Rapid Diagnosis of Induction Motor Electrical Faults using Convolutional Autoencoder Feature Extraction

Arild Bergesen Husebø[1], Surya Teja Kandukuri[2], Andreas Klausen[3], Van Khang Huynh[4], Kjell Gunnar Robbersmyr[5]

[1,2,3,4,5] *University of Agder, Department of Engineering Sciences, Jon Lilletuns vei 9, 4879 Grimstad, Norway*
arild.b.husebo@uia.no[1] surya.kandukuri@uia.no[2] andreas.klausen@uia.no[3] huynh.khang@uia.no[4] kjell.g.robbersmyr@uia.no[5]

ABSTRACT

Electrical faults such as stator turns fault and broken rotor bars are among the frequently occurring failure modes in induction motors. This article presents a novel deep learning-based approach for the rapid diagnosis of these electrical faults within a short time window of 200 milliseconds. The extended Park's vector, calculated using three-phase supply currents, is chosen as the medium for fault detection. An unsupervised convolutional autoencoder is designed to detect features distinguishing healthy and faulty conditions. The developed features are supplied to a support vector machine to classify the fault conditions. The proposed approach is validated in a laboratory setup consisting of an inverter-fed induction motor operating under time-varying load and speed conditions with an accuracy $> 95\%$.

1. INTRODUCTION

Induction motors (IMs) are ubiquitous in industrial applications, accounting to 67% of total industrial power consumption (Boldea & Nasar, 2002). Although these IMs are robust and reliable, they are prone to failures under prolonged operation due to wear, improper ventilation, excessive loading and environmental conditions (Bonnett, 2000). Especially in the case of inverter-fed IMs, the electrical faults such as the stator turns fault (STF) and broken rotor bars fault (BRB) contribute to $> 40\%$ of the frequent failures (Yeh et al., 2008). Thermal stresses in the IMs caused due to overloading or poor ventilation are common reasons for these failures. An electrical fault such as STF can rapidly progress into a hazardous catastrophic failure besides causing loss of productivity. Therefore these electrical faults are the focus of this article.

Fault diagnosis of IMs has been an active area of research for the past few decades, using vibrations (Tsypkin, 2017), currents (Nandi, Toliyat, & Xiaodong, 2005), acoustic signals (Glowacz, 2019) and stray flux measurements (Frosini, Harlisca, & Szabo, 2015). Of these approaches, the motor current signature analysis (MCSA) has been proven successful both in academia and in industrial applications (Thomson & Fenger, 2003). While single-line current MCSA was capable of providing reliable fault detection on supply-fed IMs, they proved to be inadequate in the inverter-fed IMs operating under closed-loop control (Bellini, Filippetti, Franceschini, & Tassoni, 2000). The extended Park's vector modulus (EPVM) was proven to be a reliable medium for fault detection in these machines (Cruz & Cardoso, 2001). In the case of machines operating under steady-state, the fast Fourier transform (FFT) for the spectral analysis of single-line currents or EPVM is reliable to identify the characteristic fault frequencies for detection of the IM faults. However, fault diagnosis becomes challenging in time-varying speed and load conditions. While this problem is typically addressed using time-frequency transforms such as short-time Fourier transform (STFT) (Burriel-Valencia, Puche-Panadero, Martinez-Roman, Sapena-Bano, & Pineda-Sanchez, 2017), wavelet transform (WT) (Ameid, Menacer, Talhaoui, & Azzoug, 2018) and Wigner-Ville distributions (Climente-Alarcon, Antonino-Daviu, Riera-Guasp, & Vlcek, 2014), there remain two important challenges; 1) rarely are these methods tested under multiple fault conditions. This is important as characteristic fault frequencies that are close to each other are difficult to distinguish as the time-frequency approaches can only provide either frequency or time resolution and 2) they involve significant pre-processing of the measured signals. To overcome these challenges and to perform a rapid diagnosis within a short time window such that the diagnosis is insensitive to non-stationarity, a deep learning-based diagnostics approach is proposed in this article. A very short time window of 200 ms is sufficient to perform the diagnosis, and once trained, the diagnostics method is directly scalable to a fleet of IMs. Such a short time window circumvents the problems associated with signal processing in transient conditions. Furthermore, based on training, the convolutional autoencoder (CAE) can detect features beyond the classical characteristic frequencies and thus, can be more robust in detecting multiple known fault conditions in transient operations.

The remainder of the paper is organized as follows; The pro-

posed diagnostics approach is detailed in Section 2. A description of the laboratory tests used for validation of the proposed approach is detailed in Section 3 along with a discussion on the results. Finally, the paper is concluded in Section 4.

2. DIAGNOSTICS APPROACH

An overview of the proposed diagnostics approach is shown in Figure 1. Three channel (one channel per phase) current signals are drawn from datasets containing data from various motor operating conditions. Each three channel signal is transformed using the extended park vector modulus, where the output is a single channel signal - the signal that goes into the CAE model. The single channel signal is then normalized. We will define this as the model signal. From here, the CAE estimates a set of features from the model signal. During training of the CAE, the model signal is used as both model input and output, as will be explained later. To verify the relevance of the features to any operating condition, the features are fed into a supervised classifier. During training, the condition labels are drawn from the datasets and used as classifier targets. The classifier predicts the condition based on the features from the CAE.

