



# LASPA: Language Agnostic Speaker Disentanglement with Prefix-Tuned Cross-Attention

Aditya Srinivas Menon\*, Raj Prakash Gohil\*, Kumud Tripathi, Pankaj Wasnik

Media Analysis Group, Sony Research India

{aditya.menon, raj.gohil, kumud.tripathi, pankaj.wasnik}@sony.com

## Abstract

Speaker recognition models face challenges in multi-lingual settings due to the entanglement of linguistic information within speaker embeddings. The overlap between vocal traits such as accent, vocal anatomy, and a language’s phonetic structure complicates separating linguistic and speaker information. Disentangling these components can significantly improve speaker recognition accuracy. To this end, we propose a novel disentanglement learning strategy that integrates joint learning through prefix-tuned cross-attention. This approach is particularly effective when speakers switch between languages. Experimental results show the model generalizes across monolingual and multi-lingual settings, including unseen languages. Notably, the proposed model improves the equal error rate across multiple datasets, highlighting its ability to separate language information from speaker embeddings and enhance recognition in diverse linguistic conditions.

**Index Terms:** Language agnostic, Speaker recognition, Multi-lingual speaker diarization, speaker representation, prefix-tuners

## 1. Introduction

Speaker recognition in multi-lingual contexts presents significant challenges due to the complex interaction between language-specific characteristics, speaker identity, and acoustic variability. When speakers use different languages, their phonetic and prosodic patterns are influenced by the linguistic structure of each language and can vary considerably leading to these challenges. As a result, traditional speaker recognition systems often misinterpret these variations as different speakers, failing to recognize the same individual across languages. This issue arises due to the entanglement of language-dependent and speaker-dependent acoustic features within the speech signal [1, 2]. Linguistically, each language imposes unique articulatory and prosodic requirements, which cause systematic variation in a speaker’s voice. Features such as phoneme structure, intonation, and co-articulation strategies differ significantly across languages, leading to distinct acoustic behavior. These differences, while linguistically driven, are not indicative of a change in speaker identity. However, speaker recognition models often fail to distinguish between language-induced variability and speaker-specific traits, reducing accuracy when identifying speakers across multiple languages.

Disentangled representation learning has been proposed as an effective approach to address this issue. The goal is to separate speaker identity-related features from language-dependent acoustic characteristics, which should remain invariant across

languages. This method improves speaker recognition in multi-lingual settings by learning representations robust to linguistic variation, thereby mitigating the impact of language on recognition performance [1]. Specifically, disentangled representation learning allows the model to isolate speaker-specific traits—such as vocal-tract characteristics, prosody, and voice timbre while suppressing language-dependent features like phonetic and prosodic variation. This separation helps the system focus on what remains consistent across languages, enabling more accurate cross-lingual speaker recognition.

Traditionally, adversarial learning strategies based on the Gradient Reversal Layer (*GRL*) have been proposed to disentangle speaker-specific information from language-specific acoustic features [3, 4, 5, 6, 7]. However, *GRL*-based approaches are prone to hyperparameters and can frequently lead to unstable training [1]. To address these challenges, previous work has proposed *GRL*-independent approaches for effective disentangled representation learning [8, 9, 10]. Motivated by these approaches, we introduce LASPA, a multi-task learning strategy utilizing prefix-tuned cross attention, designed to achieve robust separation of speaker and language features while ensuring stable and efficient training. Our key contributions are summarized as follows:

1. The paper proposes a joint learning-based approach to derive language-independent speaker embeddings for multi-lingual speaker recognition and speaker diarization.
2. Speaker and language encoders, along with prefix-tuners are employed to fuse embeddings, ensuring accurate signal reconstruction by the decoder.
3. To enhance the learning of speaker embedding, this work uses multiple loss functions like Mean Absolute Pearson’s Correlation (*MAPC*) [9], Additive Angular Margin Softmax (*AAM Softmax*) [11], Mean Squared Error (*MSE*), and Negative Log Likelihood (*NLL*).

