International Information and Engineering Technology Association
Advancing the World of Information and Engineering

Journal Européen des Systèmes Automatisés
Vol. 59, No. 1, January, 2026, pp. 1-8
Journal homepage: http://iieta.org/journals/jesa

A Hybrid Deep Learning Framework for Real-Time Condition Monitoring of Wind Turbine Bearings Using Edge-Computing and Multimodal Data Fusion

Hakim Jebari[1*], Amina Eljyidi[2], Siham Rekiek[3], Kamal Reklaoui[2]

[1] Artificial Intelligence, Data Science, and Innovation Research Team, LaBEL, National School of Architecture of Tetouan Innovative Systems Engineering Laboratory, University Abdelmalek Essaâdi, Tétouan 93000, Morocco
[2] Innovative Systems Engineering Research Team, University Abdelmalek Essaâdi, Tétouan 93000, Morocco
[3] Intelligent Automation and BioMedGenomics Laboratory, University Abdelmalek Essaâdi, Tétouan 93000, Morocco

Corresponding Author Email: hjebari1@yahoo.fr

https://doi.org/10.18280/jesa.590101

Received: 31 October 2025
Revised: 22 December 2025
Accepted: 5 January 2026
Available online: 31 January 2026

Keywords:
wind turbine prognostics, predictive maintenance, deep learning, edge AI, multimodal data fusion, bearing fault diagnosis, renewable energy optimization

ABSTRACT

The global energy sector's pivot towards renewables has cemented wind power's role as a foundational technology. However, the economic sustainability of wind energy is critically contingent upon maximizing turbine operational availability and minimizing unscheduled downtime. Drivetrain failures, particularly in rolling-element bearings, represent a predominant source of revenue loss, with incident costs ranging from €250,000 to over €500,000 due to repair expenses and forgone energy production. This research presents a novel, integrated condition monitoring framework that leverages a hybrid deep learning architecture—fusing a lightweight 1D Convolutional Neural Network (CNN) with a Long Short-Term Memory (LSTM) network—for high-fidelity fault diagnosis in wind turbine bearings. The system is engineered for deployment on edge-computing platforms within the turbine nacelle, processing multimodal sensor data streams (vibration, acoustic emission, temperature, and SCADA) in real-time. Validated on a hybrid dataset encompassing real-world operational data and physics-based fault simulations, the model achieves a state-of-the-art diagnostic accuracy of 98.9% F1-score. Crucially, the optimized network demonstrates a latency of under 8 ms and a memory footprint of 4.8 MB, fulfilling the stringent computational constraints of embedded systems. This work signifies a substantial departure from conventional vibration-only techniques, offering a robust predictive maintenance solution that enhances operational reliability, optimizes logistical planning, and improves the economic viability of wind farms, especially in logistically challenging offshore environments.

1. INTRODUCTION

The decarbonization of the global energy matrix is an imperative driven by climate change mitigation and energy security concerns. Wind energy is a vanguard of this transition, with cumulative installed capacity surpassing 900 GW and projections indicating it will supply over one-fifth of global electricity by 2030 [1, 2]. The economic competitiveness of wind power, quantified by the Levelized Cost of Energy (LCOE), is inherently tied to the operational reliability and availability of wind turbines [3]. These complex electromechanical systems operate under highly stochastic and non-stationary loading conditions, which accelerate the degradation of critical components [4].

The drivetrain, encompassing the main bearing, gearbox, and generator, is a primary focus of maintenance activities due to its high failure rate and the exorbitant cost of repairs. Within this assembly, rolling-element bearings are particularly susceptible to failure, accounting for nearly one-third of all significant turbine downtime events. The financial repercussions of drivetrain failures are a primary concern for operators. Recent assessments, including those by industry consortia and academic cost models, highlight that the total cost associated with a major bearing replacement on an offshore turbine—factoring in repair, ultra-expensive jack-up vessel hire, and lost production—typically falls within the range of €250,000 to €500,000 [5, 6]. Concurrently, with the deployment of 8-12 MW turbines, studies show that the daily energy revenue loss from an outage can exceed €25,000, making downtime minimization critically urgent [7].

Traditional maintenance strategies are ill-suited to address these challenges. Reactive (run-to-failure) maintenance leads to catastrophic failures and prolonged downtime. Preventive (time-scheduled) maintenance, while reducing unexpected failures, often results in the premature replacement of healthy components, incurring unnecessary costs and resource expenditure [8, 9]. Consequently, Condition-Based Maintenance (CBM) and its advanced evolution, Predictive Maintenance (PdM), have emerged as industry standards. These paradigms utilize continuous sensor data to assess the health of assets and schedule interventions precisely when needed, thereby optimizing maintenance resources and maximizing availability [9].

However, implementing effective PdM in wind turbines is fraught with challenges. The operational environment is inherently non-stationary; rotational speed and load vary continuously with wind speed, while environmental factors like temperature gradients and tower shadow effects introduce significant noise into sensor signals [10]. Furthermore, the remote and often offshore location of wind farms creates a data transmission bottleneck. The limited bandwidth and high cost of transmitting high-frequency vibration and acoustic emission data preclude a reliance on cloud-centric analytics, necessitating robust, on-site processing capabilities [11, 12].

This research directly addresses these challenges by proposing a novel, edge-deployable hybrid deep learning framework for real-time condition monitoring of wind turbine bearings. The principal contributions of this work are fourfold:

•Architectural innovation: The development of a hybrid 1D Convolutional Neural Network and Long Short-Term Memory (CNN-LSTM) model that synergistically extracts spatial features from high-frequency sensor data and models temporal dependencies from operational history, specifically optimized for embedded deployment.

