



Intra-class variation reduction of speaker representation in disentanglement framework

YooHwan Kwon, Soo-Whan Chung and Hong-Goo Kang

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

yhkwon@dsp.yonsei.ac.kr

Abstract

In this paper, we propose an effective training strategy to extract robust speaker representations from a speech signal. One of the key challenges in speaker recognition tasks is to learn latent representations or embeddings containing solely speaker characteristic information in order to be robust in terms of intra-speaker variations. By modifying the network architecture to generate both speaker-related and speaker-unrelated representations, we exploit a learning criterion which minimizes the mutual information between these disentangled embeddings. We also introduce an identity change loss criterion which utilizes a reconstruction error to different utterances spoken by the same speaker. Since the proposed criteria reduce the variation of speaker characteristics caused by changes in background environment or spoken content, the resulting embeddings of each speaker become more consistent. The effectiveness of the proposed method is demonstrated through two tasks; disentanglement performance, and improvement of speaker recognition accuracy compared to the baseline model on a benchmark dataset, VoxCeleb1. Ablation studies also show the impact of each criterion on overall performance.

Index Terms: speaker verification, disentanglement, mutual information

1. Introduction

Speaker recognition systems have been studied for many years due to their usefulness in various applications. Recently, the accuracy of speaker recognition has dramatically improved due to advances in deep learning and the availability of large-scale datasets for training. The main objective of deep learning-based speaker recognition is to extract a high dimensional embedding vector such that it uniquely represents the characteristic of each speaker. The d-vector [1, 2] and x-vector [3] are typical examples, where they are estimated via an identity classification task with an encoder style network. The detailed extraction process differs with respect to the type of network structure and the criterion of the objective function such as softmax, triplet, and angular softmax [4]. However, given that the extracted embeddings also include speaker-unrelated information, there remains room for further improvement.

To overcome the aforementioned limitation inherent to the encoder style framework, a method for disentangling the embeddings with the use of relevant and irrelevant speaker information was proposed [5]. The method consists of two encoders, a speaker purifying encoder and a dispersing encoder, as well as a decoder for reconstruction. While the speaker purifying encoder is trained by the original speaker classification scheme, the dispersing encoder is trained by an adversarial training scheme designed to fool it from correctly classifying the speaker identity. Later, two encoded features are concatenated, following which they are fed to the decoder, which uti-

lizes a reconstruction loss to the original input so that all information is embedded within the representative features. In other words, they decompose the entirety of the speech information into speaker identity-related and -unrelated information. Although the speaker and non-speaker embeddings are learned effectively using the adversarial classifier, the method does not directly address the task of dispersing both embeddings simultaneously in disentanglement. There is an opportunity to improve the disentanglement performance by adopting a method which considers the relation of embeddings simultaneously.

In this paper, we propose a method to effectively disentangle speaker identity-related and identity-unrelated information using various types of criteria. We first introduce a criterion for minimizing mutual information between speaker-related and -unrelated representations, which is beneficial due to that it directly considers the relation between those features. We also propose a novel identity change criterion which measures the difference between the input and generated mel-spectrums. The reconstructed mel-spectrum used for the identity change loss is generated via a speaker embedding from one utterance and a residual embedding from the other utterance possessing the same speaker identity. Since the criterion enforces speaker embeddings to be similar to a different set of utterances, it reduces intra-variation within each speaker's cluster. The main contributions of this paper are as follows: (1) we propose an effective method for disentangling identity-related and identity-unrelated information using a mutual information criterion through an auto-encoder framework; (2) we introduce a speaker identity change loss criterion to further enhance the performance of speaker embeddings; (3) we use this framework to improve speaker verification performance on benchmark datasets.

The remainder of the paper is organized as follows. Section 3 presents a brief overview of related works on speaker embedding and disentanglement. In Section 4, we present the details of the proposed method such as network architectures and loss functions. Experimental results are presented in Section 5, and the conclusion follows in Section 6.

2. Related works

2.1. Speaker embedding strategy

Speaker embedding vectors are high level representations (typically obtained via deep neural networks) that aim to compactly represent a speaker's identity. They are very important for many applications such as speaker recognition and diarization. There are various speaker embedding methods that differ in terms of the type of network architecture, feature aggregation, and training criteria. Deep learning architectures such as DNN- [1, 6, 7], CNN- [2, 8–11], or LSTM-based ones [12] first extract the frame-level features from a variable length of utterances. Then, a pooling method [13–16] is used to aggregate the frame-level features to a fixed length of utterance-level. In terms of the ob-

jective function, they are trained by performing a classification task with a criterion of softmax, angular softmax or a metric learning task using a contrastive loss [2, 8], triplet loss [9] and etc [17, 18]. Nevertheless, there is still room for improvement if we introduce the concept of target-unrelated information to the extracted embedding features.

