



Multi-Task Learning with High-Order Statistics for X-vector based Text-Independent Speaker Verification

Lanhua You, Wu Guo, Lirong Dai, Jun Du

National Engineering Laboratory for Speech and Language Information Processing
University of Science and Technology of China, Hefei, China

lhyou@mail.ustc.edu.cn, {guowu, lrdai, jundu}@ustc.edu.cn

Abstract

The x-vector based deep neural network (DNN) embedding systems have demonstrated effectiveness for text-independent speaker verification. This paper presents a multi-task learning architecture for training the speaker embedding DNN with the primary task of classifying the target speakers, and the auxiliary task of reconstructing the first- and higher-order statistics of the original input utterance. The proposed training strategy aggregates both the supervised and unsupervised learning into one framework to make the speaker embeddings more discriminative and robust. Experiments are carried out using the NIST SRE16 evaluation dataset and the VOiCES dataset. The results demonstrate that our proposed method outperforms the original x-vector approach with very low additional complexity added.

Index Terms: Speaker verification, High-order statistics, X-vector, Multi-task learning, Unsupervised learning

1. Introduction

Speaker verification (SV) is the task of verifying a person's claimed identity from speech signals. Converting speaker-specific characteristics from variable length utterances into fixed length vectors is a key component of these SV systems. Over the years, low-dimensional embedding based systems have been widely used for text-independent speaker verification. Usually, the embeddings are represented by i-vectors [1]. Combined with the probabilistic linear discriminant analysis (PLDA) [2] backend, the i-vector/PLDA framework has become the dominant approach for the last decade.

With the great success of deep neural networks (DNNs) in machine learning fields, many novel DNN frameworks have been proposed to extract the speaker embeddings which can achieve comparable or even better performance compared with the traditional i-vector/PLDA system [3, 4, 5, 6, 7]. Typically, the utterance-level optimized speaker embeddings, such as x-vectors [5, 8], are used to replace i-vectors as the front-end and a trainable backend (such as PLDA) is employed for further modeling.

Most speaker embedding extraction only utilizes speaker labels and do not consider other information in the model training. Since the phonetic information is the predominant component in speech signals, some researchers have tried to add this information in model training using multi-task learning (MTL) frameworks [9, 10, 11]. However, most utterances do not have phonetic labels at hand and an additional automatic speech recognition (ASR) system is always required. This requirement will result in longer training times. Furthermore, the recognition errors in some utterances with low signal-to-noise ratios (SNR) may limit the use of the phonetic labels for speaker verification.

In this paper, we propose a novel multi-task learning strategy based on the x-vector architecture with very low complexity added. We first compute the first- and higher-order statistics (including the mean, standard deviation, skewness and kurtosis) of an input utterance, and these four statistics are then used as the reconstruction targets of the auxiliary task. The high-order statistics (HOS) contain statistical characteristics of the input signal and are more easily obtained compared with other auxiliary information. With the MTL, the deep embedding network is trained not only to classify the target speakers but also to reconstruct both the low- and high-order statistics from input features. The resulting speaker embeddings can benefit from both the supervised and unsupervised learning. We evaluate our experiments using the NIST SRE16 evaluation dataset [12] and the VOiCES dataset [13]. The experimental results show that our proposed method achieves better performance compared with the original x-vector approach and requires very little extra computation load.

The remainder of this paper is organized as follows. Section 2 gives an introduction to our x-vector baseline. Section 3 presents the high-order statistics, as well as the proposed MTL strategy. Section 4 presents the experimental setup and the results of this study. In Section 5, we summarize our work and discuss future work.

2. Baseline network architecture

The network architecture of our x-vector baseline system is the same as that described in [5]. As depicted in Figure 1, the first five TDNN (or 1-dimensional dilated CNN) layers l_1 to l_5 are stacked for extracting the frame-level features. More specifically, the TDNN layers with dilation rates of 2 and 3 are used for the second and third layers respectively, while the others retain the dilation rate of 1. The kernel sizes of the five layers are 5, 3, 3, 1 and 1, respectively.

The final frame-level output vectors of the whole variable-length utterance are aggregated into a fixed segment-level vector through the statistics pooling layer. The mean and standard deviation are calculated and then concatenated together for the statistics pooling. Two additional fully connected layers l_6 and l_7 are added to obtain a low-dimensional utterance-level representation that is finally passed into a softmax output layer.

