

TACKLING THE SCORE SHIFT IN CROSS-LINGUAL SPEAKER VERIFICATION BY EXPLOITING LANGUAGE INFORMATION

Jenthe Thienpondt, Brecht Desplanques, Kris Demuynck

IDLab, Department of Electronics and Information Systems, Ghent University - imec, Belgium

ABSTRACT

This paper contains a post-challenge performance analysis on cross-lingual speaker verification of the IDLab submission to the VoxCeleb Speaker Recognition Challenge 2021 (VoxSRC-21). We show that current speaker embedding extractors consistently underestimate speaker similarity in within-speaker cross-lingual trials. Consequently, the typical training and scoring protocols do not put enough emphasis on the compensation of intra-speaker language variability. We propose two techniques to increase cross-lingual speaker verification robustness. First, we enhance our previously proposed Large-Margin Fine-Tuning (LM-FT) training stage with a mini-batch sampling strategy which increases the amount of intra-speaker cross-lingual samples within the mini-batch. Second, we incorporate language information in the logistic regression calibration stage. We integrate quality metrics based on soft and hard decisions of a VoxLingua107 language identification model. The proposed techniques result in a 11.7% relative improvement over the baseline model on the VoxSRC-21 test set and contributed to our third place finish in the corresponding challenge.

Index Terms— speaker recognition, speaker verification, language identification, score calibration, cross-lingual trials

1. INTRODUCTION

The goal of speaker verification is to determine if two utterances are uttered by the same person. Currently, typical speaker verification systems use low-dimensional speaker embeddings extracted from speaker identification models based on Time Delay Neural Networks (TDNNs) [1, 2, 3] or ResNet [4, 5, 6] architectures. The advent of margin- and angular-based loss functions such as Additive Margin (AM) [7] and Additive Angular Margin (AAM) [8] softmax enables the use of cosine similarity between embeddings to score speaker similarity. These neural network based speaker identification models are trained on large datasets of labelled speech utterances to create robust speaker embeddings. A popular dataset is the development part of the VoxCeleb2 [9] corpus, which contains over 1 million utterances from 5994 speakers.

Speaker verification systems should be robust against cross-lingual trial conditions and discriminate between speakers independently of the language spoken. However, spoken language or dialect could be modelled as a speaker characterizing feature by the neural network. As a result, speaker verification systems are prone to underestimating the speaker similarity in positive (within-speaker) cross-lingual trials. This effect is enhanced by the domination of speakers from the Anglosphere and limited intra-speaker linguistic variability in current popular speaker identification datasets.

The VoxCeleb Speaker Recognition Challenge 2021 (VoxSRC-21) [10] aims to provide a challenging speaker verification test set with an emphasis on cross-lingual trials. The competition rules allow

to incorporate information from a pre-trained language classification model to improve robustness against cross-lingual conditions. In this paper we analyze and further develop the cross-linguality compensation techniques we used in our VoxSRC-21 track 1 submission [11].

We propose two enhancements to increase intra-speaker cross-lingual robustness. Both techniques exploit information from a language classification model. First, we propose a cross-lingual fine-tuning stage to make the speaker embedding extractor more robust against varying phonetic content. Second, we introduce and analyse the addition of language information in our previously proposed quality-aware score calibration stage [12].

The paper is organized as follows: Section 2 describes the baseline speaker verification system. Section 3 and 4 outline our proposed cross-lingual fine-tuning stage and language-aware calibration system, respectively. Section 5 describes the experimental setup we use to validate our proposed enhancements. Subsequently, Section 6 will discuss the results of the experiments. Finally, Section 7 will give some concluding remarks.

2. BASELINE SYSTEM

We choose the best performing single system from our final submission on the VoxSRC-21 validation set [11] as our baseline. The architecture of this fwSE-ResNet model is inspired by [5] and incorporates frequency-wise Squeeze-Excitation (fwSE) and frequency positional encodings [6]. The topology is defined in Table 1. Standard ResNet models are based on 2D convolutions, resulting in frequency- and time-equivariance properties of the model. However, speaker-specific speech patterns are expected to be different across lower and higher frequency regions. This makes the addition of frequency positional encodings in the network beneficial. The ResNet architecture is further enhanced to process speech by modifying the Squeeze-Excitation module to rescale activations frequency-wise instead of using the standard channel-wise rescaling. More information can be found in [6].