2.2. Disentangled feature learning

Disentanglement is a learning technique that represents the input signal’s characteristics through multiple separated dimensions or embeddings. Therefore, it is beneficial for obtaining representations that contain certain attributes or for extracting discriminative features. Adversarial training [19–23] and reconstruction based training [24–28] are widely used to obtain disentangled representations.

Tai et al. [5] proposed a disentanglement method for speaker recognition that is the baseline for our work. By constructing an identity-related and an identity-unrelated encoder, they trained each encoder to represent only speaker-related and -unrelated information using speaker identification loss and adversarial training loss. They also adopted an auto-encoder framework to maintain all input speech information within output embeddings. The information contained in the output embeddings is preserved using spectral reconstruction approaches.

2.3. Mutual Information Neural Estimator

Mutual information (MI) based feature learning methods have been popular for a long time, but they are often difficult to apply for deep learning-based approaches because it is not easy to calculate the MI for high dimensional continuous variables. Recently, a mutual information neural estimator (MINE) [29] was proposed to estimate mutual information with a neural network architecture.

By definition, the MI is equivalent to the Kullback-Leibler (KL) divergence of a joint distribution, $P_{X,Y}$, and the product of marginals, $P_{X \otimes Y}$. According to the Donsker-Varadhan representation [30], the lower bound of mutual information can be represented by:

$$I(X, Y) \geq \sup_T \mathbb{E}_{P_{X,Y}}[T\theta] - \log(\mathbb{E}_{P_{X \otimes Y}}[e^{T\theta}]). \quad (1)$$

The T function is trained by a neural network with the parameter θ , for which the output can be considered to be an approximated value of mutual information between X and Y . It has been widely used in recent works on feature learning [31–33].

3. Proposed Method

The main goal of the proposed algorithm is to extract a high-level latent embedding that contains only speaker-related information. To achieve this goal, we propose a disentanglement method to decouple speaker information from an input signal such that the embedding represents the speaker’s identity being robust to the variation of linguistic information.

3.1. Overview of the proposed algorithm

Figure 1 illustrates the proposed training strategies in our disentanglement method. Our network consists of three modules: a speaker encoder E_{spk} , a residual encoder E_{res} , and a decoder D_r . f_{spk} and f_{res} are respectively the output features of encoders E_{spk} and E_{res} . Our method reconstructs mel-scaled spectrum instead of the magnitude spectrum so that it efficiently disentangles embeddings without speaker information loss.

The network is trained in various learning criteria used in the baseline model, depicted in Figure 1a with auxiliary loss which minimizes intra-variance of clusters; speaker loss, disentanglement loss, reconstruction loss, and our novel criterion – *identity change* loss. Also, we modify disentanglement loss, which uses the adversarial classifier on the residual embedding in the baseline method, into the mutual information between f_{spk} and f_{res} . Details of each criterion are described in the Section 3.2.

3.2. Training Objective

In this section, we demonstrate the details of the proposed method with objective functions for training; speaker loss L_S , disentanglement loss L_{MI} , reconstruction loss L_R and identity change loss L_{IC} . The total objective function of the proposed method consists of four loss functions:

$$L_{total} = \lambda_1 L_S + \lambda_2 L_{MI} + \lambda_3 L_R + \lambda_4 L_{IC}. \quad (2)$$

The hyper-parameters are set based on experimental results, $[\lambda_1, \dots, \lambda_4] = [1, 0.1, 0.1, 0.1]$.

Speaker loss. The objective of the speaker loss is embedding speaker representation f_{spk} into the latent space using the encoder E_{spk} as done in [4, 8, 9, 12]. Following the baseline model, the speaker encoder is trained in a speaker label classification task using a cross-entropy criterion. The loss function is denoted as:

$$L_S = - \sum_{i=1}^C t_i \log(\text{softmax}(f_{spk})_i), \quad (3)$$

where C is the number of speakers and t is the label index.

