



# Interpretable Acoustic Representation Learning on Breathing and Speech Signals for COVID-19 Detection

Debottam Dutta<sup>1</sup>, Debarpan Bhattacharya<sup>1</sup>, Sriram Ganapathy<sup>1</sup>,  
Amir H. Poorjam<sup>2</sup>, Deepak Mittal<sup>2</sup>, Maneesh Singh<sup>2</sup>.

<sup>1</sup>LEAP Lab, Indian Institute of Science, Bangalore, India.

<sup>2</sup>Verisk Analytics, Inc., Jersey City, NJ, USA

sriramg@iisc.ac.in

## Abstract

In this paper, we describe an approach for representation learning of audio signals for the task of COVID-19 detection. The raw audio samples are processed with a bank of 1-D convolutional filters that are parameterized as cosine modulated Gaussian functions. The choice of these kernels allows the interpretation of the filterbanks as smooth band-pass filters. The filtered outputs are pooled, log-compressed and used in a self-attention based relevance weighting mechanism. The relevance weighting emphasizes the key regions of the time-frequency decomposition that are important for the downstream task. The subsequent layers of the model consist of a recurrent architecture and the models are trained for a COVID-19 detection task. In our experiments on the Coswara data set, we show that the proposed model achieves significant performance improvements over the baseline system as well as other representation learning approaches. Further, the approach proposed is shown to be uniformly applicable for speech and breathing signals and for transfer learning from a larger data set.

**Index Terms:** Learnable filterbank, representation learning, self-supervised learning, COVID-19 diagnosis

## 1. Introduction

The biomedical applications, like automatic diagnosis of diseases using speech and audio, have a strong requirement to explain away the basis the model is using to arrive at a specific decision about a sample [1–3]. The early approaches to facilitate this need involved the design of knowledge-driven features [4, 5]. A wide range of acoustic features that can reflect different voice disorders have been designed, such as jitter, which indicates the instability in the fundamental frequency [6], shimmer, which is a measure of deviations in amplitude [7], harmonic-to-noise ratio, which is an indication of hoarseness in voice [8], and maximum phonation time that indicates the lung capacity [9]. However, these approaches require extensive experiments to understand their impact on the downstream application.

In the recent years, with advancements in the deep learning approaches, hand-crafted feature extraction algorithms have been increasingly replaced by data-driven methods. This field of study, called representation learning [10], is the one in which a deep neural network learns representations from raw data without any prior assumptions. Representation learning is a well-explored field in computer vision [11] and natural language processing [12], where the representation of the data is learned by directly feeding pixels or one-hot encoded words to the neural network, respectively. Recently, there has also been

a great interest in learning general purpose representations for speech and audio [13, 14]. However, the high temporal variability and the curse of dimensionality of the acoustic signals, makes the representation learning challenging [15]. In this paper, we explore representation learning for speech directly from the raw-waveform using a modeling approach that consists of a parametric filter learning module and a self-attention style relevance weighting module. The representation learning approach pursued in this work aims at developing robust front-ends for the task of detecting COVID-19 from audio sounds of speech and breathing nature.

In audio processing, the most common approach for representation learning has been the direction of learning the time-frequency decomposition from the raw audio. These approaches can be categorized as supervised and unsupervised methods. Examples for supervised learning include phoneme classification [16], acoustic modeling for speech recognition [17], acoustic event classification using learnable audio front-end (LEAF) [18], noisy speech recognition [19] and music genre classification using deep scattering transform [20]. Using parameterized sinc functions, Ravanelli and Bengio proposed a SincNet architecture [21] in which the first layer of the convolutional neural network (CNN) learns meaningful filters directly from the raw audio signals.

For unsupervised learning of audio representations, the use of restricted Boltzmann machines [22, 23] and variational learning [24] have been explored. The wav2vec method proposed in [25] explored self-supervised pre-training using an autoregressive loss function. Further, the generative adversarial networks [26] have been explored for acoustic/modulation filter learning.

