



Article citation info:

Piechocki M, Pajchrowski T, Kraft M, Wolkiewicz M, Ewert P, Unraveling Induction Motor State through Thermal Imaging and Edge Processing: A Step towards Explainable Fault Diagnosis, *Eksploracja i Niezawodność – Maintenance and Reliability* 2023; 25(3) <http://doi.org/10.17531/ein/170114>

## Unraveling Induction Motor State through Thermal Imaging and Edge Processing: A Step towards Explainable Fault Diagnosis

Indexed by:



Mateusz Piechocki<sup>a,\*</sup>, Tomasz Pajchrowski<sup>a</sup>, Marek Kraft<sup>a</sup>, Marcin Wolkiewicz<sup>b</sup>, Paweł Ewert<sup>b</sup>

<sup>a</sup> Institute of Robotics and Machine Intelligence, Poznań University of Technology, Poland

<sup>b</sup> Department of Electrical Machines, Drives and Measurements, Wrocław University of Science and Technology, Poland

### Highlights

- Fault diagnosis is necessary to ensure the electromechanical system's reliability.
- Thermal imaging can be utilized to diagnose the induction motor state.
- Explainability methods provide insights into decision-making process of neural algorithms.
- Convolutional neural networks can efficiently operate on resource-constrained hardware.
- There is no universal approach for deploying deep learning algorithms on edge devices.

### Abstract

Equipment condition monitoring is essential to maintain the reliability of the electromechanical systems. Recently topics related to fault diagnosis have attracted significant interest, rapidly evolving this research area. This study presents a non-invasive method for online state classification of a squirrel-cage induction motor. The solution utilizes thermal imaging for non-contact analysis of thermal changes in machinery. Moreover, used convolutional neural networks (CNNs) streamline extracting relevant features from data and malfunction distinction without defining strict rules. A wide range of neural networks was evaluated to explore the possibilities of the proposed approach and their outputs were verified using model interpretability methods. Besides, the top-performing architectures were optimized and deployed on resource-constrained hardware to examine the system's performance in operating conditions. Overall, the completed tests have confirmed that the proposed approach is feasible, provides accurate results, and successfully operates even when deployed on edge devices.

### Keywords

thermal imaging, fault diagnosis, squirrel-cage induction motor, convolutional neural networks, explainability, edge processing

This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>)

### 1. Introduction

In modern industry, fault detection and diagnosis are critical for sustainable production systems. In recent years, with the expansion of Industry 4.0, predictive maintenance has become the dominant trend in manufacturing, focusing on applications for specific machines or industrial equipment [1]. In areas like mechanical power generation, transportation, or manufacturing, induction motors play a substantial role, generally applied in different types of pumps, fans, blowers, compressors, or conveyors. In the industrial setting, they are

continuously exposed to various stresses that may cause unexpected faults in the induction motor elements. In the initial stage, these deteriorations are unnoticeable but eventually may result in catastrophic failure and high financial costs [7]. According to [37], an hour of downtime may vary from 39 thousand dollars in factories producing Fast Moving Consumer Goods to 2 million dollars in the automotive industry. In addition to financial damage, there is also the possibility of losing reputation and even customers [50].

(\* ) Corresponding author.

E-mail addresses:

M. Piechocki, (ORCID: 0000-0002-3479-0237) [mateusz.piechocki@put.poznan.pl](mailto:mateusz.piechocki@put.poznan.pl), T. Pajchrowski (ORCID: 0000-0002-0002-7161) [tomasz.pajchrowski@put.poznan.pl](mailto:tomasz.pajchrowski@put.poznan.pl), M. Kraft (ORCID: 0000-0001-6483-2357) [marek.kraft@put.poznan.pl](mailto:marek.kraft@put.poznan.pl), M. Wolkiewicz (ORCID: 0000-0003-1197-8517) [marcin.wolkiewicz@pwr.edu.pl](mailto:marcin.wolkiewicz@pwr.edu.pl), P. Ewert (ORCID: 0000-0001-9760-7092) [pawel.ewert@pwr.edu.pl](mailto:pawel.ewert@pwr.edu.pl)

Therefore, the condition monitoring concept was introduced to lower repair-related expenses and minimize the probability of machinery downtime [33].

Nowadays, the most widely used strategies for anomaly detection in industrial applications based on induction motors are current signals analysis [31, 11] and mechanical vibration diagnosis [27, 10, 49]. They are commonly used and well-explored techniques that provide good results. However, these approaches require direct access to device components and additional electrical and/or mechanical connection, which is not always possible due to the hard-to-reach location of the examined object. As another factor, one can mention exorbitant requirements related to computation resources and energy efficiency, which are essential for the real-time operation of algorithms [21]. Moreover, continuous monitoring, data storage, and processing costs might be significant, especially in an environment where many machines are running. In this case, external and non-invasive observation methods might be preferable, even considering the limitations caused by the lack of direct access to the analyzed equipment. The best-known contactless measurement method is imaging. Its basic form uses visible light range and represents it through 3 channels as RGB values. However, in industrial fields, particularly in equipment monitoring, infrared thermography (IRT) is a more popular technique. Thermal imaging allows the detection of thermal changes occurring in the machine and caused by equipment deterioration, environmental conditions, or human factors and comparison with the results to ones characteristic of normal operating conditions [52, 4]. Despite their constraints, such as low resolution and field of view, limited temperature range, limited measurement frequency, sensitivity to environmental factors, and calibration requirements when changing the operating environment, methods based on thermal imaging are more and more willingly applied in the equipment condition monitoring task. Their main advantages are lack of impact on the device under observation, monitoring efficiency, portability, and operational simplicity [13], making these techniques a good choice when a general analysis of active machine is sufficient and exact device parameters, as in the case of current or vibration measurements, are not required.

Over the last few years, several works have employed

thermal imaging for fault diagnosis tasks in induction motors and their external components. The first to be mentioned is [17], in which the authors focus on rotor imbalance and the bearings of rotating machinery. The approach uses two separate image-processing pipelines, first checks the rotor imbalance by differencing the consecutive images and extracting features from the obtained differential frame. The other is based on three characteristics - the standard deviation of the pixel values, a factor related to spatial temperature distribution, and the Gini coefficient, which distinguishes between one of the four bearing states. Overall, the pipelines use classic machine learning algorithms, support vector machine (SVM), or random decision forest (RDF), and achieve 88.25% accuracy. The authors of [40, 41] have chosen another solution to deal with inter-turn faults and cooling system failure. They monitor the trend of stator temperature rise and analyze thermal distribution in a motor's transient state. On this basis, the authors have determined the severity of the occurrence and proposed a segmentation method based on sophisticated temperature thresholding, useful for visualization of the machinery in a steady state. The authors of [18] have moved away from the feature extraction approach in favor of convolutional neural networks (CNNs). The proposed solution is based on two combined VGG models [39], for respectively spatial and temporal characteristics, and can distinguish between eight classes with 95% accuracy. Another method adopting convolutional neural networks to thermal image analysis is [23], whose authors tackle the bevel gearbox defects problem with correctness of nearly 100%. Authors of [22] compared several classification methods on two experimental, ten-class datasets involving bearing and rotor faults. The experiments have proven that convolutional neural networks outperform other algorithms, reaching accuracies of above 98% on both test sets. Karabacak et al. [19] presented a worm gear condition monitoring and fault detection method based on thermal images and modified GoogLeNet [43] architecture. Their collected dataset consists of four classes, which include three defect types: wear, pitting, and tooth breakage. The authors of the proposed algorithm declare a correct classification of all samples from the test set. Mian et al. [29] proposed a redundant multi-sensor pipeline for misalignment and unbalanced state diagnosis. The approach

based on vibration signals and infrared thermal imaging independently extracts features and classifies the condition of the induction motor. In the presented results, the vision part performs better and achieves perfect accuracy. Experiments presented in [35] allow distinguishing between four states of an induction motor: healthy condition, a broken rotor bar, bearing damage, and misalignment. This methodology takes meaningful statistical characteristics that represent the thermal behavior of the electromechanical system directly from the thermogram. Then, using linear discriminant analysis (LDA), the number of features is reduced and fed into a k-nearest neighbor (k-NN) classification algorithm, which flawlessly classifies them. In [14], authors focused on bearing faults, broken rotor bar defects, and stator failures. In their study, the grey-level co-occurrence matrices (GLCM) were adopted for the feature extraction process. Using the XGBoost classifier, the solution achieved 93% accuracy in evaluation.

