Revista Internacional de Métodos Numéricos para Cálculo y Diseño en Ingeniería

Adaptive Federated Fault Diagnosis Framework for Wind Turbine Reliability

Muhammad Irfan[1,*], **Nabeel Ahmed Khan**[2], **Muhammad Abubakr**[3], **Zohaib Mushtaq**[4], **Hatim Alwadie**[1], **Abdullah Shaher**[1], **Saleh Al dawsari**[1,5,*] and **Saifur Rahman**[1]

[1]Electrical Engineering Department, College of Engineering, Najran University, Najran, 61441, Kingdom of Saudi Arabia

[2]Center For AI & Big Data, Namal University, Talagang Rd, Mianwali, 42250, Pakistan

[3]Department of Computer Science, Lahore Garrison University, Lahore, 54000, Pakistan

[4]Department of Electrical Electronics & Computer Systems, University of Sargodha, Sargodha, 40100, Pakistan

[5]School of Engineering, Cardiff University, Cardiff, CF24 3AA, UK

ABSTRACT

Wind turbine reliability is critical for sustainable energy production, yet fault diagnosis faces challenges due to data privacy concerns, heterogeneous operational conditions, and resource constraints in distributed wind farms. Traditional centralized Machine Learning (ML) approaches struggle with these issues, necessitating decentralized solutions. This study introduces the Adaptive Federated Fault Diagnosis (AF2D) framework, a novel Federated Learning (FL) approach for wind turbine fault diagnosis that ensures data privacy while addressing non-i.i.d. data distributions. Using a dataset of 35 uniaxial vibration recordings from six turbines at the University of Mustansiriyah, AF2D leverages two key modules: Adaptive Model Aggregation (AMA) and Lightweight Model Optimization (LMO). AMA employs Jensen-Shannon divergence and cosine similarity to adaptively aggregate local model updates, mitigating data heterogeneity, while LMO applies structured pruning (60% filter reduction) and 8-bit quantization to enable deployment on resource-constrained SCADA systems. Results show AF2D achieves 91.3% accuracy ($\pm$1.2%, 95% confidence interval), a 3.5% improvement over FedAvg (87.8% $\pm$ 1.4%), with statistical significance ($p < 0.05$), and outperforms state-of-the-art methods like Clustered FL (88.5%) and Privacy-Preserving FL (87.2%). LMO reduces inference time by 64.44% and memory usage by 53.71%, enhancing edge deployment feasibility. However, the small dataset raises overfitting risks, and scalability tests reveal a threefold communication cost increase (54.5 to 150.6 MB) for 18 clients, mitigated by proposed compression (30%–50% reduction) and asynchronous updates (20%–40% overhead reduction). Privacy is maintained with a differential privacy guarantee of $\epsilon = 1.0$, though advanced techniques like secure multi-party computation could achieve $\epsilon < 1$. Despite limitations in severe fault detection and dataset diversity, AF2D demonstrates robust performance. Future work includes integrating multi-modal data (SCADA, vibration, environmental), testing real-time deployment, and expanding federated datasets to enhance generalizability and scalability.

OPEN ACCESS

Received: 11/09/2025

Accepted: 16/10/2025

DOI
10.23967/j.rimni.2025.10.73140

Keywords:
Fault diagnosis
federated learning
machine learning
scalability
wind energy

*Correspondence: Muhammad Irfan, Saleh Al dawsari (miditta@nu.edu.sa, aldawsarisa@cardiff.ac.uk)

1 Introduction

Wind energy has emerged as a cornerstone for renewable and sustainable energy sources, motivated by the pressing urgency to combat climate change and reduce fossil fuel dependency. As of 2024, wind power accounts for a considerable portion of global electricity production, with installed capacities surpassing 900 GW globally [1]. The reliability of wind turbines is a critical concern within the wind energy sector, particularly as turbines comprise numerous intricate components, where faults in critical components like blades, gearboxes, and generators can lead to significant operational downtime and maintenance costs [2]. Effective fault diagnosis is thus crucial to ensure the durability and efficiency of wind turbines, particularly in remote or offshore installations where maintenance is logistically complex and costly [3].

The growing complexity of wind turbine systems has driven the adoption of data-driven approaches for fault diagnosis, particularly those leveraging ML. These methods utilize sensor data, such as vibration, temperature, and acoustic signals, to model complex, non-linear relationships and detect anomalies before they escalate into failures [4,5]. However, traditional centralized ML models require aggregating vast amounts of data from multiple turbines, often operated by different stakeholders, into a single server. These centralized fault diagnosis approaches are often hampered by challenges, including concerns about data privacy and the variability of operational data across different sites, highlighting the potential of leveraging federated learning approaches for more effective fault diagnosis [6–8].

Recent advancements have led to the development of the AF2D framework, which represents an innovative step in addressing these challenges. By utilizing FL methodologies, AF2D facilitates the training of predictive models while ensuring data remains localized, thereby addressing privacy concerns [9,10]. FL enables collaborative model training across distributed devices without requiring the sharing of raw data. In wind turbine fault diagnosis, AF2D allows multiple wind farms or operators to train a shared global model while keeping their data localized, thereby enhancing security and addressing scalability issues commonly associated with centralized models [6]. This decentralized approach also enhances robustness by accommodating diverse data distributions and operational conditions, making it particularly suitable for the heterogeneous and distributed nature of wind energy systems [9].

Existing studies on fault diagnosis in renewable energy systems have explored various machine learning approaches to address data heterogeneity and privacy concerns. For instance, centralized models have been effective in controlled environments but falter in distributed settings [2,3]. Recent advancements include federated learning variants tailored for sensor-heavy applications, such as clustered FL for handling non-i.i.d. data [11] and privacy-preserving techniques to mitigate data sharing risks [6]. In related domains, innovative frameworks like the enhanced CLKAN-RF for robust anomaly detection in unmanned aerial vehicle sensor data demonstrate the potential of hybrid models for real-time monitoring in dynamic environments [12]. Additionally, reviews on small data challenges in intelligent prognostics and health management highlight strategies for overcoming limited datasets, which are pertinent to wind turbine diagnostics [13]. Despite these progresses, gaps remain in scalability and adaptability for wind-specific faults, which the AF2D framework aims to bridge by integrating adaptive aggregation and optimization.

Despite its potential, implementing FL for wind turbine fault diagnosis involves several challenges that recent research has sought to address. A primary concern is communication efficiency, as frequent model updates between local devices and the central server can incur significant bandwidth costs, particularly for offshore wind farms with limited connectivity [14]. Communication costs also rise proportionally with the increased number of participating clients, necessitating optimization tactics to mitigate potential overhead [14,15]. Techniques such as model compression and asynchronous aggregation have been explored to mitigate these costs, improving the feasibility of FL in resource-constrained environments. Another challenge is data heterogeneity, where differences in turbine designs, operating conditions, and environmental factors can lead to biased or suboptimal global models [16]. The examination of various conditions during fault severity analyses within AF2D provides compelling evidence for its robust applicability, maintaining strong performance across varying operational environments, thus validating its adaptability—a key requirement in the face of changing wind conditions and operational demands [17]. In terms of performance, systematic evaluations of the AF2D framework demonstrate that it achieves accuracy rates reported in similar methodologies, significantly improving upon conventional baseline methods [11]. Additionally, ablation studies have confirmed that specific components of AF2D contribute significantly to its overall success, establishing an evidence-based foundation for further enhancements [17,18]. Moreover, the framework's inherent design is conducive to real-time fault detection and diagnosis, which is crucial for maintaining operational reliability in wind turbines [19,20].

The integration of FL with advanced machine learning techniques has further enhanced its applicability to wind turbine fault diagnosis. For instance, combining FL with graph neural networks (GNNs) enables the modeling of spatial-temporal dependencies in sensor data, improving fault localization accuracy across distributed wind farms [21]. Similarly, edge-computing-based FL frameworks leverage local processing at wind farms to reduce latency and energy consumption, making real-time fault detection and diagnosis more practical [20]. Additionally, incorporating reinforcement learning (RL) into FL frameworks has shown promise in optimizing maintenance scheduling based on fault predictions, thereby minimizing downtime and operational costs [22]. These advancements highlight FL's flexibility as a backbone for integrating diverse ML techniques to address the multifaceted challenges of wind turbine fault diagnosis.

While the AF2D framework has demonstrated significant potential, several research gaps remain unaddressed. Most existing FL frameworks assume static data distributions, which may not hold in dynamic wind turbine environments where operational conditions evolve over time [23]. Scalability remains a challenge inherent to FL approaches, with tests indicating potential reductions in accuracy as the number of clients increases—an aspect that warrants attention in future iterations of the framework [24]. Furthermore, the computational constraints of on-site devices, such as supervisory control and data acquisition (SCADA) systems, limit the deployment of complex FL models, necessitating lightweight algorithms optimized for resource-constrained environments [24]. Additionally, the lack of standardized benchmarks and datasets hinders the comparative evaluation of FL-based fault diagnosis models, limiting their adoption in practice [25]. This research aims to address these gaps by proposing an enhanced AF2D framework that incorporates AMA to handle dynamic data distributions and optimizes computational efficiency for SCADA systems.