## 2. Previous Work

Previous studies have shown that speaker representations are often entangled with factors such as emotions [12, 13], accent [14], age [15], and environmental conditions like noise and reverberation [16]. More recently, research has validated that linguistic features embedded in speaker representations make them susceptible to language variability, leading to inaccurate recognition of the same speaker across different languages [1, 2]. While state-of-the-art (SOTA) models such as ResNet [17, 18] and ECAPA (Emphasized Channel Attention, Propagation, and Aggregation) [19] perform well in general speaker recognition tasks, they struggle in multi-lingual settings where a single speaker switches between languages, primarily due to the presence of linguistic information in the speaker rep-

\*These authors contributed equally.

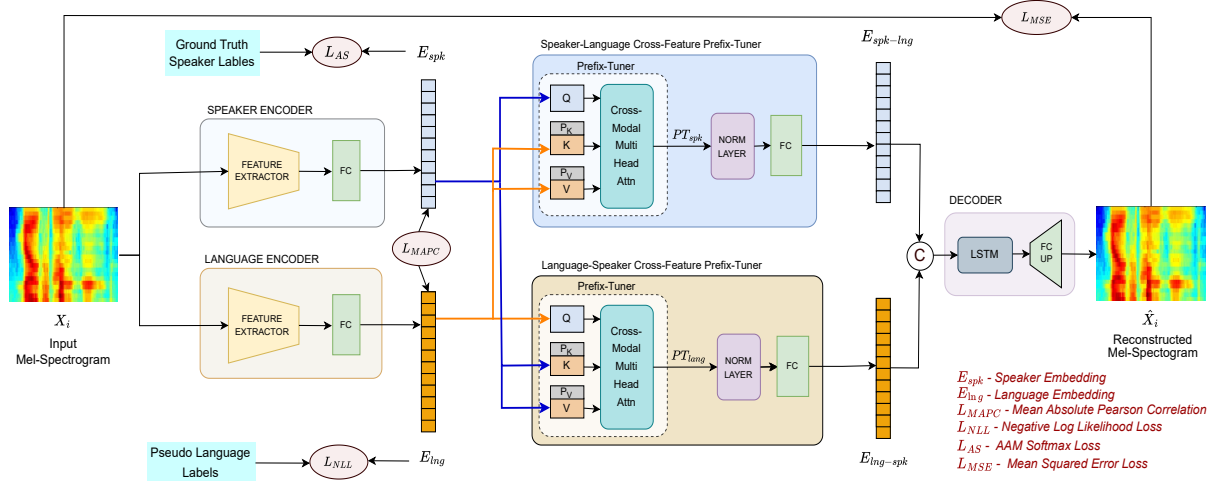


Figure 1: Block diagram of the proposed architecture. Speaker and Language embeddings are extracted using the speaker and language encoder respectively and then after it is fused through prefix-tuners, and then passed to the decoder for mel-spectrogram reconstruction.

resentation.

Recently, there has been growing interest in speaker recognition tasks in multi-lingual scenarios. In this direction, [1] addressed the challenge of speaker recognition in bilingual scenarios by introducing VoxCeleb1-B, a large-scale evaluation set derived from VoxCeleb1, which is designed to test models in bilingual settings. They also propose a novel disentangled representation learning strategy that combines *GRL* with *MAPC* minimization. This approach shows significant improvements in performance for bilingual speaker recognition. Alongside the *GRL*, the latest SOTA ReDimNet [20] uses the advantages of mapping 2D-1D convolution blocks to get more robust embeddings across different scenarios.

With the growing need for more efficient models, [21] introduced prefix-tuning as a lightweight alternative to traditional fine-tuning methods in large pre-trained models like GPT and BERT. Instead of updating all model parameters, prefix-tuning adds task-specific "prefix vectors" to the input sequence, keeping the original model weights frozen. Following its success in NLP, prefix-tuning has been extended to cross-modal applications. [22] claims that the prefix-tuned bottleneck attention helps in efficient multi-modal interaction between different modalities. It takes into account the importance of one modality to the other while performing fusion. These advances highlight the potential for prefix-tuning in multi-task learning.

### 3. Model Architecture

The architectural overview of the proposed system is shown in Fig. 1. The architecture comprises a Speaker Encoder, a Language Encoder, two Prefix-Tuners, and a Decoder. The input waveform is initially resampled to a 16 kHz sampling rate. Subsequently, mel-spectrogram is computed using a Hamming window with a 25 ms window size and a 10 ms stride. These mel-spectrograms  $X_i$  serve as the input for both the Speaker Encoder and Language Encoder.