The works mentioned above focus on developing new algorithms or applying existing ones to collected data, completely ignoring aspects related to the implementation of the solution, its efficiency, deployment, or performance under target operating conditions. Only the [2] study introduces the concept of an intelligent sensor, the perception unit integrated directly with processing hardware. The authors combined a low-cost infrared sensor with a Raspberry Pi 4 tiny computer. This efficient and powerful edge device performs calculations on the platform without unnecessary data transfer to an external computing unit. In their processing pipeline, thermal images are upscaled ten times, and then for each initially defined region of interest, 15 statistical indicators are calculated. In the next step, these features are grouped in a matrix, and their dimensionality is reduced using principal component analysis (PCA). Finally, the reduced representation is given to the feed-forward neural network (FFNN), responsible for fault detection in the gearbox and rolling bearings. Although the authors of the aforementioned manuscript included the hardware implementation of the proposed solution in their study, it neither contains a comprehensive analysis of the device's performance capabilities nor a comparison of different classification methods.

This paper presents a contactless, non-invasive fault

diagnosis and state monitoring application for squirrel-cage induction motors. The proposed approach processes thermal images and decides whether the machine captured in the input frame belongs to the healthy class or represents one of the predefined damage types. The concept of this solution has a foundation in convolutional neural networks. Accordingly, various architectures were investigated to select the most suitable structure regarding metrics, efficiency, and performance. Considering that the designed method is intended for online, near real-time operation, mostly under changing conditions, selected neural models were optimized and deployed on resource-constrained hardware. This approach allows onboard computing without transmitting measurements to the server and waiting for a response. It significantly reduces operating costs because cloud instance is no longer necessary. Moreover, the device itself performs all calculations, so there are no issues related to communication with a server, and the designed module, due to its compact size, can be freely transferred between different locations. It makes the proposed method applicable not only in the controlled laboratory setting. Last but not least, to not treat the algorithm as a black box, layers and trained weights of the best-performing model were examined utilizing interpretability methods to verify the decision-making context and information taken into account when performing the predictions. Table 1 compares and summarizes studies considering the used IRT resolution, applied methodology, and hardware application.

Sections in this article are organized as follows: Section 2 describes the experimental setup used in experiments, the collected thermal image set, the methodology of the conducted research, and elements of hardware evaluation. Section 3 presents the achieved metrics and measured performance of architectures on edge devices. Additionally, this chapter pays attention to the interpretability of the results, demonstrating the outcomes of a range of explainability analysis algorithms applied to the developed models. Finally, Section 4 discusses the findings, whereas Section 5 summarizes the work and outlines the possibilities for further research.

Tab. 1. Comparison of applications that use thermal diagnosis.

Types of Faults	IRT Resolution	Method	Hardware Application	Article
Bearing damage and rotor imbalance	640×480	SVM + RDF	None	[17]
Inter-turn	320×240	Temperature Thresholding	None	[40]
Inter-turn and cooling system	320×240	Temperature Thresholding	None	[41]
Bearing	640×480	Two combined VGG CNNs	NVIDIA GeForce GTX TITAN X	[18]
Bevel gearbox	384×288	CNN	None	[23]
Bearing and rotor	384×288	LeNet-5	None	[22]
Worm gear	160×120	GoogLeNet CNN	None	[19]
Misalignment and unbalance detection	464×348	SVM	None	[29]
Broken rotor bar, bearing damage,	320×240	LDA + k-NN	None	[35]
Bearings, broken rotor bar, and stator	80×60	GLCM + XGBoost	None	[14]
Rolling bearings and gearbox	160×120	PCA + FFNN	Raspberry Pi 4	[2]
Broken rotor and misalignment	640×512	CNN	Raspberry Pi 4B	Proposed

## 2. Material and methods

This section describes the experimental setup used to simulate the potential states of the squirrel-cage induction motor and the gathered data utilized in the study. Moreover, this chapter includes a comprehensive description of the research methodology. In addition, the last subsection contains a detailed description of evaluation hardware and frameworks, which are a meaningful aspect of the proposed approach considering its potential applicability.

### 2.1. Experimental Setup

The performed experiments have been divided considering the type of fault. Therefore, in the experimental setup, three classes have been distinguished:

- Healthy - an induction motor works correctly under normal operating conditions.
- Misalignment - a shift had been added between shafts making them no longer parallel.
- Broken Rotor - squirrel-cage rotor with broken cages.

The presented research includes four varied misalignments, whereas, in the case of the broken rotor, three

versions have been considered - with 1, 3, and 6 cages broken. All of the above configurations were tested with different levels of current load - starting with no load and ending with 6A, with a 2A step. Additionally, to increase the diversity of data and introduce hardware variations, the examinations were conducted using several couplings.

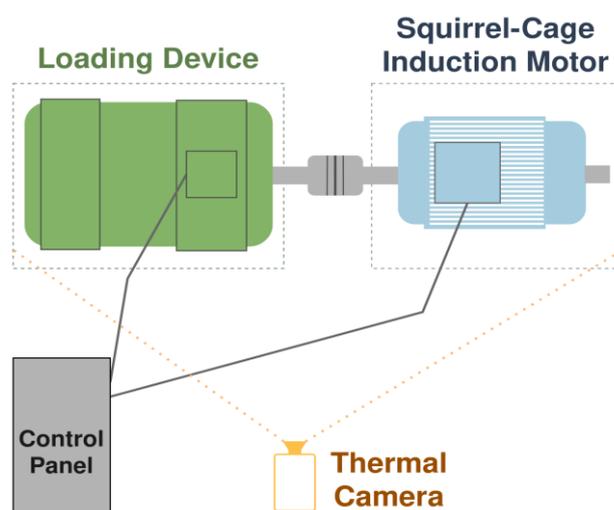


Fig. 1. The schematic diagram of the experimental test stand.

All in all, the entire study contains 42 independent trials with different hardware setups and current loads. The

illustrative diagram of the test bench is depicted in Figure 1. Meanwhile, Table 2 contains detailed configurations of the experimental setup.

Tab. 2. Configuration of examined Squirrel-Cage Induction Motor.

Parameter	Value
Producer	Celma Indukta
Type	Sh 90L-4
Phase	3-phase
Rated power	1.5 kW
Input current	6.1 / 3.5 A
Rated voltage	230 / 400 V
Rated frequency	50 Hz
Efficiency	79%
Speed	1410 RPM
Shaft diameter	24 mm

## 2.2. Dataset

In accordance with Figure 1, the operation of the system was registered using Workswell InfraRed Camera (WIC) [54]. It is a highly accurate sensor, with thermal sensitivity of up to  $\leq 0.03^{\circ}\text{C}$ , designed for industrial inspections and R&D applications. In this study, WIC 640 thermal camera was utilized. It has  $640 \times 512$ px resolution, and the detailed specification is presented in Table 3. Images were captured in series lasting about 30 seconds, with an average of four frames per second, which results in approximately 120 unique images per experiment. The collected data, 5653 frames in

total, were grouped into three classes corresponding to the induction motor state or damage and then randomly divided into five folds to perform cross-validation in the next stage. The dataset distribution is presented in Table 4, whereas Figure 2 shows sample thermal frames from the collected image set, respectively a healthy system, test with added misalignment, and instance with a broken rotor. In addition, Table 5 presents measured temperature ranges for respective classes from the dataset. Figure 3 depicts the temperature distribution for each experiment. The presented values from conducted trials were clustered to highlight the differences between groups, for example, due to different couplings or by increasing the shift between shafts in subsequent trials.

Tab. 3. Thermal camera specification.

Parameter	Value
Manufacturer	Workswell
Image resolution	$640 \times 512$
Frame rate	4 FPS
Temperature measurement range	$-40^{\circ}\text{C} \text{ --- } +550^{\circ}\text{C}$
Thermal sensitivity	$0.03^{\circ}\text{C}$

Tab. 4. Dataset distribution for 5-fold cross-validation.

Class	Fold					Total
	0	1	2	3	4	
Healthy	186	879	560	306	313	2244
Misalignment	311	314	417	381	376	1799
Broken Rotor	335	502	211	279	283	1610

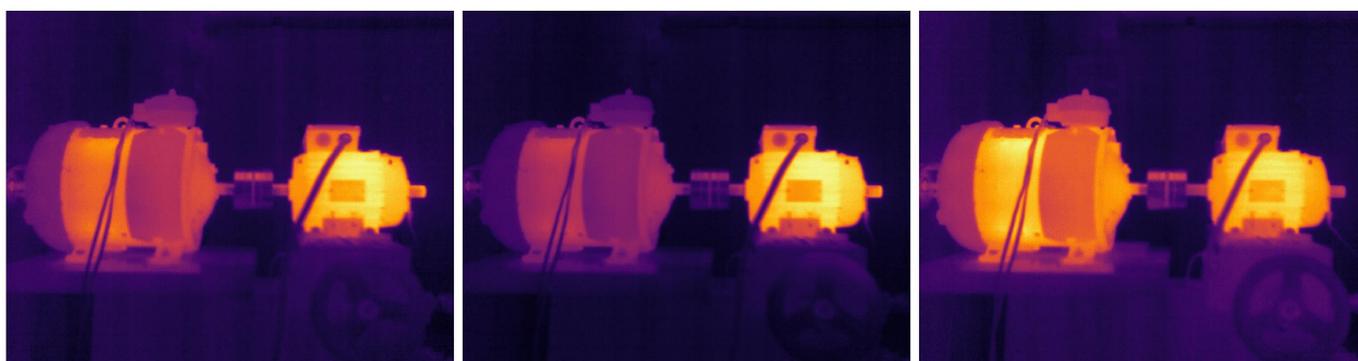


Fig. 2. Sample thermal images captured by Workswell WIC 640 InfraRed Camera in experiments with 4A current load. (left) properly functioning system; (center) misalignment between two shafts; (right) motor with a broken rotor.

