

Leveraging ASR Pretrained Conformers for Speaker Verification through Transfer Learning and Knowledge Distillation

Danwei Cai, *Student Member, IEEE* and Ming Li, *Senior Member, IEEE*

Abstract—This paper focuses on the application of Conformers in speaker verification. Conformers, initially designed for Automatic Speech Recognition (ASR), excel at modeling both local and global contexts within speech signals effectively. Previous research has established that ASR and speaker verification tasks can naturally complement each other. Building on this synergistic relationship, this study introduces three strategies for leveraging ASR-pretrained Conformers in speaker verification: (1) **Transfer learning:** We use a pretrained ASR Conformer encoder to initialize the speaker embedding network, thereby enhancing model generalization and mitigating the risk of overfitting. (2) **Knowledge distillation:** We distill the complex capabilities of an ASR Conformer into a speaker verification model. This not only allows for flexibility in the student mode’s network architecture but also incorporates frame-level ASR distillation loss as an auxiliary task to reinforce speaker verification. (3) **Parameter-efficient transfer learning with speaker adaptation:** A lightweight speaker adaptation module is proposed to convert ASR-derived features into speaker-specific embeddings, without altering the core architecture of the original ASR Conformer. This strategy facilitates the concurrent execution of ASR and speaker verification tasks within a singular model. Experiments were conducted on VoxCeleb datasets. The results are compelling: models employing ASR pretraining and knowledge distillation significantly outperform standard Conformers. Specifically, the best model using the ASR pretraining method achieved a 0.43% equal error rate (EER) on the VoxCeleb1-O test trial, while the knowledge distillation approach yielded a 0.38% EER. Furthermore, by adding a mere 4.92 million parameters to a 130.94 million-parameter ASR Conformer encoder, the speaker adaptation approach achieved a 0.45% EER, enabling parallel speech recognition and speaker verification within a single ASR Conformer encoder. Overall, our techniques successfully transfer rich ASR knowledge to advanced speaker modeling.

Index Terms—Speaker recognition, automatic speech recognition, Conformer, transfer learning, knowledge distillation

I. INTRODUCTION

SPEAKER verification, which analyzes speech signals to verify the speaker’s identity, has many applications, from voice assistants to security systems. Over the past five years, the performance of speaker verification systems has improved remarkably due to the application of deep neural networks (DNN) [1], [2]. Numerous innovations have been introduced

in network architecture [3], [4], [5], [6], training objectives [7], [8], [9], and training strategies [10], [11] specifically tailored to speaker verification models.

Prevalent network architectures in speaker verification systems are convolutional neural networks (CNNs) and time-delay neural networks (TDNNs). The key strength of CNNs and TDNNs lies in their ability to model local feature patterns effectively, which is crucial in identifying speaker-specific vocal traits. These networks have been further advanced through variants of CNN and TDNN that incorporate residual connections [12], squeeze and excitation operations [13], [6], Res2Net blocks [14], [5], [6], and ResNeXt blocks [15], [5]. These modifications have significantly improved speaker verification performance.

Despite their successful applications, TDNNs, CNNs, and their variants face limitations in extracting long-range global context, especially without deep layers. As an alternative, Transformers, with their multi-head attention mechanism, have demonstrated a more robust ability to capture global context with less fine-grained local patterns [16]. To bridge this gap, Conformer combines the convolution module with Transformer to effectively capture local and global contextual information, leading to promising results in end-to-end automatic speech recognition (ASR) [17]. Recently, Zhang *et al.* introduced multi-scale feature aggregation Conformer (MFA-Conformer) for speaker verification [18]. MFA-Conformer concatenates frame-level outputs from all Conformer blocks to enhance speaker trait extraction in speaker verification. Liao *et al.* equipped the Conformer encoder with length-scaled attention and sharpness-aware minimization training for speaker verification [19]. However, despite their strengths, Conformers are susceptible to overfitting, particularly when faced with limited data or when employing large model parameters. This challenge is acute in speaker verification, where the diversity and amount of training data may be constrained [18], [20].

The Conformer model’s ability to capture both local and global contexts is leveraged in ASR and speaker verification. ASR focuses on recognizing the linguistic content of the speech, with a higher emphasis on frame-level details. In contrast, speaker verification targets identifying speaker-specific traits derived from the speech, centering on utterance-level context. Despite these differences, the two tasks can complement each other. For instance, the frame-level phoneme modeling undertaken in ASR could support speaker verification by aiding the detection of unique speaker-specific articulation patterns. Prior studies provide evidence of this

D. Cai, and M. Li are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC, 27705, USA, e-mail: {danwei.cai, ming.li369}@duke.edu.

M. Li is also with the Suzhou Municipal Key Laboratory of Multimodal Intelligent Systems, Data Science Research Center, Duke Kunshan University, Kunshan, China.

Corresponding author: Ming Li.

synergy, showing that phoneme modeling improves speaker verification in speaker embedding networks [21] as well as the i-vector statistical model [22], [23].

In light of the above, our research aims at leveraging ASR Conformers for speaker verification in three distinct ways. This builds upon our prior research on transfer learning using a pretrained ASR Conformer, which forms our first proposed method in this paper [20]. The technique involves initializing the speaker embedding network with a Conformer pretrained on a large-scale ASR dataset. This approach addresses the tendency of Conformers to overfit with limited data [18], [20] by leveraging a model pretrained on extensive ASR data. The pretrained ASR Conformer, which learns rich features from a large ASR dataset, reduces the data requirements for the speaker verification task and enhance the model’s generalization ability. Experimental results indicate that our ASR-pretrained method outperforms alternatives across various model sizes. Notably, the best system with ASR pre-training achieved an EER of 0.48% on the VoxCeleb 1-O trials, marking a 50% relative improvement compared to its counterpart without ASR pretraining.

Second, we propose using knowledge distillation [24] to transfer knowledge from the ASR task to the speaker verification task. One challenge with straightforward transfer learning is its inherent constraint on network architecture. When using a pretrained ASR Conformer for speaker verification, the speaker model is often constrained to adopt the same network architecture as the pretrained ASR model. To overcome this limitation, we use knowledge distillation. In this process, a student model, a simpler neural network, is trained to mimic the behavior of the more complex, pretrained teacher ASR Conformer. Rather than directly replicating weights and structure, knowledge distillation transfers the functional knowledge from the teacher to the student model. This not only retains the flexibility of network architecture for the speaker verification model but also harnesses the rich information in the pretrained ASR Conformers. Furthermore, our tailored knowledge distillation procedure, bridging ASR to speaker verification, integrates phoneme recognition as an auxiliary task. This alignment reinforces the synergy between ASR and speaker verification tasks, ensuring the speaker verification model captures the nuanced phonetic differences recognized by the ASR Conformer. Experimental results prove the efficacy of our method: it consistently improves speaker verification performance over the baseline method across various architectures and frequently surpasses the ASR-pretrained approach.

Finally, we propose an adaptation mechanism to unify the tasks of ASR and speaker verification within a single Conformer model. The motivation for this approach lies in tackling the inherent inefficiency of maintaining separate models for ASR and speaker verification tasks. Such a unified Conformer has diverse applications. For example, our unified model streamlines the process in scenarios where ASR and speaker verification are sequentially needed, such as voice assistants authenticating a user and then transcribing their commands. To achieve this goal, we introduce the speaker adaptation method to transform the features learned from the ASR task into those suitable for speaker verification without changing

the inputs and outputs of the ASR Conformer. The viability of this approach is supported by the speaker information preserved in the layer outputs of the ASR Conformer encoder. Our exploratory linear probe experiments indicate that the lower layers of the ASR Conformer retain more speaker information than the upper layers. This speaker adaptation approach, therefore, represents a resource-efficient strategy that allows for the simultaneous and efficient execution of both ASR and speaker verification tasks using a single Conformer. Experiments demonstrate that incorporating a speaker adaptation module (4.92 million parameters) into a pretrained ASR Conformer encoder (130.94 million parameters) allows for parallel execution of speech recognition and speaker verification, achieving an EER of 0.45%.

II. RELATED WORKS

A. Pretrained models for speaker verification

Several studies have explored the application of self-supervised pretrained Transformers for speaker verification tasks. Fan *et al.* [25], and Vaessen *et al.* [26] adopted a direct fine-tuning approach on the pretrained model by incorporating an additional pooling layer on top of the model’s output. However, this method did not surpass the performance of CNN- or TDNN-based speaker verification models, which typically have fewer parameters than the pretrained Transformer. Novoselov *et al.* [27] fine-tuned wav2vec 2.0 by integrating two simple TDNN layers and a statistic pooling layer. Their findings suggested that utilizing the entire deep pretrained encoder architecture was unnecessary, as earlier layers potentially provided more speaker information.

Another prevalent method replaces the handcrafted feature with the pretrained frame-level feature to train TDNN- or CNN-based speaker embedding networks [28], [29]. This approach, employing a layer-wise weighted average to aggregate features from different Transformer layers, has improved performance over models using handcrafted spectral features. However, this comes at the cost of using a large number of pretrained parameters alongside a full TDNN- or CNN-based speaker embedding network. Expanding on the concept of layer-wise weighted average as a feature aggregation method, Peng *et al.* [30] proposed multi-head factorized attentive pooling, which can be viewed as a fusion of layer-wise weighted average and multi-head attentive pooling.

In this paper, instead of self-supervised pretrained Transformers, an ASR-pretrained Conformer is used as the network backbone for the speaker embedding network since there are already many large-scale publicly open ASR datasets available. We directly apply to fine-tune the pretrained Conformer with a multi-scale feature aggregation module, eliminating the need for an additional TDNN- or CNN-based speaker network. This transfer learning strategy allows the knowledge learned from ASR to be effectively transferred to speaker verification tasks.

B. ASR guided speaker verification

ASR or phonetic information plays an essential role in speaker verification. In the statistical i-vector framework,

substituting a Gaussian mixture model (GMM) with an ASR-trained DNN to gather sufficient statistics for i-vector extraction results in significant performance improvement [23], [31]. Alternatively, some researchers utilize a tandem feature that merges spectral and ASR-derived features for GMM modeling [22], [32].

