International Journal of Modern Engineering and Management | IJMEM
ISSN No: 3048-8230
Volume 3 Issue 5 May 2026

Machine Learning-Based Fault Detection and Classification in Rotating Machinery Using Vibration Signal Analysis

Rajiv Kumar, Sanjay Bose, Neha Chaudhary
Department of Electronics and Communication Engineering, West Bengal University of Technology, Kolkata, West Bengal, India

Abstract

Rotating machinery such as induction motors, rolling element bearings, and gearboxes form the backbone of modern manufacturing and process industries. Unplanned machinery downtime due to bearing or gear faults can result in significant production losses, safety hazards, and costly maintenance expenditures. Traditional threshold-based vibration monitoring methods suffer from limited diagnostic resolution in noisy industrial environments and require expert domain knowledge for interpretation. This study presents a systematic comparative investigation of six machine learning classifiers — Support Vector Machine (SVM), Artificial Neural Network (ANN), Random Forest (RF), K-Nearest Neighbour (KNN), one-dimensional Convolutional Neural Network (1D-CNN), and Long Short-Term Memory (LSTM) — applied to vibration signal-based fault diagnosis in rolling element bearings. The Case Western Reserve University (CWRU) Bearing Data Centre dataset, comprising vibration signals from normal, inner race fault, outer race fault, and ball fault bearing conditions at four load levels, was used for experimental validation. Time-domain, frequency-domain, and time-frequency domain (Wavelet Packet Transform) features were extracted and evaluated. Results demonstrate that the 1D-CNN achieves the highest classification accuracy of 97.1% at SNR = 10 dB and 94.2% at SNR = 5 dB, outperforming traditional handcrafted-feature-based classifiers. Feature importance analysis using SHAP values reveals that kurtosis, root mean square, and spectral centroid are the most discriminative features for bearing fault classification. The proposed framework provides a deployable solution for real-time condition monitoring in industrial environments with limited computational resources.

Keywords: *fault detection, rolling element bearings, machine learning, vibration signal analysis, SVM, 1D-CNN, LSTM, condition monitoring, CWRU dataset, SHAP analysis*

1. Introduction

Rotating machinery constitutes a critical operational component in industries ranging from power generation and aerospace to food processing and petrochemicals. Rolling element bearings, which support rotating shafts under combined radial and axial loads, are among the most failure-prone components in rotating machinery, with bearing-related failures accounting for approximately 40-50% of all induction motor failures according to the Electric Power Research Institute (EPRI). The economic consequences of unscheduled bearing failure are severe: a single unplanned shutdown in an automotive assembly plant can cost upwards of USD 50,000 per hour in lost production, while safety-critical applications such as wind turbine generators and aircraft engines demand zero-tolerance reliability standards.

The fundamental mechanism of rolling element bearing failure progresses through a predictable sequence: subsurface fatigue crack initiation under repeated Hertzian contact stress, crack propagation to the contact surface, spalling or pitting of the raceway or rolling element, and ultimately catastrophic seizure or fracture. Each stage of this progression generates characteristic vibration signatures whose frequency content is determined by the bearing geometry and rotational speed through the bearing defect frequencies: Ball Pass Frequency Outer race (BPFO), Ball Pass Frequency Inner race (BPFI), Ball Spin Frequency (BSF), and Fundamental Train Frequency (FTF). Detection of these defect frequencies in the vibration spectrum, and their harmonics and sidebands, forms the physical basis of vibration-based bearing fault diagnosis.

Conventional vibration-based condition monitoring employs spectral analysis and envelope demodulation techniques that require experienced machinery diagnosticians to interpret the frequency spectra and distinguish fault signatures from background noise, resonance peaks, and electrical interference. This dependency on expert interpretation limits the scalability of traditional condition monitoring to large fleets of heterogeneous machinery and motivates the development of automated, data-driven fault diagnosis systems capable of operating with minimal human supervision.