Tab. 5. Registered temperature ranges for individual classes.

Class	Min Value	Max Value	Mean Value	Standard Deviation
Healthy	23.00°C	82.43°C	38.62°C	8.86°C
Misalignment	25.52°C	104.99°C	40.31°C	12.04°C
Broken Rotor	24.86°C	83.30°C	41.24°C	11.09°C

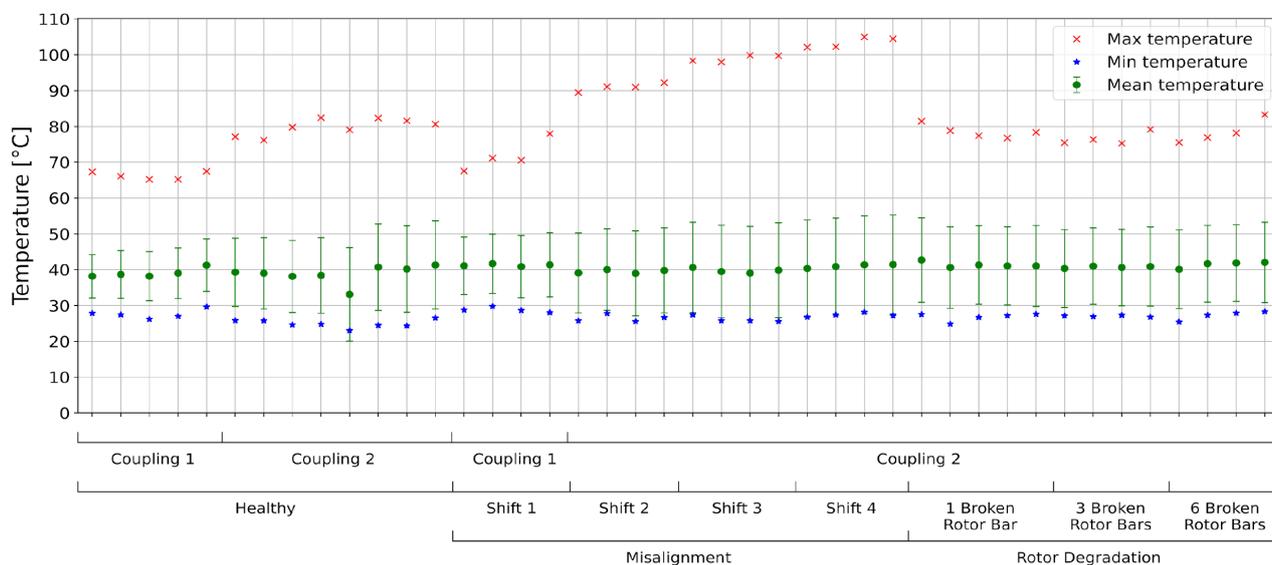


Fig. 3. Registered temperature ranges for individual experiments.

### 2.3. Methodology

In most cases, stopping the running machinery for inspection is not an option, especially in industrial applications. Moreover, induction motors are often in places where direct access is impossible or highly limited. Therefore, the presented work focuses on a contactless, non-invasive solution that allows near real-time measurement and on-device analysis using edge devices. As a starting point, PPLCNet [6] was chosen due to its high efficiency on CPU-based hardware. Additionally, the MNASNet [44] in the variant with Squeeze-and-Excitation (SEMNASNet), in three variants, was taken into account. ResNet18 and ResNet34 [12] were considered for their generalization capabilities, proven in a wide range of computer vision applications [56, 25, 32]. Along with them, ResNet10 [8] architecture was included. It is tailored to resource-constrained platforms with low computing power, using 10 instead of 18 residual layers and owing ten times fewer parameters than ResNet18. Another choice was EfficientNet [46] in versions B0-B2. Unlike conventional architectures, these models uniformly scale all network dimensions - width, depth, and resolution. This approach allows for achieving good performance in terms of

accuracy while maintaining high performance. The last architecture selected for evaluation is MixNet [45], in the smallest one, the S option. This architecture utilizes a variety of layers called MixConv, which combine multiple, different size kernels into a single depthwise convolution operation, enhancing the process of feature extraction from the input. The MixNet structure starts with small kernels to save computational cost and gradually increases their size to improve final accuracy. This approach allows taking advantage of large kernel sizes, and picking up high-resolution patterns, without accuracy and efficiency deterioration. The discussed network architectures were implemented utilizing the PyTorch Image Models library [53], and the whole training pipeline was prepared using PyTorch [47] package.

In the proposed process, thermal images were a single-channel input to the convolutional neural networks with unchanged image resolution. The only preprocessing steps used were conversion from 16-bit unsigned integer format to 32-bit floating-point data type, followed by min-max scaling to the 0-1 range. Figure 4 illustrates the flow chart of the designed approach.

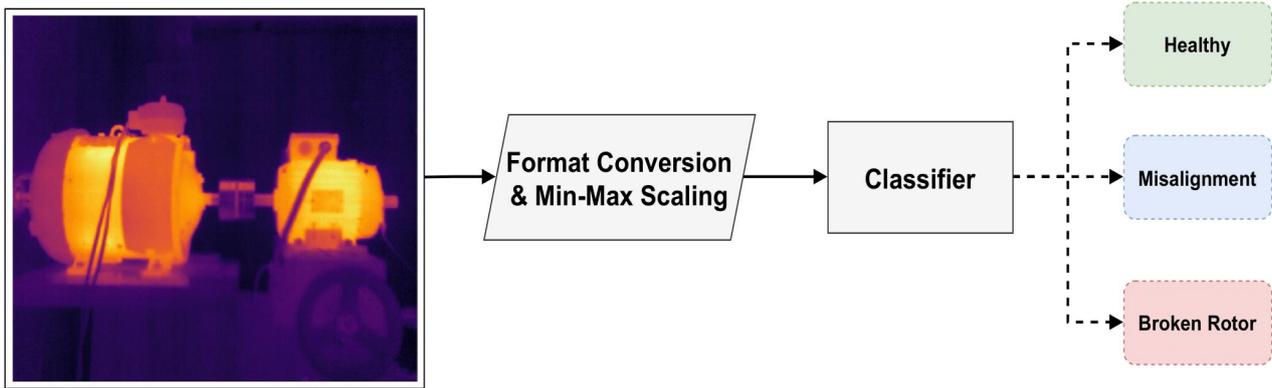


Fig. 4. Developed pipeline for the state classification of a squirrel-cage induction motor.

The dataset has been randomly split into five folds to perform cross-validation. For every split, three folds were concatenated and marked as the training set, and the other two were used respectively as validation and test set. As an optimizer, AdamW [26] was selected due to its computational efficiency, simplicity of configuration, and use of decoupled weight decay instead of L2 regularization, which results in substantially better generalization performance. The initial learning rate was set to  $3e-5$ . Since the problem is a 3-class classification, cross-entropy was chosen as a loss function. The batch size was set depending on architecture size and number of parameters, respectively 16 for simpler ones and 8 for more complex ones. The maximum number of epochs was 300. However, the early stopping callback, with patience equal to 10, was employed to monitor changes in the loss function on the validation set and stop training if there is no improvement. The NVIDIA TITAN Xp GPU with 12 GB memory and CUDA 11.8 was configured for all experiments.

#### 2.4. Evaluation Hardware and Frameworks

Considering the potential application of the proposed method, the neural network architectures described in Section 2.3 were optimized, quantized, deployed, and benchmarked on resource-constrained hardware. The following edge devices were selected for this comparison due to their recognition and common use in IoT applications, market availability, as well as producer and community support:

- **Raspberry Pi 4B (RPi)** [48] is the most used single-board computer with a quad-core ARM Cortex A72 processor. The board had 2GB of RAM and a 64-bit operating system in the experimental setup. This device is a good representation of modern low-

power, small-size embedded computer performance.