•Multimodal data fusion: The integration and synergistic fusion of heterogeneous data modalities—vibration, acoustic emission, temperature, and SCADA data—to augment diagnostic accuracy and robustness under variable operating conditions.

•Rigorous validation: A comprehensive evaluation using a hybrid dataset that merges real-world SCADA and vibration data from operational wind farms with high-fidelity, physics-based simulation models of fault progression.

•Practical implementation: A tangible implementation and performance benchmarking on an edge-computing platform (NVIDIA Jetson Xavier NX), conclusively demonstrating the feasibility of real-time, in-nacelle inference under stringent resource constraints.

2. LITERATURE REVIEW AND THEORETICAL FRAMEWORK

2.1 Economic and operational imperatives of wind turbine maintenance

The maintenance of wind turbines is distinguished by unique logistical and economic challenges. Their placement in remote onshore or harsh offshore environments complicates access, while the scale of components necessitates specialized equipment [12]. The financial model of a wind farm is exceptionally sensitive to turbine availability. As noted by the study [13], even a minor increase in availability can translate to significant revenue over a turbine's lifespan. Empirical analyses of failure rates consistently identify the drivetrain as a critical subsystem. Studies [14, 15] have meticulously documented that bearing failures are among the most common and costly events, underscoring the necessity for advanced diagnostic solutions.

2.2 Evolution of condition monitoring techniques

The foundation of wind turbine condition monitoring is rooted in vibration analysis, leveraging techniques well-established in general machinery diagnostics [16]. Standard practices include spectral analysis, envelope (demodulation) techniques, and the extraction of time-domain features (e.g., root mean square (RMS), kurtosis, crest factor) to identify incipient faults [17]. However, the variable-speed operation of wind turbines presents a fundamental challenge, as it smears frequency spectra and obfuscates characteristic bearing fault frequencies, necessitating complex order-tracking methods [18]. Complementary techniques, such as oil debris analysis [19] and acoustic emission (AE) monitoring [20], provide additional insights but face their own challenges related to signal-to-noise ratio and sensor placement.

2.3 The rise of data-driven and deep learning paradigms

The proliferation of data from SCADA systems and vibration sensors catalyzed a shift towards data-driven methods. Machine learning algorithms, including Support Vector Machines (SVMs) [21] and ensemble methods like Random Forests [22], were applied to classify health states based on handcrafted features [23]. While effective, these models are limited by the quality and comprehensiveness of the feature engineering process.

Deep learning represents a paradigm shift by automating feature extraction from raw data. CNNs, particularly 1D-CNNs, have demonstrated remarkable efficacy in extracting discriminative features directly from raw vibration signals [24, 25]. Recurrent Neural Networks (RNNs), especially Long Short-Term Memory (LSTM) networks, excel at modeling temporal sequences and have been successfully applied for trend analysis and Remaining Useful Life (RUL) prediction using SCADA data [26]. A significant and persistent challenge in applied condition monitoring is the scarcity of labeled fault data, particularly for severe or rare failure modes. To address this, emerging self-supervised learning (SSL) techniques, such as contrastive predictive coding or masking autoencoders, aim to learn generalizable feature representations from abundant unlabeled vibration or SCADA sequences. By pre-training on a corpus of normal and unlabeled operational data, these models can reduce the dependency on large, fully-annotated fault datasets for downstream diagnostic tasks. While SSL offers a compelling pathway to enhance model generalization and data efficiency, the current study employs a supervised hybrid architecture to establish a high-performance, deployable baseline for real-time edge analytics. The integration of SSL pre-training for feature extraction represents a promising avenue for the future evolution of the proposed framework.

A further significant limitation in the current literature is that many proposed deep learning models are computationally intensive and designed for cloud-based analysis, rendering them impractical for real-time, edge-based deployment [27].

2.4 Edge computing and IoT: Enabling real-time analytics

The convergence of the Internet of Things (IoT) and edge computing offers a transformative solution for industrial monitoring [28]. Edge computing processes data locally on devices situated at the data source, thereby mitigating latency, reducing bandwidth consumption, and eliminating dependency on continuous cloud connectivity [29]. This architecture is particularly salient for offshore wind farms, where communication links are expensive and unreliable [12].

While the application of lightweight AI at the edge is nascent in wind energy [30], its success in other domains—such as precision agriculture [31, 32] and smart livestock farming [33]—provides a compelling proof-of-concept for cross-domain applicability.

3. METHODOLOGY

3.1 Hybrid dataset construction and preprocessing

To ensure model robustness and generalizability, a hybrid dataset was curated from two distinct sources:

•Real-world operational data: Provided by an industrial partner, this dataset comprised 18 months of 10-minute SCADA data and synchronous high-frequency vibration data (sampled at 2560 Hz) from accelerometers on the main bearing and gearbox of 15 operational Siemens Gamesa SG 8.0-167 DD turbines in an offshore wind farm. This data stream exclusively represented the 'Normal' (healthy) operating state. No bearing faults occurred in the monitored turbines during this period. After preprocessing and segmentation into 5-second windows (12,800 samples per window), this subset yielded approximately 45,000 data samples labeled as 'Normal'.