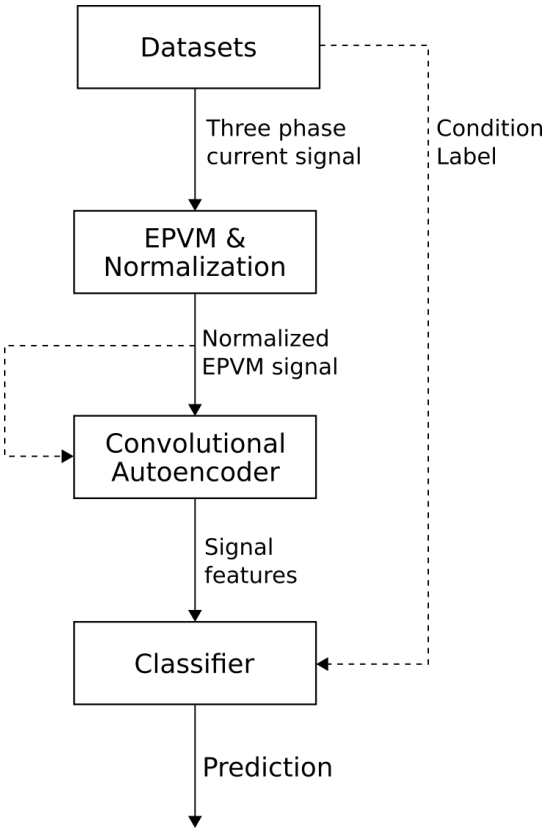

Figure 1. Method overview

2.1. EPVM & Normalization

A healthy IM is a symmetric and balanced machine wherein all the three phase currents are identical, differing only by the phase angle. However, electrical faults in IM produce a variation in the electrical circuit, while mechanical faults such as bearing fault produce a variation in the airgap. All of these faults result in periodic disturbances in the current that varies with the supply frequency f_s, rotor speed f_r and load on the system. The MCSA focuses on spectral evaluation of single line currents to detect these periodic disturbances. The EPVM on the other hand, utilizes all the three phase currents to detect these faults.

As part of the initial processing step, the EPVM is calculated using the direct and quadrature axis currents $(i_d,\ i_q)$ given by

$$i_d = \sqrt{\frac{2}{3}} i_a - \sqrt{\frac{1}{6}} i_b - \sqrt{\frac{1}{6}} i_c \tag{1}$$

where $\vec{i_{abc}} = [i_a, i_b, i_c]$ are the three phase stator currents of the motor and

$$i_q = \sqrt{\frac{1}{2}} i_b - \sqrt{\frac{1}{2}} i_c \tag{2}$$

and the EPVM $|i_P|$ is given by

$$|i_P| = \sqrt{(i_d)^2 + (i_q)^2} \tag{3}$$

the EPVM is a DC offset in the case of healthy machine. However, it holds the characteristic fault frequencies in the case of fault conditions. The main advantage of using EPVM is that it automatically removes the supply frequency which is the dominant frequency in the current spectrum, thus making distinctive features easily differentiable (Cardoso, Cruz, & Fonseca, 1999).

The EPVM transformation is followed by a normalization step where the signal is divided by its infinity norm, defined as

$$||\boldsymbol{x}||_\infty = \max\left(|x_1|, |x_2|, ..., |x_n|\right) \tag{4}$$

2.2. Convolutional Autoencoder

An autoencoder is a neural network designed to replicate an input using a set of latent representations. A common autoencoder structure resembles an hourglass, with the latent representations in the middle and a network structure mirrored about the middle. The latent representations are usually a vector of scalars. During training, the loss function is a metric of difference between the input and the reconstruction. In practice, the training target is the same as the input.

The research presented in this paper aims to adapt the convolutional properties of a convolutional neural network into an autoencoder, hence convolutional autoencoder (Mao, Shen, & Yang, 2016). Therefore, the outermost layers of the model are convolutional layers, as seen in Figure 1.

The input to the encoder part is passed through four consecutive convolutional blocks. The input to each convolutional block is a convolutional layer. Convolutional layers convolve a set of trainable kernels across the input to produce a filtered output. The convolution operation has a fixed stride, so the output dimensionality is lower than the input dimensionality. Traditionally, pooling operations are used for this dimensionality reduction, as pooling also acts as a receptive field. However, depending on the pooling technique, one must keep track of the pooling indices when upsampling the signal again in the decoder part. This makes it impossible to use the decoder on its own when using pooling techniques. Fortunately, using convolution strides instead of pooling operations has shown to be sufficient (Springenberg, Dosovitskiy, Brox, & Riedmiller, 2014). The output from each convolutional layer is sent through a rectified linear unit (ReLU) activation function, to force a positive linear output. Finally, the output is centered and re-scaled through a batch normalization operation (Ioffe & Szegedy, 2015).

The outputs from the fourth convolutional layer are concatenated into a single vector through an operation normally referred to as flattening. This vector is connected to two consecutive densely connected layers. These layers are the foundation of traditional neural networks as they consists of densely connected neurons, meaning that each neuron in the layer has connections to all the activations of the previous layer. The output from each dense layer is activated by a sigmoid function. This activation introduces additional non-linearities into the network.