In the realm of deep learning, the integration of ASR and phoneme information into speaker verification is gaining increasing attention. Three main strategies have been investigated for such integration, each with merits and challenges.

The first strategy involves applying frame-level phonetic features from an ASR to a speaker verification model. In this context, Rahman *et al.* used bottleneck phonetic features from an ASR acoustic model to replace spectral features in speaker network training, indicating the potential of phonetic features to carry speaker-specific information [33]. In similar efforts, researchers have also incorporated phonetic features alongside spectral features for speaker modeling. Zheng *et al.* used separate network stems to model these two types of features [34], while Zhou *et al.* processed these features jointly by concatenating them [21]. Depending on the modeling stage, phonetic features can be incorporated at the input of the speaker network [34], [21] or before the pooling layer [21], [35]. These research indicate that incorporating auxiliary phoneme information benefits speaker modeling. Besides, Chen *et al.* proposed to model speaker characteristics in phoneme units, termed as phoneme-unit-specific network [36]. This method can be considered as modeling speaker characteristics using multi-phonetic-head attention, which has the attention weight of phoneme posterior probability.

The second strategy employs a multi-task learning approach, leveraging phoneme recognition as an auxiliary task alongside the primary task of speaker recognition. Studies have shown that frame-level phoneme modeling enhances speaker verification performance [35], [37], [38].

The last strategy involves employing phonetic information as a guided signal to be removed from speaker modeling. A study by Wang *et al.* suggested that adversarial training to remove phonetic information at the segment level can boost speaker verification performance [38]. In contrast, Tawara *et al.* found that removing phonetic information at the frame level is beneficial for extremely short utterances of less than 1.4 seconds [39]. Hong *et al.* introduced a self-constraint learning and reconstruction strategy that eliminates phonetic information in lower layers, thereby allowing subsequent layers to capture speaker-specific features more efficiently [40].

In our study, we extend the benefits of the second approach through knowledge distillation from the ASR Conformer to the speaker verification model. This method aligns with the recognized advantages of employing phoneme recognition as an auxiliary task, thus aiming to improve speaker verification performance.

C. Parameter-efficient transfer learning with adaptors

The concept of adaptors stems from the idea of fine-tuning large pre-trained models using lightweight neural modules, which can be considered a parameter-efficient transfer

learning technique [41]. This approach incorporates trainable lightweight neural modules into a large pre-trained model while keeping the pre-trained parameters frozen during fine-tuning. This technique has seen successful applications across various domains, including computer vision [42], natural language processing [41], [43], and machine translation [44].

While adaptors have been successful in different domains, their integration into speech-processing tasks presents multiple applications. For example, adaptors are applied to self-supervised pre-trained models for speech recognition [45]. In the context of multilingual ASR, language-specific adaptors have been used to adapt a pre-trained ASR model to various languages [46], [47]. In speech translation, adaptors enable a pre-trained model to specialize in specific language pairs [48]. Additionally, adaptors have been employed to connect an ASR encoder with a multilingual denoising auto-encoder for multilingual speech translation [48]. Other applications of adaptors include speaker verification [49], [50] and other speech processing tasks [50].

Most existing applications of adaptors focus on self-supervised pre-trained models for specific downstream tasks [41], [43], [49], [50]. Moreover, adaptors have been employed to perform domain adaptation for the same task, as seen in multilingual ASR [46], [47] and multilingual speech translation [48]. These methods usually incorporate adaptor modules within the network architecture, altering the output of the pre-trained model.

In contrast, our study motivated from the application of the adaptor mechanism. We apply a similar idea to transfer knowledge across different tasks: from ASR to speaker verification. We uniquely position an adaptation module on top of the original model, ensuring that the output of the ASR Conformer remains unchanged. This design enables the simultaneous execution of ASR and speaker verification tasks within a single Conformer model.

III. METHODS

Our research explores three distinct approaches for leveraging an ASR Conformer in speaker verification. First, we utilize a pre-trained ASR Conformer to initialize the speaker embedding network, which mitigates the risk of overfitting and enhances generalization in the speaker Conformer. Second, we employ knowledge distillation from the ASR Conformer to the speaker verification model. Lastly, we introduce an adaptation mechanism that unifies ASR and speaker verification tasks within a single Conformer model. The adaptation efficiently transforms features learned by the ASR to suit speaker verification tasks, all without altering the original ASR Conformer outputs. This section elaborates on these three methodologies, starting with the architecture of the Conformer encoder.

A. Conformer

Developed primarily for ASR tasks, the Conformer encoder is adept at modeling both local and global dependencies within speech signals [17]. It improves upon the Transformer encoder [16] by incorporating a CNN to capture local spectral feature




Fig. 1. Conformer encoder architecture (left) and a Conformer building block (right) [17].

information. The Conformer consists of a convolutional subsampling layer, which reduces the length of input sequences, and a series of Conformer blocks that transform the input signal into higher-level representations. Fig. 1 presents the Conformer encoder structure.

A Conformer block consists of two feed-forward networks (FFNs) flanked by a multi-head self-attention (MHSA) module and a convolution (Conv) module. In the Conformer, the MHSA employs relative sinusoidal positional encoding [51], allowing for efficient sequence handling at unseen lengths. The convolutional module features a point-wise convolution followed by a gated linear unit, succeeded by a one-dimensional depthwise convolution. Batch normalization and Swish activation are subsequently applied. The feed-forward network contains two linear layers separated by a nonlinear activation, with dropout applied after each linear transformation. As illustrated in Fig. 1, residual connections are used between the modules, while half-step residual connections are utilized within feed-forward modules, akin to a Macaron-Net [52]. Layer normalization is applied prior to the output. Mathematically, for a given input $\mathbf{h}_{i-1} \in \mathbb{R}^{d \times T}$, the output $\mathbf{h}_i \in \mathbb{R}^{d \times T}$ of the i -th Conformer block is represented as follows:

$$\begin{aligned}
 \mathbf{h}'_i &= \mathbf{h}_{i-1} + \frac{1}{2} \text{FFN}(\mathbf{h}_{i-1}) \\
 \mathbf{h}''_i &= \mathbf{h}'_i + \text{MHSA}(\mathbf{h}'_i) \\
 \mathbf{h}'''_i &= \mathbf{h}''_i + \text{Conv}(\mathbf{h}''_i) \\
 \mathbf{h}_i &= \text{LayerNorm}(\mathbf{h}'''_i + \frac{1}{2} \text{FFN}(\mathbf{h}'''_i))
 \end{aligned} \tag{1}$$

where d denotes the dimension of the input and the output sequences, and T represents the length of the time sequence.

B. MFA-Conformer for speaker verification

Multi-scale feature aggregation (MFA) is a technique that concatenates output feature maps from all frame-level modeling modules in a speaker embedding network before utterance-level pooling. This approach has been shown to improve performance for TDNN-based networks, suggesting that lower-level features can contribute useful speaker information [6].

To apply the Conformer encoder in the speaker verification task, MFA-Conformer proposed to integrate an MFA module into the Conformer encoder [18]. Specifically, this MFA module concatenates the frame-level outputs from all Conformer blocks prior to the pooling layer:

$$\begin{aligned}
 \mathbf{H}' &= \text{Concat}(\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_L) \\
 \mathbf{H} &= \text{LayerNorm}(\mathbf{H}')
 \end{aligned} \tag{2}$$

where L is the number of Conformer blocks in the Conformer encoder, and $\mathbf{H}, \mathbf{H}' \in \mathbb{R}^{D \times T}$ with $D = L \times d$.

With this concatenated frame-level feature map \mathbf{H} , attentive statistics pooling is applied to produce an utterance-level representation [53]. Finally, the speaker embedding is extracted by applying batch normalization and a fully-connected layer to this utterance-level representation. During training, an additional fully-connected layer is applied to classify speakers in the training set from speaker embeddings.

C. Transfer learning with the ASR pretrained Conformer

While deeper Transformers are known to yield superior results as more training data become available [29], [54], training these models from scratch often requires large datasets [55]. Further, research indicates that increasing the number of layers in Conformer architectures can result in a performance drop in speaker verification tasks, suggesting potential issues of overfitting [18].

To mitigate the risks of overfitting, we employ an ASR pretrained Conformer to initialize the MFA-Conformer-based speaker embedding network. The pretraining on ASR tasks affords several advantages, such as faster convergence and enhanced generalization capabilities in the speaker verification domain.

In our approach, the parameters of the ASR pretrained Conformer encoder are used to initialize the MFA-Conformer speaker embedding network. During the early training phases, we keep these encoder parameters frozen and allow only the pooling and subsequent linear layers to be updated for a few epochs. In later stages, we proceed to fine-tune the parameters across the entire MFA-Conformer architecture to better align it with the specific needs of speaker verification. By limiting updates to the pooling and linear layers initially, these layers are tailored to adapt the frame-level feature maps derived from the ASR model to the speaker verification objective. This structured training approach ensures that the pretrained Conformer transitions smoothly to the speaker verification objective without being significantly disrupted by the random initialization of these layers.