- **Intel Neural Compute Stick 2 (NCS2)** [15] is a deep learning algorithms accelerator. It is a small, fanless USB module designed for IoT or robotics applications, especially employing computer vision models. The device's high performance is owed to 16 efficient SHAVE (streaming hybrid architecture vector engine) VLIW cores (very long instruction word), and a dedicated hardware neural accelerator, which is the leading part of the module. This chip is named Intel Movidius Myriad X Vision Processing Unit (VPU) and clocked at 700 MHz base frequency. The NCS2 is supported directly by the Intel OpenVINO toolkit, developed to optimize, deploy, and infer deep learning models, specifically convolutional neural networks. The Intel Movidius Myriad X VPU supports solely floating-point operations. Therefore, executed models can have float 32 or float 16 data format. In our application, the accelerator is a USB coprocessor and requires a host device running an operating system, e.g., a Raspberry Pi.

A detailed description of tools, libraries, and frameworks facilitating optimization, acceleration of computing, and inference on the hardware used in the analysis is shown below, whereas a brief comparison is given in Table 6.

- **OpenVINO** [16] is an open-source set of tools designed and supported by Intel to optimize and deploy AI algorithms on Intel-manufactured hardware, ranging from edge devices to cloud centers. The toolkit was developed to deal with a wide range of hardware, such as central processing

units, graphics processing units, vision processing units, and field-programmable gate arrays. OpenVINO aims to boost the performance of algorithms in computer vision, natural language processing, and automatic speech recognition applications. The principal advantage of the toolkit is the possibility to convert neural networks from one of the mainstream deep learning frameworks such as, e.g., ONNX, TensorFlow to the Intermediate Representation format and then use this model among all supported platforms without significant code changes.

- **tfLite-runtime** (TF Lite) [9] is a high-performance package designed for on-device machine learning algorithms inference and intended for mobile, embedded, and even IoT devices. The library utilizes models in the TensorFlow Lite FlatBuffers format, which is efficient, portable, and supported by several programming languages. This format is available directly through conversion from TensorFlow, and by additional optimization and quantization, aspects such as latency, size, and power consumption can be improved.
- **PyArmNN** [3] is an official Python binding for Arm NN - a machine-learning inference engine for Android and Linux platforms. The package accelerates computations on ARM Mali GPUs and CPUs from the Cortex-A family. In benchmarks, the Arm NN outperforms generic runtime libraries using architecture-tailored optimizations from Arm Compute Library. Arm NN inference engine supports TF Lite models through TF Lite Delegate.

**ONNX Runtime** [30] is an inference framework linking various libraries, which supports hardware acceleration on diverse platforms. This package simplifies the usage of deep learning models in the ONNX format, offering a standardized way of model loading, allowing, at the same time, the use of one of the numerous execution providers created to maximize the performance of the computing environment. The main advantage of ONNX Runtime is the high portability of the designed approach and benefiting from integrated access to stand-alone tools like e.g., OpenVINO

Tab. 6. Brief overview of utilized inference packages.

Framework	Version	Utilized Hardware
OpenVINO	2022.1	Intel NCS2
tfLite-runtime	2.11.0	RPi CPU
PyArmNN	32.0.0	RPi CPU
ONNX Runtime	1.14.1	RPi CPU / Intel NCS2

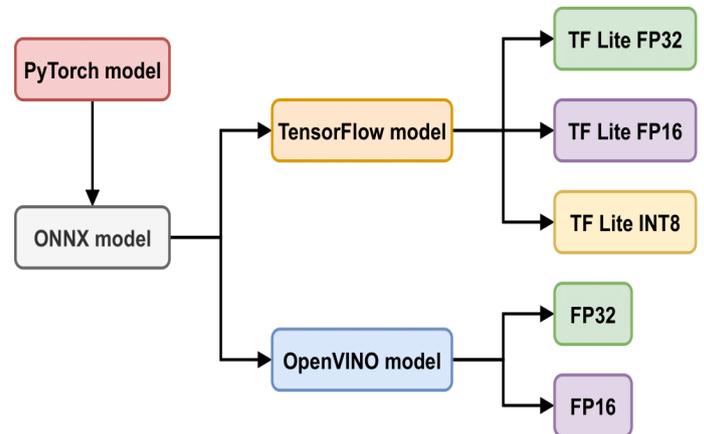


Fig. 5. Model conversion pipeline.

In order to convert the developed neural network architectures from the PyTorch format used in the experimental environment to the suitable form supported by the inference frameworks, it was necessary to perform several actions. The base models were exported to the ONNX and then to TensorFlow Lite and OpenVINO. OpenVINO instances were optimized and quantized directly from the ONNX file, whereas TF Lite requires conversion via TensorFlow. Additionally, weights and activations of final models were quantized to 16-bit floating point format and 8-bit integer data representation. Figure 5 depicts the transformation pipeline.

### 3. Results

In the proposed approach, the state of the squirrel-cage induction motor could be classified into one of three categories: healthy, misalignment, and broken rotor. In the training phase, for research purposes, the 5-fold cross-validation was added to increase the reliability of the neural algorithms, calculating statistics individually for every split. As evaluation metrics for classification models, the accuracy (1), f1 score (2), precision (3), and recall (4) metrics have been selected. Their equations in baseline representations are as follows:

$$Accuracy = \frac{TP + TN}{TP + FN + TN + FP} \quad (1)$$

$$F1Score = \frac{2TP}{2TP + FP + FN} \quad (2)$$

$$Precision = \frac{TP}{TP + FP} \quad (3)$$

$$Recall = \frac{TP}{TP + FN} \quad (4)$$

where  $TP$ ,  $FN$ ,  $TN$ , and  $FP$  refer directly to elements of the commonly used confusion matrix and denote, as follows, true positive, false negative, false positive, and true negative. The true positive and true negative describes correctly classified positive and negative samples. False positive represents a situation when output indicates the presence of a condition, whereas false negative incorrectly signifies its absence.

Considering the presented research problem is a multiclass classification and the distribution of samples in folds is not balanced, the weighted average has been applied after calculating statistics for each label. The final metric from the folds' results was obtained using the standard mean. Table 7

presents the achieved results of 5-fold cross-validation for benchmarked neural architectures.

The growing popularity of methods based on deep neural networks, paired with increasing model complexity, results in a deepening lack of operational transparency. In safety-critical applications, such as healthcare, autonomous driving, or industry, using solutions of a black-box nature may not be an option. Therefore, in recent years, model explainability techniques have become increasingly meaningful, while verification of neural networks has become a crucial part of the development process [24, 5]. The presented case is no different. To get more insights into the mode's operation, instead of treating it as a black box, the best-performing model from Table 7, EfficientNet B1, was examined with the interpretability methods from Captum [20] library. The following methods have been selected: Occlusion [55], Grad-CAM [36], Deconvolution [28], GuidedBackpropagation [42], and Saliency Map [38]. Figure 6 includes the output of explainability methods obtained on sample thermal images from the gathered dataset.

Tab. 7. Average metrics achieved for benchmarked neural network architectures in 5-fold cross-validation.

Model Architecture	Accuracy	F1 Score	Precision	Recall
PP-LCNet 50	0.924 ± 0.078	0.939 ± 0.066	0.979 ± 0.045	0.924 ± 0.078
PP-LCNet 75	0.920 ± 0.096	0.928 ± 0.086	0.954 ± 0.072	0.920 ± 0.096
PP-LCNet 100	0.880 ± 0.102	0.892 ± 0.088	0.930 ± 0.053	0.880 ± 0.102
SEMNASNet 50	0.901 ± 0.091	0.920 ± 0.080	0.964 ± 0.075	0.901 ± 0.091
SEMNASNet 75	0.913 ± 0.054	0.935 ± 0.045	0.979 ± 0.043	0.913 ± 0.054
SEMNASNet 100	0.878 ± 0.108	0.911 ± 0.071	0.978 ± 0.043	0.878 ± 0.108
ResNet10t	0.831 ± 0.167	0.866 ± 0.136	0.952 ± 0.076	0.831 ± 0.167
ResNet18	0.850 ± 0.172	0.879 ± 0.136	0.960 ± 0.061	0.850 ± 0.172
ResNet34	0.925 ± 0.093	0.929 ± 0.086	0.957 ± 0.061	0.925 ± 0.093
EfficientNet B0	0.896 ± 0.141	0.907 ± 0.131	0.946 ± 0.078	0.896 ± 0.141
EfficientNet B1	0.932 ± 0.086	0.941 ± 0.076	0.969 ± 0.046	0.932 ± 0.086
EfficientNet B2	0.854 ± 0.137	0.864 ± 0.129	0.899 ± 0.103	0.854 ± 0.137
MixNet S	0.904 ± 0.088	0.916 ± 0.077	0.947 ± 0.050	0.904 ± 0.088

The primary advantage of explainability methods is the ability to transparently identify features and regions of interest in the image that impact the final predictions made by the neural network. According to all presented algorithms in Figure 6, the main contribution to the misalignment class has an area representing the induction motor. This observation has its reflection in measured temperatures in thermal images, depicted in Figure 3, which were significantly higher compared to others. Regarding samples representing broken rotor bar class, the wires of the load device catch attention, as well as the upper right part of the induction motor. In the case of healthy data, methods point out the contours and edges of the operating test stand, both the squirrel-cage induction motor and the loading device. Overall, all listed approaches turned out to be helpful in the process of understanding the decision-making context of the neural classifier. They confirm

that the utilized algorithm extracts key features from thermal images by focusing entirely on the testbed and ignoring the surrounding background. Such a finding leads to an assumption that the presented approach is robust enough to work in varying industrial conditions and does not make decisions based on spurious correlations.

