Article:
Plavos, D., Tsialiamanis, G., Karnezis, A. et al. (2026) Graph-based convolutional neural networks for the classification of induction motor rotor bar faults using stator current & stray flux. IEEE Transactions on Industry Applications. ISSN: 0093-9994

https://doi.org/10.1109/tia.2026.3685589

Graph-based Convolutional Neural Networks for the Classification of Induction Motor Rotor Bar Faults Using Stator Current & Stray Flux

Dimosthenis Plavos, Georgios Tsialiamanis, Aristeidis Karnezis, Nikitas Tsinnas, Zihao Song, Lazaros S Sofikitis, Tiantai Deng, and Panagiotis A. Panagiotou, *Member, IEEE*

***Abstract*—Fault detection for induction motor diagnostics is crucial to ensure system reliability in electric drives. In recent years, research has shifted to combining signal processing and machine learning techniques to improve diagnostic accuracy under various machine states, notably the working conditions of induction machines with different speeds or load profiles. Modern approaches in diagnostic methods such as artificial neural networks rely on manual feature extraction. On the other hand, deep learning such as convolutional neural networks (CNNs) can automatically extract features from raw data in the learning process but are limited to grid-structured data. To this end, graph neural networks (GNNs), and particularly graph convolutional networks (GCNs), have become emerging solutions. This paper proposes a fault detection framework that integrates graph-theory-based learning and time-series analysis. The time segments of stator current and stray magnetic flux signals are constructed as graph nodes, and connections are established based on similarity metrics. Additionally, spatial features are extracted using GCNs, and time-series dynamics are modeled in combination with long short-term memory networks (LSTMs). The method is initially demonstrated using an extensive datasets of transient electromagnetic 2D finite element simulations of two induction motors of the same power rating and different rotor bar number. Then, it is verified experimentally via a range of laboratory measurements at several levels of loading, thus achieving non-invasive and highly accurate rotor fault detection across a range of rotor fault scenarios in various loading levels.**

***Index Terms*—induction motor fault diagnostics, graph convolutional neural networks, deep learning.**

I. Introduction

Electrical machines are key devices in modern industry. Among them, induction motors are the most widely used electromechanical conversion devices. However, due to various reasons such as manufacturing defects or errors introduced by related mechanical equipment, induction motors will inevitably experience faults [1]. These will eventually lead to abnormal operations and operational interruptions, resulting in safety risks and economic losses [2]. Therefore, the development of fault diagnosis methods has become increasingly important. In a cage induction motor, rotor electrical faults usually manifest as broken/cracked rotor bars/end rings. This type of fault does not immediately cause catastrophic consequences, but when a single bar breaks, adjacent or nonadjacent bars will break due to higher current density distribution, causing the fault severity to continue to deepen [3]-[4]. For non-adjacent broken bar faults, when the two breakages are half-pole pitch apart, the asymmetry of the two broken bars will cancel out, therefore causing the fault characteristic signature to be weakened or even missing. The discussed fault mechanism makes traditional diagnostic methods likely to face false diagnostic outcomes [5].

Moreover, diagnostic reliability is strongly affected by operating conditions. Broken rotor bar components in the stator current are slip-dependent [6]. Under no-load or light-load operation the slip approaches zero, so the faulty components become weak and may lie close to the dominant supply frequency component, where they can be obscured by background noise and spectral leakage [6], [7], [8]. Small speed fluctuations can further shift or broaden the spectral components, reducing separability and making broken rotor bar signatures harder to detect reliably. This can be more problematic for machine learning models trained using a limited set of operating conditions, which motivates training and evaluation across multiple load conditions to improve algorithmic robustness [9], [10], [11].

Recent research in induction motor fault detection focuses on combining signal processing with machine learning techniques to improve diagnostic accuracy under real-world conditions. Methods based on time-frequency [12], deep learning [13], and graph-based approaches [14] have shown strong potential for detecting rotor faults in non-intrusive ways. Early

This work was partially supported by Innovate UK / UKRI within activities of the project Practical Power Electronics, Machines, and Drives For All (Practical PEMD4ALL), IUK Grant 10033254, under the funding scheme "Driving the Electric Revolution – Building Talent for the Future 2".

Corresponding author: Panagiotis A. Panagiotou.

D. Plavos is with the Department of Information Technology, Uppsala University, 751 05, Uppsala, Sweden (e-mail: dimosplavos@gmail.com).

G. Tsialiamanis is with the School of Mechanical, Aerospace, and Civil Engineering, University of Sheffield, S1 3JD, Sheffield, UK (e-mail: g.tsialiamanis@sheffield.ac.uk).

A. Karnezis is with the School of Electrical and Electronic Engineering, University of Sheffield, S1 3JD, Sheffield, UK (e-mail: a.karnezis@sheffield.ac.uk).

N. Tsinnas is with the Department of Electrical and Photonics Engineering, Technical University of Denmark (DTU), 2800 Kongens Lyngby, Denmark (e-mail: s253629@dtu.dk).

Z. Song is with the School of Electrical and Electronic Engineering, University of Sheffield, S1 3JD, Sheffield, UK (e-mail: zsong15@sheffield.ac.uk).

Lazaros S. Sofikitis is with the Department of Informatics and Computer Engineering, University of West Attica, Athens, Greece (e-mail: aivc24017@uniwa.gr).

T. Deng is with the School of Electrical and Electronic Engineering, University of Sheffield, S1 3JD, Sheffield, UK (e-mail: t.deng@sheffield.ac.uk).

P. A. Panagiotou is with the School of Electrical and Electronic Engineering, University of Sheffield, S1 3JD, Sheffield, UK (e-mail: p.panagiotou@sheffield.ac.uk).

applications in this direction for motor fault detection utilised classifiers such as Artificial Neural Networks (ANNs) [15], which relied heavily on manual feature extraction and domain expertise [16]. The introduction of deep learning models, particularly Convolutional Neural Networks (CNNs), marked a significant advancement by enabling automatic feature learning from raw or minimally processed signals [17].

CNNs are primarily suited for grid-structured data such as images and for ordered sequences such as time series, but CNNs do not explicitly encode relations between multiple heterogeneous sensors. Many existing neural methods on fault-diagnosis add extra steps making them more difficult to deploy. Signals are often converted into other domains by transformations, increasing preprocessing and computational burden [12], [17]. Sequence models such as recurrent networks, including LSTMs, can learn temporal patterns in ordered data, but the representation still depends on a fixed input order and does not directly capture cross-sensor interactions unless extra fusion design is added [18]. Attention-based Transformer models can also capture long-range dependencies, but standard self-attention becomes expensive for long input windows because its cost grows quadratically with sequence length [19]. This limitation becomes more pronounced when combining measurements from different sensor types. Multi-sensor information is essentially handled through explicit fusion designs at the data, feature or decision level requiring alignment and tuning across heterogeneous channels [14], [20].

In the current work, the proposed approach avoids such complexities by using purely time-domain stator current and stray flux measurements, while learning pairwise relationships between sensors directly. Graph neural networks are well suited for this setting because the output is order-invariant to sensor, which also supports future extensions such as sensor-dropout without redesigning the input format. In induction motor fault diagnosis, informative patterns can also appear in the relationships between multiple signals, especially for non-linear or non-stationary signals. For instance, mutual information (MI) has been used to capture non-linear dependence between variables, and entropy-based measures have been applied to quantify changes in coupling or complexity between time series [21], [22]. To capture such dependencies more effectively, research has explored Graph Neural Networks (GNNs), particularly Graph Convolutional Networks (GCNs) [23], which generalise the concept of convolution to graph-structured data. In GCNs, each node aggregates information from its neighbours, capturing complex spatial dependencies and global structures within irregular datasets [23]-[24]. This capability proves effective for tasks requiring structure representations, such as fault detection, where graphs can be constructed from signal features and the physical mechanisms underpinning their relationships [25]-[26].

This paper proposes a machine learning tool for rotor fault detection. The proposed deep learning framework is a combination of the GCNs and Long Short-Term Memory networks (LSTMs), which can effectively distinguish the healthy and faulty states of asynchronous motors under various rotor cage fracture conditions independently from the machine load level [27]. The implementation of the proposed neural network architecture is performed on datasets from extensive transient electromagnetic 2D Finite-Element Analysis (FEA) simulations of two $4\ kW$ motors at various loading conditions and then validated with experiments on a $1.1\ kW$ induction motor, with the experimental results being in agreement with the theoretical ones acquired via the FEA simulations. Besides the accuracy achieved in the classification of the fault, the proposed method provides a rigorous and robust classification of the fault cases and is independent of the load condition, as the presented framework is tested on several values of loading levels from no-load to full-load, and performs equally well in terms of classification accuracy for the whole range of all the loading levels. The remainder of the paper is structured as follows: Section II provides a brief overview of graph-based neural networks and related metrics, while Section III discusses the methodology for the proposed GCNN architecture. Section IV presents the results from FEA simulations and experimental measurements evaluating its applicability and generalisation at the no-load condition, while Section V provides the conclusions of this work.

