

A MULTI-HEAD RELEVANCE WEIGHTING FRAMEWORK FOR LEARNING RAW WAVEFORM AUDIO REPRESENTATIONS

Debottam Dutta*, Purvi Agrawal[§], Sriram Ganapathy*,

*Learning and Extraction of Acoustic Patterns (LEAP) Lab, Indian Institute of Science, Bangalore, 560012

[§]Microsoft India Development Center, Hyderabad, India

ABSTRACT

In this work, we propose a multi-head relevance weighting framework to learn audio representations from raw waveforms. The audio waveform, split into windows of short-duration, are processed with a 1-D convolutional layer of cosine modulated Gaussian filters acting as a learnable filterbank. The key novelty of the proposed framework is the introduction of multi-head relevance on the learnt filterbank representations. Each head of the relevance network is modelled as a separate sub-network. These heads perform representation enhancement by generating weight masks for different parts of the time-frequency representation learnt by the parametric acoustic filterbank layer. The relevance weighted representations are fed to a neural classifier and the whole system is trained jointly for the audio classification objective. Experiments are performed on the DCASE2020 Task 1A challenge as well as the Urban Sound Classification (USC) tasks. In these experiments, the proposed approach yields relative improvements of 10% and 23% respectively for the DCASE2020 and USC datasets over the mel-spectrogram baseline. Also, the analysis of multi-head relevance weights provides insights on the learned representations.

Index Terms— Audio representations, relevance weighting, raw waveform modeling, sound event classification.

1. INTRODUCTION

Recent advances in deep learning have shown substantial improvements in various domains including vision, language and speech [1]. The study of deep networks to learn meaningful representations from data, termed as deep representation learning, has received wide interest in the last few years [2]. Specifically natural language [3] and computer vision tasks [4] have shown promising methods for representation learning. While similar efforts have also been attempted for speech and audio signals [5], knowledge driven representations, like the mel-spectrogram, continue to be the most dominant approach in speech and audio related tasks.

The conventional hand engineered speech features were inspired by multiple psycho-acoustic experiments [6]. Because of the consistent performances shown by mel-filterbank features in speech and speaker recognition tasks, they also find applications in many of the audio-classification tasks. For representation learning from raw waveforms, early works in [7], [5] used raw waveforms and power spectra respectively to learn filterbank parameters. Recently, unsupervised learning of filters were also explored by Agrawal et. al. [8, 9]. Schneider et. al. [10] made use of dilated convolutions to learn pre-trained representations from raw audio in a self-supervised framework. Further, many speech related tasks such as speech separation [11], have successfully adopted learnable front-

ends. Further, learnable front-ends have been explored recently by Zeghidour et. al. [12].

In this work, we propose a multi-head relevance weighting based framework to learn audio representations directly from the raw waveform. This work advances the previous work on speech representation learning [13, 14, 15]. The raw audio signal is windowed into short-time regions and a learnable 1-D convolutional layer is applied. The kernels of this layer are parameterized as cosine modulated Gaussian filters [16]. Following the 1-D convolution and average-pooling operation, a multi-head relevance weighting network is designed which generates self-attention mask for parts of the representation. The relevance weighted representations from the proposed framework (learned 2-D time-frequency representations) are used in the sound classification network. The entire network of filterbank parameters, relevance weights, and event classification network are learned jointly for the task.

Experiments are performed on acoustic scene classification in DCASE2020 challenge Task 1A [17] as well as the urbansound classification [18] task. In these experiments, we show that proposed approach yields representations that improve the performance relatively by up to 23 % in terms of classification accuracy over the baseline log-mel filter representation. We also compare the proposed approach with SincNet representations [19].

