

# Deep Neural Network Embeddings with Gating Mechanisms for 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

In this paper, gating mechanisms are applied in deep neural network (DNN) training for x-vector-based text-independent speaker verification. First, a gated convolution neural network (GCNN) is employed for modeling the frame-level embedding layers. Compared with the time-delay DNN (TDNN), the GCNN can obtain more expressive frame-level representations through carefully designed memory cell and gating mechanisms. Moreover, we propose a novel gated-attention statistics pooling strategy in which the attention scores are shared with the output gate. The gated-attention statistics pooling combines both gating and attention mechanisms into one framework; therefore, we can capture more useful information in the temporal pooling layer. Experiments are carried out using the NIST SRE16 and SRE18 evaluation datasets. The results demonstrate the effectiveness of the GCNN and show that the proposed gated-attention statistics pooling can further improve the performance.

**Index Terms:** Speaker verification, Gating mechanism, Gated convolution neural network, Attention mechanism

## 1. Introduction

Speaker verification (SV) is a task to verify a person's claimed identity from speech signals. During the last decade, the i-vector [1] combined with the probabilistic linear discriminant analysis (PLDA) framework [2] has become the dominant approach in this field. The i-vector algorithm serves as the front-end, which converts a variable-length utterance to a low dimensional vector, and the PLDA algorithm is an effective backend classifier.

With the great success of deep neural networks (DNNs) in machine learning fields, more attention has been drawn to the use of DNNs to extract i-vector similar vectors, known as speaker embeddings. Many novel DNN embedding-based systems have been proposed, and they have achieved comparable or even better performance compared with the traditional i-vector paradigm [3, 4, 5, 6, 7, 8, 9, 10].

In most DNN embedding systems [5, 7, 8, 9, 10], an input utterance with a variable length is first fed into several frame-level layers to obtain high-level feature representations. The frame-level layers are usually modeled by recurrent neural networks (RNNs) [9], convolution neural networks (CNNs) [7, 10] or time-delay neural networks (TDNNs) [5, 8]. Next, a pooling layer maps all frames of the input utterance into a fixed-dimensionality vector, and the speaker embedding is generated from the following stacked fully connected layers. Furthermore, the structure optimization of the networks can help to extract better speaker-discriminant embedding. Snyder et al. employed statistics pooling to replace the commonly used average pooling in x-vector systems [5, 8] in which the mean and standard de-

viation are connected together to form the utterance-level representation. The attentive statistic pooling [11, 12] is further proposed to improve the performance of the x-vector system.

Recently, the gating mechanism has proven to be useful for the natural language processing (NLP) task. Inspired by the gated linear units (GLU) [13] and the long-short term memory (LSTM) network [14], Peixin et al. proposed a gated convolution neural network (GCNN) [15] for sentence matching and demonstrated its superiority over the CNN and RNN on sentence matching tasks. The gating mechanism combined with the memory cell is introduced in the GCNN embedding layers to capture representations that contain important information. On the other hand, the gating mechanism combined with the attention mechanism also demonstrates its effectiveness for spoken-language understanding [16].

In this paper, we investigate gating mechanisms for the x-vector embedding system. More specifically, the stacked GCNN embedding layers are employed as the frame-level layers for extracting more expressive feature representations. In addition, we also apply the gating mechanism in the attention pooling layer and propose a novel gated-attention statistics pooling layer in which the attentive weights of the input frames are modeled with gating mechanism for attentive statistic pooling. We evaluate our experiments on the NIST SRE16 evaluation dataset [17] and the Call My Net2 (CMN2) part of NIST SRE18 evaluation dataset [18]. The experimental results show the application of the GCNN for text-independent speaker verification can improve the performance of the DNN embedding system. Additionally, the proposed gated-attention statistics pooling offers further improvement over the conventional attentive statistics pooling approach.

The remainder of this paper is organized as follows. Section 2 introduces our x-vector baseline and the attentive statistics pooling. Section 3 describes the application of the GCNN, as well as the proposed gated-attention statistics pooling strategy. Section 4 presents the experimental setup and the results of this study. In Section 5, we summarize our work.

