

# The DKU-MSXF Diarization System for the VoxCeleb Speaker Recognition Challenge 2023

Ming Cheng<sup>1</sup>, Weiqing Wang<sup>1</sup>, Xiaoyi Qin<sup>1</sup>, Yuke Lin<sup>1</sup>, Ning Jiang<sup>2</sup>,  
Guoqing Zhao<sup>2</sup>, and Ming Li<sup>1\*</sup>

<sup>1</sup> Data Science Research Center, Duke Kunshan University, Kunshan, China

<sup>2</sup> Mashang Consumer Finance Co., Ltd.  
`ming.li369@dukekunshan.edu.cn`

**Abstract.** This paper describes the DKU-MSXF submission to track 4 of the VoxCeleb Speaker Recognition Challenge 2023 (VoxSRC-23). Our system pipeline contains voice activity detection, clustering-based diarization, overlapped speech detection, and target-speaker voice activity detection, where each procedure has a fused output from 3 sub-models. Finally, we fuse different clustering-based and TSVAD-based diarization systems using DOVER-Lap and achieve the 4.30% diarization error rate (DER), which ranks first place on track 4 of the challenge leaderboard.

**Keywords:** Speaker diarization · Target-speaker voice activity detection · VoxSRC-23.

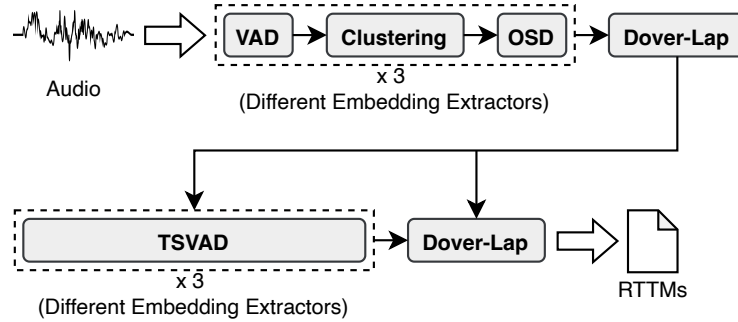
## 1 Introduction

Speaker diarization is the task of breaking up multi-party conversational audio into speaker-homogeneous segments, which aims to solve the problem of "Who spoke when." In this paper, we focus on the speaker diarization task and present the details of our submitted system to track 4 of the VoxCeleb Speaker Recognition Challenge 2023.

Fig. 1 depicts the framework of our developed system. First, the voice activity detection (VAD) module removes non-speech regions from the input audio. The remaining is split into multiple short segments, followed by speaker embedding extraction. Then, the initial clustering-based diarization results can be obtained by agglomerative hierarchical clustering (AHC) with overlapped speech detection (OSD) as the post-processing. We replace speaker embedding models trained under different conditions to repeat the above process three times and obtain the fused clustering-based results by Dover-Lap [14]. Next, target-speaker voice activity detection (TSVAD) models with different speaker embedding models are adopted to refine the clustering-based results. In the end, clustering-based and TSVAD-based results are fused again to obtain the final prediction.

---

\* Corresponding Author



**Fig. 1.** Framework of The Developed System.

In general, the framework integrates the advantages of clustering-based and TSVAD-based methods, which is similar to our previous submissions in VoxSRC-21 [18] and VoxSRC-22 [20]. The main differences are improved speaker embedding models and Seq2Seq-TSVAD models from our recent works [2, 1].

## 2 Dataset Description

According to challenge rules, any data except the test set is allowed in this task. The datasets used to train each model in our system are described as follows.

- Voice activity detection (VAD) and overlapped speech detection (OSD): VoxCeleb 1&2 [11, 4] for data simulation and VoxConverse [3] for adaptation and validation.
- Speaker embedding: VoxCeleb 1&2 [11, 4] and VoxBlink-Clean [10] for training and evaluation.
- Clustering-based diarization: VoxConverse [3] for hyper-parameter tuning.
- TSVAD-based diarization: VoxCeleb 1&2 [11, 4] for online data simulation and VoxConverse [3] for adaptation and validation.
- Data augmentation: MUSAN [15] and RIRs[9] corpora.

## 3 Model Configuration

This section describes each model in our submitted system. If not specified, the input acoustic features of all models are 80-dim log Mel-filterbank energies with a frame length of 25 ms and a frameshift of 10 ms. MUSAN [15] and RIRs [9] are applied as the data augmentation.

### 3.1 VAD & OSD

As the implementations of voice activity detection (VAD) and overlapped speech detection (OSD) tasks are very similar, we utilize the same neural networks to train these two tasks. Adopted model architectures are described as follows.

