

# The HCCL system for VoxCeleb Speaker Recognition Challenge 2022

Zhenduo Zhao<sup>1,2\*</sup>, Zhuo Li<sup>1,2\*</sup>, Wenchao Wang<sup>1</sup>, Pengyuan Zhang<sup>1,2</sup>

<sup>1</sup>Key Laboratory of Speech Acoustics and Content Understanding, Institute of Acoustics, Chinese Academy of Sciences, Beijing, China

<sup>2</sup>University of Chinese Academy of Sciences, Beijing, China

li\_zhuo@foxmail.com, {zhaozhenduo, wangwenchao, zhangpengyuan}@hccl.ioa.ac.cn

## Abstract

This report describes our submission to track1 and track3 for VoxCeleb Speaker Recognition Challenge 2022(VoxSRC2022). Our best system achieves minDCF 0.1397 and EER 2.414 in track1, minDCF 0.388 and EER 7.030 in track3.

**Index Terms:** speaker recognition, VoxSRC2022, domain adaptation, clustering algorithm, label correction.

## 1. System Description for Track1

### 1.1. Data

**Training Data:** We use VoxCeleb2-dev(vox2dev) [1] as training data, which consists of 1092009 utterances from 5994 speakers. To augment data, we first use the SoX speed function with speeds 0.9 and 1.1 to generate extra twice speakers [2]. In total, there are 17982 speakers and 3276027 utterances. Then, we use MUSAN [3] and RIRs noises [4] to perform online data augmentation. Similar to SpeakIn systems for VoxSRC2021 [5], we used a chain augment pipeline to generate samples:

- MUSAN noise with probability 0.2
- MUSAN music with probability 0.2
- MUSAN speech with probability 0.2
- RIRs noises with probability 0.6

SpeechBrain [6] was used to build the pipeline.

**Developing Data:** We use official validation sets [1, 7–9] to evaluate our models: Vox1-O, Vox1-E, Vox1-H, Vox20-dev, Vox21-dev and Vox22-dev.

**Features:** We extract 80-dimensional log-Mel Filter Banks (Fbank) as input features without any voice activity detection(VAD). The frame length is 25ms and the frame shift is 10ms. Cepstral mean normalization(CMN) is applied. Our implementation was based on TorchAudio [10].

### 1.2. Model Structures

Two main stream of current most popular model structures was used for the challenge: 1D-convolution-based ECAPA-TDNN [11] and its variants, 2D-convolution-based ResNet [12] series and RepVGG [13].

**ECAPA-TDNN:** We trained ECAPA-TDNN large with 1024 channels and its two variants, ECAPA-TDNN-X3(EX3) and ECAPA-TDNN-X4(EX4). We used SpeechBrain implementation with 1024 channels. To further boost its performance, we made it deeper and added branches in res2block to enhance representational ability. For EX3, we combine two seres2block as one basic block. For EX4, we added one more

basic block, and meanwhile, to restrict the receptive field, we did not use dilation in the first few blocks. Besides, we used group convolution to preserve the original feature map.

**ResNet/SE-ResNet:** ResNet is one of the most popular model structures currently. Here we used standard ResNet34 with 64 channels. Squeeze excitation module [14] uses an attention mechanism to re-weight feature map channels. Besides, we use a modified version of ResNet as described in [15] except that the SE module was added only in the first two blocks. We trained SE-ResNet101 with 32 channels and 64 channels.

**HS-ResNet & Res2Net** In order to model multi-scale features, we used Res2Net [16] and its variants hs-resnet [17] with dssa module. Both models are 50 layers deep, scale 8 and width 14 in Res2Net and scale 8 and width 6 in HS-ResNet.

**RepVGG** RepVGG use 3 branches of convolution and batch normalization when training and re-parameterize them as one for inference. Branch greatly boost capability to model multi-level features. We use the official version with configuration b1.

### 1.3. Pooling

Deep neural network-based systems use a pooling layer to aggregate frame-level features into segment-level embeddings. We used two pooling methods: attentive statistics pooling(ASP) [18] and channel-wise correlation pooling(CWC) [19].