The approach of self-supervised learning from a number of tasks has been explored by Pascual et al. [27]. In these applications, while representation learning has achieved performance improvements for the downstream tasks considered, they have been limited in terms of the interpretability. The key exception is the SincNet filterbank [21] where the authors used parametric sinc filters to improve interpretability.

In this work, we extend the previous work by Agrawal et al. [28] on filterbank learning using cosine modulated Gaussian filter-banks. We explore the learning of representations on breathing and speech signals for the task of COVID-19 detection. The raw audio samples are processed through a bank of 1-D convolutional filters. A self-attention based sub-band weighting approach, called relevance weighting, provides rich representations of the audio signal for the downstream recurrent neural network layers. While the architecture is similar to the prior work [28], the key contributions from the proposed work over the previous work in [28] are,

This work was supported by the grants from Verisk Analytics Inc.

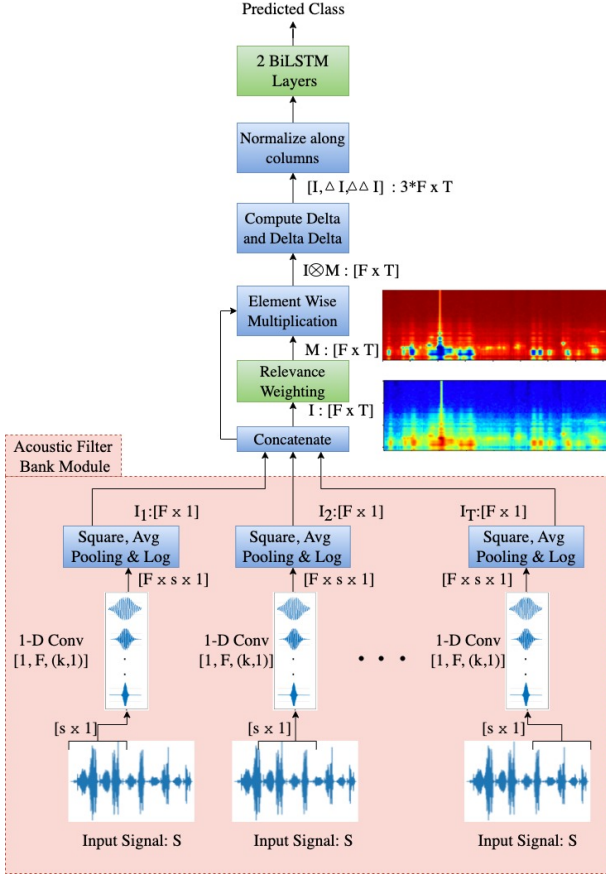


Figure 1: Overall pipeline for filter learning. The input signal is passed through a 1-D convolutional layer where we use Gaussian set of filters. The relevance weighting module consists of a fully connected layer followed by a sigmoid function. Two maps are also shown as output of convolution operation and the corresponding learned mask. Blocks in green depict a learnable function whereas the blue blocks depict fixed functions.

- A unified filter learning paradigm that provides meaningful representations from both speech and breathing signals.
- Transfer learning of the filter representations from a large scale data set of speech/breathing sounds to the final target data set.
- Exploring supervised and self-supervised representation learning objectives for acoustic filter learning.
- Experimental evidences on COVID-19 detection task to highlight the advantages of the proposed representation learning framework.

## 2. Acoustic Filterbank with Relevance Weighting

The block schematic of the proposed front-end is shown in Figure 1. In the following subsections, we explain the main components of the proposed method.

