



A Deep Neural Network for Short-Segment Speaker Recognition

Amirhossein Hajavi¹, Ali Etemad^{1,2}

¹Department of Electrical and Computer Engineering, Queens University, Kingston, Canada

²Ingenuity Labs, Queens University, Kingston, Canada

{a.hajavi, ali.etemad}@queensu.ca

Abstract

Today's interactive devices such as smart-phone assistants and smart speakers often deal with short-duration speech segments. As a result, speaker recognition systems integrated into such devices will be much better suited with models capable of performing the recognition task with short-duration utterances. In this paper, a new deep neural network, UtterIdNet, capable of performing speaker recognition with short speech segments is proposed. Our proposed model utilizes a novel architecture that makes it suitable for short-segment speaker recognition through an efficiently increased use of information in short speech segments. UtterIdNet has been trained and tested on the VoxCeleb datasets, the latest benchmarks in speaker recognition. Evaluations for different segment durations show consistent and stable performance for short segments, with significant improvement over the previous models for segments of 2 seconds, 1 second, and especially sub-second durations (250 ms and 500 ms).

Index Terms: Speaker Verification, Deep Neural Network, Short Segments

1. Introduction

Speaker recognition has seen profound improvements due to the recent advancements in deep learning. Accordingly, accuracy levels of proposed deep neural networks (DNN) for speaker recognition (both verification and identification) are far surpassing previous state-of-the-art techniques. Recent examples include the use of embeddings obtained from convolutional neural networks (CNN) for speaker recognition in [1, 2, 3], the use of auto-encoder models for speaker identification in [4, 5], and a number of cases utilizing ResNet for both speaker recognition and identification in [6, 7].

As discussed in the review [8], initial attempts at speaker recognition were performed under highly controlled conditions with very limited size of vocabulary. More recently, challenging factors such as environment noise, impersonation, and different ethnic diversities, are being tackled using techniques such as UBM-GMM and Joint Factor Analysis [9], as well as hybrid GMM and support vector machines [10]. Emergence of challenging datasets such as Speaker In The Wild (SITW) [11] and its extended variations such as VoxCeleb1 and VoxCeleb2 [1, 3], with more than 5,000 speakers and one million utterances, have enabled the opportunity to tackle speaker recognition in real-world scenarios. As a result, a number of deep learning solutions have been proposed for this purpose, including the models studied in [12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28].

Current solutions for speaker recognition either rely on the use of an entire utterance for carrying out the recognition task [20, 29] or require a segment of speech, often two seconds or greater, for accurate performance [1, 2, 3, 22]. However, with the emergence of voice-based interactive devices

such as smart-phone assistants, smart home devices such as smart speakers, in-vehicle entertainment and navigation systems, and other consumer electronics, it is imperative that accurate speaker recognition be performed using short-duration speech segments. Despite the fact that most existing speaker recognition solutions target medium or long utterances, namely 6 seconds or more, the mentioned target environments often deal with short-duration commands, such as "Hey Siri", "Okay Google", "Volume Up", or others. As a result, the aim of this work is to develop a DNN architecture capable of accurate speaker verification with short-segment utterances, specifically under 2 seconds.

In this study, a novel architecture has been proposed to tackle the problem of speaker verification using short-duration segments. Our proposed model is applied over a windowed duration of utterance, where very short segments can be used as inputs. The model creates an embedding of the received utterance, which can be accurately verified with respect to different individuals. Our proposed architecture is designed with the aim of preserving the information often lost through DNN pipelines by means of utilizing the information at multiple junctions and feeding that information to three fully connected networks (FCN). We use the VoxCeleb2 dataset for training, and evaluate our model with respect to state-of-the-art techniques such as [3], [22].

The rest of the paper is structured as follows. First, we go through the related work, followed by a detailed description of our proposed model architecture. Next, we evaluate the performance of the proposed method with respect to recent works in the field.