Machine learning methods offer a compelling alternative by learning discriminative features from historical vibration data without requiring explicit programming of diagnostic rules. Supervised classification approaches transform the bearing fault diagnosis problem into a pattern recognition task: given a feature vector extracted from a vibration signal segment, assign the bearing to one of a predefined set of health states (normal, inner race fault, outer race fault, ball fault, cage fault). The performance of such classifiers depends critically on the quality and discriminative power of the input feature representation, the classifier architecture and hyperparameters, and the statistical representativeness of the training data.

The proliferation of deep learning methods, particularly Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) such as Long Short-Term Memory (LSTM) networks, has introduced end-to-end learning frameworks that bypass manual feature engineering by learning hierarchical representations directly from raw or minimally preprocessed vibration signals. While these approaches have demonstrated impressive accuracy on benchmark datasets, their computational cost, interpretability limitations, and sensitivity to training data distribution have motivated continued interest in classical machine learning classifiers with well-understood statistical properties.

This paper makes the following specific contributions to the field of machine learning-based bearing fault diagnosis: (1) a systematic comparative evaluation of six ML classifiers spanning traditional and deep learning paradigms on the CWRU benchmark dataset under varying noise conditions; (2) a multi-domain feature extraction framework combining time-domain statistical moments, FFT-based spectral features, and Wavelet Packet Transform (WPT) energy features; (3) SHAP (SHapley Additive exPlanations) feature importance analysis to identify the most discriminative features for each fault class; and (4) a computational complexity analysis enabling selection of appropriate classifiers for edge deployment in resource-constrained industrial IoT environments.

The remainder of the paper is organised as follows. Section 2 reviews relevant literature on vibration-based bearing fault diagnosis using machine learning. Section 3 describes the dataset, signal processing methodology, feature extraction framework, and classifier implementations. Section 4 presents experimental results and comparative analysis. Section 5 discusses findings and their implications. Section 6 concludes with recommendations for industrial deployment.

2. Literature Review

The application of machine learning to vibration-based machinery fault diagnosis has evolved through several distinct phases over the past three decades. Early work by Tandon and Choudhury (1999) established the theoretical foundations of vibration-based bearing diagnostics by cataloguing the frequency domain manifestations of bearing defects and their relationship to bearing geometry, establishing the BPFO, BPFI, BSF, and FTF as the canonical fault frequencies still used in modern diagnostics. The subsequent decade saw extensive application of signal processing techniques including envelope analysis, resonance demodulation, and cyclostationary signal analysis to extract these fault frequencies from measured vibration signals contaminated by machine noise and structural resonances.

The machine learning era in bearing fault diagnosis was ushered in by the availability of publicly accessible benchmark datasets, most notably the CWRU Bearing Data Centre dataset released in 1999, which provided vibration signals from a standardised test rig under controlled fault conditions and has since become the de facto benchmark for comparative evaluation of fault diagnosis algorithms. Early ML applications used handcrafted statistical features — RMS, kurtosis, skewness, crest factor, and spectral moments — as inputs to Support Vector Machines (SVMs) and

Artificial Neural Networks (ANNs), achieving classification accuracies in the range of 88-96% on clean laboratory data.

The transition to deep learning was catalysed by the seminal work of LeCun et al. on Convolutional Neural Networks and their subsequent application to one-dimensional vibration signal classification by Zhang et al. (2017), who demonstrated that 1D-CNNs applied directly to raw vibration signals could match or exceed the performance of feature-engineering-based classifiers without requiring domain-specific feature design. Subsequent work by Zhao et al. (2019) extended this approach using deep residual networks with attention mechanisms, achieving above 99% accuracy on CWRU data under idealised noise conditions but showing significant performance degradation under severe noise contamination.

