

MIRNet: Learning multiple identities representations in overlapped speech

Hyewon Han, Soo-Whan Chung and Hong-Goo Kang

Department of Electrical & Electronic Engineering, Yonsei University, Seoul, South Korea

[hwhan, jsh6293]@dsp.yonsei.ac.kr, hgkang@yonsei.ac.kr

Abstract

Many approaches can derive information about a single speaker's identity from the speech by learning to recognize consistent characteristics of acoustic parameters. However, it is challenging to determine identity information when there are multiple concurrent speakers in a given signal. In this paper, we propose a novel deep speaker representation strategy that can reliably extract multiple speaker identities from an overlapped speech. We design a network that can extract a high-level embedding that contains information about each speaker's identity from a given mixture. Unlike conventional approaches that need reference acoustic features for training, our proposed algorithm only requires the speaker identity labels of the overlapped speech segments. We demonstrate the effectiveness and usefulness of our algorithm in a speaker verification task and a speech separation system conditioned on the target speaker embeddings obtained through the proposed method.

Index Terms: speaker separation, speaker representation, multi-talker background.

1. Introduction

Speech is one of the most widely used media for recognizing peoples' identities due to its distinctive characteristics. Typically, an individual's vocal identity is represented by acoustic characteristics such as pitch, formant, and speaking styles that are consistently observed in speech signals [1, 2, 3]. In early studies, speaker representations were obtained using statistical models such as Gaussian mixture models [4] with Joint Factor Analysis or Support Vector Machines [5, 6, 7]. By removing rapidly varying acoustic features caused by linguistic changes, these methods obtained normalized speaker-discriminative features with clustering or classification methods [8].

Thanks to advances in deep learning and the availability of large-scale datasets [9, 10], deep neural network-based speaker modeling strategies have recently shown great success in speaker recognition. For example, in [11, 12], neural network models were trained to perform classification given the reference speaker labels, from which representations for speaker information were obtained from the output of the last hidden layer. These methods showed better performance than previous statistical approaches such as i-vectors [7].

Speaker identity information obtained from such methods can be applied to a variety of speech interface related applications. For example, the accuracy of automatic speech recognition (ASR) systems can be improved by using speaker identity information to reduce the bias caused by speaker-dependent characteristics [13, 14]. The performance of tasks such as speech enhancement and separation can also be improved when this information is available, which can be used to specify or represent the desired voice components in input signals distorted by noise, reverberation or interfering speech [15, 16, 17].

However, the aforementioned approaches are often inconvenient for real applications because target speakers must first

enroll their information before post-tasks can be performed. Therefore, they cannot be used with any unenrolled speakers. It is possible to directly extract speaker representations from a mixed input signal, but these representations are prone to error; thus, the overall performance of the post-task can degrade as a result of incorrect guidance on speaker identity.

In this paper, we propose a novel deep learning-based speaker representation method that extracts multiple identities in the case of simultaneous speech. Our proposed method includes two contributions: a speaker model and its training strategy. First, we construct a speaker model with three modules: *speech analysis*, *spectral attention*, and *speaker embedding*. The speech analysis module transforms the input spectrum into a latent domain and the spectral attention module estimates each speaker's spectral information in the latent space using a temporal attention layer, after which the speaker embedding module finds a representation for each speaker's identity. Using this structure, the network is able to directly estimate speaker identities from mixed speech. Second, we propose an effective training strategy for this network that only uses speaker identity information. It utilizes a decision method for speaker permutations to alleviate the ambiguity of speaker assignment. Our approach is novel in that it models speaker identities from the mixed speech with a single network without requiring any clean spectral information. Using this method, we can acquire each speaker's identity information even in cases in which clean utterances are not enrolled.

To verify the effectiveness of our proposed approach, we perform several experiments that utilize the extracted identity information. We first measure speaker verification performance in an overlapped speech scenario. By measuring the similarity between obtained speaker identities from pairs of mixed speech, we show that our network effectively models speaker identity information. We also show that our extracted speaker representations from mixed speech contain distinctive speaker information by applying them to a speaker-conditioned speech separation task.

The rest of the paper is organized as follows. Section 2 describes prior research that is related to our work. The proposed network structure and training strategy for learning speaker identity are shown in Section 3. Experimental results and applications for our proposed method are shown in Section 4, with final conclusions in Section 5.