## 2. Baseline network architecture

### 2.1. X-vector baseline

The network architecture of our x-vector baseline system is similar to that described in [8]. The first five TDNN (or 1-dimensional dilated CNN) layers are stacked to extract the frame-level DNN features. The TDNN layers with dilation rates of 2 and 4 are used for the second and third layers, respectively, while the others retain the dilation rate of 1. The kernel sizes of these 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 for statistics pooling. Two additional fully connected layers followed with a softmax layer are used to predict speaker labels. Once the DNN is trained, the output of the linear affine layer on top of the statistics pooling is extracted as the speaker embedding.

## 2.2. Attentive statistics pooling

Instead of treating the output representations from the last frame-level layer equally in the statistics pooling, the attention mechanism is applied to weight the more speaker-discriminative frames in the input utterance [11, 12].

A single-head attention [11] strategy is applied in the x-vector baseline system. Suppose  $\mathbf{h}_t$  is the hidden representation of the  $t$ -th input frame below the attention layer. The attention weight  $\alpha_t$  for  $\mathbf{h}_t$  can be calculated as

$$\begin{aligned} e_t &= \mathbf{w}_2^T f(\mathbf{W}_1 \mathbf{h}_t) \\ \alpha_t &= \frac{\exp(e_t)}{\sum_{\tau} \exp(e_{\tau})} \end{aligned} \quad (1)$$

where the  $\mathbf{W}_1$  and  $\mathbf{w}_2$  are the learned parameters, and  $f$  is the ReLU function. Then, the weighted mean vector  $\mathbf{u}$  and the standard deviation vector  $\delta$  are computed as follows

$$\begin{aligned} \mathbf{u} &= \sum_t \alpha_t \mathbf{h}_t \\ \delta &= \sqrt{\sum_t \alpha_t \mathbf{h}_t \odot \mathbf{h}_t - \mathbf{u} \odot \mathbf{u}} \end{aligned} \quad (2)$$

where  $\odot$  is the elementwise multiplication. Finally, we concatenate the  $\mathbf{u}$  and  $\delta$  to perform the attentive statistics pooling.

## 3. DNN with gating mechanisms

In this section, we will employ the gating mechanisms for the x-vector system in two ways. As depicted in Figure 1, we first use four GCNN [15] layers to replace the first four TDNN layers for extracting the frame-level representations. On the other hand, we propose the gated-attention statistics pooling as an alternative attention method for aggregating the frame-level vectors. The remaining part is similar to that of the x-vector baseline.

### 3.1. Frame-level GCNN layers

The GLU [13] uses the output gate to control what information should be propagated through the convolutional layers. To better control the path through which information flows in the hierarchical structure, we adopt both output and forget gates in the GCNN. In more detail, for the  $L$ -th frame-level layer, the GCNN can be described as follows:

$$\begin{aligned} \mathbf{o}_t^L &= \sigma(\mathbf{h}_t^{L-1}(c) * \mathbf{W}_o + \mathbf{b}_o) \\ \mathbf{f}_t^L &= \sigma(\mathbf{h}_t^{L-1}(c) * \mathbf{W}_f + \mathbf{b}_f) \\ \mathbf{g}_t^L &= \tanh(\mathbf{h}_t^{L-1}(c) * \mathbf{W}_g + \mathbf{b}_g) \\ \mathbf{c}_t^L &= \mathbf{f}_t^L \odot \mathbf{c}_t^{L-1} + (1 - \mathbf{f}_t^L) \odot \mathbf{h}_t^{L-1} \\ \mathbf{h}_t^L &= \mathbf{o}_t^L \odot \mathbf{g}_t^L + \mathbf{c}_t^L \end{aligned} \quad (3)$$

