



Build a SRE Challenge System: Lessons from VoxSRC 2022 and CNSRC 2022

Zhengyang Chen^{1,*}, Bing Han^{1,*}, Xu Xiang², Houjun Huang², Bei Liu¹, Yanmin Qian^{1,†}

¹MoE Key Lab of Artificial Intelligence, AI Institute
X-LANCE Lab, Department of Computer Science and Engineering, Shanghai Jiao Tong University
²AISpeech Ltd, Suzhou, China
{zhengyang.chen, hanbing97, yanminqian}@sjtu.edu.cn

Abstract

Many speaker recognition challenges have been held to assess the speaker verification system in the wild and probe the performance limit. Voxceleb Speaker Recognition Challenge (VoxSRC), based on the voxceleb, is the most popular. Besides, another challenge called CN-Celeb Speaker Recognition Challenge (CNSRC) is also held this year, which is based on the Chinese celebrity multi-genre dataset CN-Celeb. Last year, our team participated in both speaker verification closed tracks in CNSRC 2022 and VoxSRC 2022, and achieved the 1st place and 3rd place respectively. In most system reports, the authors usually only provide a description of their systems but lack an effective analysis of their methods. In this paper, we will outline how to build a strong speaker verification challenge system and give a detailed analysis of each method compared with some other popular technical means.

Index Terms: speaker recognition challenge, VoxSRC, CNSRC, large margin finetuning

1. Introduction

In order to evaluate how well current speaker recognition technology is able to identify speakers in unconstrained or ‘in the wild’ data and explore the performance limitation of the speaker recognition system, many speaker recognition challenges have been held in recent years [1, 2, 3]. Among these challenges, the most prestigious one is the Voxceleb Speaker Recognition Challenge (VoxSRC) [2], which is held once a year since 2019 and based on the very popular public speaker recognition dataset Voxceleb [4, 5]. All the speech segments in the Voxceleb dataset are ‘real world’ utterances that are collected from YouTube for several thousand individuals. Similar to Voxceleb, there is another dataset called CN-Celeb [6, 7], which is downloaded from Chinese social media websites. And the data collectors of CN-Celeb also held the first CN-Celeb Speaker Recognition Challenge (CNSRC) ¹ in 2022.

This year, our team participated in the speaker recognition closed track in both the CNSRC and VoxSRC where only specified data can be used for training in the closed track. During these challenges, we adopted very similar strategies and achieved the 1st place in CNSRC 2022 challenge and 3rd place in VoxSRC 2022 challenge. After comparing the top systems of these challenges in recent years [8, 9, 10], we find that they

*Equal Contribution

†Corresponding Author

This work was supported in part by China NSFC projects under Grants 62122050 and 62071288, and in part by Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0102.

¹<http://www.cnceleb.org/competition>

all have a lot in common. However, the system descriptions of these top systems are usually not detailed enough and lack further analysis of the technologies they used.

In this paper, we will first outline our challenge systems and show how to build a strong speaker verification system that can be used in the challenge. Then, from the perspectives of data augmentation, model backbone, loss functions and back-end scoring, we will give a further analysis of the technologies used in each module to verify their necessity and effectiveness.

2. System Pipeline

The outline of our system is shown in Figure 1. The system can be roughly divided into three parts: data processing, system training, and back-end scoring. Each part in the figure has a great influence on the final system performance. In this section, we will introduce a detailed description of each part.

2.1. Data Processing

Speed perturbation, additive noise, and reverberation augmentation are the most popular data augmentation methods in the challenge [11, 12, 9], and we perform them in an online manner:

1. Randomly sample a speed perturbation ratio from $\{1.0, 1.1, 0.9\}$ and do speed perturbation augmentation (ratio 1.0 means no speed perturbation). The speed perturbation will change the pitch of the audio and we consider the processed audio from a new speaker.
2. Decide whether to do noise augmentation with probability 0.6. If doing noise augmentation, randomly sample a noise type from $\{\text{babble, noise, music, reverberation}\}$, which are from MUSAN [13] and RIR² dataset.

For feature extraction, we use torchaudio [14] toolkit to extract the 80-dimensional Fbank and then do mean-normalization on it.

