

ENHANCING SPEAKER VERIFICATION WITH W2V-BERT 2.0 AND KNOWLEDGE DISTILLATION GUIDED STRUCTURED PRUNING

Ze Li^{1,3}, Ming Cheng^{1,3}, Ming Li^{2,3†}

¹School of Computer Science, Wuhan University, Wuhan, China

²School of Artificial Intelligence, Wuhan University, Wuhan, China

³Suzhou Municipal Key Laboratory of Multimodal Intelligent Systems,
Duke Kunshan University, Kunshan, China

{lize389, ming.cheng}@whu.edu.cn, ming.li369@dukekunshan.edu.cn

ABSTRACT

Large-scale self-supervised Pre-Trained Models (PTMs) have shown significant improvements in the speaker verification (SV) task by providing rich feature representations. In this paper, we utilize w2v-BERT 2.0, a model with approximately 600 million parameters trained on 4.5 million hours of unlabeled data across 143 languages, for the SV task. The MFA structure with Layer Adapter is employed to process the multi-layer feature outputs from the PTM and extract speaker embeddings. Additionally, we incorporate LoRA for efficient fine-tuning. Our model achieves state-of-the-art results with 0.12% and 0.55% EER on the Vox1-O and Vox1-H test sets, respectively. Furthermore, we apply knowledge distillation guided structured pruning, reducing the model size by 80% while achieving only a 0.04% EER degradation. Source code and models are released at https://github.com/ZXHY-82/w2v-BERT-2.0_SV.

Index Terms— speaker verification, w2v-BERT 2.0, LoRA, knowledge distillation, structured pruning

1. INTRODUCTION

Speaker verification (SV) aims to authenticate the identity of a speaker by analyzing the voice signal. In recent years, significant advancements in deep learning, coupled with large-scale labeled datasets [1–5], have led to substantial improvements in the performance of deep neural network-based SV systems [6–8].

However, the scale of existing labeled datasets remains insufficient to meet the increasing complexity of model architectures. As a result, researchers have turned to large-scale Pre-Trained Models (PTMs) [9–13], which are typically trained on hundreds of thousands or even millions of hours of unlabeled speech data. These models offer powerful feature representations that can significantly enhance performance on downstream tasks. Chen et al. [13, 14] employ a layer-wise weighted average of PTM’s features, followed by a speaker model such as ECAPA-TDNN [6] for the SV task. Kim et al. [15] introduces Layer-wise Attentive Pooling, which applies time-dynamic weighting to multi-layer representations, overcoming the limitation of conventional weighted summation that ignores

certain layers. Peng et al. [16] introduces Context-Aware Multi-Head Factorized Attentive Pooling, which incorporates contextual information and grouped queries, thereby obtaining more robust utterance-level representations. Zhao et al. [17] and Cai et al. [18] build upon the concept of MFA-Conformer [19], concatenating all or part of the features from different PTM layers to capture richer and more comprehensive speaker representations.

Previous studies focused on Transformer-based self-supervised PTMs for the SV task [13–16]. In contrast, w2v-BERT 2.0 [12] is a self-supervised PTM built on a Conformer-based architecture, which has been demonstrated by MFA-Conformer [19] to be effective for SV and superior to the Transformer-based architecture. Furthermore, w2v-BERT 2.0 adopts a training strategy that optimizes both a contrastive loss and a masked prediction loss simultaneously, and trained on 4.5 million hours of unlabeled audio covering 143 languages, leading to strong performance on audio classification tasks.

In this work, we utilize w2v-BERT 2.0 as the encoder for the SV task. Speaker embeddings are extracted using the MFA [19] structure, and a Layer Adapter [18] module is introduced for each layer’s output before concatenation, enabling better adaptation of the PTM’s output to the specific task domain. Additionally, Low-Rank Adaptation(LoRA) [20] is employed for efficient fine-tuning. To enhance the model’s practicality for real-world deployment, we apply a knowledge distillation guided structured pruning technique [21], which prunes the PTM while minimizing performance degradation.