2. Related Work

2.1. Speaker identity representation

Speaker identity representation is the task of deriving representative embeddings for this information by mapping acoustic features into a latent space. Various deep learning-based approaches have been proposed to improve the performance of speaker representations. D-vectors [11] are a high level speaker representation embedding extracted from the last hidden layer

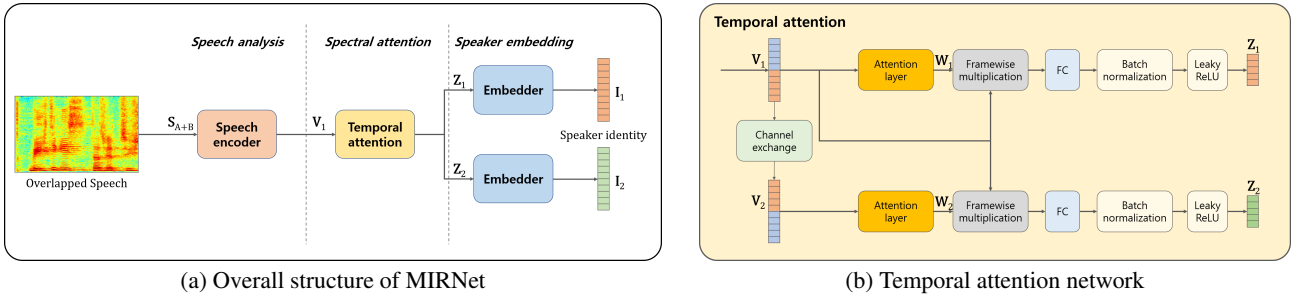


Figure 1: *Demonstration of the proposed speaker representation method.*

of a deep neural network architecture that is designed to perform a speaker identification task. To capture the sequential information of speech signals, x-vectors [12] exploit a time-delay neural network (TDNN) architecture that consists of frame-wise outputs and statistical pooling at the sentence level. Also, various methods are proposed to model effective speaker representation by modifying the aggregation method [18, 19].

Various training criteria have been investigated to further improve the uniformity of output speaker representations. The most well-known loss function is cross-entropy loss, which facilitates the categorization of speaker identity. Recently, metric-based learning criteria such as triplet loss [20, 21] and prototypical loss [22, 23] have been used to enlarge the similarity between same speaker pairs but minimize the similarity between different speaker pairs. In this work, we apply cross-entropy loss for a learning network to categorize the speaker identities.

2.2. Permutation invariant training

Speech separation is the task of estimating individual signals for each speaker from a mixed speech signal. One of the most difficult issues in this work is how to avoid the speaker permutation problem, *i.e.* how to correctly assign the speaker ID of the separated signal in each processing frame. To alleviate this problem, the permutation invariant training (PIT) criterion was proposed [24, 25], which simply considers all of the loss terms by calculating every possible permutation of candidate pairs. From those candidates, the network determines optimal speaker pairs that minimize the error. Inspired by this approach, our training criterion also assigns a loss function for all possible pairs and finds the optimal assignment. However, we compute the PIT loss on estimated speaker identity information such as embeddings or the distribution of speaker classes, not the separated spectral bins as the conventional method did.

3. Multiple identities representation network (MIRNet)

Although speaker information is mingled in an overlapped speech signal, it is possible to extract a target speaker’s identity if an appropriate target-related condition vector is given [26]. Conventionally, a reference speech signal is used as a condition vector, which necessitates a cumbersome pre-enrollment process. In this work, we separate speaker identity information in an overlapped signal without providing any reference signal. Figure 1 illustrates a block diagram of the proposed multiple identities representation network (MIRNet). Although the MIRNet can be generalized to an arbitrary number of speaker inputs, we fix the number of speakers to two for simplicity in this paper. The proposed network consists of three stages: speech analysis, spectral attention and speaker embedding.

3.1. Speech analysis

In the speech analysis stage, an input speech signal is transformed into a latent domain representation, \mathbf{V}_1 .

$$\mathbf{V}_1 = E_S(S_{A+B}), \quad (1)$$

where E_S is a speech encoder network and \mathbf{V}_1 is an embedding with $2D$ -channels. The input to the encoder, S_{A+B} , denotes a linear magnitude spectrum on a logarithm scale. We do not use the mel-spectrum that is popularly used in speaker embedding tasks because of its over-smoothed spectral characteristics, which makes it difficult to distinguish one speaker from another. We construct the architecture of the speech encoder based on 1-D convolution layers. The detailed parameter settings of the speech encoder are described in Table 1. The speech encoder emits spectral embeddings with $2D$ channels, where we assume that each D channel contains the information for one speaker.

