



# What is Learnt by the LEArnable Front-end (LEAF)? Adapting Per-Channel Energy Normalisation (PCEN) to Noisy Conditions

Hanyu Meng, Vidhyasaharan Sethu, Eliathamby Ambikairajah

University of New South Wales, Australia

{hanyu.meng, v.sethu, e.ambikairajah}@unsw.edu.au

## Abstract

There is increasing interest in the use of the LEArnable Front-end (LEAF) in a variety of speech processing systems. However, there is a dearth of analyses of what is actually learnt and the relative importance of training the different components of the front-end. In this paper, we investigate this question on keyword spotting, speech-based emotion recognition and language identification tasks and find that the filters for spectral decomposition and the low pass filter used to estimate spectral energy variations exhibit no learning and the per-channel energy normalisation (PCEN) is the key component that is learnt. Following this, we explore the potential of adapting only the PCEN layer with a small amount of noisy data to enable it to learn appropriate dynamic range compression that better suits the noise conditions. This in turn enables a system trained on clean speech to work more accurately on noisy test data as demonstrated by the experimental results reported in this paper.

**Index Terms:** learnable audio front-end, adaptive front-end, pre-channel energy normalization, speech signal classification

## 1. Introduction

The speech front-end is a crucial component in speech signal classification systems, and has been the focus of research for many decades. The Mel filterbank [1] leads to representations inspired by our understanding of human perception, and is arguably the most widely used front-end across a range of different speech applications. With the advent of deep learning systems, there has been an interest in end-to-end systems that attempt to learn the optimal transformations to extract information from speech waveforms for any target application [2–5]. More recently, focus has shifted to learnable front-ends that constraints the architecture of the front-end but allows the model parameters to be learnt in conjunction with the back-end [6–12]. They have been employed in a wide range of applications, including but not limited to speaker verification [7], spoofing detection [9], and emotion recognition [8].

A number of general-purposed learnable front-ends such as Time-Domain Filterbank (TD-Fbank) [12], SincNet [10], CGCNN [11] and LEArnable Front-end (LEAF) [6] have been developed recently. Among these, LEAF stands out as having fewer parameters and higher reported accuracy in tasks such as audio events classification. Its universal representation generated from raw speech signals has also made it applicable in other audio-related tasks, such as medical acoustic signal feature learning [13], analog acoustic recognition [14], bird activity detection [15], speaker verification [16], and limited-vocabulary speech recognition tasks [17]. Despite these successes, little is known about what is exactly learnt by LEAF from speech signals.

Analyses of the LEAF model have suggested that one of its components, Per-Channel Energy Normalization (PCEN) plays an important role in effectively compensating for the impact of environmental noise on speech intelligibility [18–21]. PCEN has also been widely applied in acoustic scene classification [22] and long-distance bioacoustic event detection [23]. However, beyond this, there have been limited analyses and insights in the operation of LEAF. To address this shortcoming, we investigate which components of LEAF learn during model training and to what extent.

In this paper, we demonstrate that only the PCEN layer of the broader LEAF model learns during training and there is no observable change in the characteristics of any of the other components away from their initial values. Following this, we leverage our findings to develop a noise adaptation strategy whereby only the PCEN layer of LEAF is adapted using a small amount of noisy data to enable the LEAF model to be used under noisy conditions.

## 2. What is learnt by LEAF?

### 2.1. LEAF

LEAF is a general-purpose audio front-end designed for audio event classification [6]. It mainly comprises three learnable and one non-learnable elements: Spectral Decomposition, Energy Estimation (non-learnable), Smoothing, and Dynamic Range Compression. As depicted in Figure 1, a frame of speech with  $M$  samples passes through a parallel filterbank of  $N$  Gabor filters [24], which is initialised to be equally spaced along the mel-frequency scale. During training, both the centre frequencies ( $f_i$ ) and bandwidths ( $BW_i$ ) of all filters can be learnt. The filterbank is followed by energy estimation implemented by a sample-wise squaring operation, which in turn is smoothed using a low pass filter. The low pass filtering is implemented as a pooling operation comprising of a Gaussian Low-pass Filters (LPF), an approach that has been shown to be effective with 2D features [25]. The learnable parameter is the standard deviation  $\sigma_i$  (equivalently its bandwidth) of the Gaussian LPF in each frequency channel.

