

Self-Supervised Training of Speaker Encoder with Multi-Modal Diverse Positive Pairs

Ruijie Tao, *Student Member, IEEE*, Kong Aik Lee, *Senior Member, IEEE*, Rohan Kumar Das, *Senior Member, IEEE*, Ville Hautamäki, *Member, IEEE*, and Haizhou Li, *Fellow, IEEE*

Abstract—We study a novel neural architecture and its training strategies of speaker encoder for speaker recognition without using any identity labels. The speaker encoder is trained to extract a fixed dimensional speaker embedding from a spoken utterance of various length. Contrastive learning is a typical self-supervised learning technique. However, the contrastive learning of the speaker encoder depends very much on the sampling strategy of positive and negative pairs. It is common that we sample a positive pair of segments from the same utterance. Unfortunately, such poor-man’s positive pairs (PPP) lack the necessary diversity. In this work, we propose a multi-modal contrastive learning technique with novel sampling strategies. By cross-referencing between speech and face data, we find diverse positive pairs (DPP) for contrastive learning, thus improving the robustness of the speaker encoder. We train the speaker encoder on the VoxCeleb2 dataset without any speaker labels, and achieve an equal error rate (EER) of 2.89%, 3.17% and 6.27% under the proposed progressive clustering strategy, and an EER of 1.44%, 1.77% and 3.27% under the two-stage learning strategy with pseudo labels, on the three test sets of VoxCeleb1. This novel solution outperforms the state-of-the-art self-supervised learning methods by a large margin, at the same time, achieves comparable results with the supervised learning counterpart. We also evaluate our self-supervised learning technique on LRS2 and LRW datasets, where the speaker information is unknown. All experiments suggest that the proposed neural architecture and sampling strategies are robust across datasets.

Index Terms—Self-supervised learning, speaker recognition, diverse positive pairs, multi-modal, progressive clustering

I. INTRODUCTION

SPEAKER recognition (SR) seeks to authenticate an identity claim by using the speaker’s voice [1]–[3]. It typically relies on a speaker encoder that transforms a speech sample into a speaker embedding vector. For supervised learning of

This work is supported by 1) Huawei Noah’s Ark Lab; 2) National Natural Science Foundation of China (Grant No. 62271432); 3) Internal Project Fund from Shenzhen Research Institute of Big Data (Grant No. T00120220002); 4) Human-Robot Collaborative AI for Advanced Manufacturing and Engineering (Grant No. A18A2b0046), Agency of Science, Technology and Research, Singapore. (*Corresponding author: Ruijie Tao*).

Ruijie Tao, and Ville Hautamäki are with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 119077 (e-mail: ruijie.tao@u.nus.edu; villeh@cs.joensuu.fi).

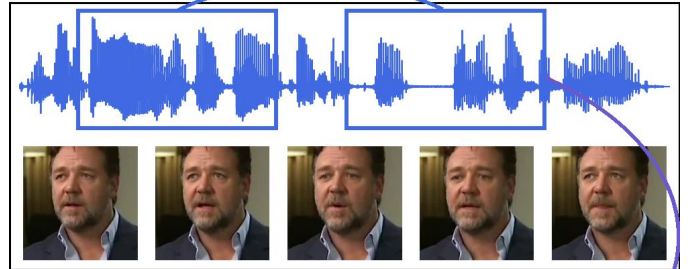
Kong Aik Lee is with the Institute for Infocomm Research, A*STAR, Singapore 138632 (e-mail: lee_kong_aik@i2r.a-star.edu.sg).

Rohan Kumar Das is with Fortemedia, Singapore 138589 (e-mail: ecerohan@gmail.com).

Ville Hautamäki is also with the School of Computing, University of Eastern Finland, Joensuu 80101, Finland.

Haizhou Li is with Shenzhen Research Institute of Big Data, School of Data Science, The Chinese University of Hong Kong, Shenzhen, China, and Kriston AI, China (e-mail: haizhouli@cuhk.edu.cn).

Poor-man's positive pairs: segments from the same utterance



Diverse positive pairs: segments from different utterances

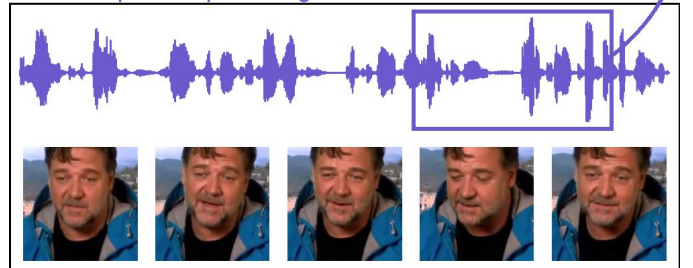


Fig. 1. A segment is an excerpt from a video clip. In contrastive learning, a pair of two segments, either positive or negative, forms a training data point. A poor-man’s positive pair (PPP) in the upper panel is made up by two segments from the same utterance, that represents the same speaker, acoustic environment, speaker state, and discussion topic, therefore, under a homogeneous acoustic condition. A diverse positive pair (DPP) is made up by two segments from two distinct utterances of the same speaker, one from the upper panel and another from the lower panel, where only their speaker identity is in common. DPP is more effective for comparison.

speaker encoder, a large-scale dataset with manually annotated speaker labels is required [4]–[6]. As manual annotation is labour intensive and costly, the *self-supervised learning* (SSL) learning technique becomes a competitive alternative by solving a pretext task on unlabelled data [7]. It has shown promising results in many areas, such as GPT [8] and BERT [9] in natural language processing (NLP), MOCO [10], BYOL [11] and DINO [12] in computer vision (CV), wav2vec [13] and HuBERT [14] in speech processing. With these success, we are prompted to investigate the training of speaker encoder on the abundantly available unlabelled data.

Contrastive learning [7], [15] is a successful implementation of self-supervised learning. It forces the encoder to produce similar representations between a pair of positive samples, i.e., speech samples by the same speaker. A positive pair contains an anchor segment and a positive counterpart, which are typically two disjoint segments in the same utter-

ance [15], [16], while a negative pair consists of two speech segments from different speakers, typically from two distant utterances. For each anchor segment, the speaker encoder learns to discriminate the positive pair from all negative pairs in the mini-batch.

It is efficient to sample negative pairs from two distant utterances. However, we believe that the positive pairs from the same utterance are not the best learning target as they lack sufficient diversity. While contrastive learning encourages the speaker encoder to learn the speaker voice characteristic [17], the resulting encoder is also affected by other confounding factors. For instance, in an utterance from an indoor talk show in the upper panel of Fig. 1, the speaker encoder may also learn the spoken content, the speaker emotion, the speaker state, the acoustic environment, and the recording channel, if the positive pairs are always extracted from the same utterance during comparison. We refer to such positive pairs as the *poor-man's positive pairs* (PPP).

In contrast, if a positive pair is extracted from two distant utterances of the same speaker, for example, an indoor and an outdoor interview of the same person in Fig. 1, we can take one segment from each of the two utterances to form a positive pair. In this way, the non-identity information is very different between the two samples, thus greatly reducing the effect of the confounding factors. We refer to such positive pairs as the *diverse positive pairs* (DPP). As opposed to the poor-man's positive pairs (PPP), we have good reason to expect that the DPP will serve better contrastive learning than the PPP counterpart.

In general, prior studies also suggest that contrastive learning benefits from diverse and varied training samples. The study on the prototypical loss in the supervised learning paradigm shows that speaker recognition benefits from varied positive samples generated from the ground-truth speaker labels across utterances [18]. The similar idea has been validated in computer vision. In [19], it is suggested to find the nearest neighbour for each anchor image as the positive counterpart rather than the augmented anchor image. In SCAN [20] and CoCLR [21], a fixed number of positive pairs for each image are discovered after one round of contrastive learning. The newly found positive pairs are then used for a new round of contrastive training. These studies all point to the direction that DPP lead to better models. To the best of our knowledge, there is no study of DPP in self-supervised learning of speaker encoder yet.

In this work, we hypothesize that DPP will outperform PPP in the self-supervised learning of the speaker encoder. The question is how to sample the DPP such that they are both accurate, i.e., from the same speaker, and diverse, i.e., varying across different acoustic conditions. One way is to use the anchor utterance to search for positive utterances of the same speaker in the database. However, it can hardly guarantee the accuracy and diversity of found positive pairs. From the biometric recognition study, we know that facial image and voice constitute complementary biometric evidence [22], [23]. So we are motivated to apply both audio and visual data to find positive counterparts that are accurate and diverse.

We are inspired by the co-training technique to construct the

framework, which describes a data sample from two different views and enhances two encoders gradually [24], [25]. We involve a face encoder and train it with the speaker encoder together. To ensure that the found positive pairs are truly positive, we make use of the complementary nature of the two modalities, then exploit both the audio and visual modalities to search for positive pairs of video clips. This complementary effect improves the quality of the found positive pairs. As far as diversity is concerned, the cross-modal co-reference allows us to find positive speech pairs from very different acoustic conditions, and positive pairs of facial images from very different photographic environments.

We make the following contributions in this paper.

- For the first time, we hypothesize and validate the idea of diverse positive pairs (DPP) for self-supervised learning of speaker encoder.
- We propose a multi-modal contrastive learning (MCL) framework with diverse positive pairs (MCL-DPP) via a novel neural architecture and formulate its self-supervised learning strategies.
- We successfully implement MCL and MCL-DPP frameworks and achieve state-of-the-art performance for self-supervised learning, that is comparable with its supervised learning counterpart.

II. RELATED WORK

A. Speaker encoder and speaker recognition

A neural network solution to speaker recognition typically consists of a speaker encoder and a speaker comparison module mechanism.

The speaker encoder learns to convert a time-domain speech signal or its spectral features, i.e., spectrograms, filter banks, and Mel-frequency cepstral coefficients (MFCCs) [26] into an utterance-level speaker embedding. The examples of speaker encoders include time-delay neural networks (TDNN) based x-vector [27], convolutional neural network (CNN) based ResNet [18]. Recently, the emphasized channel attention, propagation and aggregation in time-delay neural network (ECAPA-TDNN), has attracted much attention [28] that adopts many advanced designs, such as Res2Net blocks [29], squeeze-and-excitation blocks [30] and multi-layer feature aggregation. As a speaker characterization frontend, the speaker encoder usually be trained in a supervised manner with classification objectives [27] or metric learning objectives [18].

The speaker comparison module is designed to decide if two speaker embeddings are from the same speaker. At run-time for speaker verification, the cosine similarity [18] or probabilistic linear discriminant analysis (PLDA) [31] backend can be used to calculate the similarity between the test and target speaker embeddings. It is noted that speaker embeddings are also widely used in related areas, such as speaker diarization [32], speaker extraction [33], text-to-speech synthesis (TTS) [34] and voice conversion [35]. Therefore, the quality of the speaker encoder is all important across many studies.

