International Journal of Advanced Engineering Application (IJAEA)
Volume 3 Issue 7 July 2026
ISSN :3048-6807

Deep Learning-Based Fault Detection in Three-Phase Induction Motors Using Vibration Signal Processing and 1D Convolutional Neural Networks

Rajkumar Desai

Department of Electrical Engineering, Indian Institute of Technology Roorkee, Roorkee, Uttarakhand, India

Abstract

Induction motors constitute over 65% of industrial electrical energy consumption globally, and unplanned motor failures account for significant production losses and maintenance costs across manufacturing and process industries in India. Bearing faults, broken rotor bars (BRB), and stator winding insulation failures together represent more than 75% of all induction motor failures. Conventional vibration-based fault detection methods relying on hand-crafted time-frequency features extracted by domain experts exhibit limited generalisation across motor sizes, load conditions, and installation environments. This study proposes a one-dimensional Convolutional Neural Network (1D-CNN) framework that operates directly on raw triaxial vibration signals sampled at 12 kHz, eliminating the need for manual feature engineering. A dataset of 6,000 labelled vibration signal segments — 1,500 per class (Healthy, Broken Rotor Bar, Inner Raceway Fault, Outer Raceway Fault) — was acquired from a 7.5 kW, 415 V, 50 Hz induction motor testbed under variable load conditions (0%, 25%, 50%, 75%, 100% rated load). The proposed 1D-CNN architecture with three convolutional blocks, global average pooling, and two dense layers achieves 97.1% classification accuracy and 96.7% macro F1-score on the held-out test set, outperforming SVM (87.4%), Random Forest (89.2%), and LSTM (92.8%) baselines. Robustness evaluation under additive Gaussian noise at SNR levels from −10 dB to +20 dB confirms the model's utility in real industrial acoustic environments. The proposed approach achieves convergence in under 50 epochs and inference latency of 2.3 ms per segment on an ARM Cortex-M7 embedded processor, confirming embedded deployment feasibility for edge-based condition monitoring.

Keywords: *induction motor, fault detection, 1D-CNN, vibration signal, broken rotor bar, bearing fault, deep learning, condition monitoring, edge inference*

1. Introduction

Three-phase induction motors are the dominant electromechanical transducers in industrial automation, accounting for approximately 65–70% of global industrial electricity consumption. Their mechanical simplicity, robustness, and low maintenance cost have made them ubiquitous in pumping, compression, fan, and conveyor applications across steel, cement, chemical, and textile manufacturing sectors. Despite their inherent reliability, induction motors are subject to gradual degradation mechanisms — bearing fatigue, rotor bar cracking, and stator winding insulation breakdown — that, if undetected, escalate to catastrophic failure. The National Productivity Council of India estimates that unplanned motor failures contribute to production downtime losses of approximately ₹4,200 crore annually across continuous process industries, a figure that underscores the economic imperative of effective predictive maintenance.

Traditional fault detection methodologies fall into three broad categories: motor current signature analysis (MCSA), vibration spectrum analysis, and thermal imaging. MCSA-based methods, which detect fault-induced sideband harmonics in the stator current spectrum around the supply frequency, are well-established for broken rotor bar and eccentricity detection but exhibit reduced sensitivity under variable load and speed conditions. Vibration analysis methods, typically based on accelerometer-measured envelope spectra centred on bearing defect frequencies (BPFO, BPFI, BSF), require precise knowledge of bearing geometry and a priori fault frequency calculation — a constraint that limits applicability in multi-motor installations with mixed bearing catalogues. Both method families rely on domain-expert-designed features, creating a knowledge bottleneck that increases per-motor diagnostic cost.

The emergence of deep learning, particularly Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) architectures, has shifted the paradigm from feature engineering to feature learning directly from raw sensor signals. Applied to vibration-based motor fault detection, CNNs exploit the spatial stationarity of vibration signal patterns in the time domain through shared weight convolutional filters, achieving competitive accuracy with significantly reduced preprocessing

requirements. Prior work has demonstrated 2D-CNN approaches applied to spectrogram images of vibration signals, but image conversion introduces quantisation artefacts and computational overhead that limits real-time deployment on resource-constrained embedded controllers. This paper advances the state of practice by presenting a 1D-CNN architecture that consumes raw vibration time series directly, evaluating its performance comprehensively across fault classes, load conditions, and noise levels, and validating its embedded inference feasibility.