The latent encodings are represented by a dense layer, where the amount of neurons corresponds to the amount of features. The activation of this layer is linear, so the output is just the value of each neuron output.

The decoder part has similar structure as the encoder part, except it is mirrored. The convolutional layers are replaced by transposed convolutional layers, also known as deconvolution layers, with the same stride size as their encoder counterparts. Finally, the output of the decoder is the reconstructed signal, either reconstructed from an input signal using the entire autoencoder model, or reconstructed from a set of features using only the decoder part. Table 1 shows the complete structure of the autoencoder.

The outermost convolutional and deconvolutional layers have a stride of 4, while the remaining ones have a stride of 5 in order to reduce the dimensions. Dimensionality after the four convolutional layers is therefore $1000/(4\times5^3) = 2$, and vice-versa after the deconvolutional layers. The reasoning behind this drastic reduction in dimensionality is to capture the entire signal in the models receptive field.

Table 1. Model composition

Block	Layer	Output Dimension
Input	-	[1000]
Conv A1	Conv	[100, 250]
	ReLU	[100, 250]
	BatchNorm	[100, 250]
Conv A2	Conv	[100, 50]
	ReLU	[100, 50]
	BatchNorm	[100, 50]
Conv A3	Conv	[100, 10]
	ReLU	[100, 10]
	BatchNorm	[100, 10]
Conv A4	Conv	[100, 2]
	ReLU	[100, 2]
	BatchNorm	[100, 2]
Flatten	Flatten	[200]
Dense A1	Dense	[100]
	Sigmoid	[100]
Dense A2	Dense	[50]
	Sigmoid	[50]
Encodings	Dense	[5]
	Linear	[5]
Dense B2	Dense	[50]
	Sigmoid	[50]
Dense B1	Dense	[100]
	Sigmoid	[100]
Reshape	Dense	[200]
	Reshape	[100, 2]
ConvTransp B4	ConvTransp	[100, 10]
	ReLU	[100, 10]
	BatchNorm	[100, 10]
ConvTransp B3	ConvTransp	[100, 50]
	ReLU	[100, 50]
	BatchNorm	[100, 50]
ConvTransp B2	ConvTransp	[100, 250]
	ReLU	[100, 250]
	BatchNorm	[100, 250]
ConvTransp B1	ConvTransp	[1, 1000]
Output	-	[1000]

The CAE has two modes: training and evaluation. During training, the pre-processed current signals from the previous section is fed to the network as both input and target (stippled arrow in Figure 1). The training objective is to minimize the mean squared reconstruction error

$$\text{MSE} = \frac{1}{N}\sum_{n=0}^{N-1}\left(f(x_n) - x_n\right)^2 \tag{5}$$

where $f(x_n)$ is the neural network transformation of input x_n which is a vector containing the preprocessed current signal segment n. The optimization algorithm used is called Adam

(**Ada**ptive **M**oment Estimation) (Kingma & Ba, 2014).

2.3. Classifier

The CAE is an unsupervised learning algorithm that identifies distinctive features for various health conditions. However, a classifier is necessary to diagnose and label the fault conditions. Thus, a classifier is used to determine a machine health condition based on the features extracted from the encoder. Although there exists various classification algorithms, the support vector machine (SVM) is used in this paper. The SVM uses a radial basis kernel function (RBF) to transform data into higher dimensions, enabling hyperplanes that best separate data of different labels to be fit with greater ease (Vapnik, 2000). This is known as the kernel trick. During training, the features from the CAE are used as inputs and condition labels from the datasets are used as targets. During the training phase, the SVM is supplied with features and corresponding labels (healthy, STF and BRB). Further in testing phase, the SVM is tasked to classify fault conditions based on new sets of features.

3. LABORATORY VALIDATION

The laboratory test bench used to evaluate the proposed approach is shown in Figure 2. It consists of two three-phase IMs coupled with a two-stage planetary gearbox and a hybrid bevel-planetary helical gearbox. The test motor ('pitch drive' in Figure 2) is a 1.1 kW, three-phase, 4 pole IM that is connected to a commercial FOC controller under speed-control paradigm. The load is supplied by a 3 kW, three-phase, 8 pole IM also connected to a commercial FOC controller under torque-control paradigm. The test motor is equipped with LEM LTS-6NP Hall effect current sensors on all the three phases. An NI USB-6215 DAQ is used for data acquisition in conjunction with MATLAB Data Acquisition Toolbox.
Using this test setup, the test motor was first operated in healthy condition under various (steady-state) load and speed conditions and three-phase currents were collected. Further, artificially seeded faults, STF and BRB are introduced sequentially, as shown in Figure 2 (a) and (b) and the test sequence was repeated.