2. RELATED WORK

Sainath et. al. [5] proposed a framework to learn filters from the input power spectrum of the signal. In the followup efforts, Hoshen et. al. [20] investigated the learning of CNN filters directly from raw waveform which are initialized with Gammatone filters. In the direction of interpretable filterbank learning, Zeghidour et. al. [21] initialized filterbanks with Gabor wavelets. Recently, the SincNet approach proposed by Ravanelli et. al. [22] considers the convolution kernels as sinc filters and learns only the low and high frequency cut-off. The proposed work in this paper is inspired by Agrawal et. al. [13], where the approach explored cosine modulated Gaussian filterbank. The Gaussian kernel has better time-frequency localization compared to sinc filters. In this work, we also consider a more generalized relevance weighting approach through multi-head sub-networks. The motivation comes from prior works on Mixture of Experts (MoE) models [23, 24].

3. FILTERBANK LEARNING WITH MULTI-HEAD RELEVANCE WEIGHTING

The block schematic of the model is shown in Figure 1. The proposed architecture is motivated by prior work on speech representation learning [13, 16].

3.1. Acoustic filterbank layer

For this layer, we consider the 1-D convolution layer consisting of cosine-modulated Gaussian kernels. The input to this layer are the T frames (windowed segments of audio) of raw audio files where each frame contains S samples. This matrix of size $S \times T$ is passed through a 1-D convolutional layer where the kernels are parameterized by cosine modulated Gaussian filter [13],

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

where $g_i(n)$ is the i -th kernel ($i = 1, 2, \dots, F$) and μ_i is the centre frequency of the i -th kernel. The kernels perform convolution in each frame and generate F feature maps. These feature maps are then squared, average pooled and log transformed. This way, from each frame, a F dimensional feature is obtained. For the T frames, the $F \times T$ shaped representation \mathbf{x} is generated which we call the learned time-frequency (t-f) representation. The learned t-f representation is also ordered in the increasing order of the center frequency μ_i before the relevance weighting.

3.2. Relevance weighting

The multi-head relevance weighting method is proposed to enhance different parts of the learned t-f representation. Each head of the multi-head relevance network is a 2 layer fully connected network with sigmoidal activation in the hidden and output layers. Each head generates a soft mask for different parts of the learned spectrogram. We use a frequency based splitting of the t-f presentation to generate parts of the spectrogram used in the heads. For example, the model with h heads splits \mathbf{x} along frequency axis into h non-overlapping segments, $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(h)}$. Let, $f_i \times T$ be the shape of i -th segment, $\mathbf{x}^{(i)}$ which is fed as an input to i -th head.

The relevance network for each head receives one sample from each time-frequency bin along with a temporal context of $(2c + 1)$ from that sub-band to generate the relevance weight for that sample. Here, the context window size, c is a hyper parameter in our experiments. The output of the relevance weighting network is,

$$W_{k,j}^{(i)} = \sigma(\Omega_2^{(i)}(\sigma(\Omega_1^{(i)}y_{k,j} + b_1^{(i)}) + b_2^{(i)})) \quad (2)$$

where $W_{k,j}^{(i)}$ is the entry in k^{th} row and j^{th} column of the weight mask $W^{(i)}$, $y_{k,j}$ is the $2c + 1$ dimensional input vector $[\mathbf{x}_{k,(j-c)}, \dots, \mathbf{x}_{k,j}, \dots, \mathbf{x}_{k,(j+c)}]^T$, $\Omega_1^{(i)}$, $b_1^{(i)}$ and $\Omega_2^{(i)}$, $b_2^{(i)}$ are weight matrices and biases for the first and second layer respectively. The same network is shared across all the elements of $\mathbf{x}^{(i)}$ to generate $W^{(i)}$ of the same shape $f_i \times T$. The weighted i -th segment, $\mathbf{x}_W^{(i)}$ is obtained by element-wise multiplication of the weight mask $W^{(i)}$ with $\mathbf{x}^{(i)}$. At this stage, we call $\mathbf{x}_W^{(i)}$ the enhanced t-f representation segment for the i -th head. In total, the h enhanced t-f segments will be generated from each of h heads. Then, skip-connection is used to add each part with the corresponding regions of the input representation, $\mathbf{x}^{(i)}$. Finally, all the representation segments are spliced to form the inputs to the audio classification network.

