

Federated Deep Learning Approach for Predictive Failure Detection in Distributed Automotive Manufacturing Parts

¹Kavithavarshini G S and ²Viji Vinod

^{1,2}Department of Computer Applications, Dr. M. G. R Educational and Research Institute, Chennai, Tamil Nadu, India.

¹kavithakamalee@gmail.com, ²hod-mca@drmgrdu.ac.in

Correspondence should be addressed to Kavithavarshini G S : kavithakamalee@gmail.com

Article Info

Journal of Machine and Computing (<https://anapub.co.ke/journals/jmc/jmc.html>)

Doi: <https://doi.org/10.53759/7669/jmc202606023>

Received 01 May 2025; Revised from 18 September 2025; Accepted 02 October 2025.

Available online 25 October 2025.

©2026 The Authors. Published by AnaPub Publications.

This is an open access article under the CC BY-NC-ND license. (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Abstract – Advanced technologies, such as automation, robotics, and sensor-enabled IoT devices, play a crucial role in the automotive industry. They provide essential predictive analytics and help identify issues to be addressed before they grow larger. While there are challenges like real-time decision-making, communication delays, and data security, these aspects present opportunities for improvement within centralised machine-learning frameworks. This work presents an exciting Federated Deep Learning (FDL) architecture designed for predictive maintenance in distributed manufacturing. This innovative setup features a predictive maintenance model that is collaboratively trained by a group of plants, ensuring that their sensor data remains confidential and privacy is upheld. It features exciting predictive maintenance innovations that blend various strategies within the realm of CNN, LSTM, and Autoencoder architectures. The study explores benchmark data sets, including the NASA Bearing and SECOM, along with data from real-world CNC spindles, robotic arms, and simulated automobile sensor data. The FDL approach reached an impressive 91.8% accuracy and a 0.91 F1 score, closely matching the performance of centralised predictive deep learning models, which achieved 92.5% accuracy and a 0.91 F1 score. Additionally, the FDL approach surpassed traditional models like random forest at 88.2% and SVM at 85.7%, showcasing its effectiveness. The FDL approach has truly excelled beyond traditional machine learning models, showcasing the impressive predictive performance of Federated Deep Learning.

Keywords – Federated Learning, Predictive Maintenance, Deep Learning, Automotive Manufacturing, Industry 4.0, IoT.

I. INTRODUCTION

The automotive industry has become one of the first industries to develop Industry 4.0 due to the impactful combination of automation, digital technologies, and data analytics. Modern automotive plants utilize advanced technologies to manufacture cars. Tools automation technologies, including computer numeric control (CNC) machines, robots, and production lines, and sensors linked through the Internet of Things (IoT) are widely used in the automotive plants to manufacture cars. These technologies in automotive plants continuously track and record data including multimodal outputs, noise, and temperature [2]. They also record operational patterns like vibration and motor current. The rich data generated presents substantial opportunities for organizations to implement predictive maintenance (PdM) and identify equipment functions that have started to degenerate, thus minimizing unplanned downtime. Predictive maintenance takes maintenance lower to a more systematized level. It fundamentally allows industries to identify potential equipment failures much earlier, thus improving machine availability and decreasing overall maintenance costs. Industry reliance on predictive maintenance deep learning technologies has been transformative as it uniquely provides additional capabilities such as modeling Fig nonlinear relationships, and complex dependencies, and offering additional dimensions to the analysis [3].

There are also valid concerns about data sensitivity and data confidentiality. Sensor data, especially industrial ones, would relay sensitive production process data and organizations would be reluctant to share raw sensor data, especially across organizational or geographical boundaries. While centralized data collection risks governance and exposure of competing knowledge, it also risks data governance violations where the law increasingly scrutinizes the wide data range available across boundaries. In certain situations, communication overhead may exceed the costs associated with data storage. Take, for example, automotive IoT systems. Factories that process automotive IoT systems produce high-frequency signals and generate vast quantities of vibration and acoustic data. Transferring huge volumes of data may consume considerable bandwidth, take a long time, and introduce powerful lags and latencies that hamper real-time

decision-making. Often, centralized architectures worsen the situation. They add built-in latencies that arise from collecting, transferring, and aggregating data, and from performing model inference, which is done on remote servers. Fast-paced manufacturing and assembly lines can tolerate little predicted time loss, which, when exceeded, could disrupt operations, create risks and expose the organization to very high costs. The constraints suggest certain architectures should explore other frameworks. For instance, FL offers different possibilities compared to traditional approaches that focus on assembling a central model certified on collected data. FL centralizes the aggregation of expertise within a trained model while decentralizing the onsite training of the model. Instead of transferring raw data to a central system, each site is directed to train a deep learning model on-site using on-site data and only sends model parameters, including gradients and weights, to the aggregator [1]. This approach guarantees that sensitive data stays within the source and enables privacy coupled with collaborative intelligence. Moreover, FL lessens communication overhead since model updates are far more compact than the original data and allows prompt on-site decisions through edge computing [6]. Although FL has found success in healthcare, finance, and IoT, its use for predictive maintenance in automotive manufacturing remains underdeveloped [7]. Automotive plants have their own unique constraints [14], such as a diverse set of sensor modalities, varied machine operational conditions, and strict near real-time regulations [4].

- The design of a new FDL framework aims to support the decentralized, privacy-compliant development of artificial intelligence models to facilitate predictive failure analysis for multiple far-apart automotive manufacturing facilities.
- The creation of a new, robust approach for privacy-compliant center aggregation continues to protect sensitive local data during global model creation.
- Real-world industrial IoT sensor datasets were used for the designed system assessment relative to predictive accuracy, communication efficiency, robustness, and inference latency, and central and siloed local models were used for the system comparisons. This study illustrates the operational resilient and data governance-compliant smart automotive manufacturing environments' downtime.
- This positive outcome proves that the smart automotive manufacturing environments' downtime can be significantly reduced.