3.1. Datasets

The data collected in three different health conditions: healthy, STF and BRB is further processed for feature extraction. Each dataset contains data from a varying range of motor supply frequencies as shown in Table 2. As an initial processing step, all the data is processed into one large data file in order to make data allocation faster during training. The original datasets were sampled with a sampling frequency of 25 kHz, which is a lot more than what is needed to detect fault signatures considering the supply frequencies. Therefore, all datasets are decimated with a factor of five using a FIR filter,

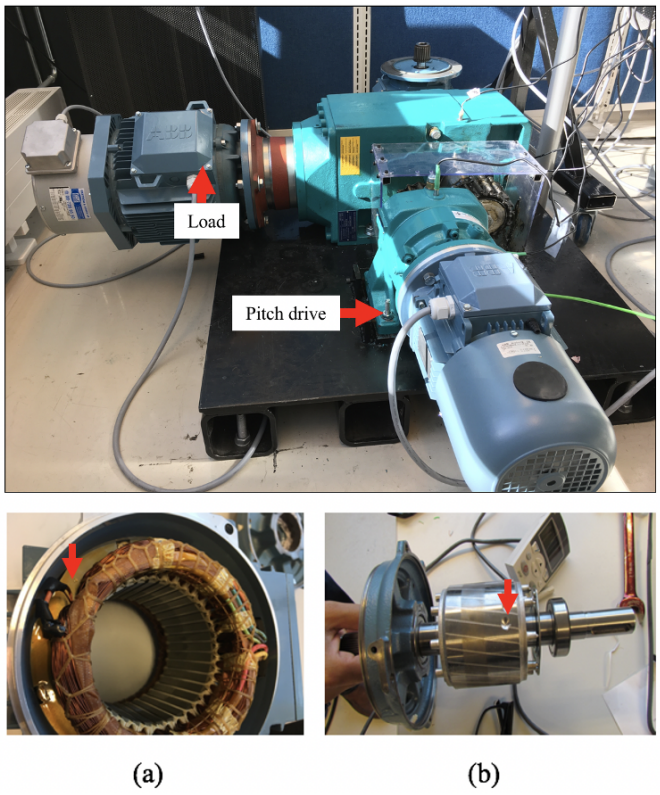

Figure 2. The laboratory test bench with three-phase IM along with the artificially seeded faults, (a) STF and (b) BRB.

effectively resulting in a sample frequency of 5 kHz, or a sample period of 200 μs. After this step, each signal is segmented into smaller signals to reduce the input size of the model for easier training. These segments are of 1000 data points in size (per phase), effectively making one model-ready signal 200 ms long. The choice of signal length is based on two criteria; one, the length of signal should be long enough to ensure that the fault phenomenon is distinguishable while two, is not too long to make the convolutional autoencoder computationally expensive for the available hardware.

3.2. Results

The main objective of the proposed method is to extract features from time-domain signals in very short time windows, where the features are used to classify health conditions based on learning from historic data. The secondary objective is to verify the generalization of the method with regards to supply frequency, that is, variations in speed and load. As the supply frequency of the machine is either increasing or decreasing, both time and frequency domain signals will see either a stretching or a shrinking along the time or frequency axis. In addition to this, depending on exactly when a measurement is started, the supply frequency phase shift will vary from one signal to another. The EPVM method removes the fun-

Table 2. Number of signals per health condition w.r.t. supply frequency

Freq	Healthy	STF	BRB
38.3			6600
38.5		9000	
38.8	3150		
39.2	13500		
43.5		4500	
45.5	13050		
46.8		9000	
47.5		9000	
48.0	13050		
48.28			6300
48.82			7200
49.1			3300
Total	42750	31500	23400

damental frequency, which is the supply frequency, but other signatures will remain at the same phase shift in the signal.

Therefore, the method will first be validated on data captured during supply frequencies similar to those it has seen during training. The validation data is not used during training, but stem from measurements using the same supply frequency as the training data. Secondly, the model is trained on data from supply frequencies different to those used to validate it. Three different cases are presented, as seen in Section 3.2.2 to show differences in performance depending on motor speed during training.

3.2.1. Similar supply frequency

The CAE training process using all available supply frequencies can be seen in Figure 3. Only 10% of the signals are used for training while the remaining data is reserved for validation. The horizontal axis shows the epoch number. An epoch represents one iteration of training using all the training data. The model was trained for a total of 100 epochs. At this point, the training loss reduction per epoch is minimal. The vertical axis shows the training loss function, which is the reconstruction MSE as described in Eqn. 5. At the end of every epoch, the state of the model at that time is validated using a subset of the data reserved for validation. As shown in Figure 3, the validation graph fluctuates around the training curve, although it mostly stays approximately the same as the training loss. This is a sign of good model generalization.

Signal reconstruction examples are shown in Figure 4 and 5. To generate the reconstruction, an EPVM-transformed signal is first encoded by the encoder part of the trained CAE. The encoder returns 5 features, which are then decoded by the decoder part of the trained CAE. The output is the reconstructed signal. In Figure 4 it can be seen that the reconstruction is far from a perfect match of the input signal. This result is anticipated considering signal noise and the fact that 1000 data points are compressed into 5 scalar features. The condition

