

# Octave-band Filtering for Convolutional Neural Network-based Diagnostics for Rotating Machinery

Namkyoung Lee<sup>1</sup>, Michael H. Azarian<sup>1</sup>, and Michael Pecht<sup>1</sup>

<sup>1</sup>*Center for Advanced Life Cycle Engineering (CALCE), University of Maryland, College Park, MD, 20742, USA*

*nkleee@umd.edu  
mazarian@umd.edu  
pecht@umd.edu*

## ABSTRACT

The performance of a machine learning model depends on the quality of the features used as input. Research into feature extraction methods for convolutional neural network (CNN)-based diagnostics for rotating machinery remains in a developmental stage. In general, the input to CNN-based diagnostics consists of a power spectrum without significant pre-processing. This paper introduces octave-band filtering as a feature extraction method for pre-processing a spectrum prior to use with a CNN. This method is an adaptation of a feature extraction method originally developed for noise measurement. The method developed for diagnosis of machinery faults differs from filtering methods applied to the noise measurement in that its filter design reflects the knowledge of a bearing's characteristic frequencies. The effectiveness of octave-band filtering is demonstrated through a case study. The method reduces the CNN size and improves the diagnostic accuracy in a noisy environment.

## 1. INTRODUCTION

Deep learning-based prognostics and health management systems have been inspired by audio-visual applications of deep learning such as image classification and speech recognition. The process of detecting faults for rotating machinery using vibration data is similar to the process of transcribing a speech. They both process one-dimensional time-domain data and find patterns in the data to determine outputs.

A speech recognition system comprises two models: an acoustic model and a language model. An acoustic model extracts acoustic characteristics from the system input and a language model estimates sequences of words using those characteristics. State-of-the-art speech recognition systems

(He et al., 2019, Synnaeve et al., 2019; Saon et al., 2017) use a combination of a convolutional neural network (CNN) and a recurrent neural network (RNN) to generate a combination of an acoustic and a language model.

Many deep learning-based diagnostics for rotating machinery that use CNNs to extract vibrational characteristics are an adaptation of the CNN structure of speech recognition systems. Janssens et al. (2016) used a one-dimensional CNN to diagnose machinery faults using a power spectrum of vibration data. The diagnostics often incorporate domain knowledge into a CNN structure. Zhao et al. (2017) developed a CNN structure having dynamic weighting layers that are applied to wavelet coefficients. The structure enables automatic feature selection while other systems select features manually during the training process. The usage of an RNN is optional when determining machinery health state as some machinery data may not have sequential information useful for diagnosing health.

In comparison with advancements in the structure of CNNs, feature extraction methods for diagnostics have room for improvement. Jiang et al. (2019) used the mel frequency cepstral coefficients (MFCCs) (Davis & Mermelstein, 1980) as inputs for CNN-based bearing diagnostics and demonstrated the robustness of the diagnostics to Gaussian noise. As the MFCCs are derived from the human auditory system, their direct use may result in the loss of useful information for detecting machinery faults. Although Sun et al. (2017) introduced compressed sensing to compress raw vibration data using a random Gaussian matrix, this method utilized the characteristics of generic vibration data. Previous researchers did not incorporate domain knowledge into the way they process vibration data. Therefore, their feature extraction methods may omit defect frequencies of rotating machinery such as characteristic frequencies of a bearing.

State-of-the-art speech recognition systems use an established domain-specific input, the mel frequency cepstrum (MFC). The MFC is a power spectrum that is

---

Namkyoung Lee et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

filtered by a mel filter bank. A mel filter bank consists of bandpass filters that analyze logarithmically scaled frequency bands. The filtering compensates for discrepancies between human nonlinear auditory perception and the spectral representation of sound on a linear scale. As a result, the MFC result encompasses a relatively large number of frequency components in low-frequency bands, and a small number of frequency components in high-frequency bands compared to frequency components' composition of a power spectrum that uses a linear frequency scale. The logarithmic scale, a mel scale, models human pitch interpretation. By infusing auditory knowledge into the feature extraction method, the performance of speech recognition models improves while reducing the size of original inputs.

This paper introduces octave-band filtering as a feature extraction method that is inspired by the development of the MFC. The method analyzes a power spectrum using an octave filter bank. The frequency bands of filters that comprise the bank are designed with the knowledge of a bearing's characteristic frequencies such as inner race ball pass frequency, and outer race ball pass frequency. The filtering minimizes the loss of useful information for diagnosis and improves diagnostic accuracy. The filtering also reduces the CNN size as a result of input size reduction.