The primary contributions of this paper can be summarized as follows:

- We are the first to apply the w2v-BERT 2.0 PTM to the SV task, achieving state-of-the-art(SOTA) results of 0.12% and 0.55% EER on the Vox1-O and Vox-H test sets, respectively.
- We employ the MFA structure, combined with the Layer Adapter and LoRA modules, to efficiently adapt the model to the SV task.
- We utilize a knowledge distillation guided structured pruning strategy, reducing the model size by 80% with only a 0.04% EER degradation.

2. METHODS

2.1. Pre-trained Model: w2v-BERT 2.0

w2v-BERT 2.0 [12] is a large-scale multilingual self-supervised model, designed for speech representation and introduced in the SeamlessM4T framework [12]. Building on the w2v-BERT [22] architecture, it consists of 24 conformer layers and integrates both

[†]Corresponding Author: Ming Li.

This research is funded in part by the National Natural Science Foundation of China (62571223), Yangtze River Delta Science and Technology Innovation Community Joint Research Project (2024CSJGG01100) and Guangdong Science and Technology Plan (2023A1111120012). Many thanks for the computational resource provided by the Advanced Computing East China Sub-Center.

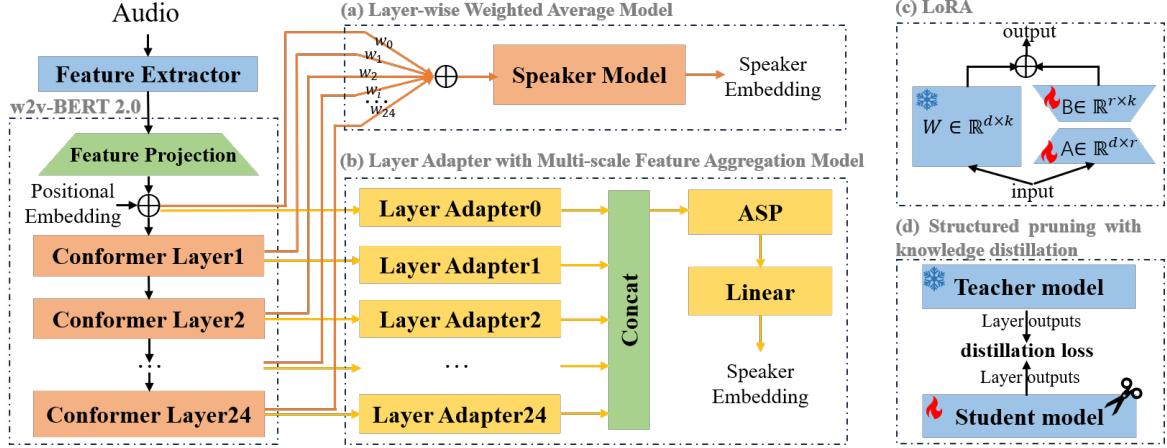


Fig. 1: Module architecture for speaker verification with w2v-BERT 2.0 and knowledge distillation guided structured pruning.

contrastive learning and masked language modeling. The model is trained on 4.5 million hours of unlabeled audio data, covering 143 languages. In this paper, we apply w2v-BERT 2.0 for the SV task. Given a speech utterance x , we first extract its fbank features, and then input them into the PTM to obtain the features of each layer:

$$[h_0, h_1, \dots, h_L] = \text{W2v-BERT-2.0}(\text{Fbank}(x)) \quad (1)$$

where $h_i \in \mathbb{R}^{D \times T}$ represents the output of the i -th conformer layer, with D as the hidden dimension and T as the number of frames.

2.2. Layer-wise Weighted Average Model

The layer-wise weighted average approach [13] is currently one of the most widely used and effective methods for fine-tuning a PTM on the SV task. In this method, each layer of the PTM is assigned a learnable weight, which is updated during training. The final frame feature H is obtained by computing a weighted average of all layer outputs as shown in Fig.1(a), replacing the Fbank feature fed into the speaker model for extracting the speaker embedding.

$$H = \sum_{i=0}^L \frac{e^{w_i}}{\sum_{j=0}^L e^{w_j}} \cdot h_i \quad (2)$$

where w_i is the weight of the i -th layer, and h_i is the feature output from the i -th layer.