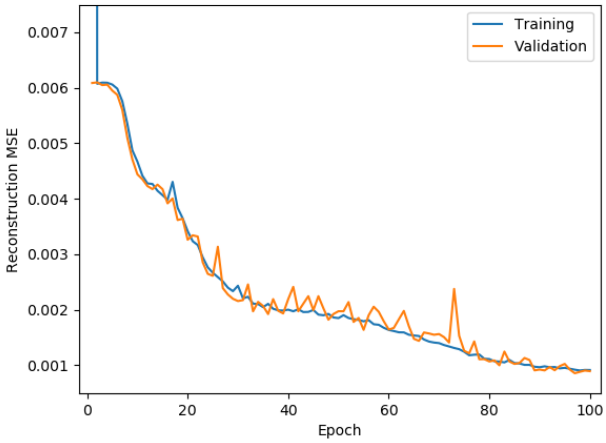

Figure 3. Training history when trained on similar speeds

type with overall best reconstruction is STF. STF condition signals distinguish themselves from healthy condition very well, compared to BRB which are very similar. The main signature of the STF condition has a high amplitude compared to the rest of the signal. This is not the case for BRB where the fault signatures blend more into the signal itself. This makes the MSE loss a good candidate for extracting STF related information, but less so for extracting healthy and BRB related information. This can be seen especially in Case 1 and 2 of section 3.2.2, where several healthy condition signals are predicted to be of BRB condition and vice-versa. A possible solution to this problem was tested using a different loss function based on the correlation between input and output signal, given as

$$\text{Correlation loss} = 1 - |\text{corr}(f(x), x)| \tag{6}$$

where $\text{corr}(f(x), x)$ is the Pearson correlation coefficient between $f(x)$ and x. However, using this loss function resulted in overall worse performance, and was hence dropped in favor of the MSE loss function. Training the model on 90% of the data instead yields a better reconstruction of the different condition signals (Figure 5). When using this amount of training data, SVM predictions mentioned later can reach 100% accuracy in all conditions.

To understand how the CAE accounts for variation in phase and speed, we need to look at the variations in the encoded features. Figure 4 shows that signals of STF condition has the most distinctive periodic waveform. We collect features generated from test data of all frequencies and calculate the variance of each feature over the data. Assuming that phase shift differences between signals are the biggest contributor to the feature variance, we plot the three features of highest variance in Figure 6. Four circular shapes appear, corresponding to all the supply frequencies of the test data. The plot shows

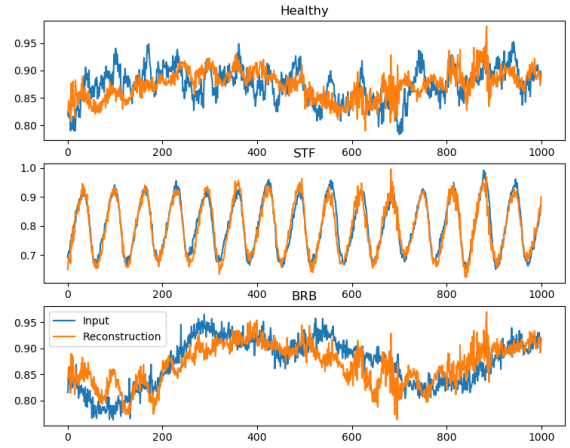

Figure 4. Signal reconstruction example

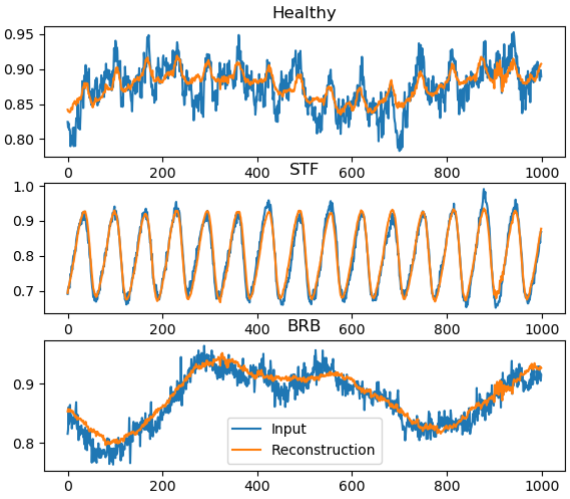

Figure 5. Signal reconstruction when trained on more data

how different features work together to compensate for phase differences. The orientation of the circular shapes also seem to be related to the supply frequency. The same effects are observed for healthy and BRB conditions, although the circular shapes are far less prominent. A drawback of the proposed method is that the features may be highly correlated and not individually informative.

The SVM was trained on features collected from the training data subjected to the trained CAE. It was then used to make predictions based on features collected from the testing data subjected to the CAE. The prediction results are shown in the confusion matrix of Figure 7. All the signals of STF condition was correctly classified, while more than 99% of the healthy

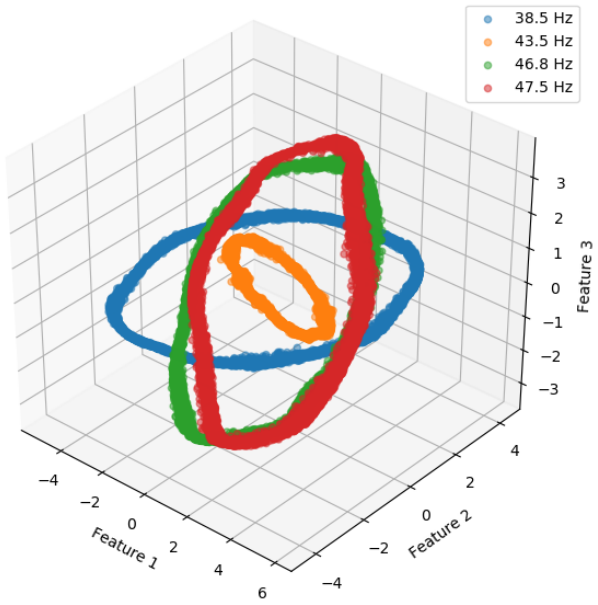

Figure 6. The three STF signal features of highest variance

condition signals where correctly classified. However, only 90% of the BRB condition signals were correctly predicted.