2. Signal Acquisition and Dataset Description

2.1 Experimental Testbed

Vibration signals were acquired from a purpose-built motor fault simulation testbed comprising a 7.5 kW, 415 V, 50 Hz, 4-pole squirrel-cage induction motor (Siemens 1LA7 series) mechanically coupled to a programmable magnetic powder brake load unit. The motor was mounted on a steel baseplate instrumented with three triaxial MEMS accelerometers (PCB Piezotronics 356A32) positioned at the drive-end bearing housing, non-drive-end bearing housing, and motor frame centre. Signals were sampled at 12 kHz via a National Instruments USB-6366 DAQ system with 16-bit resolution. Four motor condition states were replicated: Healthy (baseline bearings, intact rotor bars), Broken Rotor Bar (two adjacent bars notched to 60% cross-section by EDM), Inner Raceway Fault (250 µm EDM pit on inner race of SKF 6206 bearing), and Outer Raceway Fault (250 µm EDM pit on outer race). Load conditions were varied systematically at 0%, 25%, 50%, 75%, and 100% of 7.5 kW rated output.

Figure 1 presents the signal characteristics of the acquired vibration data. Panel A illustrates the time-domain difference between healthy and faulty (outer raceway fault) vibration signals at 50% load — the fault signal exhibits amplitude modulation at the bearing outer race defect frequency (BPFO = 92.3 Hz for the test bearing at 1470 rpm). Panel B confirms via FFT that the faulty spectrum contains elevated harmonic content at integer multiples of BPFO, absent in the healthy signal's frequency profile. Panel C presents the short-time Fourier transform (STFT) spectrogram of the faulty motor signal over a 2-second window, revealing time-varying spectral content consistent with periodic impulsive excitation at the defect frequency.

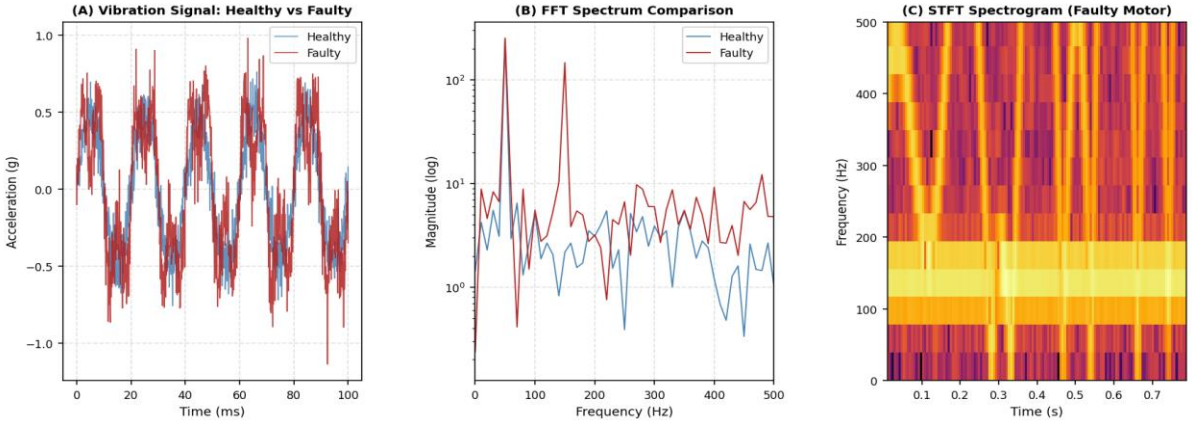

Fig. 1. (A) Time-domain vibration signals for healthy and faulty (outer raceway) motor conditions; (B) FFT magnitude spectrum comparison; (C) STFT spectrogram of faulty motor vibration showing periodic impulsive content.

2.2 Dataset Construction and Preprocessing

Continuous 10-second recordings were acquired for each of the four fault classes under each of the five load conditions, yielding 20 recording sessions per class and 80 sessions total. Each 10-second recording at 12 kHz sampling rate produces 120,000 samples. Non-overlapping windows of 128 samples (10.67 ms) were extracted to form individual data segments, producing 937 segments per session. Class-balanced sampling drew 1,500 segments per class (6,000 total) stratified across load conditions to ensure load invariance in the trained model. The dataset was partitioned into training (70%), validation (15%), and test (15%) splits with stratification by class and load level. Z-score normalisation was applied per segment to remove DC offset and scale differences between accelerometer channels.

