

# Unsupervised Voice-Face Representation Learning by Cross-Modal Prototype Contrast

Boqing Zhu, Kele Xu\*, Changjian Wang, Zheng Qin, Tao Sun, Huaimin Wang and Yuxing Peng

College of Computer,  
National University of Defense Technology, Changsha, China  
{zhuboq, kelele.xu}@gmail.com

## Abstract

We present an approach to learn voice-face representations from the talking face videos, without any identity labels. Previous works employ cross-modal instance discrimination tasks to establish the correlation of voice and face. These methods neglect the semantic content of different videos, introducing false-negative pairs as training noise. Furthermore, the positive pairs are constructed based on the natural correlation between audio clips and visual frames. However, this correlation might be weak or inaccurate in a large amount of real-world data, which leads to deviating positives into the contrastive paradigm. To address these issues, we propose the cross-modal prototype contrastive learning (CMPC), which takes advantage of contrastive methods and resists adverse effects of false negatives and deviate positives. On one hand, CMPC could learn the intra-class invariance by constructing semantic-wise positives via unsupervised clustering in different modalities. On the other hand, by comparing the similarities of cross-modal instances from that of cross-modal prototypes, we dynamically recalibrate the unlearnable instances' contribution to overall loss. Experiments show that the proposed approach outperforms state-of-the-art unsupervised methods on various voice-face association evaluation protocols. Additionally, in the low-shot supervision setting, our method also has a significant improvement compared to previous instance-wise contrastive learning.

## 1 Introduction

Neuroimage studies [Joassin *et al.*, 2011] suggest that the interactions between voice and face are not only at a perceptual level, but also in the person recognition process. Intuitively, when hearing a person's voice, even if we don't see the person, we immediately have its appearance in our minds. Even for a stranger we have never seen before, we could 'construct' what he looks like through his voice. While, whether voice and face are relevant and whether they can be perceived by

\*Corresponding author

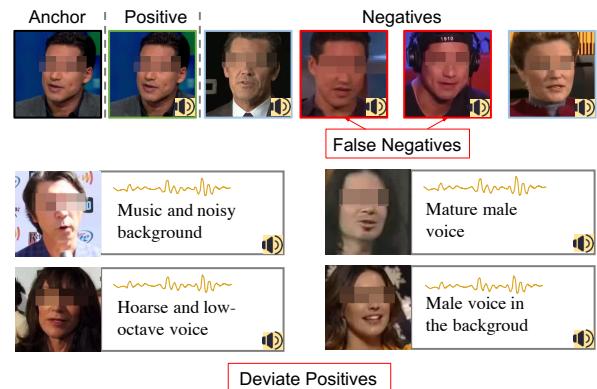


Figure 1: False Negatives and Deviate Positives examples in Cross-modal Instance Discrimination. (1) Without identity labels, negative pairs from different videos could actually belong to the same identity, such as the anchor face and negative voice in red border, which leads to false negatives. (2) The voice-face correlation might be personalized or irrelevant caused by the noise in the videos, which leads to deviate positives. The voice modality is displayed by marking a trumpet symbol in the image.

the computational models? In recent years, researchers made a sustainable effort on an enduring neural network the ability to recognize faces with the same identity only by hearing the voice or recognizing the voice from the face. It could be formulated as a voice-face association problem. The problem is an interesting scientific study on its own, and it could also benefit practical applications such as virtual reality, criminal investigation, multimodal information retrieval, *etc.*

Psychological and behavioral studies show that voice and face integration skills could emerge very early in our life by processing a number of auditory-visual examples [Bahrack *et al.*, 2005], without intended discipline. This indicates that our computational models may be feasible to learn the association between voice and face in an unsupervised manner. More recently, contrastive learning has become a prominent tool in unsupervised representation learning, leading to state-of-the-art results in many fields such as vision, language, and speech. A rock line of papers [Morgado *et al.*, 2021b; Zolfaghari *et al.*, 2021; Morgado *et al.*, 2021a; Alayrac *et al.*, 2020] have employed the *Cross-modal Instances Discrimination* (CID) method to learn the multi-modal representations.