**Table 1.** False alarm (FA) and miss detection (MISS) rates of different VAD and OSD models on the VoxConverse test set.

Task	Model	FA (%)	MI (%)	Total (%)
VAD	Conformer	2.84	1.09	3.93
	ResNet34	3.20	1.02	4.22
	ECAPA-TDNN	2.70	1.51	4.21
	Fusion	2.83	1.14	3.97
OSD	Conformer	0.59	1.41	2.00
	ResNet34	0.54	1.51	2.05
	ECAPA-TDNN	0.52	1.45	1.97
	Fusion	0.44	1.45	1.89

Conformer [7] is the first backbone network. All encoder layers share the same settings: 256-dim attentions with 4 heads and 1024-dim feed-forward layers with a dropout rate of 0.1. The kernel size of convolutions in Conformer blocks is 15. Finally, a linear layer with sigmoid activation is adopted to transform the dimension of Conformer outputs into one, representing the frame-level posterior probability of VAD or OSD.

ResNet34 [8] is the second backbone network, where the widths (number of channels) of the residual blocks are  $\{64, 128, 256, 512\}$ . At the end of convolutional blocks, the spatial average pooling extracts frame-level features from the convolutional outputs. Finally, a linear layer with sigmoid activation predicts the frame-level posterior probability of VAD or OSD.

ECAPA-TDNN [6] is the third backbone network. The number of filters in the convolutional layers is set to 1024. The scale dimension in the Res2Block is set to 8. The dimension of the bottleneck in the SE-Block and the attention module is set to 128. Finally, a linear layer with sigmoid activation predicts the frame-level posterior probability of VAD or OSD.

VAD or OSD models with different network architectures are fused by averaging their predictions at the score level. Table 1 shows the VAD and OSD models' performance on the VoxConverse test set, respectively. We adopt the fused results as the final predictions in this part.

### 3.2 Speaker Embedding

To ensemble learning, we train three speaker embedding models with diverse network architectures and training data. The first one is SimAM-ResNet34 [13] with statistics pooling (SP) [16]. The second one is ResNet101 [8] with attentive statistics pooling (ASP) [12]. The third one is SimAM-ResNet100 with ASP. The first two models are trained on the VoxCeleb2 [4] dataset. For the last model, we additionally mix the VoxBlink-Clean [10] dataset into the training.

All speaker embedding models utilize the same back-end part. After the pooling layer projects the variable-length input audio to the fixed-length vector,

**Table 2.** Equal error rates (EERs) of different speaker embedding models on the Vox-O trial.

#	Model	Training Data	EER (%)
Spk1	SimAM-ResNet34+SP	Vox2	0.81
Spk2	ResNet101+ASP	Vox2	0.49
Spk3	SimAM-ResNet100+ASP	Vox2+VoxBlink	0.44

**Table 3.** Thresholds (THR) of different AHC-based diarization systems.

#	Segment THR	Stop THR	Speaker THR
Ahc1	0.54	0.60	0.20
Ahc2	0.62	0.62	0.20
Ahc3	0.66	0.68	0.30

a 256-dim fully connected layer is adopted as the speaker embedding layer, and the ArcFace ( $s=32, m=0.2$ ) [5] is used as a classifier. The detailed configuration of model training is the same as [13]. The performance of different trained speaker embedding models is shown in Table 2.

### 3.3 Clustering-based Diarization

We adopt agglomerative hierarchical clustering (AHC) to implement the clustering-based diarization system, which is the same as we used in previous years [18, 20]. First, speaker embeddings are extracted from the uniformly segmented speech with a length of 1.28s and a shift of 0.32s, pairwise measured by cosine similarity. Two consecutive segments are merged into a longer segment if their similarity exceeds a segment threshold. Then, we perform a plain AHC on the similarity matrix with a relatively high stop threshold to obtain clusters with high confidence. These clusters are split into "long clusters" and "short clusters" by the total duration in each cluster, and the central embedding of each cluster is the mean of all speaker embeddings within the cluster. Finally, each short cluster is assigned to the closest long cluster by the similarity between their central embeddings. If a short cluster is too different from all long clusters, the similarity between them is lower than a speaker threshold, and then we treat it as a new speaker. As post-processing based on the OSD model, overlapped speech regions are assigned to the two closest speaker labels.

Based on speaker embedding models in Table 2, we develop three AHC-based diarization systems following the above approach. All the hyper-parameters are directly tuned on the VoxConverse test set by grid search. The duration for classifying long and short clusters is 6s for all models. The other thresholds are shown in Table 3.