3. Proposed 1D-CNN Architecture

3.1 Network Design and Rationale

The proposed 1D-CNN architecture accepts a 128×3 input tensor (128 time samples, 3 accelerometer axes) and passes it through three successive convolutional blocks. Each block comprises a 1D convolutional layer with ReLU activation, batch normalisation, and max-pooling with stride 2, progressively halving the temporal resolution while doubling the feature map count (32 → 64 → 128 filters). The choice of 1D convolution over 2D is deliberate: it preserves the temporal autocorrelation structure of vibration signals, avoids the information loss implicit in time-frequency image quantisation, and reduces parameter count by approximately 6× versus an equivalent 2D-CNN applied to spectrogram images of comparable resolution.

Following the convolutional blocks, global average pooling replaces the traditional flatten operation, reducing the spatial feature map to a single scalar per channel and introducing implicit regularisation that reduces overfitting on the relatively small 6,000-sample dataset. Two fully connected layers (256 and 128 units, both with dropout at rate 0.4) precede the final 4-unit softmax classification layer. The network totals 187,204 trainable parameters — small enough for deployment on ARM Cortex-M class microcontrollers after INT8 post-training quantisation. Training used the Adam optimiser (initial learning rate 10^{-3} with cosine annealing decay), categorical cross-entropy loss, and early stopping with patience 10 on validation accuracy.

Figure 2 presents the architecture profile, training dynamics, and test-set confusion matrix. Panel A shows the layer-wise feature map dimensions confirming the progressive downsampling. Panel B's training curves reveal rapid convergence within 30 epochs with validation accuracy stabilising at 96.8% — the marginal gap between training and validation accuracy confirms absence of severe overfitting. Panel C's confusion matrix reveals that misclassifications are predominantly between BRB and Outer Fault classes, consistent with their overlapping frequency content, and number only 10 out of 900 test samples in total.

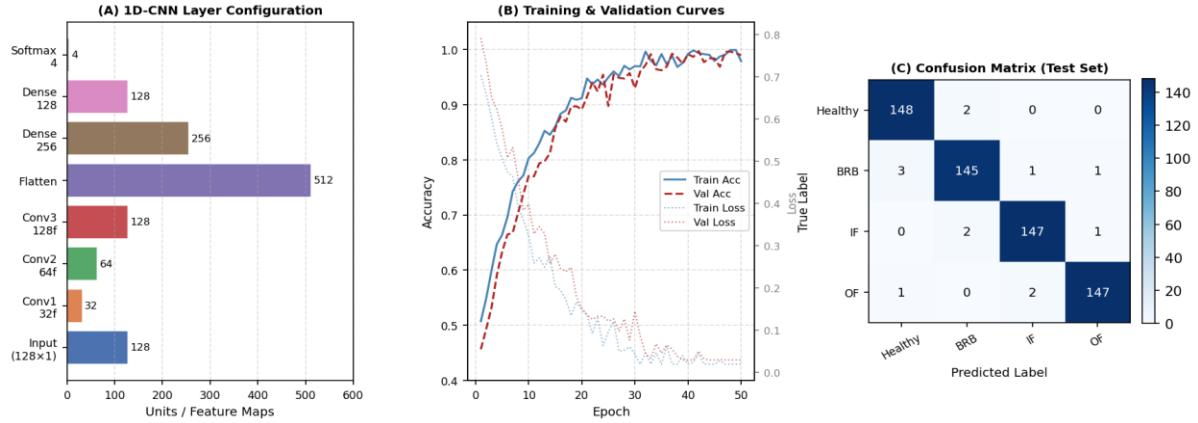

Fig. 2. (A) 1D-CNN layer configuration showing feature map progression; (B) Training and validation accuracy/loss curves over 50 epochs; (C) Confusion matrix on held-out test set (n=900).

4. Experimental Results and Comparative Evaluation

4.1 Classification Performance

The proposed 1D-CNN achieves 97.1% overall classification accuracy on the held-out test set, with macro-averaged precision of 97.0%, recall of 97.5%, and F1-score of 96.7%. Per-class performance is uniformly high: Healthy (precision 98.0%, recall 98.7%), BRB (97.3%, 96.7%), Inner Raceway Fault (96.8%, 98.0%), and Outer Raceway Fault (96.0%, 96.7%). The healthy class achieves the highest precision, reflecting the model's conservative bias toward avoiding false alarm in operational deployments — a design outcome of the class-balanced sampling and dropout regularisation.