## 2. METHOD

Octave-band filtering is spectral analysis that uses octave filters as bandpass filters for noise measurement. This paper uses the filtering to pre-process a power spectrum of vibration data prior to using a CNN for rotating machinery diagnosis. The vibration data are collected from an accelerometer, which is mounted near or at an element of rotating machinery. An overview of the developed pre-processing method that includes octave-band filtering is shown in Figure 1.

Vibration data are divided into several segments by sliding a Hann window. The windowing provides a certain length of time-domain data for Fourier transformation and compensates for spectral leakage. The magnitude squared of the Fourier transform generates a power spectrum, which is fed to a modified octave filter bank. The filter bank bins the spectrum at certain frequency bands, multiplies weights, and sums the weighted bands. The output of the filter bank, an octave frequency cepstrum (OFC), is converted to decibels and has a reference power of  $1 \text{ g}^2/\text{Hz}$ .

An octave filter bank is an array of octave bandpass filters that are defined based on the ANSI S1.11-2004 standard (2004). Similar to bandpass filters in a mel filter bank, the range that each octave filter handles, the bandwidth, is logarithmically scaled.

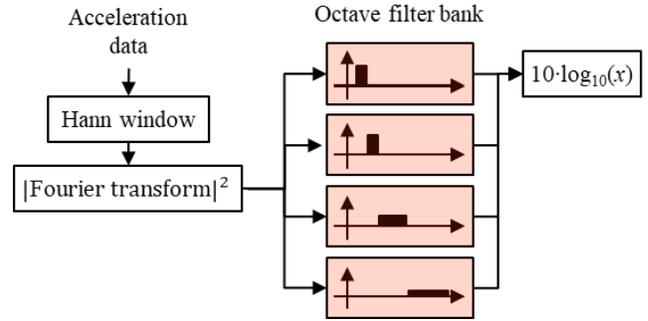


Figure 1. An overview of the developed feature extraction method.

The range of the  $i$ th filter's frequency band ( $f_i, f_{i+1}$ ) is defined as follows:

$$\frac{f_{i+1}}{f_i} = 2^{\frac{1}{l}}, i = 1, 2, 3, \dots, N \quad (1)$$

where  $l$  is the number of filters in one octave and determines the frequency resolution.

The bandwidth setting helps to identify and isolate characteristic frequencies in rotating machinery. In general, the characteristic frequencies can be categorized into two types. The peaks in a power spectrum that are located below 1 kHz are associated with the defective components' impact frequencies. These frequencies include inner race ball pass frequency (BPFI), outer race ball pass frequency (BPFO), and ball spin frequency (BSF). The locations of the frequencies are determined by the structural features of the defective components, and the distance between the frequencies is relatively short. Since octave filters have a narrow bandwidth in this frequency region, these frequencies can be efficiently isolated.

Another type of characteristic frequency is located at above 3 kHz. The peaks in the region are related to the natural resonances of the machinery components. Compared to the defect frequencies, the peaks have a broad width and small magnitude. A noisy environment that adds white noise to the vibration can attenuate the peaks' height and hamper finding the resonant frequencies. The octave filter's wide bandwidth in this frequency region helps identify these peaks.

Although the octave filter bank's logarithmic frequency scale is suitable for identifying the characteristic frequencies, the bandwidth in the low frequency region can be either too narrow or too broad to detect the defect frequencies. In order to resolve the issue, a modified octave filter bank is introduced that has a minimum bandwidth according to the minimum distance among the defect frequencies. The graph in Figure 2 shows an exemplary profile of the modified octave filter bank.

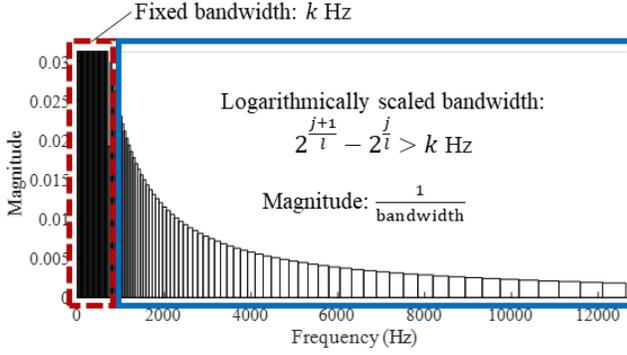


Figure 2. The profile of a modified octave filter bank.