Our baseline speaker verification system is fine-tuned using the Large-Margin Fine-Tuning (LM-FT) protocol [12]. This secondary training stage increases the margin penalty of the AAM-softmax criterion to enforce greater inter-speaker distances and decrease the intra-speaker variability of the embeddings. The increased training difficulty caused by the higher margin configuration is compensated by taking longer fixed-length crops of the training utterances during fine-tuning.

Finally, the speaker verification trial scores are calibrated by the quality-aware score calibration backend described in [12]. This calibration stage converts the raw trial scores to proper log-likelihood ratios. This post-processing step also increases the speaker discriminatory ability of the system by compensating for varying quality conditions of the recordings in the trials.

Layer name	Structure	Output $C \times F \times T$
log Mel-FBE	-	$1 \times 80 \times T$
Conv2D	3×3 , stride=1	$128 \times 80 \times T$
ResBlock1	$\begin{bmatrix} \text{pos. enc., } 80 \\ 3 \times 3, 128 \\ 3 \times 3, 128 \\ \text{fwSE, [128, 80]} \end{bmatrix} \times 12, \text{ stride}=1$	$128 \times 80 \times T$
ResBlock2a	$\begin{bmatrix} \text{pos. enc., } 80 \\ 3 \times 3, 128 \\ 3 \times 3, 128 \\ \text{fwSE, [128, 40]} \\ 1 \times 1, 128 \end{bmatrix} \times 1, \text{ stride}=2$	$128 \times 40 \times T/2$
ResBlock2b	$\begin{bmatrix} \text{pos. enc., } 40 \\ 3 \times 3, 128 \\ 3 \times 3, 128 \\ \text{fwSE, [128, 40]} \end{bmatrix} \times 15, \text{ stride}=1$	$128 \times 40 \times T/2$
ResBlock3a	$\begin{bmatrix} \text{pos. enc., } 40 \\ 3 \times 3, 256 \\ 3 \times 3, 256 \\ \text{fwSE, [128, 20]} \\ 1 \times 1, 256 \end{bmatrix} \times 1, \text{ stride}=2$	$256 \times 20 \times T/4$
ResBlock3b	$\begin{bmatrix} \text{pos. enc., } 20 \\ 3 \times 3, 256 \\ 3 \times 3, 256 \\ \text{fwSE, [128, 20]} \end{bmatrix} \times 11, \text{ stride}=1$	$256 \times 20 \times T/4$
ResBlock4a	$\begin{bmatrix} \text{pos. enc., } 20 \\ 3 \times 3, 256 \\ 3 \times 3, 256 \\ \text{fwSE, [128, 10]} \\ 1 \times 1, 256 \end{bmatrix} \times 1, \text{ stride}=2$	$256 \times 10 \times T/8$
ResBlock4b	$\begin{bmatrix} \text{pos. enc., } 10 \\ 3 \times 3, 256 \\ 3 \times 3, 256 \\ \text{fwSE, [128, 10]} \end{bmatrix} \times 2, \text{ stride}=1$	$256 \times 10 \times T/8$
Flatten (C, F)	-	$2560 \times T/8$
CAS pooling	-	5120
Linear (emb.)	-	256
AAM-softmax	-	#Speakers

Table 1. The fwSE-ResNet architecture based on [5] with frequency-wise Squeeze-Excitation and frequency positional encodings [6]. C , F and T are the channel, frequency and time dimensions, respectively. The pooling is realized by Channel-dependent Attentive Statistics (CAS) [3]. The 1×1 convolutions are used in the residual connections to match dimensions of the activation maps.

3. CROSS-LINGUAL FINE-TUNING

We want the speaker embeddings to be invariant to varying phonetic content and variation in spoken language. However, most speakers in the dataset will have a limited amount of spoken language variability. Subsequently, the model will likely interpret the spoken language or dialect as a speaker characterizing feature. We argue this could make the model underestimate speaker similarity in cross-lingual trials.

To mitigate this, we propose a cross-lingual fine-tuning stage. In this training stage we increase the intra-speaker language variability on the mini-batch level. Instead of sampling utterances randomly, we iteratively construct cross-lingual mini-batches.

