Integrating Frequency Translational Invariance in TDNNs and Frequency Positional Information in 2D ResNets to Enhance Speaker Verification

Jenthe Thienpondt, Brecht Desplanques, Kris Demuyne

IDLab, Department of Electronics and Information Systems, Ghent University - imec, Belgium

jenthe.thienpondt@ugent.be, brecht.desplanques@ugent.be

Abstract

This paper describes the IDLab submission for the text-independent task of the Short-duration Speaker Verification Challenge 2021 (SdSVC-21). This speaker verification competition focuses on short duration test recordings and cross-lingual trials, along with the constraint of limited availability of in-domain DeepMine Farsi training data. Currently, both Time Delay Neural Networks (TDNNs) and ResNets achieve state-of-the-art results in speaker verification. These architectures are structurally very different and the construction of hybrid networks looks a promising way forward. We introduce a 2D convolutional stem in a strong ECAPA-TDNN baseline to transfer some of the strong characteristics of a ResNet based model to this hybrid CNN-TDNN architecture. Similarly, we incorporate absolute frequency positional encodings in an SE-ResNet34 architecture. These learnable feature map biases along the frequency axis offer this architecture a straightforward way to exploit frequency positional information. We also propose a frequency-wise variant of Squeeze-Excitation (SE) which better preserves frequency-specific information when rescaling the feature maps. Both modified architectures significantly outperform their corresponding baseline on the SdSVC-21 evaluation data and the original VoxCeleb1 test set. A four system fusion containing the two improved architectures achieved a third place in the final SdSVC-21 Task 2 ranking.

Index Terms: speaker verification, Time Delay Neural Network (TDNN), ResNet, frequency positional encodings

1. Introduction

Speaker verification determines whether the speaker in a test utterance matches with a claimed enrolled speaker. Current state-of-the-art speaker verification systems use a deep neural network trained on utterance classification according to speaker identity. The current most popular neural network topologies are based on Time Delay Neural Network (TDNN) [1, 2], also referred to as x-vectors, and ResNet [3] architectures. The converged network is subsequently used to extract low-dimensional speaker embeddings from a bottleneck layer in the final part of the network. The speaker similarity is expressed by the cosine similarity between the speaker embeddings.

In this paper we propose architectural extensions to the current state-of-the-art ECAPA-TDNN and ResNet based speaker embedding extractors. We enhance the ECAPA-TDNN [2] architecture with a network stem based on 2D convolutions. This allows the network to initially construct local, frequency-invariant features before applying 1D convolutions which explicitly integrate the frequency positional information of the features. We also propose the addition of absolute positional frequency encodings in the ResNet based architectures. These learnable feature map offsets enable the ResNet architecture to exploit frequency positional information. Additionally, we in-

troduce a frequency-wise Squeeze-Excitation (SE) block. This module improves upon the standard SE-block [4] by enforcing the global utterance information that is injected in the local convolutions to be estimated per frequency component.

The proposed modifications are evaluated in the context of the text-independent speaker verification Task 2 of the Short-duration Speaker Verification Challenge 2021 (SdSVC-21) [5]. Building upon our previously introduced large margin fine-tuning procedure [6], we evaluate a similar fine-tuning strategy to adapt the models to the in-domain Farsi speech of this challenge. Finally, we use a quality-aware score calibration stage [6] that uses duration- and language-based quality measurements to generate speaker similarity scores that are more consistent across the varying trial conditions.

2. Baseline system architectures

2.1. ECAPA-TDNN

The ECAPA-TDNN [2] architecture is an enhanced version of the popular x-vector topology [1, 7]. The use of hierarchical grouped convolutions [8] reduces the model parameter count. It also introduces 1-dimensional TDNN-specific SE-blocks [4] which rescale the intermediate time context bound frame-level features per channel according to global utterance properties. The pooling layer uses a channel- and context-dependent self-attention mechanism to attend to different speaker characterizing properties at different time steps for each feature map. Finally, Multi-layer Feature Aggregation (MFA) [9] provides additional complementary information for the statistics pooling by concatenating the final frame-level features with the intermediate features of previous layers. The architecture is trained with the AAM-softmax [10] loss. More details about the architecture can be found in [2]. Similar to [6], we increase the intermediate channel dimension to 2048 and add a fourth SE-Res2Block to the network.