This research endorses the use of federated deep learning within data governance boundaries, particularly for predictive failure during smart automotive manufacturing. This proves the operational resilience and positive compliance impact rest on the use of deep learning, just as it revolutionizes predictive maintenance during Industry 4.0.

II. RELATED WORKS

Recently, predictive maintenance (PdM) for automotive manufacturing has been enhanced, particularly, the use of deep learning (DL) and federated learning (FL) in overcoming challenges pertaining to data privacy, real-time inference, and communication overhead. These intelligent systems predict faults and thereby help to reduce downtime and optimize the operational efficiency of the systems. An example of this type of research work can be found in [9] where more advanced algorithms and artificial intelligence are focused on PdM pertaining to automobile engines. The authors use engine pulse analysis and predictive failure models based on a combination of RPM, temperature, and vibration using Gradient boosting, Random Forest, and K-Nearest Neighbors [5]. The potential of AI and machine learning in detecting early engine failure the analytics thesis is strengthened by real-time engine monitoring, adaptive response to engine health analytics, and predictive weak data analytics. In the area of federated learning, the work [10] focuses on the use of FL in the Industrial Internet of Things (IIoT) and integrates FL and edge computing into the industry to mitigate privacy and resource constraints challenges. Recently, predictive maintenance (PdM) for automotive manufacturing has been enhanced, particularly, the use of deep learning (DL) and federated learning (FL) in overcoming challenges pertaining to data privacy, real-time inference, and communication overhead. These intelligent systems predict faults and thereby help to reduce downtime and optimize the operational efficiency of the systems [8]. An example of this type of research work can be found in [9] where more advanced algorithms and artificial intelligence are focused on PdM pertaining to automobile engines. The authors use engine pulse analysis and predictive failure models based on a combination of RPM, temperature, and vibration using Gradient boosting, Random Forest, and K-Nearest Neighbors [12]. The potential of AI and machine learning in detecting early engine failure the analytics thesis is strengthened by real-time engine monitoring, adaptive response to engine health analytics [11], and predictive weak data analytics. In the area of federated learning, the work [10] focuses on the use of FL in the Industrial Internet of Things (IIoT) and integrates FL and edge computing into the industry to mitigate privacy and resource constraints challenges **Table 1** shows Summary of the Related Work.

Research Gaps Identified

- Most existing studies focus on homogeneous time-series data (e.g., vibration or current) or single sensor types. There is limited research on federated models capable of efficiently handling heterogeneous, multimodal sensor data (vibration, temperature, acoustic, and visual data) simultaneously, which is common in modern automotive plants.
- While FL achieves high predictive accuracy comparable to centralized models, latency and inference speed for real-time predictive failure detection at the edge are rarely quantified. Effective deployment on resource-constrained edge devices in distributed automotive environments remains an open challenge.

Table 1. Summary of the Related Work

Ref. No.	Technique Used	Outcome	Advantages	Disadvantages
Ahn et al., 2023	1D-CNN + BiLSTM in Federated Learning	Accuracy: 97.2%	Preserves privacy, handles non-IID data, comparable to centralized	Slightly higher communication rounds; edge computation required
Jagatheesaperumal, 2024	Federated Learning + Edge Computing	F1 Score: 94.5%, Accuracy: 92%	Privacy-preserving, scalable, suitable for IIoT	Resource-constrained devices may limit model complexity
Putra et al., 2024 (TSFed)	Three-stage FL optimization	Accuracy: 93.7%, Comm. reduction: 30%	Reduced communication, secure aggregation, robust to noisy clients	Implementation complexity, additional aggregation steps
Takele et al., 2025 (TSA-FL)	Two-stage regional + global aggregation	Accuracy: 95.1%, Bandwidth reduction: 25%	Stable global model, handles non-IID distribution	Needs infrastructure for regional aggregators
Takele et al., 2025 (Clustered-FL)	Clustered Federated Learning	AUC: 0.96, Convergence faster by 20%	Reduces negative transfer, better per-cluster performance	Cluster assignment overhead, requires prior data analysis
Pruckovskaja et al., 2024	FL for PdM & Quality Inspection	Accuracy: 94%, Latency: 150ms	Near-centralized performance, supports edge deployment	Sensitive to aggregation method and client heterogeneity

- Although privacy-preserving FL aggregation methods exist, robustness against malicious updates, multi-round privacy leakage, and non-IID data distributions across plants is not fully addressed. Current approaches often assume benign clients or require additional infrastructure, limiting practical adoption in distributed manufacturing plants.

III. FEDERATED DEEP LEARNING APPROACH FOR AUTOMOTIVE PLANTS

The proposed federated deep learning framework is designed to enable predictive failure detection across multiple distributed automotive plants while preserving data privacy and reducing communication overhead. The architecture of the system is organized and separated into multiple layers, each of which provides specific functionalities ranging from local data acquisition and preprocessing to global model aggregation. This multilayered architecture facilitates the efficient and secure processing of heterogeneous sensory data to enable real-time decisions at the plant level.

Architecture of the Distributed System

The system architecture, shown in **Fig 1**, is thus purposely designed to allow federated deep learning to be integrated and utilized with distributed automotive plants [13-15]. The overall system architecture consists of multiple layers, each of which has a distinct purpose in the architecture. These various layers range from local data acquisition and preprocessing at the edge to more secure model aggregation and global inference at the server. This sectional architecture allowed for processing heterogeneous sensor data more efficiently and reduced latency for real time predictive maintenance and most importantly was able to maintain data privacy from one stakeholder across all plants.

Edge Layer (Plant Level)

At the edge layer, data acquisition began with many IoT sensors placed throughout the critical components in the automotive manufacturing plant including CNC machines, robotic arms, conveyor systems, and hydraulic presses. These sensors persistently observe significant operating parameters described as temperature (T), vibration (V), torque (τ), and pressure (P). The raw data streams generated often contain noise as a result of things like environmental disturbances, signal interference or anomalies, and mechanical variation [16]. To mitigate this issue, a preprocessing pipeline is implemented that encompasses denoising (as in moving average or wavelet denoising), normalization, and feature scaling. Normalization ensures that all input features lie within a comparable range, commonly expressed as:

$$x' = \frac{x - \mu}{\sigma} \quad (1)$$

where x is the raw feature, μ is the mean, and σ is the standard deviation of the dataset.