### 2.1. Acoustic Filterbank Layer

The acoustic filterbank layer receives raw audio samples which are windowed into a sequence of  $s$  samples. The windowed




Figure 2: (a) Mel spectrogram, (b) the learned spectrogram and (c) relevance weighted spectrogram of a breathing sample.

samples are processed using a bank of 1-D convolutional filters which are parameterized as a cosine modulated Gaussian function [28], where the kernels  $g_i(n)$ , corresponding to the  $i$ th filter, are denoted as

$$g_i(n) = \cos 2\pi\mu_i n \times \exp(-n^2 \mu_i^2 / 2), \quad (1)$$

where ( $i = 1, 2, \dots, F$ ) and  $\mu_i$  is the centre frequency of the  $i$ -th kernel. Here,  $F$  denotes the number of acoustic filters. The Gaussian function, given in the above equation, represents a low-pass filter while the cosine modulation makes the kernel band-pass. The center frequencies  $\mu_i$  are the only learnable parameters of the filter. The band-width of the filters is directly proportional to the center frequency (inverse of the variance in the time-domain kernel function in the above equation). This is to enable the learning of constant-Q filters.

The given audio signal is first segmented into short-time frames of  $s$  samples. Each frame is then convolved with the  $F$  filters. Finally, the output of the filters are squared, average pooled and log-transformed. This processing module generates an  $F$  dimensional output for each frame. Further, splicing the frame-level  $F$  dimensional representations from all  $T$  frames of the given audio file, gives the learned time-frequency representation,  $I$ , of dimension  $F \times T$ .

### 2.2. Relevance Weighting

The generated spectrogram of dimension  $F \times T$  is passed to a self-attention-based neural network consisting of two layers of a feed forward network. For each time-frequency bin ( $t, f$ ), we input a vector of 102 dimensions obtained for the  $i$ -th filterbank representation with  $\pm 51$  frame window on either side of the current frame. The output of the relevance network is a scalar that is passed through a sigmoid function to generate a relevance weight for the ( $t, f$ ) bin. The relevance weights generated for all the ( $t, f$ ) bins (denoted as  $M$  in Figure 1) are then element-wise multiplied with the learned representation  $I$  to generate the relevance weighted ( $t, f$ ) representation ( $J = I \otimes M$ ). This final representation  $J$  is used in the subsequent model for the task for COVID-19 detection.

The spectrogram of a breathing signal, generated using the mel-scale filterbank and the learned filters, are shown in Figure 2 (a) and (b), respectively. We can observe that the model distributes the filters such that low- and high-frequency bands are more stretched, compared to the mel spectrogram, to provide more detailed representation of those regions. The weighted spectrogram is presented in Figure 2 (c). Compared with (b), we observe that the high frequency regions are more enhanced than other regions in the relevance weighted output.

Table 1: Comparison of the performance of the COVID-19 detection models using different signal representations on breathing and speech modalities in terms of the AUC(%).

Representation Methods	Breathing						Speech					
	Fold1	Fold2	Fold3	Fold4	Fold5	Avg.	Fold1	Fold2	Fold3	Fold4	Fold5	Avg.
Mel-Spectrogram [29]	73.8	76.4	75.6	77.1	84.2	77.4	74.8	85.1	79.1	78.6	80.4	79.6
SincNet [21]	72.4	69.6	<b>81.9</b>	76.4	77.1	75.4	75.0	84.6	77.4	76.7	<b>85.2</b>	79.7
LEAF [18]	71.5	62.9	66.1	66.6	73.6	68.2	60.5	59.1	63.8	57.4	62.0	60.6
CosGauss	75.8	74.3	79.4	78.3	78.6	77.3	77.0	81.5	83.5	77.5	82.9	80.4
CosGauss-relev	75.4	76.0	79.7	78.1	82.8	78.4	76.5	<b>83.6</b>	83.3	<b>79.6</b>	84.5	<b>81.5</b>
CosGauss(maxpool)-relev	75.0	72.2	78.0	76.6	79.7	76.3	77.4	80.7	81.2	78.1	82.6	80.0
CosGauss-relev-pretr	76.1	75.9	77.7	76.9	83.2	78.0	77.1	79.7	<b>84.7</b>	77.0	81.3	80.0
CosGauss-relev-pretr-fine	<b>81.2</b>	77.0	80.9	78.4	<b>84.9</b>	<b>80.5</b>	78.3	81.5	83.5	78.0	81.5	80.6
CosGauss-SSL-pretr	77.3	73.9	79.6	79.3	81.6	78.3	78.7	80.5	81.9	75.8	81.8	79.7
CosGauss-SSL-pretr-fine	78.5	<b>77.9</b>	77.9	<b>80.2</b>	81.3	79.2	<b>80.7</b>	81.7	82.7	76.8	78.3	80.0