2.2. SE-ResNet34

Our second baseline system is based on the ResNet [11] architecture. We use the same network topology as defined in [3], with the addition of an SE-block at the end of each residual block. We also incorporate the same attentive statistics pooling layer as the ECAPA-TDNN baseline. A final modification is the incorporation of sub-center AAM [12]. This is an extension of the AAM-softmax loss function and it defines multiple prototype embeddings per training speaker. This should make the model less susceptible to potential label noise or severely corrupted data in the training set. In this paper the number of sub-centers per speaker is set to two. This system was the best scoring single system on the VoxSRC-20 validation set in our VoxSRC submission [13].

2.3. Standard speaker embedding training procedure

All baseline neural networks are trained on all of the allowed training corpora in the SdSVC-21 setting. The input features are 80-dimensional log Mel-filterbank energies extracted with a window length of 25 ms and a frame-shift of 10 ms. We apply an online augmentation strategy on the 2 s random training crops using the MUSAN [14] corpus (additive music, noise & babble) and the RIR [15] corpus (reverb). SpecAugment [16] is applied to the input features with a frequency and temporal masking dimension of 10 and 5, respectively. Finally, the input features are mean normalized across time per utterance.

The model parameter updates are determined by the Adam optimizer [17] with a cyclical learning rate schedule. We use the *triangular2* policy described in [18]. The cycle length is set to 130k iterations with a minimum and maximum learning rate of $1e-8$ and $1e-3$, respectively. The systems are trained for three full cycles. To prevent overfitting, a weight decay of $2e-4$ is applied on the weights of the AAM-softmax layer, while a value of $2e-5$ is used on all other layers of the model.

3. Proposed system architectures

TDNN and ResNet based speaker verification systems achieve similar performance [13] while being structurally very different. TDNNs depend on 1D convolutions of which the kernels cover the complete frequency range of the input features. As a consequence, the absolute frequency position of each input feature is hard-coded through the order of the filter coefficients. This makes sense in the context of speaker verification as features based on absolute frequency information such as the speaker's pitch provide crucial information. The main drawback is that many filters are needed to model the fine details of complex patterns that can occur at any frequency region.

ResNets use 2D convolutions with a small receptive field, capturing local frequency and temporal information. By exploiting local speaker-specific frequency patterns that repeat but for small frequency shifts, this model requires fewer feature channels to model fine frequency details. However, accurate absolute frequency positions of features are not explicitly encoded [19]. This can be sub-optimal as we expect significantly different patterns in the low frequency regions compared to the high frequency regions. In the next sections we enhance both architectures to incorporate the beneficial characteristics of the other model.

3.1. ECAPA CNN-TDNN with 2D convolutional stem

We want the ECAPA-TDNN to be invariant to small and reasonable shifts in the frequency domain, compensating for realistic intra-speaker frequency variability. To accomplish this, we base the initial network layers on 2D convolutions. These layers will also require less filters to model high resolution input detail. A similar modification of the standard x-vector network for text-dependent speaker verification can be found in [20].

Inspired by the powerful 2D convolutions in the ResNet architectures, we adopt a similar structure for our ECAPA-TDNN stem. The proposed network configuration is shown in Figure 1. The number of channels C in the stem is set to 128 for all experiments. To make the stem more computationally efficient we use a stride of 2 in the frequency dimension of the first and final 2D convolutional layer. The output feature map of the new stem is subsequently flattened in the channel and frequency dimensions and used as input for the regular ECAPA-TDNN network. We will refer to this network as the ECAPA CNN-TDNN.

Figure 1: The 2D convolutional stem of the ECAPA CNN-TDNN architecture. C and T indicate the channel and temporal dimension, respectively. The kernel size is denoted by k . S refers to the number of training speakers.

3.2. SE-ResNet34 with frequency position information

3.2.1. Learnable frequency positional encodings

While a certain robustness against frequency translation is beneficial, the spatial invariance of 2D convolutions can limit the ability of the network to fully exploit frequency-specific information. We argue that encoding positional frequency information in the intermediate feature representations can make the network incorporate and utilize knowledge of the feature frequency positions.

Positional encodings have been popularized with the rise of Transformer [21] models. By design, these Transformers are invariant to the re-ordering of the input sequence and positional encodings can alleviate this issue when required. These encodings can either be learnable or use a fixed representation based on the sine and cosine function [22, 21]. In this paper we focus on learnable encodings as the computational overhead is negligible.