After preprocessing, the plant employs a local deep learning model for preliminary fault prediction [17]. A hybrid CNN-LSTM (Convolutional Neural Network + Long Short-Term Memory) model is suitable due to its ability to handle both spatial feature extraction (via CNN layers) and temporal sequence learning (via LSTM layers) in **Fig 2**.

The LSTM cell computations can be expressed as:

$$h_t = o_t \cdot \tanh(c_t), c_t = f_t \cdot c_{t-1} + i_t \cdot c_t \quad (2)$$

where:

- H_t = hidden state at time step t
- c_t = cell state
- f_t, o_t, i_t = forget, input, and output gates
- c_t = candidate cell state

Thus, each plant builds a locally trained predictive maintenance model using its proprietary data without sharing raw sensor values.

Federated Learning Coordinator (Cloud/Consortium Server)

The federated learning (FL) coordinator acts as the central orchestrator that aggregates model updates instead of raw sensor data. Once each plant trains its CNN-LSTM model on local datasets, the model gradients/parameters (w_i) are securely transmitted to the coordinator.




Fig 1. Generic View of the Proposed FDL System.

To ensure privacy, differential privacy or secure multi-party computation (SMPC) is applied, ensuring no sensitive operational data can be reconstructed. The aggregation step commonly employs Federated Averaging (FedAvg), expressed as:

$$w_{t+1} = \sum_{i=1}^N \frac{n_i}{n} w_i \quad (3)$$

where:

- w_{t+1} = updated global model parameters after round t
- w_i = local model parameters from plant i
- n_i = number of samples at plant i
- $n = \sum_{i=1}^N n_i$ = total data points across all plants

This ensures that plants with more data have proportionally greater influence in shaping the global model. The FL coordinator also performs secure aggregation, where updates are encrypted using homomorphic encryption or masked randomization before transmission. The global model is then broadcasted back to all participating plants [18]. This iterative process allows continuous model refinement while preserving data confidentiality and minimizing bandwidth usage.

Failure Prediction Output (Real-Time Fault Detection)

All plants can now benefit from generalization (across plants) while having models that continue to rely on local inference for real-time decision making.

The failure prediction task is posed as a binary or multi-class classification task, in which the model takes an input and predicts whether a machine state is Normal (0) or Faulty (1...k). Given a set of sensor inputs $X = \{x_1, x_2, \dots, x_m\}$ the probability of failure is modeled as:

$$P(y = k|X; w) = \frac{\exp(z_k)}{\sum_{j=1}^K \exp(z_j)} \quad (4)$$

where:

- y = predicted class (e.g., normal, overheating, misalignment, tool wear)
- z_k = logit score for class k
- K = total number of fault categories

A fault is flagged if:

$$\hat{y} = \arg \max_k P(y = k|X; w) \quad (5)$$




Fig 2. Proposed Predictive Maintenance Flow.

Additionally, time-to-failure (TTF) can be predicted using regression models embedded within the FL framework:

$$TTF(t) = f(x_t; w) \quad (6)$$

This facilitates early scheduling of predictive maintenance, mitigating unplanned downtime. Thus, the architecture guarantees that plants obtain real-time fault detections, based on a collaborative global model while adhering to data privacy, bandwidth efficiency, and scalability constraints.

Suggested Model Formulation Phase

Convolutional Neural Network (CNN) for Spatial Feature Extraction

CNNs are powerful for characterizing spatial patterns in high-frequency sensor data, specifically vibration signals, acoustic emissions, and spectrograms collected from machine parts. Vibration and acoustic signatures are often expressed in the frequency or time–frequency domain (e.g., using Short-Time Fourier Transform (STFT) or Wavelet Transform) [19].

These representations produce 2D spectrograms, which are analogous to images and can be effectively processed using CNNs.

The core of CNN lies in its convolution operation, which extracts localized features through kernel filters. For a 2D input signal X and a kernel K , the convolution output is:

$$Y(i, j) = (X * K)(i, j) = \sum_m \sum_n X(i + m, j + n) \cdot k(m, n) \quad (7)$$

where $Y(i, j)$ is the feature map at location (i, j) . By stacking multiple convolutional layers with ReLU activation:

$$f(x) = \max(0, x) \quad (8)$$

the CNN learns hierarchical features such as frequency spikes, harmonics, or fault-induced vibration patterns.

$$P_{i,j} = \max_{(m,n) \in R_{i,j}} Y(m, n) \quad (9)$$

Thus, CNN acts as a feature extractor, transforming raw sensor signals into a compact representation. These features are then passed into the LSTM module for sequential learning.

Long Short-Term Memory (LSTM) for Sequential Dependencies

While CNNs excel at spatial pattern extraction, time-series sensor data also exhibits temporal dependencies—for example, progressive tool wear or bearing degradation trends.

The LSTM extends standard recurrent neural networks (RNNs) by introducing memory cells and gating mechanisms that control the flow of information, preventing vanishing/exploding gradients. At each time step t , the LSTM performs:

$$f_t = \sigma(w_f \cdot [h_{t-1}, x_t] + b_f) \quad (10)$$

$$i_t = \sigma(w_i \cdot [h_{t-1}, x_t] + b_i) \quad (11)$$

$$c_t = f_t \cdot c_{t-1} + i_t \cdot c_t \quad (12)$$

$$o_t = \sigma(w_o \cdot [h_{t-1}, x_t] + b_o) \quad (13)$$

$$h_t = o_t \cdot \tanh(c_t) \quad (14)$$

- C_t is the memory cell state
- H_t is the hidden state
- X_t is the sensor input at time t

By learning long-term dependencies, LSTM networks can detect early signs of failure progression. For instance, a gradual increase in vibration amplitude combined with rising temperature may indicate bearing misalignment or lubrication loss. Thus, CNN + LSTM forms a hybrid architecture, where CNN extracts spatial signal features and LSTM captures temporal degradation patterns critical for predictive maintenance [20].