Disentanglement loss. In the disentanglement mechanism, the residual embedding f_{res} contains information which is not included in the speaker vector f_{spk} . The baseline method adopts the adversarial classification to embed residual of speaker characteristics. The adversarial classification shares network parameters used in speaker loss whereas its objective is to eliminates the speaker information by fooling the classifier. The residual encoder E_{res} is trained not to estimate any speaker label by using uniform distribution, and its definition is as follows:

$$L_{adv} = \frac{1}{C} \sum_{j=1}^C \log(\text{softmax}(f_{res})_j), \quad (4)$$

where C is the number of classes.

In our strategy, we attempt disentanglement using mutual information between f_{spk} and f_{res} instead of adversarial learning. Since the *genuine* disentanglement is achieved in dispersing residual information but not in embedding features separately, we consider both f_{spk} and f_{res} in terms of disentanglement criterion. Here, we adopt the MINE method, which handles correspondence between the three embeddings using deep learning approaches. In [32], MINE controls the information differences between speakers; minimizing in the same speaker and maximizing in different speakers. MINE, in our paper, maximizes the discrepancy between disentangled features (f_{spk} , f_{res}), and minimizes between speaker representations extracted from different segments of the same speech signals as shown in Figure 1c. The criterion is designed as Equation 5.

$$L_{MI} = \mathbb{E}[T_\theta(f_{spk}^A, f_{spk}^{A'})] - \log\left(\mathbb{E}\left[e^{T_\theta(f_{spk}^A, f_{res}^A)}\right]\right) \\ + \mathbb{E}[T_\theta(f_{spk}^{A'}, f_{spk}^A)] - \log\left(\mathbb{E}\left[e^{T_\theta(f_{spk}^{A'}, f_{res}^A)}\right]\right), \quad (5)$$

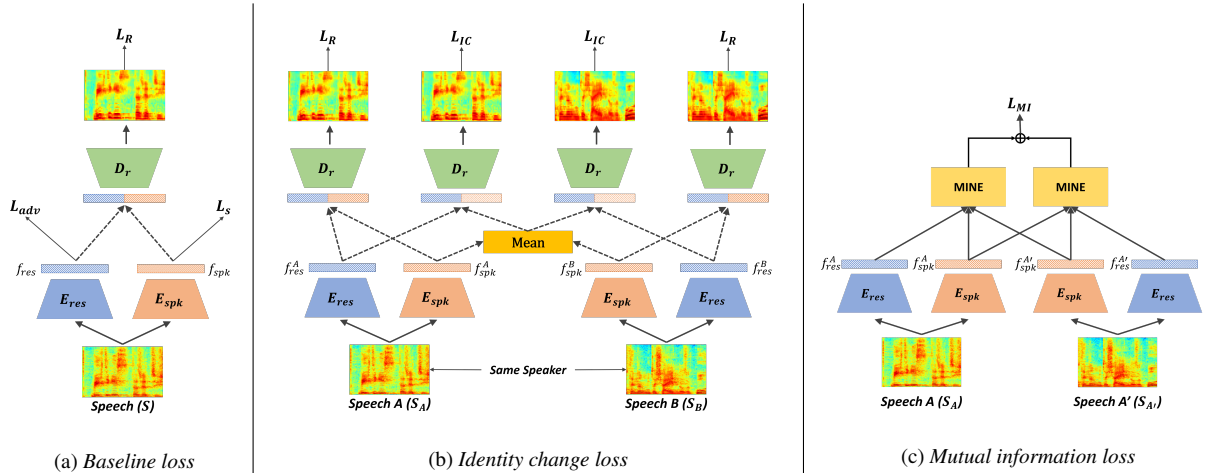


Figure 1: Overview of proposed training criteria. (a) Training criteria based on [5]: speaker loss, disentanglement loss and reconstruction loss. (b) Identity change loss: switch the speaker embedding to mean of those. (c) Mutual information loss: estimate the mutual information from speaker and residual embeddings by MINE

where f_{spk}^A and $f_{spk}^{A'}$ represent identical speaker extracted from the same speech signal with different offsets, and f_{res}^A and $f_{res}^{A'}$ are their residual embeddings. It holds the common information between speaker embeddings and disperses residuals to speaker embeddings on the other embedding.