#### 3.1. Speaker Encoder

The proposed Speaker Encoder includes a feature extractor that transforms the mel-spectrogram  $X_i$  into a  $J$  dimensional vec-

tor. The output from the feature extractor is then passed through a feed-forward layer. The final output of this process is considered as the speaker embedding  $E_{spk}$ .

#### 3.2. Language Encoder

Similarly, the proposed Language Encoder includes a feature extractor that transforms the mel-spectrogram  $X_i$  into a  $J$  dimensional vector. The final output of this process is considered as the language embedding  $E_{lng}$ .

#### 3.3. Prefix-Tuner

To facilitate interaction between speaker and language features and extract relevant information, we employ two prefix-tuners [21]. Speaker-Language cross-feature Prefix-Tuner  $PT_{spk}$  and Language-Speaker cross-feature Prefix-Tuner  $PT_{lang}$  for speaker and language information, respectively.

These prefix-tuners blend initialized parameters with learned representations to enhance feature interaction. In the Language-Speaker cross-feature Prefix-Tuner, the query  $Q$  comes from language embedding  $E_{lng}$ , while the keys  $K$  and values  $V$  are derived from speaker embedding  $E_{spk}$ , appended with prefix terms  $P_k$  and  $P_v$ . Conversely, in the Speaker-Language cross-feature Prefix-Tuner, the query comes from speaker features, and the keys and values are derived from language features  $E_{lng}$  with the corresponding prefixes  $P_k$  and  $P_v$ . Both tuners use these prefix-enhanced keys and values in cross-feature multi-head attention [23], enabling focused interaction between speaker and language information. The attention is obtained by comparing the query with the prefix-tuned keys and values. Mathematically, this is expressed in Equation 1:

$$F_{Att} = \text{Softmax} \left( \frac{Q \cdot K_p^T}{\sqrt{J_{K_p}}} \right) \cdot V_p \quad (1)$$

where  $J_{K_p}$  represents the dimensionality of the key vectors  $K_p$ . The resulting attention features  $F_{Att}$  are passed to the prefix-tuners give the output as  $E_{spk-lng}$  and  $E_{lng-spk}$

Table 1: *EER (%) and minDCF on cleaned versions of the VoxCeleb1 test sets, VoxCeleb1-B, VoxSRC 2021, VoxSRC 2020 validation set, and NISP-B. The two version of LASPA with and without prefix-tuner is compared with other state-of-the-art models*

Model	VoxSRC 2021 Val		VoxCeleb1-B		NISP-B	
	EER (↓)	minDCF (↓)	EER (↓)	minDCF (↓)	EER (↓)	minDCF (↓)
ResNet-S(Baseline)	9.22	0.500	9.69	0.618	22.33	0.993
GRL ResNet-S	8.35	0.460	8.25	0.510	–	–
<b>LASPA ResNet-S</b>	<b>6.90</b>	<b>0.415</b>	<b>5.88</b>	<b>0.319</b>	<b>14.33</b>	<b>0.668</b>
ResNet-L (Baseline)	5.16	0.308	5.96	0.397	21.13	0.973
GRL ResNet-L	4.22	0.246	3.69	0.254	–	–
<b>LASPA ResNet-L</b>	<b>3.12</b>	<b>0.212</b>	<b>2.38</b>	<b>0.102</b>	<b>14.75</b>	<b>0.558</b>
ECAPA (Baseline)	4.42	0.292	4.90	0.325	20.60	0.727
<b>LASPA ECAPA</b>	<b>3.01</b>	<b>0.195</b>	<b>2.07</b>	<b>0.086</b>	<b>11.95</b>	<b>0.544</b>
ReDimNet (Baseline)	2.15	0.086	1.66	0.083	11.90	0.607
<b>LASPA ReDimNet</b>	<b>2.10</b>	<b>0.079</b>	<b>1.62</b>	<b>0.053</b>	<b>10.22</b>	<b>0.233</b>