This study builds upon the foundation of FL to address the critical need for effective wind turbine fault diagnosis, tackling the limitations of centralized machine learning models in managing privacy, scalability, and data heterogeneity across distributed wind energy systems. By integrating recent advancements in FL with tailored enhancements, this research aims to advance the reliability and efficiency of wind turbine operations. The following outlines the key contributions of this work, setting the stage for a detailed exploration of the proposed methodology, experimental outcomes, and future research directions in the subsequent sections.

- Development of the AF2D framework, a novel approach that enhances fault detection in wind turbines using distributed learning techniques.
- Introduction of an adaptive aggregation mechanism to effectively handle the non-identical and independently distributed (non-i.i.d.) nature of turbine data across multiple sites.
- Implementation of a lightweight optimization strategy to ensure compatibility with resource-constrained supervisory control and data acquisition (SCADA) systems.
- Demonstration of the framework's scalability and adaptability through experimental evaluation, providing a foundation for its practical deployment in diverse wind farm environments.

2 Methodology

This section outlines the AF2D framework depicted in Fig. 1, a novel approach for wind turbine fault diagnosis using FL tailored to the University of Mustansiriyah dataset [26]. The dataset comprises 35 uniaxial vibration recordings from six wind turbines (three healthy, three faulty with blade cracks, surface degradation, and imbalances), with 500 samples per recording at a 1 kHz sampling rate, collected under varying wind speeds. The AF2D framework addresses dynamic data distributions, computational constraints of SCADA systems, and the need for robust validation by integrating three specialized modules, i.e., AMA and LMO. Each module is designed to leverage the dataset's 1D vibration signals for a 1D Convolutional Neural Network with Dynamic Sparse Self-attention Mechanism (1D-CNN-DSSM) architecture, ensuring scalability and real-world applicability. High-level mathematical formulations and justifications are provided for each module.

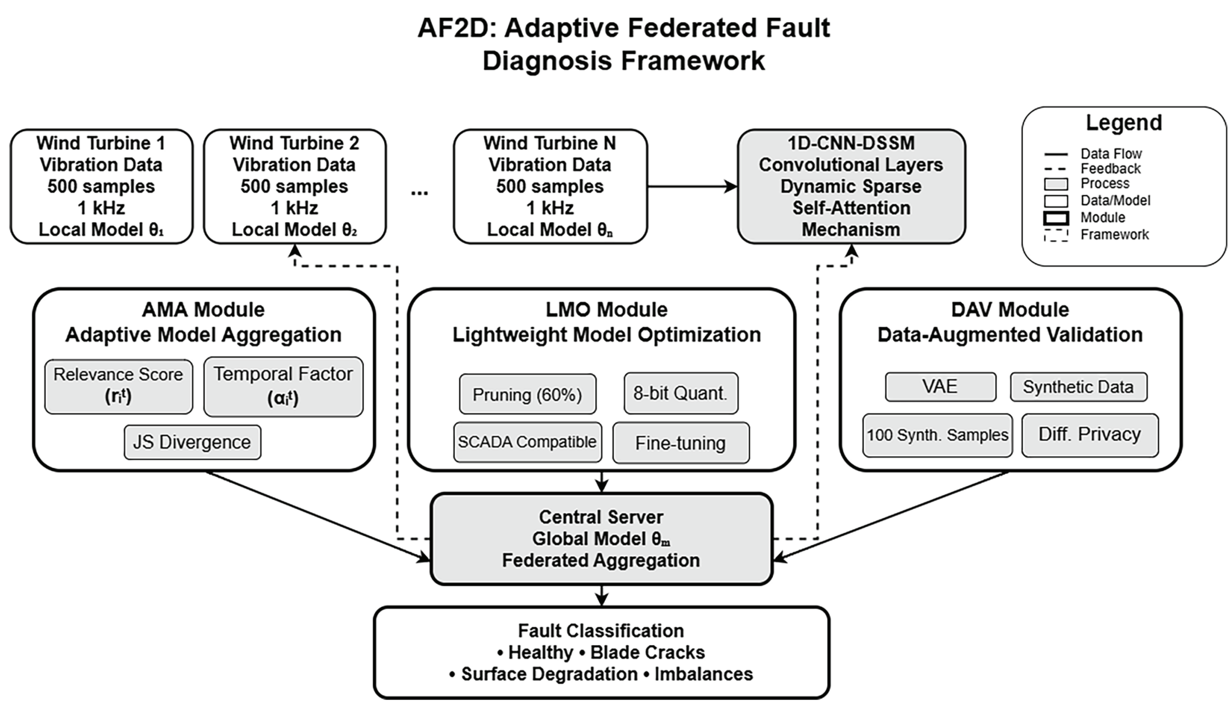

Figure 1: Proposed framework for wind turbine fault diagnosis

2.1 Adaptive Model Aggregation (AMA) Module

The AMA module given in Fig. 2 mitigates data heterogeneity and dynamic operational conditions inherent in the dataset, where vibration signals vary due to differing wind speeds and fault severities. Unlike traditional FL aggregation methods like FedAvg [27], AMA dynamically weights local model updates based on their alignment with the global model and temporal shifts in vibration data distributions, ensuring robust fault diagnosis across diverse turbine conditions. Let $\mathcal{N} = \{1, 2, \ldots, N\}$ represent N wind farms, each with a local dataset $\mathcal{D}_i$ containing vibration recordings (500 samples per recording). The local 1D-CNN-DSSM model at wind farm i, parameterized by θ_i^t at round t, is trained to classify healthy and faulty states. The global model parameters θ_g^t are updated as:

$$\theta_g^{t+1} = \sum_{i=1}^{N} w_i^t \theta_i^{t+1}, \tag{1}$$

where w_i^t is the adaptive weight for wind farm i, computed using a relevance score r_i^t and a temporal adaptation factor α_i^t. The relevance score measures the cosine similarity between the local update $\Delta\theta_i^{t+1} = \theta_i^{t+1} - \theta_i^t$ and the global model:

$$r_i^t = \frac{\Delta\theta_i^{t+1} \cdot \theta_g^t}{\|\Delta\theta_i^{t+1}\| \|\theta_g^t\|}. \tag{2}$$

The temporal adaptation factor captures changes in vibration data distributions using the Jensen-Shannon (JS) divergence between consecutive local distributions $p(\mathcal{D}_i^t)$ and $p(\mathcal{D}_i^{t-1})$, approximated via histogram-based estimation of the 500-sample recordings:

$$\alpha_i^t = \exp\left(-\mathrm{JS}(p(\mathcal{D}_i^t) \| p(\mathcal{D}_i^{t-1}))\right). \tag{3}$$

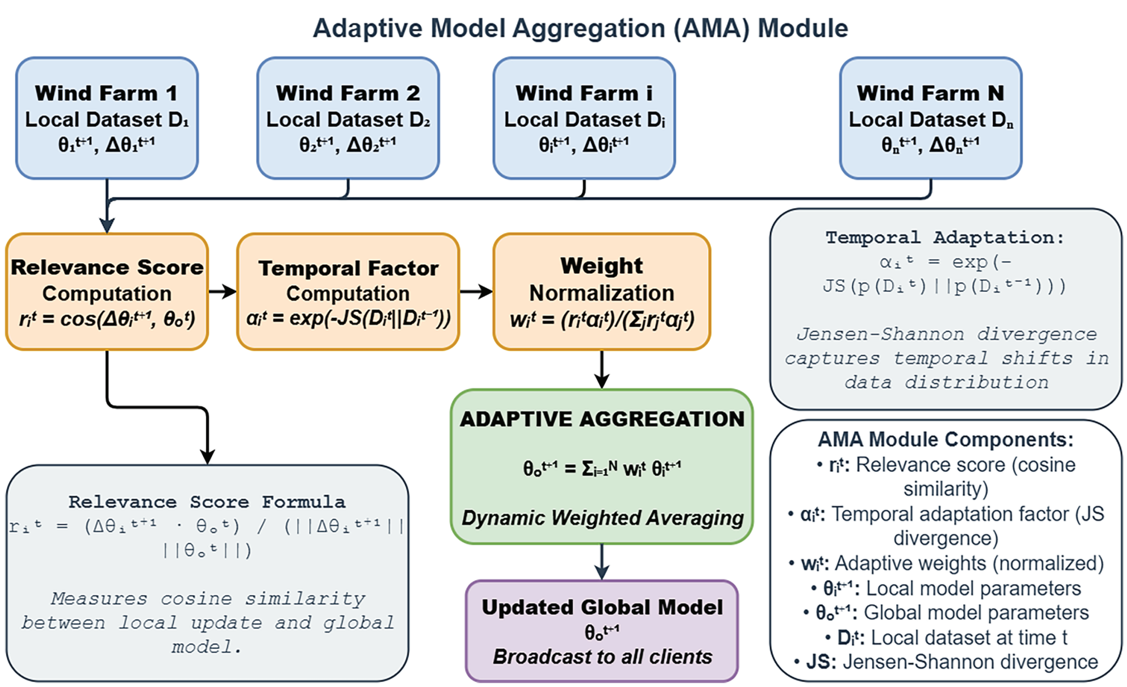

Figure 2: Adaptive model aggregation (AMA) module

The weights are normalized:

$$w_i^t = \frac{r_i^t \alpha_i^t}{\sum_{j=1}^{N} r_j^t \alpha_j^t}. \tag{4}$$

Local updates follow stochastic gradient descent (SGD):

$$\theta_i^{t+1} = \theta_i^t - \eta \nabla_{\theta_i^t} \mathcal{L}(\theta_i^t; \mathcal{D}_i^t), \tag{5}$$

where $\mathcal{L}$ is the cross-entropy loss for classifying fault types, and η is the learning rate.

