



Speaker embeddings by modeling channel-wise correlations

Themos Stafylakis¹, Johan Rohdin^{1,2}, Lukáš Burget²

¹Omilia - Conversational Intelligence, Athens, Greece

²Brno University of Technology, Faculty of Information Technology, Speech@FIT, Czechia

{tstafylakis}@omilia.com, {rohadin,burget}@fit.vutbr.cz

Abstract

Speaker embeddings extracted with deep 2D convolutional neural networks are typically modeled as projections of first and second order statistics of channel-frequency pairs onto a linear layer, using either average or attentive pooling along the time axis. In this paper we examine an alternative pooling method, where pairwise correlations between channels for given frequencies are used as statistics. The method is inspired by style-transfer methods in computer vision, where the style of an image, modeled by the matrix of channel-wise correlations, is transferred to another image, in order to produce a new image having the style of the first and the content of the second. By drawing analogies between image style and speaker characteristics, and between image content and phonetic sequence, we explore the use of such channel-wise correlations features to train a ResNet architecture in an end-to-end fashion. Our experiments on VoxCeleb demonstrate the effectiveness of the proposed pooling method in speaker recognition.

Index Terms: speaker recognition, style-transfer, deep learning

1. Introduction

Extracting low-dimensional representations from speech utterances that characterize the speaker is a challenging task with numerous applications. Apart from speaker recognition, speaker representations are used in automatic speech recognition (e.g. for speaker adaptation), in speaker diarization and separation, and in speech synthesis, for mimicking a target speaker using text-to-speech or voice conversion.

Until recently, *i*-vectors were the most commonly used means for representing speaker characteristics [1, 2]. With the advent of deep learning, neural architectures such as time-delay neural networks (TDNNs [3]), Long Short-Term Memory networks (LSTMs), 2D convolutional neural networks (such as ResNets and its extensions [4]) and Transformers have emerged, surpassing the performance of *i*-vectors in most applications and speaker recognition settings.

In order to obtain fixed-size representation vectors from utterances of variable duration and arbitrary word sequences (i.e. text-independent), a pooling method is required. The most common pooling method is a simple aggregation over the time axis, where the mean vector is concatenated with standard deviation (known as statistics pooling layer). The speaker embedding (a vector of few hundred dimensions) is obtained by projecting this pooled representation (a vector of few thousand dimensions) onto a linear layer or a shallow feed-forward neural network. Other pooling approaches have been proposed, such as single and multi-head attention, as well as the use of higher-order statistics [5, 6, 7].

In this paper we present an alternative way of modeling speaker information, that is largely inspired by style-transfer [8]. This seminal algorithm in computer vision is capable of

producing new images that combine the content of an image with the style of other images or artworks. The style of an image is modeled by the correlations between the channels of the convolutional neural network. By considering styles as broadly analogous to speakers, we experiment with modeling speakers by channel-wise correlation. This is in contrast to averaging along the time dimension of each channel frequency pair of the output tensor that is used in statistics pooling. The other sources of variability that may also be considered as style of the recording, such as noise, channel, and emotion, can be suppressed by training the model end-to-end using supervised learning with speakers as targets. The resulting embeddings are highly speaker-discriminative, yielding state-of-the-art results on VoxCeleb using *plain* cosine similarity scoring.

The rest of the paper is organized as follows. In Sect. 2 we briefly review some of the related work, such as pooling methods and style-transfer. In Sect. 3 we introduce the proposed method, and provide all the implementation and architectural details. In Sect. 4, we present our experiments on a standard speaker recognition benchmark (VoxCeleb). Finally, in Sect. 5 we provide conclusions and future work directions.

