



A Comprehensive Investigation on Speaker Augmentation for Speaker Recognition

Zhenyu Zhou^{1,2}, Shibiao Xu¹, Shi Yin³, Lantian Li¹, Dong Wang²

¹School of Artificial Intelligence, Beijing University of Posts and Telecommunications, China

²Center for Speech and Language Technologies, BNRist, Tsinghua University, China

³Huawei Technologies Co., Ltd., China

Corresponding authors: lilt@bupt.edu.cn, Wangdong99@mails.tsinghua.edu.cn

Abstract

Data augmentation (DA) has played a pivotal role in the success of deep speaker recognition. Current DA techniques primarily focus on speaker-preserving augmentation, which does not change the speaker trait of the speech and does not create new speakers. Recent research has shed light on the potential of speaker augmentation, which generates new speakers to enrich the training dataset. In this study, we delve into two speaker augmentation approaches: speed perturbation (SP) and vocal tract length perturbation (VTLP). Despite the empirical utilization of both methods, a comprehensive investigation into their efficacy is lacking. Our study, conducted using two public datasets, VoxCeleb and CN-Celeb, revealed that both SP and VTLP are proficient at generating new speakers, leading to significant performance improvements in speaker recognition. Furthermore, they exhibit distinct properties in sensitivity to perturbation factors and data complexity, hinting at the potential benefits of their fusion. Our research underscores the substantial potential of speaker augmentation, highlighting the importance of in-depth exploration and analysis.

Index Terms: speaker augmentation, data augmentation, speaker recognition

1. Introduction

In recent years, data augmentation (DA) techniques have exhibited profound efficacy in the realm of deep neural network-based speaker recognition [1, 2, 3, 4, 5, 6, 7], establishing themselves as indispensable for achieving good performance in numerous challenges and evaluations [8, 9, 10, 11]. The mainstream DA techniques involve manipulating speech samples by noise corruption [2, 12] or spectrum masking [13, 14]. These methods generate new training samples while preserving their speaker traits. The core idea of this type of DA methods is to increase the intra-speaker variations, enabling deep neural networks (DNNs) to extract more robust speaker-invariant patterns.

In contrast to speaker-preserving DA, recent studies have demonstrated the effectiveness of ‘*speaker-expansion DA*’ techniques, or speaker augmentation. Typical methods include speed perturbation (SP) [10, 15, 16] and vocal tract length perturbation (VTLP) [17]. The purpose of this type of DA methods is to create new pseudo speakers, thereby encouraging the model to establish a better speaker embedding space and gain better discriminative capacity among speakers [18]. Compared to enriching utterances for each speaker, creating new speakers seems more crucial. This is because the number of speakers within a typical speech database is far from large. For instance,

the popular VoxCeleb database contains less than 8,000 speakers. Considering the high dimensionality of typical speaker embeddings, for example, 256 in the x-vector paradigm [2], establishing a generalizable speaker embedding space with thousands of speakers is difficult. This is the reason why SP has yielded significant performance improvement in several benchmark evaluations [10, 16].

Despite the remarkable contribution to the performance, it is surprising that there has not been a comprehensive study on speaker augmentation in the speaker recognition community. Their usage is largely motivated by the success of these methods in speech recognition [19, 20, 21], lacking a deeper understanding of how they work and in which condition they work. Moreover, SP is the only speaker augmentation method for a long time, with VTLP being picked up very recently [17]. Finally, while signal processing researchers have developed numerous tools for creating speech with diverse speaker characteristics, these tools have not yet been used for speaker augmentation. Overall, speaker augmentation appears to showcase a highly promising research direction that has garnered limited attention from researchers.

This paper takes one step further towards this promising direction. We investigate two speaker augmentation methods: time-domain SP [19] and frequency-domain VTLP [20]. We conducted comprehensive studies on these two methods using two public datasets: VoxCeleb [22] and CN-Celeb [23]. The experimental results show that both SP and VTLP are effective, though SP shows marginal superiority. Secondly, the two methods show different sensitivity to the perturbation factor and data complexity, leading to different behaviours when they are employed to perform speaker augmentation. Third, a combination of SP and VTLP may lead to further performance gains with a carefully tuned fusion strategy. To the best of our knowledge, this is the first effort to systematically investigate speaker augmentation methods within the realm of speaker recognition.