2.2. System training

2.2.1. Embedding Extractor Backbone

In the challenge, we mainly focus on ResNet-based r-vector as the embedding extractor which is implemented in [15]. In addition, we also tune the channel or depth of ResNet to make it wider or deeper including ResNet101, ResNet152, ResNet221 and ResNet293 [10] for further performance improvements. The pooling function plays a vital role in embedding extractor [16]. We also adopt the multi-query multi-head attention (MQMHA) pooling [17] in part of our systems.

²<https://www.openslr.org/28>

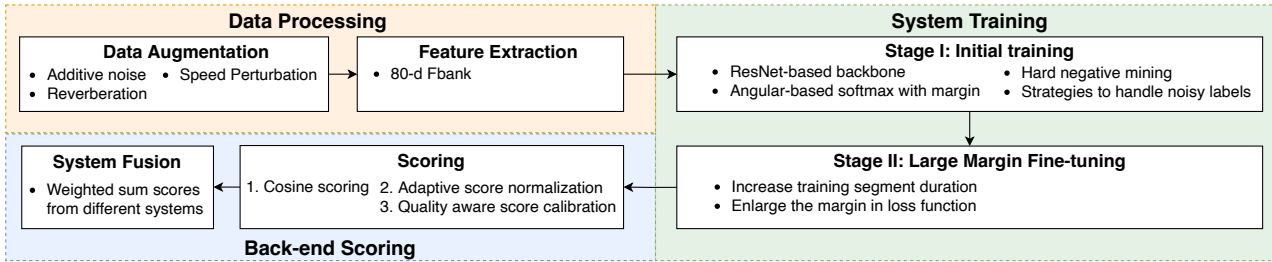


Figure 1: *The outline of our challenge system.*

2.2.2. Loss Function

In our experiment, Additive Angular Margin (AAM) loss [18, 19] is used as the primary training objective to maximize the between-class distance and minimize the within-class distance. Besides, we also involve the Additive Margin (AM) loss [20, 19] in our experiment for comparison. Besides, to alleviate noisy labels influence in the training set, we combine the Sub-center [21] with the AAM for robust training. The evaluation trial focuses more on the hard verification pairs for the speaker recognition challenge. Here, we also involve the Inter-TopK loss [17] to add an extra penalty for k closest centers to the example x_i . Besides, we also propose another hard example mining (HEM) loss which is defined as:

$$L_{HEM} = -\log \frac{e^{s \cos(\theta_{i, y_i} + m)}}{e^{s \cos(\theta_{i, y_i} + m)} + \sum_{j=1, j \neq y_i}^N e^{s \phi(\theta_{i, j})}} \quad (1)$$

where θ_{i, y_i} is the angle between the embedding and target classification center, and $\theta_{i, j}$ is the angle between the embedding and non-target classification center. To do hard sample mining, we detect the sample which has high similarity with the non-target classification center and adds an extra margin m' to $\theta_{i, j}$:

$$\phi(\theta_{i, j}) = \begin{cases} \cos(\theta_{i, j} - m') & \text{If } \cos \theta_{i, j} > \cos(\theta_{i, y_i} + m) \\ \cos \theta_{i, j} & \text{Otherwise.} \end{cases} \quad (2)$$

m' here is set to 0.1 in our experiment.

2.2.3. Training Strategy

The training process in our challenge system is divided into two stages.

Stage I: Initial Training: In this stage, we use 2 seconds training segment. We set the margin and scale in AAM (or AM) to 0.2 and 32.0 respectively, and set the margin and k in Inter-TopK loss to 0.06 and 5 respectively. For sub-center loss, we use 3 sub-centers for each class. In this stage, we iterate the training set for 150 epochs with SGD as the optimizer. During the training process, the learning rate is exponentially decreased from the initial 0.1 to the final $1e-5$.

Stage II: Large Margin Fine-tuning: Large Margin fine-tuning (LM-FT) is first proposed in [22], which aims to minimize the duration mismatch between the training and evaluation stage, and further optimize the within-class and between-class distance by enlarging the margin of loss function. In our experiment, we set the training segment duration to 6s in this stage and increase the margin to 0.5. The speed perturbation augmentation is abandoned here. Like the setup in [9], we disable the Inter-TopK loss in this stage to make the training more stable. The learning rate here is initialized as $1e-4$ and decreases

to $2.5e-5$ finally. Based on the pre-trained model from Stage I, we fine-tune it for only five epochs because the longer training segments make the model easy to overfit in this stage.

