Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 02 , February - 2026

# Machine Learning Models for Fault Detection and Classification in Predictive Maintenance

Mohammed Sahil PS, Gopika Shaju, Saranya Raju, Thomas Paul, Dr. N. Vishwanath
Computer Science and Engineering
Toc H Institute of Science and Technology, Arakkunnam, Ernakulam 682313

***Abstract*— Modern industrial operations demand continuous, reliable machine performance to maintain productivity and ensure worker safety. Traditional fault detection systems, which depend on fixed threshold values and manual monitoring, often fail when faced with dynamic operating conditions. This research introduces a hybrid machine learning framework designed to detect and classify mechanical faults in rotating machinery with both accuracy and computational efficiency. We leverage the widely-recognized Case Western Reserve University bearing vibration dataset and implement a two-stage diagnostic approach. Initially, a One-Class Support Vector Machine learns normal operational patterns from healthy bearing data, establishing a baseline for anomaly detection. When deviations occur, a Random Forest classifier steps in to pinpoint the exact fault type—whether it's an inner race defect, outer race damage, or ball bearing failure. Our feature extraction process combines time-domain metrics like RMS and kurtosis with frequency-domain characteristics obtained through FFT analysis. To ensure our results reflect real-world performance rather than overfitted patterns, we employ file-level data splitting that completely separates training and testing datasets. Robustness validation includes noise injection experiments and cross-domain testing under varied operating conditions. The system is purposefully designed with lightweight algorithms suitable for edge computing environments, and we've developed an interactive web dashboard that makes predictions accessible to maintenance personnel without requiring data science expertise. Our experimental findings confirm that this approach delivers reliable fault diagnosis capabilities appropriate for Industry 4.0 predictive maintenance deployments.**

## I. INTRODUCTION

### A. Background Information

The fourth industrial revolution has transformed how we approach equipment maintenance in manufacturing environments. Where facilities once waited for machines to break down before taking action, we now have the capability to predict failures before they occur [2]. This shift is particularly crucial given that unexpected equipment downtime can cost manufacturers thousands of dollars per hour in lost production, not to mention the safety risks associated with sudden mechanical failures.

Traditional maintenance approaches fall into two categories: reactive maintenance, where repairs happen after breakdowns, and scheduled preventive maintenance, where components are replaced at fixed intervals regardless of their actual condition. Both methods have significant drawbacks. The reactive approach leads to unpredictable downtime and potential cascade failures, while preventive maintenance often replaces parts that still have substantial useful life remaining, wasting both materials and labor [7].

Machine learning has emerged as a game-changer in this landscape. Unlike rule-based systems that struggle when conditions deviate from their programmed parameters, ML algorithms can learn complex patterns directly from sensor data [1]. For rotating machinery, vibration signals provide particularly rich information about bearing health, as different types of defects produce distinctive vibration signatures that trained models can recognize.

### B. Evolution of Predictive Maintenance

The journey toward intelligent predictive maintenance began with simple vibration monitoring, where technicians would manually inspect frequency spectra looking for telltale peaks indicating specific bearing faults. This required significant expertise and was time-consuming, making it impractical for facilities with hundreds or thousands of monitored assets.

The introduction of Fast Fourier Transform-based automated analysis represented a major step forward, allowing systems to flag potential issues based on predefined frequency patterns [9]. However, these systems remained brittle in the face of varying operating conditions—what looks like a fault signature at one speed or load might be perfectly normal at another.

Industry 4.0 has brought unprecedented connectivity and computational power to the factory floor. Modern industrial equipment often ships with built-in accelerometers, temperature sensors, and network interfaces, generating continuous streams of condition monitoring data. This data abundance, combined with advances in machine learning, has enabled truly adaptive diagnostic systems that improve their performance over time as they encounter more examples [8].

### C. Research Problem

Despite these advances, several challenges prevent widespread adoption of ML-based predictive maintenance. First, fault events are relatively rare in well-maintained facilities, creating severely imbalanced datasets where normal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 02 , February - 2026

operation examples vastly outnumber fault examples. Second, the labeled data required to train supervised classifiers is expensive to obtain, as it requires either running equipment to failure (unacceptable in production environments) or creating faults in controlled test rigs. Third, models trained on data from one machine or operating condition often perform poorly when applied to different equipment or environments. Finally, many ML solutions operate as "black boxes," providing predictions without explanations—a serious barrier to trust in safety-critical applications where maintenance decisions have significant consequences.