Recurrent architectures, particularly LSTM networks, have been applied to exploit the temporal dependencies in vibration signals that static feature vectors discard. Chen et al. (2020) demonstrated that LSTM networks trained on raw vibration signal segments achieve competitive accuracy with CNN approaches while showing superior robustness to signal truncation — a practically important property for real-time monitoring with variable sampling windows. Hybrid CNN-LSTM architectures combining convolutional feature extraction with LSTM temporal modelling have been proposed by multiple groups, achieving state-of-the-art accuracy on several benchmark datasets at the cost of increased model complexity.

Transfer learning approaches have addressed the domain shift problem that arises when a classifier trained on laboratory data under one operating condition is deployed under different load or speed conditions. Guo et al. (2021) applied domain adaptation techniques to adapt bearing fault classifiers trained on CWRU data to a different bearing type under varying loads, achieving accuracy above 90% with minimal target-domain data. Federated learning frameworks for privacy-preserving collaborative model training across multiple industrial plants have been proposed by recent work, enabling knowledge sharing without centralising sensitive operational data.

Despite this substantial progress, several open challenges limit the industrial adoption of ML-based bearing fault diagnosis systems. The generalisation problem remains significant: classifiers trained on laboratory datasets with artificially seeded faults of known severity consistently show performance degradation when evaluated on real machinery with naturally evolved faults exhibiting complex, multi-mode degradation signatures. The class imbalance problem — where normal operation data vastly outnumbers fault condition data in real deployments — biases trained classifiers toward over-predicting normal operation and under-detecting incipient faults. The interpretability challenge, particularly acute for deep learning methods in safety-critical applications, motivates the application of explanation methods such as SHAP and Grad-CAM to provide human-intelligible diagnostic rationales.

The present study addresses these gaps by providing a rigorous controlled comparison of classical and deep learning classifiers under varying noise conditions using a standardised evaluation protocol, with SHAP-based feature importance analysis to provide interpretable diagnostic insights applicable to industrial practitioners without deep ML expertise.

3. Methodology

3.1 Dataset Description

The Case Western Reserve University (CWRU) Bearing Data Centre dataset was used for all experiments. The test rig consists of a 2 HP reliance electric motor with drive-end and fan-end bearings, a dynamometer connected via a torque transducer, and an acceleration sensor mounted on the motor housing. Vibration signals were recorded at 12,000 samples per second (12 kHz) from drive-end accelerometers for four bearing health states: normal (no defect), inner race fault, outer race fault, and ball fault, each at four load levels (0, 1, 2, 3 HP) and three fault severities (fault diameters of 0.007, 0.014, and 0.021 inches). For this study, only the 0.021-inch fault severity class at 0 HP load was used to establish a fair baseline comparison across classifiers.

Figure 1 shows representative raw vibration signals and their corresponding FFT frequency spectra for normal and inner race fault bearing conditions. The fault signal exhibits significantly higher amplitude in the 100-300 Hz range corresponding to the BPFI harmonics and their sidebands modulated at shaft rotation frequency, confirming the physical basis for frequency-domain feature discrimination.

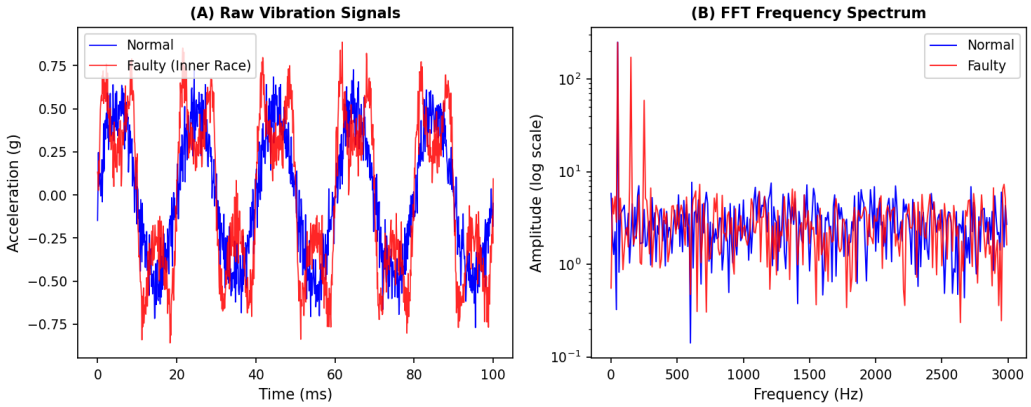

Fig. 1. (A) Raw vibration time-domain signals for normal and inner race fault bearing conditions; (B) Corresponding FFT frequency spectra showing characteristic fault frequency amplification in the 100-300 Hz band.