The frequency splitting method introduced in this section is similar to the generation of weights in a smaller dimensional representation space [25] and late fusion of separate frequency paths [26]. The frequency splitting method described here directly splits the time-frequency representation into multiple parts and uses a shared relevance sub-network for all the sub-bands within a split.

The data augmentation methods such as SpecAug [27], random cropping, scaling etc. which are performed over spectrogram,




Figure 1: Proposed multi-head relevance weighting framework for audio representation learning.

are applied to individual heads separately. The augmented features are then fed to the rest of the classifier network for downstream task. We use number of Gaussian filters $F = 80$ and kernel size $k = 705$ (16ms for audio signal sampled at 44.1kHz). The window length $S = 2048$ (46ms) and $S = 1102$ (25ms) are used for DCASE2020 Task1A challenge and USC tasks respectively. The rest of the classification network used is either a fully convolutional network (FCN) for the DCASE2020 dataset or a convolutional long short-term memory (CLSTM) based network for the USC dataset.

4. EXPERIMENTS AND RESULTS

4.1. Acoustic scene classification

All the experiments for acoustic scene classification (ASC) task are performed on DCASE2020 Task 1A development dataset [17]. This dataset contains around 14000 training samples and around 3000

Table 1: Class-wise accuracy(%) of mel, learned features over validation data. Different classes are Airport(AP), Bus(BS), Metro(MT), Metro station(MS), Urban park(PK), Public square(PS), Shopping mall(SM), Street pedestrian(SP), Street traffic(ST) and Tram (TM).

Representations	Accuracy (%)											
	AP	BS	MT	MS	PK	PS	SM	SP	ST	TM	Avg	Avg. held-out
Mel	55.7	84.5	63.6	71.0	82.4	54.8	71.0	47.8	89.5	75.0	69.5	85.2
SincNet [22]	46.9	83.8	70.3	71.0	85.5	55.5	65.3	49.1	85.8	70.9	68.5	86.5
Cos-Gauss [8]	54.3	84.8	77.4	71.4	82.9	62.3	69.4	47.5	87.2	74.3	71.1	88.5
Cos-Gauss + 2head-rel. (Prop)	53.0	86.1	76.8	72.0	84.8	58.9	69.7	56.6	90.6	76.7	72.6	89.6

Table 2: Classifier accuracy (%) in UrbanSound8K database for different sound categories: Air conditioner (AI), Car horn (CA), Children playing (CH), Dog bark (DO), Drilling (DR), Engine idling (EN), Gun shot (GU), Jack hammer (JA), Siren (SI) and Street music (ST).

Representations	Accuracy (%)									
	AI	CA	CH	DO	DR	EN	GU	JA	SI	ST
Mel	32	61	50	65	45	44	54	44	66	61
Cos-Gauss [13]	32	67	60	68	63	55	82	48	71	65
Cos-Gauss + 2head-rel. (Prop)	37	37	62	69	70	64	79	71	60	68

test samples which are recorded over 12 European cities in 10 different acoustic scenes. We use the officially provided train and validation sets for all the experiments and these sets are balanced across all the classes. The audio samples come from 3 real devices A, B and C and 6 simulated devices, s1-s6 of which s4-s6 are only part of the test set. A separate set of 6,000 recordings are also used as a blind held-out set to compare the models (reported in last column of Table 1). Along with the original train set, we perform different augmentations to increase the size of the dataset. The augmentation strategies include (i) Spectrum correction [28] (ii) pitch shift, (iii) speed change, (iv) addition of random noise and (v) mixing of audio signals from the same class. All these strategies are performed at the audio level and they increase the training set size to around 80K samples. Also, the SpecAug approach [27], random-cropping [26] and mix-up [29] are applied on the spectrogram level to increase the robustness of the model learning. The spectrogram level augmentations do not increase the data set size and the random-crop is performed only along time axis.