where  $*$  is the convolution operation and  $\sigma$  is the sigmoid function,  $\mathbf{W}_o$ ,  $\mathbf{W}_f$  and  $\mathbf{W}_g$  are convolution parameters, while  $\mathbf{b}_o$ ,  $\mathbf{b}_f$  and  $\mathbf{b}_g$  are bias, and  $\mathbf{o}_t^L$ ,  $\mathbf{f}_t^L$  and  $\mathbf{c}_t^L$  are the output gate, forget gate and memory cell for the  $t$ -th frame in layer  $L$ , respectively. If the dimensions of  $\mathbf{c}_t^L$  and the candidate  $\mathbf{g}_t^L$  are not matched in Eq. 3, we can perform a linear projection.  $\mathbf{h}_t^{L-1}(c)$  is the input, which is formed by concatenating  $\mathbf{h}_t^{L-1}$  with its nearby frames in the previous layer or input acoustic features.

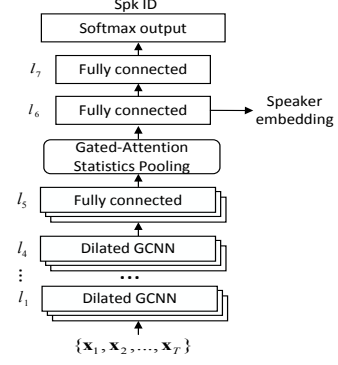


Figure 1: Network architecture of our proposed DNN.




Figure 2: Structure of the GCNN unit.

Here, we apply the GCNN in a dilated way like the TDNN in the baseline system. As shown in Figure 2, when the kernel width is 3 and the dilation rate is 2,  $\mathbf{h}_t^{L-1}(c)$  can be written as

$$\mathbf{h}_t^{L-1}(c) = \{\mathbf{h}_{t-2}^{L-1}, \mathbf{h}_t^{L-1}, \mathbf{h}_{t+2}^{L-1}\} \quad (4)$$

The GCNN optionally allows information flow through the hierarchical structure using the gating mechanism. The output gate  $\mathbf{o}_t^L$  modulates the output information of the candidate convolution  $\mathbf{g}_t^L$ , which contains a larger temporal context. On the other hand, the memory cell  $\mathbf{c}_t^L$  can be treated as a modified residual learning that contains a smaller temporal context modulated by the forget gate  $\mathbf{f}_t^L$ . Therefore, the frame-level representations will benefit from both larger and smaller temporal contexts, which are further regulated through the gating mechanism.

### 3.2. Gated-attention statistics pooling

We propose a novel gated-attention statistics pooling in which the gating mechanism is introduced into the attention pooling layer to enhance the ability to extract speaker-discriminative information from the last frame-level layer. The gated-attention statistics pooling is depicted in Figure 3. Suppose that the last frame-level layer is the  $S$ -th layer. An output gate is calculated to control the output of the  $S$ -th hidden layer.

$$\begin{aligned} \tilde{\mathbf{e}}_t^S &= \mathbf{h}_t^{S-1}(c) * \mathbf{W}_s + \mathbf{b}_s \\ \mathbf{o}_t^S &= \sigma(\tilde{\mathbf{e}}_t^S) \end{aligned} \quad (5)$$

where  $\mathbf{W}_s$  and  $\mathbf{b}_s$  are the weight and bias parameters. The  $t$ -th output vector of the  $S$ -th layer is as follows.

$$\mathbf{z}_t = \mathbf{o}_t^S \odot \mathbf{h}_t^S \quad (6)$$

Such a gating mechanism regulates each element of the  $t$ -th output  $\mathbf{z}_t$  with the corresponding element of the output gate

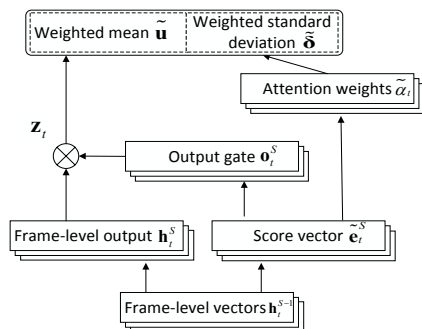


Figure 3: Gated-attention statistics pooling.