2. Related work

2.1. Style-transfer in computer vision and voice conversion

A few years ago, the seminal work of L.A. Gatys *et al.* [8] in neural style-transfer gained the attention of the computer vision community and the industry, and received wide media coverage. Since then, several works extended their method, some of which can be found in [9]. A theoretical interpretation of the way style and image texture is encoded in the channel-wise correlations (which can be considered as a Gram-Matrix) is given in [10]. The authors treat style-transfer as a domain adaptation problem, and prove the equivalence between trying to match the Gram-Matrices of feature maps and minimizing the Maximum Mean Discrepancy (MMD) with the second-order polynomial kernel.

One of the few attempts to apply neural-style transfer in speech synthesis is described in [11]. The authors employ a pretrained CNN and apply style-transfer to train a speech generation and voice conversion system. Moreover, they demonstrate the capacity of the intermediate tensors to encode speaker characteristics using the VCTK dataset. However, the scope of their work is different to ours, the network is not trained end-to-end while the representations are high-dimensional. Furthermore, their method is not tested on a speaker recognition benchmark containing noise and high intrinsic and extrinsic within-speaker session variability (such as VoxCeleb or NIST-SRE). Hence, the ability of their utterance-level representation to suppress session variability has not been examined.

2.2. Pooling layers in speaker embeddings

The vast majority of speaker embedding networks use the statistics pooling layer introduced by x-vectors, i.e. mean and standard deviation (std) for each frequency and channel combination [12]. The main family of alternative pooling methods enhance these statistics with an attention mechanism. Examples are single and multi-head attention, as well as Net and Ghost-VLAD [6, 13, 14]. Both attention and VLAD methods employ a set of heads (i.e. trainable network parameters which are broadly analogous to the means of a mixture model) to cluster the output features into broad phonetic categories in an unsupervised and end-to-end fashion. The resulting statistics pooling is defined as the concatenation of the per-head statistics. Their main difference lies in the normalization; the attention normalizes only over the time axis while the VLAD over the attentive heads. Second order statistics (i.e. std features) may also be added as in the case of average pooling. More recent approaches exist, such as vector-based attentive pooling [15], phonetically-aware attention [16], and cross-attention [17].

A second alternative to standard statistics pooling is the use of higher-order statistics (skewness and kurtosis) [7]. Skewness is a measure of the asymmetry of a distribution with respect to its mode, while the kurtosis measures its “tailedness”. The authors did not obtain any consistent improvement compared to using merely first and second statistics, as in standard statistics pooling. Recently (and concurrently to our work) a set of alternative pooling methods were examined in [18]. The authors show that the use of second order statistics is more effective compared to first order, while they also perform experiments with using the full covariance matrix, without showing improvements over the standard statistics pooling. Their work is close to ours, although we differ in several ways, such as in the suggested pooling method and in the several implementation details (e.g. instead of covariances we propose frequency-dependent channel-wise correlation matrices).

3. The proposed method

3.1. Notation

We denote by $\mathbf{X} = \{\mathbf{x}_t\}_{t=1}^{T_i}$ the sequence of acoustic features of an utterance with T_i frames, which in our method are $F_i = 80$ dimensional filterbank features. The architecture is a ResNet-34 and the output of the last ResNet block is a 3D tensor (we omit the batch axis) denoted by $\mathbf{Y} \in \mathbf{R}^{T \times F \times C}$, where T denotes the size of the temporal axis, F the size of the frequency axis, and C the number of channels. Typical values for these quantities are $(T, F, C) = (50, 10, 256)$; note though that T is variable and can be derived by T_i divided by the cumulative temporal stride (equal to $2^3 = 8$ in our architecture). Similarly, $F = 10$ since the cumulative frequency stride is again equal to 8, while $F_i = 80$.