II. Theoretical Background

In applications of graph theory for fault diagnostics, graphs are constructed with nodes representing signal features, such as wavelet coefficients or extracted descriptors. Additionally, the edges of the graph encode the similarity of interaction between them [28]. By operating over such graph structures, GCNs can model both local and global dependencies, improving fault detection performance, especially in challenging settings like few-shot learning scenarios [26] or semi-supervised environments [25]. This approach enables richer feature representations compared to traditional deep learning methods, making it highly effective for structured industrial datasets such as those encountered in rotating machinery fault detection. To effectively capture interactions between different signal segments, Granger causality analysis has emerged as a valuable tool offering a framework for ushering directional relationships between signals by quantifying how correlated they are. In fault diagnosis contexts, Granger causality has been used to construct graphs where edges represent causal influence between features, enabling more interpretable and data-driven modelling of industrial systems [29]-[30]. By applying Granger-based methods to healthy operational data, it becomes possible to define the baseline relational structure of the system. This can then be leveraged for sensitive fault detection as deviations from expected causal patterns, notionally rendering unsupervised feature extraction.

In addition to capturing spatial dependencies, modelling the temporal progression of system behaviour is equally important. LSTMs are a type of recurrent neural network designed to model temporal sequences with long-range dependencies by mitigating issues like vanishing gradients [18]. LSTMs have demonstrated strong capabilities in capturing the temporal evolution of system behaviours, enabling the detection of gradual or evolving faults over time [31]. Their ability to process sequential data makes them a complementary advantage to graph-based feature extraction methods, supporting a comprehensive spatial-temporal fault diagnosis framework.

On the other hand, GCNs provide a powerful framework for pattern recognition on multivariate time-series data by representing data within a time window as a graph [23]. The nodes of the graph correspond to sensor channels, whereas the edges indicate statistical relations among those channels. In the context of fault detection and rotating machinery, GCNs can distinguish healthy simulation data sets from faulty ones by detecting changes in the correlation among the signal features. Beyond these advantages, the application of GCNs allows to choose among different methods to define the edge weights as the main factor that differentiates the decision of healthy or faulty. Building on these concepts, the approach presented in this paper integrates graph-based learning and temporal modeling for reliable rotor fault detection. By representing time segments of stator current and stray flux signals as nodes, connections are established based on signal similarity measures. Graphs are constructed that capture both the spatial and temporal relational structures of the data. A GCN is used to extract meaningful features from these graphs, while a LSTM models the temporal progression of the motor conditions over time. This combination enables the detection of faults by a non-intrusive condition-monitoring framework and a robust machine-learning-based spatio-temporal approach.

Compared to traditional diagnostic methods for broken bars in induction motors, deep learning frameworks based on Graph Convolutional Neural Networks (GCN) and Long Short-Term Memory Networks (LSTM) effectively avoid the difficulty in diagnosing the fault when the motor is in no-load or light-load operation mode. Previous studies have improved the detectability of broken rotor bar faults when slip is low using advanced signal processing strategies [32] or specialised sensing approaches [33], but these can increase implementation complexity or instrumentation requirements. Artificial intelligence approaches such as neural networks and fuzzy logic have also been explored [34], [35], but many still aim to enhance traditional diagnostic methods rather than directly analysing the raw time domain signal to produce a diagnostic outcome [36].

III. Methodology

A. Graph Generation via Granger Tests and Entropy

GCNs [23] provide a powerful framework for pattern recognition in multivariate time-series data by representing sensor interactions within a time window as a graph. Fig. 1 presents the proposed GCN architecture. The six nodes of each graph represent the sensor channels, notably three stator line currents and three signals from search coils measuring stray flux. The objective is to train a GCN that classifies a small-size window of continuous signals as faulty or healthy. The first raw six-channel series of data are first normalised using z-score normalisation: each channel is centred by its mean (μ_i) and scaled by its standard deviation (σ_i), computed over the entire dataset, to ensure the numerical ranges are comparable across different channels. The continuous stream of each dataset is then split into non-overlapping windows of 40 rows (equivalent to 4ms). Within each window, each sensor becomes a graph node whose feature vector consists of the 40 normalised readings, as illustrated in the "Initial Graph" depicted in Fig. 1.

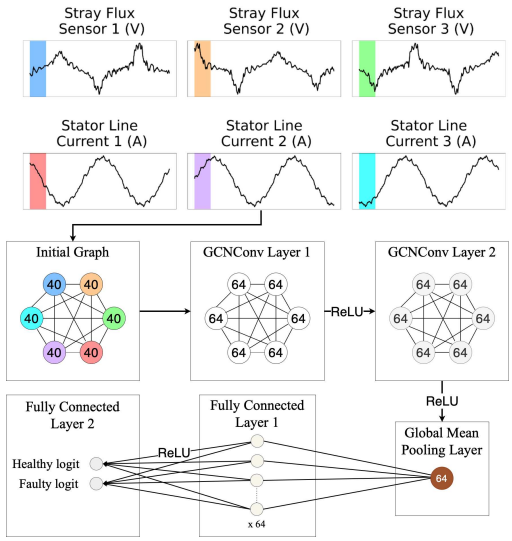

Fig. 1. Architecture of GCN model: Highlighted segments from six channels (3 stray flux and 3 stator current signals) represent a single time window used to construct a graph for spatial feature learning. This visualisation corresponds to one sliding step; subsequent steps shift the window forward in time.

The key challenge in using the GCN is defining the edge weights that reflect interdependence among acquired signals. To achieve this, using only data from healthy simulations, each simulation ($\sim 70,000$ samples) is divided into blocks of $10,000$ samples ($1s$). Within each block, Granger causality is used to quantify directed relationships between features. Specifically, for a given pair of features, the test evaluates whether the past of one feature helps predict the future of the other. To capture effects at different time delays, lags from 1 to 5 timesteps are evaluated and a p-value is computed for each lag. These p-values are then treated as a continuous correlational strength metric, where smaller values indicate stronger evidence of a directed predictive effect. For each feature pair, the smallest p-value across lags is retained and converted into an edge weight $w = 1 - p$. Repeating this process for all signal pairs produces a causality matrix whose values indicate the predictive influence between sensor channels.

However, directly applying Granger causality tests on large blocks can result in an overly dense causality matrix, because p-values tend to zero while increasing the computational cost. To mitigate this, the permutation entropy (PE) is introduced as a computationally efficient way to capture the dependencies among the six features, using an embedding dimension $m = 3$, and delay $d = 1$. In this case, each block of $10,000$ samples is divided into overlapping chunks of 100 rows with an 80-sample step. For each chunk, the permutation entropy is computed on each channel. The permutation entropy quantifies the degree of disorder in the time series of the signals and reflects its complexity. Using the entropy values instead of raw signals, Granger causality tests are run for each chunked block, generating a causality matrix per block. Averaging these matrices across blocks and all healthy simulations produces a final causal weight matrix. In every input graph, an edge from node to node is assigned a weight. To illustrate the importance of permutation entropy, Fig. 2 presents the entropy values

of a search-coil signal capturing stray flux. In healthy data, it remains high and stable, whereas in faulty conditions it exhibits a noticeable drop and fluctuations, indicating more prominent signal dynamics. These confirm the usefulness of permutation entropy as a metric for causal attribution.

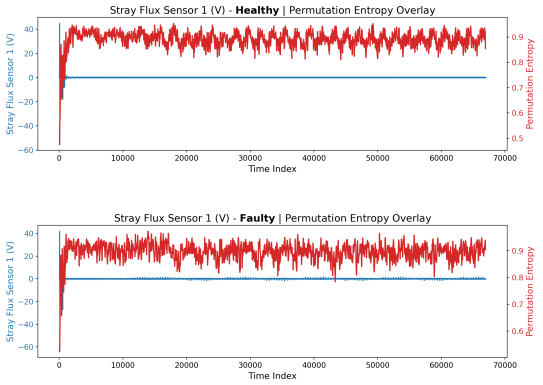

Fig. 2. Entropy values of stray flux: healthy (top) vs faulty (bottom).

B. Graph Convolutional Networks Architecture

In the proposed GCN architecture, each input graph contains six nodes. Each node carries a 40-dimensional feature vector representing a 40-sample window of continuous signal values. The graph is fully connected, resulting in up to 30 directed edges which are weighted by the infered causal matrix from the Granger causality analysis. The graph is processed through two graph convolution layers. The first layer transforms these 40-dimensional inputs into 64 hidden channels, performing weighted messages passing along the edges, followed by Rectified Linear Unit (ReLU) activation. The second GCN layer maintains the 64-dimensional representation while integrating further information across the graph. The resulting node embeddings are collapsed into a single 64-dimensional graph representation by applying a global mean-pool operation. Finally, this embedding is passed through a multilayer perceptron head; first, a fully connected layer ($64 \rightarrow 64$) with ReLU application, then a linear layer ($64 \rightarrow 2$) that produces logits for the binary classification task (healthy vs. faulty).