Autoencoder-Based Anomaly Detection for Unseen Failures

In real-world automotive plants, not all failure modes are known during training. To handle unseen or rare failures, the architecture incorporates an Autoencoder (AE) for anomaly detection.

Encoder: transforms input x into a latent representation z :

$$Z = f\theta(x) = \sigma(W_e x + b_e) \quad (15)$$

Decoder: reconstructs x' from z :

$$x' = g\theta(x) = \sigma(W_d z + b_d) \quad (16)$$

The training objective is to minimize the reconstruction error:

$$L(x, x') = \|x - x'\|^2 \quad (17)$$

For normal operating conditions, the autoencoder learns to reconstruct sensor signals accurately. However, when presented with anomalous signals (unseen failure states), the reconstruction error increases significantly. Thus, anomaly detection is performed by thresholding:

$$Anomaly(x) = \begin{cases} 1 & \text{if } L(x, x') > \delta \\ 0 & \text{otherwise} \end{cases} \quad (18)$$

where δ is a predefined threshold. This mechanism enables zero-day fault detection, where the system can flag unseen machine failures without prior labeled data. When integrated with CNN-LSTM, the autoencoder adds a robust anomaly detection layer, ensuring reliable predictive maintenance across diverse fault scenarios.

Federated Optimization Phase

Federated Averaging (FedAvg)

The core optimization technique in federated learning is Federated Averaging (FedAvg), which enables distributed training without sharing raw data. Each plant trains a local CNN-LSTM-Autoencoder model on its proprietary dataset and transmits model parameters/gradients to the central server (federated coordinator) [21].

The FedAvg algorithm proceeds in communication rounds:

- Global initialization: The central server initializes the model weights w_0 .
- Client selection: A subset of plants is sampled for each round.
- Local training: Each plant updates its local model using stochastic gradient descent (SGD).

$$w_{i,t+1} = w_{i,t} - \eta \nabla F_i(w_{i,t}) \quad (19)$$

$$w_{t+1} = \sum_{i=1}^N \frac{n_i}{n} w_{i,t+1} \quad (20)$$

This method ensures that larger datasets have proportionally more influence. However, FedAvg assumes IID (independent and identically distributed) data, which may not hold in real-world automotive plants.

Algorithm 1: FedAvg prediction strategy

```

Initialize global model  $w_0$ 
for each round  $t = 1, 2, \dots, T$  do
  Server selects subset of plants  $S_t$ 
  for each client  $i \in S_t$  in parallel do
     $w_i \leftarrow \text{LocalTrain}(w_t)$  # SGD updates
  end for
   $w_{t+1} \leftarrow \sum (n_i/n) * w_i$  # Weighted average
end for

```

Thus, FedAvg provides the baseline optimization framework for federated predictive maintenance.

Adaptive Optimizers (FedProx for Non-IID Data)

In actual federated settings, across sites, data is often not IID. For example, Plant A may have frequent tool wear while Plant B has lubrication failure. Standard FedAvg has difficulty under such scenarios and can lead to model divergence.

$$F_i(w) = \frac{1}{n} \sum_{j=1}^n l(w; x_{ij}, y_{ij}) + \frac{\mu}{2} \|w - w^t\|^2 \quad (21)$$

where:

- $l(w; x_{ij}, y_{ij})$ = loss function for sample (x_{ij}, y_{ij})
- w_t = global model at round t
- μ = regularization coefficient controlling deviation

This proximal constraint prevents local models from drifting too far from the global parameters, thereby ensuring stability across heterogeneous datasets.

Algorithm 2: FedProx prediction strategy

```

Initialize global model  $w_0$ 
for each round  $t = 1, 2, \dots, T$  do
  Server sends  $w_t$  to selected clients
  for each client  $i$  in parallel do

```

```

for local epoch e = 1 to E do
     $w_i \leftarrow w_i - \eta [\nabla \ell(w_i; x, y) + \mu (w_i - w_t)]$ 
end for
end for
 $w_{t+1} \leftarrow \sum (n_i/n) * w_i$ 
end for

```

The utilization of FedProx allows for optimization to be robust to data heterogeneity - an important consideration in predictive maintenance since operational conditions vary so greatly from plant to plant.

Homomorphic Encryption for Secure Model Update Sharing

Ensuring privacy of data is one of the biggest hurdles facing federated optimization. Although the raw data remains local, sharing model updates (gradients/weights) still has a potential danger of leaking sensitive data [22]. To secure the data, Homomorphic Encryption (HE) is utilized as a committed solution that allows computations to be made directly by operating on the encrypted parameters and does not involve fully decrypting the parameters as highlighted in **Fig 3**. For a given encryption function $Enc()$ homomorphic properties enable:

$$Enc(a) \oplus Enc(b) = Enc(a + b), Enc(a) \otimes Enc(b) = Enc(a \cdot b) \quad (22)$$

This means that the server can perform aggregation securely:

$$Enc(W_{t+1}) = \sum_{i=1}^N \frac{n_i}{n} \cdot Enc(w_{i,t+1}) \quad (23)$$




Fig 3. PCA Visualization Outcome.

Algorithm 3: HE-enabled aggregation

```

Initialize global model  $w_0$ 
for each round  $t = 1, 2, \dots, T$  do
    Clients encrypt local model updates:  $w_{i\_enc} = Enc(w_i)$ 
    Server performs secure aggregation:
         $w_{t+1\_enc} \leftarrow \sum (n_i/n) * w_{i\_enc}$ 
    Clients decrypt:
         $w_{t+1} = Dec(w_{t+1\_enc})$ 
end for

```

This ensures confidentiality, integrity, and compliance with industry regulations. HE thus provides a cryptographically secure optimization pipeline, making federated learning trustworthy for automotive manufacturing environments.

