

# Voice Conversion Augmentation for Speaker Recognition on Defective Datasets

Ruijie Tao<sup>\*</sup>, Zhan Shi<sup>†‡</sup>, Yidi Jiang<sup>\*</sup>, Tianchi Liu<sup>\*</sup>, and Haizhou Li<sup>\*‡§</sup>

<sup>\*</sup> National University of Singapore, Electrical and Computer Engineering, Singapore

<sup>†</sup> University of Houston, C.T. Bauer College of Business, USA

<sup>‡</sup> The Chinese University of Hong Kong, School of Artificial Intelligence, Shenzhen, China

<sup>§</sup> Shenzhen Research Institute of Big Data, Shenzhen, China

Email: ruijie.tao@u.nus.edu

**Abstract**—Modern speaker recognition system relies on abundant and balanced datasets for classification training. However, diverse defective datasets, such as partially-labelled, small-scale, and imbalanced datasets, are common in real-world applications. Previous works usually studied specific solutions for each scenario from the algorithm perspective. However, the root cause of these problems lies in dataset imperfections. To address these challenges with a unified solution, we propose the Voice Conversion Augmentation (VCA) strategy to obtain pseudo speech from the training set. Furthermore, to guarantee generation quality, we designed the VCA-NN (nearest neighbours) strategy to select source speech from utterances that are close to the target speech in the representation space. Our experimental results on three created datasets demonstrated that VCA-NN effectively mitigates these dataset problems, which provides a new direction for handling the speaker recognition problems from the data aspect.

## I. INTRODUCTION

Speaker recognition (SR) aims to verify the identity of a speaker based on the voice characteristic [1]–[3]. In recent years, deep learning-based SR systems have achieved significant improvement [4]–[7]. These modern solutions typically employ a classification learning pipeline, relying on datasets with abundant labelled utterances and a balanced distribution [8]–[10]. However, in real-world applications, dataset collection and speech labelling can be challenging due to diverse issues, such as minor languages, limited labor cost and data source [11]–[13]. That leads to the defective SR datasets.

We take three typical SR dataset problems for consideration. Firstly, in the partially labelled datasets, semi-supervised SR has the problem of leveraging unlabelled utterances effectively [14]. Secondly, for the small-scale dataset, it is challenging to learn a generalized model from limited data [15]. Lastly, for the dataset with an imbalanced speaker distribution, the learnt model tends to prioritize speakers who contribute the majority of utterances, while ignoring those with fewer utterances [16].

Previous studies usually tackled each issue with a specific method. For instance, contrastive learning [14] and curriculum learning [17] are studied to address semi-supervised SR. [15] proposed an adversarial few-shot learning method for the SR dataset with limited data. However, these solutions primarily focus on algorithmic aspects rather than directly addressing dataset deficiencies. Recent work [18] has demonstrated that

data augmentation with generation pipelines is also a feasible approach for speech keyword detection. [19] and [20] proved that text-to-speech (TTS) generation technology [21], [22] can be used to assist cross-domain speaker verification. Nonetheless, it remains unexplored to ensure the quality of generated utterance and provide a unified solution for all the aforementioned problems.

For these defective datasets, every utterance in the training set shares the different quality and importance [23], [24]. In other words, augmenting the dataset by making these imperfect utterances learn from the valuable utterances is an efficient direction. Motivated by that, voice conversion (VC) technique is a suitable choice to supplement the dataset. Here, VC aims to change the voice characteristic of one source speech based on the target speech while preserving linguistic information [25]–[27]. [28] proposed VC for speaker verification augmentation for the text-dependent condition but ignored the importance of the quality of the generated data.

In this paper, we propose the Voice Conversion Augmentation (VCA) strategy to address the aforementioned SR problems from the dataset perspective. It intelligently selects source and target speech from the training set, and generates the pseudo speech to involve diverse voiceprint features for augmentation. Meanwhile, to enhance the quality of the generated utterances, we propose a VCA-NN (nearest neighbours) strategy to select appropriate source speech that has the similar representation to the target speech during the VCA process.

The contribution of this paper can be summarized as follows: 1) We propose VCA, a novel and unified augmentation strategy based on voice conversion to handle diverse dataset problems in speaker recognition. 2) We design the VCA-NN strategy to search for the optimal target and source speech to generate high-quality pseudo speech. 3) We create three datasets to represent scenarios of semi-supervised, small-scale, and imbalanced learning to validate the effectiveness of VCA.