The deep embedding network is trained to predict the correct speaker labels with cross entropy (CE) loss. Once the DNN is trained, we remove the softmax layer and the last fully connected layer l_7 , and the output of the linear affine layer directly on top of the statistics pooling is extracted as the speaker embedding.

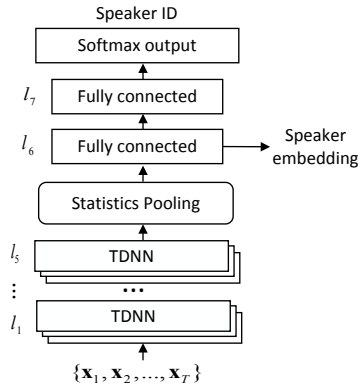


Figure 1: Network architecture of the x -vector baseline.

3. Multi-task learning with high-order statistics

3.1. High-order statistics

Higher-order statistics can be used in estimation of the shape of unimodal distributions and have been applied to many tasks [14, 15, 16]. In the original x -vector system, low-order statistics, such as the mean and standard deviation, are calculated to perform the statistics pooling and have demonstrated their effectiveness for extracting speaker-specific embeddings [5, 17]. Here, these statistics of input features are computed as the fixed-length input representations. Given an input utterance, $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$, the different order statistics can be calculated as follows

$$\boldsymbol{\mu} = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t \quad (1)$$

$$\boldsymbol{\sigma} = \sqrt{\frac{1}{T} \sum_{t=1}^T (\mathbf{x}_t - \boldsymbol{\mu})^2} \quad (2)$$

$$\mathbf{s} = \frac{1}{T} \sum_{t=1}^T \left(\frac{\mathbf{x}_t - \boldsymbol{\mu}}{\boldsymbol{\sigma}} \right)^3 \quad (3)$$

$$\mathbf{k} = \frac{1}{T} \sum_{t=1}^T \left(\frac{\mathbf{x}_t - \boldsymbol{\mu}}{\boldsymbol{\sigma}} \right)^4 \quad (4)$$

where T is the number of frames in the input utterance. $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are the mean and standard deviation vectors, respectively. In addition to the first- and second-order statistics, higher-order statistics including skewness \mathbf{s} and kurtosis \mathbf{k} , are also used to describe the statistical characteristics of the input utterance. The skewness measures the asymmetry of a distribution with respect to its mode while the kurtosis measures the “tailedness” of the data distribution. Both of these values enrich the statistical information from input features. Finally, we concatenate these four statistics into a fixed-dimensional vector, and we call this the HOS vector in this paper.

$$\mathbf{z} = [\boldsymbol{\mu}, \boldsymbol{\sigma}, \mathbf{s}, \mathbf{k}] \quad (5)$$

3.2. Multi-task learning strategy

In the original x -vector architecture, only speaker labels are considered for the DNN training. We use a MTL strategy to

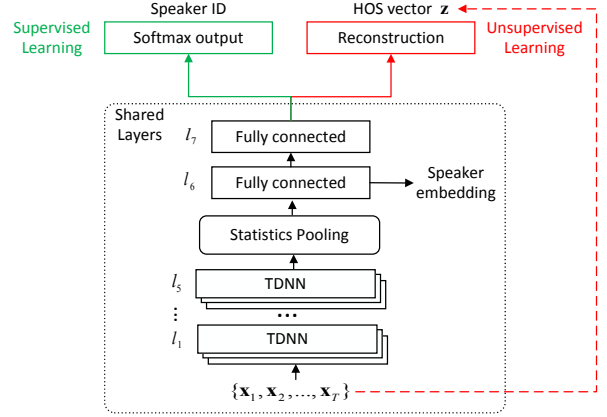


Figure 2: The proposed multi-task learning architecture.

incorporate the abovementioned HOS vector into the x -vector DNN training, where the classification of the speaker label is still the primary task, and the reconstruction of the HOS vector \mathbf{z} is the auxiliary task. As depicted in Figure 2, the speaker label is a supervised label and the HOS vector is an unsupervised label; therefore, the proposed MTL training strategy aggregates both the supervised and unsupervised learning into one framework.

In MTL, the network is trained to perform both the primary classification task and the auxiliary reconstruction task. Except for the task-specific output layers, all the layers are shared between both tasks. For unsupervised learning, we add a linear layer on top of the last shared layer l_7 to obtain the reconstructed vector \mathbf{h} . The reconstruction task aims to minimize the mean square error (MSE) loss between its output vector and \mathbf{z} , which can be formulated as

$$MSE = \frac{1}{N} \sum_{n=1}^N \left\| \mathbf{h}^{(n)} - \mathbf{z}^{(n)} \right\|^2 \quad (6)$$

where N is the total number of training samples and $\mathbf{h}^{(n)}$ is the reconstruction of statistical representation $\mathbf{z}^{(n)}$ for the n -th input utterance. Combined with the original CE loss, the final loss function can be written as

$$Loss = \alpha MSE + (1 - \alpha) CE \quad (7)$$

where α is the task weight.