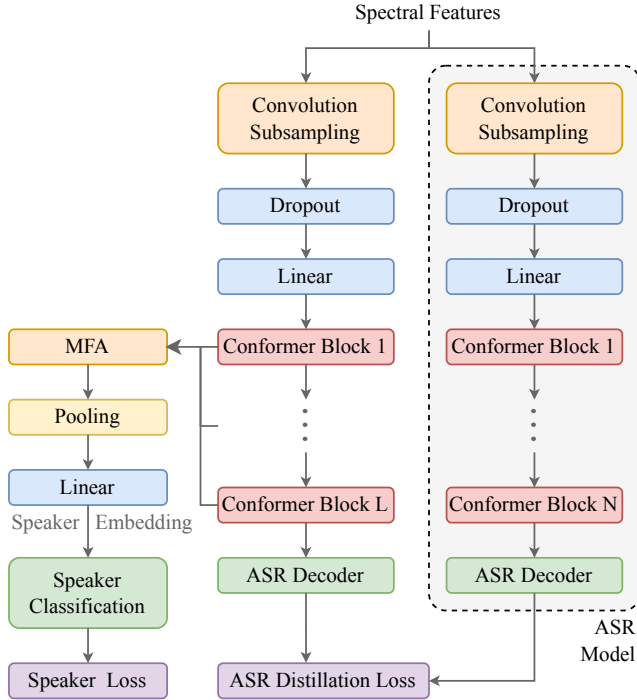


Fig. 2. Knowledge distillation from a pretrained ASR Conformer model to a MFA-Conformer-based speaker verification model.

D. Knowledge distillation from ASR to speaker verification

Knowledge distillation involves training a “student” model to reproduce the behavior of a more complex “teacher” model [24]. In our setting, an ASR pretrained Conformer acts as the teacher model, guiding the learning process of the MFA-Conformer-based speaker verification model, which serves as the student.

Given a speaker recognition dataset \mathcal{D} , the objective of a speaker verification model is to minimize the difference between its predictions and the ground-truth speaker labels. The loss function L_{spk} can be expressed as:

$$L_{\text{spk}} = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\ell_{\text{spk}}(f(\mathbf{x}), y)] \quad (3)$$

Here, $f(\cdot)$ is the MFA-Conformer speaker verification model, $f(\mathbf{x})$ is the Conformer’s prediction for the input spectral sequence \mathbf{x} , and y is the speaker label. The speaker classification loss ℓ_{spk} commonly adopts a cross-entropy format or an angular-softmax variant [7].

For distillation, the speaker MFA-Conformer student is trained to align its outputs with the ASR teacher model, as described in the loss L_{distill} :

$$L_{\text{distill}} = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\ell_{\text{distill}}(f_{\text{student}}(\mathbf{x}), f_{\text{teacher}}(\mathbf{x}))] \quad (4)$$

In this setting, $f_{\text{student}}(\cdot)$ refers to the MFA-Conformer coupled with an ASR decoder, while $f_{\text{teacher}}(\cdot)$ is the ASR model. In the distillation process, the loss function L_{distill} is formulated based on the Kullback-Leibler (KL) divergence, which quantify the divergence between the student and teacher frame-level logits outputs.




Fig. 3. Linear probe accuracy across Conformer layers for speaker identification.

The ultimate training objective combines both the speaker classification and the distillation losses:

$$L = L_{\text{spk}} + \alpha L_{\text{distill}} \quad (5)$$

where α is a hyperparameter determining the strength of the distillation effect. Fig. 2 illustrates the knowledge distillation process from ASR to speaker verification.

Our approach harnesses the strengths of both knowledge distillation and multi-task learning, offering advantages for speaker verification. Firstly, it enables the speaker verification model to utilize robust feature representations from an ASR-pretrained model, enhancing performance without extensive ASR data. This method, diverging from traditional knowledge distillation, incorporates the ASR model’s outputs as an auxiliary objective, enriching phonetic feature learning in a multi-task framework. Secondly, this synergy improves speaker discrimination by leveraging nuanced phonetic information. Lastly, our method with knowledge distillation offers more architectural flexibility, allowing for optimized designs that can cater to the specific requirements of both ASR and speaker verification tasks.

E. Speaker adaptation module: unifying ASR and speaker verification

To leverage the versatility of Conformer encoders across multiple tasks, this section explores the possibility of crafting a unified model that serves both ASR and speaker verification objectives.

1) *Inherent speaker-specific information in ASR Conformers*: Conformer encoders, originally tailored for ASR, possess innate adaptability. This flexibility is attributed to their multi-layered structure, capturing a hierarchical abstraction of speech signals. Essentially, the lower layers of the ASR Conformer capture diverse attributes of speech, such as speaker characteristics, linguistic patterns, emotional tones, and phonetic variations. In contrast, the upper layers prioritize phonetic and contextual specifics, driven by the ASR objectives.

To empirically validate this layer-wise specialization, we employed a linear probe to measure the speaker-specific information within different layers of a pretrained ASR Conformer

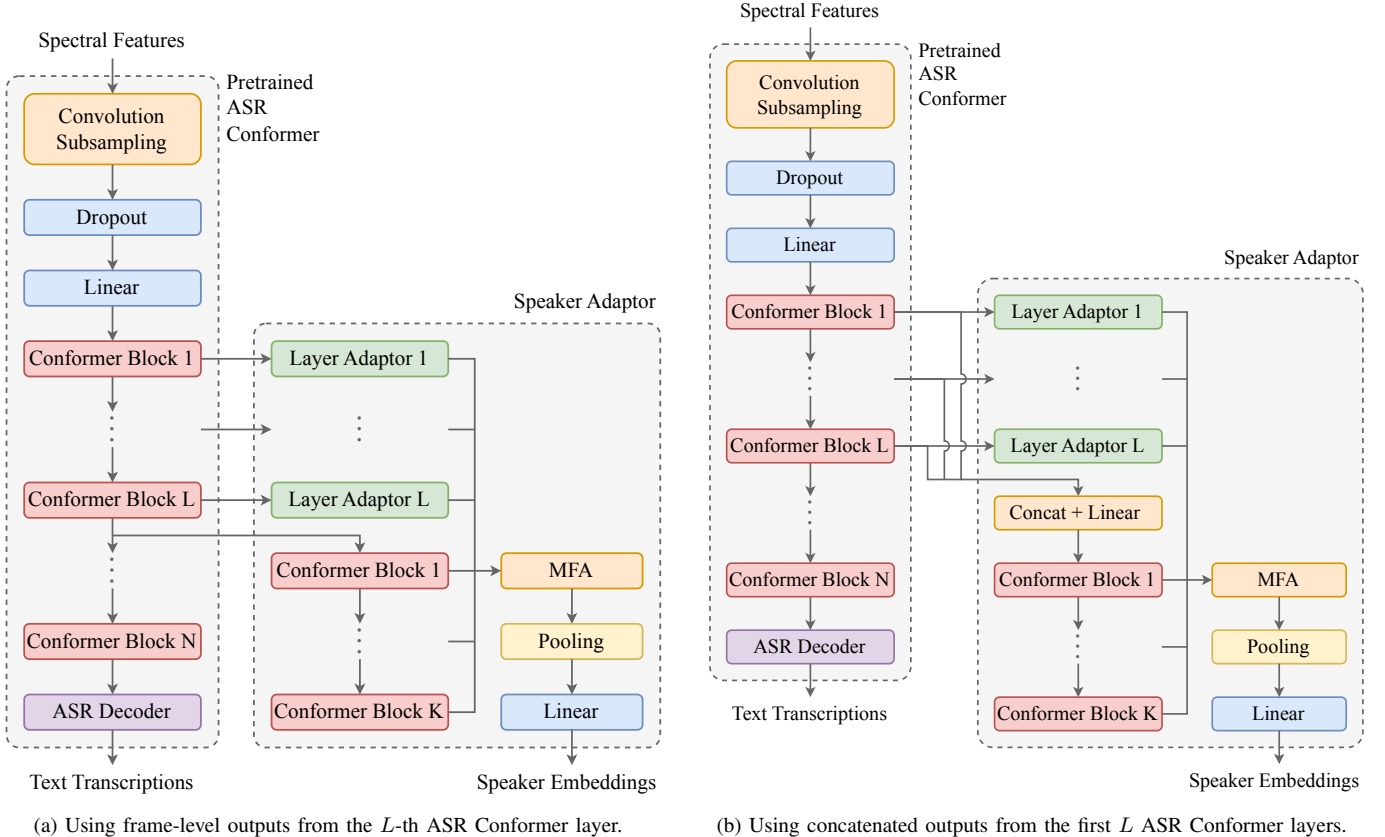


Fig. 4. Design variants of the proposed speaker adaptation module to unify ASR and speaker verification in one Conformer model.

encoder. A detailed description of the models used for this probing is provided later in section IV-C. Each Conformer layer’s output was first subjected to two linear fully-connected layers, followed by average pooling to derive speaker embeddings. These embeddings are further processed by an additional linear layer to perform speaker classification on the VoxCeleb 1 development set [56]. The results, illustrated in Fig. 3, confirm that lower layers inherently possess rich speaker-specific information. As we progress toward the upper layers, the specificity of the ASR task intensifies, diluting the speaker-specific traits.

2) *Motivation for a unified Conformer model:* The layer-wise investigation into Conformer encoders revealed an intriguing fact: despite being primarily trained for ASR, even the initial layers possess striking proficiency in speaker recognition. Remarkably, the fifth layer of a large pretrained ASR Conformer displayed an impressive training accuracy of 99.65% for speaker recognition, suggesting that ASR-trained features can effectively be used for speaker verification. This compelling evidence motivates our pursuit of a unified Conformer model that seamlessly transitions between ASR and speaker verification tasks.

3) *Speaker adaptation module:* To bridge the gap between ASR and speaker verification and unify the Conformer encoder, we introduce the speaker adaptation method. Conceptually, the speaker adaptation module is a lightweight trainable module integrated into a large-scale pretrained model [41]. Our

design operates on the intermediate representations, leaving the pretrained model’s output unchanged.

Fig. 4 visualizes the design of our proposed speaker adaptation module. It consists of three parts: L layer adaptors, K trainable Conformer layers, and a combination of a pooling layer and a subsequent fully connected layer for speaker embedding derivation.

a) *Layer adaptors:* These components work on fine-tuning the outputs from each layer of the pretrained ASR Conformer model, aligning them more closely with the objectives of speaker verification. Specifically, for a pretrained ASR Conformer, the frame-level output from the i -th Conformer layer, denoted as $\mathbf{h}_i \in \mathbb{R}^{d \times T}$, is transformed by the layer adaptor \mathbf{A}_i :

$$\mathbf{h}'_i = \mathbf{A}_i(\mathbf{h}_i) \quad (6)$$

Our layer adaptors consist of two linear layers interleaved with layer normalization and an activation function. Given our observation that deeper layers retain less speaker-centric information, these adaptors are applied only to the first L layers of the pretrained ASR Conformer.

b) *Trainable Conformer layers:* To enhance speaker feature extraction, we incorporate K additional lightweight, trainable Conformer layers within the speaker adaptation module. Inputs to these layers come from one of the two following distinct options:

- Frame-level outputs from the L -th Conformer layer of the ASR model, as illustrated in Fig. 4a.