The baseline model used for the ASC task is the fully convolutional network (FCN) classifier [30]. This FCN model is a VGG-like [31] architecture with 9 convolutional layers with small kernels (around 12M parameters) as described in Hu et. al [30]. For the log-mel representations used in the baseline system, short time processing of audio signals are performed with window lengths of 2048 and hop length of 1024 samples. The windowed signal is passed through a mel filterbank to obtain 80 log-mel energy features. These features are appended with delta, delta-delta features to obtain a final input of dimension $80 \times 423 \times 3$. A scaling operation to convert the values to the range of $[0, 1]$ is performed before feeding to FCN network.

For the representation learning experiments, all the parameters are set according log-mel baseline for fair comparison. In our proposed multi-head network approach, we calculate delta-delta features after the skip-add connection from individual representation segments and then apply SpecAug and random-crop over them. All the experiment setups are implemented in PyTorch [32]. The models are learned with stochastic gradient descent (SGD) with momentum. A cosine annealing warm restart is used as learning rate scheduler to train all the models. The maximum and minimum learning rate used for the scheduler are 0.1 and 10^{-5} respectively.

Table 1 shows the performance of different models. The log-mel baseline model (Mel) provides 69.5 % accuracy on the test data. When the log-mel features are replaced by the learnt cosine modulated Gaussian features (Cos-Gauss), the accuracy is improved to 71.1 %. Further, the 2-head relevance network accuracy improves

the performance over the baseline system with a 3.1 % absolute gain over the mel-baseline (relative improvement of 10 %). Further, on the separate held-out set (last column of Table 1), we observe that the proposed model gives consistent performance gains.

4.2. Urban sound classification

The Urban Sound classification (USC) task [18] contains urban sound events. The UrbanSound8k dataset contains 8732 sound clips sampled at 44.1 kHz, with duration up to 4s. For feature extraction, 25ms window length and 10ms hop length is used. Model evaluation is performed on the official predefined 10-fold cross validation setup. Averaging samples from all the folds we get 7.8K samples for training and around 900 samples for testing. We use the same multi-head relevance weighting setup with a different backend classifier for the USC task. Here, the backend classifier is based on CNN-LSTM network having a stack of 1 convolutional layer (with 40 filters and kernel size (5,5)), a 2-layer LSTM network (with 256 cells) followed by a fully connected layer of size 256.

Table 2 reports the performance of the mel-baseline and the proposed multi-head approach. For the proposed multi-head relevance network, we observe a 5 % absolute improvement in average accuracy over the learned acoustic filterbank features [13] and an absolute 11 % improvement over the baseline using mel-filterbank features. Also, as seen from the class-wise accuracy, the 2-head relevance network is able to achieve improved accuracy in 7 out of 10 classes over the mel baseline system. The car-horn sound had smaller number of training data (412 samples) compared to other dominant classes. In the analysis, we found that the model was substituting this class with street-music class.

4.3. Discussion

Center frequency profile - Figure 2 shows the distributions of centre frequencies of learnt Gaussian filter-bank, mel filter-bank and SincNet [22] for ASC task. The SincNet and mel filter center frequencies are similar except in very high frequencies where the SincNet approach prefers lower center frequencies than the mel-filter bank. The cosine modulated Gaussian filterbank matches the mel curve at very low frequencies, but allocates more filters in high frequency regions. This behaviour is expected as acoustic scenes contain lot of high frequency sources such as traffic sounds, beat sounds and background music. The modified center frequency profile may also explain the performance improvements seen in Table 1 (com-




Figure 2: Distribution of the centre frequencies of mel filterbank (orange), learnt Cos-Gauss filterbank (blue) and learnt SincNet filterbank (green) for DCASE challenge task.




Figure 3: (a) Mel spectrogram for audio file *airport-barcelona-0-0-a.wav* from DCASE dataset, (b) output of acoustic FB layer (learnt t-f representation \mathbf{x} in Figure 1), (c) mask produced for the last 40 frequency bands by 2nd head ($\mathbf{W}^{(2)}$), (d) mask produced for first 40 frequency bands by 1st head ($\mathbf{W}^{(1)}$), and (e) final representation with relevance weighting and scaling.

paring rows 1,2 and 3), where the Cos-Gauss filters improve by an absolute margin of 2.6 % over the SincNet approach.