In addition to the presented results and employed interpretability algorithms, and due to the potential use of the proposed approach in online, near real-time applications directly under working conditions, performance benchmarks on edge devices were carried out. As a primary device, the Raspberry Pi 4B was selected, and to explore further possibilities for increasing the throughput, an external co-processor, Intel NCS2, was added. A detailed description of evaluation hardware and frameworks is given in Section 2.4, whereas Table 8 contains measured results.

Tab. 8. Average inference time, in seconds, measured on evaluation hardware. The abbreviations FP32, FP16, and INT8 used in the table describe the format of the model's weights and parameters: 32-bit floating point, 16-bit floating point, and 8-bit integer. The following error codes are marked in the table: 1\* - Encountered unresolved custom op: PyFunc; 2\* - Aborted; 3\* - Operator not supported; 4\* - Unsupported operation HardSwish.

Model Architecture	TF Lite FP32	TF Lite FP16	TF Lite INT8	PyArmNN INT8	ONNX Runtime CPU FP32	ONNX Runtime MYRIAD FP16	OpenVINO MYRIAD FP16
PP-LCNet 50	0.336 ± 0.006	0.323 ± 0.005	0.228 ± 0.002	0.156 ± 0.003	0.156 ± 0.006	4*	4*
PP-LCNet 75	0.561 ± 0.009	0.560 ± 0.007	0.399 ± 0.003	0.292 ± 0.004	0.253 ± 0.007	4*	4*
PP-LCNet 100	0.764 ± 0.011	0.787 ± 0.009	0.581 ± 0.005	0.513 ± 0.014	0.353 ± 0.008	4*	4*
SEMNASNet 50	0.493 ± 0.009	0.449 ± 0.008	0.363 ± 0.002	0.485 ± 0.025	0.278 ± 0.113	0.084 ± 0.001	0.083 ± 0.001
SEMNASNet 75	0.927 ± 0.050	0.853 ± 0.009	0.745 ± 0.002	1.023 ± 0.016	0.438 ± 0.026	0.125 ± 0.001	0.125 ± 0.001
SEMNASNet 100	1.083 ± 0.017	1.050 ± 0.009	0.852 ± 0.007	1.102 ± 0.018	0.524 ± 0.025	0.143 ± 0.002	0.142 ± 0.001
ResNet10t	1*	1*	1*	1*	0.696 ± 0.035	0.127 ± 0.001	0.127 ± 0.001
ResNet18	1.725 ± 0.045	2.749 ± 0.021	2*	1.429 ± 0.070	0.954 ± 0.030	0.162 ± 0.002	0.161 ± 0.001
ResNet34	3.204 ± 0.018	5.511 ± 0.027	2*	2.167 ± 0.102	2.216 ± 0.134	0.274 ± 0.001	0.274 ± 0.001
EfficientNet B0	1.790 ± 0.029	1.722 ± 0.026	1.386 ± 0.006	1.577 ± 0.130	0.984 ± 0.046	0.314 ± 0.001	0.313 ± 0.002
EfficientNet B1	2.478 ± 0.030	2.487 ± 0.021	1.916 ± 0.004	2.427 ± 0.166	1.313 ± 0.041	0.456 ± 0.002	0.455 ± 0.002
EfficientNet B2	2.493 ± 0.047	2.635 ± 0.039	2.051 ± 0.008	2.135 ± 0.184	1.377 ± 0.034	0.492 ± 0.002	0.491 ± 0.002
MixNet S	1.718 ± 0.061	2.109 ± 0.032	1.150 ± 0.024	3*	0.881 ± 0.021	0.239 ± 0.001	0.238 ± 0.001