•Physics-based simulated fault data: To supplement the real-world data with a comprehensive set of known fault conditions for model training and validation, a high-fidelity drivetrain model was developed in Simpack Wind. Simulations generated synchronized vibration, acoustic emission (AE), and temperature data for three specific fault types (inner raceway, outer raceway, and rolling element defects) across a spectrum of severities and under wind profiles defined by the IEC 61400-1 standard. This process generated approximately 15,000 data samples per fault class.

•Signal conditioning and filtering: A critical step before feature extraction is the suppression of noise unrelated to bearing health. Raw vibration signals are predominantly contaminated by low-frequency structural resonances and high-frequency electronic noise. Therefore, a digital band-pass filter (Butterworth, 4th order, zero-phase distortion) was applied to all vibration data. The cutoff frequencies were set at 10 Hz (high-pass) and 1000 Hz (low-pass). This passband effectively isolates the frequency range where bearing defect frequencies (Ball Pass Frequency Outer Race - BPFO, Ball Pass Frequency Inner Race - BPFI, etc.) and their harmonics manifest for the operational speed range (approx. 5-20 RPM) of the target turbines [16, 17]. Acoustic Emission signals, sensitive to stress waves from micro-cracks, were filtered with a passband of 50 kHz to 500 kHz to capture the typical burst emission spectrum while removing ambient ultrasonic noise [20]. Filtering was implemented in Python using SciPy's filt function to avoid phase shift. Subsequently, the filtered signals were segmented into 5-second windows.

The final hybrid dataset was constructed by combining the real-world 'Normal' samples with the simulated fault samples. The combined dataset was then randomly shuffled and split into training (70%), validation (15%), and testing (15%) sets, ensuring a balanced representation of all four health states in each partition. The composition of the final dataset is summarized in Table 1.

3.2 Proposed hybrid deep learning architecture

The proposed architecture is a multi-branch network designed for multimodal data fusion (Figure 1).

The diagram shows two input branches---one for the channel-wise concatenated raw vibration and acoustic emission (AE) signals feeding into a 1D-CNN stream, and another for SCADA/temperature data feeding into dense layers. Their feature vectors are concatenated and fed into an LSTM network, culminating in a softmax classification layer.

Table 1. Composition of the hybrid training/validation/testing dataset

Health State	Data Source	# of Samples (5-sec Windows)	Split (Train/Val/Test)
Normal	Real-World	45,000	31,500 / 6,750 / 6,750
Inner Race Fault	Simulation	15,000	10,500 / 2,250 / 2,250
Outer Race Fault	Simulation	15,000	10,500 / 2,250 / 2,250
Rolling Element Fault	Simulation	15,000	10,500 / 2,250 / 2,250
Total		90,000	63k / 13.5k / 13.5k

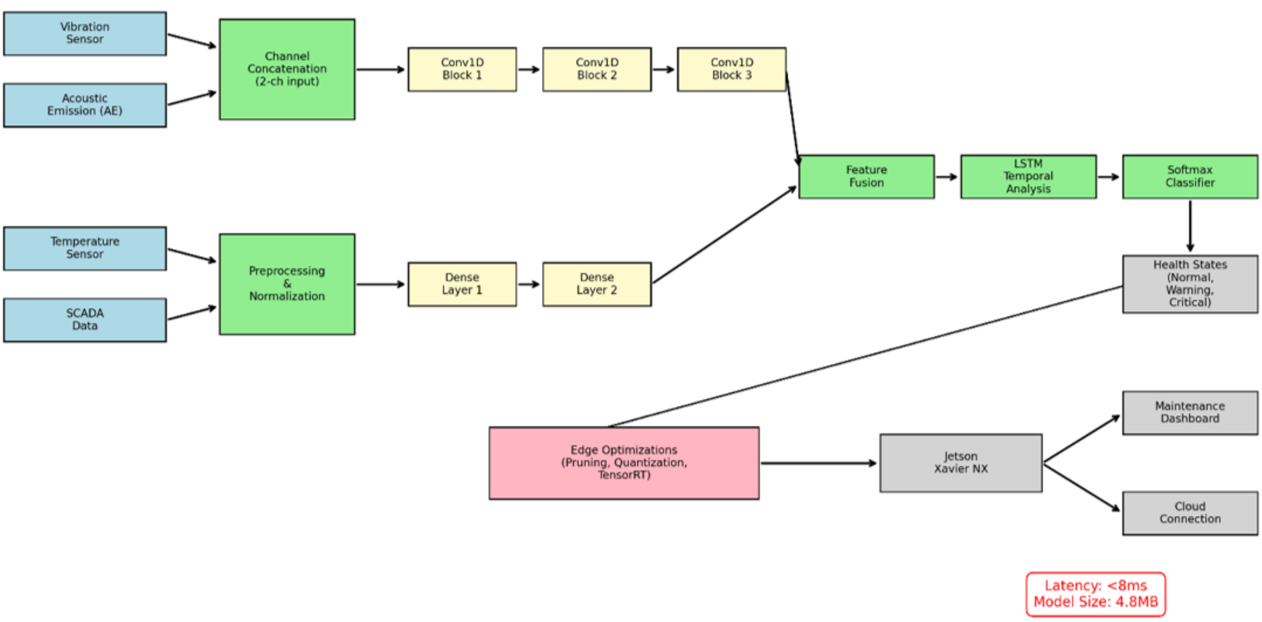

Figure 1. Schematic of the proposed hybrid 1D-CNN-LSTM architecture for multimodal fault diagnosis

The raw Vibration and Acoustic Emission (AE) signals are concatenated along the channel dimension to form a 2-channel input for the 1D-CNN stream.