After the multi-task training, the extracted speaker embedding will contain both the discriminative and unsupervised speaker information. From a model training perspective, the auxiliary task enhances the model generalization ability by introducing regularization into the shared layers. Since the vector \mathbf{z} is quite low-dimensional ($4 \times$ feature dimension), adding the auxiliary task only slightly increases the number of parameters in the output layer and requires very low extra computational cost. Moreover, the extra overhead can be neglected, since the top layers are removed when extracting speaker embeddings.

4. Experiments and analysis of results

4.1. Experimental settings

All systems are based on the TensorFlow implementation of the x -vector speaker embedding [18]. We trained the network and

Table 1: Results of the NIST SRE16. *MT-on- αp* denotes the multi-task learning system with task weight $\alpha = p/10$ and it concatenates the statistics from 1 to n orders to perform the HOS vector \mathbf{z} . For example, *MT-o3- $\alpha 3$* means $\mathbf{z} = [\boldsymbol{\mu}, \boldsymbol{\sigma}, \mathbf{s}]$ and $\alpha = 0.3$. Here, we focus on the performance with respect to weight α .

system	SRE16, Pooled			SRE16, Cantonese			SRE16, Tagalog		
	EER%	minDCF	actDCF	EER%	minDCF	actDCF	EER%	minDCF	actDCF
x-vector	8.03	0.586	0.605	4.06	0.396	0.404	12.00	0.743	0.807
MT-o4- $\alpha 1$	8.09	0.581	0.607	4.04	0.379	0.382	12.15	0.744	0.832
MT-o4- $\alpha 2$	7.91	0.571	0.583	3.83	0.390	0.398	11.99	0.723	0.768
MT-o4- $\alpha 3$	7.79	0.563	0.568	3.88	0.368	0.376	11.69	0.732	0.761
MT-o4- $\alpha 4$	8.03	0.564	0.577	3.90	0.364	0.375	12.18	0.727	0.779
MT-o4- $\alpha 10$	26.3	0.979	4.40	22.7	0.961	4.02	29.9	0.990	4.79

Table 2: Results of the NIST SRE16. Comparison results of proposed multi-task learning systems using different orders of statistics when $\alpha = 0.3$.

system	SRE16, Pooled			SRE16, Cantonese			SRE16, Tagalog		
	EER%	minDCF	actDCF	EER%	minDCF	actDCF	EER%	minDCF	actDCF
MT-o1- $\alpha 3$	7.85	0.578	0.593	3.76	0.370	0.376	11.99	0.751	0.810
MT-o2- $\alpha 3$	8.31	0.568	0.575	4.18	0.378	0.386	12.39	0.736	0.764
MT-o3- $\alpha 3$	8.05	0.572	0.585	3.89	0.376	0.386	12.21	0.730	0.783
MT-o4- $\alpha 3$	7.79	0.563	0.568	3.88	0.368	0.376	11.69	0.732	0.761

extracted x-vectors using TensorFlow¹. The other procedures (including data processing, feature extraction and PLDA backend) are implemented using Kaldi Toolkit [5, 8].

4.1.1. Training data and evaluation metric

The experiments are carried out on the NIST SRE16 evaluation dataset, and both the development and evaluation parts of the VOICES dataset. For the NIST SRE16 dataset, the training data mainly consists of the telephone speech (with a small amount of the microphone speech) from the NIST SRE2004-2010, Mixer 6 and Switchboard datasets. We also use the data augmentation techniques described in [8], and employ the babble, music and noise augmented data to increase the amount and diversity of the existing training data. In summary, there are a total of 183,457 recordings from 7,001 speakers, including approximately 96,000 randomly selected augmented recordings.

The VOICES dataset for the speaker verification task is described in the ‘‘VOICES from a Distance Challenge 2019’’ [13]. The VOICES development dataset contains 15,904 segments of noisy and far-field speech from 196 speakers. The evaluation set consists of 11,392 distant recordings from different microphone types and different rooms, both of which could be more challenging than those featured in the development set. We use both the Voxceleb1 and Voxceleb2 datasets [19, 20] as the training set for VOICES Challenge. Data augmentation techniques (including music, noise and reverberation) are also applied for model training.