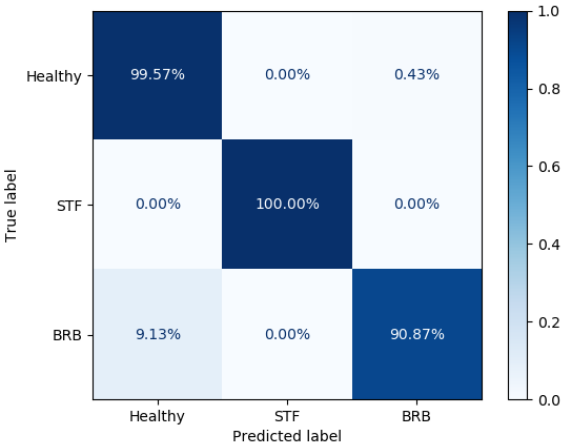

Figure 7. Classifier confusion matrix when trained on similar speeds

3.2.2. Unseen supply frequency

Testing the CAE on supply frequencies unseen during training is a benchmark of how much the model can generalize with respect to speed variations. Ideally, the receptive field of the CAE design should pick up stretching in the signals due to speed variations. However, also the signatures themselves will change.

In order to verify the viability of the model for use with unseen supply frequencies, one could argue that the model should be trained on data from supply frequencies in the middle of practical operating range. On the other hand, it might seem like a logical approach to train the model on data from the boundaries of the practical range. The paper therefore presents three cases of training data distributions to see which performs better.

Case 1 presents the former training data distribution. The model is trained on only data from the most central supply frequency available for a type of condition, and validated on data from the remaining supply frequencies. Table 3 shows how the training and test data is distributed based on supply frequency. From the training history in Figure 8, we can see

Table 3. Supply frequency distributions, case 1

Freq	Healthy	STF	BRB
38.3			Test
38.5		Test	
38.8	Test		
39.2	Test		
43.5		Train	
45.5	Train		
46.8		Test	
47.5		Test	
48.0	Test		
48.28			Train
48.82			Test
49.1			Test

that the validation error starts to increase after the fifth epoch, while the training error continues to decrease. An increasing validation error is a sign of overfitting (Rosin & Fierens, 1995), which may be a result of lack of variation in training data. The actual predictions of the SVM is shown in Figure 9. We can see a great amount of prediction errors.

Case 2 presents a training data distribution where the outermost supply frequencies of each condition is used for training as seen in Table 4. The training history in Figure 10 shows

Table 4. Supply frequency distributions, case 2

Freq	Healthy	STF	BRB
38.3			Train
38.5		Train	
38.8	Train		
39.2	Test		
43.5		Test	
45.5	Test		
46.8		Test	
47.5		Train	
48.0	Train		
48.28			Test
48.82			Test
49.1			Train

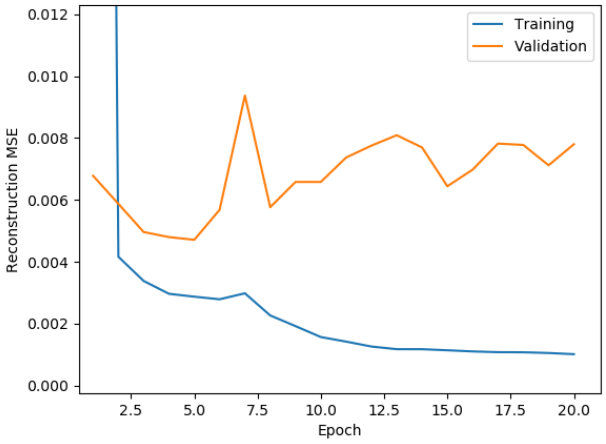

Figure 8. Training history when trained in unseen supply frequency case 1

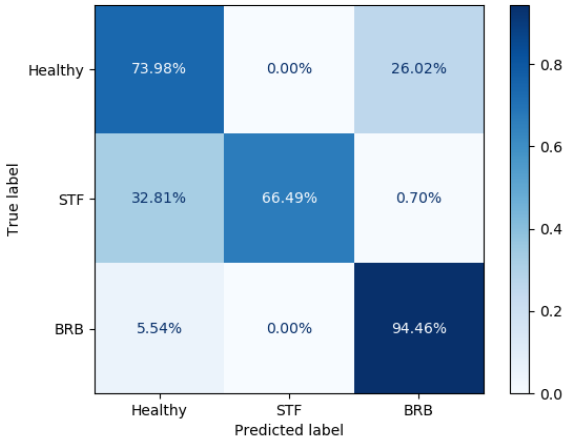

Figure 9. Classifier confusion matrix when trained in unseen supply frequency case 1

no sign of overfitting as in case 1. However, the validation error does lag behind the training error. This was not the case when trained and validated on similar supply frequencies. Both training and validation errors are decreasing at a similar rate until they converge. The SVM predictions shown in Figure 11 shows how both STF and BRB conditions have fairly high prediction accuracy, while healthy condition tends to be predicted as BRB condition.