TABLE I
THREE ASR CONFORMER ENCODERS OF DIFFERENT SIZES

Model	layers	dim	heads	hidden units	parameters
Small ¹	16	176	4	704	15.88M
Medium ²	18	256	4	1024	35.26M
Large ³	18	512	8	2048	130.94M

- Concatenated outputs from the first L layers of the pretrained ASR Conformer encoder, with a linear layer to reduce the feature dimension, as illustrated in Fig. 4b.

To maintain the efficiency of the speaker adaptation module, these trainable Conformer layers are designed to be lightweight, with reduced hyper-parameters of dimensions and hidden units.

c) Speaker embedding extraction: After the transformations brought by the layer adaptors and the trainable Conformer layers, the frame-level features are fed into the MFA module:

$$\begin{aligned} \mathbf{H}' &= \text{Concat}[\mathbf{A}_1(\mathbf{h}_1), \dots, \mathbf{A}_L(\mathbf{h}_L), \tilde{\mathbf{h}}_1, \dots, \tilde{\mathbf{h}}_K] \\ \mathbf{H} &= \text{LayerNorm}(\mathbf{H}') \end{aligned} \quad (7)$$

Here, $\tilde{\mathbf{h}}_k$ denotes the output from the k -th trainable Conformer layer. K represents the number of these layers. By design, K can be zero, indicating the absence of any new trainable Conformer layers. With these concatenated frame-level representations derived from the pretrained ASR Conformer encoder, a standard speaker verification procedure with an utterance-level pooling layer and a subsequent linear layer is used for speaker embedding extraction.

During the training phase, the pretrained ASR Conformer is kept frozen. Only speaker adaptation module components, including layer adaptors, lightweight Conformer layers, pooling, and the following linear layers, are trained under the speaker verification objective.

IV. EXPERIMENTAL SETUPS

A. Dataset

The experiments are conducted on VoxCeleb [56], [57]. For model training, we opted to employ the development set from VoxCeleb 2. This training dataset encompasses 1,092,009 audio recordings from a diverse set of 5,994 distinct speakers.

For the evaluation phase, we use both the development and test sets from VoxCeleb 1. We present the speaker verification performances based on three predefined trial lists as described in [57]:

- VoxCeleb 1-O: This represents the original trial list associated with VoxCeleb 1, encompassing 37,720 trials derived from 40 speakers.
- VoxCeleb 1-E: An expanded trial list that comprises 581,480 trials sourced from 1,251 speakers.
- VoxCeleb 1-H: A more challenging trial list with 552,536 trials from 1,190 speakers. All test pairings within this list share the same linguistic background and gender.

B. Data Augmentation

To enhance the robustness and versatility of our model, we integrated various data augmentation methodologies. First, we apply speed perturbation to the audio samples by accelerating or decelerating the content by factors of 1.1 and 0.9, respectively [58], [59]. As a result, this approach produced two supplementary replicas of each original audio, expanding the entire training dataset to include 17,982 distinct speakers and 3,276,027 unique utterances.

For the enlarged training dataset, two primary strategies were utilized:

- Additive noise augmentation: The MUSAN dataset [60] served as our noise source, enabling us to add ambient noise, musical sounds, and babble noise onto our audio files. The babble noise was generated by merging between three to eight separate speech files in the MUSAN dataset. The signal-to-noise ratios (SNR) range from 0 to 20 dB.
- Convolutional reverberation noise augmentation: We employed the collection of 40,000 simulated room impulse responses (RIR) from the study in [61]. Only simulated RIRs originating from small to medium-sized rooms are used.

To maintain variability during training epochs, we integrated on-the-fly data augmentation, applying the aforementioned noise augmentations with a likelihood of 0.6 for each training speech.

C. Pretrained ASR Conformer

We utilize pretrained ASR models from the NEMO toolkit [62]. The choice of the Conformer model from NEMO was driven by its performance and generalization capabilities, as demonstrated in various benchmarks. This ASR Conformer adopts the same encoder architecture as illustrated in [17] but uses a linear decoder and the connectionist temporal classification (CTC) for decoding.

In our experiments, we use three sizes of the NEMO ASR Conformer: small, medium, and large. Despite variations in size, each of these models shares a convolution subsampling rate of $\frac{1}{4}$, along with a consistent kernel size of 31 for their convolution modules. Table I shows the differences in Conformer layer numbers, encoder dimensions, attention heads, and linear hidden units across the three Conformer encoders.

According to the NEMO toolkit documentation, each Conformer-CTC model is trained on English corpora collated from 10 distinct datasets.⁴ In total, this collection spans approximately 10,000 hours of speech data.⁵

¹https://catalog.ngc.nvidia.com/orgs/nvidia/teams/nemo/models/stt_en_conformer_ctc_small

²https://catalog.ngc.nvidia.com/orgs/nvidia/teams/nemo/models/stt_en_conformer_ctc_medium

³https://catalog.ngc.nvidia.com/orgs/nvidia/teams/nemo/models/stt_en_conformer_ctc_large

⁴These datasets include Librispeech, Fisher Corpus, Switchboard-1, WSJ-0 and WSJ-1, National Speech Corpus (Part 1, Part 6), VCTK, VoxPopuli (EN), Europarl-ASR (EN), Multilingual Librispeech (MLS EN 2,000 hours subset), and Mozilla Common Voice (v7.0).

⁵This estimate is derived from the training data descriptions provided at the mentioned link in the previous footnote.

TABLE II
SPEAKER VERIFICATION PERFORMANCE OF ASR PRETRAINED MFA-CONFORMER ON VOXCETLEB 1.

Model	Size	Pretrained	VoxCeleb 1-O		VoxCeleb 1-E		VoxCeleb 1-H	
			EER[%]	minDCF	EER[%]	minDCF	EER[%]	minDCF
ECAPA-TDNN [11]	46.6M	×	0.68	0.0753	0.91	0.1006	1.72	0.1695
HuBERT Large [28]	316.61M+	✓	0.72	-	0.70	-	1.32	-
Wav2Vec2.0 Large (XLSR) [28]	317.38M+	✓	0.73	-	0.68	-	1.23	-
UniSpeech-SAT Large [28]	316.61M+	✓	0.63	-	0.63	-	1.29	-
WavLM Large + QMF [29]	316.62M+	✓	0.38	-	0.48	-	0.99	-
NEMO Small	15.88M	×	0.88	0.1367	1.08	0.1342	2.20	0.2245
NEMO Medium	35.26M	×	0.94	0.1200	1.26	0.1487	2.41	0.2398
NEMO Large	130.94M	×	0.96	0.1375	1.22	0.1391	2.35	0.2278
NEMO Large first 4 layers	35.02M	×	0.86	0.1051	1.03	0.1188	1.97	0.1920
NEMO Large first 6 layers	48.72M	×	0.80	0.1101	1.04	0.1202	2.04	0.2012
NEMO Large first 8 layers	62.42M	×	0.81	0.1121	1.00	0.1183	1.93	0.1904
NEMO Small	15.88M	✓	0.74	0.1101	0.90	0.1054	1.90	0.1893
NEMO Medium	35.26M	✓	0.61	0.0946	0.78	0.0891	1.67	0.1649
NEMO Large	130.94M	✓	0.48	0.0673	0.71	0.0785	1.54	0.1538
+ QMF			0.43	0.0623	0.66	0.0709	1.35	0.1350
NEMO Large first 4 layers	35.02M	✓	0.77	0.1065	1.04	0.1159	1.95	0.1862
NEMO Large first 6 layers	48.72M	✓	0.58	0.0618	0.84	0.0937	1.62	0.1571
NEMO Large first 8 layers	62.42M	✓	0.64	0.0982	0.86	0.0944	1.77	0.1732

D. Implementation details

Speech utterances are cropped to 2 seconds for training the speaker embedding network. We use a logarithmic Mel-spectrogram with 80 frequency bins as the acoustic feature, computed over Hamming windows of 20ms with a 10ms shift.

During training, the Additive angular margin (AAM) loss [7] is employed with a re-scaling factor of 32 and an angular margin of 0.2 to learn discriminative representations. The speaker embedding dimension is set to 256. We utilize the AdamW optimizer, beginning with a learning rate of 0.001. Additionally, we implement a cosine annealing learning rate scheduler, incorporating a warm-up phase spanning one training epoch. Our chosen batch size is 512, with a weight decay of 10^{-7} .

After convergence, we employ large margin fine-tuning (LMFT) [11]. Speech segments are expanded to 6 seconds, and the angular margin in the AAM loss is increased to 0.5. We turn off speed perturbation data augmentation, reverting the training data to its original set.

E. Evaluation

To generate speaker verification scores, we apply the adapted score normalization [63] after cosine similarity on two given speaker embeddings. In adapted score normalization, we utilize an imposter cohort randomly chosen from 30,000 training utterances, with an adapted cohort size of 700.

Although our standard procedure involves only this score normalization, we further calibrate the verification scores using the Quality Measure Function (QMF) [64], [11] for specific systems as per their requirements. The calibration model is trained on 30,000 trials generated from the VoxCeleb 2 development set. This model incorporates several quality metrics including the duration and SNR of the enrollment and testing utterances, the magnitudes of the embeddings, and the verification score itself.

