrics
Model effectiveness was evaluated using standard classification and system-level measures. Accuracy is the proportion of all predictions classified correctly. Precision is the proportion of predicted failures that are genuine failures, whereas recall is the proportion of genuine failures detected by the model. The F1-score is the harmonic mean of precision and recall and is useful when balancing false alarms against missed failures. Latency is the elapsed time required to preprocess an observation, perform inference, and complete the relevant model-update operation. Scalability is assessed by examining whether performance and communication behaviour remain stable as the number of participating federated clients increases.

Experimental Setup
The section provides the description of the experimental set up in which the proposed FL-EC framework will be tested under controlled and reproducible conditions. The experiments were performed based on the AI4I 2020 Predictive Maintenance dataset consisting of 10,000 cases with 22 sensor-derived and operational features of the behavior of industrial machines. The binary target variable suggests normal functioning or failure and thus, the dataset can be used in fault prediction activities in IIoT.

Client Partition Strategy (IID / Non-IID)
In order to model a federated IIoT setup, the data was distributed among several virtual edge clients. Both non-IID and IID data distributions were taken into consideration. Non-IID partitions are used to simulate practical industrial environments in which machines run with varying loads and under varying conditions, and the data heterogeneity is introduced among clients.

Baselines
The FL-EC models were compared with centralised CNN and LSTM baselines. The federated environment considered Federated Averaging (FedAvg), which computes a data-weighted average of client updates; Federated Proximal (FedProx), which adds a proximal term to limit local-model drift; Federated Normalised Averaging (FedNova), which normalises client updates to account for unequal local work; and Federated Dynamic Regularisation (FedDyn), which uses dynamic regularisation to improve optimisation under heterogeneous data.

Hyperparameters and Training Configuration
All experiments were trained on fixed hyperparameters to ensure fairness. The federated training was done on several rounds where each client had a specified number of local

epochs in each round then model aggregation was done.

Evaluation Metrics and Latency Measurement

The measures of performance were through accuracy, precision, recall, and F1-score. Latency was determined as the end-to-end response time, including the local inference and model update latencies, which allowed straightforward comparison between centralized and FL–EC systems.

Proposed FL–EC Framework

The proposed FL-EC framework is an application-specific integration artifact for predictive maintenance; it is not presented as the first general FL-EC architecture. Its purpose is to organise local sensing, edge processing, protected parameter exchange, federated orchestration, and comparative evaluation within one reproducible workflow focused on privacy, latency, scalability, and fault-detection performance in an enterprise context where monitoring is highly required.

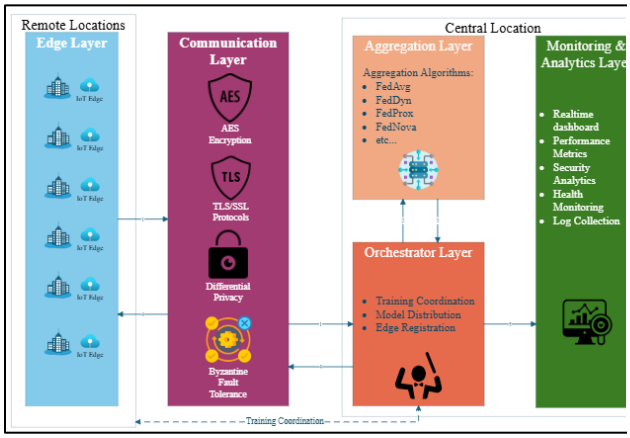


Figure 2: Proposed Federated Predictive Maintenance Framework

Figure 2 presents the end-to-end architecture. At the edge layer, IIoT devices or edge nodes acquire and preprocess operational observations, train local model copies, and perform local fault inference. Raw observations remain within the local client boundary; only model parameters or updates are prepared for transmission. This separation reduces routine raw-data movement and enables low-latency local decisions, although FL alone does not eliminate all information-leakage risks.

The communication layer connects the edge clients with the central orchestrator. It supports authenticated parameter transfer and, in the secure variant, protected aggregation or encrypted update exchange. These controls are intended to reduce unauthorised access and update reconstruction rather than to claim complete elimination of adversarial risk. Communication efficiency is evaluated because the additional protection may increase computation, payload size, or convergence time.