2. Related Work

DNNs were first used for speaker recognition in [12], and have since shown promising results by outperforming the traditional HMM-GMM techniques. Since the introduction of DNNs, attempts have made to incorporate such techniques with I-Vector methods in [13, 17]. Later on, further attempts were made to tackle speaker recognition with DNN under difficult conditions such as distant talking [14, 15]. This was followed by the integration of I-Vector with DNNs [30], which outperformed HMM-GMM-based methods, which were considered state-of-the-art at the time.

Next, novel deep neural network architectures such as auto-encoders [4, 5] found their way into speaker recognition. In these methods, auto-encoders were used to create embeddings in a lower dimension prior to reconstruction. These lower dimension embeddings, in the form of bottle-neck features, were used to discriminate speakers.

Recurrent neural networks (RNN) have been utilized in a number of studies. Recently, RNN models were employed in [23, 31, 25] with mel-frequency cepstrum coefficients (MFCC)

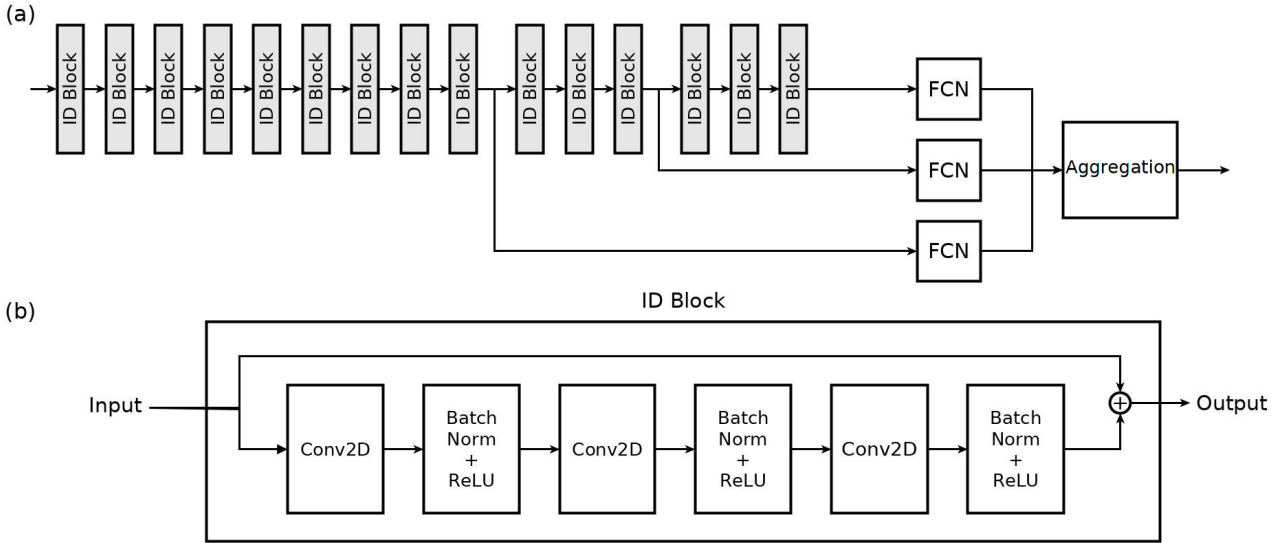


Figure 1: Architecture of our proposed model, UtterIdNet: (a) the overall scheme, followed by (b) the internal architecture of ID Blocks.

as inputs. In [25], a long short-term memory (LSTM) architecture was applied on MFCC, resulting in an embedding used to verify the speaker of the utterance by means of cosine distance. Other attempts, such as the model proposed in [23] and [31], have used the LSTM architecture as an intermediate tool in extracting i-vectors. More sophisticated models combining CNN and RNN-based solutions were proposed in [26] and [24], applying several convolution layers in between the MFCC input and the RNN.

Utilizing adversarial networks has also been explored for speaker recognition. Proposed solutions in [32, 33, 27, 27, 28] explore the use of adversarial networks and generative adversarial networks both as discriminative models for verification of the speaker as well as generative models. Such generative models were mostly used to transform the conditions of the utterance into more convenient environments in which to perform the speaker recognition task.