We evaluate speaker verification performance using two metrics: (1) Equal Error Rate (EER): This denotes the error rate at the point where the false acceptance rate equals the false rejection rate. (2) Minimum Detection Cost (minDCF): This represents the minimal value of a detection cost function. The function is a weighted sum of false-reject and false-alarm error rates for a given decision threshold [65]. The parameters for this function are set as follows: $C_{\text{Miss}} = 1$, $C_{\text{FA}} = 1$, and $P_{\text{Target}} = 0.01$.

V. EXPERIMENTAL RESULTS

A. Transfer learning with the ASR pretrained Conformer

In this subsection, we present speaker verification results using our first proposed method. Specifically, we explore the efficacy of initializing the MFA-Conformer speaker verification model with a pretrained ASR Conformer. The performance of various MFA-Conformer speaker embedding networks, both with and without ASR pretraining, are detailed in Table II.

1) *MFA-Conformer’s performance without ASR pretraining*: We first analyze the performance of the MFA-Conformer model without integrating ASR pretraining. The results indicate that increasing the trainable parameters does not yield improved speaker verification performance. Specifically, upon increasing model parameters by a factor of eight (from 15.88 million to 130.94 million), the EERs observe a decrease ranging from 7% to 13% across the three testing trials. This suggests that MFA-Conformers tend to overfit, especially in scenarios with limited data availability.

2) *MFA-Conformer’s performance with ASR pretraining*: Integrating ASR pretraining into the MFA-Conformer model leads to significant improvements across all evaluated model sizes. For example, the small MFA-Conformer with ASR pretraining recorded a relative reduction in EER of 15.9% on the VoxCeleb 1-O trials compared to its non-pretrained counterpart. This relative reduction was even more significant

TABLE III
COMPARISON OF ASR PRETRAINING METHOD AND SSL AS FRONT-END MODULE METHOD. PERFORMANCE ARE REPORTED ON EER (%).

Model	Pretrained Model			Speaker Model	LMFT	QMF	Vox1-O	Vox1-E	Vox1-H
	Size	Training Data	Usage						
HuBERT Base [29]	94.7M	960 hr	front-end module	ECAPA-TDNN	×	×	0.989	1.068	2.216
HuBERT Large [29]	316.6M	60k hr	front-end module	ECAPA-TDNN	×	×	0.808	0.822	1.678
HuBERT Large [29]	316.6M	60k hr	front-end module	ECAPA-TDNN	✓	✓	0.585	0.654	1.342
WavLM Base+ [29]	94.7M	94k hr	front-end module	ECAPA-TDNN	×	×	0.84	0.928	1.758
WavLM Large [29]	316.6M	94k hr	front-end module	ECAPA-TDNN	×	×	0.617	0.662	1.318
WavLM Large [29]	316.6M	94k hr	front-end module	ECAPA-TDNN	✓	✓	0.383	0.480	0.986
Conformer Medium	35.3M	10k hr	parameter initialization	pretrained Conformer	×	×	0.78	0.97	2.04
Conformer Medium	35.3M	10k hr	parameter initialization	pretrained Conformer	✓	×	0.61	0.78	1.67
Conformer Medium	35.3M	10k hr	parameter initialization	pretrained Conformer	✓	✓	0.52	0.72	1.48
Conformer Large	130.9M	10k hr	parameter initialization	pretrained Conformer	×	×	0.74	0.91	1.91
Conformer Large	130.9M	10k hr	parameter initialization	pretrained Conformer	✓	×	0.48	0.71	1.54
Conformer Large	130.9M	10k hr	parameter initialization	pretrained Conformer	✓	✓	0.43	0.66	1.35

for larger models, with the large MFA-Conformer recording a 50% reduction on the same trail. These results confirm the benefits of leveraging ASR pretraining with 10k hours of speech data for speaker verification models, particularly for larger Conformer models, where the risk of overfitting is higher.

3) *Benchmarking against large self-supervised speech models*: Large self-supervised speech models for speaker verification are used as feature extractors to replace the hand-crafted feature with an additional speaker embedding model append. Compared to larger self-supervised pretrained models with more than 300 million parameters (HuBERT Large, Wav2Vec2.0 Large, UniSpeech-SAT Large), the ASR pretrained MFA-Conformers achieve comparable or even better verification performance on VoxCeleb 1-O trials. For instance, while the UniSpeech-SAT large model (with 316.62 million parameters) achieved an EER of 0.63% on VoxCeleb 1-O trials, the large ASR pretrained MFA-Conformer (with 130.94 million parameters) recorded an EER of 0.48%. Such results emphasize the efficiency of ASR pretraining on the speaker MFA-Conformer model.

However, MFA-Conformers do not outperform large self-supervised models on the VoxCeleb 1-E and VoxCeleb 1-H trials. A plausible reason is the difference in the volume of training data used for pretraining. While self-supervised models utilized speech data ranging from 56,000 to 188,000 hours, the training data of the ASR Conformer used in this study are limited to approximately 10,000 hours. Nevertheless, our proposed ASR pretraining method offers flexibility. Integrating an MFA module and a pooling layer can readily transform an ASR pretrained Conformer into a speaker verification task. This eliminates the need for supplementary TDNN- or CNN-based speaker networks, which are commonly employed in large self-supervised models.

To facilitate a direct comparison between the ASR pretrained method and the large self-supervised speech model method, table III highlights various configurations, including different model sizes, training data, usage types, and additional fine-tuning techniques like Large Margin Fine-Tuning (LMFT) and Quality Measure Function (QMF). From the table, the

medium Conformer model with ASR pretraining demonstrates comparable performance to the WavLM Base+ and HuBERT Base models. Similarly, the large ASR pretrained Conformer model exhibits performance on par with the HuBERT Large and WavLM Large models using a smaller size of model and training data, making it a competitive option in the realm of speech model methods.

4) *Exploring the potential of extracting lower layers*:

We also extend our experiments by using subsets of the larger Conformer model, specifically extracting the initial 4, 6, and 8 layers, to initiate MFA-Conformer training. These truncated models perform better than the full version when ASR pretraining was not applied, which reaffirms the earlier observation regarding the overfitting tendency of Conformers with increased parameters. When ASR pretraining is applied, these truncated models outperform their counterparts without ASR pretraining, emphasizing the benefits of ASR pretraining. The experiments of the truncated Conformers present a way to balance model size and speaker verification performance.

B. Knowledge distillation from ASR to speaker verification

This section presents the results of our second proposal, which explores the application of knowledge distillation from ASR to speaker verification. For these experiments, we used the NEMO Large ASR-CTC model in Table I, as the teacher model in the knowledge distillation process. We set the hyperparameter α in equation 5 to 1. The speaker verification performance of various MFA-Conformer models, considering different training methodologies and model sizes, are shown in Table IV.

1) *Influence of ASR knowledge distillation*: The primary objective of our experiments is to determine the effectiveness of ASR distillation in enhancing the performance of MFA-Conformer models. The application of ASR distillation consistently shows promising improvements across various model scales and sampling rates. For instance, considering the NEMO Small model, the ASR distillation technique (EER of 0.54%) reduces the EER by 38.6% on the VoxCeleb 1-O trials compared to the vanilla version (EER of 0.88%). The NEMO Medium model with ASR distillation outperforms its vanilla

TABLE IV
SPEAKER VERIFICATION PERFORMANCE OF MFA-CONFORMER WITH ASR DISTILLATION ON VOXCETLEB 1.

Model	Sampling Rate	Size	MACs ⁶	Training Method	VoxCeleb 1-O		VoxCeleb 1-E		VoxCeleb 1-H	
					EER[%]	minDCF	EER[%]	minDCF	EER[%]	minDCF
NEMO Half Small	$\frac{1}{2}$	8.73M	405.18M	Baseline	0.62	0.0792	0.84	0.0907	1.67	0.1676
				ASR Distillation	0.65	0.0725	0.79	0.0881	1.50	0.1477
				+ QMF	0.56	0.0572	0.74	0.0775	1.36	0.1333
NEMO Small	$\frac{1}{4}$	15.88M	1.12G	Baseline	0.88	0.1367	1.08	0.1342	2.20	0.2245
				ASR Pretrained	0.74	0.1101	0.90	0.1054	1.90	0.1893
				+ QMF	0.61	0.0937	0.83	0.0954	1.69	0.1687
				ASR Distillation	0.54	0.0625	0.74	0.0782	1.54	0.1568
+ QMF	0.43	0.0575	0.67	0.0705	1.37	0.1429				
NEMO Half Medium	$\frac{1}{2}$	19.30M	803.04M	Baseline	0.64	0.0855	0.89	0.1020	1.74	0.1750
				ASR Distillation	0.43	0.0485	0.69	0.0727	1.37	0.1364
				+ QMF	0.38	0.0388	0.66	0.0668	1.24	0.1221
NEMO Medium	$\frac{1}{4}$	35.26M	2.31G	Baseline	0.94	0.1200	1.26	0.1487	2.41	0.2398
				ASR Pretrained	0.61	0.0946	0.78	0.0891	1.67	0.1649
				+ QMF	0.52	0.0875	0.72	0.0783	1.48	0.1538
				ASR Distillation	0.52	0.0689	0.72	0.0791	1.49	0.1429
+ QMF	0.48	0.0589	0.67	0.0711	1.34	0.1364				
NEMO Half Large	$\frac{1}{2}$	72.16M	2.52G	Baseline	0.87	0.0799	1.04	0.1145	1.93	0.1838
				ASR Distillation	0.52	0.0564	0.75	0.0808	1.55	0.1516
				+ QMF	0.48	0.0619	0.72	0.0735	1.42	0.1439
NEMO Large	$\frac{1}{4}$	130.94M	8.53G	Baseline	0.96	0.1375	1.22	0.1391	2.35	0.2278
				ASR Pretrained	0.48	0.0673	0.71	0.0785	1.54	0.1538
				+ QMF	0.43	0.0623	0.66	0.0709	1.35	0.1350
				ASR Distillation	0.53	0.0589	0.79	0.0852	1.64	0.1611
+ QMF	0.45	0.0562	0.75	0.0802	1.49	0.1475				

counterpart by approximately 44.7% relatively in EER on the same trials.