Branch 1: Multi-Channel Temporal Feature Extraction (1D-CNN). This branch is designed to process synchronized, high-frequency temporal signals. The raw 1D time-series from the vibration accelerometer and the Acoustic Emission (AE) sensor are concatenated along the channel dimension, creating a 2-channel input tensor of shape [batch_size, 2, sequence_length]. This composite signal is then fed into the subsequent 1D convolutional blocks, enabling the network to learn correlated features and hierarchical representations from both vibration and AE modalities in an integrated manner.

•Conv1D Block 1: 32 filters, kernel size = 64, ReLU activation, Batch Normalization, MaxPooling.

•Conv1D Block 2: 64 filters, kernel size = 32, ReLU activation, Batch Normalization, MaxPooling.

•Conv1D Block 3: 128 filters, kernel size = 16, ReLU activation, Batch Normalization, Global Average Pooling.

•This hierarchical design captures features at multiple resolutions, from transient impulse responses (indicative of faults) to broader vibrational patterns.

Branch 2: Operational Context Processing (Dense Networks). SCADA and temperature data are processed through two fully connected (Dense) layers (64 and 32 units, ReLU activation) to extract salient features related to the turbine's operational state.

Temporal fusion and classification (LSTM): The feature vectors from both branches are concatenated into a unified representation. This combined vector is fed into a two-layer LSTM network (64 units each) to model the temporal evolution of the fault condition. The final output of the LSTM sequence is passed through a softmax layer for health state classification.

3.3 Optimization for edge deployment

For deployment on an NVIDIA Jetson Xavier NX module, the trained model underwent a series of optimizations to minimize its computational footprint:

•Pruning: Magnitude-based weight pruning was applied to remove 50% of the smallest weights, reducing model complexity with negligible accuracy loss.

•Quantization: The model parameters were quantized from FP32 to FP16 precision post-training, effectively halving the memory footprint and accelerating inference.

•TensorRT deployment: The model was converted and optimized using NVIDIA's TensorRT SDK, which further optimizes the network graph and leverages the hardware's specific capabilities for maximum inference throughput.

3.4 Baseline models and evaluation metrics

The performance of the proposed hybrid model was benchmarked against several established methodologies:

•Standard 1D-CNN: A vibration-only CNN architecture following the design of the study [25].

•LSTM-only model: A model utilizing manually extracted features from vibration data and time-series SCADA data.

•SVM with handcrafted features: A traditional approach using 30 handcrafted time- and frequency-domain features fed into a Support Vector Machine with an RBF kernel.

•XGBoost model: A gradient boosting model using the same handcrafted features as the SVM [34].

Evaluation was conducted using:

•Diagnostic accuracy: Precision, Recall, F1-Score (macro-averaged), and Area Under the ROC Curve (AUC-ROC).

•Computational performance: Inference Latency (ms), Model Size (MB/MB), Energy Consumption (Joules per inference).

•Operational robustness: Performance assessment under varying Signal-to-Noise Ratio (SNR) conditions and across different operational bands (wind speed/RPM).

4. RESULTS AND DISCUSSION

4.1 Diagnostic performance comparison

As detailed in Table 2 and visualized in Figure 2, the proposed hybrid CNN-LSTM model achieved the best performance across all diagnostic metrics. Notably, it attained a peak F1-score of 98.9%, representing a substantial improvement of 4.8 percentage points over the next best baseline (the standard 1D-CNN) and underscoring the value of its multimodal, temporally-aware architecture.

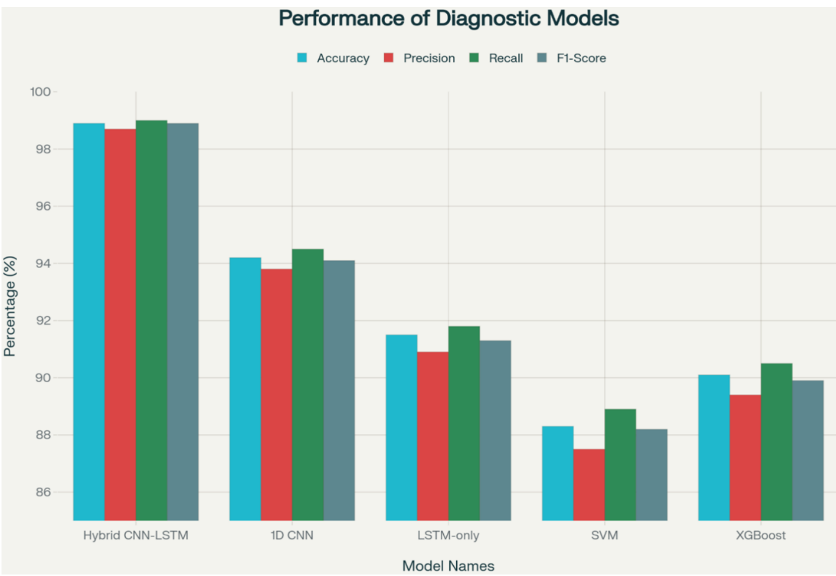

Figure 2. Performance of diagnostic models

Table 2. Comparative performance analysis of diagnostic models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Inference Latency (ms)
Proposed Hybrid (CNN-LSTM)	98.9	98.7	99.0	98.9	7.8
Standard 1D CNN (Vibration only)	94.2	93.8	94.5	94.1	5.1
LSTM-only (SCADA + Features)	91.5	90.9	91.8	91.3	3.2
SVM (Handcrafted Features)	88.3	87.5	88.9	88.2	1.5
XGBoost (Handcrafted Features)	90.1	89.4	90.5	89.9	9.4