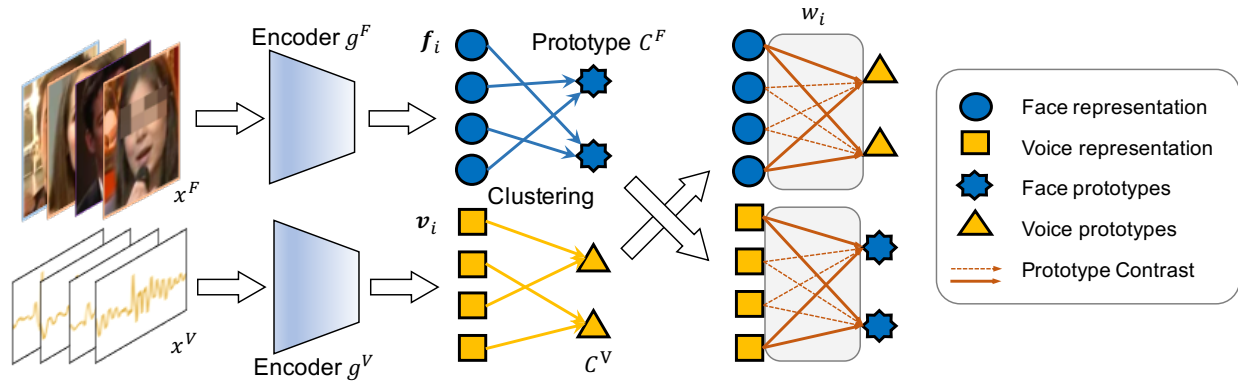


Figure 2: **Overview of the proposed method.** Face frames and voice clips are encoded into corresponding representations. To address the false negatives, clustering is conducted separately on different modalities to learn the prototypes of each instance, aiming to construct semantic-wise positives. Deviate positives are discounted by down-weighting instances.

In the CID methods, the representations are learned by two opposite forces, the repellence of negative pairs and the attraction of positive pairs. The positive pairs are generated by the audio and images extracted from the same video, and the negative pairs are those extracted from different videos. Then, a contrastive loss (e.g. InfoNCE [Oord *et al.*, 2018]) is used to pull close the positive pairs and push apart the negative pairs. Despite their improving performance, the CID methods are still confronted with two challenges: (1) False negative problem. The negative pairs are sampled from different videos regardless of their semantics. For example, two videos have the same *identity* (person), the representations of the voice and face should be pulled close, but they have undesirably pushed apart because they are sampled as negative pairs, as in Figure 1. Since the current CID methods need to utilize a large number of negatives, the problem of the false negative is non-ignorable. (2) Deviate positive problem. Previous methods have learned the representations by encouraging the network to memorize all the voice-face associations in the dataset. While they ignore the biometric specificity across different identities. Some identities hold an ‘untypical’ or personal voice face correspondence. Moreover, a large amount of video data inevitably contain noise, including blocked faces, excessive head rotation, background music, the voice from outside of the screen, *etc.* We named them ‘deviate positives’ instead of ‘false positives’, because they are all reasonable existing, and to distinguish from supervised learning where there are ground-truth labels.

We firstly reveal the impact of the false negatives and deviate positives in Section 3.1 from a theoretical view. To combat the above issues, we then propose a Cross-modal Prototype Contrastive Learning (CMPC) algorithm, which takes advantage of previous contrastive methods meanwhile resisting adverse effects of false negatives and deviate positives. Specifically, we obtain the semantic structure by clustering on each modality, and each cluster centroid is used as the prototype of the corresponding instance. The learned prototypes are used to construct the semantic positives with the instance in another modality. We use an online clustering approach, which keeps the consistency of instance representations and

reduces computation. Additionally, we dynamically down-weight the deviate positives’ contribution to overall loss by an Instance Recalibrate method. By combining our methods, we can outperform previous state-of-the-art unsupervised methods on various voice-face association evaluation protocols. Moreover, with low-shot supervision, our method achieves significant improvement compared to previous instance-wise contrastive learning, even outperforms fully supervised models.

The rest of this paper is organized as follows. We first introduce the related work in the following Section 2, then we analyze the problems and present the details of the proposed approach in Section 3. The experiments are given in the Section 4. Finally, we conclude our paper in Section 6.