Finally, the per channel energy compression (PCEN) acts as an input energy level-dependent gain controller that computes appropriate gains to enhance or attenuate signals in each frequency channel [26]. The PCEN layer for the  $i^{th}$  frequency channel takes as input the smoothed energy estimates produced by the low pass filter layer,  $E[n, i]$ , and is formulated as:

$$PCEN[n, i] = \left( \frac{E[n, i]}{(M[n, i] + \epsilon)^{\alpha_i}} + \delta_i \right)^{\gamma_i} - \delta_i^{\gamma_i}, \quad (1)$$

where  $M[n, i]$  is the smoothed version of the input representation achieved by a first-order infinite impulse response (IIR)

filter as expressed in Equation 2 with a learnable smoothing factor,  $s_i$ .

$$M[n, i] = s_i E[n, i] + (1 - s_i) M[n - 1, i]. \quad (2)$$

As illustrated in Figure 1, the PCEN layer has four learnable parameters ( $s_i, \alpha_i, \delta_i, \gamma_i$ ) per frequency channel.




Figure 1: An overview of LEAF (the input speech frame contains  $M$  samples).

## 2.2. Analysis of LEAF Parameters: Experimental Setup

To investigate what is learnt by LEAF, we train LEAF on three of the tasks where its performance has been perviously reported: emotion recognition, keyword spotting, and language identification [6]. The tasks were chosen to represent a range of speech processing applications as well as prediction accuracies ranging from quite high to somewhat low (refer Table 1). Following this, we compare the spectral and compression characteristics of the three learnable components of LEAF before and after training.

For these analyses, we replicate the dataset settings reported in [27]. The keyword spotting task is trained on SpeechCommandsV2 dataset [28], emotion recognition experiments were carried out on CREMA-D dataset [29], and we use Voxforge [30] for language identification. All three datasets are sampled at 16kHz. For all three tasks, we utilise EfficientNetB0 [31] as the back-end, which is a lightweight Convolution Neural Network (CNN) based classification network.

Prior to training, the LEAF model was initialised as described in [6]. Specifically, we initialised the 40 Gabor filters using the mel-scale and set the input window size to 401 samples with a hop size of 160 samples, corresponding to 25ms for audio sampled at 16kHz. The Gaussian LPFs were initialised with a standard deviation of 0.4 for all channels. The initial values of the PCEN layer parameters in each frequency channel were set to  $\alpha = 0.96$ ,  $\gamma = 2$ ,  $\delta = 2$ , and  $s = 0.04$ .

For training the models, we utilised the ADAM optimiser with a fixed learning rate of  $10^{-4}$ , and employed mini-batches of size 256. We set the input sequence length for SpeechCommandsV2 and Voxforge to 1 second, and for CREMA-D as 3 seconds, based on the audio files durations in each dataset. To ensure consistent loudness range across different recordings, we rescaled the raw speech signals to a range between 15 dB Sound Pressure Level (SPL) and 30 dB SPL. In the test phase, we adopted the approach in [27], and compute predictions for non-overlapping one-second segments and averaging the logits across the entire recording.

Table 1 presents the test accuracy for all tasks with four different LEAF model settings. The ‘‘Untrained’’ settings indicate that all parts of LEAF were frozen and the initial values were not updated during training. The ‘‘Filter Trained’’ settings indicate that Gabor filters and Gaussian low pass filters were set to be trainable, and their parameters would be updated during model training (but the PCEN layer parameters would not be updated). Conversely, the ‘‘PCEN Trained’’ settings refer to the condition where only the PCEN layer parameters were trained and the Gabor filter and Gaussian low pass filter parameters were kept unchanged from their initial values. Finally, ‘‘Fully Trained’’ settings refer to the standard setting where all parameters are trainable.

From the table, the first interesting observation that stands out is that there is little difference between the various trained and untrained versions of LEAF across all three tasks. It is worth noting that for the emotion recognition task, where the untrained LEAF shows the highest accuracy, we used the data partitioning reported in [27], whereby the data was partitioned by shuffling speaker groups and segmenting the data to ensure speaker independence in each partition. For the other two speech tasks, we used the data partitions as per the original dataset release. These results prompt the question ‘‘What is learnt by the LEAF model?’’. Specifically, through our analyses, we aim to answer the question: Which element of the model has undergone the most significant changes as a result of the training?

Table 1: Classification accuracy of LEAF models (mean  $\pm$  std.dev, over three runs).