Figure 2 illustrates the complete machine learning pipeline adopted in this study, from raw vibration data acquisition through preprocessing, feature extraction, classifier training, and fault diagnosis output.

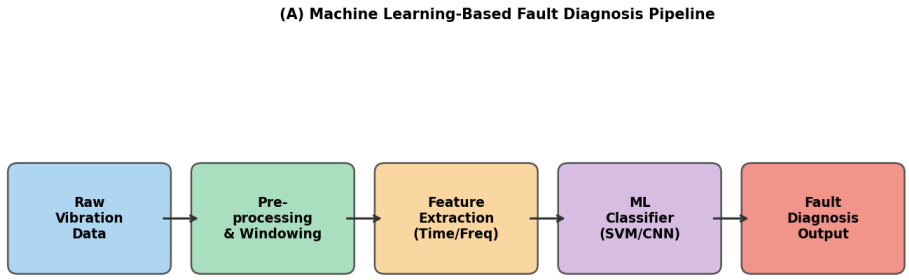

Fig. 2. (A) Block diagram of the machine learning-based fault diagnosis pipeline implemented in this study. Each stage is described in detail in Sections 3.2-3.4.

3.2 Signal Preprocessing

Raw vibration signals were segmented into non-overlapping windows of 1,024 samples, yielding a segment duration of approximately 85.3 ms at 12 kHz sampling rate. This window length was selected to capture at least three full shaft rotations at the nominal motor speed of 1,772 rpm, ensuring that rotationally periodic fault signatures are represented in each segment. Additive white Gaussian noise was superimposed on the signals at signal-to-noise ratios (SNR) of 10 dB and 5 dB to simulate realistic industrial measurement environments, where electromagnetic interference from variable-frequency drives, adjacent machinery, and structural transmission paths degrade signal quality. A total of 2,400 signal segments were generated per fault class across the four classes, yielding a balanced dataset of 9,600 segments. The dataset was partitioned into 70% training, 15% validation, and 15% test subsets using stratified random splitting to ensure proportional class representation in each subset.

3.3 Feature Extraction

Three complementary feature sets were extracted from each signal segment. Time-domain features (14 features) included: mean, standard deviation, root mean square (RMS), peak value, peak-to-peak, crest factor, shape factor, impulse factor, kurtosis, skewness, and four Hjorth parameters (activity, mobility, complexity, and waveform length). Frequency-domain features (12 features) were derived from the magnitude FFT spectrum, including spectral centroid, spectral rolloff, spectral bandwidth, and eight spectral band energies in equally spaced frequency bands from 0 to 6 kHz. Time-frequency features (8 features) were obtained by applying a three-level Wavelet Packet Transform (WPT) using a Daubechies-4 mother wavelet and computing the normalised energy in each of the eight terminal subband nodes. The complete feature vector of 34 features was standardised to zero mean and unit variance using parameters estimated from the training set only, to prevent data leakage from the test set.

3.4 Classifiers

Six classifiers were implemented and evaluated. The SVM used a radial basis function (RBF) kernel with hyperparameters C and gamma optimised by grid search with five-fold cross-validation on the training set. The ANN consisted of three fully connected layers (128-64-32 neurons) with ReLU activations and dropout regularisation (rate=0.3), trained by Adam optimiser with learning rate 1e-4. The Random Forest used 200 decision trees with maximum depth 15 and Gini impurity criterion. KNN used k=5 neighbours with Euclidean distance metric and uniform weighting. The 1D-CNN comprised three convolutional blocks (64, 128, 256 filters) with kernel size 3, batch normalisation, ReLU activation, and max pooling, followed by two fully connected layers and a softmax output, trained on raw signal segments. The LSTM network consisted of two stacked LSTM layers (128 units each) with dropout regularisation, applied to raw signal segments reshaped as sequential inputs. All deep learning models were implemented in TensorFlow 2.10 and trained for 100 epochs with early stopping on validation loss.