2.3. Back-end Scoring

Because the angular-based softmax directly optimizes the speaker embedding in a hyper-sphere space, we use the cosine similarity to score the trials. Then, adaptive score normalization (as-norm) [23] is used to normalize the trial score. We average the embeddings from the same speaker in training set to construct the imposter cohort and set the imposter cohort size to 600. Finally, we use the quality-aware score calibration [22] to introduce some extra information to calibrate the score. The Quality Measuring Function (QMF) attributes used in our experiments include:

- The magnitude of enroll and test embeddings.
- The duration of enroll and test utterances.
- The mean value of imposter cohort for enroll and test utterances.

3. Experimental Setup

3.1. VoxSRC 2022

The Voxceleb2 dev [5] set is used as the training set, which contains 5994 speakers. The evaluation trials in our experiment include three cleaned version trials Vox1-O, Vox1-E and Vox1-H constructed from 1251 speakers in Voxceleb1 [4], and the validation trials from VoxSRC 2021 and VoxSRC 2022. Besides, we construct a trial with 30k pairs from the Voxceleb2 dev set following the strategy proposed in [22] to train the score calibration model. In the experimental analysis part, we use ResNet34 with statistic pooling, AAM loss function and cosine similarity scoring plus as-norm and score calibration as the default setup. We don't involve the large margin fine-tuning strategy without specific notation.

3.2. CNSRC 2022

The dev set of CN-Celeb1 [6] and CN-Celeb2 [7] are used as the training set, which contains 2787 speakers in total. Because there are many short utterances less than 2s in CN-Celeb dataset [7], we first concatenate the short utterances from the same genre and same speaker to make them longer than 5s. It should be noted that we only do this operation on the training set. Then, the training utterance number is reduced from 632,740 to 508,228. We report results on CN-Celeb evaluation trial which is also used as the challenge evaluation set. It should be noted that we abandon the score calibration here because we find that the score calibration can benefit the EER but degrade

the minDCF which is the main evaluation metric of this challenge. Besides, there are multiple utterances for each enrollment speaker in CNSRC 2022. We average all the embeddings belonging to each enrollment speaker to get the final speaker embedding.

4. Results and Analysis

4.1. Data Augmentation Methods Comparison

Table 1: *Voxceleb EER (%) results comparison between different augmentation methods.* N+R: additive noise and reverberation. Perturb (S): audio perturbation by changing audio speed. Perturb (P): audio perturbation by changing audio pitch. Perturb (S+P): audio perturbation by changing both audio speed and pitch.

Aug Type	Vox1-O	Vox1-E	Vox1-H
N + R	0.947	1.098	2.002
N + R + SpecAugment	1.064	1.091	2.004
N + R + Perturb (S)	0.973	1.113	2.033
N + R + Perturb (P)	0.867	0.984	1.781
N + R + Perturb (S + P)	0.861	0.996	1.767
N + R + Perturb (S + P) *	0.803	1.001	1.777

*: applying speed perturbation with ratio {0.8, 0.9, 1.0, 1.1, 1.2}

In this section, we first analyze the effect of different data augmentation methods. The effectiveness of the additive noise (N) and reverberation (R) have been extensively verified [24, 25] and we consider it as the baseline in this experiment. The results are shown in Table 1. Apart from the additive noise and reverberation, we also investigate the effectiveness of SpecAugment [26, 27] and speed perturbation. It should be noted that the speed perturbation used in many challenge systems [12, 9] is just one case of audio perturbation method [11]. They implement it using ‘sox speed’³ command, which will change the speed and pitch (S+P) of audio. However, there is no detailed analysis of whether speed augmentation or pitch augmentation is more useful. Here, we also apply the audio perturbation by only changing audio speed (S) using ‘sox tempo’ command and the audio perturbation by only changing audio pitch (P) using ‘sox speed’ and ‘sox tempo’ command sequentially⁴.