Model	VoxCeleb1 cl.		VoxCeleb1-E cl.		VoxCeleb1-H cl.		VoxSRC 2020 Val	
	EER (↓)	minDCF (↓)	EER (↓)	minDCF (↓)	EER (↓)	minDCF (↓)	EER (↓)	minDCF (↓)
ResNet-S (Baseline)	2.27	0.167	2.43	0.175	4.74	0.299	6.91	0.390
GRL ResNet-S	2.15	0.170	2.42	0.171	4.49	0.284	6.54	0.370
<b>LASPA ResNet-S</b>	<b>1.96</b>	<b>0.155</b>	<b>2.38</b>	<b>0.156</b>	<b>4.22</b>	<b>0.184</b>	<b>6.22</b>	<b>0.256</b>
ResNet-L (Baseline)	1.17	0.083	1.30	0.091	2.58	0.164	4.06	0.231
GRL ResNet-L	0.99	0.079	1.25	0.088	2.42	0.154	3.91	0.220
<b>LASPA ResNet-L</b>	<b>0.95</b>	<b>0.063</b>	<b>1.05</b>	<b>0.086</b>	<b>2.19</b>	<b>0.119</b>	<b>3.42</b>	<b>0.212</b>
ECAPA (Baseline)	1.24	0.089	1.46	0.093	2.66	0.162	4.42	0.233
<b>LASPA ECAPA</b>	<b>0.89</b>	<b>0.045</b>	<b>1.00</b>	<b>0.076</b>	<b>2.01</b>	<b>0.095</b>	<b>3.21</b>	<b>0.192</b>
ReDimNet (Baseline)	0.37	0.030	0.53	0.051	1.00	0.097	1.66	0.128
<b>LASPA ReDimNet</b>	<b>0.31</b>	<b>0.026</b>	<b>0.49</b>	<b>0.047</b>	<b>0.86</b>	<b>0.087</b>	<b>1.42</b>	<b>0.095</b>

### 3.4. Decoder

The prefix-tuned speaker and the language embedding,  $E_{spk-ling}$  and  $E_{ling-spk}$  respectively, are concatenated and passed to the decoder that attempts to reconstruct the mel-spectrogram input  $X_i$ . The Decoder consists LSTM [24] followed by an Multi-layer perceptron(MLP) layer [25] to upscale and reconstruct the mel-spectrogram  $\hat{X}_i$

## 4. Training and Inference

During training, the input waveform is converted into mel-spectrograms, denoted as  $X_i$ . In the Speaker Encoder, the mel-spectrogram passes through a feature extraction module and a fully connected layer. Details of the feature extractor can be found in Section 5. The output of the fully connected layer is the speaker embedding,  $E_{spk}$ . Similarly, in the Language Encoder, the mel-spectrogram passes through a feature extractor and fully connected layer to obtain the language embedding,  $E_{ling}$ .

For optimization, the *AAM Softmax* loss function,  $L_{AS}$ , is applied to the speaker embedding,  $E_{spk}$ , using the ground truth speaker label to compute the speaker classification loss. Simultaneously, the *NLL* loss function,  $L_{NLL}$ , is applied to the language embedding,  $E_{ling}$ , with pseudo-language labels to obtain the language classification loss. Furthermore, a mutual metric-based information loss,  $L_{MAPC}$ , is imposed between  $E_{spk}$  and  $E_{ling}$  to encourage disentanglement. Finally, we concatenate

these embeddings and pass them through a decoder which consists of an LSTM module and a feed-forward network, to reconstruct the mel-spectrogram. A *MSE* loss,  $L_{MSE}$ , is applied between the reconstructed mel-spectrogram,  $\hat{X}_i$ , and the input mel-spectrogram. The proposed model is optimized using the final loss defined as:

$$L = L_{MSE} + L_{AS} + L_{MAPC} + L_{NLL} \quad (2)$$

We empirically found that using equal weights led to optimal performance across evaluation metrics. During inference, we mainly use the Speaker Encoder, excluding the language encoder and cross-feature prefix-tuners. This ensures rapid embedding generation while maintaining high performance and language disentanglement.

This section provides detailed information about the implementation and datasets, followed by the evaluation and results obtained. The best results in the tables are highlighted in bold.

### 4.1. Implementation Details

We opted for two versions of ResNet i.e., ResNet-S with 1.4 million and ResNet-L with 8 million parameters, ECAPA and ReDimNet with 15 million parameters - as our backbone feature extractor in LASPA as it shows significant improvement across various datasets [20]. Our experiments were carried out using 8 NVIDIA A100 GPUs (40 GB each) server. The models were trained using an Adam optimizer with a weight decay of  $2e - 5$ , and the learning rate was set at  $1e - 3$ . Each batch consisted

of processing 400 utterances at a time and each configuration is trained for 1200 epochs.