$\mathbf{o}_t^s$ . In other words, each element of the output gate, which is normalized between 0 and 1 with the sigmoid function, can be viewed as a kind of local attention to the corresponding  $\mathbf{z}_t$ . The mean of all the elements in  $\tilde{\mathbf{e}}_t^s$  can measure the importance of the  $t$ -th frame-level representation. With the softmax function, the attention weight for  $\mathbf{z}_t$  can be calculated as

$$\tilde{\alpha}_t = \frac{\exp(\text{mean}(\tilde{\mathbf{e}}_t^s))}{\sum_{\tau} \exp(\text{mean}(\tilde{\mathbf{e}}_{\tau}^s))} \quad (7)$$

Finally, the weighted mean and the weighted standard deviation of the gated-attention pooling layer are computed by

$$\begin{aligned} \tilde{\mathbf{u}} &= \sum_t \tilde{\alpha}_t \mathbf{z}_t \\ \tilde{\delta} &= \sqrt{\sum_t \tilde{\alpha}_t \mathbf{z}_t \odot \mathbf{z}_t - \tilde{\mathbf{u}} \odot \tilde{\mathbf{u}}} \end{aligned} \quad (8)$$

In the gated-attention statistics pooling, from another viewpoint, the attention score for  $\mathbf{z}_t$  is further divided into the output gate for alleviating the unimportant information among its elements. Therefore, the gated-attention statistics pooling provides a unified framework for the attention and gating mechanisms in which the former works at the frame level, and the latter can be considered as an elementwise attention. The gated-attention statistics pooling enhances both the attention and gating mechanisms by using the shared weights and helps to extract more speaker-discriminative embeddings.

## 4. Experiments and analysis of the results

### 4.1. Experimental settings

#### 4.1.1. Training data and evaluation metric

Our experiments are carried out in both the NIST SRE16 and SRE18 evaluation datasets. For the NIST SRE18 dataset, we only consider the CMN2 portion of the evaluation 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], which employ the babble, music, noise and reverb augmented data to increase the amount and diversity of the existing training data. In summary, there are a total of 205,600 recordings, including approximately 120,000 randomly selected augmented recordings. Note that the training data for the NIST SRE16 evaluation is consistent with that for the NIST SRE18.

The performance is evaluated in terms of equal error rate (EER) and the minimal detection cost function (minDCF) calculated using the SRE16 and 18 official scoring software.

#### 4.1.2. Input features

We select the 39-dimensional perceptual linear predictive (PLP) features containing delta and delta-delta coefficients as the input acoustic features. Each PLP feature is extracted from the speech signal of a 25 ms window with 10 ms frame shift. We employ the energy-based voice activity detection (VAD) to select the valid speech frames from the utterances, and the PLP features are cepstral mean and variance normalized along the whole input utterance before being fed into the embedding network.

#### 4.1.3. Model configuration

We mainly compare six systems as stated as follows.

**X\_TDNN:** This is the x-vector baseline system without an attention mechanism. The number of hidden nodes for the first four frame-level layers is 512, while that number is 1500 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. Each of the layers from  $l_1$  to  $l_7$  is followed by batch normalization [19]. The dropouts [20] and L2 weight decay are also used to prevent overfitting.

**X\_GCNN:** The GCNN layers are employed in the x-vector system as described above. The kernel size and dilation rate of each layer in the GCNN-based system are the same as those in the baseline system. The hidden size of the first four layers is 256 instead of 512 for computational efficiency. The remaining network structure is the same as that of the X\_TDNN baseline.

**X\_TDNN+Att:** The abovementioned attentive statistics pooling in Section 2.2 is used instead of the original statistics pooling in the x-vector baseline [11], and this method achieves better performance than the original method.

**X\_GCNN+Att/GAtt:** For these two systems, the attentive statistics pooling and proposed gated-attention statistics pooling are employed in the GCNN-based system, respectively. The dimensionality of  $\mathbf{w}_2$  in Eq. 1 is set to 256 so that the number of parameters in the attentive statistics pooling is comparable to that in the gated-attention statistics pooling.