Each bar in the graph of Figure 2 represents an octave bandpass filter. Based on the original bandwidth definition in Eq. (1), the first few filters that have a narrower bandwidth than  $k$  are replaced to filters that have a fixed bandwidth  $k$ . The bandwidth of the original filters before the replacement is too narrow to capture any defect frequencies. The fixed bandwidth ensures the minimum filter resolution to isolate the defect frequencies. The bandwidth  $k$  is determined by half the minimum distance among defect frequencies.

The filters that have a logarithmic frequency scale come after fixed bandwidth filters when the logarithmic bandwidth is larger than the fixed bandwidth  $k$ . In general, the transition occurs at around 1 kHz, and the logarithmic scaling contributes to feature size reduction. The bandwidth setting follows the original formula in Eq. (1), and the resolution of the filter is determined by the number of filters in one octave  $l$ . The filter number  $l$  should be set to contain the peak in a high-frequency band in one filter.

The height of the filters is defined as the inverse of the bandwidth. For example, in Figure 2, the filter that has a fixed bandwidth of 32 Hz has a height of  $1/32 \approx 0.03$ . During the filtering process, the sum of a spectrum's magnitudes within the filter's band is multiplied by the height. Therefore, the height normalizes the sum's magnitude.

### 3. A CASE STUDY

The efficacy of the octave-band filtering is demonstrated by applying the method to bearing datasets that are provided by Case Western Reserve University (CWRU). From the datasets, three features are extracted: a power spectrum, an MFC, and an OFC, which is the result of octave-band filtering. The three features and acceleration data without pre-processing are used as inputs to train CNN-based diagnostic models, and the models' performance is evaluated.

#### 3.1. Experimental setup

The CWRU dataset simulates three bearing faults by generating a dent on a drive end bearing's component that includes an inner race, an outer race, and a ball. The faults are

diagnosed by analyzing acceleration data that are sampled at 12 kHz while operating a 2-hp Reliance electric motor at around 1750 RPM. An accelerometer that collects the data is mounted at the 12 o'clock position at the drive end of the motor housing.

Acceleration data are prepared by splitting the data into 0.25-s time frames without overlapping, and the features are extracted by following the developed pre-processing method. From the datasets, 2164 time series data are obtained. The details about the data are summarized in Table 1. Among them, 60% of the data are used for training the CNN, 20% of the data are used for validation, and the rest of them are used for testing.

Table 1. The composition of the dataset

Fault	Size	Label
Normal	564	0
Ball fault	640	1
Inner race fault	640	2
Outer race fault	320	3

The data are processed using a 0.25-s Hann window and converted to spectral data through Fourier transform. By taking the absolute square of the data, a power spectrum that has 1501 frequency components is obtained. The spectrum is used for training the first CNN model and is also processed to generate an OFC and an MFC.

In order to obtain an OFC, an octave filter bank is designed. A fixed bandwidth is set using the specification of a drive end bearing (6205-2RS JEM SKF) that is provided by CWRU. Since the bearing's rotating frequency is about 29 Hz, the bearing's BPFI, BPFO, and BSF are 158 Hz, 105 Hz, and 137 Hz, respectively. Since the minimum distance among them is 21 Hz, the linear bandwidth  $k$  should be 10.5 Hz. Because the motor's rotating speed varies over time, the bandwidth is set to 8 Hz, which gives a safety margin to segregate the bearing's defect frequencies. The number of filters in one octave  $l$  is set to 16 after examining some spectral data to determine if the setting identifies global maxima in the frequency region where elements in the electric motor resonance. The filter bank designed by this process generates an OFC that has 103 frequency components.

The process of generating an MFC is similar to the OFC generation process. The mel scale is calculated using Slaney's formula (1998), which is implemented in MATLAB. The formula creates the scale between the lowest and highest frequency to be expressed, with the number of filters as a hyperparameter. The lowest and highest frequencies are set here at 0 Hz and 6 kHz, accordingly. In

order to make a fair comparison with the CNN model that is trained using an OFC, the number of filters is set to 103.

All spectral features are rescaled from 0 to 1 using a constrained min-max scaling function as follows:

$$x' = \frac{\max(\min(x, M), m) - m}{M - m} \quad (2)$$

where  $M$  and  $m$  represent the maximum and minimum input threshold, accordingly. In this experiment,  $M$  and  $m$  are set to 0 and  $-80$ , respectively. If an input  $x$  falls outside the range defined by the thresholds, the scaling function sets the input to the nearest threshold.

The CNN models diagnose the drive end bearing's health state regardless of the severity of the damage and the motor load. The high-level schematic of the models is shown in Figure 3. The models have a base CNN structure to find patterns in spectral features that are related to the bearing's state and determine the state of the bearing. The three CNN models that have hand-crafted spectral features as inputs only use the base structure. The CNN model that has raw acceleration data as inputs extracts spectral features by inputting the data into a convolutional module.