2. Related Work

Speaker augmentation, with the idea of creating pseudo speakers, perhaps originated from Jaitly and Hinton [20]. They proposed VTLP as an accompanying technique for the new generation of speech recognition systems based on deep neural nets. Intuitively, VTLP can be regarded as a ‘reverse operation’ of Vocal Tract Length Normalization (VTLN) [24], a popular approach to removing speaker variations. This normalization was crucial in the statistical model era, but in the deep learning era, neural nets are sufficiently powerful to deal with any speaker variation. Following this insight, VTLP was designed to create new speakers to enrich the training data so the model could actively learn and handle speaker variation. Inspired by the

This work was supported by the National Natural Science Foundation of China (NSFC) under Grants No.62301075/62171250.

success of VTLP, Tom et al. [19] introduced SP as an alternative speaker augmentation method and demonstrated significant performance improvements with this method in speech recognition. Since then, SP and VTLP have been widely adopted in various speech processing tasks, including low-resource speech recognition [20, 21, 25, 26], speech synthesis [27], speech separation [28], and speaker recognition [29, 10, 16, 17].

In the realm of speaker recognition, SP has been widely acknowledged as a pivotal technique in numerous evaluations [10, 16]. Nevertheless, the exploration of VTLP remains limited. Recently, Tomoka et al. [17] experimented on VTLP-based speaker augmentation. Their findings revealed that solely applying VTLP did not lead to performance enhancements. However, when coupled with a selection process to choose ‘good’ generations and combined with noise augmentation, some improvements were observed. However, their study was conducted using a non-standard dataset, and a comparison with SP was absent.

3. Review for SP and VTLP

3.1. Speed Perturbation (SP)

Speed perturbation operates by resampling the speech signal within the time domain. For a given speech signal $x(t)$, a perturbation factor α is applied to resample the signal along the time axis, yielding the output $y(t)$:

$$y(t) = x(\alpha t). \quad (1)$$

This operation in the time domain corresponds to a specific transformation in the frequency domain, described by the following change:

$$X(f) \rightarrow \frac{1}{\alpha} X\left(\frac{1}{\alpha} f\right), \quad (2)$$

where $X(f)$ and $\frac{1}{\alpha} X\left(\frac{1}{\alpha} f\right)$ represent the Fourier transform of $x(t)$ and $y(t)$, respectively. A key observation is that SP not only changes the duration of the speech signal but also modifies the spectrum by stretching or compressing the spectral range. Two important changes caused by the spectral modification are: (1) the fundamental frequency (F0) is increased (downsampling) or decreased (upsampling); (2) the spectrum envelope is stretched (downsampling) or compressed (upsampling), leading to the change of the location and width of the formants (F1, F2, ...). All these alterations lead to deviations in speaker characteristics, akin to creating a new speaker.

3.2. Vocal Tract Length Perturbation (VTLP)

The length of the vocal tract varies across different speakers, leading to unique traits for each individual [30]. The core idea of VTLP is to modify an existing speech signal to simulate the change in its vocal tract length.

For a speech signal $x(t)$, its spectrum $X(f)$ can be obtained through the Fourier transform. VTLP employs a piece-wise linear warping in the frequency domain to approximate the change in the vocal tract length, as defined below [20]:

$$f' = \begin{cases} \alpha f & 0 \leq f \leq f_0, \\ \frac{f_{\max} - \alpha f_0}{f_{\max} - f_0} (f - f_0) + \alpha f_0 & f_0 < f \leq f_{\max}. \end{cases} \quad (3)$$

where α is the perturbation factor, f_{\max} is the maximum sampling frequency and f_0 is a boundary frequency chosen such

that the warping function αf covers significant formats. In this study, f_{\max} is 8,000 and f_0 is set to 4,800.

3.3. SP vs. VTLP

Theoretically, both SP and VTLP can be viewed as warping functions in the frequency domain, albeit with distinct forms of warping functions. Another differentiating factor is that SP alters the speaking rate, introducing additional variations in speaking styles by shortening or elongating the duration of speech patterns. Additionally, a noteworthy characteristic of SP is that when the speed is decreased (i.e., $\alpha < 1$), the energy of the speech signal concentrates in low-frequency regions near zero [19], a feature not observed in VTLP.

Given these distinctions, it is crucial to conduct a careful comparative study to analyze the behaviour of these methods in creating new speakers. Furthermore, exploring whether these two methods can be combined to achieve additional performance enhancements represents an intriguing avenue that has yet to be explored.

4. Experiments

In this section, we present comprehensive experiments to evaluate and compare the effectiveness of SP and VTLP. The objective of these experiments is to examine the specialities of these two methods and provide an in-depth understanding of their practical utility.