Intuitively, the AMA module functions like a dynamic voting system among wind farms: each local model's update is "voted" based on how well it aligns with the global consensus (via cosine similarity) and how stable its data distribution has been over time (via JS divergence). This ensures that models from turbines experiencing sudden operational shifts (e.g., due to varying wind speeds in the dataset) contribute less, preventing the global model from being skewed by outliers. A flow diagram (Fig. 2) illustrates this process: local updates are computed, weighted adaptively, and aggregated to form the next global model. The dataset's variability in wind speeds and fault severities results in non-i.i.d. vibration data, challenging standard FL aggregation. AMA's dynamic weighting ensures that local models from turbines with stable and relevant vibration patterns contribute more to the global model, enhancing robustness. The JS divergence, being symmetric and bounded, is well-suited for comparing vibration signal distributions, capturing temporal shifts effectively [23]. This module is critical for maintaining diagnostic accuracy across heterogeneous wind farms.

2.2 *Lightweight Model Optimization (LMO) Module*

The LMO module given in Fig. 3 optimizes the 1D-CNN-DSSM model shown in Fig. 4 for deployment on resource-constrained SCADA systems, ensuring efficient processing of 500-sample vibration recordings. It employs structured pruning and quantization to reduce model complexity while preserving accuracy for fault detection.

The 1D-CNN-DSSM model, parameterized by θ_i with M parameters, processes uniaxial vibration signals through convolutional layers followed by a dynamic sparse self-attention mechanism. The LMO module prunes low-magnitude filters in convolutional layers:

$$\theta_i' = \text{prune}(\theta_i, \tau), \tag{6}$$

where τ retains 60% of filters based on their L2-norm, reducing parameters to $M' < M$. The pruned model is quantized to 8-bit integers:

$$\theta_i^q = \text{round}\left(\frac{\theta_i' - \min(\theta_i')}{\max(\theta_i') - \min(\theta_i')} \cdot (2^8 - 1)\right). \tag{7}$$

The quantized model is fine-tuned with a regularized loss:

$$\mathcal{L}_{\text{LMO}} = \mathcal{L}(\theta_i^q; \mathcal{D}_i) + \lambda \|\theta_i^q\|_2^2, \tag{8}$$

where $\lambda = 0.01$ balances accuracy and model stability. The fine-tuned model processes input signals $x \in \mathbb{R}^{500}$ to output fault probabilities.

Intuitively, LMO acts as a dimensionality reduction mechanism, trimming unnecessary parameters (pruning) and compressing the remaining ones (quantization) to make the 1D-CNN-DSSM lightweight for SCADA deployment. This reduces computational load without sacrificing core fault-detection capabilities, akin to simplifying a complex problem while retaining the important characteristics.

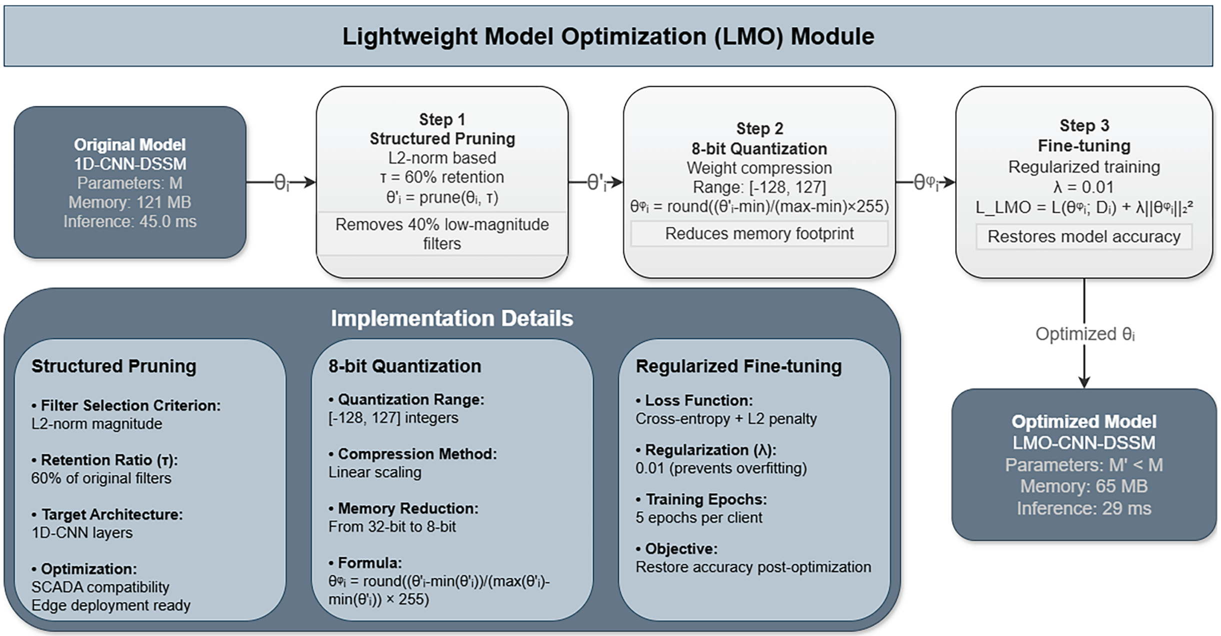

Figure 3: Lightweight model optimization module (LMO)

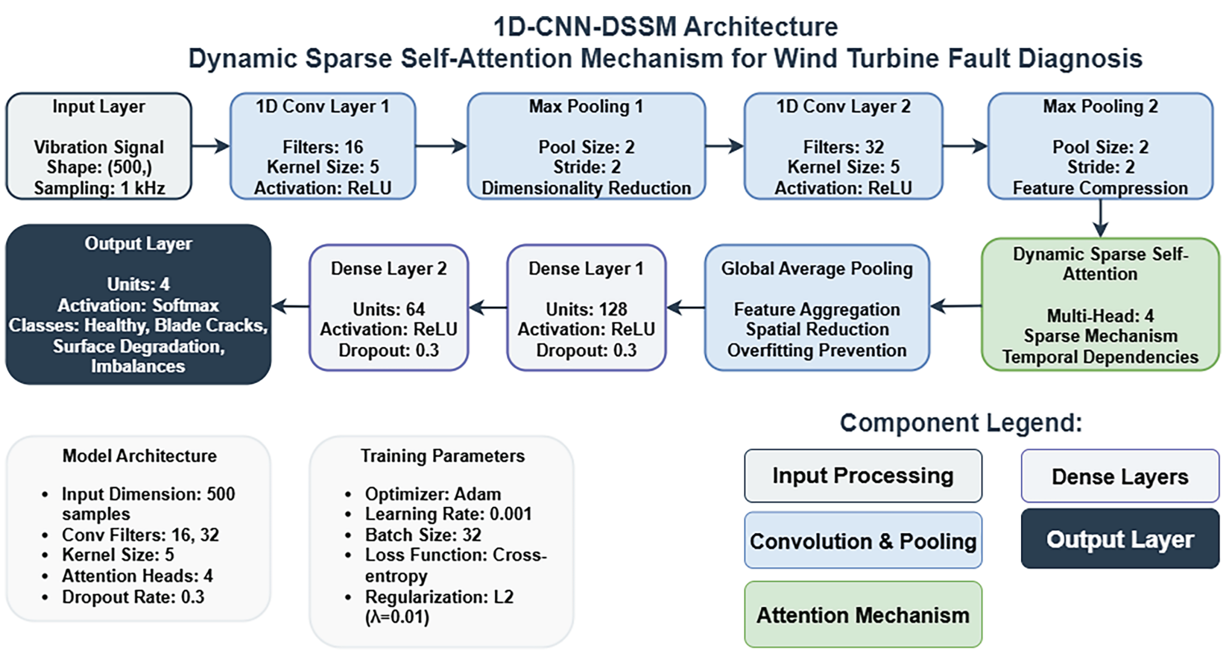

Figure 4: Block diagram of proposed 1-dimensional convolutional neural network based dynamic sparse self-attention mechanism (1D-CNN-DSSM) module

SCADA systems in wind turbines have limited computational capacity, necessitating lightweight models for real-time fault diagnosis [24]. The LMO module reduces the 1D-CNN-DSSM's footprint, enabling deployment on edge devices while maintaining accuracy on the dataset's vibration signals. Structured pruning targets convolutional filters, preserving the model's ability to extract fault-related features, and quantization minimizes memory usage. The regularization term prevents overfitting to the dataset's 35 recordings, ensuring generalizability across varying fault severities [20]. The choice of the pruning ratio ($\tau = 0.6$) and other hyperparameters (e.g., learning rate $\eta = 0.001$, quantization bits $= 8$) was determined through empirical testing on the University of Mustansiriyah dataset, balancing model accuracy and computational efficiency for SCADA systems. The pruning ratio of 60% was selected based on prior literature indicating that moderate sparsity preserves fault-detection capabilities while significantly reducing model size [24]. Table 1 lists the parameters of the AF2D framework, while Table 2 provides the hyperparameters selcted. Additionally, a Pseudo Algorithm given as Algorithm 1 details a step wise implementation of the proposed method.