The models' base structure is shown in Figure 4. The structure. The input layer size of the structure is determined by the size of spectral features to be analyzed. To avoid overfitting, the structure has only one convolutional module and two fully connected layers. The structure's inputs go through a one-dimensional convolutional layer (Conv 1D) with eight filters. The output of the Conv 1D, called a feature map, is compressed using a one-dimensional max-pooling layer (MaxPool 1D). This layer finds maximum values while sliding a window.

The compressed outputs are flattened and fed to a fully connected layer (FC) to find correlations between the health state and the outputs. The Conv 1D and FC use a rectified linear unit (ReLU) function as an activation function that is defined as follows:

$$f(z) = \max(0, z) \quad (3)$$

where  $z$  is the result of convolution. This function enables the model to express nonlinear relationships between inputs and outputs.

The FC is connected to an output layer to decide the health state using a softmax function as follows:

$$f(z_i) = \frac{e^{z_i}}{\sum_{k \in K} e^{z_k}} \quad (4)$$

where  $k$  is the order of output nodes, and  $z$  is the output of the FC's matrix multiplication.

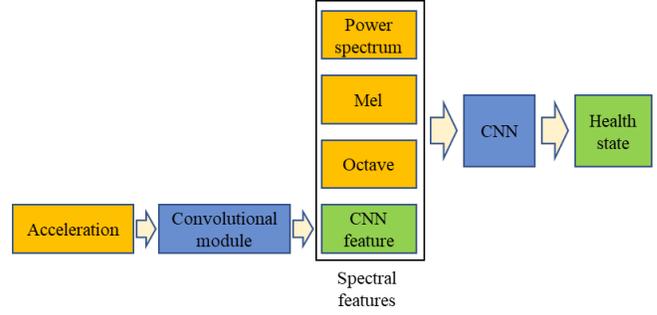


Figure 3. An overview of convolutional neural network model for rotating machinery diagnosis.

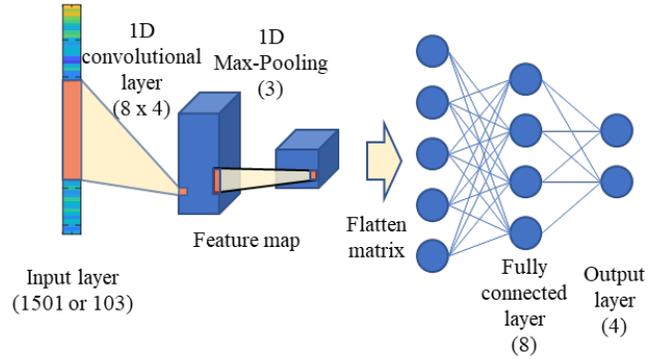


Figure 4. A convolutional neural network structure for rotating machinery diagnosis.

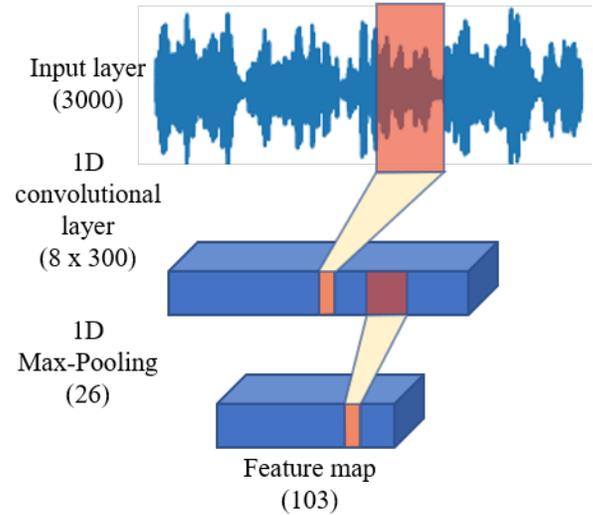


Figure 5. A convolutional module to extract spectral features from raw acceleration data.

The structure has four nodes that indicate four states: healthy, an inner race fault, an outer race fault, and a fault on a ball. The model decides the health state by selecting the node that has the maximum value among them.

The convolutional module is shown in Figure 5. This module was devised by Hoshen et al. (2015) for speech recognition. Their work showed that a convolutional layer can function as finite impulse response (FIR) bandpass filters.