2.3. Multi-scale Feature Aggregation Model

Another strategy for fine-tuning PTMs is Multi-scale Feature Aggregation. Following MFA-Conformer [19], the features of all layers are concatenated and fed into an Attention Statistics Pooling (ASP) module [23]. Unlike weighted averaging, this direct concatenation preserves the full layer information, while ASP learns the relative importance across layers and dimensions. The speaker embedding E is then obtained via a linear transformation of the ASP output.

$$E = \text{Linear}(\text{ASP}(\text{Concat}(h_0, h_1, \dots, h_L))) \quad (3)$$

2.4. Layer Adapter for Model Adaptation

Moreover, considering that directly using the raw layer features for the SV task could lead to poor generalization, we introduce a lightweight Layer Adapter [18] module for each layer output before

concatenation. The adapter structure consists of two linear layers followed by layer normalization and a rectified linear unit activation function. The first linear layer projects the input feature from dimension d to a hidden size of d' , while the second linear layer maps the d' -dimensional representation to another d' -dimensional space.

2.5. LoRA for Model Adaptation

Compared to full fine-tuning, LoRA [20] adapts PTMs efficiently by introducing a small number of trainable parameters in a low-rank space, reducing both computational and memory costs while maintaining effective task adaptation. In this paper, we apply LoRA to the query and value weights of PTM's self-attention module. The update mechanism for the model's weight matrix is described as follows:

$$W' = W + \frac{\alpha}{r} \cdot A \cdot B \quad (4)$$

where $W \in \mathbb{R}^{d \times k}$ is the original weight matrix, $A \in \mathbb{R}^{d \times r}$ and $B \in \mathbb{R}^{r \times k}$ are the low-rank matrices, r is the rank of the adaptation, and α is a scaling factor that controls the magnitude of the update.

2.6. Structured Pruning with Knowledge Distillation

The large parameter size and computational cost of PTMs pose significant challenges for deployment on resource-constrained devices. Inspired by [21], we apply knowledge distillation guided structured pruning to the w2v-BERT 2.0 model. To preserve the original representational capacity of the model, a teacher–student framework is employed, aligning the outputs of the pruned student model with those of the unpruned teacher model. The distillation loss combines L1 and cosine distances with equal weights:

$$\mathcal{L}_{\text{distill}} = \sum_{l=0}^L \sum_{t=1}^T \left(L_1(h_i^t, \hat{h}_i^t) - \text{cosine}(h_i^t, \hat{h}_i^t) \right) \quad (5)$$

where L denotes the number of layers, T denotes the number of frames, h_i^t and \hat{h}_i^t represent the t -th frame output of the i -th layer from the teacher and student models, respectively.

Pruning is achieved by optimizing the L0 regularization term $\|\theta\|_0$, where θ denotes the parameters to be pruned. However, the L0 term is discrete and non-differentiable. To address this, the parameters targeted for pruning are modeled as random variables governed by the Hard Concrete distribution, as described in [21]:

$$\theta = \{\hat{\theta}_j z_j\}_{j=1}^J, \quad z_j \sim q(z_j | \alpha_j) \quad (6)$$

where $\hat{\theta}_j$ denotes the j -th group of prunable parameters, and z_j is a stochastic binary gate sampled from the Hard Concrete distribution:

$$s_j = \sigma\left(\frac{\log u_j - \log(1 - u_j) + \log \alpha_j}{\beta}\right), \quad (7)$$

$$z_j = \min(1, \max(0, (\zeta - \gamma) \cdot s_j + \gamma))$$

where u_j is drawn from a uniform distribution, β is a temperature parameter controlling the smoothness of s_j , and ζ and γ control the upper and lower bounds of s_j . We set $\beta = 2/3$, $\gamma = -0.1$, and $\zeta = 1.1$. Finally, the expected value of the L0 norm is given by:

$$\mathbb{E}_{q(\theta|\hat{\theta}, \alpha)}[||\theta||_0] = \sum_{j=1}^{|G|} |g_j| \cdot \sigma\left(\log \alpha_j - \beta \log \frac{-\gamma}{\zeta}\right) \quad (8)$$

where $|G|$ denotes the number of groups and $|g_j|$ is the number of parameters in the j -th group.