3.2. Statistics pooling

The statistics pooling layer in speaker embeddings networks with 2D CNN architectures is a concatenation of the mean and std of each of the $F \times C$ frequency-channel pairs $[\mathbf{P}^m; \mathbf{P}^s]$, where

$$\mathbf{P}_{f,c}^m = \frac{1}{T} \sum_t \mathbf{Y}_{t,f,c} \quad (1)$$

and

$$\mathbf{P}_{f,c}^s = \sqrt{\frac{1}{T} \sum_t (\mathbf{Y}_{t,f,c} - \mathbf{P}_{f,c}^m)^2} \quad (2)$$

The speaker embeddings are extracted by projecting the (vectorized) pooling layer onto a lower dimensional space (e.g. $d_e = 256$) using a single linear layer.

3.3. Style modeling in images

The style of an image in computer vision for a given convolutional layer is modeled by the symmetric positive-semidefinite matrix \mathbf{S} with elements defined as

$$\mathbf{S}_{c,c'} = \frac{1}{TF} \sum_{t,f} \mathbf{Y}_{t,f,c} \mathbf{Y}_{t,f,c'} \quad (3)$$

Note though that images differ from spectral representations of acoustic signals. Image characteristics are considered invariant along both spatial axes, while spectral representations are invariant only along the time axis. In fact, this property explains the use of frequency-channel pairs in statistics pooling instead of channels alone, as shown in (1)-(2). Due to this property of spectral representations, pooling along both time and frequency axes as in (3) is expected to yield less discriminative speaker representations.

3.4. Pooling via frequency-dependent channel-wise correlations

A natural modification of (3) that addresses the above issue is to pool only along the time axis, i.e.

$$\mathbf{S}_{f,c,c'} = \frac{1}{T} \sum_t \mathbf{Y}_{t,f,c} \mathbf{Y}_{t,f,c'} \quad (4)$$

The 3D tensor \mathbf{S} can be viewed as a list of F symmetric and positive semidefinite matrices $\mathbf{S}_f = \mathbf{S}_{f,:,:}$ and has therefore $FC(C+1)/2$ free variables. As its size can be too large we propose to reduce it in a frequency-dependent manner as shown later.

A further proposed modification is the use of mean and variance normalization along the time axis of \mathbf{Y} for each frequency-channel pair, i.e. $\mathbf{Z}_{:,f,c} \stackrel{norm}{\leftarrow} \mathbf{Y}_{:,f,c}$. When variance normalization is applied, the variables in the diagonal of \mathbf{S}_f (i.e. the variance of each channel-frequency combination) become equal to 1, and therefore they can be removed from the pooling layer resulting in $C(C-1)/2$ free variables for each frequency f and $FC(C-1)/2$ overall. Moreover, $\mathbf{S}_{f,c,c'} \in [-1, 1]$ for all f, c, c' . The pooling layer is followed by a linear layer which reduces the size of the representation to $d_e = 256$ (the speaker embedding space), as in the standard statistics pooling.

Note that when mean and variance normalization are applied, the pooled representation encodes complementary information to that of the standard pooling; the means of each channel-frequency are being subtracted, while their variance is also set equal to one.

3.5. Additional extensions and modifications

In the pooling method described above, we add the following operations in order to decrease the size of the pooling layer and regularize the network. Recall that regularization is needed in order to extract embeddings that generalize well to speakers unseen during training.

3.5.1. Channel dropout

Dropout is a standard choice for regularizing a network. We employ it in a channel-wise fashion, that is, we sample a 1D mask

of size C' (with dropout probability $p_d = 0.25$) and we broadcast it along the time and frequency axes. Channel-wise dropout is chosen because of its Bayesian interpretation as model averaging between models with variable number of convolutional filters [19].

3.5.2. Frequency ranges

As neighboring frequency bins in \mathbf{Y} have similar statistics, we consider merging them into F_r frequency ranges. After experimentation, we concluded that merging $f_r = 2$ consecutive frequency bins in a non-overlapping manner yields more stable performance. Merging is implemented by reshaping \mathbf{Y} , so that $F_r \leftarrow F f_r^{-1}$ and $T_r \leftarrow T f_r$. For example, when experimenting with $f_r = 2$, we obtain $F_r = 5$ and $T_r = 100$ given that $F = 10$ and $T = 50$. Note that by setting $f_r = F$ we recover the original style-transfer used in computer vision, where pooling is performed along both spatial axes.