B. Self-supervised learning of speaker encoder

Self-supervised learning is achieved by deriving supervisory signals from the unlabelled data itself. It leverages the intrinsic

structural knowledge in the data. For speaker encoder training, there are two general design strategies for supervisory signals, namely single-stage learning, and two-stage learning.

1) *Single-stage learning*: Single-stage learning is a typical end-to-end training following a comparison-based pretext task. The key is to construct this pretext task effectively. *Simple contrastive learning* (SCL) technique trains the speaker encoder by attracting positive pairs (two augmented segments from the same utterance) and repelling negative pairs (two augmented segments from different utterances) [7], [17]. Others further set additional training targets to improve the comparative efficiency, such as invariance of augmentation [17], invariance of channel [16], equilibrium learning [36] and positive term regularization [37].

Besides SCL, other comparison-based self-supervised learning techniques include the MOCO framework [38], [39], which stores the negative pairs in the memory bank; the DINO framework [12], [40]–[42] that only involves positive pairs and achieves considerable improvement. For efficiency and effectiveness, we adopt the SCL framework in this study and focus on the sampling strategy of positive pairs. It is noted that our proposal can be extended to other frameworks, such as MOCO and DINO.

2) *Two-stage learning*: With a two-stage learning strategy, we view single-stage learning as the first stage and improve it with pseudo labels in the second stage. Based on the trained speaker encoder in the first stage, speaker embeddings can be derived from the unlabelled speech data for unsupervised clustering. In the second stage, the cluster identity of a speech sample serves as its pseudo speaker label for the supervised learning of the speaker encoder. The clustering-training loop is repeated to improve the speaker encoder [15], [39], [43].

In our previous work [44], we focus our study on the second stage as to how to effectively find reliable pseudo labels. The studies on two-stage learning validated that the idea of pseudo labels greatly benefits from unlabelled data. In this work, we would like to focus our study on the first stage of diverse positive pairs.

C. Multi-modal speaker recognition

Human face also provides identity information [45], that is helpful for speaker encoder training. Under the supervised learning framework, the speech-face early [22], middle [22], [46] and late [22], [23], [47], [48] fusion strategies were studied. They all concluded that human face provides complementary biometric information in speaker recognition.

In the recent study of self-supervised learning for speaker encoder [49], the visual modality is introduced in the second stage of a two-stage learning system where speech-face multi-modal embeddings are used to improve speaker clustering, thus the quality of speaker encoder. It remains a challenge how we effectively use the abundant unlabelled videos with speech-face pairs in the single-stage learning system, that motivates this study.

III. MCL: MULTI-MODAL CONTRASTIVE LEARNING

We now propose an end-to-end multi-modal contrastive learning framework (MCL) as the baseline. As shown in Fig. 2,

the MCL consists of three components: contrastive learning for speaker encoder with speech input, contrastive learning for face encoder with face frames input, and the embedding projector network with cross-modal joint loss. Here, each video clip in the training set contains only one talking face. Such data can be obtained via an audio-visual active speaker detection system [50].

We assume that the speech segments or face frames drawn from the same video clip share the same identity. On the other hand, those drawn from different video clips belong to different people. In this way, we don't rule out the possibility of having false-negative pairs. However, considering the size of the mini-batch with respect to a relatively large training set [16], the side effect of these false-negative pairs can be ignored.

A. Contrastive learning for speaker encoder

The contrastive learning scheme of the speaker encoder is similar to that in our previous work [44]. As shown in the orange box in Fig. 2. Each training video clip x_i contains the speech utterance $x_i^{(s)}$ and face frames $x_i^{(f)}$. For one utterance, we randomly consider two same-length, disjoint speech segments $x_{i,1}^{(s)}$ and $x_{i,2}^{(s)}$ after stochastic noise augmentation. When we view $x_{i,1}^{(s)}$ as the anchor segment, $x_{i,2}^{(s)}$ is the positive counterpart. They are the inputs of the speaker encoder $E^{(s)}(\cdot)$, and the outputs are speaker embeddings $y_{i,1}^{(s)}$ and $y_{i,2}^{(s)}$. As shown in (1), a contrastive loss [7] is a function whose value is low when anchor segment $x_{i,1}^{(s)}$ is similar to its positive counterpart $x_{i,2}^{(s)}$ and dissimilar to all other segments in the mini-batch (i.e., negative segments). Let $s(a, b) = \exp(\cos(a, b)) / \tau$, where \cos is the cosine similarity, and τ is the temperate parameter. **Following the previous study [7], we set τ as 0.1.** The loss function $\mathcal{L}^{(s)}$ for the mini-batch is defined as:

$$\mathcal{L}^{(s)} = \frac{1}{2M} \sum_{i=1}^M \sum_{j=1}^2 \left(-\log \frac{s(y_{i,1}^{(s)}, y_{i,2}^{(s)})}{\sum_{k=1}^M \sum_{l=1}^2 \mathbb{1}_{k \neq i, l \neq j} s(y_{i,j}^{(s)}, y_{k,l}^{(s)})} \right) \quad (1)$$

where M is the batch size, $\mathbb{1}$ is an indicator function evaluating 1 when $k \neq i$ and $l \neq j$. For each segment, there are one positive pair ($y_{i,1}^{(s)}$ with $y_{i,2}^{(s)}$) and $2(M - 1)$ negative pairs since each utterance provides two segments.

B. Contrastive learning for face encoder

There is a lack of study on self-supervised learning of face encoder. As shown in the bottom left blue box of Fig. 2, we formulate the contrastive learning for face frame sequence in the same way as that for speech signal. In other words, two speech segments are replaced by two face frames $x_{i,1}^{(f)}$ and $x_{i,2}^{(f)}$, stochastic noise augmentation is substituted by image augmentation. Similarly, the face encoder is denoted as $E^{(f)}(\cdot)$ that derives the face embeddings $y_{i,1}^{(f)}$ and $y_{i,2}^{(f)}$. We follow the same selection process for positive and negative pairs. The face contrastive loss $\mathcal{L}^{(f)}$ is also similar to that for speaker encoder training.

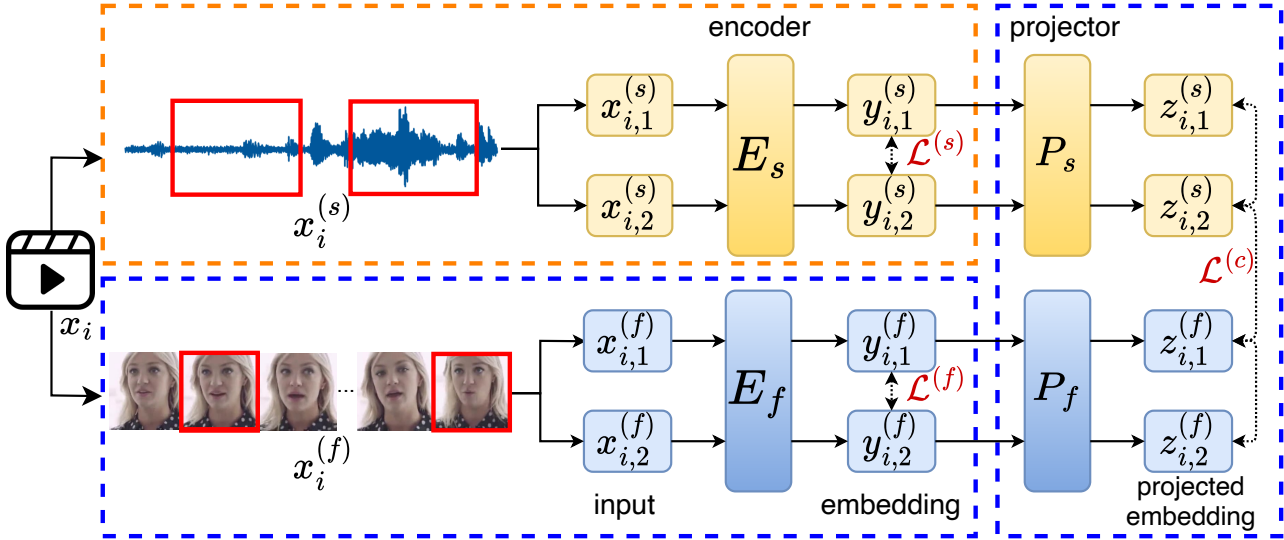


Fig. 2. The proposed baseline: *Multi-modal Contrastive Learning* (MCL) framework that consists of a speaker encoder (upper left box), a face encoder (lower left box) and two embedding projectors (right box). The contrastive learning of the speaker encoder follows that in our previous work [44]. The contrastive learning of the face encoder is formulated in this paper. We also propose a cross-modal module with two embedding projectors, to learn the cross-modal association.

C. Joint framework with cross-modal loss

A simple solution to combine multi-modal information is to train the mentioned speaker and face encoder independently and apply a score-level fusion strategy [23] for decision making. Such late fusion strategy doesn't take advantage of the interaction between two modalities. For example, as both speaker and face embedding are derived from the same video clip, they may form a positive pair. However, we can not directly measure their similarity because the speaker and face encoders are in different embedding spaces.

As shown in the right blue box in Fig. 2, for each video clip x_i , suppose that we have the speaker embedding $y_{i,1}^{(s)}$ and $y_{i,2}^{(s)}$, face embedding $y_{i,1}^{(f)}$ and $y_{i,2}^{(f)}$. We propose two projectors $P^{(s)}(\cdot)$ and $P^{(f)}(\cdot)$, that map the speaker and face embedding into the same space with the speaker projected embeddings $z_{i,1}^{(s)}$ and $z_{i,2}^{(s)}$, face projected embeddings $z_{i,1}^{(f)}$ and $z_{i,2}^{(f)}$. Here the cross-modal loss function $\mathcal{L}^{(c)}$ for the mini-batch is represented as:

$$\mathcal{L}^{(c)} = \frac{1}{4M} \sum_{i=1}^M \sum_{j=1}^2 (l_{i,j}^{(s)} + l_{i,j}^{(f)}) \quad (2)$$

Assuming that the speaker and face projected embeddings extracted from the same video clip form a positive pair, and those from different video clips form a negative pair. $l_{i,j}^{(s)}$ and $l_{i,j}^{(f)}$ in (2) are given by:

$$l_{i,j}^{(s)} = -\log \frac{s(z_{i,j}^{(s)}, z_{i,1}^{(f)}) + s(z_{i,j}^{(s)}, z_{i,2}^{(f)})}{\sum_{k=1}^M \sum_{l=1}^2 s(z_{i,j}^{(s)}, z_{k,l}^{(f)})} \quad (3)$$

$$l_{i,j}^{(f)} = -\log \frac{s(z_{i,j}^{(f)}, z_{i,1}^{(s)}) + s(z_{i,j}^{(f)}, z_{i,2}^{(s)})}{\sum_{k=1}^M \sum_{l=1}^2 s(z_{i,j}^{(f)}, z_{k,l}^{(s)})} \quad (4)$$

No indicator function is applied since this is a cross-modal loss. Finally, the combined training loss in MCL is the sum of the speaker loss $\mathcal{L}^{(s)}$, face loss $\mathcal{L}^{(f)}$ and cross-modal loss $\mathcal{L}^{(c)}$. During the evaluation, the projectors are not involved as we only compare two speaker embeddings to make a speaker verification decision. Note that the cross-modal loss is used to optimize the framework architecture.