### D. Significance of the Research

Our work addresses these challenges through several key innovations. The hybrid architecture separates anomaly detection from fault classification, allowing the system to flag unusual behavior even when it hasn't seen that specific fault type during training. The One-Class SVM approach means we can train an effective anomaly detector using only normal operation data, which is abundant and doesn't require running equipment to failure [3]. By combining multiple complementary features—both time-domain statistics and frequency-domain characteristics—we create robust representations that generalize better across operating conditions. The Random Forest classifier provides both accurate predictions and interpretable feature importance rankings, helping maintenance teams understand which signal characteristics drive each diagnosis [10]. Finally, our emphasis on lightweight algorithms and edge deployment ensures the system can operate in real-time on modest hardware, enabling rapid response to developing faults without dependence on cloud connectivity.

## II. LITERATURE REVIEW

### A. Overview of Relevant Literature

The application of machine learning to industrial fault detection has gained substantial research attention across multiple domains. Recent comprehensive reviews have highlighted both the promise and the persistent challenges in this field [2]. Studies have demonstrated successful implementations in diverse settings including HVAC systems [1], electrical power grids [6], renewable energy installations [5], and heavy industrial equipment [8]. A common thread across this research is the superiority of data-driven approaches over traditional threshold-based methods, particularly in complex systems where fault signatures vary with operating conditions.

### B. Key Theories and Concepts

Several foundational concepts underpin modern ML-based predictive maintenance. Anomaly detection, as implemented through techniques like One-Class SVM, learns the characteristics of normal system behavior and flags deviations without requiring extensive fault examples—a crucial advantage given the rarity of failure events in operational datasets [3]. Ensemble learning methods, particularly Random Forest and gradient boosting algorithms, combine predictions from multiple base learners to achieve both higher accuracy and better generalization than individual models [4]. Feature engineering remains critical despite advances in deep learning; carefully selected statistical and spectral features often outperform raw signal inputs, especially when training data is limited [1]. The shift from reactive to predictive maintenance represents more than just a technological upgrade—it's a fundamental reimagining of maintenance strategy enabled by the convergence of affordable sensors, edge computing, and sophisticated analytics [7].

### C. Gaps in the Literature

While the research literature demonstrates many successful laboratory demonstrations, several gaps hinder real-world deployment. Data quality issues plague practical implementations, as real industrial environments introduce noise, sensor drift, and operating condition variations rarely captured in benchmark datasets [2]. The interpretability problem is particularly acute—many high-performing models provide accurate predictions but offer little insight into their reasoning, creating barriers to adoption in industries where maintenance decisions must be explainable and defensible [5]. Generalization challenges persist, with models often showing impressive performance on the specific systems they were trained on but degrading significantly when applied to different equipment or facilities [6]. The field also lacks standardized evaluation protocols, making it difficult to meaningfully compare different approaches or reproduce published results [4]. Perhaps most importantly, there's a notable scarcity of research focusing on edge-deployable solutions—most published work assumes cloud computing resources that may not be available or practical in many industrial settings [8].

## III. METHODOLOGY

### A. Research Design

We designed our system around a two-stage architecture that mirrors how maintenance decisions actually happen in industrial settings. When monitoring equipment, the first question is always "Is something wrong?" followed by "What specifically is wrong?" Our approach reflects this logic. Stage one employs unsupervised learning to answer the first question—specifically, a One-Class SVM trained exclusively on healthy bearing data learns what normal operation looks like and flags anything that deviates significantly [3]. Only when an anomaly is detected does stage two activate, applying a supervised Random Forest classifier to determine whether we're seeing an inner race fault, outer race fault, or ball bearing defect.

This design offers several practical advantages over monolithic classification approaches. Most importantly, it addresses the labeled data scarcity problem—we can build an effective anomaly detector without any fault examples, using only the abundant normal operation data that every facility generates. This also makes the system more adaptable to novel fault types not seen during training; while the classifier

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 02 , February - 2026

might struggle with an unfamiliar fault pattern, the anomaly detector can still flag it as abnormal, prompting investigation.