## 2 Related Work

### 2.1 Learning Association of Voice and Face

In recent years, the interests for the human voice-face cross-modal learning tasks have surged, and the attempts have shown the feasibility of recognizing a person’s face from his voice or identifying a speaker only by giving his face image. SVHF [Nagrani *et al.*, 2018b] firstly raised voice-face matching and retrieval task. They concatenate the feature from different modalities and then learn a  $N$ -way classification task by minimizing a cross-entropy loss on the personal identity label. Instead of explicitly learning the joint embedding between voice and face, DIMNet [Wen *et al.*, 2019a] learns a shared representation under supervision from one or more covariates, such as gender, ethnicity, and age, by a multi-task object function. [Horiguchi *et al.*, 2018] and [Kim *et al.*, 2018] use  $N$ -pair loss and triplet loss to increase the voice-face similarity of the same person respectively. In [Chung *et al.*, 2020], they focus on the audio-visual speech separation task by using the still face images rather than lip movement on video clips. [Nawaz *et al.*, 2021] tries to explore the language dependent problem in voice-face association. For example, the question about: “Can a speaker be recognized irrespective of the spoken language?” In [Wen *et al.*, 2021], they used a two-level inter-modality and intra-modality alignment.

Also, they propose a reweighting scheme to better handle the hard samples. In [Nagrani *et al.*, 2018a], they learn the cross-modal representation in a self-supervised manner by the contrastive loss, and a curriculum learning schedule boosts the performance. Other works seeks the possibility of generating the human face from voice. [Wen *et al.*, 2019b] uses a generative adversarial network to reconstruct someone’s face. [Choi *et al.*, 2020] links the inference stage and generation stage by a conditional generative adversarial network.

## 2.2 Contrastive Learning

Contrastive learning has recently become a prominent tool in unsupervised representation learning, leading to state-of-the-art results in many fields such as vision, language and audio. Contrastive learning either aims to predict parts of the data, or to discriminate between individual instances. The former explicitly define a proxy tasks, such as colorization, jigsaw solving, temporal coherence and future prediction [Han *et al.*, 2020; Pan *et al.*, 2021]. The later, instances discrimination can be formed as a  $N$ -classification task which  $N$  is the number of instances. SupCon [Khosla *et al.*, 2020] extend the self-supervised contrastive approach to the fully-supervised setting, allowing us to effectively leverage label information. The main difference between different contrastive learning approaches lies in their strategy for definition of positive pairs. For example, image random cropping and rotation, or different augmentation of the same instance. After getting the positive pairs, comparing with as many negative examples is important to boost the performance. Other works focus on the hard negatives mining [Kalantidis *et al.*, 2020] or unbiased negatives sampling [Chuang *et al.*, 2020; Wu *et al.*, 2021] in contrastive learning.

Clustering-based contrastive learning [Caron *et al.*, 2018; Caron *et al.*, 2020; Zhan *et al.*, 2020] has been proposed to learn between instance embeddings and cluster assignments. Among them, PCL [Li *et al.*, 2021] formulates the clustering and representation learning as an Expectation-Maximization (EM) algorithm, which approximates and maximizes the log-likelihood function. While, whether we can learn effective representations if prototypes and instance vectors are derived from different modalities will be explored in this paper. Closer to our work, XDC [Alwassel *et al.*, 2020] performs iterative clustering and classification to learn cross-modal representation. The cluster assignments of one modality are considered as pseudo-label of another modality, following a classification objective. While, the cross-entropy classification loss is sensitive to the noisy label, the inaccurate pseudo-label make it difficult for convergence. However, our method is conceptually different from the XDC. In XDC, the cluster assignments of one modality are considered as pseudo-label of another modality, following a classification objective, which results in two weaknesses: (1) The cross-entropy classification loss is sensitive to the noisy label, the inaccurate pseudo-label make it difficult for convergence. (2) The initialization of pseudo-label gives endow the semantic information to the cluster centroid, which may lead to training bias. In our method, we actively avoid the pseudo-label and noisy label problem through the clustering and contrast iteration.

## 3 Methodology

### 3.1 Preliminary

Consider a dataset of  $N$  video instances  $\mathcal{X} = \{(x_i^V, x_i^F) \mid i = 1, \dots, N\}$  where each instance have two modalities, i.e., voice and face. The goal of voice-face association is to jointly learn high-level representations  $(\mathbf{v}_i, \mathbf{f}_i)$  for both modalities to capture cross-modal correlations. The voice and the face of the same *identity* (person) are expected to have similar representations and vice versa. To this end, let  $g^M : x^M \rightarrow \mathcal{R}^{d_m}$  be a parametrized modality specific backbone neural network that takes  $x^M$  as input from modality  $M \in \{V, F\}$  and outputs a representation of  $d_M$  dimension. All the representations are normalized onto a unit hypersphere. In this work, we follow the unsupervised setting in [Nagrani *et al.*, 2018a] to train the encoders using only video instances without the exact identity labels.