C. Pairing GCN with LSTM

To further enhance the temporal modelling capabilities and capture sequential patterns across consecutive graphs, the original GCN architecture was extend by a LSTM. In the combined GCN-LSTM model, each graph is constructed as in the original GCN model corresponding to a window of data. To model temporal evolution, sequences of 5 consecutive graphs are formed by sliding a window with a fixed length of 2 graphs. Fig. 3 visualises the extended GCN-LSTM architecture, illustrating how the graph-level embeddings from the GCN are passed through a two-layer LSTM to capture temporal dependencies. Each sequence is labelled by the type of condition case it belongs to (healthy or faulty). Within each sequence, the graphs are first processed independently through GCN layers. Each graph's node features are passed through the first GCN layer, where 40 input features result in 64 hidden channels with ReLU activation. Then, they are passed through the second GCN layer that retains the 64-channel dimensionality, eventually forming a 64-dimensional graph embedding for each graph in the sequence. These embeddings are then stacked along the temporal dimension to form a sequence tensor of shape (batch-size, sequence length, feature-dim), specifically $(\mathrm{B}, 5, 64)$. This sequence tensor is input into a two-layer LSTM network, designed to capture dependencies across continuous data. Each LSTM layer contains a hidden size of 128 units, with a dropout rate of 0.3 applied between layers to prevent overfitting. The output of the LSTM's final step is passed through a fully connected layer ($128 \rightarrow 64$) with ReLU activation, followed by a final linear layer ($64 \rightarrow 2$) that produces logits for binary classification by using class scores. In total, the model contains $246,594$ learned parameters. The rationale for deploying the augmented GCN-LSTM by a serial architecture is to use the GCN to extract latent features from the time signals, and, subsequently, to exploit the temporal characteristics of the features via the LSTM to perform the classification.

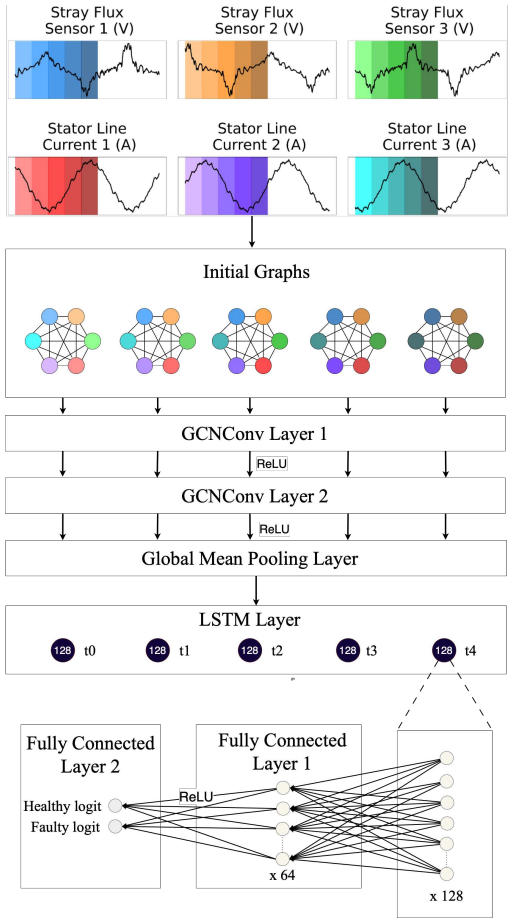

Fig. 3. Architecture of GCN - LSTM model illustrating a single inference step. The figure represents one step in a sliding window sequence (subsequent steps shift the input window forward in time).

D. Induction motor models & datasets

Two 3-phase induction motors with different rotor geometry and number of rotor bars were modelled in 2D FEA using the commercially available software package "Simcenter Magnet" in the electromagnetics environment to examine them under healthy and broken bar cases with the proposed methodology.

Table I specifies the characteristics of the simulated motors. Fig. 4a illustrates the cross-section of Motor #1 with the corresponding contour plots of the magnetic field distribution.

TABLE I
CHARACTERISTICS OF SIMULATED MOTORS

Characteristics	Motor#1	Motor#2
Supply Frequency	50 Hz	50 Hz
Stator Winding	Δ-connection	Δ-connection
Output Power	4 kW	4 kW
Rated Voltage	400 V	400 V
Rated Current	10 A	10 A
Rated speed	1450 rpm	1450 rpm
Pole pairs	2	2
Stator slots	36	36
Rotor bars	28	32

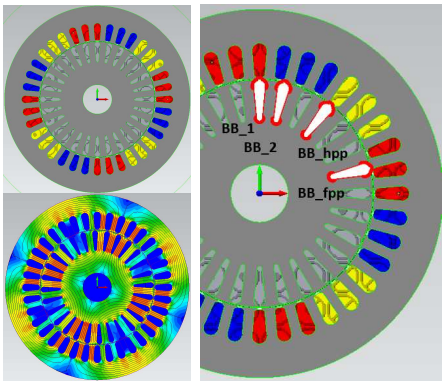

(a) (b)

Fig. 4. (a) The machine geometry (top) with the magnetic distribution (bottom) of the simulated Motor#1 model, and (b) the location of bar breakages.

Beyond the baseline (healthy) model per motor case, four modified configurations were developed for each motor to simulate bar faults at varying severity levels: single-bar breakage, two adjacent bars, two non-adjacent broken bars positioned at a half-pole pitch distance, and two non-adjacent broken bars positioned at a full-pole pitch distance. The locations of each bar breakage case are illustrated in Fig. 4b, while Table II enumerates the simulation models and the slip values under different loads levels. To balance the healthy and faulty simulation data in terms of datasets population, additional data where considered from a healthy machine model with load uniformly ranging from 10% to 90% of the rated loading.

E. Experimental Set-up & Equipment

For experimental validation, an electric machine dynamometer with electric vehicle emulation capability was utilised. This electrical machines rig is developed specifically for electric vehicle simulation systems [37] - [38]. As shown in Fig. 5a, the main components of this rig include an induction motor (IM) and a permanent magnet synchronous motor (PMSM). The induction motor is a 1.1 kW, 4-pole, 28 V, 50 Hz induction motor with 36 stator slots, 46 rotor bars, and a Δ-connected distributed winding. The specific characteristics of the induction motor are detailed in Table III.

TABLE II
SUMMARY OF FEA MODELS AND VALUES OF SLIP

Case	Load	Motor#1	Motor#2
Healthy	Healthy_fl	1.50%	1.20%
	Healthy_hl	0.74%	0.59%
	Healthy_nl	0.072%	0.054%
1 broken bar	BB_1_fl	1.56%	1.24%
	BB_1_hl	0.77%	0.61%
	BB_1_nl	0.077%	0.056%
2 adjacent broken bars	BB_2_fl	1.64%	1.30%
	BB_2_hl	0.81%	0.64%
	BB_2_nl	0.082%	0.059%
2 non-adjacent broken bars at half-pole pitch distance	BB_hpp_fl	1.09%	1.29%
	BB_hpp_hl	0.54%	0.64%
	BB_hpp_nl	0.058%	0.058%
2 non-adjacent broken bars at full-pole pitch distance	BB_fpp_fl	1.08%	1.29%
	BB_fpp_hl	0.54%	0.63%
	BB_fpp_nl	0.056%	0.057%

TABLE III
CHARACTERISTICS OF TESTED MOTORS

Characteristics	Specification
Supply Frequency	50 Hz
Stator Winding	Δ-connection
Output Power	1.1 kW
Rated Voltage	28 V
Rated Current	45.6 A
Rated speed	1450 rpm
Pole pairs	2
Stator slots	36
Rotor bars	46

To create the various rotor bar breakage scenarios, several identical rotors were drilled by precision machining at different locations of the rotor cage corresponding to the single-bar fault, two adjacent broken bars, and two non-adjacent broken bars at half-pole pitch distance, as the arrows indicate in Fig. 5b. The experimental measurements were obtained in open-loop control mode under different load conditions, being consistent with the FEA simulations, and further additional validation was performed beyond the FEA cases at other load levels. Table IV summarises the cases used in experimental measurements and their slip values in different operating loads.