4.1. Data

The datasets used in our experiments are VoxCeleb1 [31] and CN-Celeb1 [32]. VoxCeleb1 encompasses a total of 1,251 speakers, including a development set VoxCeleb1.dev for model training and an evaluation set VoxCeleb1.eval for performance assessment. CN-Celeb1 comprises 998 speakers in total, including a development set CN-Celeb1.dev for model training and an evaluation set CN-Celeb1.eval for testing.

It is noteworthy that in our experiments, we did not select the larger VoxCeleb2 and CN-Celeb2 datasets because we wanted to simulate a training condition with a limited number of speakers, thus more clearly exposing the benefits of the speaker augmentation methods. Furthermore, the MUSAN dataset [33] was used to perform speaker-preserving DA. This is intended to construct more varied acoustic conditions.

4.2. Settings

We follow the cnceleb/v2 recipe of the Sunine toolkit¹ to construct the deep speaker model. The structure of the model is ResNet34 with squeeze-and-excitation (SE) layers, accompanied by an attentive statistics pooling layer to produce utterance-level representations. These representations are then transformed by a fully connected layer to obtain a 256-dimensional x-vector. The model is trained by the Adam optimizer with AAM-softmax [34] as the training objective. The simple cosine distance is used to score the trials.

4.3. Deviation Analysis

4.3.1. Principle of speaker augmentation

In various tasks, such as speech recognition, speaker augmentation typically carries no inherent risks. However, this is not the case in the context of speaker recognition. This is because the

¹<https://gitlab.com/csltstu/sunine/>

generated speech will be assigned a new speaker label, and the accuracy of this labelling is paramount to the outcome. Clearly, the accuracy of the labels is related to how the generated speech deviates from the original speech in terms of speaker traits: the larger the deviation, the more accurate the labels.

There are three key criteria for assessing effective speaker augmentation: (1) the generated speech should not be excessively distorted and remain intelligible human speech; (2) the generated speech should exhibit clear differentiation in speaker characteristics from the original speech; (3) speech generated from the same speaker should demonstrate consistency and similarity. These criteria are all linked to the intensity of perturbation, which is controlled by the perturbation factor α for both SP and VTLP. Hence, our focus is on investigating the influence of α on the generated speech using these two methods.

4.3.2. No-distortion range

We initiated a human evaluation by listening to the generated speech. Our findings indicated that for both SP and VTLP, the optimal range for the perturbation factor α is between 0.8 and 1.2 to avoid noticeable distortion. Based on this observation, we have set α to the range of 0.8-1.2 for all subsequent analyses.

4.3.3. Deviation distribution curve

In this step, we assessed the deviation in speaker traits when a speech utterance was perturbed by SP or VTLP. The deviation in speaker traits was quantitatively assessed by calculating the cosine distance between the speaker embeddings of the speech signal before and after the perturbation. The specific steps involved in this evaluation are outlined below.

- For each utterance u_{si} in VoxCeleb1.dev and CN-Celeb1.dev, where s and i index the speaker and utterance respectively, apply SP or VTLP to modify the speech, resulting in a new utterance u'_{si} .
- Utilize a pre-trained speaker embedding model to derive the speaker embeddings e_{si} and e'_{si} for the utterances u_{si} and u'_{si} .
- Compute the cosine distance between each pair of embeddings (e_{si}, e'_{si}) to determine the utterance-level deviation. Aggregate the utterance-level deviations to compute the speaker-level deviation, denoted as $d_s = \sum_i 1 - \cos(e_{si}, e'_{si})$.
- Plot a deviation distribution curve where the x-axis represents the speaker-level deviation and the y-axis illustrates the proportion of speakers experiencing a specific deviation.

Figure 1 presents deviation distribution curves with different settings of α . Each plot corresponds to a specific perturbation approach (SP/VTLP) applied to a particular dataset (VoxCeleb/CN-Celeb). Key observations are as follows.