### 2.3. Back-end Model Architecture

The relevance weighted representations are used to compute the delta and double-delta features. For the supervised training, we base the rest of the architecture design on the baseline system provided as part of the Second DiCOVA challenge [29]. The model consists of 2 layers of a bidirectional long-short term memory (BLSTM) network followed by a two class posterior output. The model is trained using the binary cross entropy (BCE) loss<sup>1</sup>.

## 3. Transfer Learning of Representations

### 3.1. Data Resources

To develop and evaluate the proposed method, we used two data sets: The COVID-19 Sounds [30] and the Coswara [29] dataset.

#### 3.1.1. COVID-19 Sounds Data Set

This data set, which is crowd-sourced from 36, 116 participants from around the world, consists of 53, 449 audio samples of three different acoustic modalities, namely speech, cough, and breathing. In addition to the audio samples, demographic information, health conditions of the participants, and participants' self-reported COVID-19 testing status are also provided. Since the testing status is self-reported, we have selected participants who tested positive within the 14 days of the recording as the positive class, and those who never tested positive as the negative class. Moreover, as the samples were captured from a wide range of platforms, the sampling frequency and the audio format of the samples are diverse. Therefore, we only selected the samples of 44.1 kHz sampling frequency and those with "wav" audio file format. The selected samples are then split into training and validation subsets of 1, 810 and 440 samples, respectively for speech modality, and 1, 875 and 469, respectively for breathing modality.

#### 3.1.2. Coswara Data Set

The Coswara data set [31] is a crowd-sourced dataset from more than 1, 500 participants across the world. During the data collection, participants were asked to record their voice in various modalities, namely breathing, cough, sustained vowel /a/ phonation, and speech. Along with the voice samples, demographic information and self-reported COVID-19 status are also collected. The samples were captured through a website appli-

<sup>1</sup>The code and models are available at: [https://github.com/iiscleap/acoustic\\_repLearn\\_dicova2](https://github.com/iiscleap/acoustic_repLearn_dicova2)

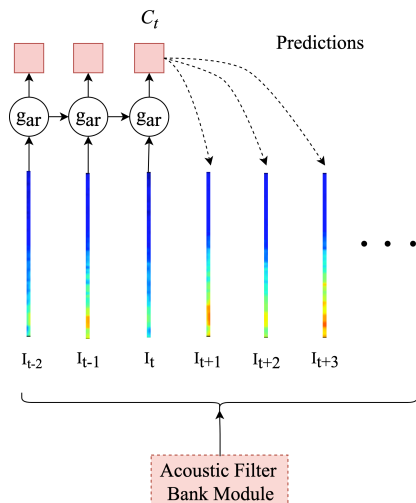


Figure 3: The proposed CPC-based architecture for self-supervised filter learning.  $g_{ar}$  and  $c_t$  represent the autoregressive module and the context vector, respectively [32]. The input to each  $g_{ar}$  is of dimension  $F \times 1$ . The detailed architecture of acoustic Filter Bank Module is shown in Figure 1.

cation and at a sampling frequency of 48 kHz. For this study, we used a subset of the Coswara data set of 965 participants split into 5 folds of 80%–20% training-validation, provided by the organizing committee of the Second DiCOVA challenge [29].