We combine this strategy with LM-FT and replace the hard sampling algorithm of LM-FT with cross-lingual sampling. First, the spoken language of an utterance is estimated using a language classification model. This enables the selection of cross-lingual utterance pairs. Subsequently, mini-batches are constructed by randomly iterating over all N training speakers in our dataset. During an iteration, each mini-batch contains samples from S speakers with each U cross-lingual utterances. The utterances are selected in cross-lingual pairs, independent of the language of previously selected utterances. This alleviates the need to have a large amount of mutually cross-lingual utterances for each training speaker in case $U > 2$. We resort to random sampling when a speaker does not have any cross-lingual utterance pairs available. A single iteration continues until all training speakers are processed, after which the procedure is repeated.

4. LANGUAGE CALIBRATION

Score calibration backends in speaker verification systems convert speaker similarity scores to well-calibrated log-likelihood-ratios [13]. Calibration based on logistic regression has recently proved to improve speaker verification performance for DNN-based systems by including Quality Metric Functions (QMFs) [12, 14, 15].

Quality-aware score calibration [12] learns a mapping l from the speaker similarity score s to produce a calibrated log-likelihood-ratio $l(s)$. The mapping is defined as $l(s) = w_s s + \mathbf{w}_q^T \mathbf{q} + b$ with w_s and \mathbf{w}_q being learnable weights for the trial score s and the quality features \mathbf{q} , respectively. Since this mapping is non-monotonic, it can improve metrics with a fixed decision threshold such as EER and MinDCF.

Figure 1 shows a histogram of s-normalized [16, 17] trial scores of the baseline fwSE-ResNet system on the VoxSRC-21 validation set. Cross-lingual trials are defined according to language labels provided by the challenge organizers. The figure clearly shows that the baseline system underestimates speaker similarity under cross-lingual trial conditions. We propose to add language features based on hard or soft output decisions of a language classification model to allow the calibration backend to compensate for the score shift induced by cross-linguality. A range of potential language features are discussed in the subsections below.

4.1. Binary cross-linguality indicator

We can use the classification output of the language classifier to determine the most probable spoken language of an utterance. Subsequently, we construct a binary feature indicating if the predicted language of the enrollment and the test side of a speaker verification trial are the same or not.

4.2. Similarity of predicted language class probabilities

A cross-lingual binary feature has some limitations. First, the language classification model is prone to errors, especially on languages with a limited amount of training data [18]. In addition, a binary feature does not express any uncertainty on the language estimation of the model. Second, it only provides information with regards to the predicted spoken language, neglecting potential information about the likeliness of the utterance to other languages.

To mitigate these issues, we construct a language feature using the output probabilities E and T of the language classification model of the enrollment and test side utterance of the speaker verification trial, respectively. In case of AAM-softmax trained models, we obtain the probabilities by scaling the output cosine distances of the




Fig. 1. Histogram of the s-normalized fwSE-ResNet trial scores on the VoxSRC-21 validation set.

language classifier by the proper AAM scale factor, followed by a softmax operation.

We want our language calibration feature to be side-independent of the trial, making most divergence-based metrics of the output probabilities unsuitable. We propose the Jensen-Shannon distance between both language classification probabilities as a calibration feature as it obeys the symmetry requirement. The Jensen-Shannon distance can be regarded as a symmetrical and smoother version of the Kullback-Leibler [19] divergence. Given both language classification output distributions, the Jensen-Shannon distance is defined as:

$$JS(E||T) = \sqrt{\frac{D_{KL}(E||M) + D_{KL}(T||M)}{2}} \quad (1)$$

with M equal to $\frac{E+T}{2}$ and D_{KL} indicating the Kullback-Leibler divergence.

4.3. Similarity of language embeddings

Language features based on the classification probabilities directly rely on the confusion of the classifier between the language classes to model similarities between languages. This could negatively impact the ability to express intra-language variability (e.g. dialects of the same language) and language information from unseen classes. However, the final linear layers of the language classification neural network potentially contain a more general and expressive representation of the spoken language. When the language classification model is trained using an angular-based loss function, such as the AAM-softmax, low-dimensional language embeddings can be extracted from the final linear projection layer. The spoken language of the utterances can be directly compared using the cosine distance between the extracted language embeddings. Scoring language embeddings should also generalize better when encountering new languages not seen during the training of the language classifier. Subsequently, we propose the cosine distance of the language embeddings of the enrollment and test side of the trial as a calibration feature.