Many studies have utilized CNN-based models for speaker recognition. In [1] a CNN architecture was introduced, which outperformed I-Vector-based methods with a large margin. Using convolution and AveragePool/MaxPool layers, the model was able to achieve an improvement in accuracy of approximately 10%. More advanced CNN-based architectures were then used for speaker recognition. For example, considering successful performance in face recognition tasks, VGGnet [7] was used to tackle both identification and verification, outperforming simple CNN architectures. The ResNet architecture [2, 3, 22] was also utilized for speaker recognition, achieving the state-of-the-art by outperforming the previous VGGnet models. Continuing the trend of employing successful models from face recognition, a ResNet architecture proposed in [22] utilizing novel aggregation methods such as GhostVlad and NetVlad, outperformed the previous use of a 50 layer ResNet model [3].

3. Method

In this section, the feature extraction process and our proposed model are described:

3.1. Feature extraction

To extract informative frequency features from the short-segment utterances, Short-Time Fourier Transform (STFT) is used in this paper. STFT, as shown in Equation 1 is the Fourier transform of the signal under a time window of τ . In Equation 1, the $w(t - \tau)$ is the window operation applied over the signal to capture the frequency features within that specific duration. We used a window of 25 ms, with a stride of 10 ms, similar to [1, 2, 3, 22].

$$X(t, f) = \int_{-\infty}^{\infty} w(t - \tau)x(\tau)e^{-j2\pi f\tau} d\tau \quad (1)$$

3.2. Model Architecture

Our proposed model is inspired by two highly popular and successful deep neural network architectures, namely ResNet [34] and DeepID3 [35]. The main contribution made in [35] was driven by the idea of preserving information often filtered out through MaxPool layers in traditional CNN models. The solution provided in the DeepID network, and later in the enhanced DeepID3, was to feed the outputs of some of the MaxPool layers to an FCN, accompanied by a softmax layer. In their model, a voting mechanism was then used to aggregate the outputs of the softmax layers. The ResNet architecture, on the other hand, reformulated the layers as learning residual mechanism given the layer inputs. These residual networks were shown to be trained easier, while performing more accurately given the increase in depth.

Speaker recognition with short-segment utterances requires maximizing the use of the limited amount of information within the available input window. Furthermore, the large scale of the VoxCeleb2 dataset requires learning methods to be efficient in terms of training epochs and resources. As a result, we decided to use the properties of the aforementioned models to create a new architecture for the task at hand.

Our proposed architecture, UtterIdNet, is presented in Figure 1 (a). Similar to the DeepID3 model, the data flow has been branched in two other location besides the output of the final block. The internal architecture of an Identity (ID) Block can

be seen in Figure 1 (b). The ID Blocks in our model transform the input spectrograms to a level of individual-specific embeddings, followed by the FCNs for a final embedding transformation. Lastly, an aggregator applies a non-linear voting scheme over the FCN outputs. Details such as number of filters and kernel sizes are summarized in Table 1.

The input utterance segments go through 9 ID Blocks to reach the required level of speaker-specific discrimination. This depth was determined through empirical evaluation. The information achieved at this point is preserved through training an FCN through a branched route. Next, successive to three new ID Blocks, another similar FCN is trained, followed by the final three ID Blocks which lead to the last FCN. All the FCNs are trained using an Adam optimizer [36]. Once the utterances have gone through the entire UtterIdNet pipeline, as discussed, three embedding vectors are generated by the FCNs, which contain discriminative speaker-specific features.

3.3. Aggregation

A non-linear aggregator combines the outputs of the three FCNs to produce a 512 dimensional representation of the individuals within the dataset. The non-linear combination used to aggregate the FCN outputs is provided in Equation 2,

$$\rho = \tanh(W_1 \times Em_1 + W_2 \times Em_2 + W_3 \times Em_3), \quad (2)$$

where W_1 , W_2 , and W_3 represent weight vectors, and Em denotes the FCN embeddings. In this paper these weights are set to be trained for each dimension within the embeddings individually through an Adam optimizer with a softmax function. If all the indices within each W were selected using a uniform function, our aggregator would act as the same voting mechanism originally implemented for DeepId3 [35].