Although some previous works [27] have shown that it is useful to use SpecAugment alone, the results in Table 1 show that combining SpecAugment with additive noise and reverberation cannot obtain further benefits. Besides, the audio perturbation by only changing the audio speed cannot improve the system performance, which demonstrates the gain of speed perturbation (S+P) comes from the audio pitch augmentation. In many challenge systems, they only speed up or slow down the audio speed with ratios of 1.1 and 0.9 respectively. Here, we additionally explore the speed ratios 1.2 and 0.8 to generate more diverse training data. However, we find that the performance from speed perturbation has been saturating and the additional speed perturbation ratio cannot bring more performance improvement. In the following experiments, we will use “N + R + Perturb (S + P)” in Table 1 as the default setup.

³<http://sox.sourceforge.net/>

⁴We can also use ‘sox pitch’ command to only change the audio pitch.

Table 2: *Voxceleb EER (%) results comparison between different acoustic features.*

Acoustic Feature	Vox1-O	Vox1-E	Vox1-H
MFCC-80d	1.292	1.351	2.436
Fbank-40d	0.941	1.147	2.104
Fbank-80d	0.861	0.996	1.767
Fbank-96d	0.931	0.953	1.741
Fbank-80d + Pitch	0.862	0.984	1.783

4.2. Acoustic Feature Comparison

In this section, we analyze the performance of different acoustic features and list the results in Table 2. The MFCC feature is actually the most widely used feature in the speech field. However, for speaker recognition challenge, more researchers turned their attention to Fbank features because the MFCC feature is transformed from the Fbank feature and some information may be lost during this transformation process. In Table 2, all the Fbank features perform better than the MFCC feature, demonstrating this point. Besides, the higher dimensional Fbank feature seems to perform better than the lower dimensional one when we compare the Fbank 40d and 80d. And if we continue to increase to 96d, although the performance will be slightly improved, it will also bring more computation. Finally, we also try to add additional pitch information with Fbank feature but there is no further gain. Considering the trade-off between performance and computation, we will use 80d Fbank in the following experiments.

4.3. Scoring Method

Table 3: *Voxceleb EER (%) results comparison between different scoring methods.* PLDA is trained on Voxceleb2 dev set.

Scoring Method	Vox1-O	Vox1-E	Vox1-H
PLDA	1.633	1.723	2.857
Cosine	1.058	1.147	2.087
+ AS-Norm	0.920	1.048	1.874
++ Score Calibration	0.861	0.996	1.767

In this section, we analyze the effect of different scoring methods and results are shown in Table 3. PLDA used to be a popular scoring method in speaker recognition field. However, with the advent of angular-based softmax [18, 20] which can optimize the speaker embedding in a hyper-sphere space, cosine scoring methods begin to dominate. The results show that the performance of PLDA scoring has lagged far behind cosine scoring. After cosine scoring, we apply as-norm and quality-aware score calibration and these two compensation strategies can make consistent improvements. We will also use the cosine scoring + as-norm + score calibration strategy in the following experiments.

4.4. Loss Function Comparison and Training Strategy

This section compares the results of different loss functions and training strategies. The results are shown in Table 5. First, we compare the AM and AAM loss and find that they both achieve comparable results. Then we choose AAM as our default loss. The Sub-center strategy is adopted here to alleviate the noisy labels in training data and obtains performance gain compared with baseline AAM. To mine the hard samples dur-

Table 4: *Voxceleb and VoxSRC results comparison between different embedding extractor backbones.* ‘-64’ denotes the ResNet with base channel number 64 and our default setup is 32. Large margin fine-tuning is applied for all the systems. The ResNet models with statistic pooling are trained with AAM loss and the ResNet models with MQMHA pooling are trained with AAM + Sub-center + Inter-Topk loss.