### 1.4. Loss Function

Margin-based loss functions have greatly improved system performance. In this challenge, we adopted circle loss [20]. Moreover, subcenter [21] and intertopk [22] are two plugin methods that could further improve the discrimination of embeddings. Circle loss is formulated as follows:

$$L_{circle} = -\log \frac{e^{s \cdot (m^2 - (1-s_p)^2)}}{e^{s \cdot (m^2 - (1-s_p)^2)} + \sum_{j=1, j \neq i}^C e^{s \cdot ((s_n^j)^2) - m^2}}$$

where  $m$  is the margin and  $s$  is the scale factor.

### 1.5. Training Protocol

We trained models with a two-stage protocol. All experiments were based on the PyTorch [23].

Adam optimizer with weight decay 5e-5 was used in the first stage. Cycle learning rate scheduler [24] was adopted, where the minimum learning rate is 1e-8 and the maximum is 1e-3, we trained 2 cycles with one cycle of 100k steps. The batch size was 1024 and the segment duration was 2s. Margin and scale were set to 0.35 and 60 for circle loss, respectively. We used subcenter k=3 and intertopk k=5, m=0.1.

The second stage was large margin fine-tuning(LMF), we expanded segment duration to 6s, only removed intertopk from

\*equal contribution

Table 1: *Res2Net50's Performance on VoxCeleb Official Evaluation Sets*

Stage	Vox1O		Vox1E		Vox1H		Vox20-dev		Vox21-dev		Vox22-dev		Vox22-eval	
	EER	DCF <sub>0.05</sub>	EER	DCF <sub>0.05</sub>	EER	DCF <sub>0.05</sub>	EER	DCF <sub>0.05</sub>	EER	DCF <sub>0.05</sub>	EER	DCF <sub>0.05</sub>	EER	DCF <sub>0.05</sub>
base	0.670	0.0483												
+LMF	0.516	0.0303	0.656	0.0386	1.172	0.0670	1.993	0.1024	1.942	0.1091	1.751	0.0977	3.126	0.1715
++AS-norm	0.532	0.0298	0.626	0.0370	1.123	0.0644	1.922	0.0980	1.891	0.1047	1.726	0.0982	3.024	0.1662
+++QMF	0.473	0.0283	0.587	0.0350	1.059	0.0610	1.792	0.0972	1.795	0.1067	1.657	0.0972	2.983	0.158

Table 2: *Experiment results on Vox22-dev and Vox22-eval*

Index	Backbone	Pooling	VoxSRC2022-dev		VoxSRC2022-eval	
			EER	minDCF <sub>0.05</sub>	EER	minDCF <sub>0.05</sub>
1	Res2Net50-32	ASP	1.657	0.0972	2.983	0.158
2	ECAPA-TDNN-large	ASP	2.402	0.1585	-	-
3	ECAPA-TDNN-X3	ASP	1.930	0.1245	-	-
4	ECAPA-TDNN-X4	ASP	1.912	0.1179	-	-
5	ResNet34-64	ASP	2.112	0.1297	-	-
6	ResNetSE101-32	CWC	1.782	0.1091	-	-
7	ResNetSE101-64	CWC	2.145	0.1437	-	-
8	Res2Net50-64	ASP	2.021	0.1356	-	-
9	HS-ResNet-DSSA	ASP	2.082	0.1279	-	-
10	RepVGG-B1	ASP	1.725	0.1007	-	-
fusion						
fusion1	1+3+4+6+8+9+10		1.484	0.0873	2.585	0.1408
fusion2	1+2+3+4+5+6+7+8+9+10		1.495	0.0874	2.538	0.1483
fusion3	1+2+3+4+5+6+7+8+9+10+LR		1.382	0.0825	2.414	0.1397

loss function, and increased weight decay from  $5e-5$  to  $4e-4$ . We used a constant learning rate of  $2e-5$  for 10k steps.