3.2. Spectral attention

In the spectral attention stage, we extract two different sets of embeddings from the spectral embedding output of the speech encoder using an attention mechanism. The rationale for the use of an attention mechanism is as follows. Some frames that are fully-overlapped with two speakers do not represent speakers’ identities well, but other frames clearly represent each speaker’s identity when they are spoken by only a single speaker. In addition, some frames contain only silence. We compute frame-wise attention weights from the spectral embeddings to estimate the amount of importance that each frame implies.

Figure 1 illustrates a detailed block diagram of the self-attention mechanism [27]. We first obtain two embedding vectors, one from the spectral embedding \mathbf{V}_1 and the other from \mathbf{V}_2 , re-structured by a channel exchange. Since the attention layer shares parameters for attention vectors, half of the channels of \mathbf{V}_1 are flipped onto the channel axis to create a different input embedding \mathbf{V}_2 , which is used to extract another attention vector W_2 .

$$\mathbf{V}_2 = \text{Concat}(\mathbf{V}_{1,D+1:2D}, \mathbf{V}_{1,1:D}). \quad (2)$$

The attention vectors are related to the framewise power of each speaker, which refers to the speaker information in each frame. They are each multiplied with the spectral embedding \mathbf{V}_1 and used as inputs to the following fully-connected layer with a non-linear activation function. From the spectral attentive stage, the network outputs \mathbf{Z}_1 and \mathbf{Z}_2 where they contain discriminative information of each speaker.

Figure 2 illustrates an example log magnitude spectrogram of mixed speech from two speakers and those for each speaker individually, as well as power contours and attention weights estimated by the self-attention mechanism.

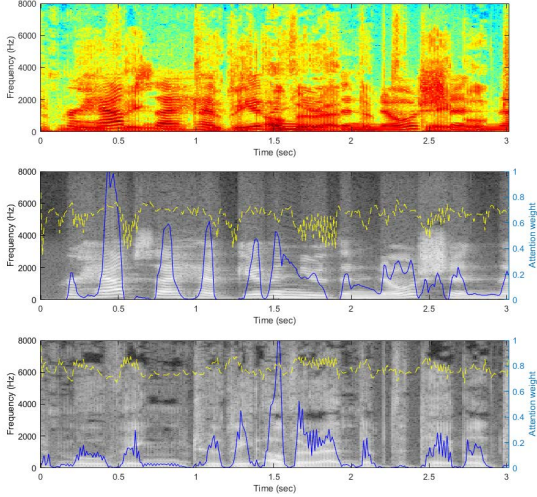


Figure 2: Visualization of analysis on temporal attention for each speaker in a mixture. **Top:** Spectrogram of input speech signal S_{A+B} . **Mid:** Spectrogram of speech S_A . **Bottom:** Spectrogram of speech S_B . Yellow line shows attention weight values, and the blue line is the power contour of speech signal.

3.3. Speaker embedding

The speaker embedding network extracts speaker vectors from the outputs of the spectral attention stage. The network architecture is similar to the one used for speaker verification, where we use ResNet-18 [28] as a backbone network while changing the pooling strategy to be temporal average pooling (TAP).

3.4. Training strategy

For training, we randomly mix signals from two different speakers. The networks are jointly trained using cross-entropy loss with a classifier for the speaker embeddings I_1 and I_2 . The permutation invariant training method is used to solve the permutation problem frequently occurring in speech separation tasks. The entire training criterion is as follows:

$$\mathcal{L} = \min(\mathcal{L}_1, \mathcal{L}_2), \quad (3)$$

$$\mathcal{L}_1 = \mathcal{L}_{CE}(y_1, y_A) + \mathcal{L}_{CE}(y_2, y_B), \quad (4)$$

$$\mathcal{L}_2 = \mathcal{L}_{CE}(y_2, y_A) + \mathcal{L}_{CE}(y_1, y_B), \quad (5)$$

where \mathcal{L}_{CE} is cross-entropy loss.

Suppose that we know the labels of speaker y_A, y_B when we generate the mixture S_{A+B} , and the classifier estimates speaker identities (y_1, y_2) . Then, the losses between the labels y_A, y_B and y_1, y_2 are computed for every pair to find the optimal speaker pairs.

4. Experiments

In this section, we describe two experiments on a speaker verification task and a speaker-conditioned speech separation task to prove the effectiveness of our method.