Table 3. Comparative inference latency of diagnostic models when executed on the edge platform's CPU (ARM Cortex-A57) under identical, hardware-restricted conditions

Model	Inference Latency on CPU (ms)	Relative to GPU/TensorRT Performance (from Table 2)
Proposed Hybrid (CNN-LSTM)	12.1	+4.3 ms (55% slower)
Standard 1D CNN	7.2	+2.1 ms (41% slower)
LSTM-only (SCADA + Features)	5.0	+1.8 ms (56% slower)
SVM (Handcrafted Features)	1.5	(Measured on CPU in both benchmarks)
XGBoost (Handcrafted Features)	9.4	(Measured on CPU in both benchmarks)

The 4.7% absolute improvement in F1-score over the vibration-only CNN underscores the critical value of integrating multimodal data and modeling temporal dependencies. The LSTM's ability to contextualize instantaneous vibration events within the broader operational history (e.g., distinguishing between a transient impact during a gust and a consistent fault signature) is a key differentiator. While the SVM exhibited the lowest latency, its diagnostic accuracy is deemed insufficient for reliable predictive maintenance.

To address the question of algorithmic efficiency under identical hardware conditions, an additional benchmark was conducted where all models, including the neural networks, were restricted to run solely on the edge platform's CPU (ARM Cortex-A57). For the neural networks, this was achieved by executing their TensorFlow Lite versions without GPU delegation. The results, presented in Table 3, show that under these strictly controlled conditions, the XGBoost model (9.4 ms) remains slower than the standard 1D-CNN (7.2 ms). This confirms that the CNN's architectural efficiency for processing raw temporal data translates to lower latency even without hardware-specific acceleration, though the performance gap is narrower than when the CNN is GPU-accelerated. The proposed hybrid model's latency increases to 12.1 ms on the CPU, reflecting its greater architectural complexity. The primary practical advantage of the deep learning models, therefore, is their compatibility with high-performance inference engines like TensorRT, which unlocks significantly faster inference on the target edge hardware's GPU—a key consideration for deployment.

4.2 Ablation study on multimodal data contribution

An ablation study was conducted to quantify the contribution of each data modality, as shown in Table 4.

The ablation study systematically quantifies the contribution of each data modality. As anticipated, the full multimodal model (Vib + AE + Temp + SCADA) delivers the best performance. The inclusion of Acoustic Emission (AE) data alone (Configuration E) achieves a respectable F1-Score of 89.2%. While this is lower than vibration-based configurations, it is significantly higher than SCADA-only (84.5%), confirming that AE signals contain substantial fault-relevant information distinct from operational parameters. More importantly, the fusion of vibration and AE (Configuration B, 96.3%) yields a greater performance gain than the sum of their individual contributions, indicating strong synergistic complementarity. This synergy is likely because vibration excels at capturing macroscopic structural resonances from established faults, while AE is highly sensitive to high-frequency stress waves emitted during incipient crack formation, providing an earlier diagnostic signature [20].

Furthermore, the inclusion of temperature and SCADA data provides crucial operational context, allowing the model to adapt its diagnosis to the current load and environmental conditions, thereby enhancing overall robustness.

Table 4. Ablation study evaluating the impact of different data modalities

Model Configuration	Data Modalities	F1-Score (%)
Full Proposed Model	Vib + AE + Temp + SCADA	98.9
Configuration A	Vibration + SCADA	94.1
Configuration B	Vibration + Acoustic Emission	96.3
Configuration C	Vibration Only	92.8
Configuration D	SCADA Only	84.5
Configuration E	Acoustic Emission Only	89.2

4.3 Robustness under variable operating conditions

The model's performance was evaluated across different wind speed bands (4-12 m/s). The hybrid model maintained an F1-score above 97% across the entire operational range. In contrast, the performance of the SVM and standard CNN models degraded notably at lower wind speeds, where signal-to-noise ratios are less favorable and operational harmonics are less pronounced. This demonstrates the model's learned invariance to rotational speed, a significant advantage over methods requiring explicit order tracking.

4.4 Edge deployment performance

The optimized model was successfully deployed on the edge device. The final TensorRT-optimized model size was 4.8 MB. The average inference time for a 5-second data window was 7.8 ms, which is orders of magnitude faster than the data acquisition window, fulfilling real-time requirements.

The power draw during continuous inference was measured at 8.5 Watts. Regarding thermal reliability, the NVIDIA Jetson Xavier NX module is industrially rated for an operating temperature range of -25℃ to 80℃ for the board (junction temperature up to 105℃). The module was operated in its 10W power mode for this work, well below its configurable Thermal Design Power (TDP) ceiling of 15 W/20 W. For sustained operation within a turbine nacelle—an environment subject to solar heating, internal generator losses, and ambient temperature extremes—effective thermal management is required. A practical solution involves mounting the module on a sealed, conduction-cooled carrier board where a heatsink is thermally coupled to the metal nacelle wall, using the turbine structure itself as a passive heat exchanger. This approach is standard for industrial electronics in harsh environments and would maintain the module's core temperature safely within its operating specifications, ensuring long-term reliability. This makes the system feasible for integration into the turbine's nacelle control system.