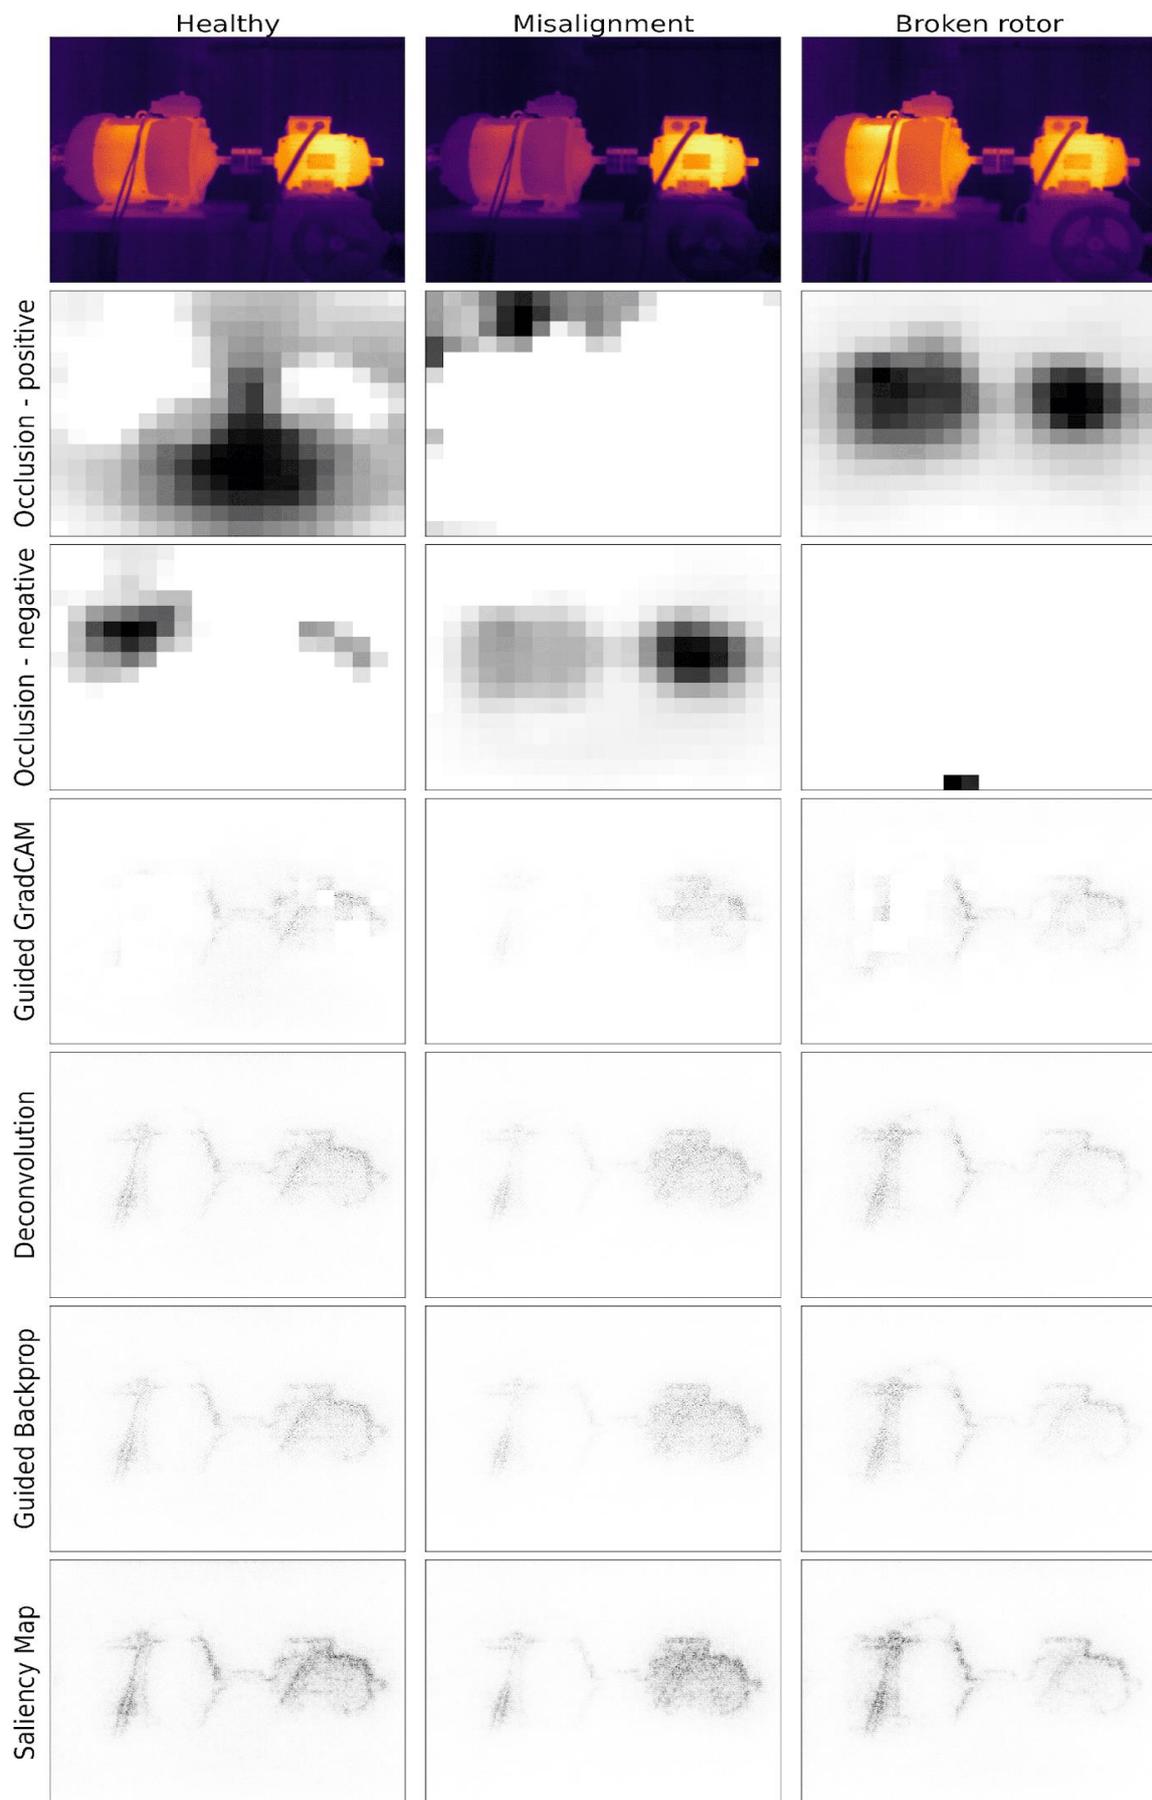


Fig. 6. Visualization of explainability methods applied to the samples from the collected dataset.

As shown in Figure 5, TensorFlow models, optimized for mobile and embedded inference, were quantized to three data representations: 32-bit and 16-bit floating-point and 8-bit signed integer format. The tflite-runtime library was used for inferencing. Models with INT8 data format were additionally evaluated using PyArmNN, a Python language binding for the Arm NN package. Meanwhile, a different model format is utilized by the ONNX Runtime package. It requires, by design, neural networks in ONNX representation. The inference library optimizes and quantizes models directly during loading by employing hardware acceleration providers. With this approach, it was straightforward to check both CPU-based Raspberry Pi 4B and an external accelerator, like Intel NCS2, without significant changes in the test source code. In addition, to maximize the performance of the co-processor

and examine the magnitude of the deceleration related to the utilization of the ONNX Runtime library, the OpenVINO toolkit, with a tailored optimization pipeline and inference engine, was used directly. Figure 7 illustrates the comparison between examined architectures in terms of obtained accuracy, the number of model parameters, and measured inference time on Raspberry Pi 4B CPU with ONNX Runtime framework. From this graph, it is clear that smaller models are processed much faster due to fewer parameters, but despite this fact, they achieve as good results as larger models. Moreover, the group of more complex architectures tends to perform worse due to overfitting to the training set caused by excessive learning capacity compared to the complexity of the task and data.

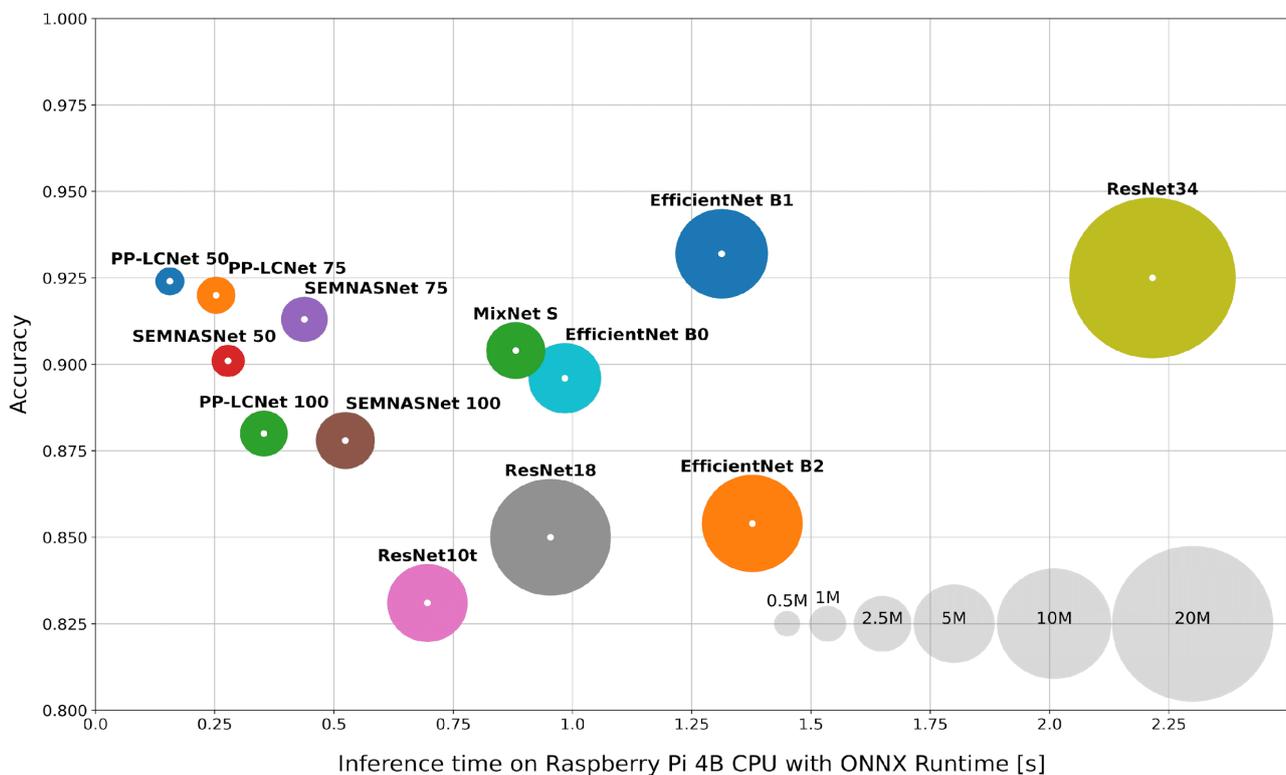


Fig. 7. Comparison of models in terms of achieved accuracy, measured inference time on CPU with ONNX Runtime framework, and the number of parameters. Light gray circles symbolize the number of model parameters in millions.

#### 4. Discussion

The study compares state-of-the-art convolutional neural network architectures utilized in a squirrel-cage induction motor's condition monitoring task. The conducted examinations include three classes with samples produced under various states and configurations. According to Table 7, EfficientNet B1 turned out to be the best-performing

architecture, achieving an average accuracy of 93%, with the lowest standard deviation among folds compared to other models.

Moreover, the interpretability techniques, depicted in Figure 6, helped decode the convolutional neural network hidden under the black box term. The employed methods proved that the presented algorithm focuses on the area

referring to the induction motor and its components. Shafts that are not properly aligned can cause friction increase, which correlates directly to higher power consumption as the machines are less efficient. Ultimately, it translates into a temperature rise within a housing, in the coupling and bearings [51]. The used explainability algorithms clearly depict that the machinery temperature notably increases along with the occurrence of misalignment. On the other hand, healthy rotor bars close to broken rotor bars carry additional current resulting in asymmetrical overheating in the rotor cages [34]. This fault is characterized by the presence of hot spots on the housing, which, with further propagation of the defect, causes thermal bending of the rotor and results in a significant decrease in the reliability of the drive system.