5. EXPERIMENTAL SETUP

To analyse the performance impact of the proposed cross-lingual fine-tuning stage and the integration of the language calibration fea-

tures, we apply our proposed enhancements on the baseline speaker verification system described in Section 2.

5.1. Training configuration

The baseline speaker embedding extractor is trained on the development part of VoxCeleb2. During training, we take random crops of two seconds of each utterance and apply a random augmentation using the MUSAN corpus [20] (babble, music, noise) and the RIR [21] dataset (reverb) to prevent overfitting. The input features consist of 80-dimensional log Mel-filterbank energies (Mel-FBE) with a window length of 25 ms and a frame shift of 10 ms. To further enhance robustness, we apply SpecAugment [22] which randomly masks 0 to 10 frequency bands and 0 to 5 frames in the time-domain. Subsequently, all filterbank energies are mean normalized per utterance. A mini-batch size of 128 is used during training.

The model is trained using the Adam optimizer [23] with a cyclical learning rate [24] using the *triangular2* policy with the minimum and maximum learning rate varying between $1e-8$ and $1e-3$, respectively. The cycle length is set to 130k. A weight decay of $2e-5$ is used to regularize the model during training. The system is trained for one cycle with the AAM-softmax loss function using a margin and scale value of 0.2 and 30, respectively.

5.2. Cross-lingual large-margin fine-tuning

After the initial training phase, we apply cross-lingual LM-FT (CL LM-FT) on the model to create more discriminative speaker embeddings. In this stage, the crop size is extended to four seconds with a simultaneous AAM-softmax margin increase to 0.4. We use these settings as opposed to the originally proposed configuration in [12] for computational reasons. Additionally, we change the random sampling of training utterances to the cross-lingual sampling strategy described in Section 3. We keep the initial batch size of 128 and vary the ratio of speakers S and cross-lingual utterances U . We do not change the augmentation strategy.

5.3. Quality and language aware calibration

After the fine-tuning stage, speaker verification trial scores are normalized using adaptive s-normalization [16, 17] with an imposter cohort size of 400 speakers. Subsequently, we apply quality-aware score calibration using the log duration QMF [11].

We apply our proposed language calibration by adding the language features from Section 4 to the calibration backend. We evaluate three types of language features based on the similarity between either output language predictions, output language probabilities or language embeddings. The language information is extracted from each utterance by an ECAPA-TDNN [3] language classifier¹ pre-trained on VoxLingual07 [18] using the AAM-softmax loss.

The calibration backend is trained on a custom VoxCeleb2 subset with half of the utterances cropped between 2 and 4 seconds. We initially select 100k trials and balance the amount of positive and negative trials. Half of the trials are cross-lingual. We discard 20% of both positive and negative trials with the least and greatest cosine distance between the trial language embeddings, respectively. We apply this selection to compensate for overfitting induced by the fact that VoxCeleb2 is also the training dataset of the speaker embedding extractor. We only generate within-gender trials and did not include positive trials with utterances originating from the same video.

¹<https://huggingface.co/TalTechNLP/voxlingual07-epaca-tdnn>

System Configuration	Cross-lingual		Standard Benchmarks					
	VoxSRC-21 Val		VoxCeleb1-O		VoxCeleb1-E		VoxCeleb1-H	
	EER(%)	MinDCF	EER(%)	MinDCF	EER(%)	MinDCF	EER(%)	MinDCF
fwSE-ResNet	2.82	0.1538	0.64	0.0489	0.84	0.0925	1.51	0.1471
fwSE-ResNet + LM-FT	2.41	0.1343	0.55	0.0383	0.76	0.0824	1.35	0.1300
fwSE-ResNet + CL LM-FT	2.25	0.1234	0.58	0.0375	0.74	0.0800	1.30	0.1228
+ log duration QMF	2.11	0.1143	0.50	0.0377	0.71	0.0777	1.26	0.1204
++ binary QMF (4.1)	1.84	0.1038	0.58	0.0639	0.78	0.0843	1.42	0.1436
++ Jensen-Shannon QMF (4.2)	1.67	0.0899	0.59	0.0586	0.77	0.0837	1.38	0.1366
++ cosine distance QMF (4.3)	1.63	0.0827	0.55	0.0539	0.74	0.0794	1.30	0.1274

Table 2. Analysis of cross-lingual fine-tuning and calibration with language information of the fwSE-ResNet system.