The aggregation and orchestration layer selects participating clients, coordinates training rounds, combines local updates using the selected federated optimiser, and redistributes the updated global model. In contrast with Fan et al. (2022), whose primary contribution is secure aggregation, and Liu et al. (2024), whose framework emphasises device selection and computation offloading, the present artifact centres on a unified predictive-maintenance comparison across model architectures, aggregation methods, security settings, latency, scalability, and monitoring. This integrative evaluation constitutes the principal contribution of the framework.

High-level algorithmic flow. (1) Initialise a global CNN, LSTM, or CNN-LSTM model and select the participating edge clients. (2) Each client preprocesses its local observations and trains a local model for the specified number of local epochs. (3) The client transmits only its model update, using the protected communication option when the secure variant is evaluated. (4) The orchestrator aggregates client updates using FedAvg, FedProx, FedNova, or FedDyn. (5) The updated global model is redistributed and the process repeats for the configured communication rounds. (6) The final model is evaluated through local inference, predictive metrics, latency, convergence behaviour, communication overhead, and client-scaling tests.

3. RESULTS AND ANALYSIS

The results are reported in distinct experimental stages to prevent values from different runs being compared as though they came from one configuration. The centralized baselines correspond to the standalone CNN, LSTM, and hybrid CNN-LSTM evaluations. The federated results then report mean performance for each model–algorithm combination across the simulated federated experiments, followed by pooled algorithm-level summaries, secure-versus-standard server effects, efficiency analysis, and the separate NetworkFed/PySyft evaluation. This ordering preserves the

experimental hierarchy and avoids mixing model-specific, pooled, and networked results.

Table 1. Centralized baseline model performance

Model	Acc.	Prec.	Rec.	F1
CNN	0.9465	0.9712	0.9465	0.9558
LSTM	0.9430	0.9718	0.9430	0.9537
Hybrid CNN-LSTM	0.9320	0.9708	0.9320	0.9465

Table 1 summarizes the centralized baseline results. CNN achieved the highest centralized accuracy (0.9465) and F1-score (0.9558), followed closely by LSTM with 0.9430 accuracy and 0.9537 F1. The hybrid CNN-LSTM remained strong but lower at 0.9320 accuracy and 0.9465 F1. These values establish the centralized reference points for the subsequent federated experiments.

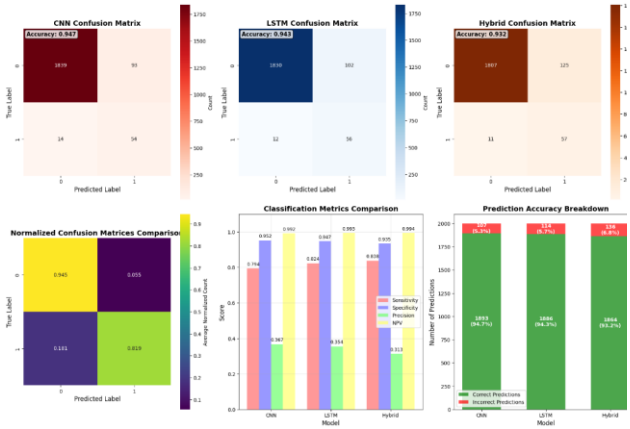


Figure 3. Centralized CNN, LSTM, and hybrid model confusion matrices and comparative performance

Figure 3 confirms the same ordering. CNN produced 1,839 true negatives and 54 true positives, LSTM produced 1,830 true negatives and 56 true positives, and the hybrid model produced 1,807 true negatives and 57 true positives. The comparative confusion matrices show that all three centralized models remain above 0.93 accuracy, while failure-class sensitivity is lower than overall accuracy because the underlying classification problem is imbalanced.

Federated simulation results were then evaluated across the four aggregation algorithms evaluated in the study: FedAvg, FedDyn, FedNova, and FedProx. Table 2 reports the mean model–algorithm values averaged across the relevant simulated distributions and server settings. These means must not be confused with the centralized baselines in Table 1 or with the pooled algorithm averages reported later.