## 4.2. Datasets

For training, we use the VoxCeleb2 [26] dev set, a benchmark dataset commonly used for speaker recognition tasks consisting of 5,994 speakers. To obtain language-agnostic speaker representations, we leverage language pseudo-labels from the VoxCeleb2 dataset using a Spoken Language Recognition (SLR) model trained on the VoxLingua107 [27] dataset. For testing, we employ three original monolingual test sets from VoxCeleb1 [28], VoxSRC2020, and the multi-lingual VoxSRC2021 validation set [18], along with VoxCeleb1-B [1]. The model is tested on the multi-lingual NISP-B dataset to further evaluate the method on out-of-distribution data.

We create the NISP-B dataset from NISP [29], a multi-lingual speaker profiling dataset that features speech recordings in English, Hindi, Kannada, Malayalam, Tamil, and Telugu. This dataset includes intra-speaker cross-lingual trials and inter-speaker monolingual trials, with balanced representation across all languages.

## 5. Experimental Analysis

### 5.1. Baselines

We use ResNet-(S [17] and L [18]), ECAPA [19] and ReDimNet [20] as speaker encoder architectures which are trained from scratch. For the Language Encoder we use ECAPA architecture and is initialized from the model pre-trained on VoxLingua107. We evaluate using Equal Error Rate (EER) and minimum decision cost function (*minDCF*).

### 5.2. Results

Our experiments demonstrate the effectiveness of the proposed architecture in disentangling language information from speaker embeddings, leading to improved performance across various test sets. As shown in Table 1, our approach LASPA, consistently outperforms baseline models on VoxCeleb1-B, VoxSRC 2021 validation, and NISP-B cross-lingual test sets. Specifically, ReDimNet, ResNet-L, and ECAPA architectures with our disentanglement method achieve the lowest EER and *minDCF* values, with LASPA ReDimNet reducing the EER to 2.10 % on VoxSRC 2021 validation and 1.62 % on VoxCeleb1-B. Moreover, LASPA ReDimNet achieves notable performance gains with an EER of 10.22 % on NISP-B, a challenging multi-lingual dataset.

In Table 1, we observe similar trends, with our models consistently achieving lower EER and *minDCF* values on cleaned versions of monolingual VoxCeleb1, VoxCeleb1-E, VoxCeleb1-H, and VoxSRC 2020 validation sets. These results confirm that the disentanglement strategy improves the robustness of speaker embeddings, enabling language-agnostic performance across multi-lingual datasets.

The SLR metric refers to the ability of a system to correctly identify the language being spoken from an audio signal. Table 2 demonstrates the observed decrease in SLR accuracy from speakers in the VoxCeleb1-B dataset, indicating less language-specific information in  $E_{spk}$ . Table 2 also shows the cosine similarity between audios of the same speaker across languages. Higher scores indicate robust and consistent speaker embeddings for individual speakers, confirming that LASPA enhances speaker distinction. Column 3 of the table shows the Diarization

Table 2: Spoken Language Recognition (%) on VoxCeleb1-B, Cosine Similarity scores on VoxCeleb1 cl. and Diarization Error Rate (%) on DISPLACE dataset across models.

Model	SLR Accuracy (↓)	Similarity (↑)	DER (↓)
DISPLACE [30]	–	–	32.5
ResNet-S	88.4	0.348	30.4
ResNet-L	89.2	0.391	29.2
ECAPA	91.1	0.531	28.7
ReDimNet	81.2	0.612	28.5
<b>LASPA</b>	<b>78.5</b>	<b>0.641</b>	<b>28.2</b>

Table 3: Ablation Study Results reporting EER (%) on monolingual VoxCeleb1 cl. and bilingual VoxCeleb1-B dataset

Configuration	VoxCeleb1 cl. (↓)	VoxCeleb1-B (↓)
LASPA (Full Model)	<b>1.96</b>	<b>5.88</b>
No Prefix-Tuning	2.14	6.44
Speaker-only	2.27	7.00

Error Rate (DER) on the DISPLACE dataset [30, 31], showing that our approach outperforms the DISPLACE baseline.