Consider the input of a residual module in our baseline ResNet architecture as $\mathbf{X} \in \mathbb{R}^{C \times F \times T}$, with C , F and T indicating the channel, frequency and time dimension, respectively. We define the positional encoding vector $\mathbf{p} \in \mathbb{R}^F$ as a trainable vector. The elements in this vector are broadcasted to match the dimensions of the targeted input feature map. The input of the residual module is now defined as $\mathbf{X} + \mathbf{p}$. We add a unique positional encoding to the input of each residual block after branching the skip connection.

3.2.2. Frequency-wise Squeeze-Excitation (fwSE)

Squeeze-Excitation (SE) [4] blocks have been successfully applied in both TDNN and ResNet based speaker verification architectures. The first stage, the squeeze operation, calculates a vector $\mathbf{z} \in \mathbb{R}^C$ containing the mean descriptor for each feature map. It is followed by the excitation operation, which calculates a scalar for each feature map given the information in $\mathbf{z}$. Subsequently, each feature map is rescaled with its corresponding scalar.

We argue that rescaling per feature map is not tailored for the processing of speech in ResNets. Instead, we propose a frequency-wise Squeeze-Excitation (fwSE) block, which injects global frequency information across all feature maps. We calculate a vector $\mathbf{z} \in \mathbb{R}^F$ containing the mean descriptor for each frequency-channel of the intermediate feature maps in the following manner:

$$\mathbf{z}^f = \frac{1}{C \times T} \sum_{i=1}^C \sum_{j=1}^T x_{i,j}^f \quad (1)$$

with x^f the elements of $\mathbf{X}^f \in \mathbb{R}^{C \times F \times T}$, the component of input feature map $\mathbf{X} \in \mathbb{R}^{C \times F \times T}$ corresponding with frequency position f . From the mean descriptors in $\mathbf{z}$ we calculate a vector $\mathbf{s} \in \mathbb{R}^F$ containing the scaling scalars for each frequency-channel in the second stage, the excitation operation:

$$\mathbf{s} = \sigma(\mathbf{W}_2 f(\mathbf{W}_1 \mathbf{z} + \mathbf{b}_1) + \mathbf{b}_2) \quad (2)$$

with $\mathbf{W}$ and $\mathbf{b}$ indicating the weights and bias of a linear layer, $f(\cdot)$ the ReLU activation function and σ the sigmoid function. Finally, $\mathbf{X}^f$ is scaled with the corresponding scalar in $\mathbf{s}$. The proposed frequency-wise SE-blocks are inserted at the end of each residual module before the additive skip connection.

4. SdSVC-21 specific modifications

In this section we describe the modifications to our speaker verification pipeline for the Short-duration Speaker Verification Challenge 2021 (SdSVC-21). The challenge creates a difficult speaker verification set by incorporating short duration cross-lingual trial pairs. All speakers in the set are native Farsi speakers with the trial test utterances containing either Farsi or English speech. The allowed training data consists of VoxCeleb1 & VoxCeleb2 [23, 24], LibriSpeech, from which we only use the *train-other-500* component [25], the Farsi subset of the Mozilla Common Voice corpus [26] and a part of the DeepMine dataset [27] with utterances from 588 in-domain speakers. As opposed to last year’s challenge, the in-domain training data is extended with English spoken utterances. More information about the challenge can be found in the SdSVC-21 evaluation plan [5].

4.1. Domain-balanced large margin fine-tuning

We fine-tune all systems with an adapted version of the Large Margin Fine-Tuning (LM-FT) strategy presented in [6]. The performance gains of regular LM-FT were much smaller on the in-domain SdSVC-21 validation set compared to the other domains in the training set. To mitigate this issue, we propose a domain-balanced and SdSVC-21 specific version of the LM-FT strategy.

To preserve the verification performance on the shortest SdSVC-21 test utterances, we increase the duration of the temporal crop to only 3 s during LM-FT. To compensate, we increase the AAM-softmax margin to a less aggressive value of 0.3. In comparison, we used an increase to 6 s and a margin of 0.5 on VoxCeleb data previously [6]. We also equalize the sampling probability of each speech domain present in our training dataset. We sample a random utterance for each of the 588 in-domain DeepMine speakers and for each speaker in a random subset of 588 speakers from each of the other domains. Once the pool of selected utterances is exhausted during batch creation, the random selection of training samples is repeated. Due to time constraints, we disabled the hard prototype mining based utterance selection criterion [28, 6].