Table 2. Federated model–algorithm mean performance

Configuration	Acc.	F1	Prec.	Rec.
CNN-FedAvg	0.5137	0.3585	0.3453	0.5137
CNN-FedDyn	0.5000	0.3333	0.2500	0.5000
CNN-FedNova	0.5179	0.3795	0.4318	0.5179
CNN-FedProx	0.5050	0.3424	0.2934	0.5050
Hybrid-FedAvg	0.5515	0.4296	0.4938	0.5515
Hybrid-FedDyn	0.4995	0.3345	0.2536	0.4995
Hybrid-FedNova	0.5056	0.3455	0.4093	0.5056
Hybrid-FedProx	0.5340	0.3954	0.4157	0.5340
LSTM-FedAvg	0.9456	0.9425	0.9554	0.9456
LSTM-FedDyn	0.4794	0.3223	0.2474	0.4794
LSTM-FedNova	0.9471	0.9453	0.9530	0.9471
LSTM-FedProx	0.9553	0.9545	0.9609	0.9553

Table 2 changes the interpretation of the federated results materially. The strongest configuration is LSTM-FedProx (accuracy 0.9553, F1 0.9545), followed by LSTM-FedNova (0.9471, F1 0.9453) and LSTM-FedAvg (0.9456, F1 0.9425). In contrast, CNN and hybrid federated configurations remain near 0.50-0.55 accuracy, and LSTM-FedDyn falls to 0.4794. The results therefore support an architecture-dependent conclusion: high federated performance is concentrated in LSTM with FedAvg, FedNova, or FedProx, rather than being a general property of every FL-EC model.

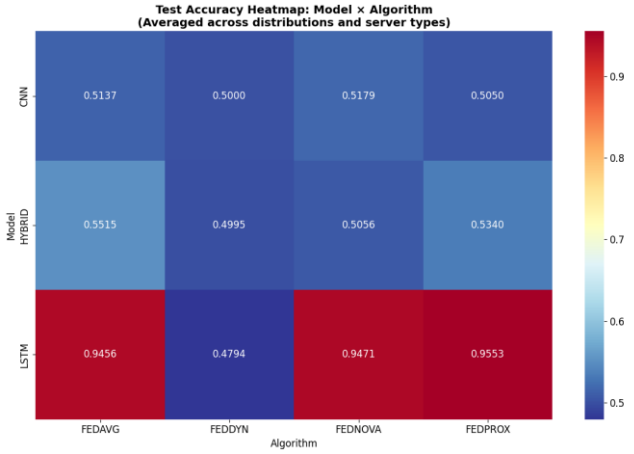


Figure 4. Federated test-accuracy heatmap by model and aggregation algorithm

Figure 4 visualizes the same model–algorithm means. The separation between LSTM and the CNN/hybrid models is pronounced, and FedDyn is the weakest algorithm for LSTM. LSTM-FedProx reaches 0.9553 test accuracy, whereas the strongest CNN and hybrid configurations remain at 0.5179 and 0.5515, respectively. This pattern demonstrates that federated performance depends strongly on the interaction between model architecture and aggregation algorithm.

Table 3. Pooled algorithm-level performance

Algorithm	Mean Acc.	Mean F1	Balance
FedAvg	0.6703	0.5769	0.6289
FedDyn	0.4929	0.3300	0.3915
FedNova	0.6569	0.5568	0.6172
FedProx	0.6648	0.5641	0.6126

Table 3 is a pooled algorithm-level summary, not a best-model table. FedAvg has the highest mean accuracy (0.6703) and balance score (0.6289), followed by FedProx and FedNova, whereas FedDyn is lowest at 0.4929 mean accuracy. The pooled means are substantially below the best LSTM values because they average over the weak CNN and hybrid runs shown in Table 2. The analysis also reports 95% accuracy confidence intervals of 0.649-0.691 for FedAvg, 0.489-0.497 for FedDyn, 0.636-0.678 for FedNova, and 0.644-0.686 for FedProx.

Security was evaluated as a pooled server-type effect across federated runs. This is another aggregation level and should not be compared directly with a single LSTM configuration. The pooled results show that secure operation produced only a small change in mean accuracy and recall while increasing mean precision.