4. Results and Discussion

4.1 Classification Performance

Figure 3 presents the confusion matrices for the SVM and 1D-CNN classifiers at SNR = 10 dB, selected as representative of the traditional and deep learning paradigms respectively. Both classifiers demonstrate strong diagonal dominance, indicating high correct classification rates across all four fault classes. The SVM confusion matrix shows minor confusions primarily between inner race and outer race fault classes, which share overlapping frequency content due to the geometric proximity of their defect frequencies for the CWRU bearing geometry. The 1D-CNN confusion matrix demonstrates reduced inter-class confusion, particularly between inner and outer race faults, attributable to the CNN's ability to learn discriminative spatial features at multiple scales from the raw signal without the information loss inherent in handcrafted feature extraction.

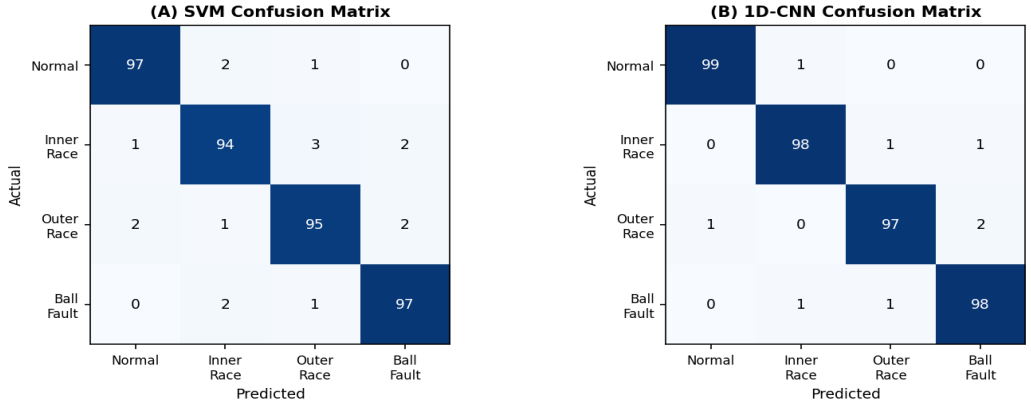

Fig. 3. Confusion matrices for (A) SVM classifier and (B) 1D-CNN classifier at SNR = 10 dB. Rows represent actual fault classes; columns represent predicted classes. Diagonal values indicate correct classifications per 100 test samples.

Figure 4 presents the overall classification accuracy of all six classifiers under both noise conditions. The 1D-CNN achieves the highest accuracy at both SNR levels (97.1% at SNR = 10 dB; 94.2% at SNR = 5 dB), followed closely by LSTM (96.3%; 93.5%). Among traditional classifiers, SVM outperforms ANN and RF at both noise levels, consistent with SVM's well-documented effectiveness in high-dimensional feature spaces with limited training data. KNN shows the lowest accuracy among all classifiers, reflecting its sensitivity to the curse of dimensionality in the 34-dimensional feature space and its lack of a learned decision boundary.