State-of-the-art unsupervised audio-visual representation learning methods [Alayrac *et al.*, 2020; Morgado *et al.*, 2021b; Patrick *et al.*, 2021] usually follow the paradigm of *Cross-modal Instance Discrimination* (CID). For a batch of different video instances, CID uses the voice-face pairs from the same instance as the positive pairs and those from different instances as the negative pairs. The representations are then trained in a contrastive learning manner [He *et al.*, 2020; Chen *et al.*, 2020] across the two modalities. For simplicity, we only describe the case from voice to face and the same goes for the other side. Given a positive pair  $(\mathbf{v}_i, \mathbf{f}_i)$ , CID establishes a negative face representation set  $\mathbb{N}_i^F = \{\mathbf{f}_j \mid j \neq i\}$  and seeks to maximize the mutual agreement between  $\mathbf{v}_i$  and  $\mathbf{f}_i$  by minimizing the InfoNCE loss [Oord *et al.*, 2018]:

$$\ell_c(\mathbf{v}_i, \mathbf{f}_i) = -\log \frac{\exp(\mathbf{v}_i^T \mathbf{f}_i / \tau)}{\exp(\mathbf{v}_i^T \mathbf{f}_i / \tau) + \sum_{\mathbf{f}_j \in \mathbb{N}_i^F} \exp(\mathbf{v}_i^T \mathbf{f}_j / \tau)}, \quad (1)$$

where  $\tau$  is the temperature hyper-parameter. This process can be interpreted as maximizing the similarity  $s_i^p = \mathbf{v}_i^T \mathbf{f}_i$  between the positive pair while suppressing the similarity  $s_{i,j}^n = \mathbf{v}_i^T \mathbf{f}_j$  between the negative pairs. Moreover, the gradients of  $s_i^p$  and  $s_{i,k}^n$  are computed as:

$$\frac{\partial \ell_c}{\partial s_{i,k}^n} = \frac{1}{\tau} \frac{\exp(s_{i,k}^n / \tau)}{\exp(s_i^p / \tau) + \sum_{\mathbf{f}_j \in \mathbb{N}_i^F} \exp(s_{i,j}^n / \tau)}, \quad (2)$$

$$\frac{\partial \ell_c}{\partial s_i^p} = -\frac{1}{\tau} \frac{\sum_{\mathbf{f}_j \in \mathbb{N}_i^F} \exp(s_{i,j}^n / \tau)}{\exp(s_i^p / \tau) + \sum_{\mathbf{f}_j \in \mathbb{N}_i^F} \exp(s_{i,j}^n / \tau)}. \quad (3)$$

As mentioned in Section 1, CID suffers from the false negative and the deviate positive problems. On one hand, for a false negative  $(\mathbf{v}_i, \mathbf{f}_k)$ , their similarity  $s_{i,k}^n$  could be high as they share the same identity. From Eq. 2, this will impose a heavy penalty which pushes  $\mathbf{v}_i$  away from  $\mathbf{f}_k$ . This is suboptimal as  $\mathbf{v}_i$  and  $\mathbf{f}_k$  should be matched together. On the other hand, for a deviate positive, the positive similarity  $s_i^p$  tends to be small. According to Eq. 3, there will be a heavy penalty which pulls the deviate positive close, which could lead to local optima and hamper the generalizability of the representations.

### 3.2 Cross-modal Prototype Contrast

To tackle false negative problem, we propose a cross-modal prototype-based contrastive learning method. Our motivation is to extract a set of prototypes which encodes the underlying semantic structure of the whole dataset by clustering the video instances and contrast the instances with the prototypes. For each instance, its representation in one modality is pulled close to the prototype where it belongs in the other modality and pushed away from other prototypes. As false negatives usually belong to the same prototype, they will be optimized to have similar representations.