In addition, the proposed solution was successfully evaluated on edge devices showing its ability to operate under time constraints. Moreover, the approach demonstrates its feasibility for industrial applications through the practical use of the presented method, along with the hardware implementation, and not just an evaluation of the algorithms. According to the provided results in Table 8, the ONNX Runtime framework has proved to be the most efficient option for CPU-based processing on Raspberry Pi 4B. On the other hand, if additional computing power is needed or the application requires a higher time regime, the Intel NCS2 co-processor will be a good choice. Carried-out experiments have shown that utilization of ONNX Runtime with MYRIAD FP16 provider, built on the top of the OpenVINO toolkit, at least gives three times better performance than CPU-based computation. It is worth noting that the time difference between the performance of a plain OpenVINO model and one invoked using ONNX Runtime is negligible.

Compared to other applications summarized in Table 1, the proposed approach offers a significant advantage in induction motor fault diagnosis by thermal image analysis. The presented method is contactless and does not require additional modification of the existing electromechanical system setup. Moreover, unlike previously presented diagnostics solutions based on machine learning, it is well suited for deployment using edge processing devices, which was verified in a series of experiments. By efficiently processing higher-resolution thermal images, it is possible to

extract finer details and more precise thermal characteristics of the object under study. This advantage is particularly valuable in equipment condition monitoring applications, where accurate fault diagnosis is critical to ensure operational efficiency and prevent costly failures. What distinguishes this research is the comprehensive evaluation of classification models, focusing on leading neural network architectures. This comparison allows identifying the most suitable model for the task, leading to improved accuracy and performance. The selected models have been specifically optimized for operation on edge devices with low latency, enabling real-time analysis and immediate response in time-critical scenarios. In addition, the proposed solution has undergone a thorough hardware evaluation, ensuring its compatibility and readiness for implementation in real-world operating conditions. Factors like computational efficiency, latency, model size, and system stability have been carefully assessed to verify seamless operation (Fig. 7). This not only enhances the practicality of the solution but also streamlines its implementation and adoption in various industrial settings. Moreover, the incorporation of state-of-the-art interpretability methods (Fig. 6) allows for reliable verification and confirmation of the effectiveness of the employed neural networks. By providing insights into the decision-making process of the models, these methods enhance trust and understanding, making the proposed approach more robust and transparent.

## 5. Conclusions

The presented fault diagnosis method allows contactless, non-invasive induction motor monitoring, which is the main advantage over conventional methods based on current or vibration analysis. It effectively processes thermal images extracting essential information for classification. The collected dataset includes samples from 42 independent experiments, which differ in configuration, load, and components. Obtained results show potential in the squirrel-cage induction motor's state classifying task. All selected models perform well, with EfficientNet B1 architecture achieving the highest metrics. Also, this model was verified using interpretability methods confirming that its attention focuses on the relevant parts of the image - the test stand and examined motor, rather than surrounding elements in the

background. Nevertheless, it should be taken into account that thermal imaging characterizes higher measurement uncertainties compared to current or vibration readings due to its high sensitivity to environmental factors such as temperature or humidity. It would be necessary to calibrate thermal camera parameters every time the working environment changes noticeably to minimize the above limitations and maximize performance.

What sets this publication apart is an extensive performance benchmark of the proposed solution on resource-constrained hardware. The conducted tests have proven that the selected neural architectures, with appropriate optimizations, such as weights quantization and hardware acceleration, can be efficiently processed on low-cost devices like Raspberry Pi 4B or Intel NCS2 if more computing power is needed. Moreover, the examinations have shown that

ONNX Runtime, an inference library, is the most suitable choice for on-device neural network processing due to its portability and the numerous selections of hardware acceleration providers.

The research will further focus on extending and diversifying the dataset. The new thermal frames will be gathered from various workstations, under different environments and working conditions, including other induction motor types and models. Additionally, future work will include the development of a complete prototype with a thermal camera as a perception sensor, an edge device as a computing module, and with results visualization unit. These components will allow for fully practical use of the system in hard-to-reach areas without affecting the operation of the drive system.

## References

1. Achouch M, Dimitrova M, Ziane K, Karganroudi SS, Dhouib R, Ibrahim H, Adda M. On Predictive Maintenance in Industry 4.0: Overview, Models, and Challenges. *Applied Sciences*, 12(16), 2022. <https://doi.org/10.3390/app12168081>
2. Alvarado-Hernandez AI, Zamudio-Ramirez I, Jaen-Cuellar AY, Osornio-Rios RA, Donderis-Quiles V, Antonino-Daviu JA. Infrared Thermography Smart Sensor for the Condition Monitoring of Gearbox and Bearings Faults in Induction Motors Sensors, 22(16), 2022. <https://doi.org/10.3390/s22166075>
3. Arm. PyArm NN 32.0.0, 2023. Online 12.04.2023 [<https://github.com/ARM-software/armnn>]
4. Balakrishnan GK, Yaw CT, Koh SP, Abedin T, Ra AA, Tiong SK, Chen CP. A Review of Infrared Thermography for Condition-Based Monitoring in Electrical Energy: Applications and Recommendations. *Energies*, 15(16), 2022. <https://doi.org/10.3390/en15166000>
5. Buhrmester V, Münch D, Arens M. Analysis of Explainers of Black Box Deep Neural Networks for Computer Vision: A Survey. *Machine Learning and Knowledge Extraction*, 3(4):966–989, 2021. <https://doi.org/10.3390/make3040048>
6. Cui C, Gao T, Wei S, Du Y, Guo R, Dong S, Lu B, Zhou Y, Lv X, Liu Q, Hu X, Yu D, Ma Y. PP-LCNet: A Lightweight CPU Convolutional Neural Network. *arXiv*, 2021. <https://doi.org/10.48550/arXiv.2109.15099>
7. Gangsar P, Tiwari R. 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, 2020. <https://doi.org/10.1016/j.ymssp.2020.106908>
8. Gong J, Liu W, Pei M, Wu C, Guo L. ResNet10: A lightweight residual network for remote sensing image classification. *14th International Conference on Measuring Technology and Mechatronics Automation (ICMTMA)*, 975–978, 2022. <https://www.doi.org/10.1109/ICMTMA54903.2022.00197>
9. Google. tflite-runtime 2.11.0, 2022. Online 12.04.2023 [<https://www.tensorflow.org/lite>]
10. Guo D, Li X, Luo F. On-line fault diagnosis of rotating machinery based on deep residual network. *37th Youth Academic Annual Conference of Chinese Association of Automation (YAC)*, 103–108, 2022. <https://www.doi.org/10.1109/YAC57282.2022.10023904>
11. Halder S, Bhat S, Zychma D, Sowa P. Broken Rotor Bar Fault Diagnosis Techniques Based on Motor Current Signature Analysis for Induction Motor - A Review. *Energies*, 15(22), 2022. <https://doi.org/10.3390/en15228569>
12. He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778, 2016. <https://www.doi.org/10.1109/CVPR.2016.90>
13. Huda ASN, Taib S. Application of infrared thermography for predictive/preventive maintenance of thermal defect in electrical equipment. *Applied Thermal Engineering*, 61(2):220–227, 2013. <https://doi.org/10.1016/j.applthermaleng.2013.07.028>