## 1.6. Back-end

After LMF, cosine distance was used for  $4s \times 10$  scoring. Evenly cut 10 4-second long segments from utterances, the mean of score matrix with the size of  $\mathbb{R}^{10 \times 10}$  served as the score of a trial. Then, we used adaptive score normalization(AS-norm) [25] and quality measure functions(QMF) [26] to calibrate the scores. For AS-norm, speaker-wise averaged embeddings from vox2dev, leading to 5994 cohort speakers with top 400 imposter scores were used. By the way, we removed imposter variance. For QMF, we followed IDLAB’s method to generate 30k trials from vox2dev. Then we trained the logistic regression(LR) model to calibrate the AS-normed score. In the end, another LR model was used to combine all of the calibrated models to get the final fused score. While the generated trials could not perfectly fit the evaluation distribution, we manually tuned the model weights based on Vox22-dev trials.

## 1.7. Results

### 1.7.1. Ablation Study

Res2Net50 is our best single system, table 1 shows its performance on all of the evaluation sets at different stages. Equal Error Rate(EER) and minimum Decision Cost Function(minDCF) with  $C_{FA} = 1, C_M = 1, P_{target} = 0.05$  was reported. After the first stage training, model achieves  $EER = 0.67\%, minDCF = 0.0483$  on Vox1O. EER improved from 0.67% to 0.516% and minDCF improved from 0.0483 to 0.0303 after LMF. AS-norm further decreased EER

to 0.532 and minDCF to 0.0298. QMF finally push the limit of the model to  $EER=0.473\%$  and  $minDCF=0.0283$ . In total, with these stacked methods, the performance got relative improvement of 29.4% and 41.4% on EER and minDCF, respectively.

### 1.7.2. System Performance

Table 2 shows our 10 subsystems performance on Vox22-dev and Vox22-eval. We found that for TDNN-based model, the larger model is, the better performance we got. It was half true for ResNet-series models, as we found simply double the channels cannot bring improvement.

We fused ResNet-series models firstly with equal weights, got  $EER=2.585, minDCF=0.1408$ . Adding TDNN-series models only got improvement on EER but degradation on minDCF. Then, we introduced LR model to train on generated QMF set and tuned the weights based on model coefficients, finally achieves  $EER=2.414$  and  $minDCF=0.1397$ .

## 2. Semi-Supervised Domain Adaptation

Semi-supervised speaker recognition attempts to automatically exploit a large amount of target or source unlabeled data in addition to a large amount of source or target labeled data to improve performance. There are three general goals, one is to obtain better performance on the target domain data, the other is to improve the performance on all domains, that is, to improve the domain robustness, and the third is to achieve better performance on the source domain data. The goal of this competition is to achieve the performance of the target domain.

We attempt two frameworks, one is pseudo labeling, and the other is self-supervised learning. The pseudo-label solution

contains five stages: **1** source label data model training, **2** embeddings domain adaptation, **3** pseudo-label generation, **4** Supervised training on target domain data with pseudo-labels and source domain label data, **5** pseudo label correction and re-train.

### 2.1. Pre-processing

Because of the lack of filtering when constructing the CN-Celeb2 [27, 28], it contains much noisy audio. One of the most intuitive manifestations is that there is much-repeated audio in CN-Celeb2, and some are given different labels. We directly used md5sum to de-duplicate the speech and the number of audio decreased from 455,946 to 409,628.

### 2.2. Base Model Training

To obtain high confidence edges by using the voting strategy, it is beneficial to select models with as much variance as possible, from the model structure to the training Protocol. We selected the following five models: (1) *SE-ResNet34* with 32 channel; (2) *ECAPA-TDNN* with 1024 channel; (3) *Conformer-MFA* with 256 hidden dim [29]; (4) *SE-ResNet101* with 32 channel; (5) *Cot-Net* [30] with 32 channel; Others settings are shown in Table 3.