	Keyword Spotting	Emotion Recognition	Language Identification
Untrained	94.78 $\pm$ 0.1	<b>44.01<math>\pm</math>8.6</b>	91.73 $\pm$ 0.4
PCEN Trained	95.07 $\pm$ 0.2	39.66 $\pm$ 1.5	89.95 $\pm$ 2.1
Filters Trained	94.62 $\pm$ 0.1	40.11 $\pm$ 1.3	<b>95.13<math>\pm</math>1.4</b>
Fully Trained	<b>95.18 <math>\pm</math> 0.3</b>	41.10 $\pm$ 1.4	91.03 $\pm$ 0.6

## 2.3. Analysis of LEAF Parameters: Results

Extracting the weights from the fully trained and untrained LEAF models, we compute the centre frequencies and bandwidths of the Gabor filters as well as magnitude responses of the Gaussian low pass filters and compare them to each other (trained vs untrained across all three tasks). We also reconstruct the compression function of the PCEN layers and compare them to each other. These comparisons are shown in Figure 2(a)-2(f).

**a) Gabor Filterbank:** Figures 2(a) and 2(b) represent the changes in centre frequencies and bandwidths of each filter of the 40 Gabor filters across three tasks as well as the initial values. No appreciable deviations from the initial values can be observed for any of the trained models in any of the three tasks. These results strongly indicate that the initial Gabor filters may be optimal and learning does not help.

**b) Gaussian Lowpass Filterbank:** In Figure 2(c), we plot the frequency response of all 40 Gaussian lowpass filters for all three tasks. Once again, there appears to be no appreciable deviation from the initial values in the pass band of the low pass filters. This again suggests there may be no benefit to learning these filters.

**c) PCEN Range Compression:** In Figures 2(e) to 2(f), we present input level dependent gain imparted by the PCEN layer

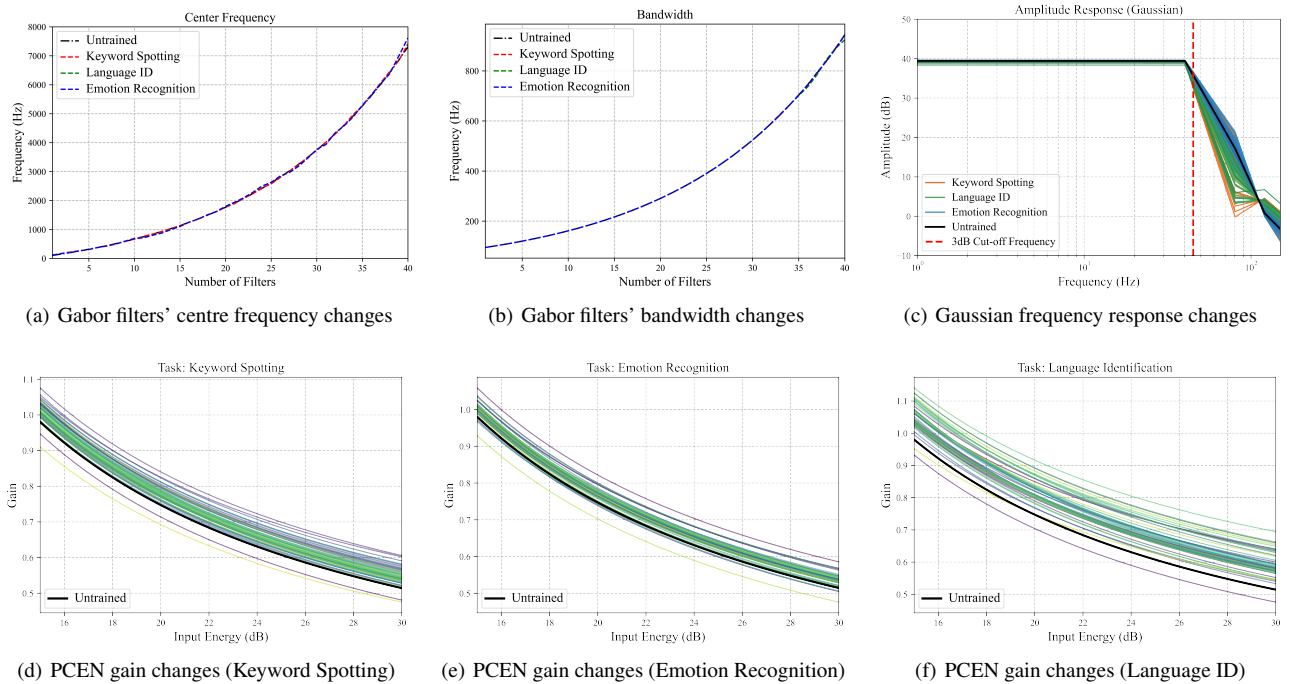


Figure 2: Visualising learnt information of LEAF across three tasks.