The performance is evaluated in terms of equal error rate (EER), the minimal detection cost function (minDCF) and actual detection cost function (actDCF) calculated using the SRE16 and VOICES official scoring softwares. For the NIST SRE16, the equalized results are used.

4.1.2. Input features

For the NIST SRE16, our input acoustic features are 23-dimensional MFCC features with a frame-length of 25 ms that are mean-normalized over a sliding window of up to 3 s. An energy-based voice activity detection (VAD) is used to filter out nonspeech frames. Instead of 23-dimensional MFCCs, we use the 30-dimensional same type of MFCCs for VOICES.

4.1.3. Model configuration

In all x-vector based systems, for both the SRE16 and VOICES, the number of hidden nodes for the first four frame-level layers is 512, while that number is 1536 for the last frame-level layer. Each of the two fully connected layers l_6 and l_7 has 512 nodes. All the nonlinear activation functions of hidden layers are ReLU. We use the same type of batch normalization and L2 weight decay as in [18] to prevent overfitting.

4.1.4. PLDA backend

For the NIST SRE16, the DNN embeddings are centered using the unlabeled development data and are projected using LDA, which reduces the dimensionality of x-vectors to 100. For training the PLDA model, all the training data (except the Switchboard data) and their corresponding augmented versions are used. Finally, the PLDA model is adapted to the unlabeled data through the unsupervised adaptation in Kaldi. For VOICES, no adaptation technique is used. We select the longest 200,000 recordings from the training set to train the backend, and the best LDA dimension is 110.

4.2. Results and analysis

4.2.1. NIST SRE16

Table 1 presents the performance of MTL systems with different task weights in NIST SRE16. It can be observed that the proposed MTL systems with $\alpha = 0.2 \sim 0.4$ outperform the

¹<https://github.com/hsn-zeinali/x-vector-kaldi-tf/tree/master/local/tf>

x-vector baseline system (i.e., $\alpha = 0$ in multi-task learning). When $\alpha = 0.3$, we can obtain the best performance, which is better than the x-vector baseline in terms of all evaluation metrics. On Cantonese, this strategy provides a 4% relative improvement in terms of EER and a 7% improvement in terms of both minDCF and actDCF over the original x-vector. The results demonstrate the effectiveness of the proposed MTL strategy. Note that we still can obtain an EER of 22.7% on Cantonese even when $\alpha = 1.0$. This result shows that embeddings still contain speaker-discriminative information using only our unsupervised learning and shows the importance of unsupervised information.

In Table 2, we investigate the effect of different orders of statistics in multi-task learning where $\alpha = 0.3$. On average, adding both the skewness and kurtosis can achieve better performance. This result demonstrates that speaker embeddings can benefit from the higher-order statistics.

4.2.2. VOICES

For VOICES, we directly use $\alpha = 0.3$, which is tuned in the NIST SRE16. The results of VOICES are reported in Table 3 and Table 4.

Table 3: Results on the VOICES development where $\alpha = 0.3$.

system	EER%	minDCF	actDCF
x-vector	3.36	0.387	0.515
MT-o1- α 3	3.22	0.369	0.495
MT-o2- α 3	2.97	0.335	0.409
MT-o3- α 3	3.13	0.369	0.443
MT-o4- α 3	3.37	0.373	0.477

Table 4: Results on the VOICES evaluation where $\alpha = 0.3$.

system	EER%	minDCF	actDCF
x-vector	8.25	0.628	0.779
MT-o1- α 3	7.90	0.595	0.709
MT-o2- α 3	8.02	0.590	0.702
MT-o3- α 3	7.89	0.603	0.687
MT-o4- α 3	7.66	0.572	0.639

We can observe a more obvious improvement provided by our proposed MTL strategy on both development and evaluation. It is interesting to see that the MTL system using only the mean and standard deviation can achieve the best performance for the development set and can improve on the x-vector baseline system by 12% in EER, 13% in minDCF and 21% in actDCF. For the more challenging evaluation set, using all four orders statistics in multi-task learning can provide the largest improvement and outperform the original x-vector by 7% in EER, 9% in minDCF and 18% in terms of actDCF. These results make clear that our MTL strategy makes the speaker embedding more discriminative and robust by joint learning speaker classification and reconstruction of the high-order statistics.