## II. BACKGROUND: DIFFUSION-BASED VOICE CONVERSION

Voice conversion (VC) transforms the source speech  $s$  into the new speech, i.e., ‘pseudo speech’, preserving the linguistic content of  $s$  and the speaker characteristic of the target speech  $t$ . Based on that, as an augmentation proposal, VC can

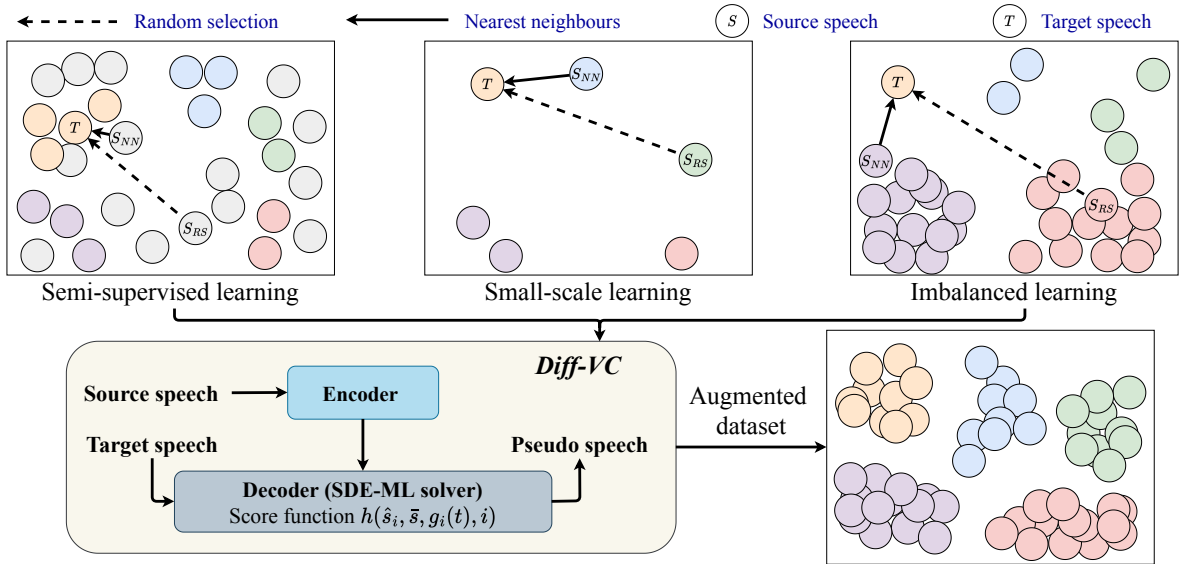


Fig. 1. Data distribution illustration for three kinds of typical defective speaker recognition datasets and the augmented dataset with the Voice Conversion Augmentation (VCA) strategy. For augmentation, the Diff-VC module transfers the source speech into the pseudo speech with the target speech’s voice characteristic. In data distribution illustration, circles with the same colour represent the utterances from the same speaker, and the grey circles denote the unlabelled utterances. For each target speech, VCA-RS selects random utterance as the source speech, while VCA-NN searches nearest neighbours to boost the quality of the pseudo speech.

involve multiple speech representations for the same person with different prosody, content, and emotions. Our proposed VCA employs the diffusion-based voice conversion (Diff-VC) method [29] due to its stable performance.

Diff-VC follows the idea of the auto-encoder. The encoder is designed to capture speaker-independent speech characteristics. It transforms the mel-spectrogram of  $s$  into a speaker-independent representation, denoted as  $\bar{s}$ , which is further fed into the decoder. The decoder is tasked with the reverse diffusion process [30], which aims to convert the speech representation into a new mel-spectrogram. This output reflects the voiceprint of the target speaker instead of the source speaker. It is formalized through a reverse stochastic differential equation (SDE) [31], given by:

$$d\hat{s}_i = \left( \frac{1}{2}(\bar{s} - \hat{s}_i) - h(\hat{s}_i, \bar{s}, i) \right) \beta_i di + \sqrt{\beta_i} d\overleftarrow{W}_i, \quad (1)$$

where  $\hat{s}_i$  represents the evolving speech during reverse diffusion at iteration time  $i$ , with  $\bar{s}$  as the encoder-derived representation. The score function  $h(\hat{s}_i, \bar{s}, i)$  is trained to incorporate the target speaker’s characteristic  $g_i(t)$ . In addition,  $\beta_i$  modulates the noise level, and  $d\overleftarrow{W}_i$  adds randomness through the backward Wiener process in SDEs. The resultant mel-spectrogram can closely mirror the target speaker’s voice, and it can be recovered into the 1D speech waveform as the pseudo speech.