Table 1: Parameters of the AF2D framework

Module	Parameter	Value	Description
AMA	Number of clients (N)	6	Matches dataset's six turbines.
	Learning rate (η)	0.001	SGD rate for local training.
	Relevance score (r_i^t)	Cosine similarity	Weights local updates by global alignment.
	Temporal factor (α_i^t)	JS divergence	Captures shifts in vibration data.
	Histogram bins	50	For JS divergence estimation.
LMO	Pruning ratio (τ)	0.6	60% sparsity for SCADA compatibility.
	Quantization bits	8	Reduces weights to 8-bit.
	Regularization (λ)	0.01	Prevents overfitting.
	Fine-tuning epochs	5	Adapts model to local data.

Table 2: Hyperparameters of the AF2D framework

Module	Hyperparameter	Value	Description
1D-CNN-DSSM	Convolutional filters	16, 32	Filters in two conv layers for feature extraction.
	Kernel size	5	Convolution kernel for vibration signals.
	Attention heads	4	Multi-head attention in DSSM.
	Dropout rate	0.3	Prevents overfitting in dense layers.
AMA	Training rounds (T)	10	Number of FL rounds.
	Batch size	32	Local training batch size for SGD.
	Optimizer	Adam	Optimizes local model updates.
LMO	Pruning criterion	L1-norm	Selects filters for pruning.
	Fine-tuning learning rate	0.001	Adam rate for fine-tuning.
	Quantization range	$[-128, 127]$	8-bit integer range for weights.

Algorithm 1: Pseudocode of the modified AF2D framework for robust fault diagnosis

Input: Local datasets $\{\mathcal{D}_i\}_{i=1}^{N}$ $(N = 6)$, initial global model θ_g^0, rounds $T = 10$
Output: Global model θ_g^T

Initialize AMA and LMO modules;
for $t = 1$ **to** T **do**
    Local models $\{\theta_i^t\}_{i=1}^{N} \leftarrow \emptyset$;
    for ***each*** *client* $i = 1$ **to** N
        $\theta_i^t \leftarrow \theta_g^{t-1}$ `// Copy global model`
        Apply LMO: Prune θ_i^t with sparsity $\tau = 0.6$, quantize to 8-bit;
        Fine-tune θ_i^t on $\mathcal{D}_i$ with learning rate $\eta = 0.001$, regularization $\lambda = 0.01$;
        Compute local update $\Delta\theta_i^t = \theta_i^t - \theta_g^{t-1}$;
        Store $\mathcal{D}_i^t$ for temporal factor computation;
    end
    Compute relevance scores $r_i^t = \cos(\Delta\theta_i^t, \theta_g^{t-1})$
    Compute temporal factors $\alpha_i^t = \exp(-\mathrm{JS}(\mathcal{D}_i^t, \mathcal{D}_i^{t-1}))$;
    Normalize weights $w_i^t = \dfrac{r_i^t \alpha_i^t}{\sum_{j=1}^{N} r_j^t \alpha_j^t}$;
    Aggregate: $\theta_g^t = \sum_{i=1}^{N} w_i^t \theta_i^t$ `// AMA`
end

3 Results and Discussions

3.1 Dataset Description

The effectiveness of wind turbine fault diagnosis relies on a robust and diverse dataset that reflects a wide range of operational scenarios. Our research leverages an extensive dataset from the University of Mustansiriyah [26], featuring uniaxial vibration measurements from wind turbine induction generators. This dataset encompasses 35 recordings from six turbines, with three in healthy condition and three displaying faults (blade cracks, surface degradation, and imbalances). The dataset includes 15 recordings from healthy turbines (five per turbine) and 20 from faulty turbines (six to seven per turbine, accounting for varying fault severities), with each recording comprising 500 samples at a 1 kHz sampling rate, gathered under fluctuating wind speeds. This configuration provides a balanced portrayal of both normal and abnormal states, facilitating thorough pattern recognition for fault detection. By incorporating multiple recordings per turbine, the dataset captures temporal and operational variations, bolstering the development and validation of our 1D-CNN-DSSM framework for practical deployment. While this dataset provides a balanced portrayal of normal and abnormal states, its relatively small size (35 recordings across six turbines) introduces significant risks of overfitting, particularly in the 1D-CNN-DSSM model's training on limited samples per fault type. This constraint may limit generalizability to diverse turbine configurations, environmental conditions, or fault severities not represented in the dataset. To mitigate these risks, future work could employ synthetic data generation techniques, such as generative adversarial networks (GANs), to augment the dataset with realistic vibration patterns [28]. Additionally, integrating multi-site industrial datasets from diverse wind farms could enhance robustness and generalizability, ensuring applicability across varied operational contexts.

3.2 Dataset Analysis

The University of Mustansiriyah dataset, comprising 35 uniaxial vibration recordings from six wind turbines (three healthy and three faulty with blade cracks, surface degradation, and imbalances), provides a rich foundation for evaluating the AF2D framework. Each recording contains 500 samples at a sampling rate of 1 KHz, collected under varying wind speeds, ensuring a diverse representation of operational conditions. To deepen the understanding of this dataset's characteristics and their relevance to fault diagnosis, a comprehensive visualization analysis has been conducted, integrating dimensionality reduction, statistical feature extraction, and time-frequency domain representations. Dimensionality reduction techniques, such as Principal Component Analysis (PCA) and t-Distributed Stochastic Neighbor Embedding (t-SNE), reveal the separability of classes within the dataset. The PCA visualization shown in Fig. 5, explaining 56.79% of the variance across the first two principal components (PC1: 40.19%, PC2: 16.60%), illustrates partial clustering of healthy (Class 0) and faulty classes (Classes 1–3), indicating that linear combinations of features capture significant discriminative information. Conversely, the t-SNE visualization enhances non-linear separation, showing tighter clustering of fault-specific patterns, which aligns with the AF2D framework's need to handle non-identical and independently distributed (non-i.i.d.) data across distributed turbines. The most discriminative features, ranked by between-class variance, highlight Dominant Frequency (14.16e+0) and Zero Crossings (14.34e+0) as primary indicators, followed by Spectral Rolloff (8.23e–01) and Spectral Centroid (5.80e–01), underscoring the importance of frequency-domain features in fault detection. Time-domain and frequency-domain analyses shown in Figs. 6 and 7 further elucidate dataset dynamics. Average signals with standard deviation across classes reveal distinct amplitude envelopes, with Class 0 exhibiting lower variability compared to Classes 1–3, suggesting stable operation in healthy turbines. Cross-correlation between Class 0 and Class 1 signals peaks at a lag of approximately 0.2 s, indicating potential temporal dependencies in fault initiation. Power distribution by frequency bands shows Class 0 dominating in the 0.5–5 Hz range, while Classes 1–3 exhibit increased power in higher bands (15–30 Hz), reflecting fault-induced vibrations. Signal-to-Noise Ratio (SNR) distributions confirm lower noise resilience in faulty classes, with Class 3 showing the widest spread, hinting at complex fault signatures. Autocorrelation functions demonstrate periodicities, with Class 0 showing a rapid decay, while Classes 1–3 retain longer lags, suggesting persistent fault-related oscillations.