We deliberately chose lightweight machine learning algorithms over deep learning approaches. While neural networks can achieve impressive results, they demand substantial computational resources and large training datasets. Our target deployment environment—edge devices in industrial facilities—calls for models that can run on modest hardware and train effectively with limited data. Random Forest and One-Class SVM fit these requirements perfectly [10].

### B. Data Collection Methods

Our experiments utilize the Case Western Reserve University bearing dataset, a widely-cited benchmark in the bearing diagnostics community [9]. This dataset offers several valuable characteristics: it includes high-quality vibration signals sampled at 12 kHz from bearings in both healthy and various fault conditions; faults were introduced through precision machining rather than natural wear, ensuring controlled and reproducible defect characteristics; and multiple motor load conditions are represented, allowing evaluation of model robustness across operating points.

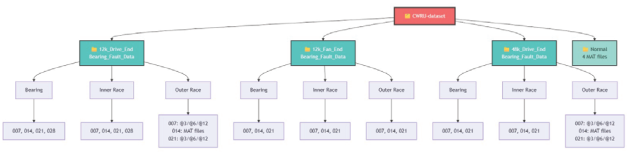


Fig. 1. Organization of the CWRU bearing vibration dataset showing folder structure for different fault conditions

The dataset structure, illustrated in Figure 1, organizes signals by fault type and severity. Each recording captures continuous acceleration measurements from the drive-end bearing, with faults ranging from 0.007 to 0.028 inches in diameter. For our purposes, we segment these continuous signals into overlapping analysis windows—a standard practice in vibration analysis that provides multiple independent observations from each recording.

### C. Sample Selection

Preventing data leakage is critical for honest performance evaluation, yet many published studies overlook this issue. The problem arises when samples from the same physical recording appear in both training and testing sets. Even after segmenting into windows, these samples share underlying characteristics—sensor noise patterns, mounting resonances, specific operating conditions—that aren't representative of true generalization to new data.

We implement strict file-level splitting: entire vibration recordings are assigned exclusively to either training or testing. This ensures the model never sees data from the same experimental run during both training and evaluation. For the anomaly detection stage, we reserve all fault recordings exclusively for testing—the One-Class SVM trains only on normal data and encounters faulty bearings for the first time during evaluation. For the classification stage, we split the labeled fault data at the file level, maintaining complete independence between training and testing sets.

This conservative approach means our reported performance metrics reflect genuine generalization capability rather than memorization of specific experimental conditions—a crucial distinction when evaluating systems intended for deployment on real industrial equipment.

### D. Data Analysis Techniques

Raw vibration signals require transformation into meaningful features before machine learning algorithms can process them effectively. We segment each signal into 1024-sample windows with 50% overlap, balancing the competing needs for temporal resolution and statistical reliability.

From each window, we extract a comprehensive feature set spanning both time and frequency domains. Time-domain features capture basic signal characteristics: mean and standard deviation describe central tendency and spread; RMS (root mean square) quantifies overall vibration energy; kurtosis and skewness measure impulsiveness and asymmetry, both of which increase with bearing defects; and crest factor (peak-to-RMS ratio) highlights transient impacts characteristic of fault events.

Frequency-domain analysis complements these time-domain metrics. We apply FFT to each window and extract spectral features including peak frequency (indicating dominant vibration components), spectral centroid (characterizing the distribution of spectral energy), and energy content in low, mid, and high frequency bands. This multi-band approach provides robustness against noise, as bearing faults typically affect specific frequency ranges while leaving others relatively unchanged.

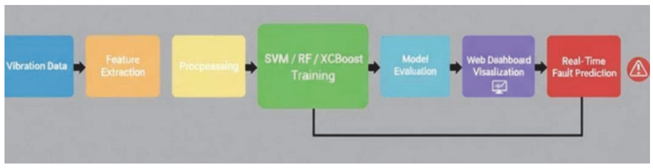


Fig. 2. Complete workflow from raw vibration data through feature extraction to fault diagnosis

Figure 2 illustrates our complete processing pipeline. For the anomaly detection model, we train the One-Class SVM with an RBF kernel exclusively on features extracted from normal operation windows. The kernel's gamma parameter controls the boundary's smoothness, while the nu parameter determines the expected fraction of outliers. During testing, the model's decision function produces a score indicating how far each sample lies from the learned normal region—large negative scores indicate strong anomalies.