### 2.3. domain adaptation

During the evaluation process, domain adaptation is necessary due to a large domain mismatch. An important manifestation of domain mismatch on embedding is the difference between the mean and variance. Thus, the simplest approach is to align the centers of the different domains directly, and experiments show significant improvements. Further, aligning the variance can also achieve improvements in theory. We attempt to apply CORAL [31], CORAL+ [32] and CORAL++ [33] into embeddings of target domain directly and use cosine similarity for scoring. However, there is no performance improvement unless we use back-ends, LDA, and PLDA [34, 35]. Due to time constraints and inconvenient operations, we do not do this work systematically and will do these in the future.

### 2.4. cluster

Because AHC is computationally infeasible and k-means depend on the estimate of k, we propose a novel cluster algorithm, a progressive sub-graph clustering algorithm based on two Gaussian fitting and multi-model voting, denoted as GMVPG clustering. The key points of this algorithm are as follows: First, finding high-confidence positive trials, that is, edges, using a multi-model voting strategy based on the KNN affinity graph. Secondly, utilizing connected sub-graphs to obtain pseudo labels Then, using iterative top-k information to gradually merge sub-classes to prevent super-classes. Finally, two Gaussian distributions are introduced to fit the intra-class score distribution to further check for high-confidence edges. The detailed algorithm is shown as follows:

### 2.5. Supervised training and fine-tune

#### 2.5.1. Training stage1

The training data for track3 contains VoxCeleb2, the unlabeled target domain dataset from Cnceleb2, and the small amount of labeled set. Speed perturbation augmentation is used in all data, and other augmentations are the same as track 1. In the first training stage for track 3, we explore two training strategies, one is to train the model from scratch, and the other is to utilize

models from Track1 as the pre-trained model. For the latter, it is necessary to make the new model be converged before starting training, freezing the extractor and training the classification layer first is an effective approach. In addition, we also explore two different training protocols, one is Adam optimizer with cycle learning rate scheduler, as shown in Section 1, and the other is SGD optimizer with ReduceLRonPlateau scheduler.

---

#### Algorithm 1: GMVPG clustering algorithm

---

**Input** : Embeddings of target domain audio data from multi models after domain adaptation  $X^t, t = 0, 1, \dots, T$ ;

**Output**: pseudo label of target domain audio data  $D$ , here,  $Ans = (d_i, y_i), i = 0, 1, \dots, M, y_i = -1, 0, 1, \dots, N$ ;

- 1 Collect information about KNN affinity graph for  $D$ ,  $K$  is set to 500;  $sim(x_i^t, k)$  represents the top  $kst$  similarity of  $x_i^t$ ;
  - 2 filter out partial audio  $d_i$ , if  $\exists t, sim(x_i^t, K) > th_{high}$ ;
  - 3 Construct initial KNN affinity graph for  $D$ , initial  $k$  is set to 10;
  - 4 preserve edges between  $d_i$  and  $d_j, e_{ij}$  if  $sim(x_i^t, x_j^t) \geq sim(x_i^t, k) \quad \forall t$ , other edges are deleted;
  - 5 Utts are deleted if no edges are connected. For convenience, we denote  $E_k = e_{ij}$  for all preserved edges and  $utt_k = d_i$  for all preserved utts when  $k$ , respectively;
  - 6 Obtain initial labels by searching connected sub-graph (SCSG),  $G(k)$ , one sub-graph means one class;
  - 7 **repeat**
  - 8  $E_{add} = E_{k+5} - E_k, U_{add} = U_{k+5} - U_k$ ;
  - 9 Split  $E_{add}$  to  $E_{old.old}, E_{new.old}$  and  $E_{new.new}$  according to whether utt is in  $U_k$ ;
  - 10 Generate temp pseudo labels  $tmpG_{new}(k+5)$  for  $U_{add}$  based on  $E_{new.new}$  only;
  - 11 Combine class for  $U_k$  based on  $E_{old.old}$  only;  $tmpG_{old}(k+5) = subspkCombine(E_{old.old}, G(k))$ ;
  - 12 Process the relationship between  $U_{add}$  and  $U_k$ ;  $G(k+5) = CbNewOld(tmpG_{new}(k+5), G(k), E_{new.old})$ ;
  - 13  $k = k + 5$ ;
  - 14 **until**  $k=50$ ;
  - 15 Throw out classes with fewer than 10 utts;
- 

#### 2.5.2. Training stage2

In this stage, we only use CN-Celeb dataset without speed perturbation to finetune all systems. But, the VoxCeleb weights of the classification layer are preserved to prevent overfitting. Two fine-tuning training protocols are utilized, one is SGD with a  $2e-5$  learning rate, and the other is Adam with a  $2e-5$  learning rate. Others are the same as Track 1.