Time-frequency representations - Figure 3 provides the time-frequency representations from the proposed approach. The proposed approach to filterbank learning emphasizes higher frequencies (more filters on the high frequency region). This attributes to broader blue upper region in Figure 3(b). The two relevance weights (masks) for the learned representation from a 2-head setting are shown in Figure 3(c) and (d). The two masks clearly show the time adaptivity of the weights. The audio file comes from an airport scene and around 3,5,6 and 7 sec regions, the audio contains sound of metal box closing and opening. The mask in Figure 3 (c) asserts more weight on these acoustic events and suppresses the rest. The final weighted representations (Figure 3 (d)) is seen to provide sharp time frequency localization of the events.

Hyper-parameter selection - We experiment with various design aspects and hyper-parameters of our model. Starting with the one relevance head (1H-L) model, from Table 3 it can be seen that context window length $c = 10$ (context length used in relevance weight network) offers the best accuracy. When the number of heads is increased, we see a decrease in accuracy if the data augmentation (SpecAug and random-crop) from the heads are not aligned (NA), i.e., when the time and frequency mask posi-

Table 3: Analysis of DCASE model performance for different number of heads (*H), context window size c , aligned (A) and non-aligned (NA) feature in SpecAug (SA), type of relevance weighting architecture (R.Arch), different ways of network sharing for sub-band splitting (Split) and presence of skip-add connection (skipAdd).

Model	c	SA	R.Arch	Split.	skipAdd	Acc(%)
1H-L	5	-	FC	-	no	71.4
	10	-	FC	-	no	71.5
	20	-	FC	-	no	70.7
2H-L	10	NA	FC	40-40	no	71.2
	10	A	FC	40-40	no	72.1
	10	A	FC	even-odd	no	69.9
	10	A	FC	40-40	yes	72.6
	10	A	Conv	40-40	no	72.0
	10	A	FC	26-27-27	no	69.6
4H-L	10	A	FC	20-20-20-20	no	69.3
80H-L	10	A	FC	No sharing	yes	69.5

tions and cropping of time frames in random-crop are not aligned amongst the heads. Hence, with the aligned version of SpecAug and random-crop, the model gains by an absolute margin of 0.9 % in accuracy. We also experiment with the even-odd splitting (odd frequency bands as input to first head and even frequency bands as input to second head in the relevance network). These experiments show that relevance head splitting using 40-40, where the spectrum of first 40 bands is used in the first relevance head and the next 40 bands used in the second head, provides the best accuracy. We also considered using a 1-D convolution (Conv) layer (with 8 kernels of size 1×3) for the first layer of relevance head sub-network. Finally, results from 3-head and 4-head experiments are shown with roughly equal sub-band splits. A separate relevance head for all the 80 sub-bands is also shown in the Table 3. Here, the relevance network for each sub-band is not shared with other sub-bands. The best performance is achieved with 2-head relevance weighting and skip-add.

Model complexity analysis - For DCASE2020 Task1a dataset, the FCN classifier has around 11.78M parameters. Our proposed front-end with learnable acoustic filterbank and a single relevance sub-network adds only 1.1k parameters to the model (Cosine-Gaussian filterbank has only 80 mean trainable parameters). Thus in terms of total parameters involved, the best performing 2-head relevance net approach adds only 0.02 % extra parameters over the FCN baseline classifier. For the Urban sound classification task, the total number of parameters in mel baseline is 1.18M while in the 2-head proposed approach it increases to 1.19M. This leads to a negligible increase of 0.06 % in the number of parameters.

5. SUMMARY

In this paper, we explore audio representation learning with an acoustic filterbank learning and multi-head relevance weighting. We introduce a novel frequency splitting method to learn and enhance different parts of the time-frequency representation. The proposed multi-head framework generates representations that emphasizes high frequency regions of the audio signal. The multi-head framework is shown to provide significant improvements in both DCASE2020 Task 1A and UrbanSound8k datasets.

6. REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [2] 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.
- [3] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [4] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proc. of ICML*, 2020, pp. 1597–1607.
- [5] T. N. Sainath, B. Kingsbury, A.-r. Mohamed, and B. Ramabhadran, "Learning filter banks within a deep neural network framework," in *Proc. of ASRU*, 2013, pp. 297–302.
- [6] 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. [Online]. Available: <http://www.jstor.org/stable/1417526>
- [7] Z. Tüske, P. Golik, R. Schlüter, and H. Ney, "Acoustic modeling with deep neural networks using raw time signal for LVCSR," in *Proc of Interspeech*, 2014.
- [8] P. Agrawal and S. Ganapathy, "Unsupervised Raw Waveform Representation Learning for ASR," in *Proc. Interspeech 2019*, 2019, pp. 3451–3455.
- [9] —, "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.
- [10] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," *CoRR*, vol. abs/1904.05862, 2019.
- [11] 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.
- [12] N. Zeghidour, O. Teboul, F. de Chaumont Quitry, and M. Tagliasacchi, "Leaf: A learnable frontend for audio classification," *ICLR*, 2021.
- [13] 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.
- [14] —, "Unsupervised raw waveform representation learning for ASR," in *INTERSPEECH*, 2019, pp. 3451–3455.
- [15] —, "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.
- [16] P. Agrawal, "Neural representation learning for speech and audio signals," *PhD Thesis*, 2021.
- [17] A. Mesaros, T. Heittola, and T. Virtanen, "A multi-device dataset for urban acoustic scene classification," in *Proceedings of DCASE2018*, 2018, pp. 9–13.
- [18] J. Salamon, C. Jacoby, and J. P. Bello, "A dataset and taxonomy for urban sound research," in *22nd ACM International Conference on Multimedia (ACM-MM'14)*, Orlando, FL, USA, Nov. 2014, pp. 1041–1044.
- [19] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sincnet," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 1021–1028.
- [20] Y. Hoshen, R. J. Weiss, and K. W. Wilson, "Speech acoustic modeling from raw multichannel waveforms," in *Proc. of ICASSP*, 2015, pp. 4624–4628.
- [21] N. Zeghidour, N. Usunier, I. Kokkinos, T. Schaiz, G. Synnaeve, and E. Dupoux, "Learning filterbanks from raw speech for phone recognition," in *Proc. of ICASSP*, 2018, pp. 5509–5513.
- [22] M. Ravanelli and Y. Bengio, "Interpretable convolutional filters with SincNet," in *Proceedings of Neural Information Processing Systems (NIPS)*, 2018.
- [23] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton, "Adaptive mixtures of local experts," *Neural computation*, vol. 3, no. 1, pp. 79–87, 1991.
- [24] M. I. Jordan and R. A. Jacobs, "Hierarchical mixtures of experts and the em algorithm," *Neural computation*, vol. 6, no. 2, pp. 181–214, 1994.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *CoRR*, vol. abs/1706.03762, 2017.
- [26] M. D. McDonnell and W. Gao, "Acoustic scene classification using deep residual networks with late fusion of separated high and low frequency paths," in *Proc. of ICASSP*, 2020, pp. 141–145.
- [27] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *Proc. of Interspeech*, 2019.
- [28] T. Nguyen, F. Pernkopf, and M. Kosmider, "Acoustic scene classification for mismatched recording devices using heated-up softmax and spectrum correction," in *Proc. of ICASSP*, 2020, pp. 126–130.
- [29] H. Zhang, M. Cissé, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *CoRR*, vol. abs/1710.09412, 2017. [Online]. Available: <http://arxiv.org/abs/1710.09412>
- [30] H. Hu, C. H. Yang, X. Xia, X. Bai, X. Tang, Y. Wang, S. Niu, L. Chai, J. Li, H. Zhu, F. Bao, Y. Zhao, S. M. Siniscalchi, Y. Wang, J. Du, and C. Lee, "A two-stage approach to device-robust acoustic scene classification," *CoRR*, vol. abs/2011.01447, 2020. [Online]. Available: <https://arxiv.org/abs/2011.01447>
- [31] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. of ICLR*, 2015.
- [32] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, 2019, pp. 8024–8035.