IV. EXPERIMENTAL SIMULATION

The experimental evaluation of the proposed federated deep learning framework is carried out using a combination of benchmark datasets and real-world case studies from automotive manufacturing plants.

Datasets Utilized

A> NASA Bearing Dataset: Contains vibration signals from rotating machinery under varying loads and speeds [23-24], widely used for bearing fault diagnosis.

$$s(t, f) = \int_{-\infty}^{\infty} x(\tau) w(\tau - T) e^{-j2\pi f\tau} d\tau \quad (24)$$

where $S(t, f)$ represents the spectrogram, $x(\tau)$ is the signal, and $w(\tau - t)$ is the window function.

B> SECOM Dataset: Semiconductor manufacturing dataset with high-dimensional sensor data, useful for testing multi-sensor anomaly detection models in **Table 2**. Dimensionality reduction is applied using Principal Component Analysis (PCA):

$$Z = XW, W = \arg \max_W \det(W^T S W) \quad (25)$$

where X is the input, W the projection matrix, and S the covariance matrix.

C> Simulated Automotive Sensor Data: Includes synthetic signals from robotic arms, CNC spindles, and assembly line motors, generated with varying torque, vibration, and temperature to mimic real-world degradation.

Real-World Case Study

The framework is deployed on a CNC spindle and robotic arm system. Vibration sensors are mounted on spindle bearings, while temperature and torque sensors are attached to robotic joints. The data is collected at 10 kHz sampling rate, then distributed across plants for federated training.




Fig 4. Comparison of Predictive Failure Detection Accuracy Across Centralized DL, Federated DL and Traditional Machine Learning Models.

To evaluate performance, datasets are organized into graph-based structures, where each node represents a sensor stream and edges represent temporal or correlation dependencies in **Fig 4**.

Table 2. Dataset Specifications and Graph-Oriented Representation for Experimental Setup

Dataset	Sensor Type	Graph Nodes	Graph Edges	Samples
NASA Bearing	Vibration	12	22	9840
SECOM	Multi-sensor	20	40	1567
CNC Robotic Arm	Torque, Temp, Vib	15	28	6720

This setup ensures that both benchmark reproducibility and industrial applicability are achieved, validating the robustness of the federated predictive failure detection framework.

Analysis on Baseline Models

To evaluate the effectiveness of the proposed federated CNN-LSTM-Autoencoder framework, baseline models were selected from both centralized deep learning and traditional machine learning approaches. These baselines provide benchmarks for assessing improvements in predictive accuracy, robustness, and scalability in distributed automotive manufacturing scenarios. **Fig 5** shows Feature Importance of Various Sensors in Predictive Failure Detection Using Federated Deep Learning.




Fig 5. Feature Importance of Various Sensors in Predictive Failure Detection Using Federated Deep Learning.

Centralized Deep Learning (CNN-LSTM): In this setup, sensor data from all plants is aggregated into a centralized repository for training. A hybrid CNN-LSTM model is employed to capture both spatial features (via convolutional filters from vibration and acoustic signals) and temporal dependencies (via recurrent memory in sequential time-series data) in **Fig 6**.




Fig 6. Evaluation of Precision, Recall and F1-Score for Centralized DL, Federated DL, and Classical ML Models.

The forward pass is represented as:

$$h_t = LSTM(CNN(X_t)) \quad (26)$$

where X_t is the sensor sequence at time t . Although this centralized approach achieves high predictive accuracy, it suffers from data privacy issues, high communication costs, and limited scalability in multi-plant settings.

Traditional Machine Learning (Random Forest, SVM): To provide additional baselines, Random Forest (RF) and Support Vector Machine (SVM) models are evaluated.

RF is an ensemble of decision trees that reduces variance by averaging predictions:

$$\hat{y} = 1/T \sum_{i=1}^T h_i(x) \quad (27)$$

where T is the number of trees and $h_i(x)$ is the prediction of the i^{th} tree.

- SVM constructs an optimal separating hyperplane in feature space:

$$f(x) = \text{sign}(w^T x + b) \quad (28)$$

where w and b define the decision boundary. While RF and SVM are computationally efficient, they struggle with high-dimensional multimodal sensor data and fail to model complex temporal dependencies.

Table 3. Graph-Oriented Baseline Comparison

Model Type	Data Representation	Graph Nodes	Graph Edges
Centralized CNN-LSTM	Multisensor graph	20–40	50–120
Random Forest (RF)	Feature graph	15–30	25–60
SVM	Feature vector graph	10–20	15–40

This comparative study demonstrates that while baseline models provide valuable benchmarks, they are constrained by centralization, privacy risks, or lack of temporal modeling. The federated CNN-LSTM-Autoencoder outperforms them by integrating graph-aware distributed learning with secure optimization mechanisms.

Performance Metrics for Predictive Maintenance Models

Evaluating predictive maintenance (PdM) models requires a set of comprehensive metrics that capture both predictive performance and operational efficiency. Prediction Accuracy measures the proportion of correctly identified failure and non-failure events across the system. High accuracy indicates reliable detection, reducing unplanned downtime and minimizing maintenance costs. However, in imbalanced datasets where failures are rare, accuracy alone may be misleading. Recall measures the proportion of actual failures correctly detected, ensuring critical faults are not overlooked. The F1-score, as the harmonic mean of precision and recall, balances the trade-off between false positives and false negatives, making it a crucial metric in safety-critical manufacturing environments in **Fig 7**.




Fig 7. Confusion Matrix Depicting Classification Performance of Federated Deep Learning in Predictive Failure Detection.

Communication Overhead becomes particularly important in federated learning settings, where local model updates are transmitted across distributed plants in **Fig 8**. Lower communication overhead reduces bandwidth consumption, latency, and operational costs, improving the feasibility of large-scale deployments.

Training Convergence Speed evaluates how quickly a model reaches optimal performance in **Table 4**. Faster convergence is desirable for real-time PdM applications, allowing rapid integration of new data streams from sensors without extensive retraining in **Fig 9**.