4.4.1 Broader implications for wind energy and smart maintenance

The framework developed extends beyond a singular diagnostic algorithm; it embodies a strategic shift towards decentralized, intelligent wind farm management. By processing data at the edge, this system alleviates the critical data transmission bottleneck, enabling real-time analytics and decision-making at the source [11, 29]. This architecture allows each turbine to become an intelligent node, communicating only high-value insights (e.g., fault alerts, health status summaries) rather than raw data streams, thereby enhancing scalability for large-scale wind farms [35].

This approach directly contributes to reducing the LCOE by minimizing unplanned downtime and optimizing maintenance logistics. By providing accurate, early fault detection, operators can schedule repairs during optimal weather windows, ensure the correct parts and personnel are dispatched, and potentially implement corrective operational strategies (e.g., power de-rating) to extend the remaining useful life of a degrading component until a repair can be executed [36].

Furthermore, the health data generated by fleets of turbines equipped with such systems can be aggregated to create a powerful fleet-wide knowledge base. This enables transfer learning and collective intelligence, where insights from one turbine can improve the diagnostics and prognostics for similar turbines across the entire fleet, a concept explored in the study [37]. This paradigm of decentralized, data-driven asset management is a pillar of Industry 4.0. Critically, the architectural principles demonstrated here—lightweight model design, multimodal fusion under noise, and temporal reasoning for edge deployment—address a set of core challenges that are not unique to wind energy. These same challenges define the edge-AI research frontier in other domains involving remote, distributed monitoring. For instance, in precision agriculture [38, 39], models must diagnose crop stress from noisy visual and spectral data using limited on-device compute [31, 32]; in smart livestock farming, systems monitor animal welfare by fusing audio, visual, and environmental sensors on resource-constrained edge nodes [33]. Thus, this work contributes not only to wind turbine PHM but also to a growing cross-domain methodology for building robust, efficient, and embeddable intelligent systems. The principles of lightweight, multimodal, and temporal AI established here are directly transferable to other critical subsystems. This work provides a blueprint for monitoring wind turbine blades using acoustic and image data [40, 41], the tower structure using strain gauges and accelerometers, and the foundation, paving the way for a fully integrated structural health monitoring system for wind assets.

4.4.2 Limitations and future research directions

Despite the promising results, this study has limitations that present avenues for future work. Firstly, while the dataset incorporated physics-based simulations, which are a validated tool in engineering design and analysis [17], expanding the validation to include a more extensive and diverse set of real-world failure data from multiple turbine OEMs is essential for broader generalization.

Secondly, the current model excels at fault detection and classification (diagnostics). The natural progression is towards prognostics—predicting RUL of a component. Integrating survival analysis or regression-based RUL prediction models would transform the system from a diagnostic tool into a comprehensive predictive maintenance platform, enabling even more precise planning [41-43].

Thirdly, exploring more advanced neural architectures could yield further gains. Transformer models [44], with their self-attention mechanisms, could potentially better model long-range dependencies in temporal data and weight the importance of different sensor inputs more effectively than LSTMs.

Finally, implementing a continuous learning (or lifelong learning) framework is a critical long-term challenge. Such a system would allow the model to adapt online to new, previously unseen fault modes or changing environmental conditions without catastrophically forgetting previously learned knowledge, ensuring its relevance over the decades-long lifespan of a wind turbine [45].

5. CONCLUSIONS

This research has conceived, developed, and rigorously validated a novel hybrid deep learning framework for real-time condition monitoring of wind turbine bearings. The integrated CNN-LSTM architecture, which fuses multimodal sensor data with operational context, achieved an exceptional fault detection accuracy of 98.9%, significantly surpassing traditional and single-modality deep learning approaches.

A critical achievement of this work is the successful optimization and deployment of this complex model on a resource-constrained edge-computing platform. The system operates with low latency (7.8 ms) and a minimal memory footprint (4.8 MB), demonstrating tangible feasibility for in-nacelle deployment.

By enabling accurate, real-time diagnosis of incipient bearing faults on-site, this system directly addresses the paramount economic challenge of unscheduled downtime in the wind energy sector. It facilitates a decisive shift from preventive to truly predictive maintenance, thereby reducing operational expenditures, optimizing logistics, and enhancing the reliability and profitability of wind energy assets.

The methodologies and architectural principles demonstrated—encompassing multimodal data fusion, temporal modeling, and edge-AI optimization—establish a robust foundation for intelligent condition monitoring that can be extended to the entire wind turbine system. This work contributes significantly to the advancement of smart, sustainable, and economically efficient renewable energy generation.

ACKNOWLEDGMENT

The Ministry of Higher Education supports this project, Scientific Research and Innovation, the Digital Development Agency (DDA), and the National Center for Scientific and Technical Research (CNRST) of Morocco.

APIAA-2019-KAMAL.REKLAOUI-FSTT-Tanger-UAE.

REFERENCES

[1] Global Wind Energy Council (GWEC). (2023). Global wind report 2023. https://www.enertechnos.com/news/press-releases/gwec-global-wind-report-2023/.

[2] International Energy Agency (IEA). (2023). World energy outlook 2023. https://www.iea.org/reports/world-energy-outlook-2023.