4.2. Score estimation and normalization

The speaker enrollment models are constructed by averaging the corresponding L_2 -normalized enrollment embeddings. The verification trials are scored by calculating the cosine similarity between the enrollment model and the test utterance embedding. Scores are normalized using top-2000 adaptive s-normalization [29, 30]. The imposter cohort consists of top imposters selected from the pooled SdSVC-21 Farsi training data and the Farsi component of Common Voice. Each imposter speaker is represented by the average of their length-normalized training embeddings.

4.3. Score fusion and quality-aware score calibration

Score calibration maps the trial scores to log-likelihood-ratios that can be converted to interpretable probabilities [31]. It also allows for easy score fusion by producing a weighted average across system scores [31]. In addition, the score calibration stage can be used to compensate for variability in trial conditions by including Quality Measurements (QMs) [31, 6] such as duration metrics. This is very effective for evaluation sets with varying duration conditions, since most of the embedding extractors are trained with fixed-length audio crops.

The calibration is based on logistic regression with learnable weights and bias to scale and shift the original output score together with the QMs to obtain a calibrated log-likelihood-ratio [6]. This procedure makes the evaluation decision thresholds implicitly condition dependent. To train the parameters we create our own speaker verification trials from the in-domain training set. We generate 100 within-gender verification trials per training speaker. We select between 1 and 10 Farsi enrollment utterances per trial. The number of target and non-target trials as well as the number of Farsi-only and cross-lingual trials is balanced.

For our fusion submission, we first take a weighted score average according to the performance of the systems on the SdSVC-21 validation set. Subsequently, we apply the quality-aware calibration on these averaged scores. The logistic regression uses the quality metrics described below.

4.3.1. Duration-based quality measures

The SdSVC-21 trials are asymmetric in the sense that there is significantly more enrollment speech than test speech. We introduce $\log(d_t - d_{min})$ as a QM with d_t the duration of the test utterance and d_{min} the minimum expected test utterance duration in the DeepMine corpus (i.e. 1 s).

It is harder to anticipate the impact of the multiple enrollment durations on the embedding quality. We introduce $\log(n_e)$ as a basic second QM, where n_e is the number of enrollment utterances associated with the trial. We cap n_e to a maximum of three, as we expect the quality improvement on the average enrollment embedding to saturate.

4.3.2. Cross-lingual likelihood-based quality measures

Due to the availability of in-domain English training speech in SdSVC-21, the performance impact of cross-lingual trials seems rather limited. We notice small gains by introducing the language log-likelihood-ratio QM on the test utterance produced by a Farsi vs. English language classifier. The language identification system is a Gaussian Backend (GB) trained on the output of the neural network pooling layer on the SdSVC-21 training set. The GB achieves 99.8% language classification accuracy on the SdSVC-21 validation test utterances.

Table 1: EER and MinDCF of the baseline and proposed architectures on VoxCeleb1 and SdSVC-21 validation and test sets.

Architecture	Variant	VoxCeleb1-O			SdSVC-21 Val		SdSVC-21 Test	
		EER(%)	MinDCF	MinDCF2	EER(%)	MinDCF	EER(%)	MinDCF
ECAPA-TDNN	Standard	0.85	0.0498	0.1227	1.86	0.0783	1.35	0.0632
	LM-FT	0.76	0.0481	0.1010	1.64	0.0631	1.20	0.0548
	LM-FT + QMs	0.73	0.0466	0.0961	1.50	0.0625	1.14	0.0520
ECAPA CNN-TDNN	Standard	0.72	0.0477	0.1193	1.43	0.0683	1.21	0.0568
	LM-FT	0.66	0.0391	0.0820	1.43	0.0664	1.08	0.0502
	LM-FT + QMs	0.61	0.0372	0.0903	1.45	0.0584	1.01	0.0471
SE-ResNet34	Standard	0.78	0.0432	0.1008	1.70	0.0872	1.39	0.0601
	LM-FT	0.70	0.0428	0.0884	1.57	0.0720	1.22	0.0534
	LM-FT + QMs	0.65	0.0389	0.0877	1.43	0.0653	1.17	0.0510
fwSE-ResNet34 w/ positional encodings	Standard	0.70	0.0410	0.0856	1.59	0.0776	1.23	0.0526
	LM-FT	0.63	0.0354	0.0776	1.36	0.0652	1.13	0.0483
	LM-FT + QMs	0.61	0.0335	0.0812	1.43	0.0632	1.08	0.0463

5. Results

We evaluate the systems on the original VoxCeleb1 test set and the SdSVC-21 validation and test set. We report the EER and MinDCF metric using a P_{target} value of 10^{-2} with $C_{Miss} = 10$ and $C_{FA} = 1$, as this is the main metric for SdSVC-21. For VoxCeleb1-O, we also report the MinDCF2 value with $C_{Miss} = 1$. All reported scores are s-normalized according to Section 4.2 for the SdSVC-21 results. For scores on the VoxCeleb1-O trials, we use the training part of the VoxCeleb1 dataset to calculate the imposters in the s-norm cohort with a size of 120. For system discussions we only focus on the large SdSVC-21 test set containing 4.6M trials due to its reliability.