The final loss function employs the augmented Lagrangian method [24] for more effective fine-grained control of the sparsity in the pruned model:

$$\max_{\lambda_1, \lambda_2} \min_{\theta, \alpha} \mathbb{E}_{q(\theta|\hat{\theta}, \alpha)} [\mathcal{L}_{\text{distill}} + \lambda_1(||\theta||_0 - t) + \lambda_2((||\theta||_0 - t)^2)] \quad (9)$$

where λ_1 and λ_2 are learnable Lagrange multipliers and t represents the predefined target sparsity.

3. EXPERIMENTS

3.1. Datasets

The experiments are conducted using the VoxCeleb1&2 [1, 2], VoxBlink2 [3] and CN-Celeb1&2 [4, 5] datasets. For VoxCeleb model training, we utilize the VoxCeleb2 development set and the VoxBlink2 dataset. During the evaluation phase, both the VoxCeleb1 development and test sets are used. The SV performance is evaluated based on three official trial lists: Vox1-O, Vox1-E and Vox1-H. For CN-Celeb, only the development sets of CN-Celeb1 and CN-Celeb2 are used for training. We choose to average all the embeddings that belong to the same enrollment speaker to get the final speaker embedding for the CN-Celeb test set evaluation.

3.2. Model Configuration

We use w2v-BERT 2.0 as the encoder to extract features from each layer, followed by the design of four distinct modules for speaker embedding extraction, as described below:

Layer-wise Weighted Average Model: Similar to [13], we utilize the small ECAPA-TDNN model as the speaker model to process the weighted average of all layer outputs.

MFA Model: This model consists of an ASP module and a Linear module, where the outputs from all layers are concatenated and directly fed into the ASP module. The speaker embedding dimension is set to 256, and the hidden dimension of the ASP module is matched to the layer dimension.

Layer Adapter with MFA model: Building upon the MFA model, we add a Layer Adapter module after each layer. The hidden dimension d' of the Layer Adapter is set to 128, while the ASP module's hidden dimension is also set to d' .

LoRA with Layer Adapter and MFA model: In this model, LoRA is applied to the query and value linear layers of the self-attention modules in each Conformer layer of the PTM. The rank r is set to 64, and the weight scaling factor α is set to 128.

3.3. Training Details

Our training process is divided into three stages as follows:

i) PTM freeze training: In this phase, the PTM is frozen. The acoustic features are 80-dimensional fbank coefficients with a frame length of 25ms and a hop size of 10ms. These features are normalized using mean and standard deviation before being fed into the PTM. On-the-fly data augmentation [27] is applied by adding background noise or convolutional reverberation noise. The MU-SAN [28] and RIR Noise [29] datasets are used as noise sources and room impulse response functions, respectively. The speed perturbation [30], which speeds up or down each utterance by a factor of 0.9 or 1.1, is applied to yield shifted pitch utterances that are considered from new speakers, but is not utilized during training with the VoxBlink2 dataset. AdamW [31] optimizer with weight decay of 1e-4 is used, along with a StepLR scheduler with 5 epochs decay. The learning rate starts at 1e-4 and decreases to 1e-5, with a decay factor of 0.1. The margin and scale of ArcFace [32] are set to 0.2 and 32, respectively. A linear warm-up learning rate schedule is used for the first 5 epochs to stabilize training. The input frame length is set between 200 and 300.

ii) Joint fine-tuning: Subsequently, the PTM is unfrozen for fine-tuning. The learning rate starts at 1e-5 and decays to 5e-6 using a cosine decay schedule over 2 epochs, with a total of 4 epochs dedicated to fine-tuning.

iii) Large margin fine-tuning and score calibration: In this stage, the Large-Margin Fine-Tune (LMFT) [33] strategy is introduced, using only the VoxCeleb2 dataset. All data augmentation strategies are stopped. The input frame length is set between 500 and 600. For ArcFace, a margin of 0.5 is applied. The learning rate starts at 1e-5 and decays to 5e-6 using a cosine decay schedule over 1 epoch, with a total of 2 epochs dedicated to fine-tuning. Additionally, AS-norm [34] and QMF [35] are used for scoring calibration.