As shown in Fig. 5a, the stator current signals are captured by three identical current sensors (one for each phase terminal). Also, three identical flux sensors, which are search coils with a 1500 turns copper winding, are placed on the surface of the induction motor with a displacement of 120 electrical degrees capturing the stray flux signals outside the yoke of the machine. All measured signals were logged onto an oscilloscope with 8 channels, 12-bit resolution, 20 MHz bandwidth and a serial bus decoding with 256 MS buffer memory. Each of the six waveforms (three current signals and three stray flux

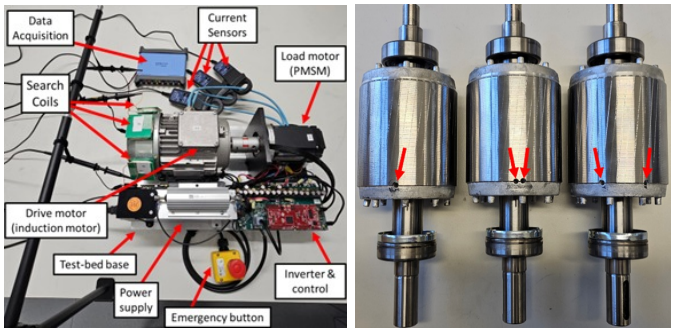

(a) (b)

Fig. 5. Hardware used for experimental measurements: (a) electric vehicle emulator system and instrumentation for data acquisition, and (b) rotors with broken bars at different locations: 1 broken bar (left), 2 adjacent broken bars (middle), 2 non-adjacent broken bars at half pole pitch distance (right).

TABLE IV
SUMMARY OF EXPERIMENTAL MEASUREMENTS - CASES & VALUES OF SLIP PER LOAD LEVEL

Load	H	BB1	BB2	BB_hpp
0%	2.47%	2.13%	2.47%	2.60%
20%	3.67%	3.47%	3.67%	4.00%
35%	5.33%	5.20%	5.47%	5.93%
50%	7.47%	7.27%	7.87%	8.53%
65%	8.80%	8.73%	9.47%	10.20%
70%	9.67%	9.40%	10.20%	12.33%
75%	10.20%	10.27%	11.13%	13.20%
80%	12.27%	12.13%	13.47%	15.40%
85%	13.20%	13.27%	15.20%	17.67%
90%	14.87%	14.80%	17.07%	19.47%
95%	17.00%	16.47%	18.07%	21.80%
100%	18.60%	19.20%	21.40%	25.67%

signals) occupies a separate channel, and is captured at a 20 kHz sampling rate, each lasting 50 seconds, which ensure the measurement adequately covers the signal when the motor is operating at the steady state.

IV. RESULTS

A. Results from FEA simulations

1) Results of GCNN by FEA Simulation Data:

To evaluate rigorously the methods employed in a comprehensive manner, the performance of the standard GCN model is assessed first, using two evaluation strategies. As illustrated in Fig. 6, these are: (i) a random train-test split, and (ii) a generalisation test, where simulations from specific machine fault scenarios were excluded from training and only used in testing. In the first approach, $\approx 80\%$ of the entire corpus of healthy and faulty data was randomly selected for training, while the remaining $\approx 20\%$ formed the test set. This setup resulted in the creation of $33,500$ training graphs ($17,420$ healthy and $16,080$ faulty) and $8,375$ test graphs ($4,355$ healthy and $4,020$ faulty). The GCN was trained for 20 epochs using the Adam optimiser (learning rate: 1×10^{-3}, batch size: 128 graphs) with $11,074$ learned parameters. All neural network architecture implementations were conducted on the Apple M1 CPU, with an average training time per epoch ~ 1.846s, whereas the average inference time reached ~ 0.273s. The CPU time needed for the causal matrix creation was 89.01s. The GCN achieved a classification accuracy of $97,72\%$, indicating high reliability in distinguishing healthy and faulty data windows, as illustrated in Fig. 7(a).

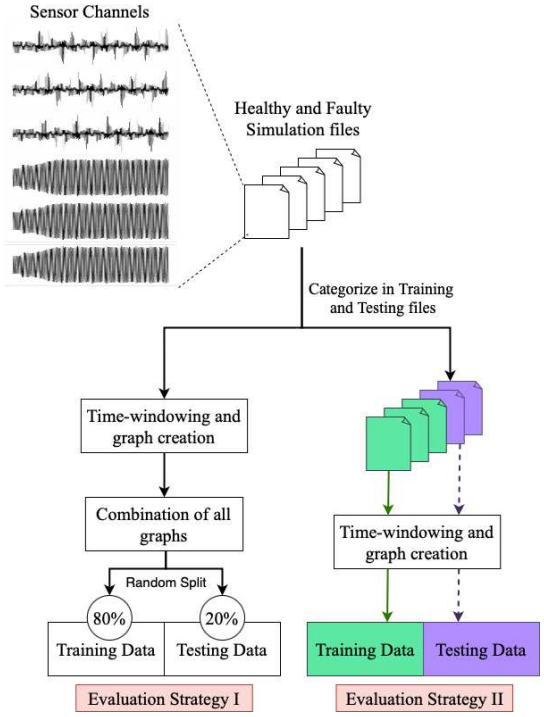

Fig. 6. The two evaluation strategies for the proposed models: "Strategy I" involves combining all data, creating graphs through time-windowing and then randomly splitting into train and test. "Strategy II" assigns entire simulation scenarios exclusively to either training or testing, ensuring the model is evaluated on "unseen" operational conditions.

To evaluate the model's generalisation ability, three independent runs were conducted, where three healthy and two faulty simulation datasets were excluded from training. In each run, the training graphs are equal, $16,750$ for both healthy and faulty data, and the training graphs are $5,025$ and $3,359$, correspondingly. Each accuracy discussed below is reflecting the average of the accuracies on the diagonal of the corresponding confusion matrix, as the dataset is balanced regarding the number of healthy and faulty instances in it. The model was trained only on the remaining simulations and then evaluated on the two "unseen" simulations. Despite never having "seen" these scenarios during training, the GCN maintained a high average generalisation accuracy of 98.08%, with the results presented in Table V showing that the model's detection ability is not limited to specific simulation cases.

Despite the strong performance of the initial GCN, the Granger causality tests led to dense adjacency matrices that do not highlight the dependencies between sensor channels. Replacing the causality graph construction with the PE-based yielded better results with lower computational cost. In terms

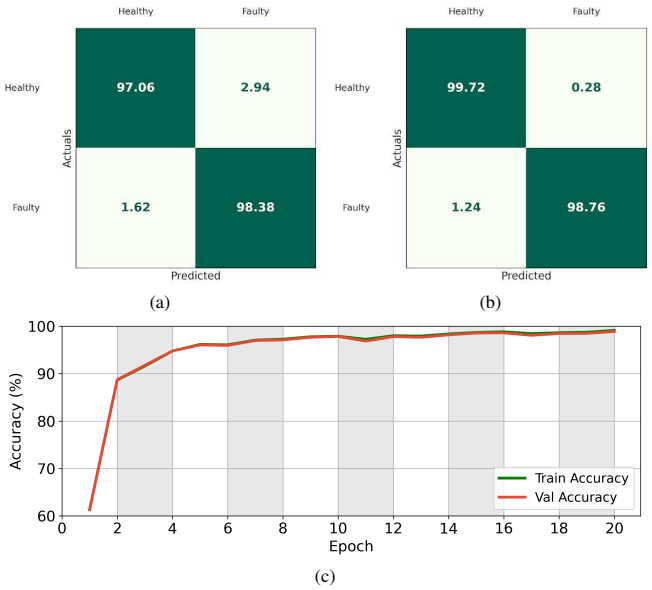

Fig. 7. Confusion matrices of the implemented NN architectures: a) GCN, b) GCN-PE and (c) learning curve of PE-enhanced GCN model.

of accuracy, it achieved an improved value of 99.4%, Fig. 7(b), on the random train-test split, and an average value of 99.01% in the generalisation scenario. Table V gives the average accuracy score by three different generalisation runs using the PE-enhanced GCN. Moreover, Fig. 7(c) presents the learning curve for the Random Train-Test Split scenario, which illustrates that the PE-enhanced GCN converges rapidly. It surpasses accuracy of 90% by the second epoch and then steadily increases to approximately 99%. The train and validation curves show the rapid convergence and strong generalisation performance.

From the perspective of computational cost and burden, the time needed to form the final causal matrix dropped from 89.01s to 14.19s, resulting in a $\times 6.27$ speed-up. This efficiency is crucial for cases of scaling to longer simulations that require the computation of multiple Granger causality matrices. As outlined in the methodology section, the GCN-LSTM model allows the system to learn the temporal dependencies and evolution patterns from the graphs created and the signal dynamics, yielding even higher accuracy. In the random train-test split setting, the model achieved an accuracy of 99.93% (Fig. 8(b)), outperforming the GCN with PE. Remarkably, high accuracy results were observed in the generalisation scenario, where the GCN-LSTM reached accuracies in the range of $99.66\% - 99.95\%$ on the three entirely "unseen" sets of healthy and unhealthy simulations (Table V). Fig. 8(d) visualises the learning curve of the GCN-LSTM method for the random train-test split scenario, showing that it increases quickly to over 95% after epoch 2 and reaches an accuracy near 100%.