- For both SP and VTLP, a α value further from 1 results in a greater deviation compared to a α value closer to 1, as evidenced by a rightward shift on the distribution curve. This outcome is to be expected.
- For both SP and VTLP, deviations induced by a specific α form a compact distribution with a variance below 0.01. Furthermore, the distributions corresponding to different α values are distinctly separated. These findings suggest the feasibility of generating additional pseudo speakers by aggregating perturbed speech with varying α values.
- In comparison to VTLP, the impact of SP on deviation is more responsive to the α value, evident in the minimal overlap between the pair of deviation distribution curves corre-

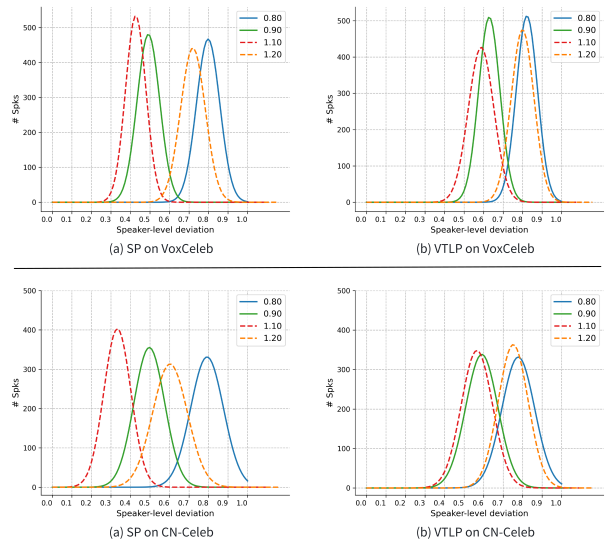


Figure 1: *The deviation distribution curves. Curves with $\alpha > 1$ are plotted in dotted lines, while curves with $\alpha < 1$ are plotted in solid lines.*

sponding to (1.10, 1.20) or (0.80, 0.90).

- A comparison between VoxCeleb and CN-Celeb reveals a greater distribution curve overlap for both SP and VTLP in CN-Celeb. This could be attributed to the more complex acoustic conditions and speaking styles present in CN-Celeb. The behaviour of any perturbation method may vary inconsistently with such complex data, resulting in dispersed deviation distributions. This observation underscores the necessity for meticulously designed perturbations for complex data, and speech produced by different perturbation factors should not be simply regarded as two different persons. It also suggests the need for the development of selection strategies, as discussed in [17].

4.3.4. Deviation-Perturbation curve

We plot a deviation-perturbation curve that reflects the average deviation across all speakers against α values, ranging incrementally from 0.8 to 1.2, as shown in Figure 2. These curves provide a clearer comparison of the influence of α on SP and VTLP.




Figure 2: *The deviation-perturbation curve with (a) SP and (b) VTLP.*

It is evident that as α deviates from 1 in both directions, the curves for VTLP exhibit a noticeable symmetry on both

VoxCeleb and CN-Celeb datasets. This symmetry is less pronounced for SP, particularly on the more complex CN-Celeb dataset, aligning with the findings in Fig. 1. Additionally, the distribution curves for SP appear smoother than those for VTLP, suggesting that VTLP may induce more rapid changes in speaker traits. This factor should be considered when utilizing these methods to generate new speakers, especially when creating multiple speakers from the same individual using different α values.

4.4. Speaker Recognition Results

4.4.1. Results with SP and VTLP

This section presents the speaker recognition results with speaker augmentation by SP and VTLP. Table 1 and Table 2 present the performance with SP and VTLP, respectively, and in each table, results on both VoxCeleb1 and CN-Celeb1 are reported, with various configurations of α .

Table 1: Results in EER(%) with SP-based speaker augmentation.

Method	α	# Spks	EER(%)	
			Vox	CNC
Baseline	-	$\times 1$	2.457	11.067
SP	0.95, 1.05	$\times 3$	2.776	12.171
	0.90, 1.10	$\times 3$	1.760	10.307
	0.85, 1.18	$\times 3$	1.776	9.592
	0.80, 1.20	$\times 3$	1.824	9.327
	0.80, 0.90, 1.10, 1.20	$\times 5$	1.627	9.868

Table 2: Results in EER(%) with VTLP-based speaker augmentation.

Method	α	# Spks	EER(%)	
			Vox	CNC
Baseline	-	$\times 1$	2.457	11.067
VTLP	0.93, 1.07	$\times 3$	2.080	11.276
	0.90, 1.10	$\times 3$	2.010	10.397
	0.83, 1.17	$\times 3$	2.069	10.115
	0.80, 1.20	$\times 3$	2.149	10.234
	0.80, 0.90, 1.10, 1.20	$\times 5$	1.883	9.930