### 3.4 TSVAD-based Diarization

Given speaker profiles (e.g., x-vector [16]), the TSVAD method can estimate each speaker’s frame-level voice activities and perform robustly even in complex acoustic environments. We adopt the Seq2Seq-TSVAD [2] consisting of the front-end extractor with segmental pooling [19], Conformer [7] encoder, and proposed speaker-wise decoder. The front-end extractor is initialized by the pre-trained speaker embedding model, the same as the one for extracting speaker profiles.

Each model training starts from BCE loss & Adam optimizer with a learning rate of  $1e - 4$  and a linear warm-up of 2,000 iterations. The whole process is described as follows.

- First, the model with a frozen front-end extractor is trained on simulated data until back-end convergence.
- Second, all model parameters are unfrozen to train on both simulated (Vox-Celeb 1&2) and real data (VoxConverse dev set) at the ratio of 0.8/0.2. Also, the learning rate decreases to  $1e - 5$  in the second half of this phase.
- Third, all data simulation and augmentation are removed, inspired by the large margin finetuning (LMFT) [17] in speaker verification. Meanwhile, to use much training data as possible, we mix the first 186 samples of the VoxConverse test set into finetuning and leave the last 46 samples as validation, namely the VoxConverse test46 set.

During inference, a clustering-based diarization is required first to extract speaker profiles from each speaker’s speech segments. Then, each test audio is cut into fixed-length chunks with a stride of 1s and fed into the TSVAD model with extracted speaker profiles. Chunked predictions are stitched by averaging the overlapped predicted regions, which can also be viewed as a score-level fusion.

Using speaker embedding models in Table 2 as different profile extractors, we develop three TSVAD-based diarization systems. The first one equipped with SimAM-ResNet34 is trained and inferred under audio chunks of 64s. The last two equipped with ResNet101 and SimAM-ResNet100 are trained and inferred under audio chunks of 16s, limited by higher GPU memory costs of the deeper networks. The other configurations share the same settings: 512-dim attentions with 8 heads, convolutions with a kernel size of 15, and 1024-dim feed-forward layers with a dropout rate of 0.1. The speaker capacity and output VAD resolution are set to 30 and 0.08s, respectively.

## 4 Experimental Results

Table 4 illustrates the performance of our developed diarization systems on the VoxConverse test set, VoxConverse test46 set, and VoxSRC-23 challenge test set. As the TSVAD models utilize part of the VoxConverse test set in training, their performance is only evaluated on the other two datasets.

Systems #1-3 represent the AHC-based diarization with different speaker embedding models, obtaining the DERs of 5.51%, 5.32%, and 5.36% on the

**Table 4.** Diarization error rates (DERs) of different systems on the VoxConverse test set, VoxConverse test46 set, and VoxSRC-23 challenge test set.

#	Method	DER (%)		
		VoxConverse Test	VoxConverse Test46	VoxSRC-23 Challenge Test
1	Ahc1	4.83	4.14	5.51
2	Ahc2	4.49	3.92	5.32
3	Ahc3	4.55	3.92	5.36
4	Dover-Lap (#1-3)	-	3.81	5.19
5	+ TSVAD with Spk1	-	2.85	4.49
6	+ TSVAD with Spk2	-	2.93	4.57
7	+ TSVAD with Spk3	-	2.91	4.53
8	Dover-Lap (#4-7)	-	2.73	4.30

challenge test set. Then, we adopt the Dover-Lap method to fuse them and obtain system #4 to achieve a 5.19% DER on the challenge test set.

Systems #5-7 represent the TSVAD-based diarization with different speaker embedding models. Using the fused clustering-based results to extract speaker profiles, these three models have very close DERs varying from 4.49% to 4.57% on the challenge test set. Finally, systems #4-7 are fused by the Dover-Lap again to obtain system #8, achieving the best 4.30% DER on the challenge test set.

## 5 Conclusions

This paper describes our system development for track 4 of the VoxSRC-23. This year, we mainly focus on improving the speaker embedding, AHC-based diarization, and TSVAD-based diarization models. To achieve the best performance, we also train diverse sub-models for each part of the whole framework. The finally fused method shows significant improvement, obtaining the DERs of 2.73% on the VoxConverse test46 set and 4.30% on the VoxSRC-23 challenge test set, respectively.

## Acknowledgments

This research is funded in part by the National Natural Science Foundation of China (62171207) and Science and Technology Program of Suzhou City(SYC2022051). Many thanks for the computational resource provided by the Advanced Computing East China Sub-Center.