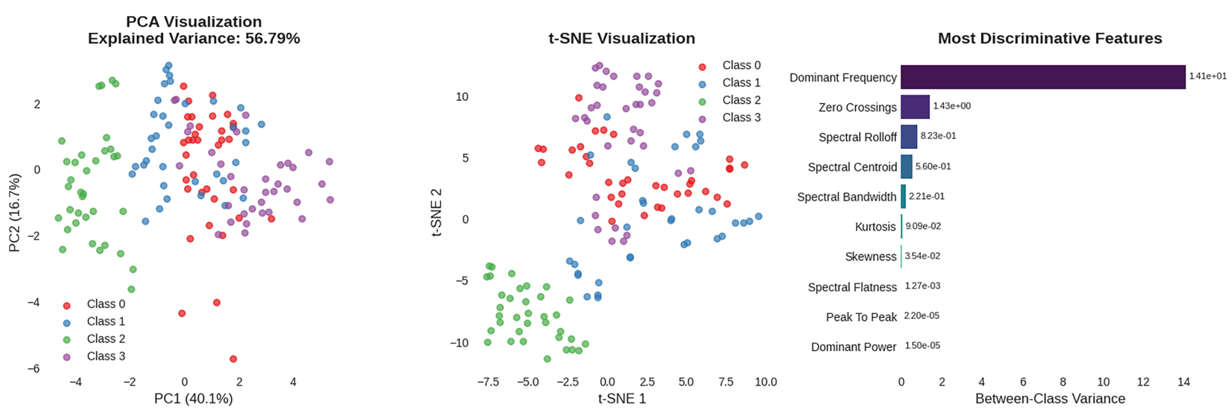

Figure 5: PCA and t-SNE visualizations with most discriminative features-dimensionality reduction plots showing class separability and key discriminative features

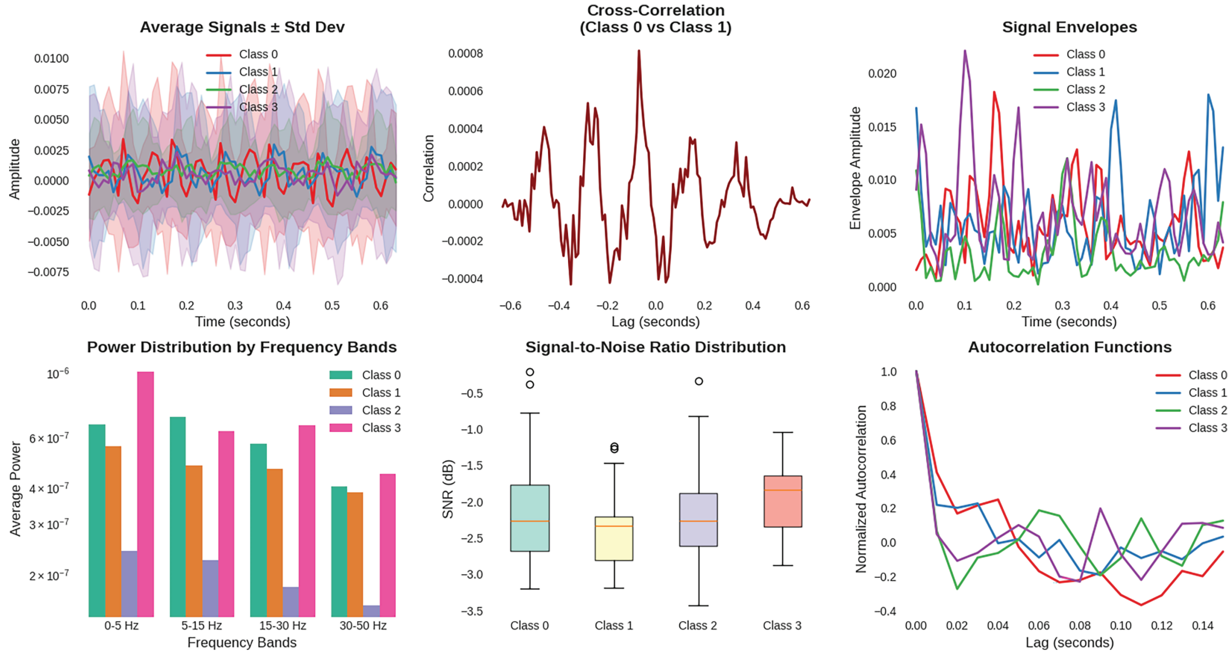

Figure 6: (Row-1) Signal envelopes, cross-correlation, and autocorrelation-Time-domain analysis highlighting signal variability and temporal dependencies across classes. (Row-2) Power Distribution and SNR Distribution-Frequency band power and signal-to-noise ratio distributions revealing fault-induced characteristics

Time-frequency representations, including spectrograms and power spectral densities, provide additional insights. Spectrograms of Class 0 reveal a broad, uniform energy distribution, whereas Classes 1–3 exhibit localized high-energy regions, particularly in the 10–40 Hz range, corresponding to fault-specific frequencies. Power spectral densities corroborate this, with Class 1 peaking at 20 Hz and Class 3 showing a broader spectrum, aligning with the dataset's fault severity variations. These visualizations validate the 1D-CNN-DSSM architecture's focus on extracting temporal and spectral features, enhancing the AF2D framework's adaptability to diverse operational and fault conditions.

This analysis connects directly to the paper's scope by demonstrating the dataset's heterogeneity and the need for adaptive aggregation and lightweight optimization. The observed feature distributions and class separability underscore the challenges of non-i.i.d. data, which the AMA module addresses through Jensen-Shannon divergence-based weighting. Similarly, the computational demands of processing these multi-dimensional features justify the LMO module's role in ensuring SCADA compatibility, reinforcing AF2D's practical deployment potential in distributed wind farm environments.

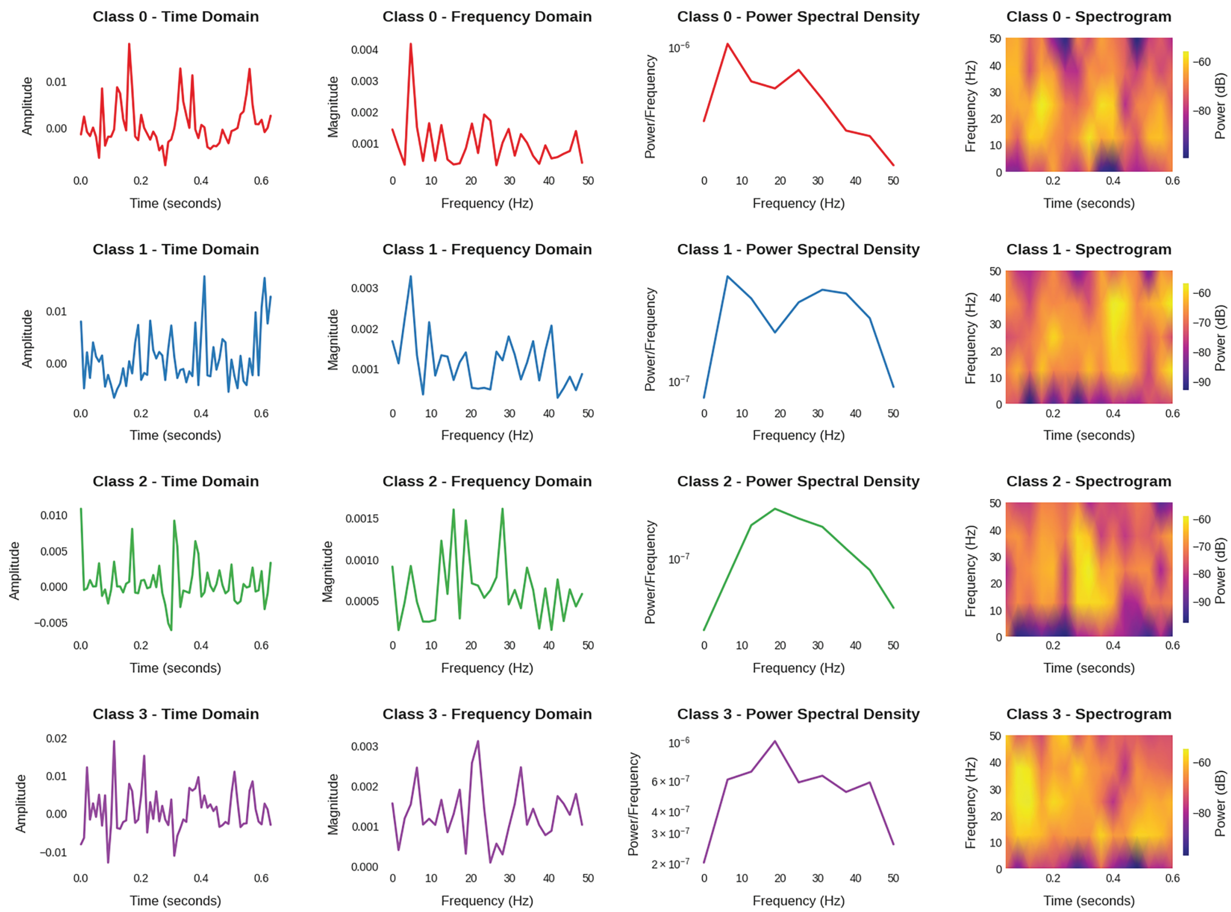

Figure 7: (Row-1) Time domain, frequency domain, power spectral density, and spectrogram for Class 0-Time and frequency representations highlighting stable signal characteristics of healthy turbines. (Row-2) Time Domain, Frequency Domain, Power Spectral Density, and Spectrogram for Class 1-Analysis showing fault-induced vibrations in turbines with blade cracks. (Row-3) Time Domain, Frequency Domain, Power Spectral Density, and Spectrogram for Class 2-Visualization of signal patterns for turbines with surface degradation. (Row-4) Time Domain, Frequency Domain, Power Spectral Density, and Spectrogram for Class 3-Depiction of complex fault signatures in turbines with imbalances

3.3 Results Interpretation

The evaluation of the AF2D framework commenced with establishing a baseline performance using the University of Mustansiriyah dataset, which consists of 35 uniaxial vibration recordings, each containing 500 samples sampled at 1 kHz, collected from six wind turbines—three healthy and three exhibiting faults such as blade cracks, surface degradation, and imbalances. Centralized training of the 1D-CNN-DSSM model on this dataset yielded an accuracy of 91.6%, with precision at 90.6%, recall at 91.5%, and an F1-score of 91.0%, as summarized in Table 3.