We observe that: (1) With an appropriate setting of α , both SP and VTLP significantly enhance performance compared to the baseline system. This underscores the effectiveness of the two speaker augmentation methods. (2) Comparing SP and VTLP, SP generally outperforms VTLP. This may be attributed to the complex effect of SP on both speed and spectral warping. (3) The choice of perturbation factor α is critical. Setting α close to 1 is risky and may result in a severe performance drop. This is not surprising as this setting may lead to insufficient speaker deviation so the perturbed data should not be regarded as a new speaker. In contrast, setting α to far from 1 is also not optimal, probably due to the potential distortion caused by the strong perturbation. (4) Comparing VoxCeleb and CN-Celeb, VoxCeleb requires a smaller α to reach optimal performance, while CN-Celeb demands a larger α . This observation is consistent with the phenomena illustrated in Figure 1 and Figure 2:

both figures show that the deviation on the speech of VoxCeleb is more sensitive to the change of α .

The key observation is the ability to generate more than two new speakers by varying α values. For example, when producing four new speakers using VTLP with four distinct α values (0.80, 0.90, 1.10, 1.20), improved performance was achieved on both VoxCeleb and CN-Celeb compared to generating only two new speakers. While a similar performance enhancement was observed for SP on VoxCeleb, the same effect was not observed on CN-Celeb. These findings align perfectly with the deviation-perturbation curves depicted in Figure 2, where the VTLP curves exhibit sharper delineations, allowing for segmentation into multiple segments with varying deviations. This segmentation facilitates the creation of diverse new speakers. Conversely, the SP curves are smoother and less conducive to segmentation. Particularly on CN-Celeb, the downward slope on the right side of the LP curve indicates reduced differentiation among speakers with varying α values.

4.4.2. Results of SP + VTLP

We combine SP and VTLP to get more pseudo speakers by choosing the best #Spks \times 3 configuration for SP and VTLP shown in Table 1 and 2 respectively and pool the created data to train the model. The results are reported in Table 3.

Table 3: Results in EER (%) with SP + VTLP.

Method	α	# Spks	EER(%)	
			Vox	CNC
Baseline	-	$\times 1$	2.457	11.067
SP	0.90, 1.10	$\times 3$	1.760	-
VTLP	0.90, 1.10	$\times 3$	2.010	-
SP + VTLP	Fusion	$\times 5$	1.744	-
SP	0.80, 1.20	$\times 3$	-	9.327
VTLP	0.83, 1.17	$\times 3$	-	10.115
SP + VTLP	Fusion	$\times 5$	-	9.665

The results reveal that combining the two perturbation methods leads to further performance improvements on the VoxCeleb dataset, but not as good as the SP perturbation on the CN-Celeb dataset. This suggests that the two augmentation methods exhibit complementarity, but a more careful exploration of the fusion scheme is required.

5. Conclusion

This study delves into a comprehensive analysis of two prominent speaker augmentation techniques, Speed Perturbation (SP) and Vocal Tract Length Perturbation (VTLP), within the realm of speaker recognition tasks. Through extensive analyses conducted on the VoxCeleb1 and CN-Celeb1 datasets, we unveil distinct properties of SP and VTLP. In addition to explaining the speaker recognition outcomes using these augmentation methods, our analysis offers insights into the development of safer and more diverse speaker augmentation techniques, particularly the creation of more distinct speakers from a single speaker. We also observed the feasibility of combining SP and VTLP, highlighting the importance of carefully designing the fusion approach. Future research endeavours will focus on exploring optimal fusion strategies for SP and VTLP, as well as analyzing and employing other speaker augmentation methods, such as various voice morphing and conversion techniques.