IV. MCL-DPP: MCL WITH DIVERSE POSITIVE PAIRS

The success of multi-modal contrastive learning relies on an effective sampling of positive and negative pairs. We argue that the poor-man's positive pairs (PPP) sampling technique as discussed in Section III doesn't provide the necessary diversity for robust self-supervised learning. We propose the MCL-DPP framework with a progressive clustering algorithm to improve the diversity of positive pairs.

A. Trade-off between diversity and accuracy of positive pairs

The quality of contrastive learning depends on the training samples. When sampling positive pairs, we would like the pairs to be truly positive, at the same time, diverse enough. The PPP sampling technique in speaker encoder training ensures that two segments are positive because they are from the same utterance. However, they are generally homogeneous in terms of acoustic conditions, linguistic content, and speech prosody among others. They are far from diverse.

To improve the diversity of positive pairs, we propose to relax the restriction by allowing two positive segments to come from different utterances. Without the ground-truth labels, the relaxed condition may lead to false-positive pairs. In other words, there is a trade-off between diversity and accuracy.

Here we use both the speaker and face modality to improve the accuracy and diversity of found positive pairs. **For each video clip x , we extract the speaker projected embedding**

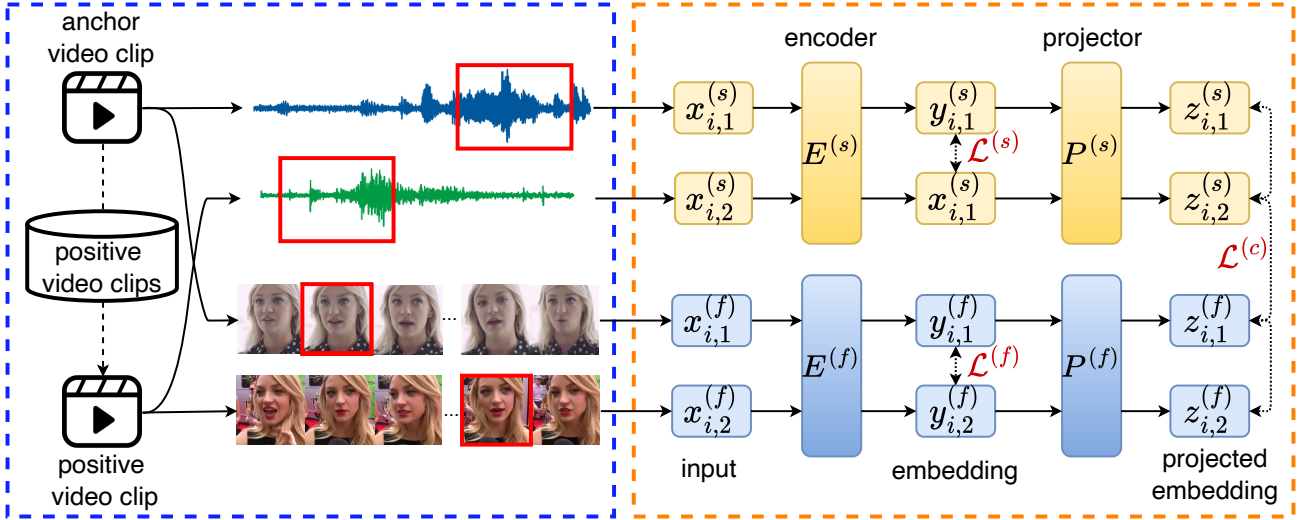


Fig. 3. The proposed framework: *Multi-modal Contrastive Learning with Diverse Positive Pairs* (MCL-DPP). The module in the orange box is the same as that in MCL. The difference is the data sampling module in the blue box. For each anchor video clip, we select one positive video clip to provide positive speech segment and face frame. The positive video clips are found from progressive clustering in our proposed algorithm.

$z^{(s)}$ and face projected embedding $z^{(f)}$ from the whole clean speech utterance and one single face frame to guarantee the robust representation. We concatenate them as the multi-modal projected embedding, $z = z^{(s)} \oplus z^{(f)}$, to serve as a query to retrieve its diverse positive video clips from the training database beyond the current video clip. It is apparent that two multi-modal projected embeddings of higher cosine similarity score are more likely to be from the same speaker, thus, forming a positive pair. These projected embeddings can also be replaced by the un-projected speaker embedding $y^{(s)}$ and face embedding $y^{(f)}$. For the fusion approach, previous audio-visual person recognition works have shown that concatenation is an efficient method for merging identity information from two modalities [46], [49]. The reason is that temporal information is not crucial in this task.

B. Search DPP by clustering

Here we introduce a clustering-based approach to obtain diverse positive pairs. By clustering, we assume that the video clips in the same cluster are likely to be from the same speaker. The video clips, as represented by their projected embeddings, are assigned to clusters in their projected domain. Let's set the number of clusters as C . Any two video clips in the same cluster form a positive pair since they are likely to be from the same person. Given an anchor video clip x_a as the query, we denote the found positive video clip as x_p :

$$x_p \in \{x_i \mid c(x_a) = c(x_i)\}, i \in [1, N] \quad (5)$$

where $c(\cdot)$ is the cluster identity for a video clip. N is the total number of training video clips.

When applying the sampling by clustering, an appropriate clustering technique is required. We adopt the k -means algorithm for its computational efficiency [51]. It is not trivial to estimate the actual number of speakers in an unlabelled dataset. Note that the objective of MCL-DPP is to find

Algorithm 1 Pseudocode for the progressive clustering algorithm used in MCL-DPP framework

```
# dimension of  $X_s$  and  $X_f$ : (M, 2, L) and (M, 2, S)
# L: length of speech segment, S: size of face frame
# M: batch size, C: number of clusters
# C starts from amount of training data
# l(): contrastive loss function
# sample(): randomly select one video clip
# k-means(): get the DPP dictionary by clustering.
# In this dictionary, each key is one video clip and corresponding values are clips in the same cluster.
```

```
dic_DPP = k-means(C) # DPP dictionary initialization.
```

```
def train_one_epoch():
```

```
    for  $X_s, X_f$  in loader: # load a minibatch
         $Y_s, Y_f = E_s(X_s), E_f(X_f)$  # embeddings
         $Z_s, Z_f = P_s(Y_s), P_f(Y_f)$  # projected embed
        loss = l( $Y_s$ ) + l( $Y_f$ ) + l( $Z_s, Z_f$ )
        loss.backward() # back-propagate
        update( $E_s, E_f, P_s, P_f$ ) # update network
    res_this_e = validate() # validation result in this epoch
    # validation result needs to be improved every epoch
    if res_this_e is worse than res_last_e:
        C = C / 2 # divide number of clusters by two
        dic_DPP = k-means(C) # update DPP dictionary
```

```
def load_one_data():
```

```
    anchor = sample(all_data) # pick an anchor video clip
    positives = dic_DPP[anchor] # its positive video clips
    positive = sample(positives) # pick a positive video clip
     $x_s = [anchor_s, positive_s]$  # speech segments
     $x_f = [anchor_f, positive_f]$  # face frames
    return  $x_s, x_f$ 
```

the correct positive pairs rather than the correct number of speakers. It could still serve our purpose if the clustering results in more clusters than the actual number of speakers.

As shown in Fig. 4, we use an example to illustrate the effect of the number of clusters. We randomly select 100 video clips from 2 speakers and show their natural grouping on the left panel of Fig. 4. With 10 k -means clusters, most of the video clips in each cluster come from the same people. For the anchor video clip, all of the found positive counterparts are from the same speaker. With 5 k -means clusters, we see larger clusters than 10 k -means clusters, and more positive counterparts for each anchor video clip. Such positive counterparts are further away from the anchor video clip, thus more diverse.

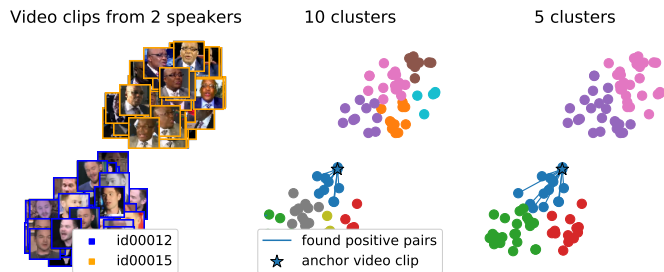


Fig. 4. Explanation for reasonability of progressive clustering. Left Panel: The natural grouping of a set of randomly selected 100 video clips for 2 speakers; Middle Panel: 10 k -means clusters, each has few video clips; Right Panel: 5 k -means clusters, each has more video clips. The clusters are colored differently.

C. Progressive clustering

Based on this finding, we propose the progressive clustering algorithm in MCL-DPP. Progressive clustering is motivated by the curriculum learning principle, an efficient strategy in machine learning by increasing the difficulty of training gradually [52]. We propose to start the clustering with a high number of clusters and progressively reduce the number of clusters. Let’s set the starting C to be the number of training video clips. In this case, each cluster has only one video clip. By sampling a positive segment from the same utterance, we form a poor-man’s positive pair. Therefore, the MCL-DPP training actually starts from the MCL training setup.

We use a small labelled validation set to monitor the progressive clustering process, which is valid and commonly used in related studies [15], [39], [40], [49] and the competition¹². During training, the validation performance (i.e., EER) needs to be improved every epoch; otherwise, we halve the number of clusters C and repeat the clustering. As C decreases, we force more video clips into one cluster, making more diverse positive counterparts available. We note that with a large cluster, utterances from different speakers may be assigned to the same cluster, leading to low cluster purity, but more diverse positive pairs. Since we obtain these samples, As shown in Fig. 3, for an anchor video clip, we randomly sample

one positive counterpart utterance from the same cluster, and extract a positive speech segment and face image from the sampled utterance. In this way for speech modality, a positive pair is made up by two segments from two distinct utterances. On the other hand, the negative pair sampling strategy is the same to that in MCL. It is apparent that the found positive pairs in MCL-DPP are more diverse than those in MCL.

