TRAINING MULTI-TASK ADVERSARIAL NETWORK FOR EXTRACTING NOISE-ROBUST SPEAKER EMBEDDING

Jianfeng Zhou¹, Tao Jiang², Lin Li^{*1}, Qingyang Hong^{*2}, Zhe Wang³, Bingyin Xia³

¹College of Electronic Science and Technology, Xiamen University, China ²School of Information Science and Engineering, Xiamen University, China ³Media Coding Technology Lab, Huawei Media Technology Institute lilin@xmu.edu.cn, qyhong@xmu.edu.cn

ABSTRACT

Under noisy environments, to achieve the robust performance of speaker recognition is still a challenging task. Motivated by the promising performance of multi-task training in a variety of image processing tasks, we explore the potential of multitask adversarial training for learning a noise-robust speaker embedding. In this paper, we present a novel framework that consists of three components: an encoder that extracts the noise-robust speaker embeddings; a classifier that classifies the speakers; a discriminator that discriminates the noise type of the speaker embeddings. Additionally, we propose a training strategy using the training accuracy as an indicator to stabilize the multi-class adversarial optimization process. We conduct our experiments on the English and Mandarin corpuses and the experimental results demonstrate that our proposed multi-task adversarial training method could greatly outperform the other methods without adversarial training in noisy environments. Furthermore, the experiments indicate that our method is also able to improve the speaker verification performance under the clean condition.

Index Terms— multi-task, speaker embedding, adversarial training, speaker verification

1. INTRODUCTION

The task of speaker verification is to verify the identity of speaker from a given speech utterance. In the past decade, the i-vector system has achieved significant success in modeling speaker identity and channel variability in the i-vector space [1], which maps variable-length utterances into fixed-length vectors. Then the fixed-length vectors will be fed to a back-end classifier such as probabilistic linear discriminant analysis (PLDA) [2].

Recently, with the rise of deep learning [3] in various machine learning applications, the works [4, 5, 6] focused on using neural network to verify speakers have explored its potential capability in speaker recognition tasks. More recently, many studies [7, 8, 9] have concentrated on extracting utterance-level representation, which is known as speaker

embedding, using neural networks combined with a pooling layer. This utterance-level representation can be further processed by fully-connected layers.

Since proposed by Goodfellow et al. [10], generative adversarial networks (GAN) have become the focus of many studies in recent years. Its great success in image processing has inspired people to consider whether it can also be applied into the field of speech processing. In the paper [11], Zhang et al. attempted to use conditional GAN to solve the impact of performance degradation caused by the variable-duration of utterances. Ding et al. [12] proposed a multi-tasking GAN framework to extract the more distinctive speaker representation. And Yu et al. [13] proposed to train an adversarial network for front-end denoising.

In the field of speaker recognition, there is a large quantity of literature concerning the sharp degradation of performance in the noisy environments. A common way to improve the robustness of the system is to train the system using a dataset consisting of clean and noisy data [14]. Speech enhancement is another way of denoising such as short-time spectral amplitude minimum mean square error (STSA-MMSE) [15] and many DNN-based enhancement methods [16, 17, 18]. Unlike previous works denoising in the front-end, we plan to use a multi-task adversarial framework to extract the noise-robust speaker representation directly.

In this paper, we borrow the adversarial training idea of GAN [10] and use the multi-task adversarial network (M-TAN) structure to extract a noise-robust speaker embedding. The entire framework consists of three parts: an encoder that extracts the noise-robust speaker embeddings; a classifier that classifies the speaker; a discriminator that discriminates the noise type of the speaker embeddings, which also plays the adversarial role combined with the encoder. In addition, we propose a new loss function, namely AL-Loss (anti-label loss), to realize the multi-class adversarial training. Furthermore, in order to balance the adversarial training process, a new training strategy has been presented by employing the training accuracy as an indicator to judge whether the adversarial training has reached a balance.



Fig. 1. The framework of our proposed multi-task adversarial network.