3.4. Pruning Details

We perform structured pruning on the joint fine-tuned w2v-BERT 2.0 model, focusing on the FFN intermediate dimensions, convolution channels, and the number of attention heads of each Conformer layer. First, the teacher-student framework is initialized with the w2v-BERT 2.0 model, the teacher model remains frozen. The target sparsity increases linearly to the pre-set value over the first 10,000 items. The total number of epochs is 20. AdamW optimizer is used, with a learning rate of 2e-4 and 2e-2 for the student model parameters and sparsity-related parameters, respectively. After pruning, the pruned student model is further distilled by an additional 20 epochs from the teacher model. Finally, the pruned student model replaces the initial joint fine-tuned w2v-BERT 2.0 model, and further fine-tuning on the SV task based on the stages outlined in Section 3.3.

4. RESULTS

Table 2 shows the SV performance of different w2v-BERT 2.0 based model architectures, as described in Section 3.2. The results show that the features extracted from the w2v-BERT 2.0 provide substantial benefits for the SV task. Even with the simplest MFA structure, the model achieves an EER of 0.26% on the Vox1-O test set. The introduction of the Layer Adapter effectively transforms the raw features from the PTM to better suit the SV task, while significantly reducing the number of parameters from 65.6M to 6.2M through dimensionality scaling, resulting in an improved EER of 0.18%. For the model using the layer-wise weighted average method, although

Table 1: Performance comparison of the w2v-BERT 2.0 based SV model with other SV models.

Frontend	Model	Params	LMFT	Score calibration	Vox1-O		Vox1-E		Vox1-H		CN-Celeb Test	
					EER(%)	mDCF	EER(%)	mDCF	EER(%)	mDCF	EER(%)	mDCF
Fbank	ECAPA-TDNN(C=1024) [6]	14.7M	×	×	0.87	0.107	1.12	0.132	2.12	0.210	-	-
	CAM++ [7]	7.2M	×	×	0.73	0.091	0.89	0.100	1.76	0.173	6.78 [‡]	0.383 [‡]
	ReDimNet-B6 [8]	15.0M	✓	✓	0.37	0.030	0.53	0.051	1.00	0.097	-	-
	ERes2NetV2 [25]	17.8M	✓	×	0.61	0.054	0.76	0.082	1.45	0.143	6.04 [‡]	0.362 [‡]
	ResNet221 [26]	23.8M	✓	✓	0.51	-	0.68	-	1.21	-	5.66 [‡]	0.330 [‡]
	ResNet293 [3]	98.9M	✓	✓	0.17*	0.006*	0.37*	0.037*	0.68*	0.070*	-	-
HuBERT Large	ECAPA-TDNN(C=512) [14]	317+8.8M	✓	✓	0.59	-	0.65	-	1.23	-	-	-
Wav2Vec2.0 Large	ECAPA-TDNN(C=512) [14]	317+8.8M	✓	✓	0.59	-	0.63	-	1.14	-	-	-
UniSpeech-SAT Large	ECAPA-TDNN(C=512) [14]	317+8.8M	✓	✓	0.54	-	0.57	-	1.18	-	-	-
WavLM Large	ECAPA-TDNN(C=512) [13]	317+8.8M	✓	✓	0.38	-	0.48	-	0.99	-	-	-
	CA-MHFA [16]	317+2.3M	✓	✓	0.42	-	0.48	-	0.96	-	-	-
	LAP+ASTP [15]	317+2.3M	✓	✓	0.37	0.059	0.50	0.055	1.01	0.099	-	-
Nemo Large	MFA [18]	131M	✓	✓	0.43	0.062	0.66	0.071	1.35	0.135	-	-
w2v-BERT 2.0	LoRA_Adapter_MFA	580+6.2M	×	×	0.23*	0.029*	0.38*	0.040*	0.81*	0.082*	4.67[‡]	0.297[‡]
			✓	×	0.14*	0.020*	0.31*	0.032*	0.73*	0.071*	-	-
			✓	✓	0.12*	0.025*	0.27*	0.028*	0.55*	0.051*	-	-