After the compelling results on Motor #1, the evaluation was extended to the second machine (Motor #2) to further confirm the validity of the method. Since the value of incorporating Permutation Entropy is evident, the work focused on the GCN-PE results and used the generalisation-based evaluation, which best reflects signal dynamics and indicates system robustness. As shown in Fig. 8(a), the GCN-PE method reached an accuracy of 94.74%, when using whole files (65% for training with balanced healthy and faulty data, 15% for validation and 20% unseen files for testing). Using the same split, the GCN-LSTM method achieved an accuracy score of 98.83% (Fig. 8(c)), confirming the performance gain when adding LSTM. Three runs using different test files were also carried out for both methods and the initial GCN, with all results presented in Table V. Overall, the experiments underline the benefits of PE, both in accuracy and computational efficiency, and the impact of LSTM on improving the final model.

TABLE V
PREDICTION ACCURACY

Test Files	**Average Accuracy % in 3 runs**		
	Initial GCN	***GCN & PE***	***GCN-LSTM***
Motor#1	99.08%	99.01%	99.79%
Motor#2	94.94%	95.30%	98.83%
Experimental	87.91%	89.71%	98.38%

The GCN-LSTM method progressively distinguishes the healthy from faulty data, while Principal Component Analysis (PCA) is applied in three stages of the architecture implementation. As shown in Fig. 9, these are the raw input windows, the hidden-state embeddings produced by the GCN-LSTM model, and the final 2D logits before softmax. By projecting each representation into its top two or three principal components, this approach facilitates the visual inspection and assessment of the model's ability to separate the classes at each step. Regarding Motor #1, in the 3D PCA of the raw inputs, the faulty and healthy data windows are intermixed indicating that there is not a single feature informative enough to discriminate. The next 3D PCA shows a more compelling separation with only a few intermixed points, confirming that the model has captured the discriminative patterns of the data. The last 2D diagram shows two non-overlapping clusters of points, indicating that the model was able to establish a decision boundary that distinctively separates the two classes effectively. A similar trend appears in Motor #2, though with slightly greater overlap between the two classes, which is consistent with the smaller accuracy score in the case of the Motor #2 datasets. The performance drop observed for Motor #2 links to inherent properties such as Principal Slot Harmonics (PSH) [39], [40], [41] as well as to speed- and torque-ripple effects relating to the combination of stator slot-rotor bar number which are to some extent amplified by the 2D geometry of the FEA setting.

2) Generalisation in No-Load Operation via FEA Data

In this work, an evidence-informed decision was made to intentionally choose to evaluate the model's generalisation ability exclusively under no-load scenarios. The rationale for this decision was based on the challenges associated with fault detection in induction motors operating at no-load. In no-load driving conditions, the motor stator current becomes sensitive to small speed changes and small load variations, which makes it difficult for traditional methods such as those based on fast Fourier transform (FFT) to distinguish differences in spectra due to spectral leakage and other

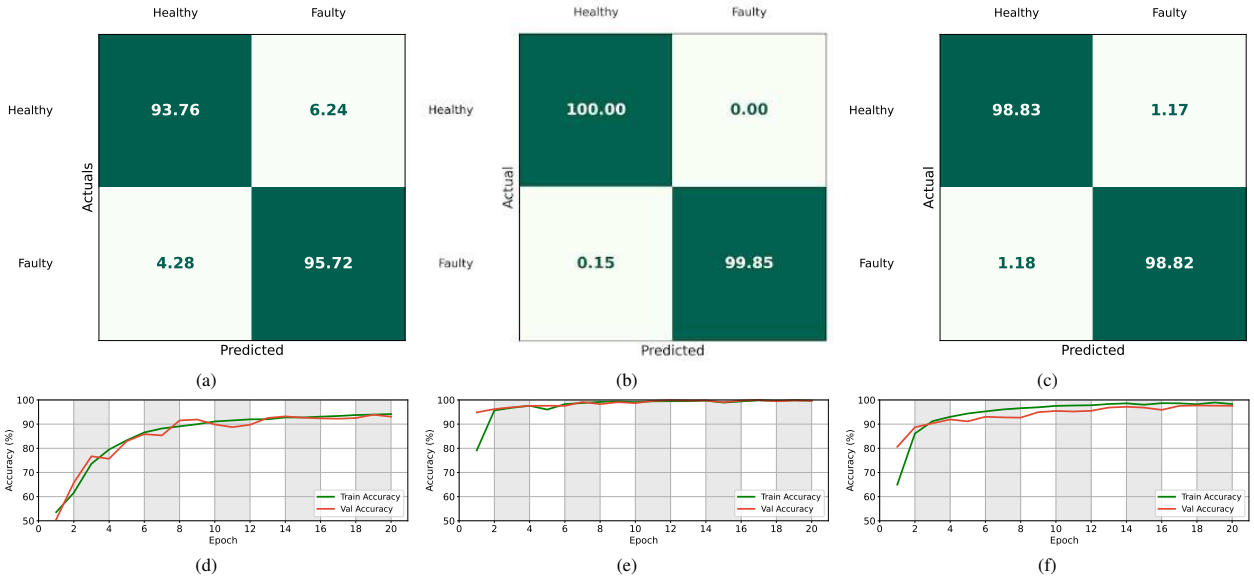

Fig. 8. FEA results - Confusion matrix for: a) GCN-PE for Motor #2 b) the GCN-LSTM model (Motor #1), c)the GCN-LSTM model (for Motor #2), and Learning Curves of: d) the GCN-PE architecture for Motor #2, e) of GCN-LSTM architecture (Motor #1), f) GCN-LSTM architecture (Motor #2).

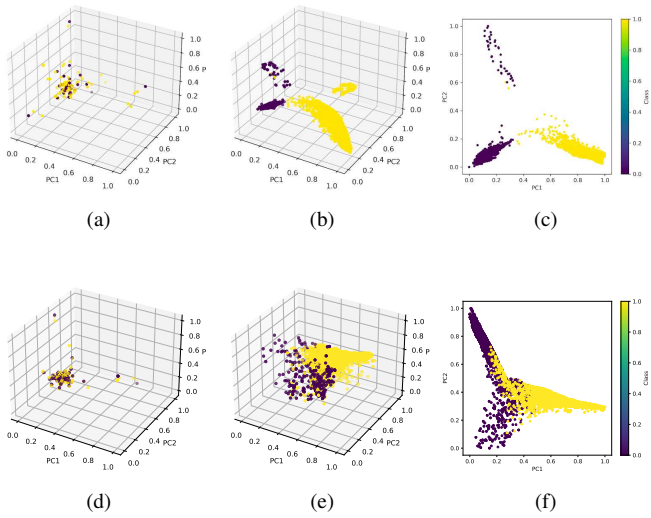

Fig. 9. (a–c) PCA in three stages for the dataset of Motor #1: (a) raw input windows (3D), (b) hidden-state embeddings produced by the GCN-LSTM model, (c) final two-dimensional logits before softmax; (d–f) corresponding PCA results on the dataset of Motor #2.

reasons [9]. Moreover, the currents caused by faults are much weaker under no-load conditions and easily get masked in the background noise, making steady-state Motor Current Signature Analysis (MCSA) susceptible to misdiagnosis [8]. By learning from clearer, more reliable examples and then testing in data from no-load conditions, it is ensured that the proposed model is trained in well-defined fault patterns and generalises in other scenarios. Examining this scenario where the model uses no-load datasets only during the neural network testing process, the proposed GCN-LSTM architecture reached an accuracy score of 97.87%, thus being slightly lower compared to the value of 99% when trained and tested across all loading conditions. Nevertheless, these results still indicate the robustness of the proposed approach. As such, by combining permutation entropy with Granger causality, the proposed framework captures informative relations among signals that traditional approaches such as frequency-based, time-frequency representations, or steady-state methods cannot detect. Therefore, a trial study was conducted excluding the stator line current signals, where it was observed that the model still retains its ability to detect faults with high accuracy. Specifically, it achieved an accuracy of 97.32% on a test set including all no-load simulations. This demonstrates that the model is not dependent on load-related features and that the causal relationships among signals contain sufficient information for reliable diagnosis, even under no-load operation. In the simulation dataset of Motor#2, the model reached an accuracy of 87.30%. Although this accuracy level is lower than the result acquired for the simulation cases of Motor#1, it still demonstrates the model's ability to identify previously unseen no-load faults which is poses a significant diagnostic challenge.