Briefly speaking, along the progressive clustering, we decrease the number of clusters to improve the diversity of the found positive counterparts, therefore, improve the quality of the speaker and face encoders. By iterating the training of encoders and the clustering steps, we gradually boost the encoders and improve the diversity of clusters. As C becomes close to the actual number of speakers, it is expected that we achieve a trade-off between the cluster purity and the diversity of the found positive pairs. It is easy to understand that when C goes below the actual number of speakers, the cluster purity deteriorates, that we should avoid. Also, our method can roughly estimate the number of speakers by capturing C at the best validation performance.

The pseudo-code of the progressive clustering algorithm is summarized in Alg. 1. We consider the progressive clustering algorithm of MCL-DPP as a single-stage learning strategy, as the learning of speaker encoder takes place end-to-end.

V. MCL-DPP WITH TWO-STAGE LEARNING

MCL-DPP is a single-stage learning framework with a focus on DPP sampling techniques. In the context of the two-stage learning as discussed in Section II-B, MCL-DPP is the single-stage learning strategy to obtain the speaker and face encoders. We are motivated to study a second stage learning to improve the resulting MCL-DPP framework. In the second stage, we perform clustering with a fixed number of speakers to obtain pseudo speaker labels and train the speaker and face encoders using these labels. The clustering-training loop is repeated iteratively in a process. As we use the speaker and face encoders from the MCL-DPP as the initial encoders, this two-stage learning is referred to as the MCL-DPP-C framework.

We next discuss how the clustering-training loop works. We first concatenate the speaker and face projected embedding to form a speech-face representation and perform a multi-modal clustering over all training data, which is called speaker clustering. As in [44], [49], the number of clusters is fixed and decided by the elbow estimation method. The resulting clusters are then used as pseudo speaker labels for supervised classification learning. In the classification learning, we train the speaker and face encoder separately by applying the *additive angular margin softmax* (AAM-softmax) loss [45]. This efficient loss function emphasises the distance between different classes and within the same class. The speaker and face recognition performance can be improved in this stage by directly learning the class distribution [43]. This clustering-training loop is repeated for several iterations. We follow a setting similar to that in [49] for a fair comparison.

We shouldn’t confuse between the progressive clustering in the first stage, which is part of the MCL-DPP algorithm, and the speaker clustering in the second stage. The progressive

¹<https://www.robots.ox.ac.uk/~vgg/data/voxceleb/competition2020.html>

²<https://www.robots.ox.ac.uk/~vgg/data/voxceleb/competition2021.html>

TABLE I

FOUR TRAINING SETS IN OUR EXPERIMENTS. VOXCELEB2 AND VOXMINI ARE WITH SPEAKER LABELS, WHILE LRS2 AND LRW ARE NOT. NEITHER MCL NOR MCL-DPP USES ANY SPEAKER IDENTITY INFORMATION FOR TRAINING.

Name	# Video clips	# Hours	# Speakers
VoxCeleb2	1,091,724	2,360	5,994
VoxMini	100,000	218	4,081
LRS2	96,318	196	N.A.
LRW	538,766	182	N.A.

clustering in the first stage seeks to obtain accurate and diverse positive pairs for contrastive learning. It doesn't seek to estimate the number of clusters. However, the speaker clustering in the second stage aims to obtain the high-quality pseudo labels as the supervisory signals, therefore, the number of clusters, i.e. the number of speakers, matters.

VI. EXPERIMENTS

A. Dataset

We train our systems on the VoxCeleb2 dataset [5], an audio-visual dataset derived from YouTube interviews and used for speaker recognition. Each video clip is of more than 4 seconds and features only one visible talking face. VoxCeleb2 contains 1,092,009 utterances from 5,994 speakers, out of which 1,091,724 utterances are accompanied by matching faces in the video.

While the VoxCeleb2 dataset is popular, its information about the number of speakers and their distributions is known. In this study, we also train our framework on two audio-visual speech recognition datasets, LRS2 [53] and LRW [54], where speaker identity information is unknown. LRS2 contains long video clips with variable length, and LRW contains short video clips with a fixed duration. These two datasets provide a perfect setup for self-supervised learning.

We also create the VoxMini dataset as a subset of VoxCeleb2 for ablation study³. To ensure a similar speaker distribution to the original VoxCeleb2 dataset, we randomly select 12.5% video clips for each person in VoxCeleb2. Furthermore, we drop the data from speakers who have less than five video clips to have the VoxMini dataset of 100,000 video clips from 4,081 speakers.

A summary of the training set can be found in Tab. I. In all datasets, the speech utterances and face frames are synchronized. We extract one face frame every 0.4 seconds. Notice that speaker identity labels are not used in training. They are only used for post-analysis. We evaluate the system on the VoxCeleb1 dataset [4] as follows:

- 1) VoxCeleb1-O (Vox-O): the original list of VoxCeleb1 with 37,720 trials from 40 speakers and 4,708 video clips.
- 2) VoxCeleb1-E (Vox-E): the extended list of VoxCeleb1 with 581,480 trials from 1,251 speakers and 145,160 video clips.
- 3) VoxCeleb1-H (Vox-H): the hard list of VoxCeleb1 with 552,536 trials from 1,190 speakers and 137,924 video clips.

³https://github.com/TaoRuijie/VoxMini/blob/main/train_mini.txt

Here, Vox-O set is used for validation, and Vox-E and Vox-H are used for testing. There is no overlap of speakers across the training, validation and test sets. All the datasets contain audio-visual data.

B. Data augmentation

We perform speech and face augmentation (mentioned in Section III-A and III-B) for contrastive learning in both MCL and MCL-DPP system, to improve the diversity of training samples, thus the robustness of speaker and face embedding.

1) *Speech augmentation*: We apply an online augmentation strategy with two datasets: RIR [55] corpus contains reverberation signals collected from the different rooms or spaces; MUSAN corpus [56] contains a variety of ambient sounds, such as nature noises (the sounds from thunder, rain, ball, etc.), background music (playing the instrument or singing) and babble (a group of people talking at the same time).

2) *Face augmentation*: Facial images are usually distracted by non-identity information, such as colour, background, and image layout. Data augmentation is effective in visual representation learning [57], [58], an adequate face augmentation technique will help the encoder to learn the individual facial characteristics. We first crop a small facial image region with a random scale from 0.4 to 1, then apply the horizontal flip, colour jitter and grey image process with a probability of 0.5, 0.8 and 0.2, respectively. Finally, we apply the randaugment [59] and Gaussian blur techniques.

C. Model

1) *Speaker encoder*: The speaker encoder is an ECAPA-TDNN [28], with a channel size of 512. The input is an 80-dimensional log mel-spectrogram from the speech segments, while the output is a 192-dimensional speaker embedding. This encoder structure is commonly used for speaker characterization due to its remarkable performance and fast training speed.

2) *Face encoder*: The 2D ResNet34 network with squeeze-and-excitation (SE) module is used as a face encoder [60]. The input is the face image, which has been reshaped into $3 \times 112 \times 112$. The channel size is set as 512, and the dimension of the output face embedding is set as 512.

3) *Projector*: We adopt a 4-layer multi-layer perceptron (MLP) [12] to project the speech and face embeddings to the common space. The speech and face projectors share the same architecture. Each layer contains a linear layer followed by a Gaussian error linear unit (GeLU). The output of each layer is of 1024, 1024, 256 and 512 dimensions, respectively. The projected embedding is L2-normalized.

D. Implementation Details

1) *Training*: We train our systems with the Adam optimizer [61]. The initial learning rate is 10^{-4} , and we decrease it by 5% for every 5 epochs. The batch size is set to 180. In MCL and MCL-DPP frameworks, we use the 2-second speech segment to train the speaker encoder and a single face image extracted from the 2-second video to train the face encoder.

During training MCL-DPP with the progressive clustering strategy, if the validation performance does not improve in the

last 3 training epochs, we halve the number of clusters C . For instance, when training on VoxCeleb2, the initial number of clusters is 1,091,724, which is the same as the number of training video clips. C is halved to 545,862 for clustering in the subsequent batch. We repeat this process to achieve the best performance on the validation set.

We further train the MCL-DPP-C by implementing a two-stage learning strategy. For a fair comparison, we follow the previous work and set the number of clusters to 6,000 on the VoxCeleb2 dataset. This number is determined by the elbow method [49]. It should be noted that the MCL-DPP with the progressive clustering strategy also reaches a similar number of clusters.

2) *Evaluation*: The test set consists of target trials, i.e., two video clips from the same person, and the imposter trials, i.e., two video clips from different persons. We extract the speaker embedding and face embedding via the encoders without involving the projectors at run-time inference. For speaker verification, we compare two trials by the cosine similarity between the two speaker embeddings. For face verification, we randomly select five faces from each video clip and generate their face embeddings. We compare two trials by the mean cosine similarity between the two set of face embeddings. For multi-modal speaker-face verification, we fuse the speech and face similarity scores for the detection decision [23]. We report all results in terms of equal error rate (EER).

VII. RESULTS AND ANALYSIS

We report our results in four subsections: Section VII-A compares our MCL-DPP and MCL-DPP-C with the corresponding SOTA self-supervised training methods. Section VII-B presents MCL-DPP results on four different training sets, including two real-world unlabelled datasets. In Section VII-C, we present a post-analysis to show that the MCL-DPP framework benefits from diverse positive pairs and study the efficiency and robustness of MCL-DPP. Finally, we discuss the other two sampling approaches in Section VIII.

A. Comparison with the state-of-the-art frameworks

1) *MCL-DPP with progressive clustering*: First, we train the MCL-DPP framework under the progressive training strategy on the VoxCeleb2 dataset with speech and face data and compare it with other single-stage self-supervised learning results. We report the evaluation for speaker verification, face verification, and speaker-face verification in Tab. II. For speaker verification, MCL-DPP achieves an EER of 2.89%, 3.34% and 6.47% on Vox-O, Vox-E and Vox-H, respectively, that outperform the best prior work, i.e., Cho et al. [40] by 40.17% on Vox-O. For face verification, it also achieves an EER of 1.74% in Vox-O. Finally, for speaker-face verification, it achieves an EER of 0.49%, 0.89% and 1.77% EER on the Vox-O, Vox-H and Vox-E, which outperforms all competitive systems by a large margin. Here we also provide the performance of MCL for comparison. MCL-DPP exhibits notable improvement due to the utilization of diverse positive pairs.