### III. VOICE CONVERSION AUGMENTATION (VCA)

As shown in Figure 1, VCA employs the voice conversion technique to synthesize augmented data and mitigate three typical dataset problems. Both the source speech set  $\mathcal{S}$  and target speech set  $\mathcal{T}$  come from the training set  $\mathcal{R}$ .

#### A. Semi-supervised learning

In the semi-supervised learning scenario,  $\mathcal{R}$  consists of  $N$  labelled utterances to form the target speech set in VCA, denoted as  $\mathcal{T} = \{(x_i, y_i)\}_{i=1}^N$ , while unlabelled utterances form the source speech set in VCA, denoted as  $\mathcal{S} = \{(x_j)\}_{j=N+1}^{|\mathcal{R}|}$  ( $y_j$  is unknown). The training challenge arises since unlabelled utterances cannot directly contribute to the classification learning. To leverage the unlabelled utterances effectively, we implement voice conversion for each utterance in the labelled target set  $\mathcal{T}$ , selecting source speech from the unlabelled set  $\mathcal{S}$ . After conversion, the synthetic utterance exhibits similar voice characteristic to the target speaker, while the speech content and prosody differ. This process is repeated  $K$  times for each target speech by selecting the different source speech.  $K$  is named as ‘generation coefficient’, which means the number of augmented samples for each speech utterance. The augmented dataset provides a more comprehensive speech representation for each speaker, which can benefit the SR training.

#### B. Small-scale learning

In the small-scale learning scenario,  $\mathcal{R} = \{(x_i, y_i)\}_{i=1}^{|\mathcal{R}|}$ , where  $|\mathcal{R}|$  is small. The limited dataset size can adversely affect the generalizability of the learned speaker model. In VCA,  $\mathcal{R}$  plays the role of the source speech set  $\mathcal{S}$  and the target speech set  $\mathcal{T}$  at the same time. For each sample in  $\mathcal{T}$ , we treat it as the target utterance and perform voice conversion with the sample from the source set  $\mathcal{S}$ . This process is repeated  $K$  times to achieve various speaker representations. The target and source speech should come from different speakers to guarantee the diversity of the generated utterances.

### C. Imbalanced learning

The imbalanced dataset contains the minority part and the majority part. The minority part has  $M$  utterances and each speaker in it has a small number of utterances, this part be viewed as the target speech set  $\mathcal{T} = \{(x_i, y_i)\}_{i=1}^M$  in VCA. On the other hand, each speaker from the majority part has a large number of utterances, used as the source speech set  $\mathcal{S} = \{(x_j, y_j)\}_{j=M+1}^{|\mathcal{R}|}$ . This imbalanced distribution makes the trained SR model more biased towards the speakers in the majority part, resulting in poor representation learning. For each target utterance of minority part, VCA is applied  $K$  times to generate several pseudo utterances for supplementary. This process can achieve a more balanced speaker distribution to enhance the SR system.

### IV. NEAREST NEIGHBOURS SELECTION

To study the strategies of selecting the optimal utterance as the source speech, we take semi-supervised learning as an example, and other problems can apply analogous methods.

For each target speech, a straightforward strategy is **VCA-RS (random selection)**, which randomly chooses one sample from the entire source speech set  $\mathcal{S}$  as the source speech (represented as  $S_{RS}$  in Figure 1). However, VCA-RS lacks control over the quality of the generated utterance. For instance, transferring an utterance from a female’s voice to a male’s voice is challenging and may result in low fluency. Based on that, we propose a two-stage searching strategy named **VCA-NN (nearest neighbours)** to obtain proper source speech (represented as  $S_{NN}$  in Figure 1).

1) *Stage 1:* We train an SR model  $\Phi$  with the classification loss based on all given labelled utterances. As a basic SR model with the distinguishing ability,  $\Phi$  can extract the speech features of all utterances in  $\mathcal{R}$ .