### 5.3. Ablation Study

To further analyze the contribution of each component in LASPA, we conducted an ablation study, focusing on the impact of the language encoder and prefix-tuning. We compared the full LASPA model with two reduced configurations: one without prefix-tuning (but with the language encoder and decoder), and a "Speaker-only" model (excluding language-related components). We performed this experiment using ResNet-S as the speaker encoder model.

Table 3 presents the Equal Error Rate (EER) results on the VoxCeleb1-C cleaned and VoxCeleb1-B datasets for these configurations. As shown in Table 3, the LASPA model achieves the lowest EER on VoxCeleb1 cl. and VoxCeleb1-B. Both models that integrate language information significantly outperform the "Speaker-only" baseline. Additionally, prefix-tuners account for just 1.16 % of the model's total parameters, making them a highly efficient way to enhance performance.

## 6. Conclusion

In this work, we introduced LASPA, a novel approach for language-agnostic speaker disentanglement using prefix-tuned cross-attention. Our method effectively separates speaker identity from linguistic information, addressing key challenges in multi-lingual speaker recognition. By integrating prefix-tuning, we achieve efficient adaptation while using only a small fraction of the model's total parameters. Experimental results demonstrate that LASPA consistently improves speaker recognition performance across both monolingual and multi-lingual scenarios, including unseen languages. The reduction in equal error rate (EER) across multiple datasets validates the effectiveness of our approach in mitigating the influence of language on speaker embeddings. Given its efficiency and strong generalization capability, LASPA offers a promising direction for robust speaker recognition in diverse linguistic environments. Future work will explore further refinements and applications of prefix-tuning in speaker verification and related tasks.