TABLE II

THE EER OF THE MCL-DPP AND MCL FRAMEWORK TRAINED ON THE VOXCELEB2 DATASET WITH BOTH SPEECH AND FACE DATA UNDER THE SINGLE-STAGE PROGRESSIVE CLUSTERING STRATEGY IN COMPARISON WITH OTHER COMPETITIVE SINGLE-STAGE SYSTEMS, WHERE ‘S’ DENOTES THE SPEECH MODALITY, ‘F’ DENOTES THE FACE MODALITY, WHILE ‘S+F’ DENOTES THE COMBINATION OF BOTH.

System	Test Modality	Vox-O	Vox-E	Vox-H
Nagrani et al. [62]	S	22.09	-	-
Chung et al. [63]	S	17.52	-	-
Inoue et al. [64]	S	15.26	-	-
Cai et al. [49]	S	8.86	10.15	16.20
Huh et al. [17]	S	8.65	-	-
Zhang et al. [16]	S	8.28	-	-
Xia et al. [38]	S	8.23	-	-
Mun et al. [36]	S	8.01	-	-
Tao et al. [44]	S	7.36	7.90	12.32
Sang et al. [37]	S	6.99	-	-
Heo et al. [41]	S	6.70	-	-
Jung et al. [42]	S	5.40	-	-
Cho et al. [40]	S	4.83	-	-
MCL	S	7.60	8.50	14.26
MCL-DPP	S	2.89	3.17	6.27
MCL	F	6.21	7.16	10.12
MCL-DPP	F	1.74	2.30	3.55
MCL	S+F	3.33	3.76	6.61
MCL-DPP	S+F	0.49	0.89	1.77

2) *MCL-DPP-C with pseudo labels*: We further train the MCL-DPP-C framework as described in Section V [49]. In Tab. III for speaker verification, MCL-DPP-C achieves an EER of 1.44%, 1.77% and 3.60% on the Vox-O, Vox-H and Vox-E, respectively. To the best of our knowledge, it is the first time the self-supervised learning speaker verification system achieves this performance. For speaker-face verification, our method achieves an EER of 0.33% on the Vox-O test set. In addition, we report the performance of the supervised speaker recognition system with the same speaker encoder and the ground-truth labels. By exploring the additional visual information, our MCL-DPP-C significantly reduces the gap between supervised and self-supervised learning (1.01% EER versus 1.44% EER). We also provide MCL-C for comparison. Here MCL-C starts from the MCL instead of MCL-DPP for the second training stage. As the results show, MCL-DPP-C performs better than MCL-C (e.g., 1.44 % VS 1.95 % EER for speaker recognition in Vox-O). Hence, we conclude that the initialization with MCL-DPP contributes to the improved performance of MCL-DPP-C.

It is noted that we did not use other optimization approaches in MCL-DPP and MCL-DPP-C, such as large-margin fine-tune [65], score normalization [66] and LGL in our previous work [44]. Therefore, there is room for improvement further.

B. MCL vs MCL-DPP across different training sets

We now train the MCL-DPP framework on four different datasets: VoxCeleb2, VoxMini, LRS2 and LRW, and report the results in Fig. 5, and compare with the baseline MCL framework. Both MCL and MCL-DPP use multi-modal information, and the difference is that MCL-DPP utilizes diverse positive pairs.

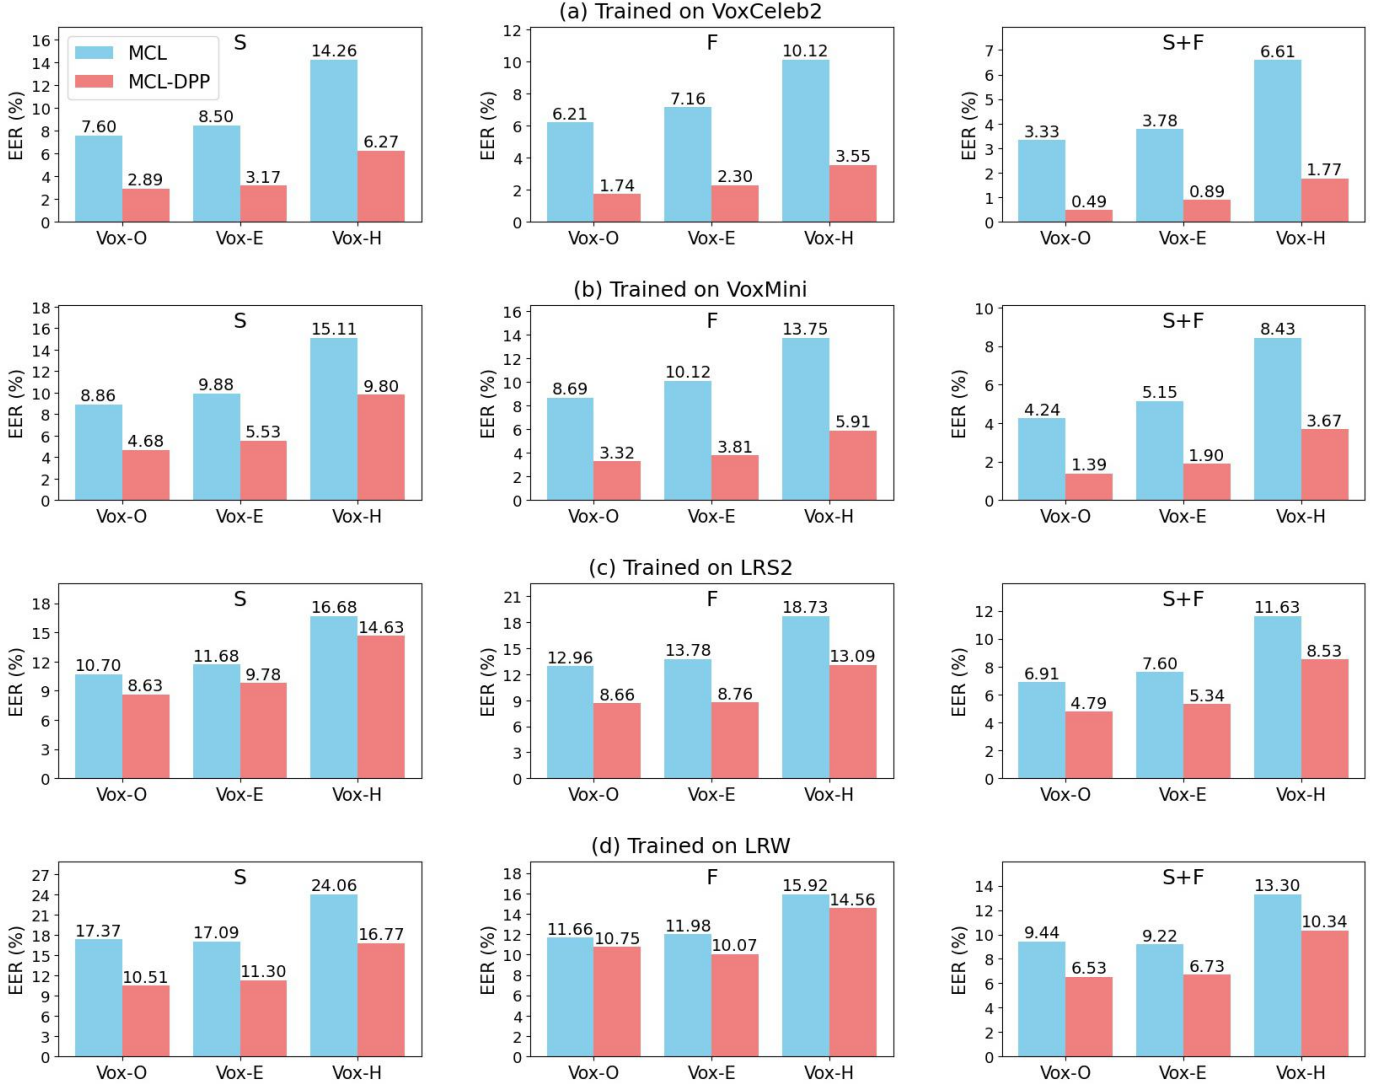


Fig. 5. Comparison of two frameworks, i.e. MCL and MCL-DPP, trained on different datasets and evaluated on Vox-O, Vox-E and Vox-H. The left column figures with ‘S’ denotes the speaker verification, the middle column figures with ‘F’ denotes the face verification, and the right column figures with ‘S+F’ denotes the combination of both for verification.

In Fig. 5, MCL trained on the VoxCeleb2 dataset achieves 7.60% EER for speaker verification on Vox-O. This result is similar to the previous contrastive learning-based results in Tab. II. The MCL-DPP framework achieves an EER of 2.89%, which confirms the contribution of the diverse positive pairs. On the small-scale dataset VoxMini, the MCL-DPP framework achieves an EER of 4.68% for speaker verification on Vox-O. It is worth noting that it even outperforms the SOTA method trained on the entire VoxCeleb2 dataset, despite that VoxMini only contains 9.2% video clips of the VoxCeleb2 dataset.

We further evaluate the MCL-DPP framework that is trained on the real-world unlabelled dataset, i.e., LRS2 and LRW, and evaluated on Vox-O, MCL achieves an EER of 10.70% and 17.37% for speaker verification, respectively. Notice that LRW contains short utterances, thus the EER is higher than that of LRS2. In comparison, MCL-DPP achieves an EER of 8.63% and 10.51% when trained on LRS2 and LRW datasets, respectively, which outperforms MCL by 19.35% and

39.49%. Similar results are observed when evaluated for face verification and speaker-face verification. All results on Vox-O, Vox-E and Vox-H suggest that the MCL-DPP framework is generic and works well for real-world unlabelled datasets.

C. Measurement of diversity

In view of the fact that VoxCeleb2 and VoxMini come with ground-truth speaker labels, we would like to study what makes MCL-DPP work. First, we define the diversity measurement in speech modality. For two utterances $x_i^{(s)}$ and $x_j^{(s)}$, we use a well pre-trained speaker encoder⁴ to extract the speaker embedding $y_i^{(s)}$ and $y_j^{(s)}$, then the diversity of these two utterances is defined by their L_2 distance $d = L_2(y_i^{(s)}, y_j^{(s)})$. The greater the L_2 or d is, the more distant the two utterances are. Here cosine distance is also valid but will result in an inverse relationship. With L_2 distance, a poor-man’s positive

⁴<https://github.com/TaoRuijie/ECAPA-TDNN>

TABLE III

THE EER OF THE MCL-DPP-C AND MCL-C FRAMEWORK TRAINED ON THE VOXCELEB2 DATASET WITH BOTH SPEECH AND FACE DATA UNDER THE TWO-STAGE LEARNING STRATEGY WITH PSEUDO LABELS IN COMPARISON WITH OTHER TWO-STAGE SYSTEMS INCLUDING THE SUPERVISED LEARNING COUNTERPART. ‘S’ DENOTES THE SPEECH MODALITY, ‘F’ DENOTES THE FACE MODALITY, AND ‘S+F’ DENOTES THE COMBINATION OF BOTH.