Case 3 presents a case where the training data is selected with respect to all the supply frequencies except one, which is reserved for validation. The data is distributed as shown in Table 5. The training history of this case (Figure 12) show similar trends to that of case 2. The validation reconstruction error still lags behind the training reconstruction error,

but there are no signs of overfitting. The fluctuations seem slightly smaller than case 2. The fact that the validation error stays consistent can also be a sign of good generalization. The SVM predictions in case 3 (Figure 13) results in much higher accuracy than the earlier cases. Note that from Table 5, STF condition had the greatest supply frequency difference between training and testing data, although it performed very well during classification.

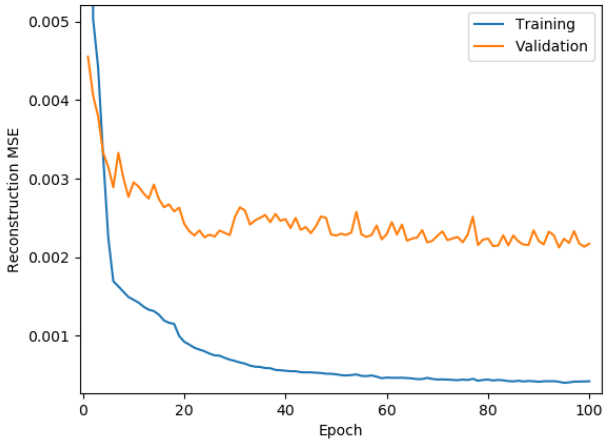

Figure 10. Training history when trained in unseen supply frequency case 2

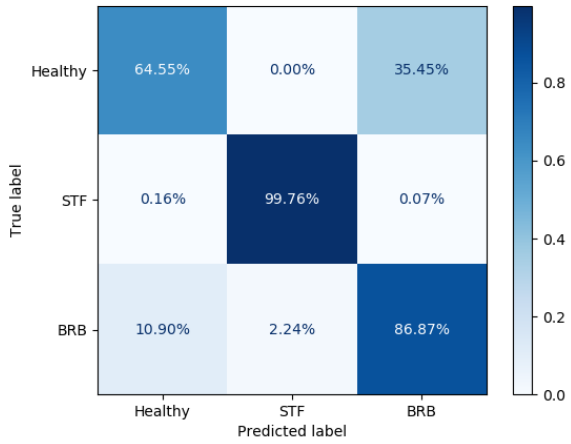

Figure 11. Classifier confusion matrix when trained in unseen supply frequency case 2

Table 5. Supply frequency distributions, case 3

Freq	Healthy	STF	BRB
38.3			Train
38.5		Train	
38.8	Train		
39.2	Test		
43.5		Test	
45.5	Train		
46.8		Train	
47.5		Train	
48.0	Train		
48.28			Test
48.82			Train
49.1			Train

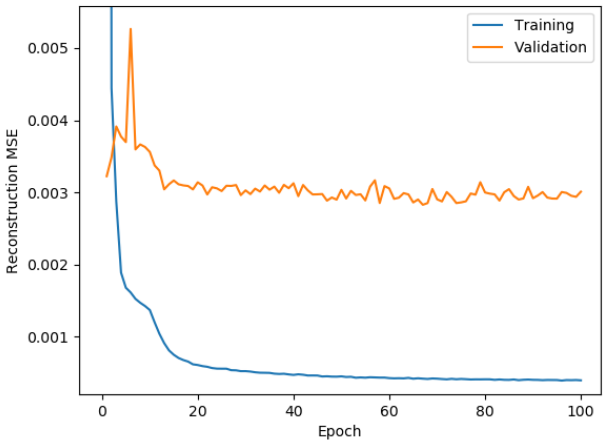

Figure 12. Training history when trained in unseen supply frequency case 3

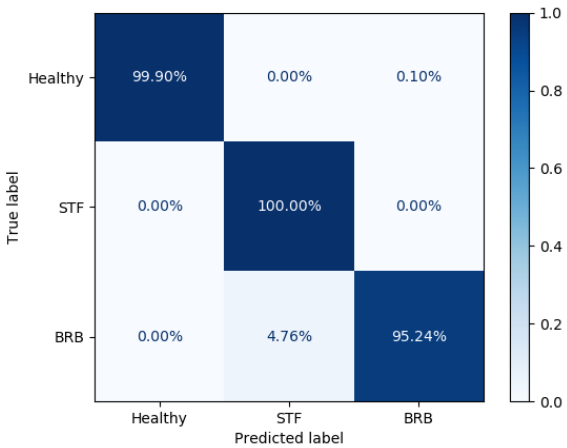

Figure 13. Classifier confusion matrix when trained in unseen supply frequency case 3

4. CONCLUSION

This research presents an approach for rapid fault diagnosis in IMs based on features extracted from EPVM. The method is based on unsupervised training using a CAE. The model is designed to work with a current signal of 200 ms length, and outputs 5 feature representations of that signal. The short time window alleviates problems related to transient conditions. The 5 features are used with an SVM classifier to diagnose the health of an IM under varying operating conditions. It is shown that prediction accuracy can be higher than 95% when working with supply frequencies similar to those the CAE has seen during training. Additionally, the classifier also works well on validation data recorded during supply frequency input in between training data supply frequencies. Typically, the motors in industrial applications operate under predefined speed and load profiles. In such cases, the proposed approach may be trained on a prototype motor and scaled to be utilised on a fleet of machines. The rapid diagnosis gives particular advantage in quick evaluation for maintenance tasks and planning.