### 2.6. pseudo label correction

Since the GMVPG method brings some mislabeled and noisy samples, it is important to correct labels after the initial model training. The key points of this algorithm are as follows:

- i. Split audio into three types according to its similarity to centers, high/median/low-confidence;
- ii. Stas the correlation

between each class based on audio with median-confidence **iii**. Integrate the results from multi models to correct the labels.

The detailed algorithm is shown as follows: **1.** calculate the similarity of all audio in CN-Celeb to the two most similar class centers, denoted as  $sim_i^{top1}$  and  $sim_i^{top2}$ , the labels of most similar class centers are denoted as  $y_i^{top1}$  and  $y_i^{top2}$ ; **2.** Split audios according to  $sim_i^{top1}$  and  $sim_i^{top2}$ :  $sim_i^{top1} > 0.5$  and  $sim_i^{top2} < 0.4$  high-confidence;  $sim_i^{top1} > 0.5$  and  $sim_i^{top2} > 0.4$  median-confidence;  $sim_i^{top1} < 0.5$  low-confidence; **3.** Use audio with median-confidence to find labels of samples, which comes from the same speaker but are given labels of multiple speakers by the GMVPG clustering algorithm, **4.** Two classes are merged into one when multiple models all show that they need to be merged. **5.** Filter out audio that is low confidence, other audio is labeled by using predicted posterior probability.

**Algorithm 2:** Functions of GMVPG clustering

```

1 Function CheckCombine ( $U$ ):
2   Calculate similarity between all  $d$ , for  $d \in U$ ;
3   Using two-Gaussian distribution to fit the scores,
    $\mu_1, \sigma_1, w_1, \mu_2, \sigma_2, w_2$  represent the parameters
   of max and min Gaussian;
4   if  $\mu_2 > th_{nm}$  OR  $w_1 \geq 0.5$  OR
    $(\mu_1 - \sigma_1) \leq (\mu_2 + \sigma_2) + \epsilon$  then
5     return Yes
6   else
7     return No
8   end
9 End Function
10 Function subspkCombine ( $E, G$ ):
11   for each  $e_{ij} \in E$  do
12      $tmp_D = G_i 2utts + G_j 2utts$ ;
13     if CheckCombine( $tmp_D$ ) == No then
14       delete  $e_{ij}$  from  $E$ 
15     end
16   end
17   SCSG based on  $E$  to get  $G$ ;
18   return  $G$ 
19 End Function
20 Function CbNewOld ( $G_{new}, G_{old}, E_{new.old}$ ):
21    $U^{old} = G_{old}.nodes$ ;
22    $U^{new} = G_{new}.nodes$ ;
23   for each  $u_i \in U_{new}$  do
24      $utt2subgraph_i =$ 
      $\{u_j^{old} \text{ if } E_{new.old}(u_i, u_j^{old})\}$  for
      $u_j^{old} \in U^{old}$ ;
25      $U_i = utt2subgraph_i.nodes$ ;
26     if CheckCombine( $U_i$ ) == No then
27       delete  $u_i$  from  $G_{new}$  and  $E_{new.old}$ 
28     end
29   end
30   SCSG based on  $\{G_{new}, E_{new.old}, G_{old}\}$  to get  $G$ ;
31   return  $G$ 
32 End Function

```

## 2.7. Score calibration

Score calibration has been an essential part in recent VoxSRCs. However, unlike track1&2, track3 provided only 50 labeled speakers. To build a matched developing set as far as

we can, we filtered 70 speakers with 20 segments each from unlabeled data according to their clustering purity. We gave labeled speakers more weight when generating developing trials as pseudo-label could be wrong. Finally, we got 40000 trials which is equal to the validation trials. The results show that our developing set is slightly easier but still brings performance gain.