2. MULTI-TASK ADVERSARIAL NETWORK

2.1. CNN Based Embedding Learning

CNN-based neural network architecture has proved its superior performance in speaker verification tasks [7, 12]. In this work, we use the CNN-based architecture for speaker embedding learning which includes the encoder and classifier of the framework shown in the dotted line of Fig. 1 (a). The details of the architecture are as follow. Four one-dimensional convolutional layers with 1*1 filter, 1 stride and 256 channels followed by an average pooling layer which maps the frame-level feature to an utterance-level representation. Then, the speaker representation will be fed to the next two fullyconnected layers with 256 and 1024 nodes in sequence. Finally, the output layer with N_s (the number of speakers in training data) nodes will take the speaker embeddings as input. The output of last hidden layer is extracted as utterance-level speaker embedding. Besides, batch normalization and RELU activation function are applied to all layers except the output layer. And the verification back-ends are shown in Fig.1 (b).

2.2. Multi-Task Adversarial Network

The entire architecture of MTAN is shown in Fig.1 (a). And the implementation details of the encoder and classifier have been demonstrated in Section 2.1. As to the discriminator, it is just an output layer with M (the number of noise types in training data) nodes. The arrows indicate the forward propagation direction.

Given an input $x \in \mathbb{R}^{t*m}$ where t and m refer to the frame number and acoustic feature dimension of the utterance respectively, the encoder maps it to a speaker embedding $E(x) \in \mathbb{R}^n$, where n is the dimension of latent embedding. Then the classifier and the discriminator try to predict the classes of E(x). Since our goal is to encode speaker information while eliminating performance degradation caused by noise, the encoder should extract a latent representation that is more discriminative for speaker and robust for noise. In order to achieve this goal, we use the multi-task adversarial network to learn discriminative speaker feature and simultaneously improve its noise robustness. Specifically, we train the classifier cooperated with the encoder to extract discriminative speaker feature. Besides, we play a minimax game by training discriminator to maximize the probability of assigning the correct noise label to the embedding extracted from the encoder and simultaneously training the encoder to maximize the probability of assigning the wrong noise label to the embedding.

2.3. Loss Function

In this work we consider cross entropy loss function and its two variants. For the cooperative training of the classifier and encoder, we directly minimize the cross entropy loss l_s^C (the superscript C means classifier). For multi-class adversarial training, the output of the discriminator will be fed to a cross entropy loss function l_s^D (the superscript D means for discriminator) and its variants including FL-Loss (fixed label loss) proposed in [13] and AL-Loss. The details of loss functions will be addressed in Section 2.3.1 and Section 2.3.2. Then a minimax game will be executed with the value function l_{adv} , which can be formulated as follow:

$$\max_{E} \min_{D} l_{adv} = \gamma l_s^D - \beta l_{var} \tag{1}$$

where γ and β are scale parameters and l_{var} could be FL-Loss or AL-Loss. When training an adversarial network, rather than directly using the minimax loss, we split the optimization into two independent objectives, one for encoder and one for discriminator. Therefore, we train the encoder, discriminator and classifier by $\min_{E}(l_s^C + \beta l_{var})$, $\min_{D} \gamma l_s^D$ and $\min_{C} l_s^C$ respectively.

2.3.1. FL-Loss

Compared with the cross entropy loss function, FL-Loss uses the fixed label "clean speech" [13] for all inputs to train the encoder. It can be formulated as follow:

$$l_{fl} = -\frac{1}{N} \sum_{i=1}^{N} log \frac{e^{W_{y_c}^T y_i + b_i}}{\sum_{j=1}^{M} e^{W_j^T y_j + b_j}}$$
(2)

where N is the training batch size, y_c is the label of clean speech and y_i is the output of discriminator corresponding to y_c . Besides, W and b are the weights and biases of the output layer. By assigning all data to clean speech label, the embedding from noisy speech will be close to the embedding from clean speech, since the constraint of FL-Loss will regularize the encoder to learn a map function from noisy data distribution to clean data distribution.

2.3.2. AL-Loss

Inspired by the FL-Loss function, we propose the AL-Loss function combined with the cross entropy loss function for the multi-class adversarial task, which is formulated as follow:

$$l_{al} = -\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1, j \neq m_c}^{M} \log \frac{e^{W_j^T y_j + b_j}}{\sum_{k=1}^{M} e^{W_k^T y_k + b_k}}$$
(3)

where m_c is the corresponding ground truth label of the *i*th sample. Unlike FL-Loss, we use the anti-label to calculate the loss value, where the anti-label means flipping the value of each bit in one hot vector of the ground truth label. $\min_E l_{anti}$ means that the encoder would be trained to assign the output of encoder to a wrong noise label equally, i.e., after adversarial training, the embedding extracted from encoder will be invariant to the clean and noisy speech.