Our results also enable a direct comparison between the ASR distillation and ASR pretraining techniques. Notably, in most cases, models trained with ASR distillation outperform or come close to their ASR pretrained counterparts. For instance, the EER in the VoxCeleb 1-O trial for the NEMO Medium model decreases by 14.8% with ASR distillation compared to ASR pretraining. The improvements from ASR distillation primarily come from two factors. First, the student model benefits from the robustness of the larger ASR teacher model trained on extensive ASR datasets, exposing the student model to a wide range of speech patterns and accents. Second, the auxiliary task of ASR at frame-level modeling enhances the student model’s ability to capture fine-grained, speaker-specific features, which is critical for speaker verification.

However, the NEMO Large model with ASR distillation does not consistently outperform the ASR pretraining method. This might be due to the shared model architecture between the student and teacher models, as the ASR-pretrained NEMO Large model was used as the teacher. This outcome suggests no one-size-fits-all answer, and the best approach could depend on the specific model architecture or data constraints.

2) *Reduced Conformer layers with increased convolution subsampling rate*: To explore the impact of model size and sampling rate, we reduced the number of Conformer layers by half and increased the convolution subsampling rate from $\frac{1}{4}$ to $\frac{1}{2}$ for the three Conformer models. The ASR teacher model remained the same as in previous experiments. To match the convolution subsampling rate between the teacher and student

models for the KL divergence loss at frame level, we added a convolutional layer to the student model with a kernel size of 3, padding of 1, and stride of 2, increasing the student model’s convolution subsampling rate from $\frac{1}{4}$ to $\frac{1}{2}$.

Our results show that MFA-Conformer models with a $\frac{1}{2}$ convolution subsampling rate, even with nearly half the number of parameters, achieve comparable or better verification performance with ASR distillation compared to those with a $\frac{1}{4}$ convolution subsampling rate. For example, the NEMO Half Medium model with ASR distillation achieved EERs of 0.43%, 0.69%, and 1.37%, while the NEMO Medium model’s EERs were 0.52%, 0.72%, and 1.49% for VoxCeleb 1-O, VoxCeleb 1-E, and VoxCeleb 1-H trials, respectively.

The integration of ASR distillation into the MFA-Conformer model training presents a promising direction in speaker verification. Our results demonstrate consistent improvements across different model scales, indicating the robustness and versatility of this method. Moreover, the potential to achieve similar or even better results than ASR pretraining further highlights the efficacy of ASR distillation.

C. Speaker adaptation: unifying ASR and speaker verification

In this section, we delve into the effectiveness of our proposed speaker adaptation approach in bridging the gap between ASR and speaker verification tasks.

⁶MACs (Multiply-Accumulate Operations) are calculated based on a 5-second speech input.

⁷#ASR param indicates the model size (in million parameters) of the ASR Conformer encoder when it has L layers. #adap param represents the speaker adaptation modules’ total model size including L adaptor layers (in million parameters).

TABLE V
THE NETWORK CONFIGURATION OF THE SPEAKER ADAPTATION MODULE.

Layer	Structure
Layer Adaptor	$\begin{bmatrix} \text{Linear}(D, 128) \\ \text{LayerNorm}(128) \\ \text{ReLU}() \\ \text{Linear}(128, 128) \end{bmatrix} \times L$
Trainable Conformer	V2: $\text{Linear}(D, 176)$ V3: $\begin{bmatrix} \text{Concatenation} \\ \text{Linear}(D \times L, 176) \end{bmatrix}$ Conformer(dim = 176, head = 4, hidden = 704) $\times K$
MFA	Concatenation LayerNorm($128 \times L + 176 \times K$)
Pooling	Attentive statistics pooling
Linear	Linear($(128 \times L + 176 \times K) \times 2, 256$)

We employ three pretrained ASR Conformer encoders — small, medium, and large, as referenced in Table I. These encoders are integrated with our speaker adaptation technique. For each encoder size, we assess three distinct configurations of the speaker adaptation module, based on the one depicted in Fig. 4:

- **V1:** This version extracts directly from the first L layers of the ASR Conformer without the intervention of layer adaptors.
- **V2:** In alignment with Fig. 4a, this version integrates layer adaptors to refine the outputs from the first L ASR Conformer layers. Subsequently, K lightweight Conformer layers process the frame-level outputs derived from the L -th ASR Conformer layer.
- **V3:** As illustrated in Fig. 4b, this configuration feeds the K lightweight Conformer layers with a concatenated output from the first L Conformer layers of the pretrained ASR model. An auxiliary linear layer ensures the alignment of concatenated feature dimensions.

All configurations use a lightweight Conformer layer architecture consistent across the ASR encoders. These lightweight Conformer layers have 174 dimensions, 704 hidden units, and 4 attention heads, the same as the Conformer layer configuration in the NEMO Small ASR-CTC model. Additionally, the layer adaptor always maps the frame-level outputs from ASR Conformer layers to a 128-dimensional feature space. A detailed architecture configuration can be found in Table V. Notably, when the model lacks trainable Conformer layers (i.e., $K = 0$), the V2 and V3 configurations converge to become identical. The specific EERs for distinct model configurations, considering variations in both L and K , are outlined in Tables VII, VIII, and IX, each corresponding to a unique pretrained ASR model.

1) *Baseline - ASR Conformer without speaker adaptation:* Before introducing any adaptation method, it is crucial to understand the innate capabilities of the ASR Conformer encoder when used for speaker verification. Our baseline is free from any layer adaptor (configuration V1) and does not incorporate additional trainable Conformer layers ($K = 0$). Here, the frame-level outputs of the ASR Conformer are concatenated and subsequently routed to the pooling layer to

extract speaker embeddings. The results consistently indicate a notable trend: ASR models with a more significant number of parameters (or layers) often exhibit superior performance compared to their smaller counterparts. For instance, while the NEMO Small ASR-CTC model with 12 layers has an EER of 1.73%, its larger counterpart, the NEMO Large ASR-CTC model with 6 layers, surpasses it with a more desirable EER of 1.18%. While increasing the ASR Conformer’s layers generally leads to a decrease in the EER, the relationship is not strictly linear. For instance, in NEMO Large ASR-CTC mode, while moving from 6 to 10 layers results in an EER reduction from 1.18% to 0.97%, further increasing to 14 layers sees a slight EER increase to 1.01%.

2) *ASR Conformer with layer adaptors:* After assessing the ASR Conformer without speaker adaptation, we investigated the effect of introducing layer adaptors (configuration V2) without integrating additional trainable Conformer layers ($K = 0$). Using the layer adaptor, the ASR Conformer’s feature dimensions are reduced to 128, resulting in a smaller concatenated feature dimension after MFA concatenation. This led to a more compact speaker adaptation module in V2 compared to V1. Our findings indicate that introducing layer adaptors substantially enhances the speaker verification performance. Specifically, for the NEMO Small ASR-CTC model with $L = 12$, we observed an EER of 1.10%, marking a relative 36% reduction from the baseline’s 1.73% in the absence of speaker adaptation. Similar performance improvements are also witnessed across medium and large ASR-CTC models. The consistent performance improvement across different model sizes proves the effectiveness of layer adaptors.

3) *ASR Conformer with trainable lightweight Conformer layers:* Expanding our investigation, we delved into the impact of incorporating trainable lightweight Conformer layers into the ASR Conformer under configuration V1, specifically with $K = 2$ and $K = 4$. Adding additional trainable layers to the ASR Conformer resulted in improved performance. Compared to the baseline model, adding just two trainable layers demonstrated a marked reduction in EER across all configurations. However, these performance gains tend to plateau. For instance, while adding 2 trainable layers yields a noteworthy improvement, the benefits diminish, or in some cases even slightly reverse, with the addition of 4 layers. One plausible explanation is that the inputs to these lightweight trainable Conformer layers come from highly abstract signals from the ASR model. Therefore, an increase in their number could potentially lead to overfitting.

4) *Comparing the input of the trainable Conformer layers:* Our subsequent investigation aimed at the inputs channeled into the trainable lightweight Conformer layers. We compared configurations V2 and V3, explicitly focusing on $K = 2$ and $K = 4$. In configuration V2, the inputs to the trainable Conformer layer are sourced directly from the frame-level outputs derived from the L -th ASR Conformer layer. Conversely, in configuration V3, the trainable Conformer layer receives its inputs from a concatenation sourced from the ASR model’s first L Conformer layers. A clear distinction in performance emerged from the results: Configuration V3 consistently outperforms V2 across all ASR model sizes and

TABLE VI
EER[%] AND MINDCF OF THE SPEAKER ADAPTATION METHOD ON VOXCCEBLEB 1.