The Conv 1D in the module in Figure 5 works as FIR bandpass filters and the MaxPool 1D compresses the Conv 1D's outputs. Since the filter length determines frequency resolution, the kernel size of the Conv 1D is set to 300. This setting enables the filter to decompose accelerations with about 43 Hz frequency resolution.

In order to match the output size of the module with the feature size of the MFC and the OFC, the pooling size of the module is set to 26. The output of the module, a feature map, replaces hand-crafted spectral features and is fed to the base structure's input.

All CNN models are trained using an Adam optimizer (Kingma & Ba, 2014) with constant learning rate of 0.001. Each neural layer's weights in the models are initialized depending on its activation function type. A ReLU based layer's weights are initialized using the He initializer (He et al., 2015), and a softmax based layer's weights are initialized using the Xavier initialization (Glorot & Bengio, 2010). All neural layers' biases are initialized to zero. The training continued until the training epoch reached 80 and the model that shows the minimum validation loss during the training is selected for testing.

### 3.2. Experimental results

In order to show the characteristics of an OFC in comparison with other spectral features, a power spectrum, an MFC, and an OFC are extracted from the same acceleration data and plotted in Figure 6. The asterisk and circle markers in the graph represent the central frequencies of the MFC and the OFC, respectively.

Octave-band filtering compresses the size of a power spectrum by 93% while preserving the characteristic frequencies of a bearing. The MFC reproduces the overall shape of the power spectrum, including the approximate location of prominent minima and maxima. Although the MFC has the same number of compressed features, the OFC has higher frequency resolution in the low-frequency region. This difference impacts the identification of bearing defect frequencies.

In Figure 7, the graph's low-frequency region (0–250 Hz) is magnified, and three defect frequencies are annotated using vertical dashed lines. In the graph, the MFC has only nine frequency components in the region, while the OFC has 29 frequency components. Both the MFC and the OFC capture the peaks at the BPFO and BPF1 because both cepstral central frequencies are located near the defect frequencies.

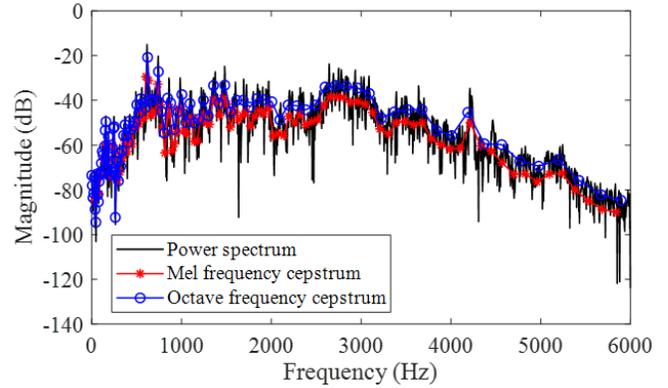


Figure 6. Spectral features of a data sample: a power spectrum, a mel frequency cepstrum, and an octave frequency cepstrum.

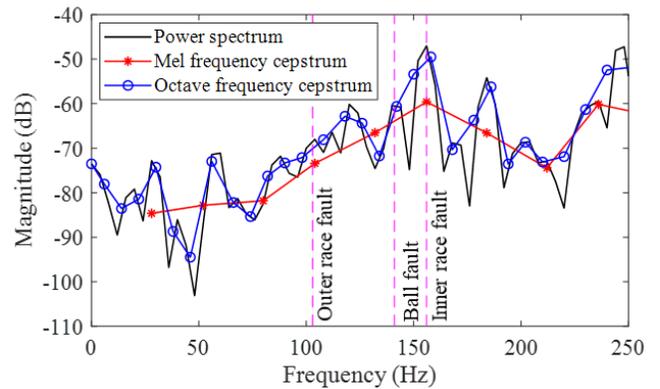


Figure 7. Bearing's defect frequencies and center frequencies for three spectral features

On the other hand, the MFC is not able to localize the BSF. The MFC's central frequencies are located away from the BSF. Moreover, if the motor speed were to decrease, the distance between the defect frequencies would get closer, so the MFC would not distinguish between a BSF and a BPF1.

In order to reveal the observed differences among the MFC and the OFC in the input analysis, the acceleration data for testing are perturbed by adding white Gaussian noise into the data. The noise level is set to a 5 dB signal-to-noise ratio. The added noise distorts the landscape of spectral features. However, the frequency components that are related to defect frequencies are not distorted.

All CNN models are trained and validated using unperturbed training data, and the perturbed data are used for testing. This setup reveals whether the CNN models interpreted the defective frequencies as important features for the models' decisions. The noise fades some patterns in the spectral features that are generated due to the dependency among data. If a trained model learned these features to decide the

health state, this manipulation can interfere with the model's decision resulting in low test accuracy.