Successive to extraction of embeddings for each short utterance segment, a simple aggregation scheme, Time-Distributed Voting (TDV) is utilized. This aggregation technique collects embedding vectors for successive durations of the input utterance, and carries out a voting mechanism to verify the speaker identity. The reason for utilizing a simple aggregator was the fact that our solution is aimed at speaker recognition for short-segment utterances. Hence, it is likely that by employing more advanced aggregation techniques, better performance is achieved on longer utterance segments.

4. Experiments and Results

In this section, we describe the datasets and the experimental setup utilized for speaker verification. Next, the results obtained by our model using different short segment durations are presented, and compared to the state-of-the-art [22]. Furthermore, in addition to short-segment windows, the performance of the proposed model on full utterances has been evaluated with respect to recent successful techniques. Lastly, the potential of UtterIdNet for real-time applications is explored, followed by a discussion on required memory resources for potential edge-device applications.

4.1. Datasets

The VoxCeleb1 [1] and VoxCeleb2 [3] datasets have been collected through automatic pipelines from open source media, and contain 1,250 and 5,994 speakers respectively. The utilized pipeline was based on computer vision techniques such as face recognition and active speaker detection. Due to the datasets

Table 1: Architectural details of our proposed model.

Module	Input Spectrogram ($257 \times T \times 1$)	Output Size
ConvModule	Conv2D, 7×7 , 64	$257 \times T \times 64$
	Maxpool, 2×2 , stride (2, 2)	$128 \times T/2 \times 64$
	conv, 1×1 , 48 conv, 3×3 , 48 conv, 1×1 , 96	$\times 2$ $128 \times T/2 \times 96$
	conv, 1×1 , 96 conv, 3×3 , 96 conv, 1×1 , 128	$\times 3$ $64 \times T/4 \times 128$
	conv, 1×1 , 128 conv, 3×3 , 128 conv, 1×1 , 256	$\times 3$ $32 \times T/8 \times 256$
	conv, 1×1 , 256 conv, 3×3 , 256 conv, 1×1 , 512	$\times 3$ $16 \times T/16 \times 512$
	Maxpool, 3×1 , stride (2, 2)	$7 \times T/32 \times 512$
Embedding	FCN $\times 3$, 512	1526
Nonlinear aggregation	FCN, 512	512

being collected from real media content, *in the wild* scenarios were maintained in these datasets.

VoxCeleb2, as discussed by the authors in [3], contains several flaws in its annotations. Therefore, its use in testing of models has not been advised, while it has been widely used for training purposes. VoxCeleb1, however, has been collected under extremely strict constraints, leading the dataset to be free of any annotation issues. As a result, models are often tested on this dataset as a benchmarking practice. Accordingly, in this paper, the VoxCeleb2 with over 1 Million utterances has been used for training, while VoxCeleb1 was used as the test set for our proposed model.

4.2. Training

Training of the proposed model was performed using a standard softmax loss. For optimization, an Adam optimizer with a decaying learning rate and an initial value of 10^{-4} was utilized. The learning rate decay was a factor of 10 for every 36 epochs. Training was done on a single Nvidia 1080 Ti GPU.

4.3. Testing and verification

Table 2 presents the results. In this table, the performance of UtterIdNet with respect to Thin ResNet34 + GhostVlad [22] is provided. In order to evaluate the impact of segment durations on performance, 250 ms, 500 ms, 1 sec, and 2 sec segment sizes were used. UtterIdNet outperforms Thin ResNet34 + GhostVlad in all the different short-segment scenarios. Notably, the performance of UtterIdNet shows reasonable amount of consistency and stability as the segment size is decreased. As a result, the proposed UtterIdNet model outperforms ResNet34 + GhostVlad by a considerable margin for sub-second segments and especially for 250 ms segments. The performance of other methods such as x-vector [29] or ResNet50 [3] have not been reported for short segments. Moreover, to the best of our knowledge, these models were not publicly available, and therefore a direct comparison was not possible.