Index	Model Type	Pooling	Param #	Vox1-O		Vox1-E		Vox1-H		VoxSRC21-val		VoxSRC22-val	
				DCF _{0.01}	EER	DCF _{0.01}	EER	DCF _{0.01}	EER	DCF _{0.05}	EER	DCF _{0.05}	EER
S1	ECAPA (c=512)	ASP	6.20M	0.0901	0.771	0.1051	0.962	0.1670	1.784	0.1911	3.540	0.1676	2.499
S2	ECAPA (c=1024)	ASP	14.7M	0.0802	0.606	0.0889	0.813	0.1556	1.594	0.1841	3.267	0.1641	2.398
S3	ResNet34	Statistic	6.63M	0.0635	0.654	0.0920	0.824	0.1424	1.456	0.1434	2.574	0.1456	2.085
S4	ResNet101	Statistic	15.9M	0.0404	0.505	0.0666	0.651	0.1101	1.163	0.1162	2.025	0.1154	1.691
S5	ResNet152	Statistic	19.8M	0.0341	0.415	0.0623	0.606	0.1027	1.101	0.1189	2.052	0.1077	1.569
S6	ResNet34	MQMHA	8.61M	0.0471	0.675	0.0843	0.771	0.1353	1.369	0.1417	2.411	0.1408	1.955
S7	ResNet34-c64	MQMHA	27.8M	0.0510	0.638	0.0771	0.756	0.1275	1.353	0.1429	2.604	0.1375	2.001
S8	ResNet101	MQMHA	23.8M	0.0442	0.425	0.0615	0.602	0.1003	1.051	0.1050	1.885	0.1029	1.551
S9	ResNet101-c64	MQMHA	68.7M	0.0335	0.388	0.0575	0.576	0.0964	1.044	0.1124	1.961	0.1034	1.566
S10	ResNet152	MQMHA	27.7M	0.0321	0.378	0.0549	0.552	0.0898	0.980	0.0967	1.765	0.1036	1.457
S11	ResNet221	MQMHA	31.6M	0.0357	0.330	0.0539	0.535	0.0855	0.966	0.1009	1.795	0.0976	1.420
S1-S11	Fusion	-	-	0.0282	0.303	0.0483	0.491	0.0795	0.892	0.0931	1.645	0.0898	1.330

ing training, our proposed HEM loss and Inter-Topk are both investigated here. The results show that the HEM loss and Inter-Topk can both improve the performance significantly, especially on the challenge’s harder trial. In addition, our proposed HEM achieves the best EER on VoxSRC2021-val trial. Finally, we apply the large margin fine-tuning (LM-FT) on the pre-trained model. It can be seen that the performance improvement brought by the LM-FT is huge. We will also apply the LM-FT in the following experiments.

Table 5: *Voxceleb and VoxSRC EER (%) results comparison between different loss functions.* The LM-FT is performed on the model with Inter-Topk loss.

Loss Functions	Vox1-O	Vox1-E	Vox1-H	VoxSRC21-val
AM	0.840	0.987	1.796	3.453
AAM	0.861	0.996	1.767	3.450
+ Sub-center	0.824	0.985	1.733	3.340
++ HEM	0.782	0.970	1.684	3.060
++ Inter-TopK	0.782	0.936	1.658	3.153
LM-FT	0.649	0.797	1.394	2.429

4.5. Embedding Extractor Backbone Comparison

Except for the ResNet-based backbone, we also involve a very popular TDNN-based model, ECAPA-TDNN [28], for comparison. Besides, we also apply different pooling methods for the ResNet-based model. From the results in Table 4, we find that the ResNet-based systems perform much better than the ECAPA-TDNN systems. Although the ResNet34 has similar performance with ECAPA-TDNN (c=1024) following the configuration in the original ECAPA-TDNN paper [28]⁵, the ResNet-based model shows higher performance upper bound after adding more extensive data augmentation and advanced training strategies. Besides, we find that making ResNet deeper can improve the performance a lot but making ResNet wider has little effect on the performance. Further, we add the MQMHA pooling and Sub-center+Inter-TopK loss functions to the ResNet training and achieve further performance improvement. Finally, we fuse all the systems’ scores in Table 4 based on the performance on trial VoxSRC22-val to get the fusion system.

⁵In our experiment, ResNet34 has 2.048% EER on Vox1-H and ECAPA-TDNN (c=1024) has 2.017% EER on Vox1-H following the setup in [28].

4.6. Results for CNSRC

Table 6: *Results on CNSRC 2022.* Without specific notation, the models using statistic pooling and are trained with AAM loss.