## 2.8. Results

Models shown in Table 3 are used to cluster. Systems in Table 4 are developed for Track3. Results we submitted are shown in Table 5. After GMVPG clustering algorithm, we obtain 1711 speakers and 348861 utts. After model training and label correction, we obtain 1760 speakers and 387700 utts.

Table 3: Results of base model before/after adaptation

mdl	loss	vox2-train	t3-dev-EER	
			ini	adapt
se-resnet34-32	circle	clean-fb64-sgd	16.86	14.29
cotnet	circle	clean-fb64-sgd	16.65	14.55
conformer	circle	aug-fb80-adam	16.95	14.14
ecapa-large	circle	aug-fb80-adam	18.02	14.62
se-resnet101-32	circle	aug-fb80-adam	14.06	11.90

Table 4: Results of systems for Track3. v0 means pseudo labels before correction, and v1 means after.

	mdl	loss	train	t3-dev-EER	
				ini	calib
S1	Res2Net50	circle	v0-sgd	<b>8.45</b>	8.20
S2	ResNet34	circle	v0-sgd	8.61	<b>7.66</b>
S3	ECAPA-X4	circle	v0-sgd	9.57	8.81
S4	ECAPA-X4	circle	v0-adam	10.47	9.65
S5	ECAPA-X4	circle	v1-sgd	8.78	8.42

Table 5: Results of systems that we submitted

mdl	mode	dev-EER	eval-EER
S1	ini	8.45	8.07
S2	ini	8.61	8.64
S1+S2	ini	8.01	7.57
S1+S2+S3+S4	ini	7.87	7.40
S1+S2+S3+S4+S5	calib	<b>6.77</b>	<b>7.03</b>

## 3. Conclusions

In this paper, we summarized our systems for VoxSRC2022 in detail. For track1, we explore various strong speaker embedding extractors and some training tricks. For Track3, we explore some domain adaptation methods firstly. Then, we propose a novel progressive sub-graph clustering algorithm based on two Gaussian fitting and multi-model voting to obtain pseudo labels. Thirdly, we explore some fine-tuning tricks to achieve better performance. Finally, we propose one label correction algorithm to correct noisy labels. Our Fusion systems achieve 6<sup>th</sup> and 1<sup>st</sup> place in Track 1 and 3 respectively.