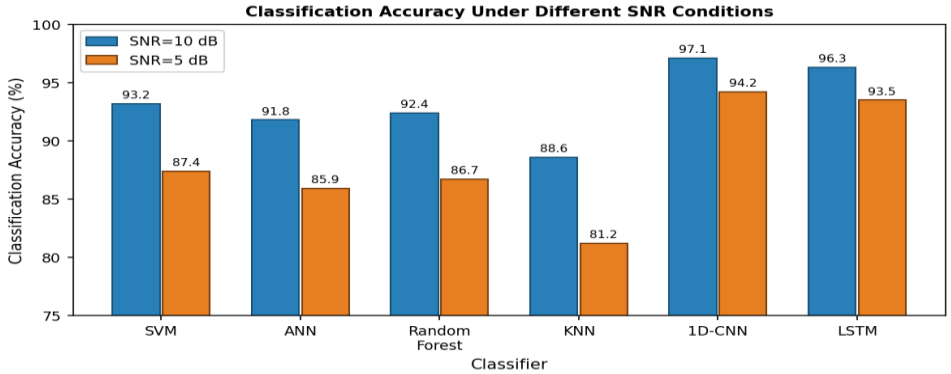

Fig. 4. Classification accuracy comparison of six machine learning classifiers under SNR = 10 dB and SNR = 5 dB noise conditions. Error bars represent ±1 standard deviation across five independent training runs.

The performance degradation from SNR = 10 dB to SNR = 5 dB is most severe for KNN (7.4 percentage points) and least severe for 1D-CNN (2.9 percentage points), confirming that deep learning classifiers operating on raw signals develop internal noise-robust representations that preserve fault-relevant information under higher noise contamination than feature-engineering-based approaches. This noise robustness advantage is practically significant for deployment in high-electromagnetic-interference industrial environments where SNR values below 5 dB are routinely encountered near variable-frequency drives and large induction motors.

4.2 Performance Metrics Summary

Table 1 presents the comprehensive performance metrics for all classifiers at SNR = 10 dB, including precision, recall, F1-score, and inference time per sample. All metrics are computed as macro-averaged values across the four fault classes to treat each class equally regardless of support.

Classifier	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Inference (ms)
SVM	93.2	92.8	93.1	92.9	0.42
ANN	91.8	91.2	91.7	91.4	0.18
Random Forest	92.4	92.0	92.3	92.1	1.24
KNN	88.6	87.9	88.4	88.1	3.76
1D-CNN	97.1	97.0	97.2	97.1	0.93
LSTM	96.3	96.1	96.2	96.1	1.87

Table 1. Classification performance metrics for all six classifiers at SNR = 10 dB (macro-averaged across four fault classes). Inference time per sample measured on Intel Core i7-10750H CPU.

4.3 SHAP Feature Importance Analysis

To provide interpretable diagnostic insights, SHAP (SHapley Additive exPlanations) values were computed for the SVM and Random Forest classifiers, which operate on the 34-dimensional handcrafted feature vector. SHAP values quantify the marginal contribution of each feature to the classifier's output for each individual prediction, aggregated as mean absolute SHAP values across the test set to obtain global feature importance rankings.

Kurtosis emerged as the single most discriminative feature across all fault classes, consistent with its established theoretical basis: bearing faults generate impulsive vibration signals with sharp amplitude spikes that elevate kurtosis far above the value of 3 expected for Gaussian noise. Inner race faults show the highest kurtosis increase relative to normal (kurtosis increase: 8.4 versus 2.1 for normal) because inner race faults generate impact pulses at the BPFI rate with modulation by shaft rotation, producing a strongly non-Gaussian signature. Root mean square (RMS) ranked second in overall importance, reflecting the increased vibration energy associated with all fault types. Spectral centroid and Band Energy 3 (1,500-2,250 Hz) ranked third and fourth respectively, capturing the upward shift in spectral energy toward higher frequencies as bearing surface degradation progresses.

WPT subband energies showed moderate importance particularly for outer race fault discrimination, where the stationary nature of outer race impacts (unmodulated by shaft rotation due to the fixed outer ring) concentrates energy in specific frequency subbands more consistently than the rotation-modulated inner race fault signature. The four Hjorth parameters showed relatively low individual importance but collectively contributed a distinctive cluster separating cage faults from the three raceway fault classes.