Index	Model	Param #	DCF _{0.01}	EER (%)
S1	ResNet34 *	6.63M	0.3958	7.981
S2	ResNet34	6.63M	0.3707	6.590
S3	ResNet34 †	6.63M	0.3593	6.702
S4	ResNet152	19.8M	0.3386	5.762
S5	ResNet221	23.8M	0.3270	5.543
S6	ResNet293	28.6M	0.3202	5.553
S7	DF-ResNet	14.8M	0.3361	6.279
S8	ResNet34 + LM-FT †	6.63M	0.3458	6.319
S9	ResNet34 + LM-FT	6.63M	0.3543	6.221
S10	ResNet152 + LM-FT	19.8M	0.3251	5.452
S11	ResNet221 + LM-FT	23.8M	0.3179	5.284
S12	ResNet293 + LM-FT	28.6M	0.3164	5.227
S13	DF-ResNet + LM-FT	14.8M	0.3185	6.117
S9-S13	Fusion	-	0.2975	4.911
-	SpeakerIn System [29]	-	0.3185	5.953
-	STAP System [30]	-	0.3399	5.728

*: did not apply speed perturbation.

†: using MQMHA pooling and AAM + Sub-center + Inter-TopK loss.

In the previous sections, we have introduced our systems in VoxSRC 2022, which have very competitive performance. Here, similar systems are applied in the CNSRC 2022 and the results are given in Table 6. Besides, we also involved the DF-ResNet [31] in this challenge. From the results, we find that the speed perturbation augmentation, deep ResNet and large margin fine-tuning all play an important role. Finally, we weighted sum the scores of all the LM-FT systems (except the system denoted with †) based on their performance to get the fusion system. Our team also achieved the 1st place in the CNSRC 2022, which further confirms that our system is not only strong but also generalizable to different scenarios.

5. Conclusion

In most system reports, the authors usually only provide a description of their systems but lack an effective analysis of their methods. In this paper, based on our winner systems in CNSRC 2022 and VoxSRC 2022, we outline how to build a strong speaker verification challenge system. Then, from the perspectives of data augmentation, model backbone, loss functions, and back-end scoring, we share some experiences learned from these challenges and give a further analysis of the technologies used to verify their necessity and effectiveness.