To show how models converge, exemplary loss curves of the four CNN models are shown in Figure 8. The graphs' titles in Figure 8 represent the trained models' input type. The models' training and validation losses were recorded at every epoch. All models converged in the training process, but their convergence rates were different depending on the number of spectral features that were used for diagnosis.

The model that uses a power spectrum had 1501 spectral features as inputs, and it converged at 10 epochs. On the other hand, the other models have 103 spectral features as inputs, and they converged at about 40 epochs. This result indicates that the redundant features in the power spectrum accelerate a model's learning curve, although this redundancy does not improve a model's accuracy.

The performance of the four models for the second experiment is listed in Table 2.

Table 2. Performance of convolutional neural network-based diagnostics

Input	Accuracy		Number of Parameters
	Mean	Standard Deviation	
Octave frequency cepstrum	98.66 %	0.69 %	2,196
Mel frequency cepstrum	91.22 %	8.63 %	2,196
Power spectrum	98.06 %	1.39 %	32,020
Acceleration	96.54 %	5.82 %	4,828

The training and testing procedure are conducted five times without replacing data, and the mean and standard deviation of the results are calculated. This experiment evaluates the CNN models' capability to learn informative features. The model that uses an MFC (MFC model) shows the worst accuracy due to the model's overfitting. Three test records out of five have accuracy below 95%, which reduces the mean accuracy.

The accuracy of the model that uses raw acceleration data showed a similar problem. One 85% accuracy test result dropped the mean accuracy. This result can be related to the model's learning process. Since the model learns spectral features by itself, the features do not represent the characteristics of the system with as much fidelity as the features obtained from the other methods.

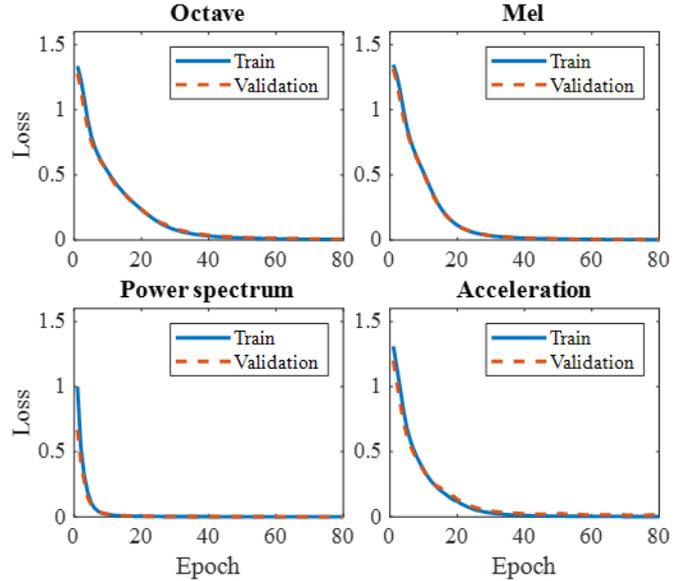


Figure 8. Loss curves of four models while training

The other two models gave consistent test accuracy, and the difference between the two models was not significant. The precisions and recalls of the results did not show significant difference (0.1% difference at most) and the values were closed to the accuracy results.

The confusion matrix in Figure 9 shows the one of the overfitted MFC model's diagnostic results. The labels in the matrix represent the states of the bearings. B, IR, and OR are the abbreviations of a defect on a ball, an inner race defect, and an outer race defect. The model is prone to detect inner race faults over other faults. This result supports the analysis of features provided above with reference to Figure 6.

The model's overfitting may result from the lack of information about the defect frequencies. The deficiency makes the model rely on frequency characteristics in high-frequency bands that may have a spurious correlation with the bearing's health.

The results of the model that uses an OFC (OFC model) in Figure 10 contrast with the results of the MFC model. Although both the MFC and the OFC have the same number of features, the OFC model is able to classify defective bearings correctly. These results show the role of defect frequencies in diagnostics.

The number of parameters of the CNN models is related to the size of the model and the execution time. The power spectrum model is about 14.5 times larger than the MFC and the OFC models. The difference between the power spectrum model and other models results in the connection between the flattened feature map and the FC.