3.5.3. Frequency-dependent channel reduction

The number of channels in \mathbf{Y} is $C = 256$ which would result in $C(C-1)/2 \approx 33K$ variables per frequency bin. So we transform \mathbf{Y} as follows

$$\mathbf{Y}_{t,f,c'} \leftarrow \sum_c \mathbf{L}_{f,c,c'} \mathbf{Y}_{t,f,c}. \quad (5)$$

where \mathbf{L} is a learnable tensor of shape (F_r, C, C') , acting on the channel axis and reducing its size from C to C' (we choose $C' = 64$). The rationale for using a 3D kernel instead of a 2D is that different linear combinations of channels should be better suited for obtaining speaker discriminative information from each frequency range.

3.6. Summary and order of operations

We summarise here the set of operations and their order we use in order to extract speaker embeddings. Recall that the output tensor of the last ResNet block is \mathbf{Y} with shape (T, F, C) .

1. **Dropout** We apply channel-wise dropout with p_d by sampling a binary mask of shape C and broadcasting it along the other two axes.
2. **Frequency ranges** We reshape the tensor \mathbf{Y} with new shape (T_r, F_r, C) , by merging f_r consecutive frequency bins into F_r frequency ranges.
3. **Channel Reduction** We reduce the channel size by applying \mathbf{L} as shown in (5), i.e. $\mathbf{Y} \xleftarrow{\mathbf{L}} \mathbf{Y}$, where the shape of \mathbf{Y} becomes (T_r, F_r, C') .
4. **Normalization** We apply mean and variance normalization along the time axis, i.e. $\mathbf{Z} \xleftarrow{norm} \mathbf{Y}$.
5. **Pooling** We apply pooling as shown in (4), but using \mathbf{Z} instead of \mathbf{Y} , by which we obtain \mathbf{S} , a 3D tensor with shape (F_r, C', C') .
6. **Flattening** We flatten the pooled representation to a vector of size $d_p = F_r C' (C' - 1)/2$, which corresponds to the number of free variables in \mathbf{S} .
7. **Embedding** We reduce the size of the pooled representation vector to d_e using a linear layer. The resulting vector is the speaker embedding.

4. Experiments

4.1. VoxCeleb 1 & 2 datasets

We evaluate the systems on the three VoxCeleb test sets, namely the original VoxCeleb 1 *test* set (composed of only 40 speakers) [20] and the (much larger and representative) *Extended* and *Hard* sets [14]. As training set, we use a subset of the VoxCeleb 2 *development* set which we augmented with babble, noise, reverberation, and music using Kaldi [21, 3]. The training set contains 5750 speakers and approximately 4.9M utterances including the augmented versions. As a *held-out* set we used 455 utterances (plus augmentations) from 90 training speakers.

4.2. Architecture, training, and loss function

The backbone of the network is a 34-layer ResNet. All convolutional kernels are 3×3 , the strides are $(1, 2, 2, 2)$ for the 4 blocks in both time and frequency axes, while the number of channels is $(64, 128, 256, 256)$ (the first convolutional layer also outputs 64 channels). The number of convolutional layers per block is $(3, 4, 6, 3)$. Finally, Squeeze and Excitation layers (with reduction ratio $r = 4$) are added only to the first 2 blocks of the ResNet [22, 23].

The network is trained using multi-speaker classification. As optimization criterion we employ the Additive Angular Margin (AAM) loss, which has shown notable improvements over the plain Cross Entropy loss [24]. We set the scale coefficient of the loss equal to 30, while for the margin coefficient we adopt a curriculum learning approach and progressively increase it as $(0.1, 0.2, 0.3)$ during training [25].