2) *Stage 2:* Then we calculate the cosine similarity among all utterances from the target set  $\mathcal{T} = \{(x_i, y_i)\}_{i=1}^N$  and source set  $\mathcal{S} = \{(x_j)\}_{j=N+1}^{|\mathcal{R}|}$  using the extracted features. These features come from the speaker recognition model developed in Stage 1. For each target utterance  $(x_i, y_i) \in \mathcal{T}$ , we intelligently sample a set  $\mathcal{P}_i$  from  $\mathcal{S}$ .  $\mathcal{P}_i$  contains  $K$  utterances that exhibit the highest similarity score to the target speech  $(x_i, y_i)$ ,  $K$  is the generation coefficient.

$$\mathcal{P}_i = \arg \text{top-}K \cos\left(\frac{\Phi(x_i) \cdot \Phi(x_j)}{\|\Phi(x_i)\| \cdot \|\Phi(x_j)\|}\right) \quad (2)$$

where  $j \in [N + 1, |\mathcal{R}|]$ . In other words, for each target speech  $(x_i, y_i)$ , VCA-NN can obtain  $K$  pseudo utterances based on the corresponding set  $\mathcal{P}_i$ . This approach can achieve higher voiceprint similarity between the source and target speech. By selecting a source speech proximate to the target speech, the generated pseudo speech can be smoother and more fluent, which further boosts SR training.

### V. DATASET CREATION AND EXPERIMENTAL SETTINGS

#### A. Dataset generation

VoxCeleb2 [8] is the large-scale real-world SR dataset. It contains 1,092,009 speech utterances from 5,994 speakers.

Based on VoxCeleb2, we create three training sets to represent the aforementioned scenarios: 1) **Vox-Semi dataset:** For semi-supervised scenario, Vox-Semi contains the labelled part (target speech set) and the unlabelled part (source speech set). The labelled part has 1,000 speakers, each with 10 utterances as target speech. The unlabelled part contains 40,000 utterances as source speech. VCA generates  $K = 9$  samples for each labelled target speech. The augmented Vox-Semi contains 1,000 speakers, each with 100 utterances. 2) **Vox-Small dataset:** For small-scale scenario, Vox-Small has 1,000 speakers, while each with 5 utterances. VCA generates  $K = 9$  samples for each speaker’s utterance. This leads to the augmented Vox-Small with 1,000 speakers, each with 50 utterances. 3) **Vox-Imb dataset:** For imbalanced scenario, Vox-Imb has minority part (target speech set) and majority part (source speech set). Majority part has 1000 speakers, each with 10 utterances as source speech, and minority part has 4,000 speakers, each with only 1 utterance as target speech. VCA generates  $K = 9$  samples for each target speech of minority part. The augmented Vox-Imb contains 5,000 speakers, each with 10 utterances.

#### B. Implementation details

During evaluation, the validation set is Vox1-O and the test sets are Vox1-E and Vox1-H. They are the subsets of VoxCeleb1 dataset, which contains speech utterances from 1,251 speakers [9]. All evaluation sets contain several trials, each has two utterances. The cosine similarity between the output features of the given trial is calculated. The metric is the equal error rate (EER) and minDCF (minC), lower is better. The additive data augmentation methods are used in all experiments to boost training [32], [33]. Diff-VC is trained on the LibriTTS dataset [34]. The baseline in our work is to train the speaker model on the original dataset, we also re-implement some previous methods on our datasets with the same experimental settings for comparison.

### VI. RESULTS AND ANALYSIS

#### A. Results comparison

The baseline in our work is to train the speaker model directly on the original datasets without VCA. In Table I, we report the results of the speaker model trained on three datasets that are augmented with our proposed VCA-RS and VCA-NN strategies. In each scenario, we also re-implement several previous methods for comparison.

1) *Results for semi-supervised learning:* The upper section of Table I illustrates the results for the semi-supervised learning scenario (Vox-Semi), where VCA methods outperform the baseline and the semi-supervised contrastive learning solution [14]. In particular, VCA-NN achieves an average improvement of 5.12% in EER compared to the baseline for three evaluation sets.