System	Test Modality	Vox-O	Vox-E	Vox-H
Supervised [67]	S	1.01	1.26	2.43
Cai et al. [15]	S	3.45	4.02	6.57
Cho et al. [40]	S	2.10	-	-
Cai et al. [49]	S	1.81	2.06	3.80
Tao et al. [44]	S	1.66	2.18	3.76
MCL-C	S	1.95	2.16	3.91
MCL-DPP-C	S	1.44	1.77	3.27
MCL-C	F	2.15	3.04	4.33
MCL-DPP-C	F	1.49	1.88	2.76
MCL-C	S+F	0.59	0.71	1.28
MCL-DPP-C	S+F	0.33	0.43	0.82

pair comes from the same utterance, thus leading to a low d . Then we define D as the average diversity for all the available positive pairs in the training set. Then N_+ represents the mean number of the available positive pairs for each sample.

D. Ablation Studies

In this section, we train MCL and MCL-DPP with different conditions on the VoxMini dataset to study how MCL-DPP works.

1) *Effect by the diverse positive pairs*: To show the effect of the diversity of positive pairs, we train the MCL framework with speech modality only (as seen in the orange box of Fig. 2), on the VoxMini dataset, by following four sampling strategies for positive counterparts. We always sample the positive counterparts from the correct target speakers to eliminate the impact of the false-positives to study the effect of diversity since the VoxMini dataset has ground-truth speaker labels.

C1: To select a positive counterpart from the anchor utterance itself, i.e., MCL (w/o face).

C2: To select a low diversity counterpart utterance for each anchor utterance and fix it across all training epochs.

C3: To select a high diversity counterpart utterance for each anchor utterance and fix it across all training epochs.

C4: To randomly select a positive counterpart from epoch to epoch.

In the four strategies, C1 basically samples the poor-man’s positive pairs. C2 and C3 samples high diverse positive counterparts than C1, but at two extremes. C4 offers the highest level of diversity among the found positive counterparts. We plot the EER for speaker verification on the Vox-O validation set as a function of the number of training epochs in Fig. 6. Comparing C1 and C4, we find that using diverse positive pairs greatly improves the results by a large margin. While C2 and C3 lead to the same number of found positive pairs, C3 performs much better than C2 because of a higher diver-

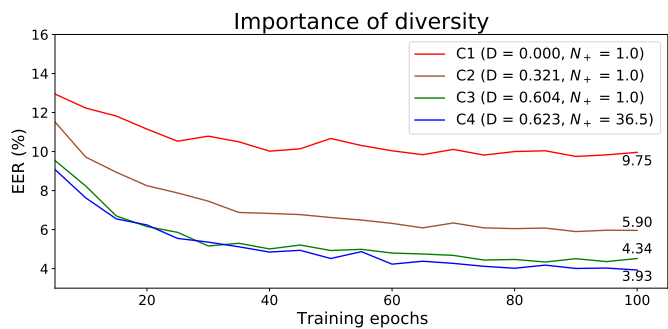


Fig. 6. Speaker verification EER curve on the Vox-O during training on the VoxMini dataset for C1 to C4. D is the average diversity for the available positive pairs, N_+ is the average number of available positive pairs per sample.

TABLE IV

COMPARISON BETWEEN MCL AND MCL-DPP THAT TRAINED WITH SINGLE OR MULTIPLE MODALITIES. (S) INDICATES TRAINING (AND CLUSTERING) WITH SPEECH MODALITY, (F) WITH FACE MODALITY, AND (S+F) WITH TWO MODALITIES TOGETHER. THE SPEAKER AND FACE VERIFICATION RESULT ARE EVALUATED ON VOX-O SET.

Index	Method	EER(%)	
		Speech	Face
S1	MCL (S)	9.64	-
S2	MCL-DPP (S)	8.81	-
S3	MCL (F)	-	10.64
S4	MCL-DPP (F)	-	9.93
S5	MCL (S+F)	8.86	8.69
S6	MCL-DPP (S+F)	4.68	3.32

sity. All results indicate that high diversity improves, which provides the theoretical basis for our MCL-DPP framework.

2) *Diversity vs quality of positive pairs*: Then we study how MCL-DPP works by reporting the change of diversity as a function of the number of clusters, the mean accuracy of found positive pairs during the progressive clustering of MCL-DPP training on the VoxMini dataset in Fig. 7. We also report the speaker verification EER curve on Vox-O.

In Fig. 7(a), we observe that the D increases from 0.00 to 0.37, as the number of clusters decreases from 100K to 12.5K. With an appropriate number of clusters, we achieve high accuracy for the found positive pairs. For instance, 89.91% accuracy for 12.5K clusters in Fig. 7(b). The contrastive learning benefits from the correct and diverse positive pairs, leading to a decrease of EER as seen in Fig. 7(c). However, when the number of clusters becomes smaller than the actual number of speakers (i.e., 4.1K speakers) in the training set, the accuracy of the found positive pairs drops. For example, we observe a 40.47% accuracy for 1.5K clusters. With a trade-off between diversity and accuracy, the MCL-DPP framework achieves the best performance. For example, when we reach 3.1K clusters in the progressive clustering, we observe a low EER of 4.68.

The experiments confirm the idea of diverse positive pairs and acknowledge that there is a trade-off between diversity and accuracy that we need to observe. We also validate the effectiveness of the proposed progressive clustering strategy.

3) *Multi-modal contrastive learning*: In this paper, we use MCL-DPP to refer to the general multi-modal MCL-DPP

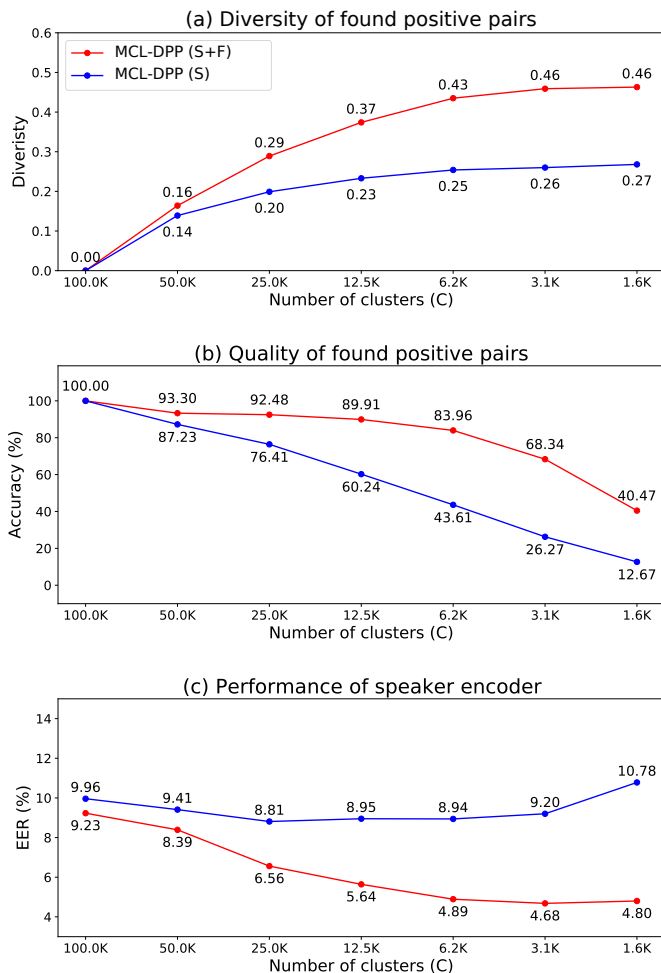


Fig. 7. The diversity, accuracy of found positive speech pairs, and the resulting EER of the MCL-DPP frameworks trained with speech modality alone and with two modalities together. The systems are trained on the VoxMini dataset. In the progressive clustering, the number of clusters C decreases during training, (a) the mean diversity D for all found positive pairs. (b) the mean accuracy for all found positive pairs. (c) the speaker verification EER on Vox-O.

framework, of which the encoders are trained on both speech segments and their accompanying face data. We believe that the face modality contributes to finding correct and diverse positive speech pairs, and vice versa. To verify that, we repeat the experiments for MCL-DPP and MCL training on the VoxMini with speech or face modality only. Here no projector is required and the clustering for MCL-DPP is based on the speaker or face embedding only. The other settings remain the same as those in the MCL-DPP with progressive clustering training.

We report the speaker or face verification results on Vox-O in Tab. IV. The usage of the cross-modal loss in the MCL framework resulted in an improvement in the EER for both speaker verification (comparing S1 to S5, 9.64% to 8.86%) and face verification (comparing S3 to S5, 10.64% to 8.69%). This demonstrates the effectiveness of the cross-modal loss outlined in Section III-C in enhancing the speaker or face representation. Meanwhile, when training with speech modality alone, using diverse positive pairs improves the EER



Fig. 8. Comparison of the found positive face pairs (a) with MCL-DPP (F), (b) with MCL-DPP (S+F). The faces in the first column are the anchor faces, while the rest are the found positive faces.

from 9.64% to 8.81% (S1 versus S2). The improvement is not as great as that in the MCL-DPP framework (S5 versus S6, from 8.86% to 4.68%). Similar findings have been observed in the experiments of face modality (S3 versus S4). These experiments verify that the MCL-DPP framework discovers diverse positive pairs by multi-model complementarity.

To understand how the MCL-DPP benefits from audio-visual information to improve the diversity, we compare the diversity of the found positive speech pairs in Fig. 7(a) between MCL-DPP (S) and MCL-DPP (S+F). Apparently, MCL-DPP (S) doesn't find the positive pairs as diverse as the MCL-DPP does. Furthermore, Fig. 7(b) shows that the progressive clustering with speech modality alone leads to low accuracy for the found positive pairs. The MCL-DPP (S) suffers from both low diversity and low accuracy, leading to worse EER. That proves multi-modality is important for robust clustering.

To explain the diversity difference between single-modal and multi-modal systems, we compare the found positive faces in MCL-DPP (F) with those of the MCL-DPP (S+F) in Fig. 8. It is noted that, the MCL-DPP (F) framework discovers the found positive faces that are similar to the anchor face (see Fig. 8(a)), while the general MCL-DPP (S+F) offers much more diverse found positive faces (see Fig. 8(b)), with the assistance of audio-visual interaction.