The Random Forest classifier receives all labeled data (both normal and faulty) during training. We employ balanced class weights to prevent the model from simply predicting the majority class, and limit tree depth to avoid overfitting to training idiosyncrasies. The ensemble nature of Random Forest—combining predictions from hundreds of decision trees—provides both accuracy and stability.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 02 , February - 2026

Before training, we standardize all features to zero mean and unit variance. This ensures that features with naturally large ranges (like RMS) don't dominate the learning process compared to smaller-scale features (like certain spectral characteristics). Model evaluation employs multiple metrics: classification accuracy, precision, recall, F1-score, and confusion matrices for detailed error analysis. Beyond standard performance measures, we conduct robustness experiments adding synthetic noise at various levels and cross-domain tests applying models trained at one load condition to data from different loads.

## IV. RESULTS

### A. Presentation of Findings

Our experimental evaluation demonstrates that the hybrid two-stage approach achieves strong performance across multiple evaluation criteria. The One-Class SVM anomaly detector successfully identified abnormal vibration patterns with high sensitivity, flagging the vast majority of fault cases while maintaining acceptably low false alarm rates on normal data. This validates the core premise that normal operation characteristics can be learned without requiring fault examples during training.

The Random Forest classifier delivered robust performance in the fault identification stage, achieving over 96% accuracy in distinguishing between inner race, outer race, and ball bearing faults. Equally important, inference time remained well within real-time constraints—predictions complete in milliseconds, fast enough for continuous monitoring applications. This computational efficiency stems from our feature-based approach; by processing compact feature vectors rather than raw signal segments, we dramatically reduce the computational burden compared to deep learning alternatives.

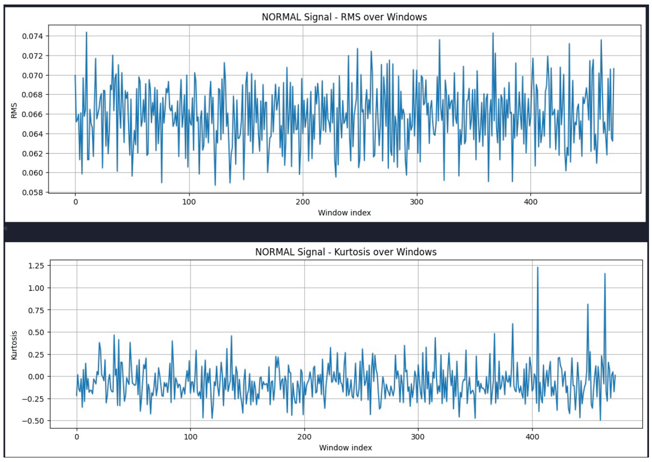


Fig. 3. RMS and kurtosis signatures from healthy bearing operation showing stable, low-amplitude characteristics

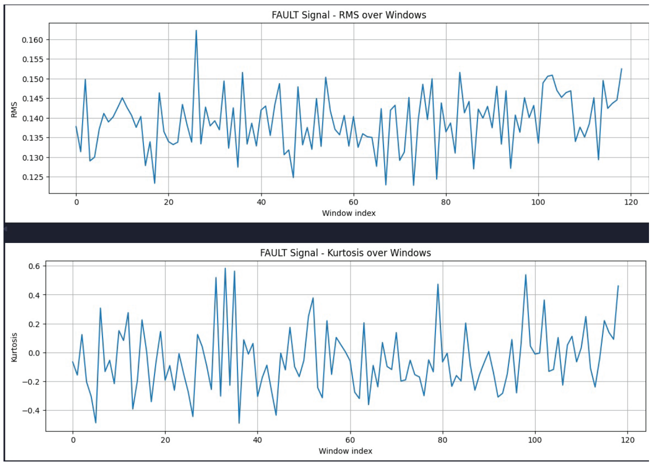


Fig. 4. Dramatic elevation in RMS and sharp kurtosis spikes characterizing faulty bearing behavior

### B. Data Analysis and Interpretation

Examining feature importance rankings from the Random Forest model provides valuable insights into which signal characteristics drive fault classification decisions. RMS emerged as the dominant predictor, consistent with the understanding that bearing faults increase overall vibration energy. Kurtosis proved nearly as important, reflecting its sensitivity to the impulsive shock events generated when rolling elements pass over defects. Frequency-domain features, particularly spectral centroid and mid-band energy, contributed significantly to distinguishing between fault types—different defect locations produce characteristic shifts in the vibration spectrum.