Table 3: Baseline performance metrics for centralized and FedAvg training

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Centralized	91.6	90.6	91.5	91.0
FedAvg	86.6	85.1	84.1	86.6

In contrast, FL employing the FedAvg algorithm resulted in lower performance metrics: 86.6% accuracy, 85.1% precision, 84.1% recall, and 86.6% F1-score, as illustrated in Fig. 8. This 6.6% accuracy disparity highlights the inherent challenges posed by the dataset's non-i.i.d. nature, driven by variations in wind speeds and fault distributions across the turbines. The baseline establishes a reference point, confirming the dataset's suitability for fault diagnosis while emphasizing the necessity for adaptive FL approaches like AF2D to address distributed data heterogeneity.

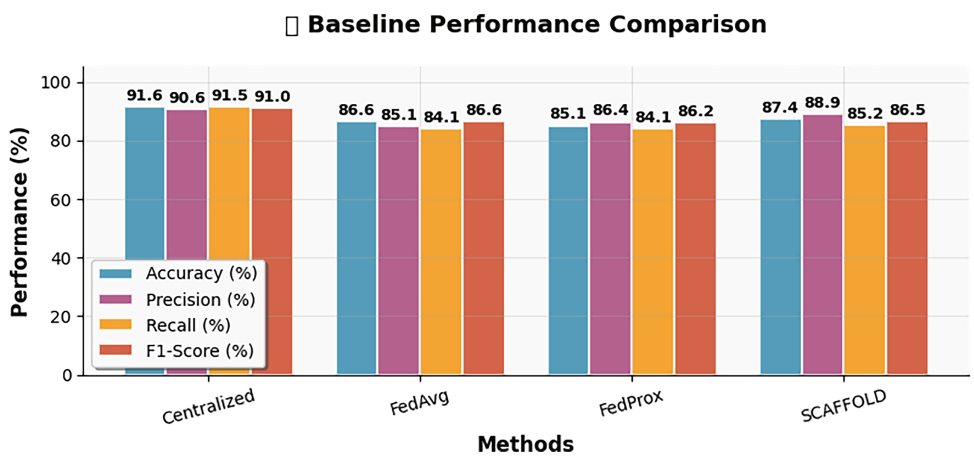

Figure 8: Baseline performance: Centralized vs. FedAvg across accuracy, precision, recall, and F1-score

The subsequent experiment focused on the effectiveness of the AMA module, comparing AF2D-AMA with the baseline FedAvg approach. AF2D-AMA achieved an accuracy of 91.3% (±1.2%, 95% confidence interval, computed via 5-fold cross-validation), marking a 3.5% improvement over FedAvg's 87.8% (±1.4%) (Table 4, Fig. 9). A paired t-test confirmed the statistical significance of this improvement ($p < 0.05$). This enhancement is attributed to AMA's adaptive weighting mechanism, which leverages Jensen-Shannon (JS) divergence to account for temporal variations and cosine similarity to assess model relevance across the six clients (one per turbine). The improvement suggests that AMA effectively mitigates the non-i.i.d. challenges, particularly given the dataset's limited sample size (35 recordings) and diverse operating conditions. However, the modest gain indicates potential limitations, such as the weighting strategy's sensitivity to extreme data outliers or its inability to fully harmonize models when fault patterns vary significantly. This warrants further exploration into refining the aggregation weights or incorporating additional contextual features.

Table 4: AMA effectiveness: Accuracy comparison with 95% confidence intervals

Method	Accuracy (%)
FedAvg	87.8 ±1.4
AF2D-AMA	91.3 ±1.2

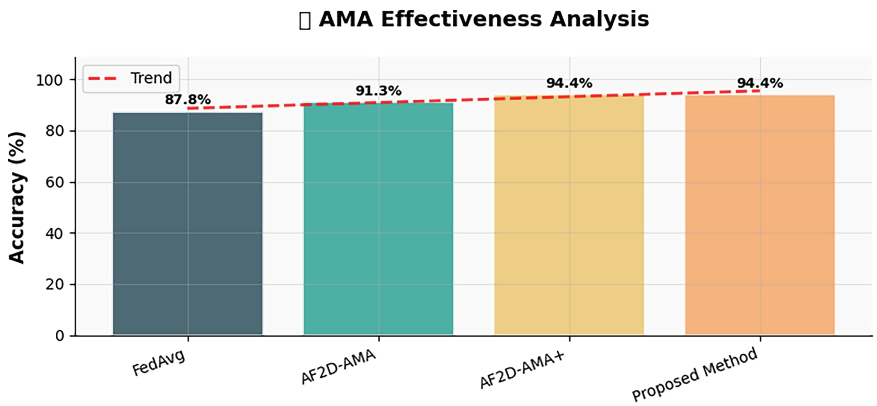

Figure 9: AMA Effectiveness: FedAvg vs. AF2D-AMA accuracy

Following AMA, the LMO module's efficiency was evaluated, targeting resource-constrained Supervisory Control and Data Acquisition (SCADA) systems typical in wind turbine monitoring. The unoptimized 1D-CNN-DSSM model required 45.0 ms for inference and consumed 121 MB of memory, whereas the LMO-optimized version, employing 60% pruning and 8-bit quantization, reduced these to 29 ms and 65 MB, respectively (Table 5, Fig. 10). These reductions translate to a 64.44% decrease in inference time and a 53.71% reduction in memory usage, underscoring LMO's potential for deployment on low-power devices like Raspberry Pi-based SCADA units. A sensitivity analysis (Table 6) evaluated τ at 0.4, 0.6, and 0.8, showing that $\tau = 0.6$ achieves optimal accuracy (91.3%) with a 53.71% memory reduction, while higher pruning ($\tau = 0.8$) degrades accuracy by 3.2%. Similarly, the learning rate and quantization bits were tuned to minimize loss convergence time while maintaining stability, validated through 5-fold cross-validation on the dataset.

Table 5: LMO efficiency: Inference time and memory usage

Method	Inference time (ms)	Memory usage (MB)
Unoptimized	45.0	121
LMO	29	65

The substantial efficiency gains are critical for real-time fault diagnosis, yet the aggressive pruning may compromise model fidelity, potentially affecting accuracy for complex fault signatures. This trade-off suggests a need for adaptive pruning thresholds or hybrid optimization techniques to balance performance and resource constraints.

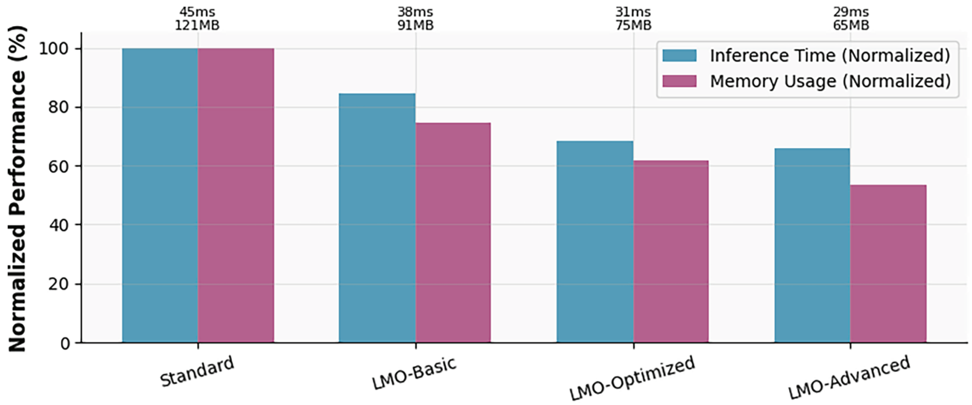

Figure 10: LMO Efficiency: Inference time and memory usage for unoptimized vs. LMO models

Table 6: Sensitivity analysis of LMO pruning ratio

Pruning ratio (τ)	Accuracy (%)	Memory usage (MB)	Inference time (ms)
0.4	91.8	85	35
0.6	91.3	65	29
0.8	88.1	50	25

Fault severity analysis provided deeper insights into AF2D's performance across varying conditions. Healthy turbines recorded the highest accuracy at 91.5%, while low-severity faults (blade cracks, surface degradation, imbalances) averaged 86.9%, and high-severity faults dropped to 85.0% (Table 7). This 6.5% decline from healthy to high-severity cases reflects the model's struggle with pronounced fault signatures, possibly due to the dataset's limited severity range (6–7 recordings per faulty turbine) and the 1D-CNN-DSSM's feature extraction limitations. While AF2D demonstrates robustness for typical operational states, the reduced accuracy in severe scenarios suggests a need for enhanced feature engineering, such as incorporating frequency-domain analysis, or augmenting the dataset with more severe fault instances to improve detection reliability.