14. Ibrahim A, Anayi F, Packianather M. New Machine Learning Model-Based Fault Diagnosis of Induction Motors Using Thermal Images. 2nd International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE), 58–62, 2022. <https://www.doi.org/10.1109/ICACITE53722.2022.9823832>
15. Intel Corporation. Intel Neural Compute Stick 2, 2019. Online 12.04.2023 [<https://intel.com/content/www/us/en/products/sku/140109/intel-neural-compute-stick-2/specifications.html>]
16. Intel Corporation. Openvino 2022.1, 2022. Online 12.04.2023 [<https://docs.openvino.ai/2022.1/index.html>]
17. Janssens O, Schulz R, Slavkovikj V, Stockman K, Loccufer M, Van de Walle R, Van Hoecke S. Thermal image based fault diagnosis for rotating machinery. *Infrared Physics Technology*, 73:78–87, 2015. <https://doi.org/10.1016/j.infrared.2015.09.004>
18. Janssens O, Van de Walle R, Loccufer M, Van Hoecke S. Deep Learning for Infrared Thermal Image Based Machine Health Monitoring. *IEEE/ASME Transactions on Mechatronics*, 23(1):151–159, 2018. <https://www.doi.org/10.1109/TMECH.2017.2722479>
19. Karabacak YE, Gürsel Özmen N, Gümüşel L. Worm gear condition monitoring and fault detection from thermal images via deep learning method. *Eksplatacja I Niezawodność – Maintenance and Reliability*, 22(3):544–556, 2020. <https://doi.org/10.17531/ein.2020.3.18>
20. Kohlikyan N, Miglani V, Martin M, Wang E, Alsallakh B, Reynolds J, Melnikov A, Kliushkina N, Araya C, Yan S, Reblitz-Richardson O. Captum: A unified and generic model interpretability library for PyTorch, 2020. <https://www.doi.org/10.48550/ARXIV.2009.07896>
21. Li W, Huang R, Li J, Liao Y, Chen Z, He G, Yan R, Gryllias K. A perspective survey on deep transfer learning for fault diagnosis in industrial scenarios: Theories, applications and challenges *Mechanical Systems and Signal Processing*, 167:108487, 2022. <https://doi.org/10.1016/j.ymssp.2021.108487>
22. Li Y, Du X, Wan F, Wang X, and Yu H. Rotating machinery fault diagnosis based on convolutional neural network and infrared thermal imaging. *Chinese Journal of Aeronautics*, 33(2):427–438, 2020. <https://doi.org/10.1016/j.cja.2019.08.014>
23. Li Y, Gu JX, Zhen D, Xu M, Ball A. An evaluation of gearbox condition monitoring using infrared thermal images applied with convolutional neural networks. *Sensors*, 19(9), 2019. <https://doi.org/10.3390/s19092205>
24. Liang Y, Li S, Yan C, Li M, Jiang C. Explaining the black-box model: A survey of local interpretation methods for deep neural networks. *Neurocomputing*, 419:168–182, 2021. <https://doi.org/10.1016/j.neucom.2020.08.011>
25. Lopes JF, da Costa VGT, Barbin DF, Cruz-Tirado LJP, Baeten V, Barbon Junior S. Deep computer vision system for cocoa classification. *Multimedia Tools and Applications*, 81(28):41059–41077, Nov 2022. <https://www.doi.org/10.1007/s11042-022-13097-3>
26. Loshchilov I, Hutter F. Decoupled weight decay regularization. *arXiv*, 2019. <https://doi.org/10.48550/arXiv.1711.05101>
27. Luczak D, Brock S, Siembab K. Cloud based fault diagnosis by convolutional neural network as time-frequency RGB image recognition of industrial machine vibration with internet of things connectivity. *Sensors*, 23(7), 2023. <https://doi.org/10.3390/s23073755>
28. Mahendran A, Vedaldi A. Salient deconvolutional networks. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, 120–135, Cham, 2016. Springer International Publishing. [https://doi.org/10.1007/978-3-319-46466-4\\_8](https://doi.org/10.1007/978-3-319-46466-4_8)
29. Mian T, Choudhary A, Fatima S. Multi-sensor fault diagnosis for misalignment and unbalance detection using machine learning. *IEEE International Conference on Power Electronics, Smart Grid, and Renewable Energy (PESGRE)*, 1–6, 2022. <https://www.doi.org/10.1109/PESGRE52268.2022.9715938>
30. Microsoft. ONNX Runtime 1.14.1, 2023. [<https://onnxruntime.ai>]
31. Orłowska-Kowalska T, Wolkiewicz M, Pietrzak P, Skowron M, Ewert P, Tarchala G, Krzysztofia M, Kowalski CT. Fault diagnosis and fault-tolerant control of PMSM drives—state of the art and future challenges. *IEEE Access*, 10:59979–60024, 2022. <https://www.doi.org/10.1109/ACCESS.2022.3180153>
32. Ramirez OJV, Cruz de la Cruz JE, Machaca WAM. Agroindustrial plant for the classification of hass avocados in real-time with ResNet-18 architecture. 5th International Conference on Robotics and Automation Sciences (ICRAS), 206–210, 2021. <https://www.doi.org/10.1109/ICRAS52289.2021.9476659>
33. Ran Y, Zhou X, Lin P, Wen Y, Deng R. A survey of predictive maintenance: Systems, purposes and approaches. *arXiv*, 2019. <https://doi.org/10.48550/arXiv.1912.07383>
34. Reljic D, Jerkan D, Marcetic D, Oros D. Broken bar fault detection in IM operating under no-load condition. *Advances in Electrical and Computer Engineering*, 16:63–70, 01 2016. <https://www.doi.org/10.4316/AECE.2016.04010>

35. Reséndiz-Ochoa E, Morales-Hernández LA, Cruz-Albarran IA, Alvarez-Junco S. Induction motor failure analysis using machine learning and infrared thermography. *IEEE International Autumn Meeting on Power, Electronics and Computing (ROPEC)*, volume 6, 1–6, 2022. <https://www.doi.org/10.1109/ROPEC55836.2022.10018653>
36. Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision*, 128(2):336–359, Oct 2019. <https://doi.org/10.1007%2Fs11263-019-01228-7>
37. Siemens. The true cost of downtime 2022. Online 12.04.2023 [<https://assets.new.siemens.com/siemens/assets/api/uuid:3d606dbe0-43e4-80b1-d04e27ada920/dics-b10153-00-7600truecostofdowntime2022-144.pdf>]
38. Simonyan K, Vedaldi A, Zisserman A. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv*, 2013. <https://www.doi.org/10.48550/ARXIV.1312.6034>
39. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. *arXiv*, 2015. <https://doi.org/10.48550/arXiv.1409.1556>
40. Singh G, Anil Kumar TC, Naikan VNA. Induction motor inter turn fault detection using infrared thermographic analysis. *Infrared Physics Technology*, 77:277–282, 2016. <https://doi.org/10.1016/j.infrared.2016.06.010>
41. Singh G, Naikan VNA. Infrared thermography based diagnosis of inter-turn fault and cooling system failure in three phase induction motor. *Infrared Physics Technology*, 87:134–138, 2017. <https://doi.org/10.1016/j.infrared.2017.10.007>
42. Springenberg JT, Dosovitskiy A, Brox T, Riedmiller M. Striving for simplicity: The all convolutional net. *arXiv*, 2015. <https://doi.org/10.48550/arXiv.1412.6806>
43. Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A. Going deeper with convolutions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1–9, 2015. <https://www.doi.org/10.1109/CVPR.2015.7298594>
44. Tan M, Chen B, Pang R, Vasudevan V, Sandler M, Howard A, Le QV. MnasNet: Platform-aware neural architecture search for mobile. *arXiv*, 2018. <https://doi.org/10.1109/CVPR.2019.00293>
45. Tan M, Le QV. MixConv: Mixed depthwise convolutional kernels. *arXiv*, 2019. <https://doi.org/10.48550/arXiv.1907.09595>
46. Tan M, Le QV. EfficientNet: Rethinking model scaling for convolutional neural networks. *arXiv*, 2020. <https://doi.org/10.48550/arXiv.1905.11946>
47. The Linux Foundation. Pytorch 1.13.1, 2022. Online 12.04.2023 [<https://pytorch.org>]
48. The Raspberry Pi Foundation. Raspberry Pi 4 tech specs, 2019. Online 12.04.2023 [<https://www.raspberrypi.com/products/raspberry-pi-4-model-b/specifications>]
49. Tran MQ, Amer M, Dababat A, Abdelaziz AY, Dai HJ, Liu MK, Elsis M. Robust fault recognition and correction scheme for induction motors using an effective iot with deep learning approach. *Measurement*, 207:112398, 2023. <https://doi.org/10.1016/j.measurement.2022.112398>
50. VertivCo. Understanding the cost of data center downtime: An analysis of the financial impact on infrastructure vulnerability. [[https://www.vertiv.com/4a3537/globalassets/images/about-images/news-and-insights/articles/white-papers/understanding-the-cost-of-data-center/datacenter-downtime-wp-en-na-sl-24661\\_51225\\_1.pdf](https://www.vertiv.com/4a3537/globalassets/images/about-images/news-and-insights/articles/white-papers/understanding-the-cost-of-data-center/datacenter-downtime-wp-en-na-sl-24661_51225_1.pdf)]
51. Verucchi CJ, Giraldo E, Meira M, Ruschetti CR, Bossio JM, Bossio GR. Efficiency assessment of induction motors drives operating under shaft misalignment conditions. *Advances in Electrical and Electronic Engineering*, 18(3):142–152, 2020. <https://www.doi.org/10.15598/aeee.v18i3.3596>
52. Wang R, Zhan X, Bai H, Dong E, Cheng Z, Jia X. A review of fault diagnosis methods for rotating machinery using infrared thermography. *Micromachines*, 13(10), 2022. <https://www.doi.org/10.3390/mi13101644>
53. Wightman R. Pytorch image models. *GitHub*, 2019. <https://www.doi.org/10.5281/zenodo.4414861>
54. Workswell. Workswell infrared camera - stationary radiometric thermal camera, 2020. Online 12.04.2023 [<https://workswell-thermal-camera.com/workswell-infrared-camera-wic>]
55. Zeiler MD and Fergus R. Visualizing and understanding convolutional networks. *arXiv*, 2013. <https://doi.org/10.48550/arXiv.1311.2901>
56. Zhang H and Wang F. Fault identification of fan blade based on improved ResNet-18. *Journal of Physics: Conference Series*, 2221(1):012046, May 2022. <https://www.doi.org/10.1088/1742-6596/2221/1/012046>