Table 4. Secure versus standard federated server impact

Metric	Standard	Secure	Difference	Significant?
Accuracy	0.6231	0.6193	-0.0038	No
F1	0.5036	0.5103	+0.0066	No
Precision	0.4846	0.5170	+0.0323	Yes (p=0.044)
Recall	0.6231	0.6193	-0.0038	No

As shown in Table 4, secure aggregation reduced mean accuracy by only 0.0038 and changed F1 from 0.5036 to 0.5103. Precision increased from 0.4846 to 0.5170, with this difference statistically significant (p = 0.044), whereas the other metric differences were not significant. The supported interpretation is that the secure server setting did not materially degrade pooled predictive performance.

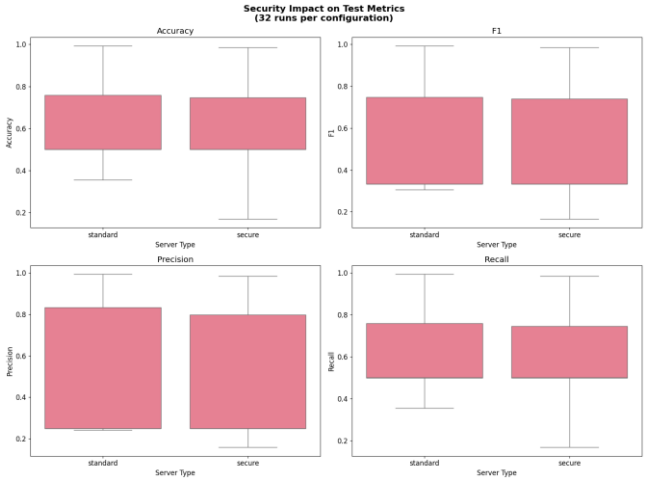


Figure 5. Security impact on accuracy, F1, precision, and recall across 32 runs per configuration

Figure 5 shows strongly overlapping standard and secure distributions for accuracy, F1, precision, and recall, consistent with the pooled statistics in Table 4. The modest precision shift is visible, but the overall distributions remain similar. This supports privacy-preserving aggregation with limited measured performance cost in the simulated experiments.

Real-time suitability is assessed using the efficiency analysis, which reports accuracy per minute for the leading model–algorithm configurations. This measure captures the predictive-performance-to-training-time trade-

off without treating efficiency as equivalent to classification accuracy or inference latency.

Table 5. Efficiency ranking based on accuracy per minute

Model	Algorithm	Score	Std. Dev.
LSTM	FedAvg	1.458	0.715
LSTM	FedNova	1.398	0.740
LSTM	FedDyn	1.394	0.513
LSTM	FedProx	0.794	0.344
CNN	FedDyn	0.257	0.088

Table 5 ranks LSTM-FedAvg highest on the reported efficiency measure (1.458), followed by LSTM-FedNova (1.398) and LSTM-FedDyn (1.394). This ranking is an efficiency measure and does not imply that FedDyn has the best predictive accuracy; Table 2 shows that LSTM-FedDyn accuracy is only 0.4794. Reporting the two measures separately is necessary because an efficiency ranking is not equivalent to a predictive-performance ranking.

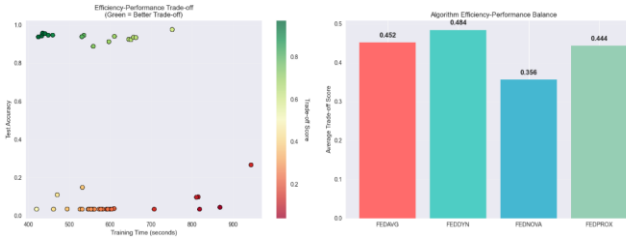


Figure 6. Efficiency-performance trade-off across federated configurations

Figure 6 further shows that computational efficiency and predictive accuracy are not interchangeable. High-accuracy configurations occupy a distinct region of the trade-off plot, while the algorithm-level trade-off score reflects the balance between runtime and performance. Consequently, the evidence supports configuration-aware efficiency assessment rather than a single universal latency advantage for FL-EC.

The NetworkFed/PySyft stage is reported separately from the 48-experiment simulation because it represents a different networked training condition and produces a distinct performance distribution. The networked analysis therefore focuses on the model–algorithm patterns observed in the corresponding experiment outputs rather than merging them with the simulated federated means.