Table 1 compares the performance of the baseline models with the proposed architectures. It also shows the impact of our proposed domain-balanced LM-FT strategy and of the calibration with the proposed QMs. We use a margin value of 0.3 and a crop size of 3 s for all systems during the LM-FT stage as this provides the best performance on the in-domain SdSVC-21 sets as shown in Table 2. This table also includes a fine-tuning strategy that only modifies the sampling to be domain-balanced. For VoxCeleb1-O trials we only include the duration-based QMs as defined in [32].

Table 2: Evaluation of the domain-balanced large margin fine-tuning strategy on the ECAPA CNN-TDNN architecture.

Method	Vox1-O	SdSVC-21 Test	
	EER(%)	EER(%)	MinDCF
Baseline	0.72	1.21	0.0568
Domain-balanced FT	0.74	1.14	0.0523
LM-FT (m=0.3, s=3)	0.66	1.08	0.0502
LM-FT (m=0.5, s=5)	0.60	1.25	0.0573

The proposed ECAPA CNN-TDNN architecture outperforms the baseline ECAPA-TDNN model significantly. After fine-tuning and applying quality metrics, the EER and MinDCF on the SdSVC-21 test set show a relative improvement of 11.4% and 9.4%, respectively. Similar gains are observed on SdSVC-21 Test for the proposed fwSE-ResNet34 with relative improvements of 7.7% in EER and 9.2% in MinDCF.

The proposed LM-FT strategy for the SdSVC-21 target domain is beneficial and produces an average relative improvement of EER and MinDCF of 10.6% and 11.1% respectively on the SdSVC-21 test set across the different systems. The SdSVC-21 specific quality-aware calibration improves results on the SdSVC-21 test set further with an average relative improvement of 5% in both the EER and MinDCF value.

Table 3: Fusion system performance on SdSVC-21 Test.

Systems	EER(%)	MinDCF
Fusion	1.06	0.0473
Fusion + LM-FT	0.92	0.0416
Fusion + LM-FT + QMs	0.86	0.0386

The results of our final fusion submission for SdSVC-21 are shown in Table 3. We also provide the impact of LM-FT and quality-aware calibration. The four system ensemble consists of the baseline ECAPA-TDNN, the proposed ECAPA CNN-TDNN along with a second variant and the proposed fwSE-ResNet34 with positional encodings. See our SdSVC-21 technical report for more details. The weighted score fusion relatively improves upon the strong ECAPA CNN-TDNN single system on the SdSVC-21 test set by 14.9% and 18.0% in EER and MinDCF, respectively.

6. Conclusions

In this paper we proposed a 2D convolutional stem for the ECAPA-TDNN speaker verification model, incorporating frequency translational invariance in the initial layers of the network. We also introduced frequency-wise Squeeze-Excitation blocks and positional encodings for ResNet architectures, allowing the model to fully exploit positional information on the frequency axis. The proposed modifications significantly improve upon the strong baseline systems. Our final SdSVC-21 fusion submission containing the proposed architectures, tweaked by a large margin fine-tuning strategy and a quality-aware calibration step, produces a MinDCF score of 0.0386 and EER of 0.86% on the SdSVC-21 test set. These promising results encourage us to further explore hybrid architectures in the future.