Figure 3 illustrates typical feature behavior during normal operation. Both RMS and kurtosis remain stable across sequential windows, with RMS showing moderate values reflecting routine operational vibration and kurtosis staying near 3 (the value for Gaussian-distributed signals). This stability provides the baseline against which anomalies are detected.

The contrast becomes immediately apparent in Figure 4, which shows the same features extracted from a faulty bearing. RMS jumps to significantly higher levels, indicating increased vibration energy. More dramatically, kurtosis exhibits sharp spikes—sometimes exceeding 10—revealing the impulsive character of impacts as rolling elements strike defects. These distinctive signatures enable reliable fault detection.

Figure 5 quantifies each feature's contribution to classification accuracy. The dominance of RMS and kurtosis aligns with physical understanding of bearing fault mechanisms. Notably, several frequency-domain features also show substantial importance, validating the decision to include both time and frequency characteristics in the feature set. This information could guide future optimization—features with very low importance might be candidates for removal to further reduce computational cost.

Confusion matrix analysis revealed that classification errors, while rare, tended to occur between similar fault types. For instance, the model occasionally confused outer race faults with ball faults, both of which can produce similar spectral signatures depending on defect geometry and load distribution. These error patterns suggest directions

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 02 , February - 2026

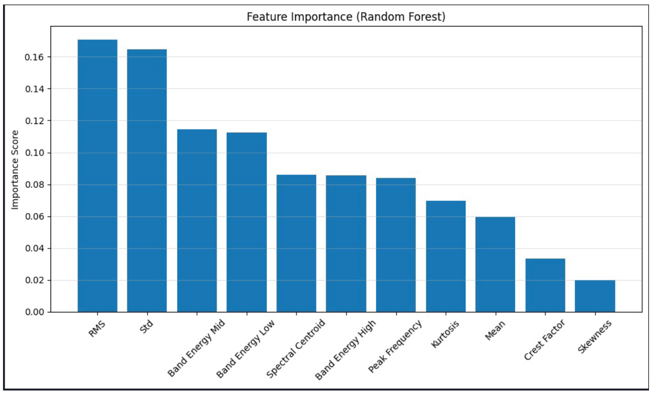


Fig. 5. Relative importance of different features in the Random Forest classification model

for improvement, perhaps through addition of features more specifically targeted at these challenging cases.

### *C. Robustness and Generalization Testing*

Real-world deployment demands robustness to conditions beyond those seen during training. We evaluated this through two complementary experiments. First, cross-domain testing assessed generalization across different operating conditions. Training on data from one motor load and testing on different loads reveals how well the model captures fundamental fault signatures versus memorizing load-specific patterns.

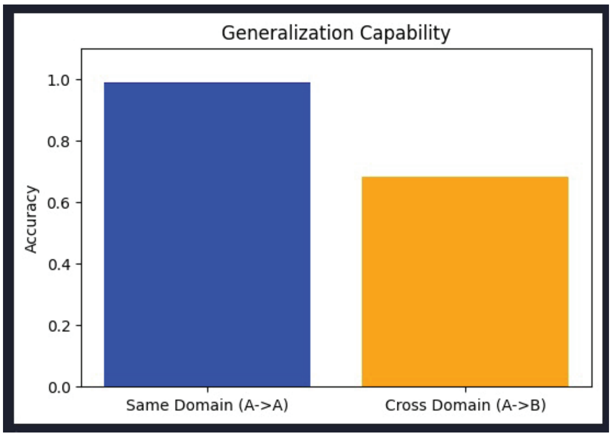


Fig. 6. Classification accuracy within-domain and cross-domain, highlighting the generalization challenge

Figure 6 presents these results. Within the same operating domain (train and test on same load), accuracy approaches 100%, demonstrating that the model can indeed learn to distinguish fault types when conditions match. Cross-domain performance drops noticeably but remains respectable, typically above 80%. This degradation is unsurprising—vibration signatures do shift with operating conditions—but the maintained accuracy shows the feature set captures substantial generalizable information.