* indicates results obtained using the VoxCeleb2 and VoxBlink2 datasets for training. [‡] indicates results obtained using only the CN-Celeb1&2 datasets for training.

Table 2: Performance comparison of different w2v-BERT 2.0 based model architectures.

Model	Data	Params	Vox1-O EER
ECAPA-TDNN(freeze PTM)	VoxCeleb2	580+8.8M	0.49%
+ LMFT (Joint Fine-tuning)			0.26%
+ Joint Fine-tuning			0.29%
++ LMFT			0.22%
MFA(freeze PTM)	VoxCeleb2	580+65.6M	0.46%
+ LMFT (Joint Fine-tuning)			0.28%
+ Joint Fine-tuning			0.38%
++ LMFT			0.26%
Adapter_MFA(freeze PTM)	VoxCeleb2	580+6.2M	0.43%
+ LMFT (Joint Fine-tuning)			0.22%
+ Joint Fine-tuning			0.28%
++ LMFT			0.18%
LoRA_Adapter_MFA(freeze PTM)	VoxCeleb2	580+12.5M	0.30%
+ LMFT (Joint Fine-tuning)			0.22%
+ Joint Fine-tuning (LoRA merge)			0.30%
++ LMFT			0.24%
LoRA_Adapter_MFA(freeze PTM)	& VoxBlink2	580+12.5M	0.27%
+ LMFT (Joint Fine-tuning)			0.15%
+ Joint Fine-tuning (LoRA merge)			0.23%
++ LMFT			0.14%
CN-Celeb Test EER			
Adapter_MFA(freeze PTM)	CnCeleb1 &2	580+6.2M	6.51%
+ Joint Fine-tuning			5.17%
LoRA_Adapter_MFA(freeze PTM)	CnCeleb1 &2	580+12.5M	4.87%
+ Joint Fine-tuning (LoRA merge)		580+6.2M	4.67%

it is followed by a powerful ECAPA-TDNN network, the weighted summation of features across layers leads to greater information loss compared to feature concatenation, ultimately achieving an EER of only 0.22%. The use of the LoRA module has significantly enhanced training efficiency. In particular, during the PTM freezing phase, the model’s performance on the Vox1-O test set improved from 0.43% to 0.30%, and on the CN-Celeb test set from 6.51% to 4.87%. Additionally, before fine-tuning with PTM unfrozen, the LoRA module’s weights are merged into the PTM. However, it is important to note that employing the LoRA module on small and simple dataset may pose a risk of overfitting. For instance, when training only on the VoxCeleb2 dataset, performance slightly declined after unfreezing the PTM. After incorporating the VoxBlink2 dataset, the overfitting

Table 3: Results of knowledge distillation guided structured pruning on a w2v-BERT 2.0 based SV model trained on VoxCeleb2 and VoxBlink2. MACs and FLOPs were measured on 1-s long segments.

Model	Sparsity	Params	MACs	FLOPs	LMFT	Vox1-O EER
LoRA_Adapter_MFA	0%	580+6.2M	28.75G	57.72G	×	0.23%
	≈ 80%	124+6.2M	6.31G	12.75G	✓	0.14%

issue was effectively mitigated.