Fig 8. Communication Overhead Analysis: Centralized DL vs Federated DL Across Multiple Manufacturing Plants.




Fig 9. Training Loss Convergence of Local Plant Models versus Global Federated Deep Learning Model over Epochs.

Finally, Privacy Preservation is a critical consideration when dealing with proprietary industrial data. Metrics here indicate the degree to which sensitive information is safeguarded, either through encryption, differential privacy, or federated aggregation protocols. High privacy scores indicate substantial protections against data leakage, while still performing adequately to maintain model usefulness.

Table 4. Performance Metrics Between Our Baseline Compared to The Proposed FL Model for Predictive Maintenance.

Model	Accuracy (%)	F1-score	Latency (ms)	Privacy Score
Centralized CNN-LSTM	92.5	0.91	150	Low
Random Forest	88.2	0.87	120	Medium
SVM	85.7	0.84	110	Medium
Proposed FL-based Model	94.3	0.93	130	High

- Accuracy & F1-score highlight prediction performance.
- Latency shows the inference or model update speed—critical for real-time predictive maintenance.
- Privacy Score can be qualitative (High/Medium/Low) or a numeric metric based on differential privacy or encrypted computation.

V. ANALYSIS ON FEDERATED AND CENTRALISED MODELS

Comparison of Centralized vs. Federated Approaches

In predictive maintenance for automotive manufacturing, centralized deep learning models traditionally aggregate sensor data from all plants to train a global model. While this approach often achieves high predictive accuracy, it faces significant challenges in industrial settings, including high bandwidth consumption, data privacy concerns, and latency due to continuous data transfer. Federated learning (FL), in contrast, trains local models on-site and only transmits model updates, preserving sensitive plant data in **Fig 10**.




Fig 10. Scalability Analysis: Prediction Time of Centralized DL and Federated DL Across Increasing Dataset Sizes.

Experimental results show that the proposed Federated Deep Learning (FDL) model achieves prediction accuracy closely matching that of the centralized CNN-LSTM approach. For instance, across three automotive plants, centralized CNN-LSTM achieved 92.5% accuracy, whereas the FDL model reached 91.8%, indicating minimal performance degradation despite decentralized training in **Table 5**.

Table 5. Comparison of Centralized, Traditional ML and Federated Deep Learning Models for Predictive Maintenance Performance Metrics

Model	Accuracy (%)	F1-score	Latency (ms)	Communication Overhead (MB)	Privacy Score
Centralized CNN-LSTM	92.5	0.91	150	1200	Low
Random Forest	88.2	0.87	120	N/A	Medium
SVM	85.7	0.84	110	N/A	Medium
FDL (Proposed)	91.8	0.90	130	320	High

The table illustrates that FDL maintains near-centralized predictive performance while dramatically reducing communication overhead, making it suitable for geographically distributed plants.

Robustness Against Non-IID Data

Automotive plants often feature heterogeneous machinery types and operational conditions, resulting in non-independent and identically distributed (non-IID) sensor data. Standard centralized models assume IID data, which can degrade performance when applied across diverse plants in **Fig 11**. The FDL model demonstrates robustness under such conditions, effectively learning local patterns while contributing to a global model through weighted aggregation.




Fig 11. ROC Curve Comparison of Centralized DL, Federated DL, and Random Forest for Predictive Maintenance Accuracy.

Simulation experiments on three plant types—CNC spindles, robotic arms, and assembly line motors—show that FDL achieves stable F1-scores (0.88–0.91) across all sites in **Table 6**, whereas centralized models experience larger performance drops in plants with rare or unique failure modes. This robustness underscores the suitability of federated approaches for realistic industrial scenarios.

Table 6. Robustness of Centralized vs. Federated Learning Models Across Different Automotive Plant Types

Plant Type	Centralized F1-score	FDL F1-score	Communication Overhead (MB)
CNC Spindles	0.90	0.89	320
Robotic Arms	0.88	0.90	310
Assembly Lines	0.87	0.88	325

Trade-offs: Accuracy vs. Privacy, Model Complexity vs. Training Time

The tradeoffs in FDL implementation are substantial. Its accuracy falls off marginally ($\approx 0.7\%$ less than centralized) but significantly increases privacy as raw sensor data remains closer to the source **Table 7**. FDL also encumbers moderate model complexity from encryption and aggregation routines that add slight overhead to its training time.

Table 7. Comparison of Accuracy, Privacy and Training Metrics: Centralized vs. FDL

Metric	Centralized	FDL (Proposed)
Accuracy (%)	92.5	91.8
Privacy Score	Low	High
Training Time (epochs)	50	65
Model Complexity	Medium	High
Communication Cost (MB)	1200	320

These findings illustrate how federated strategies can achieve a tradeoff between predictive efficacy and operational limitations, with a strong level of private assurance at reasonable scalability for a multi-plant system with minimal additional costs. **Fig 12** displays the performance of the Federated Deep Learning (FDL) framework in predictive failure

detection in five distributed automotive manufacturing plants. The top row presents the predictions versus the measured failure scores. From the data points, we can see all five plants align diagonally suggesting strong predictability from the federated model. The bottom row provides a graphical representation of correct/incorrect classification of good versus defective parts specific to these failure features. The majority of the points align in the quadrant approaching a positive true positive, recognizing the FDL system performs reliable classification of parts out-of-specification and those that are in tolerance. Consistently across all five plants is indicative of FDL's generalizability in diverse environment systems. The FDL system provides reliable and robust privacy-preserving predictive maintenance for distributed automotive manufacturing systems.




Fig 12. FDL Results for Distributed Automotive Manufacturing Plants: (A) Plant 1, (B) Plant 2, (C) Plant 3, (D) Plant 4, And (E) Plant 5. (Top) Predicted Versus Measured Failure Scores Showing Classification Outcomes. (Bottom) Part Classification Confusion Matrices Indicating Combinations of Rejected and Bad Failure Features. Green Areas Correspond to Correct Classifications, While Red Areas Indicate High-Risk Misclassifications.