Estimated state	Normal	99 22.9%	0 0.0%	0 0.0%	0 0.0%	100% 0.0%
	B	0 0.0%	97 22.4%	0 0.0%	1 0.2%	99.0% 1.0%
	IR	0 0.0%	31 7.2%	141 32.6%	47 10.9%	64.4% 35.6%
	OR	0 0.0%	0 0.0%	0 0.0%	17 3.9%	100% 0.0%
		100% 0.0%	75.8% 24.2%	100% 0.0%	26.2% 73.8%	81.8% 18.2%
	Normal	B	IR	OR		
	True state					

Figure 9. A confusion matrix that shows the diagnostic results of the MFC model.

Estimated state	Normal	99 22.9%	0 0.0%	0 0.0%	0 0.0%	100% 0.0%
	B	0 0.0%	128 29.6%	6 1.4%	0 0.0%	95.5% 4.5%
	IR	0 0.0%	0 0.0%	135 31.2%	0 0.0%	100% 0.0%
	OR	0 0.0%	0 0.0%	0 0.0%	65 15.0%	100% 0.0%
		100% 0.0%	100% 0.0%	95.7% 4.3%	100% 0.0%	98.6% 1.4%
	Normal	B	IR	OR		
	True state					

Figure 10. A confusion matrix that shows the diagnosis results of the OFC model.

As the input for a power spectrum processed by a CNN is large, the feature map is also large. Since every component in the feature map is connected to FC, the CNN model processing a power spectrum has a large number of parameters compared to the other two CNN models.

#### 4. CONCLUSIONS

Vibration-based diagnostics for rotating machinery employing a convolutional neural network (CNN) uses

domain knowledge to improve its accuracy. Creating features that reflect the machinery vibrational characteristics is a means of infusing domain knowledge into diagnostics.

This paper introduces octave-band filtering as a feature extraction method and develops an approach to incorporate knowledge about the machinery’s characteristic frequencies into the filtering. The octave-band filtering outputs spectral features by analyzing the power spectrum of vibration data. The developed approach designs bandpass filters that select the frequency bands to be analyzed during the filtering for better identification of the machinery’s characteristic frequencies.

The frequency bands in the low-frequency region, usually the bands under 1 kHz where the bearing defect frequencies are located, have a fixed narrow width to improve the separation of those frequencies from each other. The frequency bands that are above the low-frequency region are logarithmically scaled, resulting in broader bandwidth compared to the fixed width bands. This scaling helps the CNN to identify resonant frequencies of a bearing’s components because the resonances affect the magnitudes of a wide range of neighboring frequencies. The scaling is also applicable to other rotating machinery elements such as gears and shafts, since these elements’ defect frequencies are also related to the rotating speed of machinery and their resonance.

The case study demonstrated the performance improvement of CNN-based diagnostics due to octave-band filtering. The filtering reduced the spectrum data by 93%, resulting in a decrease of the CNN input size by the same amount. The decrease in the CNN input size did not affect the diagnostic accuracy. The accuracy improved slightly by 0.6%, achieving 98.66% accuracy in a 5 dB signal-to-noise ratio white noise environment, compared to a CNN using spectrum data without octave-band filtering.

The same case study compared the method with mel-scale filtering that inspired the development of octave-band filtering. The mel-scale filtering extracted the same number of spectral features, but its features reflect human auditory perception, resulting in loss of information about characteristic frequencies of machinery in low-frequency bands. The implementation of the mel-scale filtering dropped the accuracy by 7.44% compared to the CNN with octave-band filtering in the same noise environment.

The results of the case study show that logarithmic frequency scaling is suitable for compressing power spectrum of machinery vibration. In particular, octave-band filtering improves the quality of compression result since it infuses knowledge about defect frequencies of a bearing into frequency filtering process. The improved quality can reduce the risk of overfitting during the training and enhances robustness to noise.