Reconstruction loss. The disentangled embeddings, f_{spk} and f_{res} preserve the spectral information in the input spectrum when they are combined. The decoder $D_r(f_{spk}, f_{res})$ is trained to generate a reconstructed spectrum using a concatenated embedding input. The reconstruction loss L_R is defined by measuring the distance between input and the reconstructed spectrum using an MSE criterion as follows:

$$L_R = \|D_r(f_{spk}, f_{res}) - S_{mel}\|^2, \quad (6)$$

where S_{mel} is a mel-spectrum of the input speech signal S . Reconstructing the mel-spectrum instead of a magnitude spectrum can reduce the burden of the decoder during the spectrum generation process, while it still enables the generation of embeddings containing all information of input.

Identity change (IC) loss. Intra-class variance inevitable in each speaker cluster is caused by the variation of linguistic information, recording environments, and speakers' emotional or health state. To further improve speaker recognition performance by minimizing intra-class variances in speaker clusters, we propose identity change loss. Instead of minimizing intra-class variance directly, we use a reconstruction loss criterion that measures spectral distance between the reference and reconstructed one. Since the reconstructed mel-spectrum is generated by substituting the identity embedding with the one extracted from different utterances spoken by the same speaker, we may obtain perfect reconstruction only when the substitute embedding has the same distribution as the original identity. The identity change loss is described in Equation 7.

$$\begin{aligned} L_{IC} &= \|\hat{S}_A - S_A\|^2 + \|\hat{S}_B - S_B\|^2, \\ \hat{S}_A &= D_r\left(\frac{f_{spk}^A + f_{spk}^B}{2}, f_{res}^A\right), \\ \hat{S}_B &= D_r\left(\frac{f_{spk}^A + f_{spk}^B}{2}, f_{res}^B\right), \end{aligned} \quad (7)$$

Table 1: Verification results on VoxCeleb1 test set. S , C and AM are Softmax, Contrastive and Angular margin loss, respectively.

	Model	Criterion	EER
Chung <i>et al.</i> [2]	Encoder	S + C	5.04%
Xie <i>et al.</i> [16]	Encoder	S	5.02%
Tai <i>et al.</i> [5]	Enc(2)+Dec	S	3.83%
Proposed	Enc(2)+Dec	S	3.18%
	Enc(2)+Dec	AM	2.54%

where S_A and S_B are the mel-spectrum of speech signals A , B spoken by the same speaker, and \hat{S}_A and \hat{S}_B are the reconstructed mel-spectrum using substituted identities. In the proposed method, f_{spk}^A and f_{spk}^B are substituted with the mean of two identities depicted in Figure 1b; it guides the direction where speaker embeddings to be gathered to minimize intra-class variance.

4. Experiments

4.1. Dataset configuration

We train our model on VoxCeleb2 [2], which is a large-scale audio-visual dataset containing over 1 million utterances for 5,994 celebrities, extracted from YouTube videos. We evaluate our model on VoxCeleb1 [8] test set which consists of 677 clips spoken by 40 speakers. Clips are segmented into 3 seconds with a random offset from each utterance for training. They are sliced in every 10ms with 25ms window length and transformed into log-magnitude spectrum with the FFT size of 512; thus, the dimension of input speech features is 300×257 . For reconstruction, we prepare mel-spectrogram in logarithm scale using 64 mel-filterbanks as outputs.

4.2. Implementation details

The structures of the speaker encoder and the residual encoder are designed based on ResNet34 with small changes into the pooling strategy. Both encoders use a time average pooling (TAP) method to embed variable length input features into a fixed dimension of utterance level. The decoder consists of

Table 2: Ablation study of the proposed method

	L_s	L_r	L_{adv}	L_{mi}	L_{ic}	EER (%)
Baseline	✓	✓	✓	-	-	3.83%
Proposed	✓	✓	✓	✓	-	3.71%
	✓	✓	-	✓	-	3.81%
	✓	✓	✓	-	✓	3.59%
	✓	✓	-	✓	✓	3.18%

3 fully-connected layers and 9 transposed convolutional layers referenced by [34]. In the training phase, the batch size of the input is set to 32 and the model is trained with the Adam optimizer [35]. The learning rate is set to 1e-3 and reduced by half every 10 epochs until convergence.

4.3. Training strategy

Phase I. Disentanglement training. In phase I, the network is pre-trained using speaker loss, disentanglement loss and reconstruction loss, similar to the baseline strategy. According to each experimental setup, either adversarial loss or mutual information loss is used.