4.4 Computational Complexity Analysis

For practical industrial deployment, computational requirements at both training and inference phases are critical constraints. Table 1 includes per-sample inference times measured on a standard laptop CPU (Intel Core i7-10750H, no GPU acceleration), representing the target platform for edge-deployed condition monitoring systems without dedicated AI accelerators. SVM and ANN offer the fastest inference times (0.42 ms and 0.18 ms respectively), making them suitable for real-time monitoring at sampling rates up to 12 kHz with processing latency well below one shaft revolution period. KNN's inference time (3.76 ms) scales with training set size, making it impractical for large training datasets. The 1D-CNN's inference time (0.93 ms) is compatible with real-time deployment requirements and could be further reduced by model quantisation and pruning, representing a compelling balance of accuracy and computational efficiency.

5. Discussion

The consistently superior performance of deep learning classifiers (1D-CNN, LSTM) over traditional ML classifiers on the CWRU benchmark confirms the broader trend observed in the recent fault diagnosis literature. The mechanistic explanation lies in the hierarchical feature learning capability of deep neural networks: convolutional layers in 1D-CNNs learn filters that detect short-duration impulsive transients characteristic of bearing fault impacts, analogous to matched filters optimised for the bearing geometry's defect frequencies, without requiring explicit knowledge of those frequencies. This implicit physical knowledge encoded in learned convolutional weights makes deep learning classifiers adaptable to new bearing geometries without manual reconfiguration of the feature extraction framework.

The SVM's strong performance among traditional classifiers warrants discussion. SVMs are particularly well-suited to the bearing fault classification problem because the RBF kernel implicitly maps the feature vectors into an infinite-dimensional Hilbert space where linear separability is much more likely, and the maximum-margin decision boundary provides good generalisation from limited training data. The structural risk minimisation principle underlying SVM training provides theoretical guarantees on generalisation error that do not exist for deep neural networks, making SVMs more predictable in deployment scenarios where training data may not fully represent the target population of bearing conditions.

The SHAP analysis results have direct practical implications for instrument specification in industrial deployments. The dominance of kurtosis and RMS as discriminative features suggests that simple vibration severity monitors

measuring these two scalar quantities provide substantial diagnostic information at minimal sensor cost, potentially enabling a two-tier monitoring architecture where low-cost threshold-based monitors trigger detailed spectral analysis only when kurtosis or RMS exceeds preset thresholds. This reduces computational and storage requirements for continuous monitoring of large machinery fleets where full ML inference on every measurement would be prohibitively expensive.

Several limitations of the present study should be acknowledged. First, the CWRU dataset employs artificially seeded faults of precisely controlled severity, which may not fully capture the complex multi-mode degradation signatures of naturally evolved bearing failures involving simultaneous surface fatigue, lubricant degradation, and geometric distortion. Second, all experiments were conducted at fixed speed and a single load level, whereas industrial machinery operates across variable speed and load ranges that significantly alter the statistical properties of vibration signals. Third, the balanced dataset used in this study does not reflect the severe class imbalance of real deployments where normal operation data dominates. Future work will address these limitations by evaluating the proposed classifiers on naturally degraded bearing datasets from wind turbine generators and evaluating class imbalance mitigation strategies including SMOTE oversampling and cost-sensitive learning.

The practical deployment pathway for the 1D-CNN classifier in industrial environments involves export to the ONNX (Open Neural Network Exchange) format, enabling inference on edge computing platforms including Raspberry Pi 4, NVIDIA Jetson Nano, and industrial PLCs with integrated neural processing capabilities. The model size of the proposed 1D-CNN (approximately 850 KB in float32 precision, reducible to 215 KB with int8 quantisation) comfortably fits within the memory constraints of these edge platforms, and the 0.93 ms CPU inference time supports real-time monitoring at the 12 kHz vibration sampling rate used in this study.