### 3.2. Supervised Pre-training

In supervised pre-training, we first learn the filter on the larger data set along with the target labels. In our case, we use the COVID-19 Sounds data for the pre-training. The filters trained using this data set are then used to initialize filters for learning on the Coswara data. We report results in Table 1 for two settings, one in which pre-trained filters are not fine-tuned (CosGauss-relev-pretr) and one in which pretrained filters are fine-tuned (CosGauss-relev-pretr-fine) using the Coswara data.

### 3.3. Self-supervised Pre-training

For our experiments, we used the contrastive predictive coding (CPC) [32], which is an autoregressive method. The input audio signals are first converted into compact representations using the encoder module. Then, these representations are used to predict future representations using autoregressive models implemented as a recurrent layer.




Figure 4: Distribution of centre frequencies of mel and CosGauss filters learned from Coswara data with pre-training in a supervised (CosGauss-relev-pretr-fine) and self-supervised (CosGauss-SSL-pretr-fine) fashion.

The proposed architecture for self-supervised learning is shown in Figure 3. The acoustic filterbank module, described in Figure 1, generates representations at the frame-level of dimension  $F \times 1$ . These are then passed through an LSTM block. The CPC training is performed on the COVID-19 Sounds samples without using the label information. We hypothesize this scenario to reflect semi-supervised learning frameworks, where a large data set (in our case, the COVID-19 Sounds data set) is available without label information, while the target data set (in our case, the Coswara data set) is labelled but smaller in size.

We used the COVID-19 Sounds data set to pre-train the filters in self-supervised manner. Then, the pre-trained filters are used to initialize filters in our model which are then fine-tuned on the Coswara data.

## 4. Results

Table 1 compares the performance of different signal representations on breathing and speech modalities in terms of the area under the receiver operating characteristic curve (AUC). The representation learning methods above the dotted line are the baseline approaches compared in this paper. The first row shows the mel-spectrogram representation, followed by the SincNet representations [21] in the second row. The third row shows the performance of LEAF representations [18].

The fourth and fifth row shows the representations based on cosine modulated Gaussian filters without and with relevance weighting respectively. The sixth row is similar to the fifth, except that the average pooling operation in acoustic filterbank layer is replaced by max-pooling. Out of these three, cosine modulated Gaussian filters with relevance weighting (fifth row) shows the best result for which the results improve over the baseline approaches (absolute improvements of 1% in AUC for breathing modality and 1.9% in the case of speech modality) for both the modalities of speech and breathing. Further, the improvements are consistent on four of the five folds in speech and three of the five folds in breathing.

The seventh and eighth rows show the results for supervised pre-training without and with fine-tuning, respectively. The results for self-supervised pre-training without and with fine-tuning are represented in the ninth and tenth rows, respectively. These results represent the various transfer learning approaches

explored in this work. The following are the major insights derived from these experiments,

- The results with pre-training on the COVID-19 Sounds data set alone in supervised (seventh row) or self-supervised fashion (ninth row) without fine-tuning degrades the performance over the direct supervised learning of the filters on the Coswara data set (fifth row) for both the speech and breathing modalities.
- Using the fine-tuning followed the pre-training improves the performance of the transfer learning. The supervised pre-training with fine-tuning (eighth row) offers improved performance over the self-supervised pre-training (tenth row) for speech and breathing modalities.
- For the setting of supervised pre-training with fine-tuning, the transfer learning improves significantly (absolute improvements of 2.1%) in the case of breathing modality while the transfer learning approaches are inferior to the one trained directly on the Coswara data set for the speech modality. We hypothesize that speech signals are more cultural and linguistic in nature, where the speech modality of COVID-19 Sounds data set, collected primarily in UK, differs drastically from the Coswara data set collected in the Indian sub-continent. On the other hand, the breathing signals are more universal and less linguistically and geographically influenced.