As optimizer we use stochastic gradient descent with momentum equal to 0.9. The minibatch size is 256, however to train the model with a single GPU we split the minibatch into 16 ‘‘microbatches’’ of 16 examples each and use gradient accumulation. The initial learning rate (LR) is equal to 0.2 which we divide by 2 when the loss does not improve for more than 3000 model updates in the held-out set (the final LR is $0.2/64$). The experiments are conducted using TensorFlow [26].

We emphasize that the model architecture, the training method, and its hyperparameters are optimized with the baseline, where mean and std features are used for pooling. Moreover, we did not observe any differences in time required per epoch between the baseline and the proposed architectures.

4.3. Experimental Results

To evaluate our method we experiment we examine several configurations and ablations. The results we obtained are given in Table 1, and are obtained using cosine-similarity, without score or any other normalization. For the two baseline experiments we train and evaluate a ResNet-34 with the standard temporal pooling. The first includes mean and std features while the second only std (suggested in [18]). We emphasize that for the baseline experiments the tensor operations are not included, i.e. mean and std pooling is performed directly on the output of the last ResNet block.

The first experiment using a variant of the proposed pooling (P1) is essentially a direct implementation of the pooling as suggested in image style-transfer (recall that pooling with $F_r = 1$ is equivalent to (3), i.e. to pooling along both time and frequency axes.). Two such networks are examined, one with mean and a second one with mean and variance normalization, denoted by P1 and P5, respectively. Note that when mean-only normalization is applied, the number of free-parameters is $F_r C_r (C_r + 1)/2$, i.e. the size of the pooled representa-

Table 1: Results on the Original, Hard, and Extended VoxCeleb test sets.

System						VoxCeleb-O		VoxCeleb-H		VoxCeleb-E	
						minDCF	EER(%)	minDCF	EER(%)	minDCF	EER(%)
B1	Baseline	(mean & std pooling)				0.091	1.40	0.145	2.48	0.090	1.43
B2	Baseline	(std pooling)				0.104	1.56	0.154	2.61	0.099	1.54
		C'	F_r	L	Normalization						
P1	Proposed	128	1	2D	mean	0.126	1.74	0.177	2.93	0.114	1.74
P2	Proposed	64	5	2D	mean	0.106	1.66	0.166	2.80	0.104	1.67
P3	Proposed	64	5	3D	mean	0.089	1.55	0.148	2.49	0.092	1.48
P4	Proposed	64	10	3D	mean	0.109	1.70	0.159	2.70	0.103	1.62
P5	Proposed	128	1	2D	mean & var.	0.102	1.56	0.158	2.63	0.101	1.57
P6	Proposed	64	5	2D	mean & var.	0.095	1.46	0.143	2.40	0.091	1.41
P7	Proposed	64	5	3D	mean & var.	0.071	1.16	0.128	2.17	0.079	1.22
P8	Proposed	64	10	3D	mean & var.	0.085	1.41	0.137	2.33	0.087	1.39
B1s	B1 with adaptive score normalization (as-norm)					0.079	1.26	0.127	2.23	0.080	1.31
P7s	P7 with adaptive score normalization (as-norm)					0.059	1.07	0.115	1.99	0.071	1.13

tion vector is increased by $F_r C_r$. By comparing P1 and P5 we observe that variance normalization helps significantly the model to attain results that are very close to the baseline B1. We should also note that for the experiments with $F_r = 1$, we set $C' = 128$ in order to have pooling layers of same order of magnitude, since their size grows linearly with F_r .

In the next set of experiments we use either $F_r = 5$ or $F_r = 10$, where the latter indicates no reshaping of the tensor \mathbf{Y} (as it corresponds to $f_r = 1$). The results confirm that the frequency axis should be treated differently compared to the time axis, and pooling should result in a 3D tensor, as shown in (4). They moreover demonstrate that reshaping the tensor by a factor of $f_r = 2$ (i.e. $F_r = 5$) yields the best performance.