**X\_GCNN\_Fusion:** The complementarity between the above two different attention pooling methods is also investigated here. We only report the results using the score fusion of the X\_GCNN+Att and X\_GCNN+GAtt with equal weights.

#### 4.1.4. Network training

The deep embedding networks are implemented with the TensorFlow toolkit [21]. The network is trained on approximately 10-second long chunks. We use the Adam optimizer [22] with an initial learning rate of 0.00015 and decay the learning rate based on the validation set.

#### 4.1.5. PLDA backend

We use the Kaldi toolkit to implement the PLDA backend similar to that in [23]. For both NIST SRE16 and SRE18, the DNN embeddings are centered using the unlabeled development data and projected using LDA. The PLDA model is trained and then adapted to the unlabeled data through the unsupervised adaptation in Kaldi.

## 4.2. Results and analysis

Since it is meaningful to observe the performance of both systems with and without data augmentation, we will report the performance of both systems.

Table 1: Comparison results of the different systems without data augmentation

system	SRE16, Pooled		SRE16, Cantonese		SRE16, Tagalog		SRE18, CMN2	
	EER%	minDCF	EER%	minDCF	EER%	minDCF	EER%	minDCF
X_TDNN	8.04	0.635	4.66	0.429	11.41	0.802	9.91	0.622
X_GCNN	7.77	0.598	4.40	0.401	11.19	0.762	9.35	0.613
X_TDNN+Att	8.00	0.628	4.60	0.423	11.39	0.790	9.47	0.632
X_GCNN+Att	7.68	0.596	4.28	0.397	11.07	0.757	9.43	0.609
X_GCNN+GAtt	7.48	0.585	4.24	0.397	10.70	0.742	9.12	0.596
X_GCNN_Fusion	<b>7.04</b>	<b>0.568</b>	<b>3.89</b>	<b>0.374</b>	<b>10.19</b>	<b>0.734</b>	<b>8.69</b>	<b>0.581</b>

Table 2: Comparison results of the different systems with data augmentation

system	SRE16, Pooled		SRE16, Cantonese		SRE16, Tagalog		SRE18, CMN2	
	EER%	minDCF	EER%	minDCF	EER%	minDCF	EER%	minDCF
X_TDNN	7.63	0.597	4.20	0.412	11.02	0.756	8.67	0.579
X_GCNN	7.35	0.565	4.23	0.387	10.55	0.733	8.35	0.566
X_TDNN+Att	7.47	0.596	4.18	0.422	10.76	0.757	8.67	0.573
X_GCNN+Att	7.19	0.562	3.86	0.385	10.53	0.725	8.20	0.556
X_GCNN+GAtt	7.24	0.560	4.07	0.388	10.44	0.721	8.14	0.550
X_GCNN_Fusion	<b>6.82</b>	<b>0.545</b>	<b>3.67</b>	<b>0.371</b>	<b>9.99</b>	<b>0.710</b>	<b>7.83</b>	<b>0.537</b>

#### 4.2.1. Results

Table 1 presents the performance of the different DNN embedding systems without data augmentation in the NIST SRE16 and SRE18. It can be observed that our GCNN-based systems outperform the TDNN baseline in all evaluation conditions in terms of both EER and minDCF, regardless of whether the systems are equipped with an attention mechanism or not. Clearly, it demonstrates the effectiveness of our frame-level GCNN layers. Moreover, X\_GCNN+GAtt achieves a better performance compared with X\_GCNN+Att and provides an 8% relative improvement over X\_TDNN+Att on Cantonese in terms of EER. This observation shows our proposed gated-attention statistics pooling can further enhance the performance of GCNN-based systems over the original attentive statistics pooling. Among all the above systems, the fused system achieves the best performance and improves on the attention-based TDNN baseline by 15% on Cantonese and 11% on Tagalog in NIST SRE16 in terms of EER.