2) *Results for small-scale learning:* The middle part of Table I reports the performance of the Vox-Small dataset. Here, VCA-RS and VCA-NN perform competitively with the previous speaker augmentation method [35]. Due to the limited dataset scale, VCA-NN is not very stable. VCA-RS exhibits

the best results with an average deduction of 17.62% in EER relative to the baseline.

3) *Results for imbalanced learning*: The bottom portion of Table I compares the performance of VCA against weighted-learning method [36] and oversampling-based method [37] for imbalanced learning. Among these, VCA-NN methods achieve 6.99% EER and 0.435 minDCF on the Vox1-O dataset, while baseline only achieves 7.80% EER and 0.467 minDCF.

4) *Summary*: As a unified solution, VCA can effectively enhance the performance of SR model under different scenarios, and the improvements are comparable to those of the previous specific solution. Meanwhile, while there are instances where VCA-RS has better performance, VCA-NN generally outperforms VCA-RS.

TABLE I

THE COMPARISON FOR DIFFERENT AUGMENTATION METHODS UNDER SCENARIOS OF SEMI-SUPERVISED LEARNING, SMALL-SCALE LEARNING, AND IMBALANCED LEARNING.

Vox-Semi						
Method	Vox1-O		Vox1-E		Vox1-H	
	EER	minC	EER	minC	EER	minC
Baseline	8.90	0.529	9.06	0.538	13.81	0.669
[14] <sup>†</sup>	8.50	0.499	8.83	0.525	13.44	0.655
VCA-RS	8.21	0.500	8.70	0.548	13.36	0.687
VCA-NN	8.11	0.476	8.74	0.533	13.38	0.675
Vox-Small						
Method	Vox1-O		Vox1-E		Vox1-H	
	EER	minC	EER	minC	EER	minC
Baseline	11.89	0.624	12.14	0.673	17.83	0.787
[35] <sup>†</sup>	9.85	0.526	10.44	0.596	15.83	0.716
VCA-RS	9.51	0.554	9.91	0.591	15.25	0.726
VCA-NN	9.95	0.562	10.48	0.612	15.43	0.732
Vox-Imb						
Method	Vox1-O		Vox1-E		Vox1-H	
	EER	minC	EER	minC	EER	minC
Baseline	7.80	0.467	8.27	0.491	12.55	0.619
[36] <sup>†</sup>	7.67	0.453	8.00	0.475	12.18	0.597
[37] <sup>†</sup>	7.55	0.478	8.14	0.480	12.31	0.607
VCA-RS	7.13	0.427	7.41	0.446	11.20	0.571
VCA-NN	6.99	0.435	7.32	0.444	11.18	0.567

<sup>†</sup>: our re-implementation on these datasets.

## B. Ablation study

1) *Robutness for different backbones*: Figure 2 investigates the robustness of VCA by analyzing the SR performance of various backbones, including X-Vector [6], ResNet34 [38], [39], CAM++ [5], and RawNet3 [40] trained on the Vox-Imb. VCA-NN reaches the lowest EER across multiple test sets, with VCA-RS following closely. These results indicate that the integration of VCA methods significantly boosts the performance of SR systems for various backbone architectures.

2) *Efficiency of VCA with different number of pseudo utterances*: Table II compares the performance of the VCA-RS and VCA-NN for the different generation coefficients  $K$  for each target speech, where  $K = 0$  denotes the baseline. VCA-NN generally performs better than VCA-RS since it can generate pseudo speech with a high-quality representation. Meanwhile, a small  $K$  will even damage the system in some evaluation sets.




Fig. 2. Performance of baseline, VCA-RS and VCA-NN for different backbone architecture trained on the Vox-Imb dataset.

TABLE II

EER (%) OF VCA-RS AND VCA-NN FOR DIFFERENT NUMBERS OF GENERATED UTTERANCES TRAINED ON THE VOX-SEMI DATASET.  $K = 0$  DENOTES THE BASELINE METHOD.