7. References

- [1] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. ICASSP*, 2018, pp. 5329–5333.
- [2] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Proc. Interspeech*, 2020, pp. 3830–3834.
- [3] D. Garcia-Romero, G. Sell, and A. McCree, "Magnet: X-vector magnitude estimation network plus offset for improved speaker recognition," in *Proc. Odyssey 2020*, 2020, pp. 1–8.
- [4] J. Hu, L. Shen, and G. Sun, "Squeeze-and-Excitation networks," in *Proc. IEEE/CVF CVPR*, 2018, pp. 7132–7141.
- [5] H. Zeinali, K. A. Lee, J. Alam, and L. Burget, "Short-duration speaker verification (sds) challenge 2020: the challenge evaluation plan," *arXiv preprint arXiv:1912.06311*, 2019.
- [6] J. Thienpondt, B. Desplanques, and K. Demuynck, "The IDLab VoxSRC-20 submission: Large margin fine-tuning and quality-aware score calibration in DNN based speaker verification," in *Proc. ICASSP*, 2021.
- [7] D. Snyder, D. Garcia-Romero, G. Sell, A. McCree, D. Povey, and S. Khudanpur, "Speaker recognition for multi-speaker conversations using x-vectors," in *Proc. ICASSP*, 2019, pp. 5796–5800.
- [8] S. Gao, M.-M. Cheng, K. Zhao, X. Zhang, M.-H. Yang, and P. H. S. Torr, "Res2Net: A new multi-scale backbone architecture," *IEEE TPAMI*, 2019.
- [9] Z. Gao, Y. Song, I. McLoughlin, P. Li, Y. Jiang, and L.-R. Dai, "Improving Aggregation and Loss Function for Better Embedding Learning in End-to-End Speaker Verification System," in *Proc. Interspeech*, 2019, pp. 361–365.
- [10] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *2019 IEEE/CVF CVPR*, 2019, pp. 4685–4694.
- [11] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE/CVF CVPR*, 2016, pp. 770–778.
- [12] J. Deng, J. Guo, T. Liu, M. Gong, and S. Zafeiriou, "Sub-center ArcFace: Boosting face recognition by large-scale noisy web faces," in *Proc. ECCV*, 2020, pp. 741–757.
- [13] J. Thienpondt, B. Desplanques, and K. Demuynck, "The IDLAB voxceleb speaker recognition challenge 2020 system description," *arXiv preprint arXiv:2010.12468*, 2020.
- [14] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [15] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *Proc. ICASSP*, 2017, pp. 5220–5224.
- [16] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Proc. Interspeech*, 2019, pp. 2613–2617.
- [17] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *Proc. ICLR*, 2015.
- [18] L. N. Smith, "Cyclical learning rates for training neural networks," in *IEEE WACV*, 2017, pp. 464–472.
- [19] M. A. Islam, M. Kowal, S. Jia, K. G. Derpanis, and N. D. Bruce, "Position, padding and predictions: A deeper look at position information in cnns," *arXiv preprint arXiv:2101.12322*, 2021.
- [20] Z. Chen, D. Chen, H. Ding, and Y. Lin, "The NetEase Games system description for text-dependent sub-challenge of SDSVC 2020," <https://sdsvc.github.io/2020/descriptions/Team14.Task1.pdf>, accessed: 2020-03-30.
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NeurIPS*, 2017, p. 6000–6010.
- [22] J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin, "Convolutional sequence to sequence learning," in *Proc. ICML*, 2017, pp. 1243–1252.
- [23] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Proc. Interspeech*, 2017, pp. 2616–2620.
- [24] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. Interspeech*, 2018, pp. 1086–1090.
- [25] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *Proc. ICASSP*, 2015, pp. 5206–5210.
- [26] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *Proc. LREC*, 2020, pp. 4218–4222.
- [27] H. Zeinali, J. Černocký, and L. Burget, "A multi purpose and large scale speech corpus in Persian and English for speaker and speech recognition: the DeepMine database," in *Proc. ASRU*, 2019, pp. 397–402.
- [28] J. Thienpondt, B. Desplanques, and K. Demuynck, "Cross-lingual speaker verification with domain-balanced hard prototype mining and language-dependent score normalization," in *Proc. Interspeech*, 2020, pp. 756–760.
- [29] Z. N. Karam, W. M. Campbell, and N. Dehak, "Towards reduced false-alarms using cohorts," in *Proc. ICASSP*, 2011, pp. 4512–4515.
- [30] S. Cumani, P. Batzu, D. Colibro, C. Vair, P. Laface, and V. Vasilakakis, "Comparison of speaker recognition approaches for real applications," in *Proc. Interspeech*, 2011, pp. 2365–2368.
- [31] N. Brümmer and E. De Villiers, "The bosaris toolkit: Theory, algorithms and code for surviving the new dcf," *arXiv preprint arXiv:1304.2865*, 2013.
- [32] J. S. Chung, A. Nagrani, E. Coto, W. Xie, M. McLaren, D. A. Reynolds, and A. Zisserman, "Voxsrc 2019: The first voxceleb speaker recognition challenge," *arXiv preprint arXiv:1912.02522*, 2019.