Table 1 reports a comparison between our w2v-BERT 2.0 based SV model and other SV models. After LMFT and score calibration, our model achieves an EER of 0.12% on the Vox1-O test set, outperforming the SOTA ResNet293’s result of 0.17% EER [3]. Moreover, when trained only on the VoxCeleb2 dataset, the current SOTA PTM-based model yields 0.37% EER on the Vox1-O test set [15], whereas our model achieves 0.18% EER only using LMFT. Additionally, our model achieves an EER of 4.67% on the CN-Celeb test set by only using the CnCeleb data for training, further demonstrating the robustness and generalization.

Table 3 presents the results of knowledge distillation guided structured pruning applied to the w2v-BERT 2.0 based SV model. At a sparsity level of 80%, our approach achieves an EER of 0.18% on the Vox1-O test set after LMFT. Compared to the baseline system, the performance degradation is only 0.04% EER, demonstrating promising potential for practical deployment.

5. CONCLUSION

In this paper, we explore the application of the w2v-BERT 2.0 PTM in the SV task. We adopt a Layer Adapter based MFA framework, combined with efficient fine-tuning via LoRA, to aggregate multi-layer features from the PTM and extract speaker embeddings. The experimental results show that our model achieves the SOTA performance, with an EER of 0.12% on the Vox1-O test set and 4.67% on the CN-Celeb test set. Furthermore, to enhance practical deployability, we apply knowledge distillation guided structured pruning, reducing the PTM’s parameter count by 80% while incurring only a 0.04% increase in EER. Source code and models are released.