across all frequency channels in each of the three tasks. To visualise this, we reconstruct the trained and untrained PCEN functions from Equations 1 and 2 and plot the PCEN gains as a function of input energy  $E[n, i]$ . The curves in Figures 2(e) to 2(f) transition from light green (for the 1<sup>st</sup> channel) to dark purple (for the 40<sup>th</sup> channel), represent the learnt gain characteristics for all 40 channels across the three tasks. It can be seen that the gain curves for the learnt models differ from that of the untrained model across all three tasks.

Taken together these results suggest that of three potentially learnable components of the LEAF model, the Gabor filterbank for spectral decomposition, the Gaussian low pass filter for energy smoothing and the PCEN which offers dynamic range compression, only the PCEN layers appear to be actually learning anything. Consequently, constraining the learning to only this layer of the LEAF model would significantly cut down the number of trainable parameters in the model, leading to more efficient learning. This in turn suggests a noise adaptation scheme involving a small set of noisy speech samples might enable a LEAF model trained on clean speech to be employed in noisy conditions. This hypothesis is explored in the rest of the paper.

### 3. Adapting PCEN to Noisy Environments

We explore the hypothesis that adapting or tuning the PCEN layer in LEAF can enhance the accuracy of speech classification in noisy environments using a limited amount of noisy adaptation data. To test this hypothesis, we compare the performance of a system with a PCEN-adapted LEAF model to that of one trained on clean speech. For reference we also include a model trained entirely on a large amount of noisy data in the comparison<sup>1</sup>.

<sup>1</sup><https://github.com/Hanyu-Meng/Adapting-LEAF>

### 3.1. Experiment Setups

#### 3.1.1. Dataset and Partition

Of the three speech processing tasks, we choose the emotion recognition task to test the proposed PCEN adaptation scheme since it had the lowest accuracy. The CREMA-D dataset for speech emotion recognition consists of 91 speakers and 6 emotions [29], with each speaker having an almost equal number of utterances. When partitioning CREMA-D for this set of experiments, we use different partitions from those used in section 2.2. Specifically, we split the data based on sentences to ensure that each partition contained utterances from each speaker in order to minimise the impact of speaker variability on the results. As illustrated in Figure 3, we used 9 sentences for training, 1 sentence for validation, and the remaining 2 sentences for testing. Consequently, the training set contained 5811 recordings, validation set contained 545, and the test set contained 1086. For adaptation, we selected one sentence from the training set (comprising of 546 recordings) as the adaptation data to which we add different types of noise (white noise and babble noise) at different Signal-to-Noise Ratios (SNRs).

#### 3.1.2. Experimental Setup

For this experiment, we compared four models that had the same structure as the system used in the experiments reported in section 2.2, but with varying training setups as illustrated in Figure 4. These models are:

- **Clean Trained:** Trained on the entire noise-free training set and it serves as a baseline.
- **Noisy Trained:** Trained on a noisy version of the entire training data.
- **Before Adaptation (BA):** Trained on the noise-free training set without including adaptation data. This provides a reference level for performance prior to adaptation (see Figure 3).

- **PCEN Adaptation (PA):** This is the BA model with the PCEN layer adapted using the noisy adaptation data.




Figure 3: Data partitions in the CREMA-D dataset for PCEN adaptation experiments.

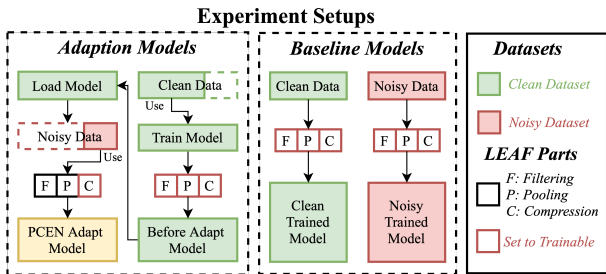


Figure 4: Overview of the PCEN adaptation experimental setup.

### 3.2. Results and Analysis

To verify the hypothesis proposed at the beginning of Section 3, we tested the adaptation of PCEN using both stationary and non-stationary noise.