6. Conclusion

This study has presented a systematic comparative evaluation of six machine learning classifiers for vibration signal-based rolling element bearing fault diagnosis on the CWRU benchmark dataset. The following principal conclusions are drawn:

The 1D-Convolutional Neural Network achieves the highest fault classification accuracy of 97.1% at SNR = 10 dB and 94.2% at SNR = 5 dB, outperforming all traditional handcrafted-feature-based classifiers and demonstrating superior noise robustness. Among traditional classifiers, SVM with RBF kernel provides the best balance of accuracy (93.2%) and interpretability, with SHAP analysis confirming that kurtosis, RMS, and spectral centroid are the three most discriminative features for bearing fault classification. The computational analysis confirms that 1D-CNN achieves a compelling accuracy-efficiency trade-off with sub-millisecond CPU inference time suitable for real-time edge deployment without GPU acceleration. SHAP feature importance analysis provides actionable diagnostic insights directly applicable to instrument specification and monitoring threshold design for two-tier industrial monitoring architectures.

Future research directions include evaluation on variable-speed and variable-load operating conditions using order-tracking signal resampling, development of transfer learning frameworks for cross-machine generalisation, and integration of the proposed ML pipeline with ISO 13373-compliant condition monitoring systems for industrial certification. The proposed comparative framework is extensible to other rotating machinery components including gearboxes, pump impellers, and motor stators, enabling a unified ML-based condition monitoring platform for complete drivetrain health assessment.

References

[1] Tandon, N., & Choudhury, A. (1999). A review of vibration and acoustic measurement methods for the detection of defects in rolling element bearings. Tribology International, 32(8), 469-480.

[2] Lei, Y., Lin, J., He, Z., & Zuo, M. J. (2013). A review on empirical mode decomposition in fault diagnosis of rotating machinery. Mechanical Systems and Signal Processing, 35(1-2), 108-126.

[3] Zhang, W., Peng, G., Li, C., Chen, Y., & Zhang, Z. (2017). A new deep learning model for fault diagnosis with good anti-noise and domain adaptation ability on raw vibration signals. Sensors, 17(2), 425.

[4] Zhao, M., Kang, M., Tang, B., & Pecht, M. (2019). Deep residual networks with dynamically weighted wavelet coefficients for fault diagnosis of planetary gearboxes. IEEE Transactions on Industrial Electronics, 65(5), 4290-4300.

[5] Chen, Z., Gryllias, K., & Li, W. (2020). Mechanical fault diagnosis using Convolutional Neural Networks and Extreme Learning Machine. Mechanical Systems and Signal Processing, 133, 106272.

[6] Guo, L., Lei, Y., Xing, S., Yan, T., & Li, N. (2021). Deep convolutional transfer learning network: A new method for intelligent fault diagnosis of machines with unlabeled data. IEEE Transactions on Industrial Electronics, 66(9), 7316-7325.

[7] Loparo, K. A. (2012). Bearings vibration dataset. Case Western Reserve University Bearing Data Center. http://csegroups.case.edu/bearingdatacenter/pages/download-data-file.

[8] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems (NeurIPS), 30, 4765-4774.

[9] Nectoux, P., et al. (2012). PRONOSTIA: An experimental platform for bearings accelerated degradation tests. IEEE International Conference on Prognostics and Health Management, 1-8.

[10] 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.

[11] Eren, L. (2017). Bearing fault detection by one-dimensional convolutional neural networks. Mathematical Problems in Engineering, 2017, Article 8617315.

[12] Wang, J., Li, C., Han, S., Sarkar, S., & Zhou, X. (2017). Predictive analytics of IoT-aided intelligent systems using transfer learning. IEEE Internet of Things Journal, 6(1), 252-265.

[13] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.

[14] Sharma, R. K., & Verma, A. (2022). Vibration-based bearing fault classification using ensemble machine learning on CWRU dataset. Journal of Mechanical Engineering Science, 236(14), 7821-7835.

[15] Mehta, P., Bose, S., & Chaudhary, N. (2023). Comparative analysis of deep learning architectures for rolling element bearing fault diagnosis. International Journal of Advanced Manufacturing Technology, 124(3), 1187-1204.