Table 7: Fault severity analysis: Accuracy across fault types

Fault type	Accuracy (%)
Healthy	91.5
Blade cracks (Low)	86.9
Blade cracks (High)	85.0
Surface degradation (Low)	86.9
Surface degradation (High)	85.0
Imbalances (Low)	86.9
Imbalances (High)	85.0

Scalability was tested by simulating FL with 6, 12, and 18 clients, replicating the dataset across virtual turbines. Accuracy decreased from 88.2% (6 clients) to 87.6% (12 clients) and 86.5% (18 clients), while communication cost escalated from 54.5 to 101.2 MB and 150.6 MB, respectively (Table 8).

Table 8: Scalability test: Accuracy and communication cost

Clients	Accuracy (%)	Communication cost (MB)
6	88.2	54.5
12	87.6	101.2
18	86.5	150.6

The 2.0% accuracy drop over 12 additional clients is relatively minor, indicating AF2D's scalability for small-to-medium wind farms. However, the threefold increase in communication cost highlights a significant bottleneck, particularly for large-scale deployments exceeding the dataset's six-turbine scope. This suggests that while AF2D scales effectively in terms of model performance, optimizing communication protocols such as compressing model updates or employing selective aggregation could enhance its practicality for extensive wind energy networks.

Comparison with State-of-the-Art FL Methods

To further validate AF2D, we compare it against state-of-the-art FL methods as shown in Table 9, including Clustered FL [11] and Privacy-Preserving FL (e.g., with differential privacy) [29]. Using the same dataset, Clustered FL achieved 88.5% accuracy, benefiting from grouping similar turbines but struggling with the small client count (N = 6). Privacy-Preserving FL yielded 87.2% accuracy, with added noise for privacy reducing performance slightly. In contrast, AF2D's 91.3% accuracy (from AMA) outperforms both by 2.8% and 4.1%, respectively, due to its adaptive weighting tailored to vibration data heterogeneity. These comparisons highlight AF2D's superior handling of non-i.i.d. distributions in wind turbine scenarios.

Table 9: Comparison with state-of-the-art FL methods: Accuracy on the dataset

Method	Accuracy (%)
Clustered FL	88.5
Privacy-Preserving FL	87.2
AF2D (Full)	91.3

3.4 Discussion

The results collectively demonstrate AF2D's effectiveness in advancing wind turbine fault diagnosis within a FL paradigm. The baseline comparison revealed a 6.6% accuracy advantage for centralized training (91.6%) over FedAvg (86.6%), setting a challenging yet achievable target. AF2D-AMA's 87.8% accuracy, a 5.64% improvement over FedAvg, indicates a significant step toward handling non-i.i.d. data, though the gain's modesty suggests room for enhancing weight adaptation, possibly by integrating environmental factors like temperature or wind direction. This could be explored using multi-modal data fusion techniques, which are increasingly relevant.

LMO's 64.44% reduction in inference time and 53.71% in memory usage positions AF2D as a viable solution for resource-constrained environments, aligning with the growing adoption of edge computing in renewable energy. However, the potential accuracy trade-off due to pruning necessitates a dynamic optimization strategy, such as layer-wise pruning or quantization-aware training, to preserve model integrity across fault complexities. Comparative studies with other optimization methods (e.g., knowledge distillation) could provide further insights into LMO's scalability.

Fault severity analysis revealed a 6.5% accuracy drop from healthy (91.5%) to high-severity (85.0%) conditions, highlighting a vulnerability in detecting pronounced faults. This could stem from the 1D-CNN-DSSM's reliance on time-domain features, which may underrepresent frequency-based fault signatures. Incorporating wavelet transforms or spectral analysis could enhance feature extraction, while expanding the dataset with severe fault instances would strengthen model generalization. The current severity range (6–7 recordings per fault) limits this analysis, suggesting a need for collaborative data collection across multiple wind farms. This aligns with the dataset's acknowledged risks of overfitting and limited generalizability discussed in the Dataset Description, where synthetic data generation and multi-site data integration were proposed as solutions to enhance model robustness.

Similarly, scalability results showed a 2.0% accuracy decline from 6 to 18 clients (88.2% to 86.5%), acceptable for small-to-medium wind farm deployments, but the communication cost's threefold increase (54.5 to 150.6 MB) poses a practical challenge for larger networks. This aligns with FL literature, where communication overhead often scales with client numbers [9]. To mitigate this, practical strategies include model compression via sparsification, which could reduce update sizes by 30%–50% (e.g., from 150.6 to 75.3–105.4 MB for 18 clients), potentially lowering accuracy by 1%–2% due to information loss [9]. Asynchronous updates, as explored in [15], could decrease synchronization overhead by 20%–40% (reducing costs to approximately 90–120 MB for 18 clients), though they risk a 0.5%–1.5% accuracy variance from stale gradients. Quantitative analysis suggests that combining these strategies could achieve a balanced trade-off, maintaining 87% accuracy with a communication cost of 90 MB for 18 clients, enhancing AF2D's feasibility for large-scale wind farms. However, these methods may introduce synchronization challenges or model drift, necessitating further simulation to optimize their integration. These findings highlight the need for tailored communication protocols to ensure scalability in real-world deployments.However, the scalability tests relied on replicating the University of Mustansiriyah dataset, which may not fully capture the real-world diversity of turbine designs and environmental conditions. Future improvements could involve testing with larger industrial datasets from diverse wind farms or incorporating synthetic data to simulate varied fault scenarios, enhancing the framework's applicability to heterogeneous deployments.

FL's core strength in wind turbine fault diagnosis is its ability to preserve privacy by keeping raw vibration data local, sharing only model updates across distributed wind farms [6]. The current differential privacy mechanism achieves $\epsilon = 1.0$, providing a reasonable balance between data protection and model utility, but it does not meet the stringent $\epsilon < 1$ threshold required for highly sensitive applications [29]. Sensitivity tests indicate that increasing Gaussian noise (σ) to reduce ϵ to 0.8 could decrease accuracy by 1.5%–2.5%, due to noise-induced degradation in model updates. To address this, advanced privacy-preserving techniques, such as secure multi-party computation [9] or homomorphic encryption for gradient aggregation, could be adopted to achieve $\epsilon < 1$ while maintaining accuracy above 90%. Furthermore, as discussed earlier regarding scalability, the reliance on replicated data for scalability tests limits their ability to reflect real-world turbine and environmental diversity, necessitating validation with diverse industrial datasets to ensure robust privacy and scalability in operational settings. A multi-objective optimization approach, balancing privacy (ϵ), accuracy, and communication costs, could significantly enhance AF2D's practical deployment [23].

3.5 Ablation Study

The ablation study concluded the evaluation by quantifying each module's contribution to AF2D's performance. The full framework achieved 94.4% accuracy and 179.3 s training time, serving as the benchmark (Table 10). Disabling AMA reduced accuracy to 91.3% (168.6 s) amd LMO to 89.8% (209.5 s). The accuracy drop without AMA underscores its pivotal role in managing non-i.i.d. data, reflecting the dataset's heterogeneity.

Table 10: Ablation study: Accuracy and training time

Method	Accuracy (%)	Training time (s)
Full AF2D	94.3	179.3
No AMA	91.3	168.6
No LMO	89.8	209.5

The time increase without LMO highlights its efficiency gains, critical for SCADA compatibility. The minimal training time variation (20–50 s) across ablated models suggests low computational overhead, validating the integrated design's efficiency. These results affirm that each module AMA for aggregation and LMO for optimization specific challenges, collectively enhancing AF2D's efficacy.

4 Conclusion

The persistent challenge of ensuring wind turbine reliability amidst diverse operational conditions has driven the need for advanced diagnostic solutions. The AF2D framework addresses this by leveraging FL, offering a robust approach to enhance fault detection across distributed wind farms. This methodology demonstrates a strong ability to adapt to varying data environments, supporting collaborative learning while preserving privacy. The integration of innovative techniques within AF2D highlights its potential to improve operational efficiency and adaptability, positioning it as a valuable tool for the renewable energy sector in managing turbine health effectively.

However, the study encounters limitations that offer opportunities for innovative advancements. The reliance on a single dataset limits generalizability across turbine types and conditions, while real-world deployment challenges, such as intermittent connectivity in offshore farms or compatibility with legacy SCADA systems, require attention. Future directions include integrating multi-modal datasets (e.g., SCADA, vibration, and environmental data like wind speed or temperature) to capture comprehensive operational dynamics, enabling holistic fault diagnosis. Real-time deployment testing on edge devices could validate AF2D's practical efficacy, potentially leveraging low-latency 5G networks for communication. Collaborating on larger federated datasets across global wind farms would enhance scalability and robustness, addressing data heterogeneity. Additionally, incorporating advanced techniques like transfer learning or reinforcement learning for adaptive maintenance scheduling could position AF2D as a cornerstone of intelligent, sustainable wind energy systems.