6. References

- [1] Z. Bai and X.-L. Zhang, "Speaker recognition based on deep learning: An overview," *Neural Networks*, vol. 140, pp. 65–99, 2021.
- [2] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [3] Y. Zhu, T. Ko, and B. Mak, "Mixup learning strategies for text-independent speaker verification," in *INTERSPEECH*, 2019, pp. 4345–4349.
- [4] C.-L. Huang, "Exploring effective data augmentation with TDNN-LSTM neural network embedding for speaker recognition," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 291–295.
- [5] S. Shahnawazuddin, W. Ahmad, N. Adiga, and A. Kumar, "In-domain and out-of-domain data augmentation to improve children's speaker verification system in limited data scenario," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7554–7558.
- [6] S. Wang, J. Rohdin, O. Plchot, L. Burget, K. Yu, and J. Černocký, "Investigation of specaugment for deep speaker embedding learning," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7139–7143.
- [7] M. Dua, S. Joshi, and S. Dua, "Data augmentation based novel approach to automatic speaker verification system," *e-Prime-Advances in Electrical Engineering, Electronics and Energy*, vol. 6, p. 100346, 2023.
- [8] S. O. Sadjadi, C. Greenberg, E. Singer, L. Mason, and D. Reynolds, "The 2021 NIST speaker recognition evaluation," *arXiv preprint arXiv:2204.10242*, 2022.
- [9] J. Huh, A. Brown, J.-w. Jung, J. S. Chung, A. Nagrani, D. Garcia-Romero, and A. Zisserman, "VoxSRC 2022: The fourth voxceleb speaker recognition challenge," *arXiv preprint arXiv:2302.10248*, 2023.
- [10] Z. Chen, B. Han, X. Xiang, H. Huang, B. Liu, and Y. Qian, "Build a SRE challenge system: Lessons from VoxSRC 2022 and CN-SRC 2022," *arXiv preprint arXiv:2211.00815*, 2022.
- [11] L. Li, T. Jiang, Q. Hong, and D. Wang, "CNSRC 2022 technical report," 2022, <http://cnceleb.org/workshop>.
- [12] H. Tak, M. Kamble, J. Patino, M. Todisco, and N. Evans, "Rawboost: A raw data boosting and augmentation method applied to automatic speaker verification anti-spoofing," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6382–6386.
- [13] 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.
- [14] H. Nam, S.-H. Kim, and Y.-H. Park, "Filteraugment: An acoustic environmental data augmentation method," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 4308–4312.
- [15] 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.
- [16] W. Wang, D. Cai, X. Qin, and M. Li, "The dku-dukeeece systems for voxceleb speaker recognition challenge 2020," *arXiv preprint arXiv:2010.12731*, 2020.
- [17] T. Wakamatsu, S. Shiota, and H. Kiya, "Vocal tract length perturbation-based pseudo-speaker augmentation for speaker embedding learning," in *2023 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2023, pp. 2228–2232.
- [18] D. Snyder, P. Ghahremani, D. Povey, D. Garcia-Romero, Y. Carmiel, and S. Khudanpur, "Deep neural network-based speaker embeddings for end-to-end speaker verification," in *2016 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2016, pp. 165–170.
- [19] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *INTERSPEECH*, 2015.
- [20] N. Jaitly and G. E. Hinton, "Vocal tract length perturbation (VTLP) improves speech recognition," in *Proc. ICML Workshop on Deep Learning for Audio, Speech and Language*, vol. 117, 2013, p. 21.
- [21] X. Cui, V. Goel, and B. Kingsbury, "Data augmentation for deep neural network acoustic modeling," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 9, pp. 1469–1477, 2015.
- [22] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, "VoxCeleb: Large-scale speaker verification in the wild," *Computer Speech & Language*, vol. 60, p. 101027, 2020.
- [23] L. Li, R. Liu, J. Kang, Y. Fan, H. Cui, Y. Cai, R. Vippera, T. F. Zheng, and D. Wang, "CN-Celeb: multi-genre speaker recognition," *Speech Communication*, vol. 137, pp. 77–91, 2022.
- [24] L. Lee and R. Rose, "A frequency warping approach to speaker normalization," *IEEE Transactions on speech and audio processing*, vol. 6, no. 1, pp. 49–60, 1998.
- [25] C. Du and K. Yu, "Speaker augmentation for low resource speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7719–7723.
- [26] M. Geng, X. Xie, S. Liu, J. Yu, S. Hu, X. Liu, and H. Meng, "Investigation of data augmentation techniques for disordered speech recognition," *arXiv preprint arXiv:2201.05562*, 2022.
- [27] E. Cooper, C.-I. Lai, Y. Yasuda, and J. Yamagishi, "Can speaker augmentation improve multi-speaker end-to-end TTS?" *arXiv preprint arXiv:2005.01245*, 2020.
- [28] K. Wang, Y. Yang, H. Huang, Y. Hu, and S. Li, "Spekeraugment: Data augmentation for generalizable source separation via speaker parameter manipulation," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [29] 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.
- [30] H. Wakita, "Direct estimation of the vocal tract shape by inverse filtering of acoustic speech waveforms," *IEEE Transactions on Audio and Electroacoustics*, vol. 21, no. 5, pp. 417–427, 1973.
- [31] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: a large-scale speaker identification dataset," *arXiv preprint arXiv:1706.08612*, 2017.
- [32] 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 *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7604–7608.
- [33] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [34] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4690–4699.