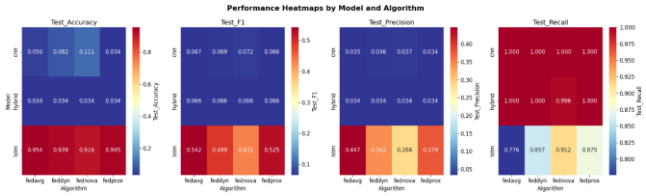


Figure 7. NetworkFed/PySyft performance heatmaps by model and aggregation algorithm

Figure 7 shows that, in the networked federated stage, LSTM test accuracy remains high across algorithms (approximately 0.916-0.954), whereas CNN and hybrid accuracies are much lower (approximately 0.034-0.111). The very high recall values for CNN and hybrid models occur alongside very low precision and accuracy, indicating a class-imbalance or prediction-bias effect rather than good overall classification. The networked evidence therefore reinforces the central result of the study: architecture and aggregation choice strongly condition federated predictive-maintenance performance.

Overall, the results do not support a blanket claim that FL-EC outperforms centralized learning for every model. Centralized CNN and LSTM are strong baselines, while federated LSTM with FedProx, FedNova, or FedAvg reaches comparable or slightly higher accuracy in the simulated experiments. Federated CNN and hybrid variants are substantially weaker, FedDyn is inconsistent across architectures, secure aggregation has little pooled performance cost, and the networked experiments remain configuration-sensitive. Taken together, these findings show that model-specific, pooled, and networked results must be interpreted within their respective evaluation settings.

Discussions, Limitations, and Future Directions

The concept of integrating Federated Learning (FL) and Edge Computing (EC) presents an interesting option to address urgent issues in the field of predictive maintenance in the Industrial Internet of Things (IIoT) related to privacy, latency, and scalability. FL can be decentralized to train models, where the data is kept at the edge devices and there is therefore a very minimal risk of data breaches. This will deal with the issue of privacy, which is key in the world of IIoT, where the data involved in operations can be very sensitive. Also, the EC component can process data locally, avoiding the transmission of data to centralized servers and greatly lowering the latency. This is essential in real time fault monitoring where delays will translate to

expensive downtime of equipment. Additionally, the FL-EC architecture is both extremely scalable and allows the processing of data in an efficient manner on a network of diverse devices that have different computing capacities. This scalability makes sure that as the number of edge devices increases the system will still work at an optimal level without affecting the accuracy or efficiency.

The FL-EC framework has various advantages over the traditional centralized systems when making comparisons. Centralized systems have the data of several devices transmitted to a central server to be processed and this presents a latency as well as privacy threat. Under FL and EC, model training is done directly on the devices and only model updates are transmitted thereby enhancing privacy maintenance (Haldikar et al., 2024). Besides, fault detection in FL-EC is also greatly improved because local processing of edge devices can provide faster decisions when compared to centralized systems where it could consume more time to upload and process data centrally (Liu et al., 2024). In general, FL and EC offer a more secure, fast, and scalable solution than traditional methods, which is why they are absolutely appropriate to IIoT-based predictive maintenance.

The findings are consistent with recent FL-EC research while adding an important qualification. Hafeez et al. (2021) focus on edge intelligence for data handling and predictive maintenance, Haldikar et al. (2024) examine FL and edge computing for real-time IIoT anomaly detection, Fan et al. (2022) concentrate on privacy-preserving aggregation, and Liu et al. (2024) emphasize edge resource selection and computation offloading. In the present experiments, centralized CNN and LSTM models remained strong, but federated performance varied sharply by architecture and optimizer. LSTM combined with FedProx, FedNova, or FedAvg produced the strongest federated results, while CNN and hybrid federated runs were markedly weaker. The contribution is therefore an integrated comparison that exposes the interaction between model architecture, aggregation, security setting, efficiency, and networked deployment rather than claiming universal FL-EC superiority.