We may also observe that the suggested frequency-dependent channel reduction (i.e. using a 3D tensor L instead of a 2D) yields consistently improved performance (by comparing P2 with P3, and P6 with P7). Allowing each frequency range to have its own weighted combination of channels increases the expressiveness of the pooled representation. We also mention that we attempted to combine our best system (P7) with B1 in the pooling layer. The improvement in performance attained was small. For example the EER dropped by 0.03% and the minDCF by 0.002 in both in VoxCeleb-E and VoxCeleb-H. Although further experimentation is required, it suggests that combining the correlations-based with mean and std features does not yield any notable improvement. We also compare B1 with P7 using score normalization (as-norm). We observe that the gains attained by the suggested pooling (P7s) over mean-std pooling (B1s) remain even when as-norm is applied.

Finally, all experiments underline the importance of mean and variance normalization (P5-8), as opposed to mean-only (P1-4). In other words, transforming the Gram-Matrices to correlation matrices results in more speaker-discriminative representations. This finding challenges the common view of frame aggregation being the natural means of extracting speaker-relevant information from utterances. Correlations, as *intrinsically normalized* statistical quantities, should be more resilient as features to the diverse factors of within-speaker variability. Temporal aggregation may not be the most effective way to remove such undesired sources of variability. Removing them via modeling correlations between different representations of the utterances (such as time-channel maps of given frequency ranges) may attain this goal more effectively and naturally.

5. Conclusions and research directions

In this paper, we introduced a new way of pooling information extracted from a CNN for modeling speaker information. The approach is inspired by neural style-transfer in computer vision, where style is modeled by a Gram-Matrix of channel-wise dot-products. We proposed a way to adapt this method to speech signals, e.g. by calculating several such Gram-Matrices, one per frequency range. We further demonstrated that transforming them to correlation matrices increases the speaker-discriminability of the extracted representations. Our experiments on VoxCeleb indicate that the proposed method yields notable improvements over the standard statistics pooling.

There are several ways by which our approach can be extended. For estimating the correlations between channels, the dot-product can be replaced by kernel functions, some of which have been examined in image style-transfer [9]. The property of Gram-Matrix being a symmetric positive semi-definite matrix can be explored in several ways, using e.g. statistical divergences (e.g. Kullback-Leibler, Jensen-Shannon, a.o.). Alternative training/finetuning schemes (e.g. prototypical, triplet, contrastive and self-supervised losses [25, 27, 28, 29]) that do not require a classification head are probably better suited for such kind of experimentation. Finally, the whole chain of transforms we apply on the tensor \mathbf{Y} is not a result of exhaustive experimentation, leaving room for further improvements.

6. Acknowledgements

The work was supported by Czech National Science Foundation (GACR) project ‘‘NEUREM’’ No. 19-26934X, and European Union’s Horizon 2020 project No. 833635 ROXANNE. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 101007666 / ESPERANTO / H2020-MSCA-RISE-2020. The authors do not see any significant ethical or privacy concerns that would prevent the processing of the data used in the study. The datasets do contain personal data, and these are processed in compliance with the GDPR and national law.