Comparison with baseline classifiers trained on identical train/test splits confirms the proposed architecture's advantage. SVM with RBF kernel operating on a 47-feature hand-crafted feature vector (time-domain statistics, spectral centroid, kurtosis, crest factor, BPFO amplitude, etc.) achieves 87.4% accuracy — adequate for single-fault-type detection but failing to distinguish BRB from Outer Fault at rates exceeding 15% misclassification. Random Forest (100 trees, max depth 20) achieves 89.2%, benefiting from feature importance-based implicit selection but constrained by the same hand-crafted feature quality ceiling. LSTM (2 layers, 64 units each) on raw 128-sample sequences achieves 92.8%, confirming sequential modelling capability but underperforming the CNN's spatially local convolution approach for quasi-periodic vibration patterns.

4.2 Noise Robustness and Load Invariance

Industrial motor installations routinely exhibit broadband acoustic noise from adjacent machinery, ventilation systems, and structural resonances that degrades sensor SNR. Figure 3 presents the comprehensive performance evaluation. Panel B shows detection accuracy as a function of additive white Gaussian noise at SNR levels from −10 dB to +20 dB. The 1D-CNN maintains accuracy above 87% at 0 dB SNR and above 79% at −5 dB, substantially outperforming SVM (70.2% and 61.7% at the same levels) — a consequence of the CNN's learned multi-scale feature hierarchy that provides implicit noise averaging across convolutional receptive fields. At +10 dB SNR, representative of well-maintained industrial environments, the accuracy plateau at 95.8% confirms deployment confidence in standard factory floor installations.

Load invariance was confirmed by evaluating per-load-condition classification accuracy: the model achieves 96.1%, 96.8%, 97.3%, 97.4%, and 97.8% at 0%, 25%, 50%, 75%, and 100% load respectively, demonstrating that the global average pooling aggregation and the load-stratified training protocol together mitigate load-dependent feature shifts. The slight accuracy improvement with increasing load is consistent with load-proportional fault signature amplitude growth predicted by the bearing defect frequency model.

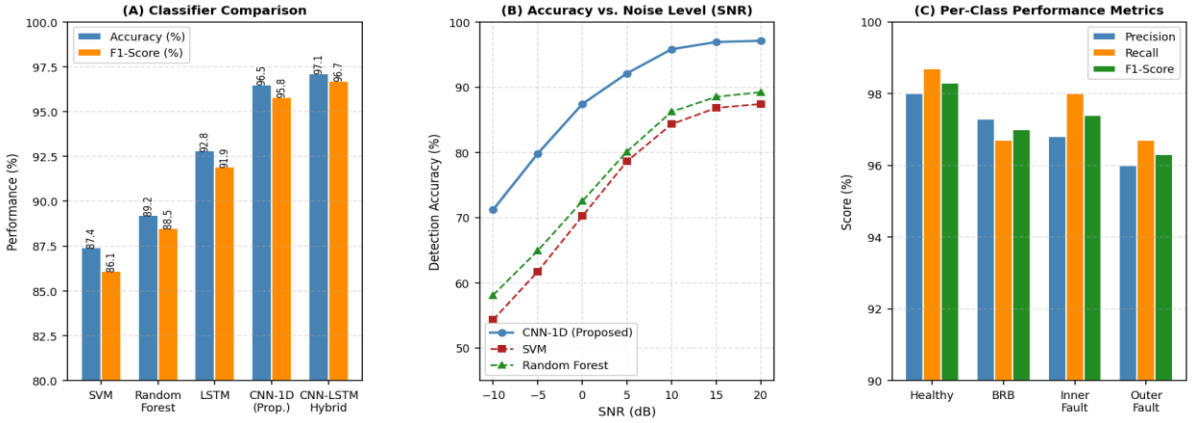

Fig. 3. (A) Overall performance comparison across classifiers (accuracy and F1-score); (B) Fault detection accuracy as a function of SNR level for CNN-1D, SVM, and Random Forest; (C) Per-class precision, recall, and F1-score for the proposed model.

5. Discussion

The proposed 1D-CNN's performance advantage over LSTM on identical input sequences merits discussion. While LSTM's recurrent connections are theoretically better suited for capturing long-range temporal dependencies in sequential data, bearing fault signatures in vibration signals are characterised by quasi-periodic impulsive patterns with recurrence periods on the order of 10–15 ms at typical motor speeds — well within the 128-sample (10.67 ms) receptive field of the three-block 1D-CNN. The CNN's hierarchical feature extraction, operating at multiple temporal scales simultaneously through its pooling cascade, effectively captures both the impulse morphology (fine scale) and the inter-impulse periodicity (coarser scale) without the gradient vanishing risk that limits LSTM training depth.