The practical uses of FL-EC framework in the automobile, oil and gas, and manufacturing industries are enormous. An example of this is in the automotive industry, in which real-time predictive maintenance can be adopted to check the health of various vital parts of a vehicle, including engines and brakes, before they fail, and to improve the safety of the vehicle. Likewise, in the oil and gas sector, predictive maintenance can be used to monitor equipment like pumps, compressors, and turbines and determine their problems before the equipment becomes catastrophic (Fausing Olesen & Shaker, 2020). Companies that save lots of money in terms of repair and inefficiencies during operations, save much on the downtime and optimization of maintenance schedules. The fact that FL-EC framework has been able to offer real-time fault detection capability with great accuracy makes it a game changer in the industries where it is important to reduce operational interruptions. In addition, the privacy-saving aspect of FL takes care of the fact that sensitive industrial information is not revealed and that is crucial to those businesses that operate with proprietary technologies and regulatory requirements (Sarraf & Pal, 2022).

The evaluation has several limitations. First, the study uses a single public predictive-maintenance dataset and simulated or controlled federated clients, so generalization to different machines, sensing environments, and physical edge hardware remains unverified. Second, the large gap between LSTM and CNN/hybrid federated performance in Table 2 shows that results are sensitive to architecture, client data composition, and optimization conditions; repeated-seed and controlled IID/non-IID partition studies are therefore necessary. Third, edge devices have restricted memory, processing capacity, and network bandwidth, which may limit complex CNN-LSTM training. Model compression, lightweight architectures, adaptive client selection, and communication-efficient updates should therefore be evaluated together with accuracy. Finally, the efficiency analysis and NetworkFed results should be complemented by direct device-level measurements of processor time, memory use, transmitted bytes, energy consumption, and end-to-end latency on physical edge hardware.

Security, personalisation, and deployment feasibility. FL reduces routine raw-data centralisation but does not prevent poisoning, backdoor, Byzantine-update, model-inversion, or membership-inference attacks. Future evaluations should include authenticated clients, update clipping, anomaly screening, robust aggregation, secure aggregation, differential-privacy controls, and explicit adversarial testing. Per-device personalisation is also a useful extension: an edge client could fine-tune selected layers of the global model on its local operating profile, while retaining the shared model as a

common starting point. Such tuning may improve performance under heterogeneous machinery, but it should be controlled to avoid overfitting, unfair performance differences, and excessive energy or computational cost.

4. CONCLUSION

This study demonstrates the potential of an integrated Federated Learning and Edge Computing (FL-EC) framework for privacy-aware predictive maintenance in IIoT systems, while also showing that performance is strongly configuration-dependent. The centralized CNN and LSTM baselines achieved 0.9465 and 0.9430 accuracy, respectively. In the simulated federated experiments, LSTM-FedProx achieved the highest model-algorithm accuracy (0.9553), with LSTM-FedNova (0.9471) and LSTM-FedAvg (0.9456) close behind, whereas CNN and hybrid federated configurations remained near 0.50-0.55 accuracy. Pooled secure aggregation introduced only a small accuracy change (0.6231 to 0.6193) and increased precision from 0.4846 to 0.5170. The NetworkFed/PySyft results likewise favored LSTM and exposed weak CNN/hybrid behavior. The evidence therefore supports FL-EC as a promising privacy-preserving approach whose benefits depend on model architecture, aggregation algorithm, data partitioning, security configuration, and deployment conditions, rather than as an unconditional replacement for centralized learning.

Future work should reproduce all comparisons across multiple seeds and controlled IID/non-IID partitions, test the framework on real industrial data and physical edge devices, and measure energy, memory, communication payload, and end-to-end latency. Robust aggregation and adversarial evaluation are required to assess poisoning and inference risks. Resource-aware personalisation, model compression, and lightweight architectures should also be investigated to adapt the global model to individual machines without compromising scalability, fairness, or deployment feasibility.

5. REFERENCES

Ali, B., Gregory, M. A., & Li, S. (2021). Multi-access edge computing architecture, data security and privacy: A review. IEEE Access, 9, 18706-18721. https://doi.org/10.1109/ACCESS.2021.3053233

Alshkeili, H. M. H. A., Almheiri, S. J., & Khan, M. A. (2025). Privacy-Preserving Interpretability: An Explainable Federated Learning Model for Predictive Maintenance in Sustainable Manufacturing and Industry 4.0. AI, 6(6), 117. https://doi.org/10.3390/ai6060117

Beltrán, E. T. M., Pérez, M. Q., Sánchez, P. M. S., Bernal, S. L., Bovet, G., Pérez, M. G., Pérez, G. M., & Celdrán, A. H. (2023). Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and challenges. IEEE Communications Surveys & Tutorials, 25(4), 2983-3013. https://doi.org/10.1109/COMST.2023.3315746