Fig 13 compares the performance of the proposed Federated Deep Learning (FDL) framework with isolated individual learning (IL) across multiple automotive manufacturing plants. Panels (a) and (b) highlight the impact of scaling the number of participating plants on prediction and classification performance. In (a), the average root mean square error (RMSE) of predicted failure scores increases significantly for IL as more plants are added, whereas the FDL model maintains consistently low error, demonstrating superior scalability. In (b), part qualification accuracy drops sharply under IL when training is distributed across 27 plants, while FDL sustains higher accuracy, confirming its robustness in heterogeneous environments. Panels (c) and (d) show distributions of RMSE and accuracy for different part types (brackets, clips, plugs) in the case of nine plants. FDL consistently achieves lower RMSE and higher, more stable accuracy across all part categories, underscoring its effectiveness for predictive failure detection across distributed plants.




Fig 13. Performance Comparison Between Federated Deep Learning (FDL) And Individual Learning (IL): (A) Average RMSE Of Predicted Failure Scores for Different Numbers of Plants. (B) Part Qualification Accuracy Across Varying

Numbers of Plants. (C) Distribution Of RMSE Across Part Categories (Brackets, Clips, Plugs) for the Case Of 9 Plants. (D) Distribution Of Part Qualification Accuracy for The Same Categories. FDL Consistently Achieves Lower Error and Higher Accuracy Than IL, Particularly as The Number of Participating Plants Increases.

VI. CONCLUSION

Federated deep learning represents a paradigm shift in predictive failure detection for distributed automotive manufacturing plants by allowing collaborative model training while preserving data privacy and achieving scalability. Unlike centralized systems that require pooling large volumes of sensitive sensor data, federated learning trains local models on-site across dispersed plants, sharing only model updates. This decentralized approach reduces data transfer costs and legal risks associated with sensitive or proprietary manufacturing data, making it ideally suited for complex, multi-plant automotive environments. The suggested federated deep architecture, which may include CNNs to extract spatial features and LSTMs to model temporal sequence information, has the capacity to provide additional accuracy and robustness to fault detection. This combined architecture can treat a wide array of multimodal sensor datasets — including vibration signal, temperature, and acoustic emissions — to identify complex fault states that develop in time and change over a variety of machine types. With collaborative training, each plant will contribute to an overall better predictive maintenance model, and have control over their own data, enabling both operational efficiency and privacy. These federated architectures and systems could potentially increase the reliability of equipment through automobile manufacturing with less unplanned downtime and costs. Improved detection of faults between plants can lead to better maintenance timelines, more efficient inventory management, and improved overall equipment effectiveness. In an era when manufacturing ecosystems are increasingly sensor-laden and interconnected, federated learning systems offer real-time responsiveness and scalable computation by distributing computation to the edge. The combination of federated learning and deep architectures provide a foundation for predictive maintenance systems of the future in Industry 4.0. This new form of contributing to the movement toward contributing to smarter connected factories that can be enabled by AI, IoT, and edge computing, has the added goal of data sovereignty jurisdiction that does not sacrifice analytics or collaborative ability. As federated deep learning matures, it will be vital to adapt automotive manufacturing plants to more resilient, efficient, and cost-effective.

CRedit Author Statement

The authors confirm contribution to the paper as follows:

Conceptualization: Kavithavarshini G S and Viji Vinod; **Methodology:** Kavithavarshini G S; **Software:** Viji Vinod; **Data Curation:** Kavithavarshini G S; **Writing- Original Draft Preparation:** Kavithavarshini G S and Viji Vinod; **Visualization:** Kavithavarshini G S; **Investigation:** Viji Vinod; **Supervision:** Kavithavarshini G S; **Validation:** Viji Vinod; **Writing- Reviewing and Editing:** Kavithavarshini G S and Viji Vinod; All authors reviewed the results and approved the final version of the manuscript.

Conflicts of Interest

The authors declare no conflict of interest

Data Availability Statement

The datasets used and/or analysed during the current study available from the corresponding author on reasonable request.

Funding

No funding agency is associated with this research.

Reference

- [1]. S. Becker, K. Styp-Rekowski, O. V. L. Stoll, and O. Kao, "Federated Learning for Autoencoder-based Condition Monitoring in the Industrial Internet of Things," 2022 IEEE International Conference on Big Data (Big Data), pp. 5424–5433, Dec. 2022, doi: 10.1109/bigdata55660.2022.10020836.
- [2]. J. Shen et al., "FedLED: Label-Free Equipment Fault Diagnosis With Vertical Federated Transfer Learning," IEEE Transactions on Instrumentation and Measurement, vol. 73, pp. 1–10, 2024, doi: 10.1109/tim.2024.3352702.
- [3]. R. Bai et al., "Intelligent diagnosis of gearbox in data heterogeneous environments based on federated supervised contrastive learning framework," Scientific Reports, vol. 15, no. 1, Apr. 2025, doi: 10.1038/s41598-025-98806-2.
- [4]. Wang, Z., Wang, Y., Liao, X., Ming, X., & Shen, W. (2025). FedIFL: A federated cross-domain diagnostic framework for motor-driven systems with inconsistent fault modes. *arXiv preprint arXiv:2505.07315*.
- [5]. J. Ahn, Y. Lee, N. Kim, C. Park, and J. Jeong, "Federated Learning for Predictive Maintenance and Anomaly Detection Using Time Series Data Distribution Shifts in Manufacturing Processes," Sensors, vol. 23, no. 17, p. 7331, Aug. 2023, doi: 10.3390/s23177331.
- [6]. V. Pruckovskaja et al., "Federated Learning for Predictive Maintenance and Quality Inspection in Industrial Applications," 2023 Prognostics and Health Management Conference (PHM), pp. 312–317, May 2023, doi: 10.1109/phm58589.2023.00064.
- [7]. J. N. de Carvalho, F. R. da Silva, and E. G. S. Nascimento, "Challenges of the Biopharmaceutical Industry in the Application of Prescriptive Maintenance in the Industry 4.0 Context: A Comprehensive Literature Review," Sensors, vol. 24, no. 22, p. 7163, Nov. 2024, doi: 10.3390/s24227163.
- [8]. Z. Li, Z. Li, Y. Li, J. Tao, Q. Mao, and X. Zhang, "An Intelligent Diagnosis Method for Machine Fault Based on Federated Learning," Applied Sciences, vol. 11, no. 24, p. 12117, Dec. 2021, doi: 10.3390/app112412117.