4) Initial number of speaker clusters: We study the effect of the initial number of clusters for the progressive clustering in the MCL-DPP training on the VoxMini dataset by setting initial C to 100K, 25.6K and 1.6K. 100K is the rough number of video clips in the dataset. We know the actual number of speakers is around 4.1K in the dataset, but we don't use the speaker information in training.

It is not a coincidence that in Fig. 9, for the initial C of 100K and 25.6K, both systems discover the number of clusters close to 4.1K. This suggests that the progressive clustering in the general MCL-DPP framework accurately estimates the number of speakers. While we don't seek to find the exact number of speakers, a near exact number guarantees high accuracy of found positive pairs. We also note that when the initial number of clusters goes below the actual number of speakers (such as 1.6k), MCL-DPP doesn't perform because the accuracy of found positive speech pairs is very low.

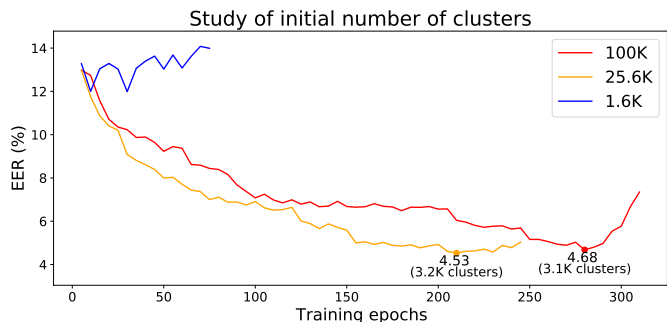


Fig. 9. MCL-DPP with the VoxMini for the different initial number of clusters: (a) 100K, (b) 25.6K, and (c) 1.6K. EER is the speaker verification performance that evaluates on Vox-O.

VIII. FUTURE WORK: OTHER DPP SAMPLING ALTERNATIVES

In this paper, we have shown that the self-supervised speaker recognition system benefits from diverse positive pairs, and studied the progressive clustering strategy. This study lays the foundation for more sampling techniques in the same direction. Next, we discuss two possible alternatives.

1) *Sampling the K -nearest neighbours*: One of the alternatives is to search for the K most similar video clips in the projected domain as the positive video clips, given an anchor video clip x_a as the query. We call this technique the *K -nearest neighbours*. We denote a found video clip as the positive counterpart x_p ,

$$x_p \in \{x_i \mid z_i \in \text{kn}(z_a)\}, i \in [1, N] \quad (6)$$

N is the total number of training video clips, z_a is the projected embedding of anchor video clip x_a .

The question is how to select a suitable K . A small K may favour the search precision, while a large K may increase the number of the false-positives. Furthermore, there could be a different number of positive counterparts available for each anchor video clip.

2) *Sampling within a threshold*: Another alternative is to set an absolute similarity threshold T in the projected domain. The video clips of cosine similarity higher than the threshold are considered positive counterparts x_p :

$$x_p \in \{x_i \mid \cos(z_a, z_i) > T\}, i \in [1, N] \quad (7)$$

However, a threshold may vary with the performance of the speaker and face encoders during the training process and require careful calibration.

The two alternatives among others will be the topics of future study.

IX. CONCLUSION

In this paper, we confirm that accurate and diverse positive pairs lead to effective self-supervised contrastive learning. We propose a positive pairs sampling approach to train the encoders without speaker labels. The visual information is also involved to guarantee the accuracy and diversity of found positive pairs. We show that diversity plays an important role in MCL-DPP and improves the results by a large margin.

In future work, we plan to leverage a large amount of real-world unlabeled data in the wild to show the potential of self-supervised learning.

REFERENCES

- [1] T. Kinnunen and H. Li, "An overview of text-independent speaker recognition: From features to supervectors," *Speech Communication*, vol. 52, pp. 12 – 40, 2010.
- [2] K. A. Lee, O. Sadjadi, H. Li, and D. Reynolds, "Two decades into speaker recognition evaluation - are we there yet?" *Computer Speech & Language*, vol. 61, p. 101058, 2020.
- [3] R. K. Das and S. R. M. Prasanna, "Investigating text-independent speaker verification from practically realizable system perspective," in *APSIPA ASC*, 2018, pp. 1483–1487.
- [4] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Interspeech*, 2017, pp. 2616–2620.
- [5] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Interspeech*, 2018, pp. 1086–1090.
- [6] M. McLaren, L. Ferrer, D. Castán, and A. Lawson, "The speakers in the wild (SITW) speaker recognition database," in *Interspeech*, 2016, pp. 818–822.
- [7] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning (ICML)*, 2020, pp. 1597–1607.
- [8] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, June 2019, pp. 4171–4186.
- [10] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 9729–9738.
- [11] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar, *et al.*, "Bootstrap your own latent—a new approach to self-supervised learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 21 271–21 284, 2020.
- [12] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 9650–9660.
- [13] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Un-supervised pre-training for speech recognition," in *Interspeech 2019*, G. Kubin and Z. Kacic, Eds. ISCA, 2019, pp. 3465–3469.
- [14] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [15] D. Cai, W. Wang, and M. Li, "An iterative framework for self-supervised deep speaker representation learning," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6728–6732.
- [16] H. Zhang, Y. Zou, and H. Wang, "Contrastive self-supervised learning for text-independent speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6713–6717.
- [17] J. Huh, H. S. Heo, J. Kang, S. Watanabe, and J. S. Chung, "Augmentation adversarial training for unsupervised speaker recognition," in *Workshop on Self-Supervised Learning for Speech and Audio Processing, NeurIPS*, 2020.
- [18] 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.
- [19] D. Dwivedi, Y. Aytar, J. Tompson, P. Sermanet, and A. Zisserman, "With a little help from my friends: Nearest-neighbor contrastive learning of visual representations," in *IEEE/CVF International Conference on Computer Vision (CVPR)*, 2021, pp. 9588–9597.
- [20] W. Van Gansbeke, S. Vandenhende, S. Georgoulis, M. Proesmans, and L. Van Gool, "SCAN: Learning to classify images without labels," in *European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 268–285.

- [21] T. Han, W. Xie, and A. Zisserman, "Self-supervised co-training for video representation learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5679–5690, 2020.
- [22] Y. Qian, Z. Chen, and S. Wang, "Audio-visual deep neural network for robust person verification," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1079–1092, 2021.
- [23] R. K. Das, R. Tao, J. Yang, W. Rao, C. Yu, and H. Li, "HLT-NUS submission for 2019 NIST multimedia speaker recognition evaluation," in *APSIPA ASC*, 2020, pp. 605–609.
- [24] A. Blum and T. Mitchell, "Combining labeled and unlabeled data with co-training," in *Proceedings of the eleventh annual conference on Computational learning theory*, 1998, pp. 92–100.
- [25] J. Slavíček, A. Swart, M. Klčo, and N. Brümmner, "The phonexia voxceleb speaker recognition challenge 2021 system description," *arXiv preprint arXiv:2109.02052*, 2021.
- [26] Z. Bai and X.-L. Zhang, "Speaker recognition based on deep learning: An overview," *Neural Networks*, vol. 140, pp. 65–99, 2021.
- [27] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: robust DNN embeddings for speaker recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5329–5333.
- [28] 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.
- [29] 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 transactions on pattern analysis and machine intelligence*, vol. 43, no. 2, pp. 652–662, 2019.
- [30] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *IEEE conference on computer vision and pattern recognition (CVPR)*, 2018, pp. 7132–7141.
- [31] S. Prince and J. H. Elder, "Probabilistic linear discriminant analysis for inferences about identity," in *In IEEE International Conference on Computer Vision (ICCV)*, Oct 2007, pp. 1–8.
- [32] I. Medennikov, M. Korenevsky, T. Prisyach, Y. Y. Khokhlov, M. Korenevskaya, I. Sorokin, T. Timofeeva, A. Mitrofanov, A. Andrusenko, I. Podluzhny, A. Laptev, and A. Romanenko, "Target-speaker voice activity detection: A novel approach for multi-speaker diarization in a dinner party scenario," in *Interspeech 2020*, 2020, pp. 274–278.
- [33] C. Xu, W. Rao, E. S. Chng, and H. Li, "SpEx: Multi-scale time domain speaker extraction network," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 28, pp. 1370–1384, 2020.
- [34] M. Chen, M. Chen, S. Liang, J. Ma, L. Chen, S. Wang, and J. Xiao, "Cross-lingual, multi-speaker text-to-speech synthesis using neural speaker embedding," in *Interspeech*, 2019, pp. 2105–2109.
- [35] Y. Zhou, X. Tian, R. K. Das, and H. Li, "Many-to-many cross-lingual voice conversion with a jointly trained speaker embedding network," in *APSIPA ASC*, 2019, pp. 1282–1287.
- [36] S. H. Mun, W. H. Kang, M. H. Han, and N. S. Kim, "Unsupervised representation learning for speaker recognition via contrastive equilibrium learning," *arXiv preprint arXiv:2010.11433*, 2020.
- [37] M. Sang, H. Li, F. Liu, A. O. Arnold, and L. Wan, "Self-supervised speaker verification with simple siamese network and self-supervised regularization," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6127–6131.
- [38] W. Xia, C. Zhang, C. Weng, M. Yu, and D. Yu, "Self-supervised text-independent speaker verification using prototypical momentum contrastive learning," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6723–6727.
- [39] J. Thienpondt, B. Desplanques, and K. Demuynck, "The IDLAB Vox-Celeb Speaker Recognition Challenge 2020 system description," *arXiv preprint arXiv:2010.12468*, 2020.
- [40] J. Cho, J. Villalba, L. Moro-Velazquez, and N. Dehak, "Non-contrastive self-supervised learning for utterance-level information extraction from speech," *IEEE Journal of Selected Topics in Signal Processing*, pp. 1–12, 2022.
- [41] H.-S. Heo, J.-w. Jung, J. Kang, Y. Kwon, Y. J. Kim, and B.-J. Lee, "Self-supervised curriculum learning for speaker verification," *arXiv preprint arXiv:2203.14525*, 2022.
- [42] J.-w. Jung, Y. J. Kim, H.-S. Heo, B.-J. Lee, Y. Kwon, and J. S. Chung, "Raw waveform speaker verification for supervised and self-supervised learning," *arXiv preprint arXiv:2203.08488*, 2022.
- [43] A. Brown, J. Huh, J. S. Chung, A. Nagrani, and A. Zisserman, "VoxSRC 2021: The third voxceleb speaker recognition challenge," *arXiv preprint arXiv:2201.04583*, 2022.
- [44] R. Tao, K. A. Lee, R. K. Das, V. Hautamäki, and H. Li, "Self-supervised speaker recognition with loss-gated learning," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6142–6146.
- [45] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4690–4699.
- [46] L. Sari, K. Singh, J. Zhou, L. Torresani, N. Singhal, and Y. Saraf, "A multi-view approach to audio-visual speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6194–6198.
- [47] G. Sell, K. Duh, D. Snyder, D. Etter, and D. Garcia-Romero, "Audio-visual person recognition in multimedia data from the IARPA Janus program," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 3031–3035.
- [48] S. Shon, T.-H. Oh, and J. Glass, "Noise-tolerant audio-visual online person verification using an attention-based neural network fusion," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3995–3999.
- [49] D. Cai, W. Wang, and M. Li, "Incorporating visual information in audio based self-supervised speaker recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2022.
- [50] R. Tao, Z. Pan, R. K. Das, X. Qian, M. Z. Shou, and H. Li, "Is someone speaking? exploring long-term temporal features for audio-visual active speaker detection," in *Proceedings of the 29th ACM International Conference on Multimedia (ACM MM)*, 2021, p. 3927–3935.
- [51] S. Lloyd, "Least squares quantization in PCM," *IEEE transactions on information theory*, vol. 28, no. 2, pp. 129–137, 1982.
- [52] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 41–48.
- [53] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, "Deep audio-visual speech recognition," *IEEE transactions on pattern analysis and machine intelligence (TPAMI)*, 2018.
- [54] J. S. Chung and A. Zisserman, "Lip reading in the wild," in *Asian conference on computer vision (ACCV)*. Springer, 2016, pp. 87–103.
- [55] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 5220–5224.
- [56] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *ArXiv*, vol. abs/1510.08484, 2015.
- [57] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. A. Raffel, "Mixmatch: A holistic approach to semi-supervised learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [58] K. Sohn, D. Berthelot, N. Carlini, Z. Zhang, H. Zhang, C. A. Raffel, E. D. Cubuk, A. Kurakin, and C.-L. Li, "Fixmatch: Simplifying semi-supervised learning with consistency and confidence," *Advances in Neural Information Processing Systems*, vol. 33, pp. 596–608, 2020.
- [59] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, "RandAugment: Practical automated data augmentation with a reduced search space," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPR)*, 2020, pp. 702–703.
- [60] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE conference on computer vision and pattern recognition (CVPR)*, 2016, pp. 770–778.
- [61] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 2015.
- [62] A. Nagrani, J. S. Chung, S. Albanie, and A. Zisserman, "Disentangled speech embeddings using cross-modal self-supervision," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6829–6833.
- [63] S.-W. Chung, H.-G. Kang, and J. S. Chung, "Seeing voices and hearing voices: Learning discriminative embeddings using cross-modal self-supervision," in *Interspeech*, 2020, pp. 3486–3490.
- [64] N. Inoue and K. Goto, "Semi-supervised contrastive learning with generalized contrastive loss and its application to speaker recognition," in *APSIPA ASC*, 2020, pp. 1641–1646.
- [65] J. Thienpondt, B. Desplanques, and K. Demuynck, "Integrating frequency translational invariance in TDNNs and frequency positional information in 2D ResNets to enhance speaker verification," in *Interspeech*, 2021.
- [66] P. Matejka, O. Novotný, O. Ploch, L. Burget, M. D. Sánchez, and J. Cernocký, "Analysis of score normalization in multilingual speaker recognition," in *Interspeech*, 2017, pp. 1567–1571.
- [67] R. K. Das, R. Tao, and H. Li, "HLT-NUS submission for 2020 NIST conversational telephone speech SRE," *arXiv preprint arXiv:2111.06671*, 2021.