Phase II. Identity change training. During phase II, we consider an efficient training strategy for identity change loss. Its motivation is based on dispersing information by setting one embedding as an anchor and stable adaptation of the other side embedding. The detailed process is shown below and the stages are processed recursively:

1. *Intra-class minimization* – The identity is replaced by the mean of two identities to generate mel-spectrogram, and its reconstruction error L_{IC} is minimized through back-propagation on the decoder and residual encoder.
2. *Adaptation* – The original identity is ingested on the decoder and the parameters of the decoder and the speaker encoder are updated to minimize reconstruction error L_R .

4.4. Experimental results

We compare the performance of our models to that of conventional models and analyze the impact of each loss function on overall performance with an ablation study under the same settings. All models for comparison are re-implemented by ours. Table 1 shows the equal error rate (EER) obtained by the Vox-Celeb1 [8] testset, where we compare our models with the encoder model [16] and the disentanglement model [5]. With the standard softmax loss and TAP aggregation, our model outperforms previous models based on the ResNet encoder by 36.6% and the disentanglement model using an adversarial method [5] by 16.9%. These results demonstrate that the represented embeddings of the proposed disentanglement approach are more informative than those of the baseline. The proposed method trained with angular margin softmax provided our best results among the experiments.

Ablation study. Table 2 shows equal error rates (EERs) obtained by ablation studies, which indicates the effectiveness of loss functions used in the proposed model. First, we trained the model using the mutual information criterion with and without the adversarial criterion. The results confirm that minimizing the mutual information between speaker and residual embeddings is effective to disentangle speaker information. Unlike adversarial training, which is applied to the encoders indepen-

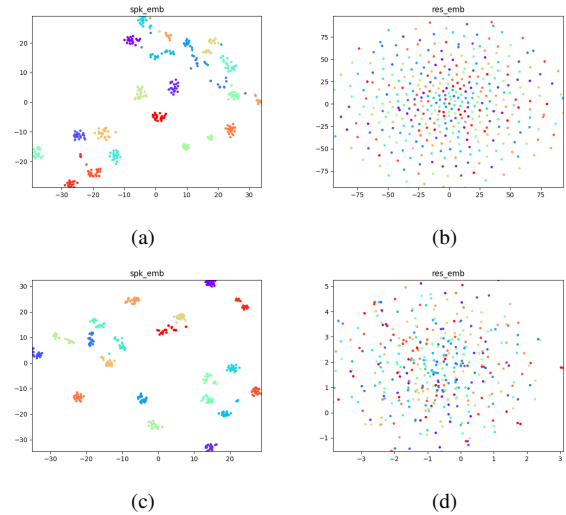


Figure 2: *t*-SNE plot of extracted embeddings: extracted from 10 speaker and 20 utterances each and each color corresponds to a different speaker. (a) and (b) are extracted from baseline model. (c) and (d) are from our proposed model.

ently, mutual information is calculated between speaker and residual embedding simultaneously, resulting in more powerful disentanglement performance. Among these experiments, the case absent adversarial criterion performs better, with an EER 3.81%. Then, the other experiments are conducted in order to investigate the effect of identity change loss. The results prove that identity change loss improves the performance of speaker embedding, and it shows the best result when it is trained using the mutual information and identity change loss criterion together, giving an EER 3.18%.

Figure 2 illustrates *t*-SNE plots [36] for visualization of the effectiveness of the proposed method more concretely. As shown in Figure 2a and Figure 2b, the proposed model also effectively disentangles speaker-related and speaker-unrelated information. Moreover, compared to the baseline with proposed model in Figure 2a and Figure 2c, our method shows more densely clustered identities with small variance.

Through the results of experiments, we proved that mutual information loss and identity change loss is helpful in learning the clearly disentangled features for speaker recognition.

5. Conclusion

In this paper, we present a novel disentanglement training scheme to estimate more informative speaker embedding vectors for robust speaker recognition. Our method is built upon auto-encoder frameworks with two encoders and trained via mutual information and identity change loss, which extracts more discriminative representations by reducing the variance in the intra-cluster. Experimental results demonstrated that our algorithm achieved improved EER compared to the baseline method. Through ablation experiments, we demonstrated the impact of each criterion to the overall performance.

Acknowledgements. This research is sponsored by Naver Corporation.