7. References

- [1] P. Kenny, G. Boulianne, P. Ouellet, and P. Dumouchel, "Joint factor analysis versus eigenchannels in speaker recognition," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 4, pp. 1435–1447, 2007.
- [2] 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.
- [3] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018.
- [4] T. Zhou, Y. Zhao, and J. Wu, "Resnext and res2net structure for speaker verification," *arXiv preprint arXiv:2007.02480*, 2020.
- [5] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive statistics pooling for deep speaker embedding," *Proc. Interspeech 2018*, pp. 2252–2256, 2018.
- [6] M. À. India Massana, P. Safari, and F. J. Hernando Pericás, "Self multi-head attention for speaker recognition," in *Interspeech 2019: the 20th Annual Conference of the International Speech Communication Association: 15-19 September 2019: Graz, Austria*. International Speech Communication Association (ISCA), 2019, pp. 4305–4309.
- [7] L. You, W. Guo, L.-R. Dai, and J. Du, "Multi-task learning with high-order statistics for x-vector based text-independent speaker verification," in *Proc. Interspeech 2019*, 2019, pp. 1158–1162.
- [8] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2414–2423.
- [9] Y. Jing, Y. Yang, Z. Feng, J. Ye, Y. Yu, and M. Song, "Neural style transfer: A review," *IEEE transactions on visualization and computer graphics*, vol. 26, no. 11, pp. 3365–3385, 2019.
- [10] Y. Li, N. Wang, J. Liu, and X. Hou, "Demystifying neural style transfer," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 2017, pp. 2230–2236.
- [11] J. Chorowski, R. J. Weiss, R. A. Saurous, and S. Bengio, "On using backpropagation for speech texture generation and voice conversion," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 2256–2260.
- [12] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for text-independent speaker verification," in *INTER_SPEECH*, 2017, pp. 999–1003.
- [13] P. Safari, M. India, and J. Hernando, "Self-attention encoding and pooling for speaker recognition," *arXiv preprint arXiv:2008.01077*, 2020.
- [14] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, "Voxceleb: Large-scale speaker verification in the wild," *Computer Speech & Language*, vol. 60, p. 101027, 2020.
- [15] Y. Wu, C. Guo, H. Gao, X. Hou, and J. Xu, "Vector-based attentive pooling for text-independent speaker verification," *Proc. Interspeech 2020*, pp. 936–940, 2020.
- [16] T. Zhou, Y. Zhao, J. Li, Y. Gong, and J. Wu, "CNN with phonetic attention to text-independent speaker verification," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 718–725.
- [17] S. M. Kye, Y. Kwon, and J. S. Chung, "Cross attentive pooling for speaker verification," *arXiv preprint arXiv:2008.05983*, 2020.
- [18] S. Wang, Y. Yang, Y. Qian, and K. Yu, "Revisiting the statistics pooling layer in deep speaker embedding learning," in *2021 12th International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2021, pp. 1–5.
- [19] Y. Gal and Z. Ghahramani, "A theoretically grounded application of dropout in recurrent neural networks," *Advances in Neural Information Processing Systems*, pp. 1027–1035, 2016.
- [20] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: a large-scale speaker identification dataset," in *INTER_SPEECH*, 2017, pp. 2616–2620.
- [21] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldi speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. EPFL-CONF-192584. IEEE Signal Processing Society, 2011.
- [22] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [23] B. Desplanques, J. Thienpondt, and K. Demuyck, "ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification," *Proc. Interspeech 2020*, pp. 3830–3834, 2020.
- [24] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4690–4699.
- [25] J. S. Chung, J. Huh, S. Mun, M. Lee, H. S. Heo, S. Choe, C. Ham, S. Jung, B.-J. Lee, and I. Han, "In defence of metric learning for speaker recognition," *arXiv preprint arXiv:2003.11982*, 2020.
- [26] H. Zeinali, L. Burget, J. Rohdin, T. Stafylakis, and J. Cernocky, "How to improve your speaker embeddings extractor in generic toolkits," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6141–6145.
- [27] C. Zhang and K. Koishida, "End-to-end text-independent speaker verification with triplet loss on short utterances," in *Proc. Interspeech 2017*, 2017, pp. 1487–1491.
- [28] M. Ravanelli and Y. Bengio, "Learning speaker representations with mutual information," in *Proc. Interspeech 2019*, 2019, pp. 1153–1157.
- [29] T. Stafylakis, J. Rohdin, O. Pichot, P. Mizera, and L. Burget, "Self-supervised speaker embeddings," in *Proc. Interspeech 2019*, 2019, pp. 2863–2867.