## 7. References

- [1] K. Nam, Y. Kim, J. Huh, H.-S. Heo, J. weon Jung, and J. S. Chung, "Disentangled representation learning for multilingual speaker recognition," in *Interspeech 2023*, 2023, pp. 5316–5320.
- [2] S. M. E. M. A. Conkie, "Generating multilingual voices using speaker space translation based on bilingual speaker data," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7624–7628. [Online]. Available: <https://arxiv.org/pdf/2004.04972.pdf>
- [3] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," 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. 1180–1189. [Online]. Available: <https://proceedings.mlr.press/v37/ganin15.html>
- [4] S. H. Mun, M. H. Han, M. Kim, D. Lee, and N. S. Kim, "Disentangled speaker representation learning via mutual information minimization," in *2022 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2022, pp. 89–96.
- [5] L. Yi and M.-W. Mak, "Disentangled speaker embedding for robust speaker verification," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 7662–7666.
- [6] F. Tong, S. Zheng, H. Zhou, X. Xie, Q. Hong, and L. Li, "Deep Representation Decomposition for Rate-Invariant Speaker Verification," in *Proc. Speaker Odyssey*, 2022, pp. 228–232.
- [7] J. Kang, J. Huh, H. S. Heo, and J. S. Chung, "Augmentation adversarial training for self-supervised speaker representation learning," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1253–1262, 2022.
- [8] M. Arjovsky and L. Bottou, "Towards principled methods for training generative adversarial networks," in *International Conference on Learning Representations*, 2017. [Online]. Available: [https://openreview.net/forum?id=Hk4\\_qw5xe](https://openreview.net/forum?id=Hk4_qw5xe)
- [9] W. H. Kang, S. H. Mun, M. H. Han, and N. S. Kim, "Disentangled speaker and nuisance attribute embedding for robust speaker verification," *IEEE Access*, vol. 8, pp. 141 838–141 849, 2020.
- [10] Y.-C. Wang, C.-Y. Wang, and S.-H. Lai, "Disentangled representation with dual-stage feature learning for face anti-spoofing," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 1955–1964. [Online]. Available: [https://openaccess.thecvf.com/content/WACV2022/papers/Wang\\_Disentangled\\_Representation\\_With\\_Dual-Stage\\_Feature\\_Learning\\_for\\_Face\\_Anti-Spoofing\\_WACV\\_2022\\_paper.pdf](https://openaccess.thecvf.com/content/WACV2022/papers/Wang_Disentangled_Representation_With_Dual-Stage_Feature_Learning_for_Face_Anti-Spoofing_WACV_2022_paper.pdf)
- [11] 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.
- [12] J. Williams and S. King, "Disentangling style factors from speaker representations," in *Interspeech 2019*, 2019, pp. 3945–3949.
- [13] R. Pappagari, T. Wang, J. Villalba, N. Chen, and N. Dehak, "X-vectors meet emotions: A study on dependencies between emotion and speaker recognition," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7169–7173.
- [14] Z. Wang, W. Ge, X. Wang, S. Yang, W. Gan, H. Chen, H. Li, L. Xie, and X. Li, "Accent and speaker disentanglement in many-to-many voice conversion," in *2021 12th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, 2021, pp. 1–5.
- [15] D. Raj, D. Snyder, D. Povey, and S. Khudanpur, "Probing the information encoded in x-vectors," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, Dec. 2019. [Online]. Available: <http://dx.doi.org/10.1109/ASRU46091.2019.9003979>
- [16] J. Campbell, "Speaker recognition: a tutorial," *Proceedings of the IEEE*, vol. 85, no. 9, pp. 1437–1462, 1997.
- [17] 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," in *Interspeech 2020*, 2020, pp. 2977–2981.
- [18] Y. Kwon, H.-S. Heo, B.-J. Lee, and J. S. Chung, "The ins and outs of speaker recognition: lessons from voxsrc 2020," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5809–5813.
- [19] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification," in *Interspeech 2020*, 2020, pp. 3830–3834.
- [20] I. Yakovlev, R. Makarov, A. Balykin, P. Malov, A. Okhotnikov, and N. Torgashov, "Reshape dimensions network for speaker recognition," in *Interspeech 2024*, 2024, pp. 3235–3239.
- [21] X. L. Li and P. Liang, "Prefix-tuning: Optimizing continuous prompts for generation," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Online: Association for Computational Linguistics, Aug. 2021, pp. 4582–4597. [Online]. Available: <https://aclanthology.org/2021.acl-long.353>
- [22] A. Ghadiya, P. Kar, V. Chudasama, and P. Wasnik, "Cross-modal fusion and attention mechanism for weakly supervised video anomaly detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 1965–1974. [Online]. Available: [https://openaccess.thecvf.com/content/CVPR2024W/MULA/papers/Ghadiya\\_Cross-Modal\\_Fusion\\_and\\_Attention\\_Mechanism\\_for\\_Weakly\\_Supervised\\_Video\\_Anomaly\\_CVPRW\\_2024\\_paper.pdf](https://openaccess.thecvf.com/content/CVPR2024W/MULA/papers/Ghadiya_Cross-Modal_Fusion_and_Attention_Mechanism_for_Weakly_Supervised_Video_Anomaly_CVPRW_2024_paper.pdf)
- [23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf)
- [24] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 11 1997.
- [25] M.-C. Popescu, V. E. Balas, L. Perescu-Popescu, and N. Mastorakis, "Multilayer perceptron and neural networks," *WSEAS Transactions on Circuits and Systems*, vol. 8, no. 7, pp. 579–588, 2009.
- [26] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *Interspeech 2018*, 2018, pp. 1086–1090.
- [27] J. Valk and T. Alumäe, "Voxlingua107: A dataset for spoken language recognition," in *2021 IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 652–658.
- [28] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: A large-scale speaker identification dataset," in *Interspeech 2017*, 2017, pp. 2616–2620.
- [29] S. B. Kalluri, D. Vijayaseenan, S. Ganapathy, R. R. M, and P. Krishnan, "Nisp: A multi-lingual multi-accent dataset for speaker profiling," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6953–6957.
- [30] S. Baghel, S. Ramoji, Sidharth, R. H, P. Singh, S. Jain, P. Roy Chowdhuri, K. Kulkarni, S. Padhi, D. Vijayaseenan, and S. Ganapathy, "The displace challenge 2023 - diarization of speaker and language in conversational environments," in *Interspeech 2023*, 2023, pp. 3562–3566.
- [31] R. Gohil, R. Viswanathan, S. Agrawal, C. M. Vikram, M. R. Kamble, K. Sabu, M. A. B. Shaik, and K. K. S. Rajesh, "Ensemble of incremental system enhancements for robust speaker diarization in code-switched real-life audios," in *Speech and Computer*, A. Karpov, K. Samudravijaya, K. T. Deepak, R. M. Hegde, S. S. Agrawal, and S. R. M. Prasanna, Eds. Cham: Springer Nature Switzerland, 2023, pp. 490–502.