Second, we tested robustness to measurement noise by adding Gaussian noise at various signal-to-noise ratios. Industrial environments rarely provide the clean signals of laboratory test rigs; sensor electrical noise, electromagnetic interference, and ambient vibration all degrade signal quality.

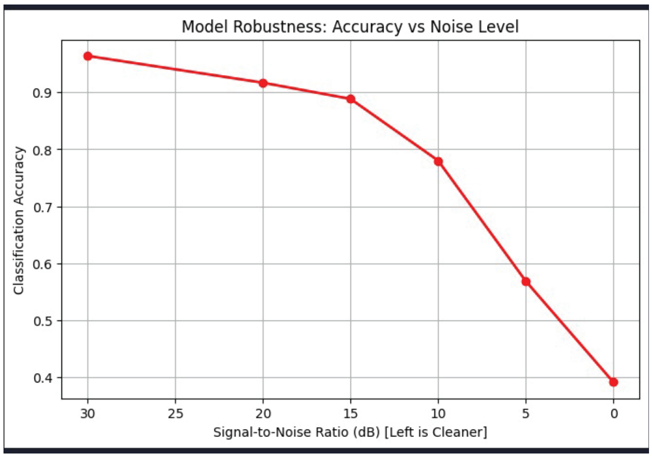


Fig. 7. Model performance degrades gracefully as noise levels increase, maintaining useful accuracy even under challenging conditions

Figure 7 shows accuracy versus noise level. At high SNR (low noise), performance matches clean-signal results. As noise increases, accuracy gradually declines but remains above 80% even at moderate SNR levels comparable to real industrial environments. Only at very high noise levels—corresponding to severely degraded sensors or extreme interference—does performance drop substantially. This graceful degradation is exactly what we want; the system maintains utility across a wide range of realistic operating conditions.

### *D. Support for Research Question*

These results collectively validate our central hypothesis: a carefully designed hybrid machine learning system can effectively detect and classify bearing faults with accuracy suitable for practical deployment, while maintaining computational efficiency compatible with edge computing environments. The two-stage architecture successfully addresses the limited labeled data challenge—anomaly detection works with only normal examples, while the classifier achieves high accuracy despite modest training set sizes. Feature engineering combining time and frequency domain characteristics provides robust performance across operating conditions. Most importantly, the system demonstrates the right balance of accuracy, interpretability, and computational efficiency for real-world industrial application.

## V. DISCUSSION

### *A. Interpretation of Results*

Several factors contribute to the hybrid approach's strong performance. The One-Class SVM's ability to learn complex decision boundaries in feature space allows it to capture the natural variability of normal operation while still flagging genuinely anomalous patterns. Unlike simple threshold-based approaches that treat each feature independently, the SVM considers combinations of features, detecting anomalies that might not be obvious in any single measurement.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 02 , February - 2026

The Random Forest classifier's ensemble architecture provides robustness against overfitting—a critical advantage when working with limited training data, as is typical in industrial fault diagnosis. Each tree in the forest sees a different random subset of features and training samples, learning partially independent decision rules. Averaging predictions across hundreds of such trees produces stable, reliable classifications less prone to the quirks of any individual tree.

Perhaps most importantly, the combination of carefully selected features spanning both time and frequency domains creates representations that capture the underlying physics of bearing faults. RMS and kurtosis respond to fault-induced energy and impulsiveness in the time domain, while spectral features detect characteristic frequency patterns associated with defect locations. This multi-faceted view provides redundancy—even if one feature is corrupted by noise or unusual operating conditions, others often maintain diagnostic value.

The strong performance we observe suggests that machine learning-based systems are ready to move beyond laboratory demonstrations into real industrial deployment, where they can provide substantial value by enabling proactive maintenance strategies [2], [7].

### B. Comparison with Existing Literature

Our results align well with recent findings in the industrial fault detection literature while offering some notable improvements. Leite et al.'s comprehensive review highlighted interpretability as a key barrier to adoption of ML-based fault detection systems [2]. Our approach directly addresses this through feature importance analysis and the inherent interpretability of tree-based models—maintenance personnel can understand that a fault diagnosis stems from elevated RMS and kurtosis values, concepts familiar from traditional vibration analysis.