6. References

- [1] S. O. Sadjadi, C. Greenberg, E. Singer, L. Mason, and D. Reynolds, "The 2021 nist speaker recognition evaluation," *arXiv preprint arXiv:2204.10242*, 2022.
- [2] A. Brown, J. Huh, J. S. Chung, A. Nagrani, and A. Zisserman, "Voxsrc 2021: The third voxceleb speaker recognition challenge," *arXiv preprint arXiv:2201.04583*, 2022.
- [3] X. Qin, M. Li, H. Bu, W. Rao, R. K. Das, S. Narayanan, and H. Li, "The interspeech 2020 far-field speaker verification challenge," *arXiv preprint arXiv:2005.08046*, 2020.
- [4] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: a large-scale speaker identification dataset," *arXiv preprint arXiv:1706.08612*, 2017.
- [5] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *arXiv preprint arXiv:1806.05622*, 2018.
- [6] Y. Fan, J. Kang, L. Li, K. Li, H. Chen, S. Cheng, P. Zhang, Z. Zhou, Y. Cai, and D. Wang, "CN-Celeb: a challenging chinese speaker recognition dataset," in *Proc. IEEE ICASSP*. IEEE, 2020, pp. 7604–7608.
- [7] L. Li, R. Liu, J. Kang, Y. Fan, H. Cui, Y. Cai, R. Vipperla, T. F. Zheng, and D. Wang, "CN-Celeb: multi-genre speaker recognition," *Speech Communication*, 2022.
- [8] J. Thienpondt, B. Desplanques, and K. Demuynck, "The idlab voxceleb speaker recognition challenge 2020 system description," *arXiv preprint arXiv:2010.12468*, 2020.
- [9] M. Zhao, Y. Ma, M. Liu, and M. Xu, "The speakin system for voxceleb speaker recognition challenge 2021," *arXiv preprint arXiv:2109.01989*, 2021.
- [10] Z. Chen, B. Liu, B. Han, L. Zhang, and Y. Qian, "The sjtu x-lance lab system for cnsrc 2022," *arXiv preprint arXiv:2206.11699*, 2022.
- [11] H. Yamamoto, K. A. Lee, K. Okabe, and T. Koshinaka, "Speaker augmentation and bandwidth extension for deep speaker embedding," in *Interspeech*, 2019, pp. 406–410.
- [12] W. Wang, D. Cai, X. Qin, and M. Li, "The dku-dukeece systems for voxceleb speaker recognition challenge 2020," *arXiv preprint arXiv:2010.12731*, 2020.
- [13] D. Snyder, G. Chen, and D. Povey, "MUSAN: A Music, Speech, and Noise Corpus," 2015, arXiv:1510.08484v1.
- [14] Y.-Y. Yang, M. Hira, Z. Ni, A. Astafurov, C. Chen, C. Puhersch, D. Pollack, D. Genzel, D. Greenberg, E. Z. Yang *et al.*, "Torchaudio: Building blocks for audio and speech processing," in *Proc. IEEE ICASSP*. IEEE, 2022, pp. 6982–6986.
- [15] H. Zeinali, S. Wang, A. Silnova, P. Matějka, and O. Plchot, "But system description to voxceleb speaker recognition challenge 2019," *arXiv preprint arXiv:1910.12592*, 2019.
- [16] S. Wang, Y. Yang, Y. Qian, and K. Yu, "Revisiting the statistics pooling layer in deep speaker embedding learning," in *Proc. ISCSLP*. IEEE, 2021, pp. 1–5.
- [17] M. Zhao, Y. Ma, Y. Ding, Y. Zheng, M. Liu, and M. Xu, "Multi-query multi-head attention pooling and inter-topk penalty for speaker verification," in *Proc. IEEE ICASSP*. IEEE, 2022, pp. 6737–6741.
- [18] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proc. CVPR*, 2019, pp. 4690–4699.
- [19] X. Xiang, S. Wang, H. Huang, Y. Qian, and K. Yu, "Margin matters: Towards more discriminative deep neural network embeddings for speaker recognition," in *Proc. APSIPA ASC*. IEEE, 2019, pp. 1652–1656.
- [20] F. Wang, J. Cheng, W. Liu, and H. Liu, "Additive margin softmax for face verification," *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, 2018.
- [21] J. Deng, J. Guo, T. Liu, M. Gong, and S. Zafeiriou, "Sub-center arcface: Boosting face recognition by large-scale noisy web faces," in *European Conference on Computer Vision*. Springer, 2020, pp. 741–757.
- [22] J. Thienpondt, B. Desplanques, and K. Demuynck, "The idlab voxsrc-20 submission: Large margin fine-tuning and quality-aware score calibration in dnn based speaker verification," in *Proc. IEEE ICASSP*. IEEE, 2021, pp. 5814–5818.
- [23] S. Cumani, P. D. Batsu, D. Colibro, C. Vair, P. Laface, and V. Vasilakakis, "Comparison of speaker recognition approaches for real applications," in *Interspeech*, 2011, pp. 2365–2368.
- [24] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for text-independent speaker verification," in *Interspeech*, 2017, pp. 999–1003.
- [25] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *Proc. IEEE ICASSP*. IEEE, 2018, pp. 5329–5333.
- [26] 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," *arXiv preprint arXiv:1904.08779*, 2019.
- [27] S. Wang, J. Rohdin, O. Plchot, L. Burget, K. Yu, and J. Černocký, "Investigation of specAugment for deep speaker embedding learning," in *Proc. IEEE ICASSP*. IEEE, 2020, pp. 7139–7143.
- [28] B. Desplanques, J. Thienpondt, and K. Demuynck, "Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," *arXiv preprint arXiv:2005.07143*, 2020.
- [29] Y. Zheng, Y. Chen, J. Peng, Y. Zhang, M. Liu, and M. Xu, "The speakin system description for cnsrc2022," *arXiv preprint arXiv:2209.10846*, 2022.
- [30] S. Jiang, J. Chen, Q. Liu, and Y. Qian, "The stap system for cn-celeb speaker recognition challenge 2022," <https://aishell-cnsrc.oss-cn-hangzhou.aliyuncs.com/T106.pdf>, 2022.
- [31] B. Liu, Z. Chen, S. Wang, H. Wang, B. Han, and Y. Qian, "DF-ResNet: Boosting Speaker Verification Performance with Depth-First Design," in *Proc. Interspeech 2022*, 2022, pp. 296–300.