4.2.3. Speed

Here, we compare the training speed of the original network with that of the multi-task training network. The networks are trained on Nvidia GeForce GTX 1080Ti GPU. Table 5 presents

the time required for each iteration when the two networks are trained using the Voxceleb datasets. It can be observed that MT-o4- α 3, in which $\mathbf{z} = [\boldsymbol{\mu}, \boldsymbol{\sigma}, \mathbf{s}, \mathbf{k}]$ ($4 \times 30 = 120$ dimensional), is only 4% slower than the original x-vector. For the other case when not all statistics are used or the dimensionality of input features is smaller than 30, our MTL systems could be faster. This result demonstrates that our multi-task learning strategy improves the performance with very low additional complexity added.

Table 5: Training speed comparison.

system	Time(min)/Iter
x-vector	10.052
MT-o4- α 3	10.476

5. Conclusions

In this study, we propose a novel MTL strategy for x-vector based architecture. The network is trained not only to classify the target speaker but also to reconstruct the HOS vector of input utterance. The experimental results demonstrate the effectiveness of our proposed strategy compared with the original x-vector architecture. The speaker embeddings will benefit from the auxiliary HOS information and can be more robust and discriminative through additional unsupervised learning. In addition, the proposed method is easy to implement and requires very low extra overhead during the training phase.

In our future studies, we will continue to focus on the use of high-order statistics and investigate other useful multi-task learning strategies for x-vector based speaker verification.

6. Acknowledgements

This work was partially funded by the National Natural Science Foundation of China (Grant No. U1836219) and the National Key Research and Development Program of China (Grant No. 2016YFB100 1303).

7. References

- [1] N. Dehak, P. 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] P. Kenny, "Bayesian speaker verification with heavy-tailed priors," in *Odyssey*, 2010, p. 14.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] G. Bhattacharya, J. Alam, and P. Kenny, "Deep speaker embeddings for short-duration speaker verification," in *Interspeech*, 2017, pp. 1517–1521.

- [7] C. Zhang and K. Koishida, "End-to-end text-independent speaker verification with triplet loss on short utterances," in *Interspeech*, 2017, pp. 1487–1491.
- [8] 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.
- [9] Y. Liu, L. He, J. Liu, and M. T. Johnson, "Speaker embedding extraction with phonetic information," in *Interspeech 2018*, 2018, pp. 2247–2251.
- [10] N. Chen, Y. Qian, and K. Yu, "Multi-task learning for text-dependent speaker verification," in *Interspeech 2015*, 2015, pp. 185–189.
- [11] Z. Tang, L. Li, D. Wang, and R. Vipperla, "Collaborative joint training with multitask recurrent model for speech and speaker recognition," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 25, no. 3, pp. 493–504, 2017.
- [12] S. O. Sadjadi, T. Kheyrkhan, A. Tong, C. Greenberg, D. Reynolds, E. Singer, L. Mason, and J. Hernandez-Cordero, "The 2016 nist speaker recognition evaluation," in *Interspeech*, 2017, pp. 1353–1357.
- [13] M. K. Nandwana, J. V. Hout, M. McLaren, C. Richey, A. Lawson, and M. A. Barrios, "The voices from a distance challenge 2019 evaluation plan," *arXiv preprint arXiv:1902.10828*, 2019.
- [14] J. Xu, P. Ye, Q. Li, H. Du, Y. Liu, and D. Doermann, "Blind image quality assessment based on high order statistics aggregation," *IEEE Transactions on Image Processing*, vol. 25, no. 9, pp. 4444–4457, 2016.
- [15] J. Richiardi and A. Drygajlo, "Evaluation of speech quality measures for the purpose of speaker verification," in *Odyssey*, 2008, p. 5.
- [16] M. K. I. Molla and K. Hirose, "On the effectiveness of mfccs and their statistical distribution properties in speaker identification," in *2004 IEEE Symposium on Virtual Environments, Human-Computer Interfaces and Measurement Systems, 2004.(VCIMS)*. IEEE, 2004, pp. 136–141.
- [17] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive s-statistics pooling for deep speaker embedding," in *Proc. Interspeech 2018*, 2018, pp. 2252–2256. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-993>
- [18] H. Zeinali, L. Burget, J. Rohdin, T. Stafylakis, and J. Cernocky, "How to improve your speaker embeddings extractor in generic toolkits," *arXiv preprint arXiv:1811.02066*, 2018.
- [19] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: A large-scale speaker identification dataset," in *Proc. Interspeech 2017*, 2017, pp. 2616–2620. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2017-950>
- [20] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *Proc. Interspeech 2018*, 2018, pp. 1086–1090. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1929>