6. REFERENCES

- [1] A. Nagrani, J. S. Chung, and A. Zisserman, “Voxceleb: A large-scale speaker identification dataset,” in *Proc. Interspeech*, 2017, pp. 2616–2620.
- [2] J. Chung, A. Nagrani, and A. Zisserman, “Voxceleb2: Deep Speaker Recognition,” in *Proc. Interspeech*, 2018.
- [3] Y. Lin, M. Cheng, F. Zhang *et al.*, “VoxBlink2: A 100K+ Speaker Recognition Corpus and the Open-Set Speaker-Identification Benchmark,” in *Proc. Interspeech*, 2024, pp. 4263–4267.
- [4] Y. Fan, J. Kang, L. Li *et al.*, “Cn-celeb: a challenging chinese speaker recognition dataset,” in *Proc. ICASSP*, 2020, pp. 7604–7608.
- [5] L. Li, R. Liu, J. Kang *et al.*, “Cn-celeb: multi-genre speaker recognition,” *Speech Communication*, vol. 137, pp. 77–91, 2022.
- [6] 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.
- [7] H. Wang, S. Zheng, Y. Chen, L. Cheng, and Q. Chen, “Cam++: A fast and efficient network for speaker verification using context-aware masking,” in *Proc. Interspeech*, 2023, pp. 5301–5305.
- [8] I. Yakovlev, R. Makarov, A. Balykin *et al.*, “Reshape Dimensions Network for Speaker Recognition,” in *Proc. Interspeech*, 2024, pp. 3235–3239.
- [9] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai *et al.*, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Trans. ASLP*, vol. 29, pp. 3451–3460, 2021.
- [10] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, 2020, pp. 12449–12460.
- [11] S. Chen, Y. Wu, C. Wang *et al.*, “Unispeech-sat: Universal speech representation learning with speaker aware pre-training,” in *Proc. ICASSP*, 2022, pp. 6152–6156.
- [12] L. Barrault, Y.-A. Chung, M. C. Meglioli *et al.*, “Seamless: Multilingual expressive and streaming speech translation,” *arXiv preprint arXiv:2312.05187*, 2023.
- [13] S. Chen, C. Wang, Z. Chen *et al.*, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE J-STSP*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [14] Z. Chen, S. Chen, Y. Wu *et al.*, “Large-scale self-supervised speech representation learning for automatic speaker verification,” in *Proc. ICASSP*, 2022, pp. 6147–6151.
- [15] J. S. Kim, H. J. Park, W. Shin, and S. W. Han, “Rethinking leveraging pre-trained multi-layer representations for speaker verification,” in *Proc. Interspeech*, 2025, pp. 3713–3717.
- [16] J. Peng, L. Mošner, L. Zhang, O. Plchot, T. Stafylakis, L. Burget, and J. Černocký, “Ca-mhfa: A context-aware multi-head factorized attentive pooling for ssl-based speaker verification,” in *Proc. ICASSP*, 2025, pp. 1–5.
- [17] Y. Zhao, S. Wang, G. Sun, Z. Chen, C. Zhang, M. Xu, and T. F. Zheng, “Whisper-pmfa: Partial multi-scale feature aggregation for speaker verification using whisper models,” in *Proc. Interspeech*, 2024, pp. 2680–2684.
- [18] D. Cai and M. Li, “Leveraging asr pretrained conformers for speaker verification through transfer learning and knowledge distillation,” *IEEE/ACM Trans. ASLP*, vol. 32, pp. 3532–3545, 2024.
- [19] Y. Zhang, Z. Lv, H. Wu, S. Zhang, P. Hu, Z. Wu *et al.*, “MFA-Conformer: Multi-scale Feature Aggregation Conformer for Automatic Speaker Verification,” in *Proc. Interspeech*, 2022, pp. 306–310.
- [20] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang *et al.*, “Lora: Low-rank adaptation of large language models.” in *Proc. ICLR*, 2022.
- [21] J. Han, P. Pálka, M. Delcroix, F. Landini, J. Rohdin, J. Černocký, and L. Burget, “Efficient and generalizable speaker diarization via structured pruning of self-supervised models,” *arXiv preprint arXiv:2506.18623*, 2025.
- [22] Y.-A. Chung, Y. Zhang, W. Han *et al.*, “W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training,” in *Proc. ASRU*, 2021, pp. 244–250.
- [23] K. Okabe, T. Koshinaka, and K. Shinoda, “Attentive statistics pooling for deep speaker embedding,” in *Proc. Interspeech*, 2018, pp. 2252–2256.
- [24] Z. Wang, J. Wohlwend, and T. Lei, “Structured pruning of large language models,” in *Proc. EMNLP*, 2020, pp. 6151–6162.
- [25] Y. Chen, S. Zheng, H. Wang, L. Cheng, T. Zhu, R. Huang *et al.*, “3d-speaker-toolkit: An open-source toolkit for multi-modal speaker verification and diarization,” in *Proc. ICASSP*, 2025, pp. 1–5.
- [26] H. Wang, C. Liang, S. Wang, Z. Chen, B. Zhang, X. Xiang *et al.*, “Wespeaker: A research and production oriented speaker embedding learning toolkit,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [27] W. Cai, J. Chen, J. Zhang, and M. Li, “On-the-Fly Data Loader and Utterance-Level Aggregation for Speaker and Language Recognition,” *IEEE/ACM Trans. ASLP*, pp. 1038–1051, 2020.
- [28] D. Snyder, G. Chen, and D. Povey, “Musan: A music, speech, and noise corpus,” *arXiv preprint arXiv:1510.08484*, 2015.
- [29] T. Ko, V. Peddinti, D. Povey, M. Seltzer, and S. Khudanpur, “A study on data augmentation of reverberant speech for robust speech recognition,” in *Proc. ICASSP*, 2017, pp. 5220–5224.
- [30] W. Wang, D. Cai, X. Qin, and M. Li, “The dku-dukeee systems for voxceleb speaker recognition challenge 2020,” *arXiv preprint arXiv:2010.12731*, 2020.
- [31] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *Proc. ICLR*, 2019.
- [32] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, “Arcface: Additive angular margin loss for deep face recognition,” in *Proc. CVPR*, 2019, pp. 4690–4699.
- [33] 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, pp. 5814–5818.
- [34] P. Matejka, O. Novotný, O. Plchot *et al.*, “Analysis of score normalization in multilingual speaker recognition.” in *Proc. Interspeech*, 2017, pp. 1567–1571.
- [35] Z. Li, Y. Lin, X. Qin, N. Jiang, G. Zhao, and M. Li, “The dku-msxf speaker verification system for the voxceleb speaker recognition challenge 2023,” *arXiv preprint arXiv:2308.08766*, 2023.