3. EXPERIMENTS

3.1. Dataset and Experimental Setting

To evaluate the effective performance of the proposed framework under the noisy environments, text-independent speaker verification (SV) experiments were conducted based on Aishell-1 [19] (a Mandarin corpus) and Librispeech [20] (an English corpus). The details of the two datasets are given as follows:

- *Aishell-1*: We use the data of all three sets of Aishell-1 as the training data which contains about 141,600 utterances from 400 speakers and use another corpus named King-ASR-L-057¹ as the test data which contains 6,167 recordings from 20 speakers.
- Librispeech: In our experiments, we use the trainclean-500 part of Librispeech as training data which

contains about 148,688 utterances from 1,166 speakers and the test-clean part as test data, which includes 2,020 recordings from 40 speakers.

We have made a noise corrupted version of the training data mentioned above by artificially adding different types of noise at different SNR levels. The original training data was divided into two parts with scale of 1:5, in which five out of six samples were added by the random noise. Specically, the noisy utterances for training were made by adding one of the five noise types (white, babble, mensa, cafeteria, callcener)² randomly on the SNR levels of 10dB or 20dB. However, the noisy utterances for the speaker verification test were obtained by adding one of the five noise types on the SNR levels of 0dB, 5dB, 10dB, 15dB and 20dB respectively.

All audios were converted to the features of 23-dimensional MFCC with a frame-length of 25 ms and the frame shift of 10 ms. Then, a frame-level energy-based voice activity detector (VAD) selection was conducted to the features.

Our implementation was based on the Tensorflow toolkit. In our experiments, Adam optimizer with a learning rate of 0.01 was used for the back propagation. We alternate between one step of optimizing the classifier and discriminator, and three steps of optimizing the encoder.

3.2. Training Stability

In this work, we use the training accuracy as an indicator to balance multi-class adversarial training. Specifically, we train the encoder to maximize the probability of assigning a speaker embedding to a wrong noise label, which means decreasing the training accuracy. However, we also train the discriminator to correctly assign the embedding to the ground truth label, which means increasing the training accuracy. So the accuracy could indicate the situation of adversarial training. The training accuracy keeping in high or low all means adversarial training doesn't get a balance. In addition, we set a lower threshold α and an upper threshold θ . When the average of the training accuracy of the latest K iterations is less than the lower threshold or higher than the upper threshold, we adjust the loss proportional factor of βl_{var} and γl_s during the training. In our experiments, the encoder is trained better than discriminator, so we just set a lower threshold ($\alpha = 0.4$) to balance the adversarial training process.

3.3. Results and Comparisons

In order to evaluate the performance of our proposed multitask adversarial network, five systems were investigated: the CNN-based architecture trained using clean data (Baseline); the CNN-based architecture trained using the noise corrupted

¹King-ASR-L-057: A Chinese Mandarin speech recognition database, which is available at http://kingline.speechocean.com

²white and babble were collected by Guoning Hu, and could be downloaded at http://web.cse.ohio-state.edu/pnl. Besides, cafeteria noise, callcener, and mensa were provided by HUAWEI TECHNOLOGIES CO., LTD.