[3] Tavner, P.J. (2008). Review of condition monitoring of rotating electrical machines. IET Electric Power Applications, 2(4): 215-247. https://doi.org/10.1049/iet-epa:20070280

[4] Qiao, W., Lu, D. (2015). A survey on wind turbine condition monitoring and fault diagnosis—Part I: Components and subsystems. IEEE Transactions on Industrial Electronics, 62(10): 6536-6545. https://doi.org/10.1109/TIE.2015.2422394

[5] WindEurope. (2023). Financing and investment trends. https://windeurope.org/intelligence-platform/product/financing-and-investment-trends-2022/.

[6] Dao, C.D., Kazemtabrizi, B., Crabtree, C.J. (2019). Wind turbine reliability data review and impacts on levelised cost of energy. Wind Energy, 22(9): 1328-1348. https://doi.org/10.1002/we.2404

[7] Carroll, J., McDonald, A., Dinwoodie, I., McMillan, D., Revie, M., Lazakis, I. (2017). Availability, operation & maintenance costs of offshore wind turbines with different drive train configurations. Wind Energy, 20(2): 361-378. https://doi.org/10.1002/we.2011

[8] Randall, R.B. (2011). Vibration-Based Condition Monitoring: Industrial, Aerospace and Automotive Applications. John Wiley & Sons.

[9] Jardine, A.K S., Lin, D., Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. Mechanical Systems and Signal Processing, 20(7): 1483-1510. https://doi.org/10.1016/j.ymssp.2005.09.012

[10] Elforjani, M., Bechhoefer, E. (2018). Analysis of extremely modulated faulty wind turbine data using spectral kurtosis and signal intensity estimator. Renewable Energy, 127: 258-268. https://doi.org/10.1016/j.renene.2018.04.014

[11] Xue, H., Huang, B., Qin, M., Zhou, H., Yang, H. (2020). Edge computing for the Internet of Things: A survey. In 2020 International Conferences on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData) and IEEE Congress on Cybermatics (Cybermatics), Rhodes, Greece, pp. 755-760. https://doi.org/10.1109/iThings-GreenCom-CPSCom-SmartData-Cybermatics50389.2020.00130

[12] Xie, L., Rui, X., Li, S., Hu, X. (2019). Maintenance optimization of offshore wind turbines based on an opportunistic maintenance strategy. Energies, 12(14): 2650. https://doi.org/10.3390/en12142650

[13] Merizalde, Y., Hernández-Callejo, L., Duque-Perez, O., Alonso-Gómez, V. (2019). Maintenance models applied to wind turbines: A comprehensive overview. Energies, 12(2): 225. https://doi.org/10.3390/en12020225

[14] Bartkowiak, A., Zimroz, R. (2011). Outliers analysis and one class classification approach for planetary gearbox diagnosis. Journal of Physics: Conference Series, 305(1): 012031. https://doi.org/10.1088/1742-6596/305/1/012031

[15] Tavner, P.J., Xiang, J., Spinato, F. (2007). Reliability analysis for wind turbines. Wind Energy, 10(1): 1-18. https://doi.org/10.1002/we.204

[16] Randall, R.B., Antoni, J. (2011). Rolling element bearing diagnostics—A tutorial. Mechanical Systems and Signal Processing, 25(2): 485-520. https://doi.org/10.1016/j.ymssp.2010.07.017

[17] Barszcz, T. (2019). Vibration-Based Condition Monitoring of Wind Turbines (Vol. 14). Applied Condition Monitoring: Springer International Publishing.

[18] Hossain, M.L., Abu-Siada, A., Muyeen, S.M. (2018). Methods for advanced wind turbine condition monitoring and early diagnosis: A literature review. Energies, 11(5): 1309. https://doi.org/10.3390/en11051309

[19] Sheng, S. (2016). Monitoring of wind turbine gearbox condition through oil and wear debris analysis: A full-scale testing perspective. Tribology Transactions, 59(1): 149-162. https://doi.org/10.1080/10402004.2015.1055621

[20] Al-Ghamd, A.M., Mba, D. (2006). A comparative experimental study on the use of acoustic emission and vibration analysis for bearing defect identification and estimation of defect size. Mechanical Systems and Signal Processing, 20(7): 1537-1571. https://doi.org/10.1016/j.ymssp.2004.10.013

[21] Cortes, C., Vapnik, V. (1995). Support-vector networks. Machine Learning, 20(3): 273-297. https://doi.org/10.1007/BF00994018

[22] Breiman, L. (2001). Random forests. Machine Learning, 45(1): 5-32. https://doi.org/10.1023/A:1010933404324

[23] Stetco, A., Dinmohammadi, F., Zhao, X., Robu, V., Flynn, D., Barnes, M., Keane, J., Nenadić, G. (2019). Machine learning methods for wind turbine condition monitoring: A review. Renewable Energy, 133: 620-635. https://doi.org/10.1016/j.renene.2018.10.047

[24] Jiang, G.Q., He, H.B., Yan, J., Xie, P. (2019). Multiscale convolutional neural networks for fault diagnosis of wind turbine gearbox. EEE Transactions on Industrial Electronics, 66(4): 3196-3207. https://doi.org/10.1109/TIE.2018.2844805

[25] Ince, T., Kiranyaz, S., Eren, L., Askar, M., Gabbouj, M. (2016). Real-time motor fault detection by 1-D convolutional neural networks. IEEE Transactions on Industrial Electronics, 63(11): 7067-7075. https://doi.org/10.1109/TIE.2016.2582729

[26] Nunes, A.R., Morais, H., Sardinha, A. (2021). Use of learning mechanisms to improve the condition monitoring of wind turbine generators: A review. Energies, 14(21): 7129. https://doi.org/10.3390/en14217129