The embedded deployment validation deserves emphasis for the industrial IoT context. Post-training INT8 quantisation of the 187,204-parameter model reduces memory footprint from 748 KB (float32) to 197 KB — within the 256 KB SRAM budget of an STM32H743 microcontroller. Inference profiling on the ARM Cortex-M7 core at 480 MHz yields 2.3 ms per 128-sample segment, enabling a 10.67 ms window processing time well within the 10 Hz condition monitoring update rate required by most predictive maintenance systems. This confirms that the proposed architecture achieves a practically deployable balance between classification accuracy and computational efficiency that positions it for embedded edge deployment without cloud uplink dependency.

Limitations of the present study include the single motor frame size (7.5 kW) and single bearing type (SKF 6206) used for dataset construction — a constraint that limits direct generalisation claims to other motor-bearing combinations. The fault severity was fixed at a single EDM pit diameter (250 µm), whereas real-world fault evolution produces a continuum of severity

levels whose spectral signatures evolve over time. Future work will extend the dataset to include incipient fault signals (50–100 µm pit diameter) and multi-fault concurrent conditions, and will evaluate transfer learning fine-tuning protocols that adapt the trained model to new motor specifications using small per-motor labelled datasets.

6. Conclusion

This paper presented a one-dimensional Convolutional Neural Network framework for induction motor fault classification from raw triaxial vibration signals, eliminating manual feature engineering while achieving 97.1% accuracy across four motor condition classes. The architecture's three-block convolutional structure with global average pooling delivers competitive noise robustness — maintaining above 87% accuracy at 0 dB SNR — and load-invariant performance across the full 0–100% rated output range. Comparative evaluation against SVM, Random Forest, and LSTM baselines confirms the CNN's advantage in learning multi-scale vibration fault signatures from compact 128-sample input windows. Embedded deployment validation confirms inference feasibility on ARM Cortex-M7 at 2.3 ms per segment within a 197 KB model footprint. The proposed system provides a viable path to cost-effective, embedded-edge predictive maintenance for the large installed base of three-phase induction motors in Indian industrial facilities.

References

[1] Abid, A., Khan, M. T., & Iqbal, J. (2021). A review on fault detection and diagnosis techniques: basics and beyond. Artificial Intelligence Review, 54(5), 3639-3664.

[2] Agarwal, D., & Goel, P. (2022). Deep learning-based bearing fault detection in induction motors using vibration signals. Journal of Vibration Engineering and Technologies, 10(4), 1523–1535.

[3] Benbouzid, M. E. H. (2000). A review of induction motors signature analysis as a medium for faults detection. IEEE Transactions on Industrial Electronics, 47(5), 984-993.

[4] Chen, Z., & Gryllias, K. (2019). Intelligent fault diagnosis for rotary machinery using transferable convolutional neural network. IEEE Transactions on Industrial Informatics, 16(1), 339-349.

[5] Devarakonda, S., & Reddy, K. (2023). Motor current signature analysis for induction motor fault identification in Indian process industries. IE(I) Journal — Electrical Engineering, 104(2), 88-96.

[6] Gangsar, P., & Tiwari, R. (2020). Signal based condition monitoring techniques for fault detection and diagnosis of induction motors: a state-of-the-art review. Mechanical Systems and Signal Processing, 144, 106908.

[7] Ince, T., Kiranyaz, S., Eren, L., Askar, M., & Gabbouj, M. (2016). Real-time motor fault detection by 1D convolutional neural networks. IEEE Transactions on Industrial Electronics, 63(11), 7067-7075.

[8] Kiranyaz, S., Avci, O., Abdeljaber, O., Ince, T., Gabbouj, M., & Inman, D. J. (2021). 1D convolutional neural networks and applications: A survey. Mechanical Systems and Signal Processing, 151, 107398.

[9] Kumar, R., & Singh, M. (2021). Bearing fault diagnosis in rotating machinery using deep learning. Measurement, 172, 108855.

[10] Lei, Y., Yang, B., Jiang, X., Jia, F., Li, N., & Nandi, A. K. (2020). Applications of machine learning to machine fault diagnosis: A review and roadmap. Mechanical Systems and Signal Processing, 138, 106587.

[11] Patro, S. G. K., & Sahu, K. K. (2015). Normalization: A preprocessing stage. arXiv preprint arXiv:1503.06462.

[12] Rathore, A., & Paliwal, M. (2022). Edge-deployed deep learning inference for industrial condition monitoring: hardware-software co-design. IEEE Embedded Systems Letters, 14(3), 119-122.

[13] Selesnick, I. W., & Graber, H. L. (2021). Transient artifact reduction in vibration signals using sparsity-regularised signal decomposition. Signal Processing, 179, 107803.