NOISE	SNR	Baseline	MIX	FL	AL	Fusion
Clean	-	6.49	7.08	5.54	5.89	5.15
White	00	39.95	30.74	30.30	30.64	27.77
	05	38.42	21.68	18.91	19.36	16.39
	10	35.69	15.25	12.23	13.07	10.35
	15	29.50	12.23	9.90	10.35	8.71
	20	24.26	10.89	8.86	9.46	7.77
	mean	33.56	18.16	16.04	16.58	14.20
Babble	00	30.74	20.05	20.00	18.71	17.72
	05	25.05	12.72	11.09	19.36	10.30
	10	19.46	10.00	8.07	13.07	7.77
	15	14.41	8.91	7.53	10.35	6.93
	20	11.09	8.07	6.49	9.46	6.09
	mean	20.10	11.95	10.64	10.50	9.76
	00	32.52	19.80	20.30	18.91	17.18
Cafeteria	05	26.73	14.36	12.03	12.72	10.74
	10	21.24	10.99	9.26	9.41	8.27
	15	16.14	8.91	7.48	7.62	6.83
	20	12.03	8.37	6.24	6.93	6.09
	mean	21.73	12.49	11.06	11.12	9.82
Callcener	00	28.81	15.79	14.85	14.31	13.27
	05	23.12	10.00	9.21	10.00	8.76
	10	17.28	8.71	7.48	7.33	6.63
	15	12.67	7.97	6.24	6.63	5.89
	20	9.90	7.72	6.49	6.29	5.89
	mean	18.36	10.04	8.85	8.91	8.09
Mensa	00	35.89	21.14	20.05	20.30	18.56
	05	31.14	14.16	11.68	13.12	10.64
	10	25.10	9.75	9.11	9.31	8.07
	15	19.21	8.71	7.23	7.67	6.68
	20	14.11	7.87	6.14	6.68	6.04
	mean	25.09	12.33	10.84	11.42	10.00

 Table 1. EER(%) of the SV system using four methods for different noise types and SNRs (dB) on Librispeech.

version of training data (MIX), which is a common method to improve the performance under noisy environments; MTAN trained using FL-Loss (FL); MTAN trained using AL-Loss (AL); the fusion system of FL and AL (Fusion). Specifically, the stabilization strategy proposed in this paper has been applied to both FL system and AL system. The equal error rate (EER) values of different methods are shown in Table 1 and Table 2.

The results show that our proposed methods achieved the best performance across all of the SNR levels on Librispeech corpus and the lowest EERs across the majority of the SNR levels on Aishell-1 corpus. We can find that both FL system and AL system outperform the baseline and MIX system which indicates the adversarial training framework truly improves the performance of SV task under the noisy environments. Besides, we have conducted score-level fusion using the weights learned by linear regression algorithm to make full use of complementary information between FL system and AL system, which could further improve the discriminative ability of the system. In addition, the results on two corpuses in clean condition show that MTAN could outperform the Baseline system and MIX system even in the clean condition.

Table 2. EER(%) of the SV system using four methods for different noise types and SNRs (dB) on Aishell-1.

NOISE	SNR	Baseline	MIX	FL	AL	Fusion
Clean	-	7.33	10.39	4.63	4.64	3.82
White	00	41.66	29.52	36.01	34.60	33.82
	05	39.54	26.51	30.83	27.42	27.03
	10	36.14	24.28	24.23	21.52	21.14
	15	31.88	20.72	19.02	17.75	16.02
	20	26.30	17.90	14.86	13.03	12.14
	mean	35.10	23.79	24.99	22.86	22.03
Babble	00	28.48	24.49	25.73	25.55	22.93
	05	22.54	18.87	17.71	17.56	15.44
	10	17.76	15.59	12.72	12.51	10.94
	15	14.10	13.64	9.35	9.81	8.86
	20	11.90	12.36	7.25	7.41	7.11
	mean	18.96	16.99	14.55	14.57	13.02
	00	29.24	24.75	25.15	25.64	22.58
Cafeteria	05	23.58	19.19	17.92	17.27	15.41
	10	18.60	15.86	12.54	12.14	10.62
	15	14.16	13.64	9.01	8.92	8.04
	20	11.44	12.23	7.17	6.88	6.62
	mean	19.40	17.13	14.36	14.17	12.65
Callcener	00	27.24	22.71	23.48	22.95	20.47
	05	21.48	17.95	15.94	15.88	13.61
	10	16.72	14.87	11.75	11.56	10.02
	15	13.16	13.11	8.50	8.42	7.83
	20	10.79	12.22	6.77	6.68	6.49
	mean	17.88	16.17	13.29	13.10	11.68
Mensa	00	33.53	25.1	26.2	25.89	23.16
	05	27.84	20.07	18.76	18.43	16.23
	10	21.90	16.59	14.24	13.69	12.07
	15	16.90	14.26	10.55	9.89	9.10
	20	13.61	12.61	8.12	7.56	7.59
	mean	22.76	17.73	15.57	15.09	13.63