B. Results from Experimental Measurements

1) Results of GCNN on experimental cases

Similar to the FEA results, each accuracy reported below is reflecting the average of the accuracies on the diagonal of the corresponding confusion matrix, as the dataset is balanced regarding the number of healthy and faulty instances in it. The proposed graph-based learning framework was tested on a laboratory-scale 3-phase induction motor. The datasets include measurements from the experimental test-rig described in Section III-E. Expectedly, such measurements include higher levels of noise and complexity because of peripheral electronics and environmental factors, making the classification task more challenging compared to the simulated case. Following the same procedure as in the FEA simulations, the first application on the experimental data involves the application of the

GCN methodology enhanced with the PE information for the causal matrix construction. A class-balanced dataset was used, comprising eight files of healthy runs and eight of faulty ones, including the transient states. From each file, 65% of the data was used for the model's training, 15% for validation and the remaining 20% for testing, with each part being randomly selected within each file. A window size of size 50 was initially used, consistent with the simulation cases, with the accuracy of the GCN-PE model being 69.70%. The reduced accuracy at this stage can be attributed to the sampling rate difference between the simulation and experimental datasets.

To address this difference in sampling rate, the window size was increased to account for the higher sampling of the experimental case; higher sampling rate indicates that longer windows are needed to capture the effects of the faults within signals. After testing different window lengths, a window of size 350 was selected as optimal, resulting in a considerable improvement of 89.71% accuracy, as can be seen in Fig. 10(a). Subsequently, the GCN-LSTM method was applied to the same experimental data. For comparison, the proposed GCNN model was first trained and tested using a window size of 50, which resulted in an accuracy of 96.73%. As expected, increasing the window size to 350 led to an increased accuracy of 98.18%, as shown in Figure 10(b), verifying the model's capability to distinguish between healthy and faulty cases. Additionally, as observed from Fig. 10(d), there is notionally a slight overfitting corresponding to the GCN-PE model – thus, highlighting the importance of incorporating the LSTM layer. As illustrated in Fig. 10(e), while the validation accuracy shows some fluctuations across epochs, it remains consistently high, indicating the model's ability to generalise well.

As in the simulations, PCA was applied to three stages of the inference procedure with the experimental data -on the raw input features, the hidden-state embeddings extracted from the GCN-LSTM network, and the final logits before the softmax layer. The first two principal components are shown in Figure 11, confirming that, similarly to the simulated case, the healthy and faulty data are entangled in their original space (raw data). The model performs the required disentanglement, exploiting the temporal characteristics and correlations of the signals (gradual disentanglement in the hidden state and, almost, complete disentanglement in the logit space).

2) Experimental Validation in No-Load Operation

To further investigate the fault-identification potential of the proposed methodology, its generalisation capability was evaluated in cases with no mechanical load applied to the motor, aligning with the no-load investigation performed with the simulation data case study. A set of experiments was performed in no-load driving and the corresponding no-load dataset was acquired. The model was trained exclusively on data derived of cases where the motor operated with with load, and it was tested on detecting the fault for the later no-load dataset. The examination of the scenario is a key finding, as the identification of faults in no-load can be challenging when using traditional signal processing approaches [6] - [42].

Similarly to the application presented in the previous section, the model's extrapolation capabilities were evaluated in no-load conditions. In this stage, the GCN-LSTM model was trained and validated based on the with-load cases. The model's accuracy was evaluated exclusively on the unseen-by-the-model no-load data, again including the motor starting transient. Using a window of size 350, the model achieved a classification accuracy of 92.79% on the no-load dataset, showing its ability to generalise and detect faulty patterns even in no-load scenarios. As shown in the confusion matrix in Fig. 10(c), the model maintained high accuracy for faulty cases, highlighting its robustness in capturing key temporal dependencies among the sensor channels. This result is particularly significant because it indicates that the temporal features captured by the model are quite general and not specific to the loads used for the training data acquisition. As mentioned, the training and testing data also include transient states, highlighting the capability of the model to identify proper temporal features which enable accurate classifications for these states. The proposed PE-enhanced GCN-LSTM framework effectively used the knowledge acquired from the loaded operating states to predict faults for no-load stages, without need for additional retraining. This result confirms the method's capability to be used in real-world applications, where no-load, low-load, and transient states are common.

V. Conclusions

This manuscript presented the application of Granger causality-informed GCNs for the reliable detection of rotor faults in induction motors. Different rotor bar breakage scenarios were evaluated, including the cases of single-bar fault and challenging cases of non-adjacent cage breakages. The healthy and rotor fault scenarios were initially evaluated with extensive data-sets from transient electromagnetic 2D finite element simulations of two $4\ kW$ induction motors with different rotor bar number. Then, the proposed methodology was verified experimentally via a range of laboratory measurements at several levels of loading, thus achieving non-invasive and highly accurate rotor fault detection across a broad range of rotor fault scenarios in various loading levels.

The proposed machine learning framework was investigated in a range of loading conditions and the Granger causality analysis was enhanced with the addition of permutation entropy to reduce the sample space used for the computation of the causal matrix. As such, the accuracy scores were increased while achieving a computational cost reduction. The architecture was extended with the incorporation of a two-layer LSTM model resulting in enhanced classification and higher accuracy. The model was assessed using both random train-test split and generalisation tests by examining the accuracy in "unseen" files during training. Both cases achieved high accuracy scores surpassing 98%. The model's generalisation ability was evaluated on a set of "No-Load" data, which is a challenging condition for traditional approaches. These data were excluded from the training process to ensure fair assessment. In a separate trial, the current signal features were removed due to the noise they potentially add to the training process. The same procedure was followed to test the proposed model's performance on experimental data and its accuracy

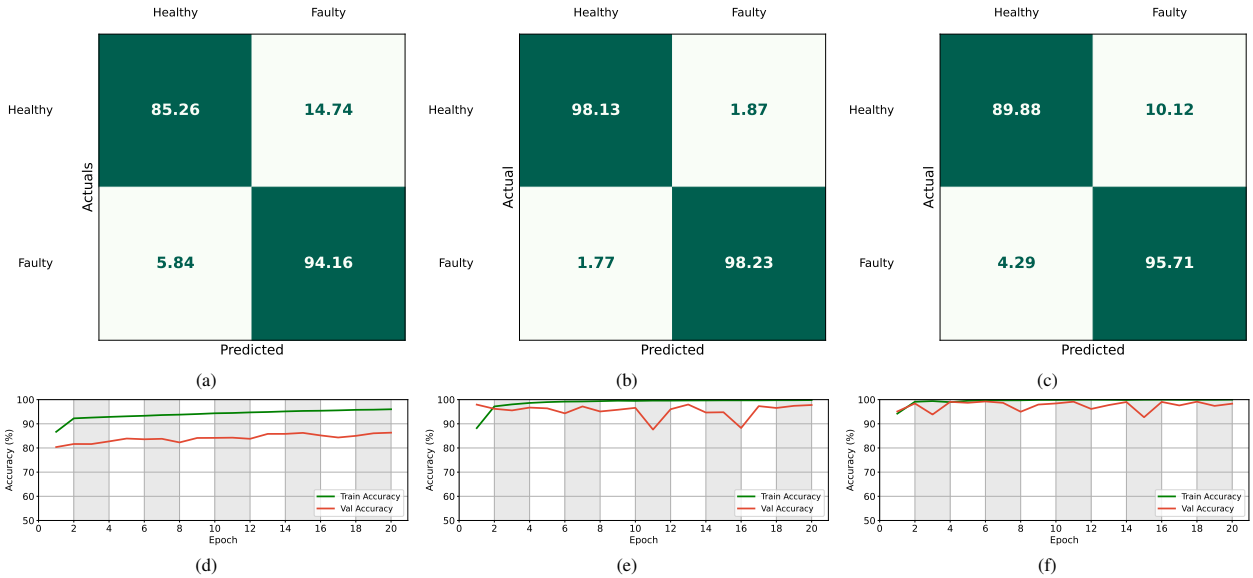

Fig. 10. Experimental results - Confusion matrix for: a) the GCN-PE model, b) the GCN-PE-LSTM model, c) generalisation to no-load data, and Learning Curves of: d) the GCN-PE architecture e) the GCN-PE-LSTM architecture, f) no-load data.

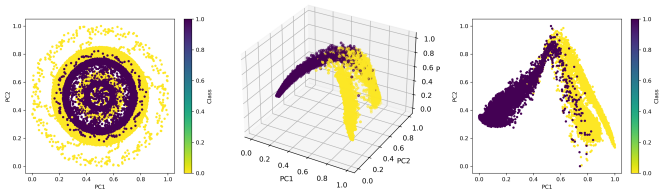

Fig. 11. PCA in three stages for experimental dataset: (a) raw input windows (3D), (b) hidden-state embeddings produced by the GCN-LSTM model, (c) final two-dimensional logits before softmax;

was higher than 98%. Similarly, the model performed classification with accuracy of 92.79% on unseen no-load data, which is quite challenging to tackle with traditional approaches. In all examined cases, the model yielded results with high classification accuracy, indicating great generalisation ability, and independence from load-related features. This architecture can be adapted to problems involving a higher number of sensors by increasing the number of nodes, modifying the Granger causality matrix computation, or adjusting the window size.