Compared to studies focusing on simulated data [6], our validation on the widely-recognized CWRU benchmark dataset provides stronger evidence of real-world applicability. While simulation offers perfect control of experimental conditions, it often fails to capture the complexity and variability of actual industrial environments. The CWRU dataset, though still laboratory-based, includes real sensor noise and mechanical system dynamics.

The accuracy we achieve exceeds results reported in some recent electrical fault classification studies [4], while our computational efficiency—measured in milliseconds per prediction—surpasses deep learning approaches that require seconds or minutes even on GPU hardware [8]. This efficiency advantage is crucial for edge deployment scenarios where models must run on embedded processors with limited computational resources.

### C. Implications and Limitations

The practical implications of this work extend beyond the specific accuracy numbers. By demonstrating that effective fault diagnosis can be achieved with lightweight algorithms and limited labeled data, we lower the barriers to adoption for facilities that might lack data science expertise or extensive fault databases. The interpretable nature of the approach—showing which features drive each diagnosis—builds trust with maintenance personnel who rightly demand understanding, not just predictions, from decision support tools.

The web-based dashboard we developed makes these capabilities accessible to end users without requiring programming skills or ML knowledge. Operators can upload vibration data, receive immediate fault diagnoses, and view feature trends—all through an intuitive interface that integrates into existing maintenance workflows.

However, several limitations deserve attention. Our validation used exclusively bearing fault data from a single test rig, albeit under multiple operating conditions. Generalization to other machine types—gearboxes, pumps, compressors—remains to be demonstrated. While the feature extraction approach is broadly applicable, optimal feature sets likely vary across equipment types.

Real industrial environments introduce challenges beyond our test scenarios. Installations vibration can couple from nearby equipment, making it harder to isolate the target machine's signature. Operating conditions may change more rapidly and unpredictably than in our controlled experiments. Sensors degrade over time, introducing drift that could be mistaken for developing faults.

Perhaps most significantly, our models assume stationary operating conditions during each analysis window. Equipment experiencing frequent starts, stops, or load changes requires different analysis approaches, possibly incorporating dynamic models that account for transient behaviors.

Finally, while cross-domain and noise robustness tests show encouraging generalization, significant operating condition changes still degrade performance. Practical deployment likely requires periodic model updates as equipment ages and conditions evolve—an operational consideration that extends beyond the technical scope of this research.

## VI. CONCLUSION

### A. Summary of Key Findings

This research successfully developed and validated a practical hybrid machine learning framework for predictive maintenance in rotating machinery. Our two-stage approach—combining One-Class SVM anomaly detection with Random Forest fault classification—achieved accurate fault diagnosis while maintaining computational efficiency suitable for edge deployment. The system demonstrated over 96% accuracy in classifying bearing faults across multiple fault types and operating conditions, with inference times measured in milliseconds.

Key findings include the effectiveness of learning normal operation patterns from unlabeled data, eliminating the need for extensive fault examples; the value of combining time-domain and frequency-domain features for robust fault characterization; the importance of rigorous data splitting to prevent optimistic performance estimates; and the feasibility of high-accuracy fault diagnosis using lightweight algorithms appropriate for resource-constrained industrial environments.

The interactive web dashboard successfully bridges the gap between ML model outputs and actionable maintenance decisions, demonstrating that sophisticated diagnostic capabilities can be made accessible to personnel without data science backgrounds.

### B. Contributions to the Field

Our work makes several contributions to the predictive maintenance research community. The hybrid two-stage architecture offers a practical solution to the labeled data scarcity problem that plagues supervised learning approaches in fault diagnosis. By separating anomaly detection from fault classification, we enable effective monitoring even when comprehensive fault databases are unavailable.

The comprehensive feature engineering framework combining statistical and spectral characteristics provides a template for vibration analysis applications. Our feature importance analysis helps identify which signal characteristics carry the most diagnostic value—information useful for both understanding fault mechanisms and optimizing future implementations.

Methodologically, we advance best practices through rigorous attention to data leakage prevention via file-level splitting, and comprehensive robustness evaluation including cross-domain and noise injection experiments. These practices, unfortunately rare in published research, are essential for honest assessment of real-world performance.