6. References

- [1] E. Variani, 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.
- [2] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *arXiv preprint arXiv:1806.05622*, 2018.
- [3] 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.
- [4] Z. Huang, S. Wang, and K. Yu, "Angular softmax for short-duration text-independent speaker verification," in *Interspeech*, 2018, pp. 3623–3627.
- [5] J. Tai, X. Jia, Q. Huang, W. Zhang, and S. Zhang, "Sef-ald: A speaker embedding framework via adversarial learning based disentangled representation," *arXiv: Audio and Speech Processing*, 2020.
- [6] G. Heigold, I. Moreno, S. Bengio, and N. Shazeer, "End-to-end text-dependent speaker verification," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5115–5119.
- [7] 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.
- [8] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: a large-scale speaker identification dataset," *arXiv preprint arXiv:1706.08612*, 2017.
- [9] 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.
- [10] M. Hajibabaei and D. Dai, "Unified hypersphere embedding for speaker recognition," *arXiv preprint arXiv:1807.08312*, 2018.
- [11] 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.
- [12] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4879–4883.
- [13] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for text-independent speaker verification," in *Interspeech*, 2017, pp. 999–1003.
- [14] W. Cai, J. Chen, and M. Li, "Exploring the encoding layer and loss function in end-to-end speaker and language recognition system," *arXiv preprint arXiv:1804.05160*, 2018.
- [15] —, "Analysis of length normalization in end-to-end speaker verification system," *arXiv preprint arXiv:1806.03209*, 2018.
- [16] W. Xie, A. Nagrani, J. S. Chung, and A. Zisserman, "Utterance-level aggregation for speaker recognition in the wild," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5791–5795.
- [17] 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," *arXiv preprint arXiv:2003.11982*, 2020.
- [18] 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.
- [19] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [20] J. Zhou, T. Jiang, L. Li, Q. Hong, Z. Wang, and B. Xia, "Training multi-task adversarial network for extracting noise-robust speaker embedding," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6196–6200.
- [21] Z. Meng, Y. Zhao, J. Li, and Y. Gong, "Adversarial speaker verification," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6216–6220.
- [22] X. Peng, Z. Huang, X. Sun, and K. Saenko, "Domain agnostic learning with disentangled representations," *arXiv preprint arXiv:1904.12347*, 2019.
- [23] G. Bhattacharya, J. Monteiro, J. Alam, and P. Kenny, "Generative adversarial speaker embedding networks for domain robust end-to-end speaker verification," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6226–6230.
- [24] J. Zhang, Z. Ling, and L.-R. Dai, "Non-parallel sequence-to-sequence voice conversion with disentangled linguistic and speaker representations," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2019.
- [25] J.-c. Chou, C.-c. Yeh, H.-y. Lee, and L.-s. Lee, "Multi-target voice conversion without parallel data by adversarially learning disentangled audio representations," *arXiv preprint arXiv:1804.02812*, 2018.
- [26] Y. Liu, F. Wei, J. Shao, L. Sheng, J. Yan, and X. Wang, "Exploring disentangled feature representation beyond face identification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2080–2089.
- [27] C. Eom and B. Ham, "Learning disentangled representation for robust person re-identification," in *Advances in Neural Information Processing Systems*, 2019, pp. 5298–5309.
- [28] A. Gonzalez-Garcia, J. Van De Weijer, and Y. Bengio, "Image-to-image translation for cross-domain disentanglement," in *Advances in neural information processing systems*, 2018, pp. 1287–1298.
- [29] M. I. Belghazi, A. Baratin, S. Rajeswar, S. Ozair, Y. Bengio, A. Courville, and R. D. Hjelm, "Mine: mutual information neural estimation," *arXiv preprint arXiv:1801.04062*, 2018.
- [30] M. D. Donsker and S. S. Varadhan, "Asymptotic evaluation of certain markov process expectations for large time. iv," *Communications on Pure and Applied Mathematics*, vol. 36, no. 2, pp. 183–212, 1983.
- [31] R. D. Hjelm, A. Fedorov, S. Lavoie-Marchildon, K. Grewal, P. Bachman, A. Trischler, and Y. Bengio, "Learning deep representations by mutual information estimation and maximization," *arXiv preprint arXiv:1808.06670*, 2018.
- [32] M. Ravanelli and Y. Bengio, "Learning speaker representations with mutual information," *arXiv preprint arXiv:1812.00271*, 2018.
- [33] E. H. Sanchez, M. Serrurier, and M. Ortner, "Learning disentangled representations via mutual information estimation," *arXiv preprint arXiv:1912.03915*, 2019.
- [34] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [36] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.