There are several aspects to be considered prior to implementing this methodology in an industrial context. Firstly, the method is only tested at single fault severity, created by artificially seeded faults. It is to be verified how the approach fares under varying degrees of fault severity. Second, although the test is performed using 3 different motors, they are all of the same size and make. Therefore, it has to be verified if the approach can be extended to different sizes and models. Finally, multiple faults scenario was not addressed here. The authors aim to address some of these shortcomings in further research.

REFERENCES

Ameid, T., Menacer, A., Talhaoui, H., & Azzoug, Y. (2018, aug). Discrete wavelet transform and energy eigen value for rotor bars fault detection in variable speed field-oriented control of induction motor drive. *ISA Transactions*, *79*, 217–231. doi: 10.1016/j.isatra.2018.04.019

Bellini, A., Filippetti, F., Franceschini, G., & Tassoni, C. (2000). Closed-loop control impact on the diagnosis of induction motors faults. *IEEE Transactions on Industry Applications*, *36*(5), 1318–1329. doi: 10.1109/28.871280

Boldea, I., & Nasar, S. A. (2002). *The Induction Machine Handbook*. CRC Press.

Bonnett, A. (2000). Root cause AC motor failure analysis with a focus on shaft failures. *IEEE Transactions on Industry Applications*, *36*(5), 1435–1448. doi: 10.1109/28.871294

Burriel-Valencia, J., Puche-Panadero, R., Martinez-Roman, J., Sapena-Bano, A., & Pineda-Sanchez, M. (2017, mar). Short-frequency fourier transform for fault diagnosis of induction machines working in transient regime. *IEEE Transactions on Instrumentation and Measurement*, *66*(3), 432–440. doi: 10.1109/tim.2016.2647458

Cardoso, A. M., Cruz, S., & Fonseca, D. (1999). Inter-turn stator winding fault diagnosis in three-phase induction motors, by park's vector approach. *IEEE Transactions on Energy Conversion*, *14*(3), 595–598. doi: 10.1109/60.790920

Climente-Alarcon, V., Antonino-Daviu, J. A., Riera-Guasp, M., & Vlcek, M. (2014, aug). Induction motor diagnosis by advanced notch FIR filters and the wigner–ville distribution. *IEEE Transactions on Industrial Electronics*, *61*(8), 4217–4227. doi: 10.1109/tie.2013.2286581

Cruz, S., & Cardoso, A. (2001). Stator winding fault diagnosis in three-phase synchronous and asynchronous motors, by the extended Park's vector approach. *IEEE Transactions on Industry Applications*, *37*(5), 1227–1233. doi: 10.1109/28.952496

Frosini, L., Harlisca, C., & Szabo, L. (2015, mar). Induction Machine Bearing Fault Detection by Means of Statistical Processing of the Stray Flux Measurement. *IEEE Transactions on Industrial Electronics*(3), 1846–1854. doi: 10.1109/TIE.2014.2361115

Glowacz, A. (2019, feb). Fault diagnosis of single-phase induction motor based on acoustic signals. *Mechanical Systems and Signal Processing*, *117*, 65–80. doi: 10.1016/j.ymssp.2018.07.044

Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Mao, X.-J., Shen, C., & Yang, Y.-B. (2016). Image restoration using convolutional auto-encoders with symmetric skip connections. *arXiv preprint arXiv:1606.08921*.

Nandi, S., Toliyat, H. A., & Xiaodong, L. (2005). Condition monitoring and fault diagnosis of electrical motors - A review. *IEEE Transactions on Energy Conversion*, *20*(4), 719-29. doi: 10.1109/TEC.2005.847955

Rosin, P. L., & Fierens, F. (1995). Improving neural network generalisation. In *International geoscience and remote sensing symposium (igarss)* (Vol. 2, pp. 1255–1257). IEEE. doi: 10.1109/igarss.1995.521718

Springenberg, J. T., Dosovitskiy, A., Brox, T., & Riedmiller, M. (2014). Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*.

Thomson, W., & Fenger, M. (2003). Case histories of current signature analysis to detect faults in induction motor drives. In *Ieee international electric machines and drives conference, 2003. iemdc'03.* (Vol. 3, p. 1459-65). IEEE.

Tsypkin, M. (2017). The origin of the electromagnetic vibration of induction motors operating in modern industry: Practical experience—analysis and diagnostics. *IEEE Transactions on Industry Applications*, *53*(2), 1669-1676.

Vapnik, V. N. (2000). *The nature of statistical learning theory*. New York, NY: Springer New York. doi: 10.1007/978-1-4757-3264-1

Yeh, C.-C., Sizov, G. Y., Sayed-Ahmed, A., Demerdash, N. A. O., Povinelli, R. J., Yaz, E. E., & Ionel, D. M. (2008, dec). A reconfigurable motor for experimental emulation of stator winding interturn and broken bar faults in polyphase induction machines. *IEEE Transactions on Energy Conversion*, *23*(4), 1005–1014. doi: 10.1109/tec.2008.2001443