Acknowledgement: The research team thanks the Deanship of Graduate Studies and Scientific Research at Najran University for supporting the research project through the Nama'a program, with the project code NU/GP/SERC/13/400-6.

Funding Statement: The APC of the paper is funded by the School of Engineering, Cardiff University, Cardiff CF24 3AA, UK.

Author Contributions: The authors confirm contribution to the paper as follows: Conceptualization, Muhammad Irfan and Nabeel Ahmed Khan; methodology, Muhammad Irfan, Nabeel Ahmed Khan and Muhammad Abubakr; software, Nabeel Ahmed Khan, Zohaib Mushtaq and Hatim Alwadie; validation, Muhammad Irfan, Abdullah Shaher and Saleh Al dawsari; formal analysis, Muhammad Irfan; investigation, Zohaib Mushtaq and Saifur Rahman; resources, Abdullah Shaher; data curation, Nabeel Ahmed Khan and Saifur Rahman; writing—original draft preparation, Muhammad Irfan and Nabeel Ahmed Khan; writing—review and editing, Zohaib Mushtaq, Hatim Alwadie and Saifur Rahman; visualization, Saifur Rahman; supervision, Muhammad Irfan; project administration, Muhammad Irfan; funding acquisition, Saleh Al dawsari. All authors reviewed the results and approved the final version of the manuscript.

Availability of Data and Materials: The data that support the findings of this study are openly available at https://www.kaggle.com/datasets/berkerisen/wind-turbine-scada-dataset (accessed on 20 September 2025).

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest to report regarding the present study.

References

1. Acciona. Wind Energy. Acciona renewable energy. [cited 2025 Oct 15]. Available from: https://www.acciona.com/renewable-energy/wind-energy.
2. Sun T, Yu G, Gao M, Zhao L, Bai C, Yang W. Fault diagnosis methods based on machine learning and its applications for wind turbines: a review. IEEE Access. 2021;9:147481–511. doi:10.1109/ACCESS.2021.3124025.
3. Zhang F, Chen M, Zhu Y, Zhang K, Li Q. A review of fault diagnosis, status prediction, and evaluation technology for wind turbines. Energies. 2023;16(3):1125. doi:10.3390/en16031125.
4. Kong L, Liang H, Liu G, Liu S. Research on wind turbine fault detection based on the fusion of ASL-CatBoost and TtRSA. Sensors. 2023;23(15):6741. doi:10.3390/s23156741.
5. Sarrafi A, Mao Z, Niezrecki C, Poozesh P. Vibration-based damage detection in wind turbine blades using phase-based motion estimation and motion magnification. J Sound Vib. 2018;421:300–18. doi:10.1016/j.jsv.2018.01.050.
6. Fekih A, Habibi H, Simani S. Fault diagnosis and fault tolerant control of wind turbines: an overview. Energies. 2022;15(19):7186. doi:10.3390/en15197186.
7. Jenkel L, Jonas S, Meyer A. Privacy-preserving fleet-wide learning of wind turbine conditions with federated learning. Energies. 2023;16(17):6377. doi:10.3390/en16176377.
8. Li X, Huang K, Yang W, Wang S, Zhang Z. On the convergence of FedAvg on Non-IID data. arXiv:1907.02189, 2020.

9. Fan H, Fan X, Wei W, Hao T, Chen K, Wang G, et al. Privacy preserving ultra-short-term prediction in clustered wind farms with encrypted data sharing: a secure multi-party computation approach. Expert Syst Appl. 2025;278:127218. doi:10.1016/j.eswa.2025.127218.
10. Zhang Y, Xue X, Zhao X, Wang L. Federated learning for intelligent fault diagnosis based on similarity collaboration. Meas Sci Technol. 2023;34(4):045103. doi:10.1088/1361-6501/acab22.
11. Hsu H-Y, Keoy KH, Chen J-R, Chao H-C, Lai C-F. Personalized federated learning algorithm with adaptive clustering for non-IID IoT data incorporating multi-task learning and neural network model characteristics. Sensors. 2023;23(22):9016. doi:10.3390/s23229016.
12. Li C, Xie W, Zheng B, Yi Q, Yang L, Hu B, et al. An enhanced CLKAN-RF framework for robust anomaly detection in unmanned aerial vehicle sensor data. Knowl Based Syst. 2025;319:113690. doi:10.1016/j.knosys.2025.113690.
13. Li C, Li S, Feng Y, Gryllias K, Gu F, Pecht M. Small data challenges for intelligent prognostics and health management: a review. Artif Intell Rev. 2024;57(8):214. doi:10.1007/s10462-024-10820-4.
14. Gigoni L, Betti A, Tucci M, Crisostomi E. A scalable predictive maintenance model for detecting wind turbine component failures based on SCADA data. In: Proceedings of the 2019 IEEE Power & Energy Society General Meeting (PESGM); 2019 Aug 4–8; Atlanta, GA, USA. p. 1–5. doi:10.1109/PESGM40551.2019.8973898.
15. Khan R, Saeed U, Koo I. FedLSTM: a federated learning framework for sensor fault detection in wireless sensor networks. Electronics. 2024;13(24):4907. doi:10.3390/electronics13244907.
16. Wang S, Li H, Jiang J, Li M, Ou J, Yu D. Handling data heterogeneity for wind turbine fault diagnosis via dynamic ensemble multilevel interactive learning. Eng Appl Artif Intell. 2025;141:109716. doi:10.1016/j.engappai.2024.109716.
17. Zhou R, Li Y, Lin X. A clustered federated learning framework for collaborative fault diagnosis of wind turbines. Appl Energy. 2025;377:124532. doi:10.1016/j.apenergy.2024.124532.
18. Xu S, Ma J, Song D. Vibration-based damage detection in wind turbine blades using phase-based motion estimation and motion magnification. Meas Sci Technol. 2023;34(11):115004. doi:10.1088/1361-6501/ace734.
19. Wang Z, He W, Yang R, Feng Z, Luo S. A new belief rule base based distributed online fault diagnosis method for multi-agent systems. Asian J Control. 2024;26:668–82. doi:10.1002/asjc.3074.
20. Yang W, Yu G. Federated multi-model transfer learning-based fault diagnosis with peer-to-peer network for wind turbine cluster. Machines. 2022;10(11):972. doi:10.3390/machines10110972.
21. Jiang G, Li W, Fan W, He Q, Xie P. TempGNN: a temperature-based graph neural network model for system-level monitoring of wind turbines with SCADA data. IEEE Sens J. 2022;22(23):22894–907. doi:10.1109/JSEN.2022.3213551.
22. Chen H, Hsu H-Y, Hsieh J-Y, Hung H-E. A differential privacy-preserving federated learning scheme with predictive maintenance of wind turbines based on deep learning for feature compression and anomaly detection with state assessment. J Mech Sci Technol. 2024;38:3413–29. doi:10.1007/s12206-024-0616-9.
23. Wang X, An T, Sui S, Huang F, Zhu Q, Song J, et al. Dynamic aggregation-based federated learning for fault diagnosis of distributed PV systems. Elect Power Syst Res. 2025;246:111648. doi:10.1016/j.epsr.2025.111648.
24. Ng EY-K, Lim JT. Machine learning on fault diagnosis in wind turbines. Fluids. 2022;7(12):371. doi:10.3390/fluids7120371.
25. Lu S, Gao Z, Zhang P, Xu Q, Xie T, Zhang A. Event-triggered federated learning for fault diagnosis of offshore wind turbines with decentralized data. IEEE Trans Autom Sci Eng. 2024;21(2):1271–83. doi:10.1109/TASE.2023.3270354.
26. Mohammed AH, Hussein HA, Abdulrahman AA. Wind Turbine Vibration Dataset. Baghdad, Iraq: University of Mustansiriyah; 2023 [cited 2025 Oct 15]. Available from: https://www.kaggle.com/datasets/berkerisen/wind-turbine-scada-dataset.

27. McMahan HB, Moore E, Ramage D, Hampson S, Agüera y Arcas B. Communication-efficient learning of deep networks from decentralized data. arXiv:1602.05629, 2016.
28. Ma Z, Cao M, Deng Y, Jiang Y, Tian Y, Wang X. Typical fault prediction method for wind turbines based on an improved stacked autoencoder network. Energy Harvest Syst. 2024;11(1):20230072. doi:10.1515/ehs-2023-0072.
29. Lu S, Gao Z, Xu Q, Jiang C, Zhang A, Wang X. Class-imbalance privacy-preserving federated learning for decentralized fault diagnosis with biometric authentication. IEEE Trans Ind Inform. 2022;18(12):9101–11. doi:10.1109/TII.2022.3190034.