## 4. References

- [1] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *conference of the international speech communication association*, 2018.
- [2] H. Yamamoto, K. A. Lee, K. Okabe, and T. Koshinaka, "Speaker augmentation and bandwidth extension for deep speaker embedding," *conference of the international speech communication association*, 2019.
- [3] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv: Sound*, 2015.
- [4] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," *international conference on acoustics, speech, and signal processing*, 2017.
- [5] M. Zhao, Y. Ma, M. Liu, and M. Xu, "The speakin system for voxceleb speaker recognition challenge 2021," *arXiv preprint arXiv:2109.01989*, 2021.
- [6] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, "SpeechBrain: A general-purpose speech toolkit," 2021, *arXiv:2106.04624*.
- [7] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: A large-scale speaker identification dataset," *conference of the international speech communication association*, 2017.
- [8] A. Nagrani, J. S. Chung, J. Huh, A. Brown, E. Coto, W. Xie, M. McLaren, D. A. Reynolds, and A. Zisserman, "Voxsrc 2020: The second voxceleb speaker recognition challenge," *arXiv preprint arXiv:2012.06867*, 2020.
- [9] 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.
- [10] Y.-Y. Yang, M. Hira, Z. Ni, A. Chourdia, A. Astafurov, C. Chen, C.-F. Yeh, C. Puhersch, D. Pollack, D. Genzel, D. Greenberg, E. Z. Yang, J. Lian, J. Mahadeokar, J. Hwang, J. Chen, P. Goldsborough, P. Roy, S. Narenthiran, S. Watanabe, S. Chintala, V. Quenneville-Bélair, and Y. Shi, "Torchaudio: Building blocks for audio and speech processing," *arXiv preprint arXiv:2110.15018*, 2021.
- [11] B. Desplanques, J. Thienpondt, and K. Demuynck, "Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," *Proc. Interspeech 2020*, pp. 3830–3834, 2020.
- [12] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [13] X. Ding, X. Zhang, N. Ma, J. Han, G. Ding, and J. Sun, "Repvgg: Making vgg-style convnets great again," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 13 733–13 742.
- [14] J. Hu, L. Shen, S. Albanie, G. Sun, and E. Wu, "Squeeze-and-excitation networks," *computer vision and pattern recognition*, 2018.
- [15] Y. Shengyu, F. Xiang, Y. Jie, L. Jingdong, and P. Yiqian, "Sogou system for the voxceleb speaker recognition challenge 2021," *arXiv: Sound*, 2021.
- [16] S.-H. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. Torr, "Res2net: A new multi-scale backbone architecture," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 2, pp. 652–662, 2019.
- [17] Z. Li, "Explore long-range context feature for speaker verification," *CoRR*, vol. abs/2112.07134, 2021. [Online]. Available: <https://arxiv.org/abs/2112.07134>
- [18] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive statistics pooling for deep speaker embedding," *Proc. Interspeech 2018*, pp. 2252–2256, 2018.
- [19] T. Stafylakis, J. Rohdin, and L. Burget, "Speaker embeddings by modeling channel-wise correlations," in *Interspeech*, 2021, Conference Proceedings.
- [20] R. Xiao, X. Miao, W. Wang, P. Zhang, B. Cai, and L. Luo, "Adaptive Margin Circle Loss for Speaker Verification," in *Proc. Interspeech 2021*, 2021, pp. 4618–4622.
- [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 *ECCV*. Springer, 2020, Conference Proceedings, pp. 741–757.
- [22] M. Zhao, Y. Ma, M. Liu, and M. Xu, "The speakin system for voxceleb speaker recognition challenge 2021," *arXiv: Sound*, 2021.
- [23] 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. Z. 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," *neural information processing systems*, 2019.
- [24] L. N. Smith, "Cyclical learning rates for training neural networks," *workshop on applications of computer vision*, 2015.
- [25] S. Cumani, P. D. Batzu, D. Colibro, C. Vair, P. Laface, and V. Vasilakakis, "Comparison of speaker recognition approaches for real applications," *conference of the international speech communication association*, 2011.
- [26] 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 *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5814–5818.
- [27] 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*, vol. 137, pp. 77–91, 2022.
- [28] 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," *international conference on acoustics, speech, and signal processing*, 2019.
- [29] Y. Zhang, Z. Lv, H. Wu, S. Zhang, P. Hu, Z. Wu, H.-y. Lee, and H. Meng, "Mfa-conformer: Multi-scale feature aggregation conformer for automatic speaker verification," *arXiv preprint arXiv:2203.15249*, 2022.
- [30] Y. Li, T. Yao, Y. Pan, and T. Mei, "Contextual transformer networks for visual recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [31] B. Sun, J. Feng, and K. Saenko, "Correlation alignment for unsupervised domain adaptation," in *Domain Adaptation in Computer Vision Applications*. Springer, 2017, pp. 153–171.
- [32] K. A. Lee, Q. Wang, and T. Koshinaka, "The coral+ algorithm for unsupervised domain adaptation of plda," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5821–5825.
- [33] R. Li, W. Zhang, and D. Chen, "The coral++ algorithm for unsupervised domain adaptation of speaker recognition," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7172–7176.
- [34] S. Ioffe, "Probabilistic linear discriminant analysis," in *European Conference on Computer Vision*. Springer, 2006, pp. 531–542.
- [35] P. Kenny, "Bayesian speaker verification with heavy tailed priors," *Proc. Odyssey 2010*, 2010.