Model	Size		MACs		Vox1-O		Vox1-E		Vox1-H	
	ASR	SpkAdap	ASR	SpkAdap	EER	minDCF	EER	minDCF	EER	minDCF
Small V3 L8 K2 + QMF	6.94M	3.49M	826.39M	116.51M	0.83 0.69	0.1223 0.1011	0.99 0.89	0.1058 0.0930	1.87 1.66	0.1798 0.1663
Medium V3 L10 K2 + QMF	17.79M	4.14M	1.79G	116.71M	0.67 0.55	0.0873 0.0807	0.88 0.80	0.0978 0.0844	1.66 1.48	0.1609 0.1494
Large V3 L10 K2 + QMF	70.85M	4.92M	7.07G	117.33M	0.57 0.45	0.0631 0.0485	0.77 0.69	0.0805 0.0727	1.52 1.35	0.1484 0.1350

TABLE VII
EER[%] OF VOXCCEBLEB 1-O OF DIFFERENT ADAPTATION METHODS APPLIED ON NEMO SMALL ASR-CTC MODEL (15.88M).⁷

#L	L	K = 0		K = 2		K = 4		
		#ASR param	EER	#adap param	EER	#adap param	EER	#adap param
V1	4	3.92M	2.49	0.73M	1.22	2.60M	1.21	4.47M
	8	6.94M	1.77	1.45M	1.26	3.32M	1.12	5.20M
	12	9.95M	1.73	2.18M	1.47	4.05M	1.34	5.92M
V2	4	3.92M	1.47	0.69M	1.13	2.56M	1.05	4.43M
	8	6.94M	1.11	1.37M	1.02	3.24M	0.94	5.12M
	12	9.95M	1.10	2.06M	1.03	3.93M	1.03	5.80M
V3	4	3.92M	—	0.98	—	2.68M	0.95	4.55M
	8	6.94M	—	0.83	—	3.49M	0.83	5.36M
	12	9.95M	—	0.79	—	4.30M	0.66	6.17M

TABLE VIII
EER[%] OF VOXCCEBLEB 1-O OF DIFFERENT ADAPTATION METHODS APPLIED ON NEMO MEDIUM ASR-CTC MODEL (35.26M).⁷

#L	L	K = 0		K = 2		K = 4		
		#ASR param	EER	#adap param	EER	#adap param	EER	#adap param
V1	6	11.44M	1.65	1.63M	1.01	3.50M	0.99	5.37M
	10	17.79M	1.40	2.69M	1.10	4.56M	1.03	6.43M
	14	24.15M	1.34	3.74M	1.20	5.61M	1.19	7.48M
V2	6	11.44M	1.08	1.14M	0.89	3.01M	0.94	4.88M
	10	17.79M	0.93	2.69M	0.84	4.56M	0.84	6.43M
	14	24.15M	0.89	3.74M	0.92	5.61M	0.86	7.48M
V3	6	11.44M	—	0.81	—	3.23M	0.90	5.10M
	10	17.79M	—	0.67	—	4.14M	0.83	6.01M
	14	24.15M	—	0.66	—	5.05M	0.77	6.92M

all values of L . For instance, considering the NEMO Large ASR-CTC model with $L = 14$ and $K = 2$, V3 achieved an EER of 0.55%, this translates to a relative reduction of 29% compared to V2. As shown in the linear probe experiments in section III-E, the early layers of the ASR Conformer model are proficient at gathering speaker-specific information. The concatenation from multiple ASR Conformer layers in V3 captures a more diverse and quality-rich set of information, which proves advantageous for the speaker adaptation module.

For a more thorough evaluation, we test the speaker adaptation method on three testing trials of VoxCeleb 1. We select one speaker adaptation module with the V3 configuration for each NEMO ASR Conformer-CTC model of varying sizes. The results can be found in Table VI. The V3 speaker

TABLE IX
EER[%] OF VOXCCEBLEB 1-O OF DIFFERENT ADAPTATION METHODS APPLIED ON NEMO LARGE ASR-CTC MODEL (130.94M).⁷

#L	L	K = 0		K = 2		K = 4		
		#ASR param	EER	#adap param	EER	#adap param	EER	#adap param
V1	6	45.55M	1.18	3.26M	0.88	5.13M	0.94	7.00M
	10	70.85M	0.97	5.37M	0.85	7.24M	0.89	9.11M
	14	96.14M	1.01	7.48M	0.89	9.35M	0.97	11.23M
V2	6	45.55M	0.86	1.38M	0.72	3.25M	0.78	5.12M
	10	70.85M	0.72	2.23M	0.78	4.11M	0.75	5.98M
	14	96.14M	0.71	3.09M	0.77	4.96M	0.76	6.84M
V3	6	45.55M	—	0.61	—	3.70M	0.69	5.57M
	10	70.85M	—	0.57	—	4.92M	0.65	6.79M
	14	96.14M	—	0.55	—	6.14M	0.65	8.01M

adaptation module with $L = 10$ and $K = 2$ achieves an 0.45% EER using the NEMO Large ASR-CTC model. In comparison, the ASR pretraining and ASR distillation techniques result in EERs of 0.43% and 0.45%, respectively, using the same Large model. While the speaker adaptation method lags slightly behind these two methods, it uniquely offers the capability of unifying ASR and speaker verification within a single Conformer model. This benefit of task unification comes with a relatively modest increase of 4.92 million parameters added to the 130.94 million parameter Large ASR Conformer encoder.

VI. CONCLUSION

This research has presented and evaluated three techniques to leverage ASR pretrained Conformers for speaker verification tasks effectively. Experiments on VoxCeleb datasets validate the efficacy of our proposed methods. First, we have shown that initializing speaker embedding networks with ASR pretrained Conformers lead to significant performance gains and generalization. The extensive ASR pretraining enables the network to extract more robust speaker representations by preventing overfitting to limited speaker data. Second, knowledge distillation from the ASR Conformer teacher to the speaker verification student model allows efficient transfer of ASR expertise. Serving as an auxiliary phonetic modeling task, this distillation approach enhances speaker modeling. Compared to direct ASR pretraining, knowledge distillation offers more flexibility in student model design. Third, our lightweight adaptation modules successfully unify ASR and

speaker verification within a single Conformer model. By refining ASR-learned features for speaker tasks, the adaptation module efficiently bridges the gap between the two modalities. This unified model delivers simultaneous ASR and speaker verification using minimal additional parameters. This research has demonstrated three promising and viable strategies to leverage ASR pretrained Conformers to advance speaker verification performance. Our methods effectively transfer rich ASR knowledge to speaker modeling. We aim to extend our approaches to multilingual models and low-resource settings for further studies.


REFERENCES

- [1] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-vectors: Robust DNN Embeddings for Speaker Recognition,” in *ICASSP*, 2018, pp. 5329–5333.
- [2] 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 TASLP*, vol. 28, pp. 1038–1051, 2020.
- [3] W. Cai, J. Chen, and M. Li, “Exploring the Encoding Layer and Loss Function in End-to-End Speaker and Language Recognition System,” in *Speaker Odyssey*, 2018, pp. 74–81.
- [4] K. Okabe, T. Koshinaka, and K. Shinoda, “Attentive Statistics Pooling for Deep Speaker Embedding,” in *Interspeech*, 2018, pp. 2252–2256.
- [5] T. Zhou, Y. Zhao, and J. Wu, “ResNeXt and Res2Net Structures for Speaker Verification,” in *SLT*, 2021, pp. 301–307.
- [6] B. Desplanques, J. Thienpondt, and K. Demuynck, “ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification,” in *Interspeech*, 2020, pp. 3830–3834.
- [7] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, “ArcFace: Additive Angular Margin Loss for Deep Face Recognition,” in *CVPR*, 2019, pp. 4685–4694.
- [8] X. Xiang, S. Wang, H. Huang, Y. Qian, and K. Yu, “Margin Matters: Towards More Discriminative Deep Neural Network Embeddings for Speaker Recognition,” in *APSPAA*, 2019, pp. 1652–1656.
- [9] 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, pp. 2977–2981.
- [10] D. Garcia-Romero, G. Sell, and A. Mccree, “MagNetO: X-vector Magnitude Estimation Network plus Offset for Improved Speaker Recognition,” in *Odyssey*, 2020, pp. 1–8.
- [11] 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 *ICASSP*, 2021, pp. 5814–5818.
- [12] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in *CVPR*, 2016, pp. 770–778.
- [13] J. Hu, L. Shen, and G. Sun, “Squeeze-and-Excitation Networks,” in *CVPR*, 2018.
- [14] S.-H. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. Torr, “Res2Net: A New Multi-Scale Backbone Architecture,” *IEEE TPAMI*, vol. 43, no. 2, pp. 652–662, 2021.
- [15] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, “Aggregated Residual Transformations for Deep Neural Networks,” in *CVPR*, 2017.
- [16] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is All You Need,” in *NeurIPS*, 2017.
- [17] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *Interspeech*, 2020, pp. 5036–5040.
- [18] Y. Zhang, Z. Lv, H. Wu, S. Zhang, P. Hu, Z. Wu, H.-y. Lee, and H. Meng, “MFA-Conformer: Multi-scale Feature Aggregation Conformer for Automatic Speaker Verification,” in *Interspeech*, 2022, pp. 306–310.
- [19] D. Liao, T. Jiang, F. Wang, L. Li, and Q. Hong, “Towards A Unified Conformer Structure: from ASR to ASV Task,” in *ICASSP*, 2023, pp. 1–5.
- [20] D. Cai, W. Wang, M. Li, R. Xia, and C. Huang, “Pretraining Conformer with ASR for Speaker Verification,” in *ICASSP*, 2023, pp. 1–5.
- [21] T. Zhou, Y. Zhao, J. Li, Y. Gong, and J. Wu, “CNN with Phonetic Attention for Text-Independent Speaker Verification,” in *ASRU*, 2019, pp. 718–725.
- [22] M. Li, L. Liu, W. Cai, and W. Liu, “Generalized I-vector Representation with Phonetic Tokenizations and Tandem Features for both Text Independent and Text Dependent Speaker Verification,” *Journal of Signal Processing Systems*, vol. 82, no. 2, pp. 207–215, 2016.
- [23] Y. Lei, N. Scheffer, L. Ferrer, and M. McLaren, “A Novel Scheme for Speaker Recognition using a Phonetically-Aware Deep Neural Network,” in *ICASSP*, 2014, pp. 1695–1699.
- [24] G. Hinton, O. Vinyals, and J. Dean, “Distilling the Knowledge in a Neural Network,” in *NeurIPS Deep Learning and Representation Learning Workshop*, 2015.
- [25] Z. Fan, M. Li, S. Zhou, and B. Xu, “Exploring Wav2vec 2.0 on Speaker Verification and Language Identification,” in *Interspeech*, 2021, pp. 1509–1513.
- [26] N. Vaessen and D. A. van Leeuwen, “Fine-Tuning Wav2vec2 for Speaker Recognition,” in *ICASSP*, 2022, pp. 7967–7971.
- [27] S. Novoselov, G. Lavrentyeva, A. Avdeeva, V. Volokhov, and A. Gusev, “Robust Speaker Recognition with Transformers Using wav2vec 2.0,” *arXiv:2203.15095*, 2022.
- [28] Z. Chen, S. Chen, Y. Wu, Y. Qian, C. Wang, S. Liu, Y. Qian, and M. Zeng, “Large-Scale Self-Supervised Speech Representation Learning for Automatic Speaker Verification,” in *ICASSP*, 2022, pp. 6147–6151.
- [29] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, “WavLM: Large-Scale Self-Supervised Pre-Training for Full Stack Speech Processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [30] J. Peng, O. Plchot, T. Stafylakis, L. Mošner, L. Burget, and J. Černocký, “An Attention-Based Backend Allowing Efficient Fine-Tuning of Transformer Models for Speaker Verification,” in *SLT*, 2022, pp. 555–562.
- [31] D. Snyder, D. Garcia-Romero, and D. Povey, “Time Delay Deep Neural Network-based Universal Background Models for Speaker Recognition,” in *ASRU*, 2015, pp. 92–97.
- [32] Y. Tian, M. Cai, L. He, and J. Liu, “Investigation of Bottleneck Features and Multilingual Deep Neural Networks for Speaker Verification,” in *Interspeech*, 2015, pp. 1151–1155.
- [33] M. H. Rahman, I. Himawan, M. McLaren, C. Fookes, and S. Sridharan, “Employing Phonetic Information in DNN Speaker Embeddings to Improve Speaker Recognition Performance,” in *Interspeech*, 2018, pp. 3593–3597.
- [34] S. Zheng, Y. Lei, and H. Suo, “Phonetically-Aware Coupled Network For Short Duration Text-Independent Speaker Verification,” in *Interspeech*, 2020, pp. 926–930.
- [35] Y. Liu, L. He, J. Liu, and M. T. Johnson, “Speaker Embedding Extraction with Phonetic Information,” in *Interspeech*, 2018, pp. 2247–2251.
- [36] X. Chen and C. Bao, “Phoneme-Unit-Specific Time-Delay Neural Network for Speaker Verification,” *IEEE/ACM TASLP*, vol. 29, pp. 1243–1255, 2021.
- [37] Z. Tang, L. Li, D. Wang, and R. Vipperla, “Collaborative Joint Training With Multitask Recurrent Model for Speech and Speaker Recognition,” *IEEE/ACM TASLP*, vol. 25, no. 3, pp. 493–504, 2017.
- [38] S. Wang, J. Rohdin, L. Burget, O. Plchot, Y. Qian, K. Yu, and J. Černocký, “On the Usage of Phonetic Information for Text-Independent Speaker Embedding Extraction,” in *Interspeech*, 2019, pp. 1148–1152.
- [39] N. Tawara, A. Ogawa, T. Iwata, M. Delcroix, and T. Ogawa, “Frame-Level Phoneme-Invariant Speaker Embedding for Text-Independent Speaker Recognition on Extremely Short Utterances,” in *ICASSP*, 2020, pp. 6799–6803.
- [40] Q.-B. Hong, C.-H. Wu, and H.-M. Wang, “Decomposition and Reorganization of Phonetic Information for Speaker Embedding Learning,” *IEEE/ACM TASLP*, vol. 31, pp. 1745–1757, 2023.
- [41] N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. de Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly, “Parameter-Efficient Transfer Learning for NLP,” in *ICML*, 2019, pp. 2790–2799.
- [42] S.-A. Rebuffi, H. Bilen, and A. Vedaldi, “Learning Multiple Visual Domains with Residual Adapters,” in *NeurIPS*, vol. 30, 2017.
- [43] J. He, C. Zhou, X. Ma, T. Berg-Kirkpatrick, and G. Neubig, “Towards a Unified View of Parameter-Efficient Transfer Learning,” in *ICLR*, 2022.
- [44] A. Bapna, N. Arivazhagan, and O. Firat, “Simple, Scalable Adaptation for Neural Machine Translation,” in *EMNLP*, 2019.
- [45] B. Thomas, S. Kessler, and S. Karout, “Efficient Adapter Transfer of Self-Supervised Speech Models for Automatic Speech Recognition,” in *ICASSP*, 2022, pp. 7102–7106.
- [46] A. Kannan, A. Datta, T. N. Sainath, E. Weinstein, B. Ramabhadran, Y. Wu, A. Bapna, Z. Chen, and S. Lee, “Large-Scale Multilingual Speech Recognition with a Streaming End-to-End Model,” in *Interspeech*, 2019, pp. 2130–2134.