5.4. Evaluation protocol

We evaluate the baseline system and proposed enhancements on the VoxCeleb test sets [25, 9] and report the EER and MinDCF metric using a P_{target} value of 10^{-2} with C_{FA} and C_{Miss} equal to 1. To analyse the proposed techniques on challenging cross-lingual data, we also evaluate the systems on the VoxSRC-21 validation and test set using the challenge MinDCF metric with a P_{target} value of 0.05.

6. RESULTS

Table 3 shows the performance impact of the proposed cross-lingual fine-tuning on the VoxSRC-21 validation set. The ratio of speakers and cross-lingual utterances (S/U) within a mini-batch during cross-lingual sampling is indicated between brackets. Standard LM-FT with random samples results in a relative performance improvement over the baseline model of 14.5% and 12.7% in EER and MinDCF, respectively. The most optimal cross-lingual sampling strategy uses 64 speakers with each two cross-lingual utterances in the mini-batch and improves performance further with a relative improvement of 6.6% in EER and 8.1% in MinDCF. As shown in Table 3, selecting more than one cross-lingual utterance pair per speaker on the mini-batch level is less effective. This is probably caused by the fact that we keep the mini-batch size constant due to computational limitations. Therefore, we have to reduce S when we increase U . Moreover, the amount of intra-speaker cross-lingual utterances is limited in the training set and setting $U > 2$ might not be optimal for every training speaker.

Method	Sampling	EER(%)	MinDCF
baseline	random	2.82	0.1538
LM-FT	random	2.41	0.1343
LM-FT	cross-lingual (16/8)	2.32	0.1241
LM-FT	cross-lingual (32/4)	2.26	0.1237
LM-FT	cross-lingual (64/2)	2.25	0.1234

Table 3. Evaluation of different configurations of cross-lingual fine-tuning on the VoxSRC-21 validation set.

In Table 2 we analyse the impact of the proposed language features in the calibration stage. In most cases, incorporating language features in the calibration backend results in a minimal performance degradation on the standard VoxCeleb test sets. However, the cali-

bration set was composed to accommodate the cross-lingual characteristics of the VoxSRC-21 validation set, as opposed to the standard VoxCeleb benchmarks. Conversely, we see a significant reduction of the cross-lingual score shift on the VoxSRC-21 validation set. Both the probability- and embedding-based features outperform the binary cross-lingual measure, showing that the system can effectively exploit the additional language information in the soft decision features. The cosine distance between the language embeddings performs the best with a relative improvement of 22.6% and 27.6% of the EER and MinDCF metric on the VoxSRC-21 validation set, respectively. These results indicate that the addition of language calibration features is currently to be considered as a performance trade-off that should be acceptable in most use-cases.

Systems	EER(%)	MinDCF
baseline + LM-FT + QMF	2.78	0.1690
baseline + CL LM-FT + lang QMF	2.72	0.1492

Table 4. Performance analysis of the proposed cross-lingual fine-tuning and language calibration on the VoxSRC-21 test set.

Finally, we evaluate the cross-lingual fine-tuning and language calibration performance impact on the VoxSRC-21 test set. Table 4 compares the LM-FT strategy with random sampling and quality-aware score calibration without language features against LM-FT with cross-lingual sampling and score calibration with language embeddings. Incorporating language information results in a relative improvement on the VoxSRC-21 test set of 11.7% on the MinDCF challenge metric. This improvement is less significant than the observed performance increases on the VoxSRC-21 validation set. We suspect this is mainly due to the significantly smaller crops (< 4 seconds) in the test set, which could deteriorate the language information extracted by the VoxLingua language classifier.

7. CONCLUSION

We proposed two enhancements in speaker verification to increase the robustness against cross-lingual trials. First, we introduced cross-lingual data sampling during fine-tuning of the embedding extractor. Second, we incorporated language information in the calibration backend to compensate for score shifts induced by cross-lingual conditions. By combining both strategies, we improved the baseline model relatively with 11.7% on the main MinDCF metric on the challenging cross-lingual VoxSRC-21 test set.