4. CONCLUSIONS

In this paper, we have explored the potential advantage of MTAN in extracting noise-robust speaker representation. The framework consists of three components: an encoder that extracts a noise-robust speaker embedding, a classifier and a discriminator that classifies the speaker and noise type of the speaker embedding respectively. Unlike the traditional multi-task learning where the encoder is trained to maximize the classification accuracy of the classifier and discriminator, MTAN is trained adversarially to the noise classification task, so that the embedding becomes speaker-discriminative and noise-robust. Experimental results on the Aishell-1 and Librispeech corpuses have shown that the proposed method could achieve dominant results in clean condition and the most noisy environments. In the future, we will conduct the experiments in lower SNR condition and other related applications.

5. ACKNOWLEDGEMENTS

This work was supported by the National Natural Science Foundation of China (Grant No.61876160).

6. REFERENCES

- 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, 2011.
- [2] S. J. D Prince and J. H Elder, "Probabilistic linear discriminant analysis for inferences about identity," in *Computer Vision*, 2007. ICCV 2007. IEEE 11th International Conference on. IEEE, 2007, pp. 1–8.
- [3] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436, 2015.
- [4] E. Variani, X. Lei, E. McDermott, I. Lopez-Moreno, and J. Gonzalez-Dominguez, "Deep neural networks for small footprint text-dependent speaker verification.," in *IEEE International Conference on Acoustics, Speech* and Signal Processing, 2014, vol. 14, pp. 4052–4056.
- [5] G. Heigold, I. Moreno, S. Bengio, and N. Shazeer, "End-to-end text-dependent speaker verification," in Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on. IEEE, 2016, pp. 5115–5119.
- [6] K. Chen and A. Salman, "Learning speaker-specific characteristics with a deep neural architecture," *IEEE Transactions on Neural Networks*, vol. 22, no. 11, pp. 1744–1756, 2011.
- [7] 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.
- [8] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive statistics pooling for deep speaker embedding," *arXiv* preprint arXiv:1803.10963, 2018.
- [9] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for textindependent speaker verification," in *Proc. Interspeech*, 2017, pp. 999–1003.
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [11] J. Zhang, N. Inoue, and K. Shinoda, "I-vector transformation using conditional generative adversarial networks for short utterance speaker verification," *arXiv* preprint arXiv:1804.00290, 2018.
- [12] W. Ding and L. He, "Mtgan: Speaker verification through multitasking triplet generative adversarial networks," *arXiv preprint arXiv:1803.09059*, 2018.

- [13] H. Yu, Z. H. Tan, Z. Ma, and J. Guo, "Adversarial network bottleneck features for noise robust speaker verification," arXiv preprint arXiv:1706.03397, 2017.
- [14] Y. Lei, L. Burget, and N. Scheffer, "A noise robust ivector extractor using vector taylor series for speaker recognition," in Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on. IEEE, 2013, pp. 6788–6791.
- [15] J. S. Erkelens, R. C. Hendriks, R. Heusdens, and J. Jensen, "Minimum mean-square error estimation of discrete fourier coefficients with generalized gamma priors," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 6, pp. 1741–1752, 2007.
- [16] Y. Xu, J. Du, L. R. Dai, and C. H. Lee, "A regression approach to speech enhancement based on deep neural networks," *IEEE/ACM Transactions on Audio, Speech* and Language Processing (TASLP), vol. 23, no. 1, pp. 7–19, 2015.
- [17] O. Plchot, L. Burget, H. Aronowitz, and P. Matejka, "Audio enhancing with dnn autoencoder for speaker recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on.* IEEE, 2016, pp. 5090–5094.
- [18] F. Weninger, H. Erdogan, S. Watanabe, E. Vincent, J. Le Roux, J. R. Hershey, and B. Schuller, "Speech enhancement with lstm recurrent neural networks and its application to noise-robust asr," in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2015, pp. 91–99.
- [19] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline," in 2017 20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA). IEEE, 2017, pp. 1–5.
- [20] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on. IEEE, 2015, pp. 5206–5210.