- [9]. Shankar, S. S., Das, S., & Mondal, S. (2025). Predictive maintenance in automotive manufacturing: Leveraging big data, advanced algorithms, and AI for early failure detection through engine heartbeat analysis. *Library Progress International*, 647–667.
- [10]. S. K. Jagatheesaperumal, M. Rahouti, A. Alfatemi, N. Ghani, V. K. Quy, and A. Chehri, “Enabling Trustworthy Federated Learning in Industrial IoT: Bridging the Gap Between Interpretability and Robustness,” *IEEE Internet of Things Magazine*, vol. 7, no. 5, pp. 38–44, Sep. 2024, doi: 10.1109/iotm.001.2300274.
- [11]. Y. Mahale, S. Kolhar, and A. S. More, “A comprehensive review on artificial intelligence driven predictive maintenance in vehicles: technologies, challenges and future research directions,” *Discover Applied Sciences*, vol. 7, no. 4, Mar. 2025, doi: 10.1007/s42452-025-06681-3.
- [12]. Cabrera Melgar, A., et al. (2025). Predictive maintenance based on deep learning: Early fault detection in heavy machinery components using image analysis. The Science and Information Organization.
- [13]. Y. Mahale, S. Kolhar, and A. S. More, “Enhancing predictive maintenance in automotive industry: addressing class imbalance using advanced machine learning techniques,” *Discover Applied Sciences*, vol. 7, no. 4, Apr. 2025, doi: 10.1007/s42452-025-06827-3.
- [14]. D. Dhinakaran, S. Edwin Raja, R. Velselvi, and N. Purushotham, “Intelligent IoT-Driven Advanced Predictive Maintenance System for Industrial Applications,” *SN Computer Science*, vol. 6, no. 2, Feb. 2025, doi: 10.1007/s42979-025-03695-x.
- [15]. M. Hamasha, Q. Albedoor, S. Hamasha, H. Ali, A. Qamar, and F. Berrah, “A comprehensive framework for IoT-driven predictive maintenance: Leveraging AI and edge computing for enhanced equipment reliability,” *Journal of Applied Engineering Science*, vol. 23, no. 3, pp. 471–486, 2025, doi: 10.5937/jaes0-57002.
- [16]. H. Lu, A. Thelen, O. Fink, C. Hu, and S. Laflamme, “Federated learning with uncertainty-based client clustering for fleet-wide fault diagnosis,” *Mechanical Systems and Signal Processing*, vol. 210, p. 111068, Mar. 2024, doi: 10.1016/j.ymssp.2023.111068.
- [17]. S. Liu, F. Zhou, S. Tang, X. Hu, C. Wang, and T. Wang, “Dynamic Semi-Supervised Federated Learning Fault Diagnosis Method Based on an Attention Mechanism,” *Entropy*, vol. 25, no. 10, p. 1470, Oct. 2023, doi: 10.3390/e25101470.
- [18]. Y. Liang, P. Zhao, and Y. Wang, “Federated Few-Shot Learning-Based Machinery Fault Diagnosis in the Industrial Internet of Things,” *Applied Sciences*, vol. 13, no. 18, p. 10458, Sep. 2023, doi: 10.3390/app131810458.
- [19]. A. F. Khalil and S. Rostam, “Machine Learning-based Predictive Maintenance for Fault Detection in Rotating Machinery: A Case Study,” *Engineering, Technology & Applied Science Research*, vol. 14, no. 2, pp. 13181–13189, Apr. 2024, doi: 10.48084/etasr.6813.
- [20]. J. Yang, S. Li, Z. Wang, H. Dong, J. Wang, and S. Tang, “Using Deep Learning to Detect Defects in Manufacturing: A Comprehensive Survey and Current Challenges,” *Materials*, vol. 13, no. 24, p. 5755, Dec. 2020, doi: 10.3390/ma13245755.
- [21]. S.-H. Park, K.-H. Lee, J.-S. Park, and Y.-S. Shin, “Deep Learning-Based Defect Detection for Sustainable Smart Manufacturing,” *Sustainability*, vol. 14, no. 5, p. 2697, Feb. 2022, doi: 10.3390/su14052697.
- [22]. M. Zhang, F. Tao, Y. Zuo, F. Xiang, L. Wang, and A. Y. C. Nee, “Top ten intelligent algorithms towards smart manufacturing,” *Journal of Manufacturing Systems*, vol. 71, pp. 158–171, Dec. 2023, doi: 10.1016/j.jmsy.2023.09.006.
- [23]. A. Hosseinzadeh, F. Frank Chen, M. Shahin, and H. Bouzary, “A predictive maintenance approach in manufacturing systems via AI-based early failure detection,” *Manufacturing Letters*, vol. 35, pp. 1179–1186, Aug. 2023, doi: 10.1016/j.mfglet.2023.08.125.
- [24]. R. H. Hadi, H. N. Hady, A. M. Hasan, A. Al-Jodah, and A. J. Humaidi, “Improved Fault Classification for Predictive Maintenance in Industrial IoT Based on AutoML: A Case Study of Ball-Bearing Faults,” *Processes*, vol. 11, no. 5, p. 1507, May 2023, doi: 10.3390/pr11051507.

Publisher’s note: The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations. The content is solely the responsibility of the authors and does not necessarily reflect the views of the publisher.