Aspects of future work intend to explore the hardware implementation of the proposed machine learning framework and the generalisation of the architecture across various machine types and power ratings. Complementary techniques in combination with Granger causality analysis may be derived to facilitate edge weights. Further, the architecture is to be expanded into the classification between multiple fault conditions and for estimation of the fault severity within classes. Another aspect of future work pertains to the real-time implementation of the proposed approach, either by tools such as FPGA or on-chip implementation of the machine learning framework.

References

[1] J. Faiz, V. Ghorbanian, and G. Joksimovic, *Fault diagnosis of induction motors*. London, UK: Institution of Engineering and Technology, 2017.

[2] P. Zhang, Y. Du, T. G. Habetler, and B. Lu, "A survey of condition monitoring and protection methods for medium-voltage induction motors," *IEEE Transactions on Industry Applications*, vol. 47, no. 1, pp. 34–46, Jan. 2011.

[3] G. Madescu, M. Biriescu, L.-N. Tutelea, M. Mot, and I. Boldea, "Fast computation of bar over-currents in the faulty cage of induction motors with experimental validation," in *2017 International Conference on Optimization of Electrical and Electronic Equipment (OPTIM) & 2017 Intl Aegean Conference on Electrical Machines and Power Electronics (ACEMP)*, Brasov, Romania, 2017.

[4] Y. Park, H. Choi, S. B. Lee, and K. N. Gyftakis, "Search coil-based detection of nonadjacent rotor bar damage in squirrel cage induction motors," *IEEE Transactions on Industry Applications*, vol. 56, no. 5, pp. 4748–4757, 2020.

[5] S. B. Lee *et al.*, "Condition monitoring of industrial electric machines: State of the art and future challenges," *IEEE Industrial Electronics Magazine*, vol. 14, no. 4, pp. 158–167, 2020.

[6] K. N. Gyftakis, J. A. Antonino-Daviu, R. Garcia-Hernandez, M. D. McCulloch, D. A. Howey, and A. J. M. Cardoso, "Comparative experimental investigation of broken bar fault detectability in induction motors," *IEEE Transactions on Industry Applications*, vol. 52, no. 2, pp. 1452–1459, 2015.

[7] A. Sapena-Bano, M. Pineda-Sanchez, R. Puche-Panadero, J. Martinez-Roman, and Ž. Kanović, "Low-cost diagnosis of rotor asymmetries in induction machines working at a very low slip using the reduced envelope of the stator current," *IEEE Transactions on Energy Conversion*, vol. 30, no. 4, pp. 1409–1419, 2015.

[8] K. C. D. Kompella, N. S. Rongala, S. R. Rayapudi, and V. G. R. Mannam, "Robustification of fault detection algorithm in a three-phase induction motor using mcsa for various single and multiple faults," *IET Electrical Power Applications*, vol. 15, pp. 593–615, 2021.

[9] R. R. Kumar, M. Andriollo, G. Cirrincione, M. Cirrincione, and A. Tortella, "A comprehensive review of conventional and intelligence-based approaches for the fault diagnosis and condition monitoring of induction motors," *Energies*, vol. 15, no. 23, p. 8938, 2022.

[10] L. P. Chisedzi and M. Muteba, "Detection of broken rotor bars in cage induction motors using machine learning methods," *Sensors*, vol. 23, no. 22, Art. no. 9079, 2023, doi: 10.3390/s23229079.

[11] S. Gundewar, P. Kane, and A. Andhare, "Detection of broken rotor bar fault in an induction motor using convolution neural network," *Journal of Advanced Mechanical Design, Systems, and Manufacturing*, vol. 16, Art. no. JAMDSM0020, 2022, doi: 10.1299/jamdsm.2022jamdsm0020.

[12] R. Yan, R. X. Gao, and X. Chen, "Wavelets for fault diagnosis of rotary machines: A review with applications," *Signal Processing*, vol. 96, pp. 1–15, 2014.

[13] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, http://www.deeplearningbook.org.

[14] X. Xiao, C. Li, H. He, J. Huang, and T. Yu, "Rotating machinery fault diagnosis method based on multi-level fusion framework of multi-sensor information," *Information Fusion*, vol. 113, p. 102621, 2025.

[15] S. Nandi, H. A. Toliyat, and X. Li, "Condition monitoring and fault diagnosis of electrical motors—a review," *IEEE Transactions on Energy Conversion*, vol. 20, no. 4, pp. 719–729, Dec. 2005.

[16] P. Kumar and A. S. Hati, "Review on machine learning algorithm based fault detection in induction motors," *Archives of Computational Methods in Engineering*, vol. 28, pp. 1929–1940, 2021.

[17] T. Ince, S. Kiranyaz, L. Eren, M. Askar, and M. Gabbouj, "Real-time motor fault detection by 1-d convolutional neural networks," *IEEE Transactions on Industrial Electronics*, vol. 63, no. 11, pp. 7067–7075, Nov. 2016.

[18] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[19] F. Duman Keles, P. M. Wijewardena, and C. Hegde, "On the Computational Complexity of Self-Attention," in *Proceedings of the 34th International Conference on Algorithmic Learning Theory (ALT)*, S. Agrawal and F. Orabona, Eds., *Proceedings of Machine Learning Research*, vol. 201, pp. 597–619, Feb. 2023.

[20] S. Gawde, S. Patil, S. Kumar, and K. Kotecha, "A scoping review on multi-fault diagnosis of industrial rotating machines using multi-sensor data fusion," *Artificial Intelligence Review*, vol. 56, no. 5, pp. 4711–4764, May 2023, doi: 10.1007/s10462-022-10243-z.

[21] Z. Huo, M. Martínez-García, Y. Zhang, R. Yan, and L. Shu, "Entropy Measures in Machine Fault Diagnosis: Insights and Applications," *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 6, pp. 2607–2620, 2020, doi: 10.1109/TIM.2020.2981220.

[22] S. Pan, T. Han, A. C. C. Tan, and T. R. Lin, "Fault Diagnosis System of Induction Motors Based on Multiscale Entropy and Support Vector Machine with Mutual Information Algorithm," *Shock and Vibration*, vol. 2016, no. 1, Art. ID 5836717, 2016, doi: 10.1155/2016/5836717.

[23] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," arXiv preprint arXiv:1609.02907, 2016.

[24] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4–24, Jan. 2021.

[25] Y. Gao, M. Chen, and D. Yu, "Semi-supervised graph convolutional network and its application in intelligent fault diagnosis of rotating machinery," *Measurement*, vol. 186, p. 110084, 2021.

[26] M. Gan *et al.*, "Few-shot fault diagnosis based on multi-scale graph convolution filtering for industry," arXiv preprint arXiv:2405.19642, 2024.

[27] D. Plavos, N. Tsinnas, L. S. Sofikitis, Z. Song, T. Deng, and P. A. Panagiotou, "Classification of Induction Motor Rotor Faults Using Graph Convolutional Neural Networks and Non-Intrusive Condition Monitoring Signals," in *Proc. 2025 IEEE Symposium on Diagnostics for Electric Machines, Power Electronics and Drives (SDEMPED)*, pp. 1–7, IEEE, 2025.

[28] X. Yu, B. Tang, and L. Deng, "Fault diagnosis of rotating machinery based on graph weighted reinforcement networks under small samples and strong noise," *Mechanical Systems and Signal Processing*, vol. 186, p. 109848, 2023.

[29] Z. Zhang and L. Wu, "Graph neural network-based bearing fault diagnosis using granger causality test," *Expert Systems with Applications*, vol. 242, p. 122827, 2024.

[30] Y. Liu and B. Jafarpour, "Graph attention network with granger causality map for fault detection and root cause diagnosis," *Computers and Chemical Engineering*, vol. 180, p. 108453, 2024.

[31] R. Sabir, D. Rosato, S. Hartmann, and C. Guehmann, "Lstm based bearing fault diagnosis of electrical machines using motor current signal," in *Proc. 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, Boca Raton, FL, USA, 2019, pp. 613–618.

[32] K. N. Gyftakis, A. J. M. Cardoso, and J. A. Antonino-Daviu, "Introducing the Filtered Park's and Filtered Extended Park's Vector Approach to detect broken rotor bars in induction motors independently from the rotor slots number," *Mechanical Systems and Signal Processing*, vol. 93, pp. 30–50, 2017.

[33] C. G. Dias and F. H. Pereira, "Broken rotor bars detection in induction motors running at very low slip using a Hall effect sensor," *IEEE Sensors Journal*, vol. 18, no. 11, pp. 4602–4613, 2018.