Bhaskaran, P. E., Chennippan, M., & Subramaniam, T. (2020). Future prediction & estimation of faults occurrences in oil pipelines by using data clustering with time series forecasting. Journal of Loss Prevention in the Process Industries, 66, 104203. https://doi.org/10.1016/j.jlp.2020.104203

Cheng, H., Xie, Z., Shi, Y., & Xiong, N. (2019). Multi-step data prediction in wireless sensor networks based on one-dimensional CNN and bidirectional LSTM. IEEE Access, 7, 117883-117896. https://doi.org/10.1109/ACCESS.2019.2937098

Civerchia, F., Bocchino, S., Salvadori, C., Rossi, E., Maggiani, L., & Petracca, M. (2017). Industrial Internet of Things monitoring solution for advanced predictive maintenance applications. Journal of Industrial Information Integration, 7, 4-12. https://doi.org/10.1016/j.jii.2017.02.003

Fan, H., Huang, C., & Liu, Y. (2022). Federated learning-based privacy-preserving data aggregation scheme for IIoT. IEEE Access, 11, 6700-6707. https://doi.org/10.1109/ACCESS.2022.3226245

Fausing Olesen, J., & Shaker, H. R. (2020). Predictive maintenance for pump systems and thermal power plants: State-of-the-art review, trends and challenges. Sensors, 20(8), 2425. https://doi.org/10.3390/s20082425

Hafeez, T., Xu, L., & Mcardle, G. (2021). Edge intelligence for data handling and predictive maintenance in IIOT. IEEE Access, 9, 49355-49371. https://doi.org/10.1109/ACCESS.2021.3069137

Haldikar, S. V., Kader, O. F. M. A., & Yekollu, R. K. (2024). Edge computing and federated learning for real-time anomaly detection in industrial Internet of Things (IIoT). In 2024 International Conference on Inventive Computation Technologies (ICICT) (pp. 1699-1703). Lalitpur, Nepal: IEEE. https://doi.org/10.1109/ICICT60155.2024.10544912

Hongbin, F., & Zhi, Z. (2023). Privacy-preserving data aggregation scheme based on federated learning for IIoT. Mathematics, 11(1), 214. https://doi.org/10.3390/math11010214

Liu, X., Dong, X., Jia, N., & Zhao, W. (2024). Federated learning-oriented edge computing framework for the IIoT. Sensors, 24(13), 4182. https://doi.org/10.3390/s24134182

Matzka, S. (2020). AI4I 2020 Predictive Maintenance Dataset UC Irvine Machine Learning Repository. https://doi.org/10.24432/C5HS5C

Naidoo, P., & Sibanda, M. (2024). Emerging trends and future directions of the industrial Internet of things. From Internet of Things to Internet of Intelligence, 91-110. https://doi.org/10.1007/978-3-031-55718-7_5

Odeh, J. O., Yang, X., Nwakanma, C. I., & Dhelim, S. (2024). Asynchronous privacy-preservation federated learning method for mobile edge network in industrial internet of things ecosystem. Electronics, 13(9), 1610. https://doi.org/10.3390/electronics13091610

Sarraf, G., & Pal, V. (2022). Privacy-Preserving Data Processing in Cloud: From Homomorphic Encryption to Federated Analytics. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 8(2), 735-706. https://doi.org/10.32628/IJSRCSEIT

Scarpa, G., Gargiulo, M., Mazza, A., & Gaetano, R. (2018). A CNN-based fusion method for feature extraction from sentinel data. Remote Sensing, 10(2), 236. https://doi.org/10.3390/rs10020236

Usman, M., Ferlin, S., Brunstrom, A., & Taheri, J. (2022). A survey on observability of distributed edge & container-based microservices. IEEE Access, 10, 86904-86919. https://doi.org/10.1109/ACCESS.2022.3193102

Zhang, J., Zeng, Y., & Starly, B. (2021). Recurrent neural networks with long term temporal dependencies in machine tool wear diagnosis and prognosis. SN Applied Sciences, 3(4), 442. https://doi.org/10.1007/s42452-021-04427-5