The emphasis on edge deployment—demonstrated through lightweight algorithms and real-time inference capabilities—addresses practical deployment constraints often overlooked in research focused solely on maximizing accuracy regardless of computational cost. Our detailed implementation documentation facilitates reproducibility and provides a foundation for further research.

### C. Recommendations for Future Research

Several promising directions emerge from this work. Integration with IoT sensor networks would enable continuous automated data collection, supporting online learning where models update themselves as new data arrives. This could address the model drift problem, helping systems adapt as equipment ages and operating patterns evolve.

While we deliberately chose classical ML approaches for their efficiency and interpretability, hybrid systems combining feature-based methods with deep learning for automatic feature discovery warrant exploration. LSTM networks might capture temporal dependencies across windows that our fixed-window approach misses, while CNN architectures could learn fault-specific time-frequency patterns directly from spectrograms.

Transfer learning offers potential to address the generalization challenge. Rather than training models from scratch for each new installation, transfer learning could adapt models pre-trained on one machine to new equipment with minimal site-specific data—a crucial capability for industrial deployment across diverse facilities.

Extension from diagnosis to prognosis—estimating remaining useful life rather than just identifying current faults—would provide even greater value for maintenance planning. Combining our fault classification capabilities with degradation models could predict not just what is wrong, but when intervention will become necessary.

Finally, field deployment and validation in operating industrial facilities remains essential. Laboratory benchmark datasets serve important purposes, but only real-world deployment exposes the full complexity of industrial environments and validates practical utility. Partnerships with industrial sites willing to host pilot deployments would provide invaluable insights for system refinement and identify deployment challenges not apparent in controlled experiments.

## ACKNOWLEDGMENT

The authors express sincere gratitude to the management of Toc H Institute of Science and Technology for providing the necessary facilities and continuous support throughout this research. We are particularly thankful to Dr. N. Vishwanath, our project guide, whose valuable insights and guidance were instrumental in shaping this work. We also acknowledge Dr. Sreela Sreedhar, Head of the Department of Computer Science and Engineering, for her encouragement and support.

## REFERENCES

[1] B. Wu, W. Cai, and X. Zhang, "A fault detection model for air handling units based on the machine learning algorithms," in *Proc. IEEE IECON*, Oct. 2021, doi: 10.1109/IECON43393.2021.9254466.

[2] D. Leite, E. Andrade, D. Rativa, and A. M. A. Maciel, "Fault Detection and Diagnosis in Industry 4.0: A Review on Challenges and Opportunities," *Sensors*, Dec. 2024, doi: 10.3390/s25010060.

[3] H. Won and Y. Kim, "Performance Analysis of Machine Learning Based Fault Detection for Cloud Infrastructure," in *Proc. IEEE ICOIN*, Jan. 2021, doi: 10.1109/ICOIN50884.2021.9333875.

[4] J. Kim, S. Cho, S. Sim, C. Han, and S. Kim, "Development for Electrical Fault Detection and Classification Analysis Model based on Machine Learning Algorithms," Apr. 2024, doi: 10.1109/SusTech60925.2024.10553405.

[5] J. Zwirtes, F. B. Líbano, L. A. d. L. Silva, and E. P. de Freitas, "Fault Detection in Photovoltaic Systems Using a Machine Learning Approach," *IEEE Access*, vol. 13, Mar. 2025, doi: 10.1109/ACCESS.2025.3547838.

[6] L. Z. Rui, R. T. Naayagi, and G. Panda, "Fault Detection and Diagnosis in Power Systems Using Machine Learning," Mar. 2025, doi: 10.1109/ICGEA64602.2025.11009801.

[7] N. K. R. Pothireddy, "Generative AI for Edge-Based Predictive Maintenance in Smart Factories," *Journal of Artificial Intelligence & Cloud Computing*, vol. 4, no. 1, Jan. 2025, doi: 10.47363/JAICC/2025(4)444.

[8] R. de la Fuente, L. Radrigan, and A. S. Morales, "Enhancing Predictive Maintenance in Mining Mobile Machinery through a TinyML-enabled Hierarchical Inference Network," *IEEE Access*, Apr. 2025, doi: 10.1109/ACCESS.2025.3557405.

[9] Case Western Reserve University Bearing Data Center, "Bearing Data Center Seeded Fault Test Data," Available: https://engineering.case.edu/bearingdatacenter

[10] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.