[34] B. Bessam, A. Menacer, M. Boumehraz, and H. Cherif, "Detection of broken rotor bar faults in induction motor at low load using neural network," *ISA Transactions*, vol. 64, pp. 241–246, 2016.

[35] W. Laala, S. Guedini, and S. Zouzou, "Novel approach for diagnosis and detection of broken bar in induction motor at low slip using fuzzy logic," in *Proc. 8th IEEE Symposium on Diagnostics for Electrical Machines, Power Electronics & Drives (SDEMPED)*, pp. 511–516, IEEE, 2011.

[36] D. A. Elvira-Ortiz, D. Morinigo-Sotelo, A. L. Zorita-Lamadrid, R. A. Osornio-Rios, and R. J. Romero-Troncoso, "Genetic Algorithm Methodology for Broken Bar Detection in Induction Motor at Low Frequency and Load Operation," in *Proc. 2019 IEEE 12th International Symposium on Diagnostics for Electrical Machines, Power Electronics and Drives (SDEMPED)*, pp. 459–465, IEEE, 2019.

[37] T. Wray, P. Panagiotou, M. Masoud, H. O'Keeffe, R. Toqeer, A. Griffo, and P. Lazari, "A compact and modular remote access platform for enhanced practical education in power electronics, machines, and drives," in *IET Conference Proceedings CP878*, vol. 2024. IET, 2024, pp. 646–652.

[38] K. N. Gyftakis, M. Salinas, N. Trachalakis, Z. Song, and P. A. Panagiotou, "Zero-sequence flux analysis aiming for reliable detection of broken rotor bars in induction motors," *IEEE Transactions on Industry Applications*, 2025.

[39] J. Bonet-Jara, J. Pons-Llinares, and K. N. Gyftakis, "Comprehensive analysis of principal slot harmonics as reliable indicators for early detection of interturn faults in induction motors of deep-well submersible pumps," *IEEE Transactions on Industrial Electronics*, vol. 70, no. 11, pp. 11692–11702, 2022.

[40] P. Vas, *Parameter Estimation, Condition Monitoring, and Diagnosis of Electrical Machines*. Oxford, UK: Oxford University Press, 1992.

[41] Z. Song, B. Morin, R. Toqeer, J. C. Mayo-Maldonado, K. N. Gyftakis, and P. A. Panagiotou, "A Comparative Investigation of Stator Fault Diagnostics in Induction Motors by Multivariable Monitoring and Finite Element Simulations," in *Proc. 2025 IEEE Symposium on Diagnostics for Electric Machines, Power Electronics and Drives (SDEMPED)*, pp. 1–7, IEEE, 2025.

[42] D. B. B. de Deus, C. A. N. Sobrinho, F. A. Belo, A. V. Brito, J. G. G. de Souza Ramos, and A. C. Lima-Filho, "Density of maxima approach for broken bar fault diagnosis in low slip and variable load conditions of induction motors," *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 12, pp. 9797–9804, 2020.

Dimosthenis Plavos received the B.Sc. degree in Informatics from the Department of Informatics at the Athens University of Economics and Buisiness, Athens, Greece, in 2024.

Since 2024, he has been working at the company Wings ICT Solutions in Athens, Greece, as an AI Research Engineer, focusing on the design of computer vision and machine learning algorithms. In 2025, he was a Visiting Researcher at the School of Electrical and Electronic Engineering, University of Sheffield, Sheffield, UK, as a secondment via the the MSCA under the IPOSEE Horizon Europe project.

His research interests include machine learning, deep learning, computer vision, graph neural networks, fault detection, and intelligent systems for biomedical applications.

Georgios Tsialiamanis received the 5-year Diploma (B.Eng. & M.Eng.) from the Department of Civil Engineering of the National Technical University of Athens, Athens, Greece, in 2017. He received his Ph.D. in machine learning for engineering from the Uni-

versity of Sheffield, Sheffield, U.K., in 2022.

In 2022, he joined the Department of Mechanical Engineering of the University of Sheffield as a postdoctoral research associate for population-based structural health monitoring. Since 2023, he is a Lecturer in Digital Twins at the School of Mechanical, Aerospace, and Civil Engineering of the University of Sheffield. His research interests include machine-learning and hybrid approaches for dynamics modelling, graph-based methods, and population-based structural health monitoring.

Dr. Tsialiamanis is a Chartered Engineer (IMechE, U.K.), and a Fellow of the UK Higher Education Academy (FHEA).

Aristeidis Karnezis received the B.Sc. degree in Applied Mathematics from the University of Crete, Greece, in 2018, the M.Sc. in Quantum Technology from the University of Sussex, Brighton, U.K., in 2019, and the Ph.D. degree in Mechanical Engineering from the University of Sheffield, Sheffield, U.K., in 2023.

In 2023, he joined the Department of Electronic and Electrical Engineering of the University of Sheffield, as a Research Associate working on multimodal sensor fusion for pipe-inspection robots. Since 2025, he is a Research Associate at the latter institution working on AI and machine-learning methods. His research interests include wave propagation, sensor fusion, simultaneous localisation and mapping. Dr. Karnezis is a member of the UK Acoustics Network (UKAN).

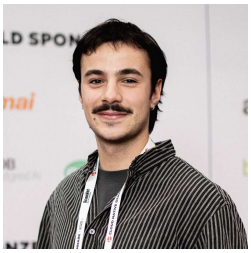

Nikitas Tsinnas received the 5-year Diploma in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 2024.

In 2024, he joined the company Wings ICT Solutions in Athens, Greece, as a Research Engineer. In 2025, he served as a Visiting Researcher at the School of Electrical and Electronic Engineering, University of Sheffield, Sheffield, U.K. He is currently pursuing the M.Sc. degree in communication technologies and system design at the Technical University of Denmark, Lyngby, Denmark. His research interests include machine learning, software development, and system design for industrial and consumer IoT.

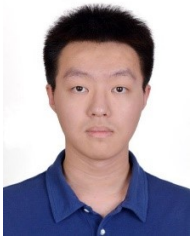

Zihao Song received the B.Eng. and the M.Sc. degrees in electronic and electrical engineering from the University of Sheffield, Sheffield, U.K., Department of Electronic and Electrical Engineering, in 2022 and 2023, respectively.

Since 2023, he is a doctoral candidate in diagnostics for electrical machines with the School of Electrical and Electronic Engineering, University of Sheffield. His research interests include condition monitoring of electrical machines, fault detection and diagnostics, finite element modelling, and development of fault detection methods.

Lazaros S. Sofikitis received the B.Sc. degree in Computer Science from the Department of Informatics of the University of Piraeus, Piraeus, Greece, in 2023.

Since 2024, he joined the company Wings ICT Solutions, Athens, Greece, as a machine learning engineer. In 2025, he was a Visiting Researcher at the School of Electrical and Electronic Engineering of the University of Sheffield, Sheffield, U.K. Currently, he is pursuing the M.Sc. degree in Artificial Intelligence from the Department of Informatics and Computer Engineering at the University of West Attica, Athens, Greece. His research interests include neural networks, deep learning, and predictive modelling for industrial and environmental systems.

Tiantai Deng received the B.Eng. degree from Harbin Institute of Technology, Harbin, China, in 2015, and the Ph.D. degree from Queen's University Belfast, U.K., in 2019. Since 2021, he is a Lecturer with The University of Sheffield, Sheffield, U.K.

Prior to his career as an academic, he was a Senior Engineer with HiSilicon and Huawei. His research interests include hardware acceleration for image processing, deep learning, and high-level design environments. Dr. Deng is a Fellow of the Higher Education Academy (FHEA), U.K.

Panagiotis A. Panagiotou received the 5-year Diploma (B.Eng. & M.Eng.) from the Department of Electrical and Computer Engineering of the University of Patras, Patras, Greece, in 2015, and the M.Sc. degree in Complex Systems and Network Theory from the Department of Mathematics of the Aristotle University of Thessaloniki, Thessaloniki, Greece, in 2017. He received the Ph.D. degree in fault diagnostics of electrical machines from Coventry University, Coventry, U.K., in 2020.

In 2021, he joined the Department of Electronic and Electrical Engineering of the University of Sheffield, Sheffield, U.K., as a postdoctoral research associate at the Rolls-Royce University Technology Centre for Advanced Electrical Machines. Since 2022, he is a Lecturer in electrical machines with the School of Electrical and Electronic Engineering of the University of Sheffield. His research interests include condition monitoring, fault detection and diagnostics, and technologies for the inspection and repair of electrical machines.

Dr. Panagiotou is a member of the IEEE, a member of the Institution of Engineering and Technology (IET, U.K.) and a Fellow of the UK Higher Education Academy (FHEA).