- [47] G. I. Winata, G. Wang, C. Xiong, and S. Hoi, "Adapt-and-Adjust: Overcoming the Long-Tail Problem of Multilingual Speech Recognition," in *Interspeech*, 2021, pp. 2451–2455.
- [48] H. Le, J. Pino, C. Wang, J. Gu, D. Schwab, and L. Besacier, "Lightweight Adapter Tuning for Multilingual Speech Translation," in *ACL-IJCNLP*, 2021, pp. 817–824.
- [49] J. Peng, T. Stafylakis, R. Gu, O. Plchot, L. Mošner, L. Burget, and J. Černocký, "Parameter-Efficient Transfer Learning of Pre-Trained Transformer Models for Speaker Verification Using Adapters," in *ICASSP*, 2023, pp. 1–5.
- [50] S. Otake, R. Kawakami, and N. Inoue, "Parameter Efficient Transfer Learning for Various Speech Processing Tasks," in *ICASSP*, 2023, pp. 1–5.
- [51] Z. Dai, Z. Yang, Y. Yang, J. G. Carbonell, Q. Le, and R. Salakhutdinov, "Transformer-XL: Attentive Language Models beyond a Fixed-Length Context," in *ACL*, 2019, pp. 2978–2988.
- [52] Y. Lu, Z. Li, D. He, Z. Sun, B. Dong, T. Qin, L. Wang, and T.-Y. Liu, "Understanding and Improving Transformer from a Multi-Particle Dynamic System Point of View," *arXiv:1906.02762*, 2019.
- [53] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive Statistics Pooling for Deep Speaker Embedding," in *Interspeech*, 2018, pp. 2252–2256.
- [54] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units," *IEEE/ACM TASLP*, vol. 29, pp. 3451–3460, 2021.
- [55] P. Xu, D. Kumar, W. Yang, W. Zi, K. Tang, C. Huang, J. C. K. Cheung, S. J. Prince, and Y. Cao, "Optimizing Deeper Transformers on Small Datasets," in *ACL IJCNLP*, 2021, pp. 2089–2102.
- [56] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: A Large-Scale Speaker Identification Dataset," in *Interspeech*, 2017, pp. 2616–2620.
- [57] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep Speaker Recognition," in *Interspeech*, 2018, pp. 1086–1090.
- [58] H. Yamamoto, K. A. Lee, K. Okabe, and T. Koshinaka, "Speaker Augmentation and Bandwidth Extension for Deep Speaker Embedding," in *Interspeech*, 2019, pp. 406–410.
- [59] W. Wang, D. Cai, X. Qin, and M. Li, "The DKU-DukeECE Systems for VoxCeleb Speaker Recognition Challenge 2020," *arXiv:2010.12731*, 2020.
- [60] D. Snyder, G. Chen, and D. Povey, "MUSAN: A Music, Speech, and Noise Corpus," *arXiv:1510.08484*, 2015.
- [61] 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 *ICASSP*, 2017, pp. 5220–5224.
- [62] O. Kuchaiev, J. Li, H. Nguyen, O. Hrinchuk, R. Leary, B. Ginsburg, S. Kriman, S. Beliaev, V. Lavrukhin, J. Cook *et al.*, "Nemo: a Toolkit for Building AI Applications Using Neural Modules," *arXiv:1909.09577*, 2019.
- [63] P. Matějka, O. Novotný, O. Plchot, L. Burget, M. D. Sánchez, and J. Černocký, "Analysis of Score Normalization in Multilingual Speaker Recognition," in *Interspeech*, 2017, pp. 1567–1571.
- [64] M. I. Mandasari, R. Saeidi, M. McLaren, and D. A. van Leeuwen, "Quality Measure Functions for Calibration of Speaker Recognition Systems in Various Duration Conditions," *IEEE TASLP*, vol. 21, no. 11, pp. 2425–2438, 2013.
- [65] "NIST 2016 Speaker Recognition Evaluation Plan," 2016. [Online]. Available: https://www.nist.gov/system/files/documents/2016/10/07/sre16_eval_plan_v1.3.pdf



Ming Li (Senior Member, IEEE) received his Ph.D. in Electrical Engineering from University of Southern California in 2013. He is currently an Associate Professor of Electrical and Computer Engineering at Duke Kunshan University. He is also an Adjunct Professor at School of Computer Science in Wuhan University. His research interests are in the areas of audio, speech and language processing as well as multimodal behavior signal processing. He has published more than 180 papers and served as the member of IEEE speech and language technical committee, APSIPA speech and language processing technical committee. He is an area chair at Interspeech 2016, 2018, 2020 and 2024, as well as the technical program co-chair of Odyssey 2022 and ASRU 2023. Works co-authored with his colleagues have won first prize awards at Interspeech Computational Paralinguistic Challenges 2011, 2012 and 2019, ASRU 2019 MGB-5 ADI Challenge, Interspeech 2020 and 2021 Fearless Steps Challenges, VoxSRC 2021, 2022 and 2023 Challenges, ICASSP 2022 M2MeT Challenge, IJCAI 2023 ADD challenge and ICME 2024 ChatCLR challenge. He received the IBM faculty award in 2016, the ISCA Computer Speech and Language 5-years best journal paper award in 2018 and the youth achievement award of outstanding scientific research achievements of Chinese higher education in 2020. He is a senior member of IEEE.



Danwei Cai is pursuing his Ph.D. degree in electrical and computer engineering at Duke University. He received his bachelor's degree in software engineering and master's degree in electronics and communication engineering from Sun Yet-Sen University in China. His primary research interests are in the area of speech processing, including speech recognition, speaker recognition, speaker diarization and computational linguistics.