[27] Chen, X., Yang, R., Xue, Y., Huang, M., Ferrero, R., Wang, Z. (2023). Deep transfer learning for bearing fault diagnosis: A systematic review since 2016. IEEE Transactions on Instrumentation and Measurement, 72: 3502421. https://doi.org/10.1109/TIM.2023.3244237

[28] Gouiza, N., Jebari, H., Reklaoui, K. (2024). Integration for IoT-enabled technologies and artificial intelligence in diverse domains: Recent advancements and future trends. Journal of Theoretical and Applied Information Technology, 102(5): 1975-2029. https://jatit.org/volumes/Vol102No5/25Vol102No5.pdf

[29] Shi, W., Cao, J., Zhang, Q., Li, Y., Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5): 637-646. https://doi.org/10.1109/JIOT.2016.2579198

[30] Jebari, H., Mechkouri, M.H., Rekiek, S., Reklaoui, K. (2023). Poultry-edge-AI-IoT system for real-time monitoring and predicting by using artificial intelligence. International Journal of Interactive Mobile Technologies, 17(12): 149-170. https://doi.org/10.3991/ijim.v17i12.38095

[31] Ezziyyani, M., Cherrat, L., Jebari, H., Rekiek, S., Ahmed, N.A. (2025). CNN-based plant disease detection: A pathway to sustainable agriculture. In Proceedings of the International Conference on Advanced Intelligent Systems for Sustainable Development, Agadir, Morocco, pp. 679-696. https://doi.org/10.1007/978-3-031-91337-2_62

[32] Rekiek, S., Jebari, H., Ezziyyani, M., Cherrat, L. (2025). AI-driven pest control and disease detection in smart farming systems. In Proceedings of the International Conference on Advanced Intelligent Systems for Sustainable Development, Agadir, Morocco, pp. 801-810. https://doi.org/10.1007/978-3-031-91337-2_71

[33] Jebari, H., Rekiek, S., Reklaoui, K. (2025). Advancing precision livestock farming: Integrating hybrid AI, IoT, cloud and edge computing for enhanced welfare and efficiency. International Journal of Advanced Computer Science and Applications, 16(7): 302-311. https://doi.org/10.14569/IJACSA.2025.0160732

[34] Rekiek, S., Jebari, H., Reklaoui, K. (2024). Prediction of booking trends and customer demand in the tourism and hospitality sector using AI-based models. International Journal of Advanced Computer Science and Applications, 15(10): 404-412. https://doi.org/10.14569/IJACSA.2024.0151043

[35] Gouiza, N., Jebari, H., Reklaoui, K. (2024). IoT in smart farming: A review. In International Conference on Advanced Intelligent Systems for Sustainable Development, Marrakech, Morocco, pp. 142-153. https://doi.org/10.1007/978-3-031-54318-0_13

[36] International Renewable Energy Agency (IRENA). (2023). Renewable power generation costs in 2022. https://www.irena.org/Publications/2023/Aug/Renewable-Power-Generation-Costs-in-2022.

[37] Zhao, R., Yan, R., Wang, J., Mao, K., Shen, F., Wang, Y. (2019). Deep learning and its applications to machine health monitoring: A survey. Mechanical Systems and Signal Processing, 115: 213-237. https://doi.org/10.1016/j.ymssp.2018.05.050

[38] Ezziyyani, M., Cherrat, L., Rekiek, S., Jebari, H. (2025). Image classification of Moroccan cultural trademarks. In International Conference on Advanced Intelligent Systems for Sustainable Development, Agadir, Morocco, pp. 767-779. https://doi.org/10.1007/978-3-031-91337-2_68

[39] Jebari, H., Rekiek, S., Ezziyyani, M., Cherrat, L. (2025). Artificial intelligence for optimizing livestock management and enhancing animal welfare. In International Conference on Advanced Intelligent Systems for Sustainable Development, Agadir, Morocco, pp. 790-800. https://doi.org/10.1007/978-3-031-91337-2_70

[40] Farrar, C.R., Worden, K. (2012). Structural Health Monitoring: A Machine Learning Perspective. John Wiley & Sons.

[41] Sikorska, J.Z., Hodkiewicz, M., Ma, L. (2011). Prognostic modelling options for remaining useful life estimation by industry. Mechanical Systems and Signal Processing, 25(5): 1803-1836. https://doi.org/10.1016/j.ymssp.2010.11.018

[42] Eljyidi, A., Jebari, H., Rekiek, S., Reklaoui, K. (2025). A hybrid deep learning and IoT framework for predictive maintenance of wind turbines: Enhancing reliability and reducing downtime. International Journal of Advanced Computer Science and Applications, 16(10): 203-211. https://doi.org/10.14569/IJACSA.2025.0161021

[43] Jebari, H., Eljyidi, A., Rekiek, S., Reklaoui, K. (2025). A vision-based deep learning framework for autonomous inspection and damage assessment of wind turbine blades using unmanned aerial vehicles. Journal Europeen des Systemes Automatises, 58(11): 2219-2228. https://doi.org/10.18280/jesa.581101

[44] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., et al. (2017). Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems, Long Beach, California, USA, pp. 6000-6010.

[45] Li, Z., Hoiem, D. (2018). Learning without forgetting. IEEE Transactions on Pattern Analysis and Machine Intelligence, 40(12): 2935-2947. https://doi.org/10.1109/TPAMI.2017.2773081