## REFERENCES

- American National Standards Institute (ANSI) (2004). Specification for octave-band and fractional-octave-band analog and digital filters. In ANSI, *ANSI S1.11-2004*, New York, United States, Acoustical Society of America.
- Case Western Reserve University. (2020, May, 14). Retrieved from <http://csegroups.case.edu/bearingdatacenter/pages/welcome-case-western-reserve-university-bearing-data-center-website>
- Davis, S. B., & Mermelstein, P. (1980). Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 28(4), 357-366.
- Glorot, X., & Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. *International conference on artificial intelligence and statistics*, Sardinia, Italy, 249-256).
- He, Y., Sainath, T. N., Prabhavalkar, R., McGraw, I., Alvarez, R., Zhao, D., Rybach, D., Kannan, A., Wu, Y., Pang, R., Liang, Q., Bhatia, D., Shangguan, Y, Li, B., Pundak, G., Sim, K. C., Bagby, T., Chang, S., Rao, K., & Gruenstein, A. (2019). Streaming end-to-end speech recognition for mobile devices. *IEEE International Conference on Acoustics, Speech and Signal Processing*, Brighton, United Kingdom, 6381-6385.
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. *IEEE international conference on computer vision*, Santiago, Chile, 1026-1034.
- Hoshen, Y., Weiss, R. J., & Wilson, K. W. (2015). Speech acoustic modeling from raw multichannel waveforms. *IEEE International Conference on Acoustics, Speech and Signal Processing*, Brisbane, Australia, 4624-4628.
- Janssens, O., Slavkovikj, V., Vervisch, B., Stockman, K., Loccufier, M., Verstockt, S., Van de Wall, R., & Van Hoecke, S. (2016). Convolutional neural network-based fault detection for rotating machinery. *Journal of Sound and Vibration*, 377, 331-345.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint*, arXiv:1412.6980.
- Saon, G., Kurata, G., Sercu, T., Audhkhasi, K., Thomas, S., Dimitriadis, D., Cui, X., Ramabhadran, B., Picheny, M., Lim, L., Roomi, B., & Hall P. (2017). English conversational telephone speech recognition by humans and machines. *arXiv preprint*, arXiv:1703.02136.
- Sun, J., Yan, C., & Wen, J. (2017). Intelligent bearing fault diagnosis method combining compressed data acquisition and deep learning. *IEEE Transactions on Instrumentation and Measurement*, 67(1), 185-195.
- Slaney, M. (1998). *Auditory Toolbox: A MATLAB Toolbox for Auditory Modeling Work*. Technical Report, version 2, Interval Research Corporation
- Synnaeve, G., Xu, Q., Kahn, J., Grave, E., Likhomanenko, T., Pratap, V., Sriram A., Liptchinsky, V., & Collobert, R. (2019). End-to-end ASR: From supervised to semi-supervised learning with modern architectures. *arXiv preprint*, arXiv:1911.08460.
- Jiang, Q., Chang, F., & Sheng, B. (2019). Bearing fault classification based on convolutional neural network in noise environment. *IEEE Access*, 7, 69795-69807.
- Zhao, M., Kang, M., Tang, B., & Pecht, M. (2017). Deep residual networks with dynamically weighted wavelet coefficients for fault diagnosis of planetary gearboxes. *IEEE Transactions on Industrial Electronics*, 65(5), 4290-4300.

## BIOGRAPHIES

**Namkyoung Lee** received the B.S. degree in electronics from Hanyang University, Seoul, Korea in 2015 and the M.S. degree in computer science from the Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea in 2017. He is currently working toward a Ph.D. degree in mechanical engineering from the University of Maryland, College Park, MD, USA. His research interests include data-driven prognostics and health management system for rotating machinery, and diagnostics and prognostics of automotive components.

**Michael H. Azarian** is a Research Scientist with the Center for Advanced Life Cycle Engineering, University of Maryland, College Park, MD, USA. He received the B.S.E. degree in chemical engineering from Princeton University, Princeton, NJ, USA, and the M.E. and Ph.D. degrees in materials science and engineering from Carnegie Mellon University, Pittsburgh, PA, USA. He spent over 13 years in industry. His research interests include the analysis, detection, prediction, and prevention of failures in electronic and electromechanical products. He holds five U.S. patents.

Dr. Azarian is chair of the SAE G-19A Test Laboratory Standards Development Committee, which is responsible for the AS6171 family of standards on detection of counterfeit electrical, electronic, and electromechanical parts. He also co-chairs the working group responsible for the IEEE 1624 standard on organizational reliability capability of suppliers of electronic products.

**Michael G. Pecht** received the B.S. degree in acoustics, M.S. degrees in electrical engineering and engineering mechanics, and Ph.D. degree in engineering mechanics from the University of Wisconsin at Madison, Madison, WI, USA, in 1976, 1978, 1979, and 1982, respectively. He is the Founder of the Center for Advanced Life Cycle Engineering, University of Maryland, College Park, MD, USA, where he is also a Chair Professor. He has been leading a research team in the area of prognostics.

Dr. Pecht is a Professional Engineer and a Fellow of the American Society of Mechanical Engineers. He was the recipient of the IEEE Undergraduate Teaching Award and the International Microelectronics Assembly and Packaging Society William D. Ashman Memorial Achievement Award

for his contributions in electronics reliability analysis. He served as the Chief Editor of IEEE Transactions on Reliability for eight years and an Associate Editor for IEEE Transactions on Components and Packaging Technologies.