## 5. Discussion

Figure 4 gives a better insight on how the center frequency of the filters are distributed for the different learning approaches explored in this work. We can observe that, for breathing, the distribution of learned filters in the low frequency regions (up to  $\sim 2$  kHz) is almost similar to that of the mel-scale filterbank. However, they take on a different distribution in the mid-range and high frequency range. Compared to the mel filterbank, the model puts fewer filters in the frequency regions that are not very informative for classifier and puts more filters in the regions that require more frequency resolutions. The filter profile of the self-supervised model is similar to the supervised models in the low frequency regions. However, unlike supervised models, it places more filters in the mid-frequency range for speech data. It is interesting that the filters of the self-supervised model that are fine-tuned on the target data, take on an almost similar distribution to that of the supervised model. For the speech, the model puts more filters than the mel filterbank in the low frequency region (0-4 kHz). Then, it becomes more sparse in the higher frequency regions.

## 6. Conclusions

In this paper, we have presented an approach for representation learning based on a parameterized convolutional neural network layer with a relevance weighting. The approach uses cosine modulated Gaussian kernels to learn the sub-band decomposition of the audio signal. Further, we have explored the transfer learning capabilities of the front-end using a supervised and self-supervised modeling framework. In our experiments on COVID-19 detection task from raw audio, we show that the proposed approach improves over the baseline mel spectrogram representations as well as the other approaches for representation learning. The improvements are also observed for both the modalities of speech and breathing, indicating the usefulness of the modeling paradigm for a broad range of acoustic signals.