4.1. Speaker verification

4.1.1. Dataset and neural network settings

The speaker representation network was trained with the LibriSpeech corpus [29], which contains 2,238 speakers. We randomly selected two speech samples from all the different pairs

Table 1: Details of model parameter settings for speech encoder and attention layer

Speech encoder			
Layer	Non-linearity	Channels	Kernel
conv1	Leaky ReLU ($\alpha = 0.2$)	512	5
conv2		512	3
conv3		512	3
conv4		512	1
conv5		1,500	1
conv6		514	1
Attention layer			
Layer	Non-linearity	Channels	
fc1	Tanh	64	
fc2	Sigmoid	1	

of speakers to make mixed input signals in every epoch. In the training stage, 3 second segments chosen at random offsets were used as inputs to the network. In each epoch, the model was trained using 93 hours of mixed signals for training and 15 hours of mixed signals for validation. The input log-scaled spectrum was calculated every 10 ms with an analysis frame length of 32 ms. The FFT size was set to 512; thus, the input dimension was 257. To measure the speaker embedding performance of the network in speaker verification, 200 evaluation pairs were generated for each acceptance and rejection scenario using the *dev-clean*, *dev-others*, *test-clean* and *test-others* subsets which contain a total of 146 speakers.

The detailed network architecture is summarized in Table 1. We use a 257-dimensional log-spectrum as the input, and the output channel dimension of the last layer in the speech encoder is 514, two times that of the input spectrum dimension. For the attention layer, the sigmoid activation function is used to limit the range of the attention weights from 0 to 1. The obtained attention vector is mapped into a 257-dimensional vector using a fully-connected layer.

4.1.2. Evaluation protocol

To prove whether the proposed model represents discriminative characteristics of speaker identities well, we measure the equal error rate (EER) for the speaker verification task using the extracted embeddings from the input mixed speech. While conventional speaker verification methods use positive and negative embedding pairs for evaluation, our method needs a new protocol setup since the output embeddings are permuted and not assigned to speaker labels.

Our evaluation method considers two scenarios to measure EER: acceptance and rejection scenarios. We prepare three mixtures S_{A+B} , $S_{A'+C}$ and $S_{C'+D}$, and each speech segment provides speaker identities (I_A, I_B) , $(I_{A'}, I_C)$ and $(I_{C'}, I_D)$, respectively. While (I_A, I_B) is set as the anchor embedding pair, we use the positive and negative samples from $(I_{A'}, I_C)$ and $(I_{C'}, I_D)$. To compute EER, d_p and d_n are the distances for the acceptance and the rejection, with equations as shown below:

$$d_p = \min_{i,j} (d(I_{i|A+B}, I_{j|A'+C})) \quad (6)$$

$$d_n = \min_{i,j} (d(I_{i|A+B}, I_{j|C'+D})) \quad (7)$$

where $i, j \in \{1, 2\}$. The distances are computed from the closest identity pairs between the anchor and positive and between the anchor and negative samples using Euclidean distance, denoted as $d(\cdot)$. $I_{i|A+B}$ denotes the identity from the i -th output from mixture S_{A+B} .

4.1.3. Experimental results

We obtained an EER of 5.00% and 6.00% with the aforementioned evaluation method for the seen and unseen speaker cases, respectively. From this result, we can see that our trained network successfully represents speaker identities from a mixed input. This result also demonstrates the efficiency of the decision method for avoiding the permutation problem.

Figure 3 depicts a t-SNE visualization of speaker embeddings extracted by the proposed network architecture. The speech samples are mixed with combinations generated from 20 different speakers, and speaker assignment is decided with our proposed permutation decision method. We can see that embeddings from the same speaker are clustered closely together. We also checked the variances of the channels containing speaker identity information to confirm that they effectively model the information without biased channel values. Therefore, we also visualized embeddings extracted from each channel

4.2. Speaker-conditioned speech separation

4.2.1. Task description

Speaker-conditioned speech separation is the task of verifying whether speaker information embeddings are informative enough for speaker separation in overlapped speech signals. In [16], VoiceFilter was proposed to improve speech separation performance using speaker identities extracted from pre-enrolled reference speech. It dramatically improved the separation performance, but it was not applicable to speech signals spoken by speakers that had never been seen before.

Here, we extract multiple identities from overlapped speech first and separate the target speakers' voices from the signal. We use VoiceFilter as our baseline for comparison, which utilizes an independent speaker model. We also compare with a separation system without speaker conditioning using uPIT [25].