The performance of the systems with data augmentation is reported in Table 2, from which we can draw a similar conclusion as from Table 1. In this case, although X\_GCNN+GAtt is comparable with X\_GCNN+Att in NIST SRE16, it is still slightly better than X\_GCNN+Att in NIST SRE18 in EER and minDCF. Again, the GCNN-based fusion system gives the largest improvement over other systems. In terms of both EER and minDCF, it is 12% better than X\_TDNN+Att on Cantonese in NIST SRE16. For NIST SRE18, a 10% relative improvement is achieved compared with the attention-based baseline system in terms of EER. These results make clear that the TDNN-based x-vector system is enhanced significantly with our introduced gating mechanisms.

#### 4.2.2. Further analysis

Here, we investigate the effect of removing different parts in the gated-attention statistics pooling. Table 3 lists the pooled results in SRE16 and the CMN2 results in SRE18 of the systems without augmentation. We set up X\_GCNN+Gonly in which the attention mechanism is removed such that only the output gate works in the pooling layer. We observe a certain decrease in the performance compared with X\_GCNN+GAtt. We can also

observe the performance degradation in the X\_GCNN+Aonly system in which the output gate is removed. This observation indicates that both the attention and gating mechanisms contribute to the final performance in the gated-attention statistics pooling layer.

Table 3: Performance comparisons of removing different parts in gated-attention statistics pooling

system	SRE16, Pooled		SRE18, CMN2	
	EER%	minDCF	EER%	minDCF
X_GCNN+Gonly	7.85	0.588	9.27	0.615
X_GCNN+Aonly	7.75	0.594	9.28	0.602
X_GCNN+GAtt	<b>7.48</b>	<b>0.585</b>	<b>9.12</b>	<b>0.596</b>

## 5. Conclusions

In this study, we employ gating mechanisms in the DNN embedding system. More specially, the GCNN is introduced into the frame-level layers where the output representations of the GCNN layer are carefully modulated through both the output gate and memory cell. Such a mechanism helps to obtain more expressive features. Furthermore, we propose a gated-attention statistics pooling in which the attention model is enhanced by the gating mechanism by sharing the same weights with the output gate. The experimental results demonstrate that GCNN-based systems outperform the TDNN-based x-vector baseline, and our proposed gated-attention statistics pooling provides further improvement over the original attentive statistics pooling in GCNN-based embedding systems. Moreover, the proposed pooling method has obvious complementarity with the attentive statistics pooling, and the fused system achieves the best performance among all the mentioned systems.

## 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. 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.
- [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] H. Bredin, "Tristounet: triplet loss for speaker turn embedding," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 5430–5434.
- [10] Z. Gao, Y. Song, I. McLoughlin, W. Guo, and L. Dai, "An improved deep embedding learning method for short duration speaker verification," in *Proc. Interspeech 2018*, 2018, pp. 3578–3582. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1515>
- [11] Y. Zhu, T. Ko, D. Snyder, B. Mak, and D. Povey, "Self-attentive speaker embeddings for text-independent speaker verification," in *Proc. Interspeech 2018*, 2018, pp. 3573–3577. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1158>
- [12] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive 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>
- [13] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. International Convention Centre, Sydney, Australia: PMLR, 06–11 Aug 2017, pp. 933–941. [Online]. Available: <http://proceedings.mlr.press/v70/dauphin17a.html>
- [14] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [15] P. Chen, W. Guo, Z. Chen, J. Sun, and L. You, "Gated convolutional neural network for sentence matching," in *Proc. Interspeech 2018*, 2018, pp. 2853–2857. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-70>
- [16] C. Li, L. Li, and J. Qi, "A self-attentive model with gate mechanism for spoken language understanding," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018, pp. 3824–3833.
- [17] 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.
- [18] NIST, "Nist 2018 speaker recognition evaluation plan," <https://www.nist.gov/document/sre18evalplan2018-05-31v6pdf>.
- [19] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 448–456. [Online]. Available: <http://proceedings.mlr.press/v37/ioffe15.html>
- [20] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [21] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, and M. Devin, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016.
- [22] D. P. Kingma and J. L. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [23] 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.