## 7. References

- [1] G. Fagherazzi, A. Fischer, M. Ismael, and V. Despotovic, "Voice for health: the use of vocal biomarkers from research to clinical practice," *Digital biomarkers*, vol. 5, no. 1, pp. 78–88, 2021.
- [2] F. Avila, A. H. Poorjam, D. Mittal, C. Dognin, A. Muguli, R. Kumar, S. R. Chetupalli, S. Ganapathy, and M. Singh, "Investigating feature selection and explainability for COVID-19 diagnostics from cough sounds," in *INTERSPEECH*, vol. 6, 2021, pp. 4246–4250.
- [3] A. Holzinger, "From machine learning to explainable AI," in *2018 world symposium on digital intelligence for systems and machines (DISA)*. IEEE, 2018, pp. 55–66.
- [4] D. M. Low, G. Randolph, V. Rao, S. S. Ghosh, and P. C. Song, "Uncovering the important acoustic features for detecting vocal fold paralysis with explainable machine learning," *medRxiv*, 2020.
- [5] R. Gupta, T. Chaspari, J. Kim, N. Kumar, D. Bone, and S. Narayanan, "Pathological speech processing: State-of-the-art, current challenges, and future directions," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 6470–6474.
- [6] A. Ozdas, R. G. Shiavi, S. E. Silverman, M. K. Silverman, and D. M. Wilkes, "Investigation of vocal jitter and glottal flow spectrum as possible cues for depression and near-term suicidal risk," *IEEE transactions on Biomedical engineering*, vol. 51, no. 9, pp. 1530–1540, 2004.
- [7] S. Alghowinem, R. Goecke, M. Wagner, J. Epps, M. Breakpear, and G. Parker, "Detecting depression: a comparison between spontaneous and read speech," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 7547–7551.
- [8] K. Shama, A. Krishna, and N. U. Cholayya, "Study of harmonics-to-noise ratio and critical-band energy spectrum of speech as acoustic indicators of laryngeal and voice pathology," *EURASIP Journal on Advances in Signal Processing*, vol. 2007, pp. 1–9, 2006.
- [9] J. Maslan, X. Leng, C. Rees, D. Blalock, and S. G. Butler, "Maximum phonation time in healthy older adults," *Journal of Voice*, vol. 25, no. 6, pp. 709–713, 2011.
- [10] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [11] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [12] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [13] J. Turian, J. Shier, H. R. Khan, B. Raj, B. W. Schuller, C. J. Steinmetz, C. Malloy, G. Tzanetakis, G. Velarde, K. McNally *et al.*, "Hear 2021: Holistic evaluation of audio representations," *arXiv preprint arXiv:2203.03022*, 2022.
- [14] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [15] M. Freitag, S. Amiriparian, S. Pugachevskiy, N. Cummins, and B. Schuller, "audeep: Unsupervised learning of representations from audio with deep recurrent neural networks," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6340–6344, 2017.
- [16] T. N. Sainath, B. Kingsbury, A.-r. Mohamed, and B. Ramabhadran, "Learning filter banks within a deep neural network framework," in *IEEE workshop on automatic speech recognition and understanding*, 2013, pp. 297–302.
- [17] Y. Hoshen, R. J. Weiss, and K. W. Wilson, "Speech acoustic modeling from raw multichannel waveforms," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2015, pp. 4624–4628.
- [18] N. Zeghidour, O. Teboul, F. d. C. Quitry, and M. Tagliasacchi, "Leaf: A learnable frontend for audio classification," *arXiv preprint arXiv:2101.08596*, 2021.
- [19] P. Agrawal and S. Ganapathy, "Modulation filter learning using deep variational networks for robust speech recognition," *IEEE journal of selected topics in signal processing*, vol. 13, no. 2, pp. 244–253, 2019.
- [20] J. Andén and S. Mallat, "Deep scattering spectrum," *IEEE Transactions on Signal Processing*, vol. 62, no. 16, pp. 4114–4128, 2014.
- [21] M. Ravanelli and Y. Bengio, "Interpretable convolutional filters with sincnet," *arXiv preprint arXiv:1811.09725*, 2018.
- [22] P. Agrawal and S. Ganapathy, "Unsupervised modulation filter learning for noise-robust speech recognition," *The Journal of the Acoustical Society of America*, vol. 142, no. 3, pp. 1686–1692, 2017.
- [23] H. B. Sailor and H. A. Patil, "Filterbank learning using convolutional restricted boltzmann machine for speech recognition," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2016, pp. 5895–5899.
- [24] P. Agrawal and S. Ganapathy, "Unsupervised raw waveform representation learning for ASR," in *INTERSPEECH*, 2019, pp. 3451–3455.
- [25] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," *arXiv preprint arXiv:1904.05862*, 2019.
- [26] P. Agrawal and S. Ganapathy, "Comparison of unsupervised modulation filter learning methods for asr," in *INTERSPEECH*, 2018, pp. 2908–2912.
- [27] S. Pascual, M. Ravanelli, J. Serra, A. Bonafonte, and Y. Bengio, "Learning problem-agnostic speech representations from multiple self-supervised tasks," *arXiv preprint arXiv:1904.03416*, 2019.
- [28] P. Agrawal and S. Ganapathy, "Interpretable representation learning for speech and audio signals based on relevance weighting," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2823–2836, 2020.
- [29] N. K. Sharma, S. R. Chetupalli, D. Bhattacharya, D. Dutta, P. Mote, and S. Ganapathy, "The second dicova challenge: Dataset and performance analysis for covid-19 diagnosis using acoustics," *arXiv preprint arXiv:2110.01177*, 2021.
- [30] T. Xia, D. Spathis, J. Ch, A. Grammenos, J. Han, A. Hasthanasombat, E. Bondareva, T. Dang, A. Floto, P. Cicutta *et al.*, "Covid-19 sounds: A large-scale audio dataset for digital respiratory screening," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [31] N. Sharma, P. Krishnan, R. Kumar, S. Ramoji, S. Chetupalli, P. Ghosh, and S. Ganapathy, "Coswara—a database of breathing, cough, and voice sounds for covid-19 diagnosis. arxiv 2020," *arXiv preprint arXiv:2005.10548*.
- [32] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.