Evaluation set	Method	Generation coefficient $K$			
		0	3	6	9
Vox1-O	VCA-RS	8.90	8.74	8.64	8.21
	VCA-NN	8.90	8.55	8.24	8.11
Vox1-E	VCA-RS	9.06	9.30	8.89	8.70
	VCA-NN	9.06	9.11	8.62	8.74
Vox1-H	VCA-RS	13.81	14.16	13.61	13.36
	VCA-NN	13.81	14.04	13.51	13.38

## VII. CONCLUSION

In this paper, we propose VCA, a new data augmentation method for defective speaker recognition datasets based on voice conversion. Our experiments prove that VCA can efficiently handle semi-supervised, small-scale, and imbalanced learning scenarios with a unified format. Meanwhile, searching for the source speech with a similar representation to the target speech is an effective strategy. Our VCA provides a new direction to solve the speaker recognition problems from the data aspect. Two points can be further studied for VCA: First, the low-quality pseudo utterances can damage the SR system. Second, due to the speed issue, we have yet to study the feasibility of extending VCA to the large-scale dataset as a general augmentation solution. The proposed technique can also be extended to other areas such as speaker anonymization.

## VIII. ACKNOWLEDGMENT

This research is supported by: Shenzhen Science and Technology Program (Shenzhen Key Laboratory, Grant No. ZDSYS20230626091302006); Shenzhen Science and Technology Research Fund (Fundamental Research Key Project, Grant No. JCYJ20220818103001002); Program for Guangdong Introducing Innovative and Entrepreneurial Teams, Grant No. 2023ZT10X044.

## REFERENCES

- [1] T. Kinnunen and H. Li, "An overview of text-independent speaker recognition: From features to supervectors," *Speech Communication*, vol. 52, pp. 12–40, 2010.
- [2] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn, "Speaker verification using adapted gaussian mixture models," *Digital Signal Processing*, vol. 10, no. 1-3, pp. 19–41, 2000, ISSN: 1051-2004.
- [3] K. A. Lee, O. Sadjadi, H. Li, and D. Reynolds, "Two decades into speaker recognition evaluation - are we there yet?" *Computer Speech & Language*, vol. 61, p. 101 058, 2020.
- [4] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized Channel Attention, propagation and aggregation in TDNN based speaker verification," in *Interspeech*, 2020, pp. 3830–3834.
- [5] H. Wang, S. Zheng, Y. Chen, L. Cheng, and Q. Chen, "CAM++: A fast and efficient network for speaker verification using context-aware masking," *arXiv preprint arXiv:2303.00332*, 2023.
- [6] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-Vectors: Robust DNN embeddings for speaker recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5329–5333.
- [7] T. Liu, R. K. Das, K. A. Lee, and H. Li, "Neural acoustic-phonetic approach for speaker verification with phonetic attention mask," *IEEE Signal Processing Letters (SPL)*, 2022.
- [8] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Interspeech*, 2018, pp. 1086–1090.
- [9] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Interspeech*, 2017, pp. 2616–2620.
- [10] M. McLaren, L. Ferrer, D. Castán, and A. Lawson, "The speakers in the wild (SITW) speaker recognition database," in *Interspeech*, 2016, pp. 818–822.
- [11] T. Liu, K. A. Lee, Q. Wang, and H. Li, "Disentangling voice and content with self-supervision for speaker recognition," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [12] R. Tao, K. A. Lee, R. K. Das, V. Hautamäki, and H. Li, "Self-supervised speaker recognition with loss-gated learning," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6142–6146.
- [13] B. Han, Z. Chen, and Y. Qian, "Self-supervised learning with cluster-aware-dino for high-performance robust speaker verification," *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*, vol. 32, pp. 529–541, 2023.
- [14] N. Inoue and K. Goto, "Semi-supervised contrastive learning with generalized contrastive loss and its application to speaker recognition," in *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2020, pp. 1641–1646.
- [15] R. Li, J.-Y. Jiang, J. Liu, C.-C. Hsieh, and W. Wang, "Automatic speaker recognition with limited data," in *Proceedings of the 13th International Conference on Web Search and Data Mining*, 2020, pp. 340–348.
- [16] V. Karthikeyan and S. Suja Priyadharsini, "A strong hybrid adaboost classification algorithm for speaker recognition.," *Sādhanā: Academy Proceedings in Engineering Sciences*, vol. 46, no. 3, 2021.
- [17] S. Zheng, G. Liu, H. Suo, and Y. Lei, "Autoencoder-Based Semi-Supervised Curriculum Learning for Out-of-Domain Speaker Verification," in *Proc. Interspeech*, 2019, pp. 4360–4364. DOI: 10.21437/Interspeech.2019-1440.
- [18] Y. A. Wubet and K.-Y. Lian, "Voice conversion based augmentation and a hybrid CNN-LSTM model for improving speaker-independent keyword recognition on limited datasets," *IEEE Access*, vol. 10, pp. 89 170–89 180, 2022.
- [19] H.-R. Hu, Y. Song, J.-T. Zhang, *et al.*, "Stargan-vc based cross-domain data augmentation for speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [20] Y. Huang, Y. Chen, J. Pelecanos, and Q. Wang, "Synth2Aug: Cross-Domain Speaker Recognition with TTS Synthesized Speech," in *2021 IEEE Spoken Language Technology Workshop (SLT)*, IEEE, 2021, pp. 316–322.
- [21] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov, "Grad-TTS: A diffusion probabilistic model for text-to-speech," in *International Conference on Machine Learning (ICML)*, 2021, pp. 8599–8608.
- [22] H. Huang, X. Xiang, F. Zhao, S. Wang, and Y. Qian, "Unit selection synthesis based data augmentation for fixed phrase speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5849–5853.
- [23] H. Song, M. Kim, D. Park, Y. Shin, and J.-G. Lee, "Learning from noisy labels with deep neural networks: A survey," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 11, pp. 8135–8153, 2023.
- [24] N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari, "Learning with noisy labels," in *Advances in Neural Information Processing Systems*, vol. 26, 2013.
- [25] B. Sisman, J. Yamagishi, S. King, and H. Li, "An overview of voice conversion and its challenges: From statistical modeling to deep learning," *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*, vol. 29, pp. 132–157, 2020.
- [26] S. H. Mohammadi and A. Kain, "An overview of voice conversion systems," *Speech Communication*, vol. 88, pp. 65–82, 2017.