To further evaluate our method, we also compared the performance on full-utterance inputs. The results of this experiment are presented in Table 3. As illustrated in this table, UtterIdNet is closely outperformed by [29], [3], and [22]. Nonetheless, given that TDV was kept simple for high effi-

Table 2: The performance of our model, UtterIdNet, with respect to [22], trained with VoxCeleb2 and tested on the VoxCeleb1 dataset. Our model performs consistently as the segment durations are decreased, outperforming [22] with considerable margins for sub-second segments.

	Model	Aggregation	EER % 250 ms	EER % 500 ms	EER % 1 sec	EER % 2 sec
Xie et al. [22]	Thin ResNet34	GhostVlad	23.23	10.58	9.25	7.97
Ours	UtterIdNet	TDV	6.88	6.46	6.41	6.33

Table 3: The result of UtterIdNet for full-length utterances. For longer segments, our model is outperformed by [29], [3], and [22].

	Model	Loss	Dims	Train set	EER %
Nagrani et al. [2]	I-Vector + PLDA	–	–	VoxCeleb1	8.80
Cai et al. [2]	ResNet34 + SAP	A-softmax + PLDA	128	VoxCeleb1	4.40
Cai et al. [2]	ResNet34 + LDE	A-softmax + PLDA	128	VoxCeleb1	4.48
Okabe et al. [29]	TDNN (X-Vector) + TAP	softmax	1500	VoxCeleb1	3.85
Hajibabai et al. [37]	ResNet20 + TAP	AM-softmax	128	VoxCeleb1	4.30
Chung et al. [3]	ResNet50 + TAP	softmax + Contrastive	512	VoxCeleb2	4.19
Xie et al. [22]	Thin ResNet34 + TAP	softmax	512	VoxCeleb2	10.48
Xie et al. [22]	Thin ResNet34 + GhostVlad	softmax	512	VoxCeleb2	3.22
Ours	UtterIdNet + TDV	softmax	512	VoxCeleb2	4.26

ciency for short-segment utterances, this margin may be reduced or eliminated with more advanced aggregation techniques.

Using the gradient sharing approach enabled through our proposed method, convergence was achieved at a much quicker rate than that of other recent DNN solutions. This was observed and evident by the performance with respect to the number of training epochs. In [22] it was reported the converged performance was obtained at 48 epochs, while UtterIdNet produced the results in Tables 2 and 3 after only 25 epochs. Moreover, for the same number of epochs (25), the full-length utterance performance of [22] appears to be 16% EER, which is considerably outperformed by UtterIdNet.

4.4. Identification

The latest reported identification rate for VoxCeleb1 and VoxCeleb2 datasets has been 89.5% [7]. This performance was achieved using *full-length utterances*, and on VoxCeleb1 which contained 1,250 speakers. We tested UtterIdNet for identification with 250 ms segments on the VoxCeleb2 validation set, which contained 5,994 speakers. Despite shorter segment sizes and a significantly larger number of speakers (approximately 4 times), we achieved an accuracy of 84.3%. These results show significant promise and the potential for considerable improvement in speaker identification in addition to verification, which we intend to study in future works.

4.5. Memory requirements

In order for the proposed model to be utilized in edge nodes such as smart-phones and Internet of Things (IoT) devices, certain memory constraints will need to be met. Given the fact that our model consists of 3 FCN components and 15 ID Blocks, to address the concerns regarding the memory requirements of such a system and the feasibility of its implementation in the

described context, we evaluated the required memory resources during run-time. The obtained estimated memory need for UtterIdNet was 268 MB. This amount of required memory is feasible for many current consumer edge devices, especially in the context of smart-phones, smart homes, and vehicles. As a result, we believe UtterIdNet can be widely employed in these contexts for low-latency recognition and verification of speakers.