Ruijie Tao (Student Member, IEEE) received the B.Eng. degree from Soochow University, China, in 2018. He received the M.Sc. degree from National University of Singapore, Singapore, in 2019. He is currently a Ph.D. candidate at National University of Singapore, Singapore. He is also the reviewer of CVPR, ICASSP, SPL, OJSP, ISCSLP. His research interests include speaker recognition, speech extraction, audio-visual representation learning.



Kong Aik Lee (Senior Member, IEEE) is currently a Senior Scientist and a Group Leader at the Agency for Science, Technology and Research (A*STAR), Singapore. He was a Senior Principal Researcher at the Data Science Research Laboratories, NEC Corporation, Japan, from 2018 to 2020. He received his Ph.D. degree from Nanyang Technological University, Singapore, in 2006. After which he joined the Institute for Infocomm Research, Singapore, as a Research Scientist and then a Strategic Planning Manager (concurrent appointment). He was the recipient of the Singapore IES Prestigious Engineering Achievement Award 2013 for his contribution to voice biometrics technology, and the Outstanding Service Award by IEEE ICME 2020. Currently, he serves as an Editorial Board Member for Elsevier Computer Speech and Language (2016 - present) and was an Associate Editor for IEEE/ACM Transactions on Audio, Speech, and Language Processing (2017 - 2021). He is an elected member of the IEEE Speech and Language Processing Technical Committee and was the General Chair of the Speaker Odyssey 2020 Workshop. His research focuses on the automatic and para-linguistic analysis of speaker characteristics, ranging from speaker recognition, language and accent recognition, diarization, voice biometrics, spoofing, and countermeasure.

recipient of the Singapore IES Prestigious Engineering Achievement Award 2013 for his contribution to voice biometrics technology, and the Outstanding Service Award by IEEE ICME 2020. Currently, he serves as an Editorial Board Member for Elsevier Computer Speech and Language (2016 - present) and was an Associate Editor for IEEE/ACM Transactions on Audio, Speech, and Language Processing (2017 - 2021). He is an elected member of the IEEE Speech and Language Processing Technical Committee and was the General Chair of the Speaker Odyssey 2020 Workshop. His research focuses on the automatic and para-linguistic analysis of speaker characteristics, ranging from speaker recognition, language and accent recognition, diarization, voice biometrics, spoofing, and countermeasure.



Rohan Kumar Das (Senior Member, IEEE) received Ph.D degree from Indian Institute of Technology (IIT) Guwahati in the year 2017 and B. Tech degree in Electronics and Communication Engineering from North-Eastern Hill University (NEHU), Shillong, India in the year 2010. His Ph.D. work focused on speaker verification using short utterances from the perspective of practical application oriented systems. Prior to his research in the field of speech processing, he has been a Project Scientist at Assam Science Technology and Environment Council

(ASTEC) from 2010 to 2011. After completing doctoral studies he worked as a Data Scientist in a multinational company called Kovid Research Labs and involved in speech analytics based application services in 2017. Later that year, he joined Human Language Technology Laboratory, National University of Singapore and as a Research Fellow and conducted postdoctoral research till March 2021. Currently, he is working as a Research and Development (R & D) Manager at Fortemedia, Singapore division to lead the research and product development unit. He was one of the organizers for special sessions on “The Attacker’s Perspective on Automatic Speaker Verification”, “Far-Field Speaker Verification Challenge 2020” in Interspeech 2020, and the Voice Conversion Challenge 2020. He served as Publication Chair of IEEE Automatic Speech Recognition Understanding (ASRU) Workshop 2019, one of the Chairs of Joint Workshop for the Blizzard Challenge and Voice Conversion Challenge 2020 and one of the Special Session Chairs of International Symposium on Chinese Spoken Language Processing (ISCSLP) 2022. He has been awarded several travel fellowships from organizations such as IEEE Signal Processing Society, International Speech Communication Association, Microsoft Research India, Xerox Research Centre India and Science and Engineering Research Board (SERB), Government of India to present research works in top tier conferences such as ICASSP and Interspeech. His research interests are speech signal processing, speaker verification, anti-spoofing, machine learning and pattern recognition.



Ville Hautamäki (Member, IEEE) received the M.Sc. and Ph.D. degrees in computer science from the University of Joensuu (currently known as the University of Eastern Finland), Joensuu, Finland, in 2005 and 2008, respectively. From 2009 to 2011, he was a Research Fellow with the Institute for Infocomm Research, A*STAR, Singapore. From 2011 to 2014, he was a Principal Investigator (PI) with the Academy of Finland Funded Project, University of Eastern Finland, funded by Academy of Finland.

In 2013, he was a Visiting Scholar for six months with the Georgia Institute of Technology, Atlanta, GA, USA. He is currently the Associate Professor with the University of Eastern Finland. In 2014, he was the recipient of the Nokia Visiting Professor grant to again visit Georgia Institute of Technology. He served in the COST Action IC1206. He is also an Associate Editor for Digital Speech Processing and IEEE Signal Processing Letters, and a Member of the IEEE Speech and Language Processing Technical Committee (SLTC).



Haizhou Li (Fellow, IEEE) received the B.Sc., M.Sc., and Ph.D. degrees in electrical and electronic engineering from the South China University of Technology, Guangzhou, China, in 1984, 1987, and 1990 respectively. He is currently a Presidential Chair Professor with the School of Data Science, The Chinese University of Hong Kong, Shenzhen, China. He is also an Adjunct Professor with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore. Prior to that, he taught with the University of Hong

Kong, Hong Kong, (1988–1990) and South China University of Technology, (1990–1994). He was a Visiting Professor with CRIN in France (1994–1995), Research Manager with the AppleSS Research Centre (1996–1998), Research Director with Lernout & Hauspie Asia Pacific (1999–2001), a Vice President with InfoTalk Corp. Ltd. (2001–2003), and the Principal Scientist and Department Head of Human Language Technology with the Institute for Infocomm Research, Singapore (2003–2016). His research interests include automatic speech recognition, speaker and language recognition, natural language processing. Dr. Li was an Editor-in-Chief of IEEE/ACM Transactions on Audio, Speech and Language Processing (2015–2018), a Member of the Editorial Board of Computer Speech and Language since 2012, an elected Member of IEEE Speech and Language Processing Technical Committee (2013–2015), the President of the International Speech Communication Association (2015–2017), the President of Asia Pacific Signal and Information Processing Association (2015–2016), and the President of Asian Federation of Natural Language Processing (2017–2018). He was the General Chair of ACL 2012, INTERSPEECH 2014, ASRU 2019 and ICASSP 2022. Dr. Li is a Fellow of the ISCA, and a Fellow of the Academy of Engineering Singapore. He was the recipient of the National Infocomm Award 2002, and the President’s Technology Award 2013 in Singapore. He was named one of the two Nokia Visiting Professors in 2009 by the Nokia Foundation, and U Bremen Excellence Chair Professor in 2019.