- [27] Y. Zhou, X. Tian, H. Xu, R. K. Das, and H. Li, “Cross-lingual voice conversion with bilingual phonetic posteriorgram and average modeling,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6790–6794.
- [28] X. Qin, Y. Yang, Y. Shi, *et al.*, “Vc-aug: Voice conversion based data augmentation for text-dependent speaker verification,” in *National Conference on Man-Machine Speech Communication*, Springer, 2022, pp. 227–237.
- [29] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, M. S. Kudinov, and J. Wei, “Diffusion-based voice conversion with fast maximum likelihood sampling scheme,” in *International Conference on Learning Representations (ICLR)*, 2021.
- [30] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [31] P. E. Kloeden, E. Platen, P. E. Kloeden, and E. Platen, *Stochastic differential equations*. Springer, 1992.
- [32] D. Snyder, G. Chen, and D. Povey, “MUSAN: A music, speech, and noise corpus,” *CoRR*, vol. abs/1510.08484, 2015. [Online]. Available: <http://arxiv.org/abs/1510.08484>.
- [33] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 5220–5224.
- [34] H. Zen, V. Dang, R. Clark, *et al.*, “LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech,” in *Proc. Interspeech*, 2019, pp. 1526–1530. DOI: 10.21437/Interspeech.2019-2441.
- [35] H. Yamamoto, K. A. Lee, K. Okabe, and T. Koshinaka, “Speaker Augmentation and Bandwidth Extension for Deep Speaker Embedding,” in *Proc. Interspeech*, 2019, pp. 406–410. DOI: 10.21437/Interspeech.2019-1508.
- [36] S. Barua, M. M. Islam, X. Yao, and K. Murase, “Mwmote—majority weighted minority oversampling technique for imbalanced data set learning,” *IEEE Transactions on knowledge and data engineering*, vol. 26, no. 2, pp. 405–425, 2012.
- [37] S. Kotsiantis, D. Kanellopoulos, P. Pintelas, *et al.*, “Handling imbalanced datasets: A review,” *GESTS international transactions on computer science and engineering*, vol. 30, no. 1, pp. 25–36, 2006.
- [38] 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 (CVPR)*, 2016, pp. 770–778.
- [39] J. S. Chung, J. Huh, S. Mun, *et al.*, “In defence of metric learning for speaker recognition,” in *Interspeech*, 2020, pp. 2977–2981.
- [40] J.-w. Jung, Y. Kim, H.-S. Heo, B.-J. Lee, Y. Kwon, and J. S. Chung, “Pushing the limits of raw waveform speaker recognition,” in *Proc. Interspeech*, 2022, pp. 2228–2232.