5. Conclusions

In this paper a new DNN, UtterIdNet, was proposed with the aim of an efficient increase of information use for short speech segments. We evaluated our method using the VoxCeleb datasets, and demonstrated that UtterIdNet outperforms the state-of-the-art for short segments. Specifically, we utilized 250 ms, 500 ms, 1 sec, and 2 sec segments, where UtterIdNet showed significant improvement in the sub-second segment range. The efficient learning of information in our proposed model was not only evident by the considerable improvement in accuracy, but also by the fact that in order to train properly, UtterIdNet took approximately half the number of epochs as the state-of-the-art. While UtterIdNet was outperformed by a small margin in full-utterance segments, we believe the choice of the simple aggregation technique for combining the different short segments within the full utterance could be a contributing factor, which we intend to investigate in future work.

6. Acknowledgements

The authors would like to thank IMRSV Data Labs for their support of this work. The authors would also like to acknowledge the Natural Sciences and Engineering Research Council of Canada (NSERC) for supporting this research (grant number: CRDPJ 533919-18).

7. References

- [1] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A Large-Scale Speaker Identification Dataset," *INTERSPEECH*, pp. 2616–2620, 2017.
- [2] W. Cai, J. Chen, and M. Li, "Analysis of Length Normalization in End-to-End Speaker Verification System," *INTERSPEECH*, pp. 3618–3622, 2018.
- [3] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *INTERSPEECH*, pp. 1086–1090, 2018.
- [4] Z. Zhang, L. Wang, A. Kai, T. Yamada, W. Li, and M. Iwahashi, "Deep Neural Network-based Bottleneck Feature and Denoising Autoencoder-based Dereverberation for Distant-talking Speaker Identification," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2015, no. 1, p. 12, 2015.
- [5] S. S. Tirumala and S. R. Shahamiri, "A Deep Autoencoder Approach for Speaker Identification," *the 9th International Conference on Signal Processing Systems*, pp. 175–179, 2017.
- [6] S. Novoselov, V. Shchemelinin, A. Shulipa, A. Kozlov, and I. Kremnev, "Triplet Loss Based Cosine Similarity Metric Learning for Text-independent Speaker Recognition," *INTERSPEECH*, pp. 2242–2246, 2018.
- [7] S. Yadav and A. Rai, "Learning Discriminative Features for Speaker Identification and Verification," *INTERSPEECH*, pp. 2237–2241, 2018.
- [8] D. A. Reynolds, "An Overview of Automatic Speaker Recognition Technology," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 4, pp. 4072–4075, 2002.
- [9] P. Kenny, "Joint Factor Analysis of Speaker and Session Variability: Theory and Algorithms," *CRIM, Montreal, (Report) CRIM-06/08-13*, vol. 14, pp. 28–29, 2005.
- [10] R. Togneri and D. Pullella, "An Overview of Speaker Identification: Accuracy and Robustness Issues," *IEEE Circuits and Systems Magazine*, vol. 11, no. 2, pp. 23–61, 2011.
- [11] M. McLaren, L. Ferrer, D. Castan, and A. Lawson, "The Speakers in the Wild SITW Speaker Recognition Database," *INTERSPEECH*, pp. 818–822, 2016.
- [12] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, and T. N. Sainath, "Deep Neural Networks for Acoustic Modeling in Speech Recognition: The Shared Views of Four Research Groups," *IEEE Transactions on Signal Processing*, vol. 29, no. 6, pp. 82–97, 2012.
- [13] Y. Song, B. Jiang, Y. Bao, S. Wei, and L.-R. Dai, "I-vector Representation Based on Bottleneck Features for Language Identification," *Electronics Letters*, vol. 49, no. 24, pp. 1569–1570, 2013.
- [14] T. Yamada, L. Wang, and A. Kai, "Improvement of Distant-talking Speaker Identification Using Bottleneck Features of DNN," *INTERSPEECH*, pp. 3661–3664, 2013.
- [15] Y. Lei, N. Scheffer, L. Ferrer, and M. McLaren, "A Novel Scheme for Speaker Recognition Using a Phonetically-aware Deep Neural Network," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1695–1699, 2014.