#### 3.2.1. Gaussian Noise Adaptation

The left pane in Figure 5 shows the classification accuracy for all four models tested by adding different levels of Gaussian noise to the clean test data in order to obtain SNR in the range of 0 dB to 20 dB. The results suggest that training the model with Gaussian noise helps the model learn the pattern of noise and improves its robustness. Further, the models trained on only clean data perform poorly when exposed to Gaussian noise. However, after adapting the PCEN layer with a small amount of noisy data, the impact of noise on accuracy can be somewhat mitigated as can be seen by comparing the performance Before Adaptation (BA model) to that after adaptation (PA model).

#### 3.2.2. Babble Noise Adaptation

We also repeated the above experiment using babble noise instead of white noise. To simulate babble noise, we followed the data augmentation approach used in [32] and used the MUSAN dataset, which contains 60 hours of speech from 12 languages [33]. Specifically, we randomly selected three speech recordings from MUSAN, mixed them together, and added them to the clean signal to simulate babble noise with SNR ranging from 0 dB to 20 dB. The right pane in Figure 5 shows how the accuracy of the four models changes under different levels of babble noise.

The results suggest that training with noisy data is not as effective under babble noise conditions. This might be due to the greater similarity between noise and speech in this case. However, we can observe from the graph that adapting PCEN with babble noise may be quite effective in allowing the model to be used under noisy conditions.

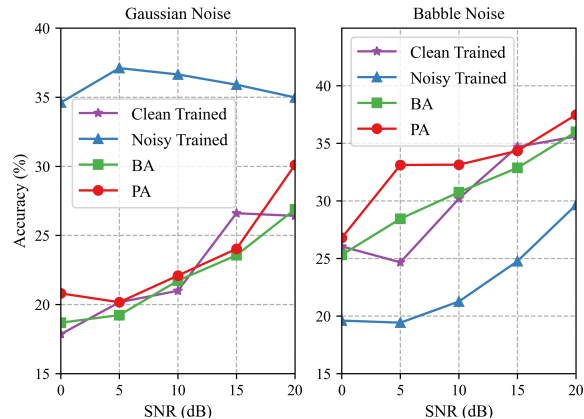


Figure 5: Accuracy of PCEN adaptation experiments on the CREMA-D dataset for different levels of Gaussian noise (left) and babble noise (right).

## 4. Conclusion

In this study, we sought to answer the question ‘What is learnt by LEAF?’, by comparing the extent of change in the characteristics of the different learnable components of the LEAF model from their initial values prior to training. This analysis was repeated on multiple speech processing tasks, and consistently our analyses revealed that only the PCEN layer changes in response to training. The filterbank and low pass filters employed for spectral decomposition and spectral energy smoothing remained unchanged for all three tasks. This suggests that the actual learning in the LEARNable Front-end occurs within a much lower dimensional subspace of the parameter space of the model. Following this, we developed a model adaptation scheme constrained to this subspace (PCEN layer only) using a small amount of noisy training data to adapt a LEAF trained on clean speech to operate more effectively in noisy conditions.

## 5. Acknowledgements

This work was funded by ARC Discovery Grant DP210101228. The authors would also like to thank UNSW, Sydney, Australia for providing PhD scholarship support.

## 6. References

- [1] S. S. Stevens and J. Volkman, “The relation of pitch to frequency: A revised scale,” *The American Journal of Psychology*, vol. 53, no. 3, pp. 329–353, 1940.
- [2] Y. Luo and N. Mesgarani, “Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [3] J. Li, Y. Wu, Y. Gaur, C. Wang, R. Zhao, and S. Liu, “On the comparison of popular end-to-end models for large scale speech recognition,” *arXiv preprint arXiv:2005.14327*, 2020.
- [4] A. Graves and N. Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *International conference on machine learning*. PMLR, 2014, pp. 1764–1772.
- [5] C. Zeng, X. Miao, X. Wang, E. Cooper, and J. Yamagishi, “Joint speaker encoder and neural back-end model for fully end-to-end automatic speaker verification with multiple enrollment utterances,” *arXiv preprint arXiv:2209.00485*, 2022.
- [6] N. Zeghidour, O. Teboul, F. d. C. Quitry, and M. Tagliasac-