4.2.2. Implementation details

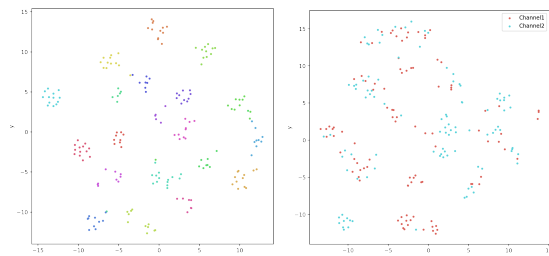
Dataset. The training dataset for our speaker model is identical to the one described in Section 4.1. The speaker model for VoiceFilter is pre-trained using the VoxCeleb2 dataset [10]. We use the WSJ0-2mix dataset [30] to train and evaluate speech separation models for the baseline and our method. Speech signals in WSJ0-2mix are sliced every 10ms with 32ms frame length, from which they are transformed into the time-frequency domain using the Fourier transform.

Model setup. We use a model pre-trained on VoxCeleb2 whose performance is 2.23% EER [31] as the baseline speaker model, which consists of 34 convolution layers with Inception. Model parameters for the speech separation model are the same as those used in [16]. For the model without speaker conditioning, we use a bidirectional LSTM (BLSTM) structure identical to VoiceFilter while doubling the dimension of the fully-connected layer and output for separate speech.

Evaluation metric. For objective evaluation, we measure signal-to-distortion ratio improvement (SDRi), which is typically used to represent speech separation performance. We also calculate the perceptual evaluation of speech quality improve-

Table 2: Performance on the speech separation task

Speaker encoder	SDRi (dB)	PESQi
Chung et al. [31]	6.147	0.480
Proposed	6.097	0.376
uPIT	5.897	0.367



(a) Speaker embedding (b) Embedding of each channel

Figure 3: t-SNE visualization of extracted speaker embedding from 20 random speakers. (a) represents speaker embeddings labeled with our permutation decision method. (b) represents the speaker embeddings from the each channel.

ment (PESQi), which quantifies the perceptual scoring for separated speech signals.

4.2.3. Separation results

Table 2 shows results comparing our method with the baseline. Our separation method results in slightly lower scores than the baseline model. Nevertheless, it is noteworthy that our method is able to achieve reliable performance on this task even with unseen speakers. It should be noted that the baseline uses speaker identities extracted from enrolled reference speech; this not only gives it a significant advantage in the speech separation task, but also means that it cannot be used at all without pre-enrolling speakers. In addition, since MIRNet was trained using a much smaller dataset compared to the baseline model, there should still be margins for it to improve.

5. Conclusion

In this work, we proposed a novel method to estimate latent embeddings from overlapped speech that reliably represent speaker identity information. The proposed network consists of speech analysis, spectral attention, and speaker embedding stages which extract information on multiple speaker identities. To make the network learn this information while finding an optimal assignment, we proposed a speaker identity decision procedure based on permutation invariant training. Experimental results showed that the proposed network can derive individual speakers' identity information from the mixtures without using acoustic information extracted from reference speech signals. In addition, the resulting embeddings showed reliable performance on a speaker-conditioned speech separation task, where it has an advantage in that it can be applied even to speakers for which clean reference speech is unavailable.

Acknowledgements. This research was sponsored by Naver Corporation.