## References

1. Cheng, M., Wang, H., Wang, Z., Fu, Q., Li, M.: The whu-alibaba audio-visual speaker diarization system for the misp 2022 challenge. In: Proc. ICASSP. pp. 1–2 (2023). <https://doi.org/10.1109/ICASSP49357.2023.10095802>
2. Cheng, M., Wang, W., Zhang, Y., Qin, X., Li, M.: Target-speaker voice activity detection via sequence-to-sequence prediction. In: Proc. ICASSP. pp. 1–5 (2023). <https://doi.org/10.1109/ICASSP49357.2023.10094752>
3. Chung, J.S., Huh, J., Nagrani, A., Afouras, T., Zisserman, A.: Spot the conversation: Speaker diarisation in the wild. In: Proc. Interspeech. pp. 299–303 (2020). <https://doi.org/10.21437/Interspeech.2020-2337>
4. Chung, J.S., Nagrani, A., Zisserman, A.: Voxceleb2: Deep speaker recognition. In: Proc. Interspeech. pp. 1086–1090 (2018). <https://doi.org/10.21437/Interspeech.2018-1929>
5. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proc. CVPR. pp. 4685–4694 (2019). <https://doi.org/10.1109/CVPR.2019.00482>
6. Desplanques, B., Thienpondt, J., Demuynck, K.: ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification. In: Proc. Interspeech. pp. 3830–3834 (2020). <https://doi.org/10.21437/Interspeech.2020-2650>
7. Gulati, A., Qin, J., Chiu, C.C., Parmar, N., Zhang, Y., Yu, J., Han, W., Wang, S., Zhang, Z., Wu, Y., Pang, R.: Conformer: Convolution-augmented transformer for speech recognition. In: Proc. Interspeech. pp. 5036–5040 (2020). <https://doi.org/10.21437/Interspeech.2020-3015>
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proc. CVPR. pp. 770–778 (2016)
9. Ko, T., Peddinti, V., Povey, D., Seltzer, M.L., Khudanpur, S.: A study on data augmentation of reverberant speech for robust speech recognition. In: Proc. ICASSP. pp. 5220–5224 (2017). <https://doi.org/10.1109/ICASSP.2017.7953152>
10. Lin, Y., Qin, X., Cheng, M., Li, M.: Voxblink: X-large speaker verification dataset on camera (2023)
11. Nagrani, A., Chung, J.S., Zisserman, A.: Voxceleb: A large-scale speaker identification dataset. In: Proc. Interspeech. pp. 2616–2620 (2017). <https://doi.org/10.21437/Interspeech.2017-950>
12. Okabe, K., Koshinaka, T., Shinoda, K.: Attentive Statistics Pooling for Deep Speaker Embedding. In: Proc. Interspeech. pp. 2252–2256 (2018). <https://doi.org/10.21437/Interspeech.2018-993>
13. Qin, X., Li, N., Weng, C., Su, D., Li, M.: Simple attention module based speaker verification with iterative noisy label detection. In: Proc. ICASSP. pp. 6722–6726 (2022). <https://doi.org/10.1109/ICASSP43922.2022.9746294>
14. Raj, D., Paola Garcia-Perera, L., Huang, Z., Watanabe, S., Povey, D., Stolcke, A., Khudanpur, S.: Dover-lap: A method for combining overlap-aware diarization outputs. In: Proc. SLT. pp. 881–888 (2021). <https://doi.org/10.1109/SLT48900.2021.9383490>
15. Snyder, D., Chen, G., Povey, D.: Musan: A music, speech, and noise corpus (2015)
16. Snyder, D., Garcia-Romero, D., Sell, G., Povey, D., Khudanpur, S.: X-vectors: Robust dnn embeddings for speaker recognition. In: Proc. ICASSP. pp. 5329–5333 (2018). <https://doi.org/10.1109/ICASSP.2018.8461375>

17. Thienpondt, J., Desplanques, B., Demuynck, K.: The idlab voxsrc-20 submission: Large margin fine-tuning and quality-aware score calibration in dnn based speaker verification. In: Proc. ICASSP. pp. 5814–5818 (2021). <https://doi.org/10.1109/ICASSP39728.2021.9414600>
18. Wang, W., Cai, D., Lin, Q., Yang, L., Wang, J., Wang, J., Li, M.: The dku-dukeeece-lenovo system for the diarization task of the 2021 voxceleb speaker recognition challenge (2021)
19. Wang, W., Lin, Q., Cai, D., Li, M.: Similarity measurement of segment-level speaker embeddings in speaker diarization. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **30**, 2645–2658 (2022). <https://doi.org/10.1109/TASLP.2022.3196178>
20. Wang, W., Qin, X., Cheng, M., Zhang, Y., Wang, K., Li, M.: The dku-dukeeece diarization system for the voxceleb speaker recognition challenge 2022 (2022)