- chi, "Leaf: A learnable frontend for audio classification," *arXiv preprint arXiv:2101.08596*, 2021.
- [7] J. Peng, R. Gu, L. Mošner, O. Plchot, L. Burget, and J. Černocký, "Learnable sparse filterbank for speaker verification," *Proc. Interspeech 2022*, pp. 5110–5114, 2022.
- [8] S. Dutta and S. Ganapathy, "Multimodal transformer with learnable frontend and self attention for emotion recognition," in *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6917–6921.
- [9] Q. Fu, Z. Teng, J. White, M. E. Powell, and D. C. Schmidt, "Fas-taudio: A learnable audio front-end for spoof speech detection," in *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 3693–3697.
- [10] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sinnet," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 1021–1028.
- [11] P.-G. Noé, T. Parcollet, and M. Morchid, "Cgcnn: Complex gabor convolutional neural network on raw speech," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7724–7728.
- [12] N. Zeghidour, N. Usunier, I. Kokkinos, T. Schaiz, G. Synnaeve, and E. Dupoux, "Learning filterbanks from raw speech for phone recognition," in *2018 IEEE international conference on acoustics, speech and signal Processing (ICASSP)*. IEEE, 2018, pp. 5509–5513.
- [13] A. M. Poirè, F. Simonetta, and S. Ntalampiras, "Deep feature learning for medical acoustics," in *Artificial Neural Networks and Machine Learning–ICANN 2022: 31st International Conference on Artificial Neural Networks, Bristol, UK, September 6–9, 2022, Proceedings; Part IV*. Springer, 2022, pp. 39–50.
- [14] B. Bergsma, M. Yang, and M. Cernak, "Peaf: Learnable power efficient analog acoustic features for audio recognition," *arXiv preprint arXiv:2110.03715*, 2021.
- [15] M. Anderson and N. Harte, "Learnable acoustic frontends in bird activity detection," *arXiv preprint arXiv:2210.00889*, 2022.
- [16] J. Li, Y. Tian, and T. Lee, "Learnable frequency filters for speech feature extraction in speaker verification," *arXiv preprint arXiv:2206.07563*, 2022.
- [17] H. Peic Tukuljac, B. Ricaud, N. Aspert, and L. Colbois, "Learnable filter-banks for cnn-based audio applications," in *Proceedings of the Northern Lights Deep Learning Workshop 2022*, no. CONF, 2022.
- [18] V. Lostanlen, J. Salamon, M. Cartwright, B. McFee, A. Farnsworth, S. Kelling, and J. P. Bello, "Per-channel energy normalization: Why and how," *IEEE Signal Processing Letters*, vol. 26, no. 1, pp. 39–43, 2018.
- [19] C. Ick and B. McFee, "Sound event detection in urban audio with single and multi-rate pcen," in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 880–884.
- [20] H.-S. Choi, S. Park, J. H. Lee, H. Heo, D. Jeon, and K. Lee, "Real-time denoising and dereverberation with tiny recurrent u-net," in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5789–5793.
- [21] J. Millet and N. Zeghidour, "Learning to detect dysarthria from raw speech," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5831–5835.
- [22] B. Kim, S. Yang, J. Kim, and S. Chang, "Domain generalization on efficient acoustic scene classification using residual normalization," *arXiv preprint arXiv:2111.06531*, 2021.
- [23] V. Lostanlen, K. Palmer, E. Knight, C. Clark, H. Klinck, A. Farnsworth, T. Wong, J. Cramer, and J. P. Bello, "Long-distance detection of bioacoustic events with per-channel energy normalization," *arXiv preprint arXiv:1911.00417*, 2019.
- [24] D. Gabor, "Theory of communication. part 1: The analysis of information," *Journal of the Institution of Electrical Engineers-part III: radio and communication engineering*, vol. 93, no. 26, pp. 429–441, 1946.
- [25] R. Zhang, "Making convolutional networks shift-invariant again," in *International conference on machine learning*. PMLR, 2019, pp. 7324–7334.
- [26] Y. Wang, P. Getreuer, T. Hughes, R. F. Lyon, and R. A. Saurous, "Trainable frontend for robust and far-field keyword spotting," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5670–5674.
- [27] J. Schlüter and G. Gutenbrunner, "Efficientleaf: A faster learnable audio frontend of questionable use," in *2022 30th European Signal Processing Conference (EUSIPCO)*. IEEE, 2022, pp. 205–208.
- [28] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," *arXiv preprint arXiv:1804.03209*, 2018.
- [29] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, "Crema-d: Crowd-sourced emotional multimodal actors dataset," *IEEE transactions on affective computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [30] S. Revay and M. Teschke, "Multiclass language identification using deep learning on spectral images of audio signals," *arXiv preprint arXiv:1905.04348*, 2019.
- [31] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International conference on machine learning*. PMLR, 2019, pp. 6105–6114.
- [32] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [33] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.