6. References

- [1] B. S. Atal, "Automatic speaker recognition based on pitch contours," *The Journal of the Acoustical Society of America*, vol. 52, no. 6B, pp. 1687–1697, 1972.
- [2] J. J. Wolf, "Efficient acoustic parameters for speaker recognition," *The Journal of the Acoustical Society of America*, vol. 51, no. 6B, pp. 2044–2056, 1972.
- [3] C.-S. Liu, W.-J. Wang, M.-T. Lin, and H.-C. Wang, "Study of line spectrum pair frequencies for speaker recognition," in *International Conference on Acoustics, Speech, and Signal Processing*. IEEE, 1990, pp. 277–280.
- [4] D. A. Reynolds and R. C. Rose, "Robust text-independent speaker identification using gaussian mixture speaker models," *IEEE transactions on speech and audio processing*, vol. 3, no. 1, pp. 72–83, 1995.
- [5] V. Wan and W. M. Campbell, "Support vector machines for speaker verification and identification," in *Neural Networks for Signal Processing X. Proceedings of the 2000 IEEE Signal Processing Society Workshop (Cat. No. 00TH8501)*, vol. 2. IEEE, 2000, pp. 775–784.
- [6] W. M. Campbell, D. E. Sturim, and D. A. Reynolds, "Support vector machines using gmm supervectors for speaker verification," *IEEE signal processing letters*, vol. 13, no. 5, pp. 308–311, 2006.
- [7] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2010.
- [8] F. K. Soong, A. E. Rosenberg, B.-H. Juang, and L. R. Rabiner, "Report: A vector quantization approach to speaker recognition," *AT&T technical journal*, vol. 66, no. 2, pp. 14–26, 1987.
- [9] M. McLaren, L. Ferrer, D. Castan, and A. Lawson, "The speakers in the wild (sitw) speaker recognition database," in *Interspeech*, 2016, pp. 818–822.
- [10] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *arXiv preprint arXiv:1806.05622*, 2018.
- [11] E. Variansi, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez, "Deep neural networks for small footprint text-dependent speaker verification," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 4052–4056.
- [12] 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.
- [13] T. Tan, Y. Qian, D. Yu, S. Kundu, L. Lu, K. C. Sim, X. Xiao, and Y. Zhang, "Speaker-aware training of lstm-rnns for acoustic modelling," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5280–5284.
- [14] G. Pironkov, S. Dupont, and T. Dutoit, "Speaker-aware long short-term memory multi-task learning for speech recognition," in *2016 24th European Signal Processing Conference (EUSIPCO)*. IEEE, 2016, pp. 1911–1915.
- [15] M. Delcroix, K. Zmolikova, K. Kinoshita, A. Ogawa, and T. Nakatani, "Single channel target speaker extraction and recognition with speaker beam," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5554–5558.
- [16] Q. Wang, H. Muckenhirn, K. Wilson, P. Sridhar, Zelin, J. R. Hershey, R. A. Saurous, R. J. Weiss, Y. Jia, and I. L. Moreno, "Voice-Filter: Targeted Voice Separation by Speaker-Conditioned Spectrogram Masking," in *Proc. Interspeech*, 2019, pp. 2728–2732.
- [17] C. Xu, W. Rao, E. S. Chng, and H. Li, "Time-domain speaker extraction network," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 327–334.
- [18] Y. Tang, G. Ding, J. Huang, X. He, and B. Zhou, "Deep speaker embedding learning with multi-level pooling for text-independent speaker verification," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6116–6120.
- [19] Y. Jung, S. M. Kye, Y. Choi, M. Jung, and H. Kim, "Improving multi-scale aggregation using feature pyramid module for robust speaker verification of variable-duration utterances," *arXiv preprint arXiv:2004.03194*, 2020.
- [20] C. Zhang and K. Koishida, "End-to-end text-independent speaker verification with triplet loss on short utterances," in *Interspeech*, 2017, pp. 1487–1491.
- [21] C. Li, X. Ma, B. Jiang, X. Li, X. Zhang, X. Liu, Y. Cao, A. Kannan, and Z. Zhu, "Deep speaker: an end-to-end neural speaker embedding system," *arXiv preprint arXiv:1705.02304*, 2017.
- [22] J. Wang, K.-C. Wang, M. T. Law, F. Rudzicz, and M. Brudno, "Centroid-based deep metric learning for speaker recognition," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3652–3656.
- [23] S. M. Kye, Y. Jung, H. B. Lee, S. J. Hwang, and H. Kim, "Meta-learning for short utterance speaker recognition with imbalance length pairs," *arXiv preprint arXiv:2004.02863*, 2020.
- [24] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 241–245.
- [25] M. Kolbæk, D. Yu, Z.-H. Tan, and J. Jensen, "Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 10, pp. 1901–1913, 2017.
- [26] Y. Shi and T. Hain, "Supervised speaker embedding de-mixing in two-speaker environment," *arXiv preprint arXiv:2001.06397*, 2020.
- [27] Z. Lin, M. Feng, C. N. d. Santos, M. Yu, B. Xiang, B. Zhou, and Y. Bengio, "A structured self-attentive sentence embedding," *arXiv preprint arXiv:1703.03130*, 2017.
- [28] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *Thirty-first AAAI conference on artificial intelligence*, 2017.
- [29] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [30] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 31–35.
- [31] J. S. Chung, J. Huh, S. Mun, M. Lee, H. S. Heo, S. Choe, C. Ham, S. Jung, B.-J. Lee, and I. Han, "In defence of metric learning for speaker recognition," in *Interspeech*, 2020.