8. REFERENCES

- [1] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-vectors: Robust DNN embeddings for speaker recognition,” in *Proc. ICASSP*, 2018, pp. 5329–5333.
- [2] David Snyder, Daniel Garcia-Romero, Gregory Sell, Alan McCree, Daniel Povey, and Sanjeev Khudanpur, “Speaker recognition for multi-speaker conversations using x-vectors,” in *Proc. ICASSP*, 2019, pp. 5796–5800.
- [3] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck, “ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification,” in *Proc. Interspeech*, 2020, pp. 3830–3834.
- [4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *IEEE/CVF CVPR*, 2016, pp. 770–778.
- [5] D. Garcia-Romero, Greg Sell, and A. McCree, “Magneto: X-vector magnitude estimation network plus offset for improved speaker recognition,” in *Proc. Odyssey 2020*, 2020, pp. 1–8.
- [6] Jenthe Thienpondt, Brecht Desplanques, and Kris Demuynck, “Integrating Frequency Translational Invariance in TDNNs and Frequency Positional Information in 2D ResNets to Enhance Speaker Verification,” in *Proc. Interspeech 2021*, 2021, pp. 2302–2306.
- [7] Feng Wang, Jian Cheng, Weiyang Liu, and Haijun Liu, “Additive margin softmax for face verification,” *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, 2018.
- [8] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, “ArcFace: Additive angular margin loss for deep face recognition,” in *2019 IEEE/CVF CVPR*, 2019, pp. 4685–4694.
- [9] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman, “VoxCeleb2: Deep speaker recognition,” in *Proc. Interspeech*, 2018, pp. 1086–1090.
- [10] *The VoxCeleb Speaker Recognition Challenge 2021 (VoxSRC-21)*, 2021 (accessed October 6, 2021), <http://www.robots.ox.ac.uk/~vgg/data/voxceleb/competition2021.html>.
- [11] Jenthe Thienpondt, Brecht Desplanques, and Kris Demuynck, “The IDLAB VoxCeleb Speaker Recognition Challenge 2021 system description,” 2021.
- [12] Jenthe Thienpondt, Brecht Desplanques, and Kris Demuynck, “The IDLab VoxSRC-20 submission: Large margin fine-tuning and quality-aware score calibration in DNN based speaker verification,” in *Proc. ICASSP*, 2021.
- [13] Niko Brümmer and Edward De Villiers, “The bosaris toolkit: Theory, algorithms and code for surviving the new dcf,” *arXiv preprint arXiv:1304.2865*, 2013.
- [14] Alexander Alenin, Anton Okhotnikov, Rostislav Makarov, Nikita Torgashov, Ilya Shigabev, and Konstantin Simonchik, “The ID R&D System Description for Short-Duration Speaker Verification Challenge 2021,” in *Proc. Interspeech 2021*, 2021, pp. 2297–2301.
- [15] Miao Zhao, Yufeng Ma, Min Liu, and Minqiang Xu, “The speakin system for voxceleb speaker recognition challenge 2021,” 2021.
- [16] Pavel Matejka, Ondřej Novotný, Oldřich Plchot, Lukas Burget, Mireia Diez, and Jan Černocký, “Analysis of score normalization in multilingual speaker recognition,” in *Proc. Interspeech*, 2017, pp. 1567–1571.
- [17] Sandro Cumani, Pier Batzu, Daniele Colibro, Claudio Vair, Pietro Laface, and Vasileios Vasilakakis, “Comparison of speaker recognition approaches for real applications,” in *Proc. Interspeech*, 2011, pp. 2365–2368.
- [18] Jörgen Valk and Tanel Alumäe, “VoxLingua107: a dataset for spoken language recognition,” in *Proc. IEEE SLT Workshop*, 2021.
- [19] Solomon Kullback and Richard A Leibler, “On information and sufficiency,” *The annals of mathematical statistics*, vol. 22, no. 1, pp. 79–86, 1951.
- [20] David Snyder, Guoguo Chen, and Daniel Povey, “Musan: A music, speech, and noise corpus,” *arXiv preprint arXiv:1510.08484*, 2015.
- [21] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, “A study on data augmentation of reverberant speech for robust speech recognition,” in *Proc. ICASSP*, 2017, pp. 5220–5224.
- [22] Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proc. Interspeech*, 2019, pp. 2613–2617.
- [23] Diederik Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *Proc. ICLR*, 2015.
- [24] L. N. Smith, “Cyclical learning rates for training neural networks,” in *IEEE WACV*, 2017, pp. 464–472.
- [25] Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, “VoxCeleb: A large-scale speaker identification dataset,” in *Proc. Interspeech*, 2017, pp. 2616–2620.