- [16] Y. Lei, L. Ferrer, A. Lawson, M. McLaren, and N. Scheffer, "Application of Convolutional Neural Networks to Language Identification in Noisy Conditions," *Odyssey*, vol. 41, pp. 1–8, 2014.
- [17] P. Kenny, V. Gupta, T. Stafylakis, P. Ouellet, and J. Alam, "Deep Neural Networks for Extracting Baum-Welch Statistics for Speaker Recognition," *Odyssey*, pp. 293–298, 2014.
- [18] P. Matejka, L. Zhang, T. Ng, H. S. Mallidi, O. Glembek, J. Ma, and B. Zhang, "Neural Network Bottleneck Features for Language Identification," *Odyssey*, vol. 2014, pp. 299–304, 2014.
- [19] I. Lopez-Moreno, J. Gonzalez-Dominguez, O. Plchot, D. Martinez, J. Gonzalez-Rodriguez, and P. Moreno, "Automatic Language Identification Using Deep Neural Networks," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5337–5341, 2014.
- [20] O. Ghahabi and J. Hernando, "Deep Learning Backend for Single and Multi-session i-Vector Speaker Recognition," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 4, pp. 807–817, 2017.
- [21] S. Shon, H. Tang, and J. Glass, "Frame-level Speaker Embeddings for Text-independent Speaker Recognition and Analysis of end-to-end model," *arXiv preprint arXiv:1809.04437*, 2018.
- [22] W. Xie, A. Nagrani, J. S. Chung, and A. Zisserman, "Utterance-level Aggregation For Speaker Recognition In The Wild," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [23] H. Park, S. Cho, K. Park, N. Kim, and J. Park, "Training Utterance-level Embedding Networks for Speaker Identification and Verification," *INTERSPEECH*, pp. 3563–3567, 2018.
- [24] J.-w. Jung, H.-s. Heo, I.-h. Yang, H.-j. Shim, and H.-j. Yu, "Avoiding Speaker Overfitting in End-to-End DNNs Using Raw Waveform for Text-Independent Speaker Verification," *INTERSPEECH*, pp. 3583–3587, 2018.
- [25] J. Wang, K.-C. Wang, M. Law, F. Rudzicz, and M. Brudno, "Centroid-based Deep Metric Learning for Speaker Recognition," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [26] R. Kumar, V. Yeruva, and S. Ganapathy, "On Convolutional LSTM Modeling for Joint Wake-Word Detection and Text Dependent Speaker Verification," *INTERSPEECH*, pp. 1121–1125, 2018.
- [27] P. S. Nidadavolu, J. Villalba, and N. Dehak, "Cycle-GANs for Domain Adaptation of Acoustic Features for Speaker Recognition," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6206–6210, 2019.
- [28] J. Rohdin, T. Stafylakis, A. Silnova, H. Zeinali, L. Burget, and O. Plchot, "Speaker Verification Using End-to-end Adversarial Language Adaptation," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6006–6010, 2019.
- [29] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive Statistics Pooling for Deep Speaker Embedding," *INTERSPEECH*, pp. 2252–2256, 2018.
- [30] T. Kinnunen, L. Juvola, P. Alku, and J. Yamagishi, "Non-parallel Voice Conversion Using i-Vector PLDA: Towards Unifying Speaker Verification and Transformation," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5535–5539, 2017.
- [31] A. Jati and P. Georgiou, "An Unsupervised Neural Prediction Framework for Learning Speaker Embeddings Using Recurrent Neural Networks," *INTERSPEECH*, pp. 1131–1135, 2018.
- [32] W. Ding and L. He, "MTGAN: Speaker Verification through Multitasking Triplet Generative Adversarial Networks," *INTERSPEECH*, pp. 3633–3637, 2018.
- [33] X. Fang, L. Zou, J. Li, L. Sun, and Z.-H. Ling, "Channel Adversarial Training for Cross-channel Text-independent Speaker Recognition," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [34] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- [35] Y. Sun, D. Liang, X. Wang, and X. Tang, "DeepID3: Face Recognition with Very Deep Neural Networks," *arXiv:1502.00873 [cs]*, 2015.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [37] M. Hajibabaei and D. Dai, "Unified Hypersphere Embedding for Speaker Recognition," *arXiv preprint arXiv:1807.08312*, 2018.