Following [Li *et al.*, 2021], we conduct clustering and contrast iteratively: the prototypes are reconstructed after each training epoch and used in the next epoch. However, as the encoders change rapidly during training, directly gathering the representations from different training batches induces severe representation inconsistency, which harms the quality of the prototypes. For this reason, we design *instance feature memories* ( $\text{Mem}_i^V, \text{Mem}_i^F$ ) for each video instance  $x_i$ . The instance feature memories are defined as the momentum update of the instance representations in both modalities:

$$\text{Mem}_i^V = m \cdot \text{Mem}_i^V + (1 - m) \cdot \mathbf{v}_i, \quad (4)$$

$$\text{Mem}_i^F = m \cdot \text{Mem}_i^F + (1 - m) \cdot \mathbf{f}_i, \quad (5)$$

where  $m \in [0, 1]$  is a momentum coefficient. Compared with  $(\mathbf{v}_i, \mathbf{f}_i)$ , the change of  $(\text{Mem}_i^V, \text{Mem}_i^F)$  are much smoother during training, which provides better consistency. After each training epoch, we cluster all instances into  $K$  distinct groups in each modality based on their instance feature memories. Here we simply adopt the  $k$ -means algorithm but we note that our method is agnostic to clustering methods. The clustering centroids  $\mathcal{C}^V = \{\mathbf{c}_i^V \mid i = 1, \dots, K\}$  and  $\mathcal{C}^F = \{\mathbf{c}_i^F \mid i = 1, \dots, K\}$  are used as the prototypes.

Next, we apply the contrastive loss between the instance representations and the semantic prototypes in a cross-modal way. For each instance representation  $(\mathbf{v}_i, \mathbf{f}_i)$ , we pull  $\mathbf{v}_i$  close to the face prototypes  $\mathbf{c}_s^F$  where  $\mathbf{f}_i$  belongs and push it away from all other face prototypes:

$$\ell_p(\mathbf{v}_i, \mathcal{C}^F) = -\log \frac{\exp(\mathbf{v}_i^T \mathbf{c}_s^F / \tau)}{\exp(\mathbf{v}_i^T \mathbf{c}_s^F / \tau) + \sum_{j \neq s} \exp(\mathbf{v}_i^T \mathbf{c}_j^F / \tau)}. \quad (6)$$

The same computation goes for  $\ell_p(\mathbf{f}_i, \mathcal{C}^V)$ . To ensure robust semantic prototype estimation, we cluster the instance feature memories for  $R$  times with different numbers of clusters and aggregate the instance-wise contrastive loss to retain local smoothness. The overall CMPC loss is defined as:

$$\begin{aligned} \mathcal{L}_{\text{CMPC}} = & \ell_c(\mathbf{v}_i, \mathbf{f}_i) + \frac{1}{R} \sum_{r=1}^R \ell_p(\mathbf{v}_i, \mathcal{C}_r^F) \\ & + \ell_c(\mathbf{f}_i, \mathbf{v}_i) + \frac{1}{R} \sum_{r=1}^R \ell_p(\mathbf{f}_i, \mathcal{C}_r^V). \end{aligned} \quad (7)$$

Compared to previous clustering-based contrastive learning methods [Caron *et al.*, 2018; Li *et al.*, 2021], our CMPC differs in the following two aspects: (1) They adopt a offline update strategy where the representations of all training instances are recomputed before clustering to achieve representation consistency. This induces heavy computation

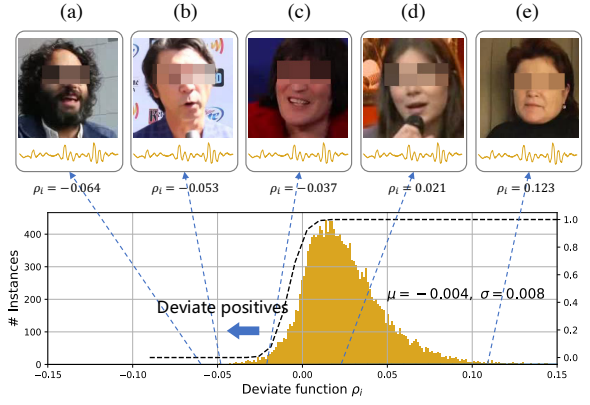


Figure 3: Deviate positives in a CMPC pretrained model. The relationship between instance numbers and  $\rho(i)$  is shown in the histogram. We show some positive pairs under different  $\rho(i)$  and describe the voice track: (a) Female voice in the background, (b) Noisy background with music, (c) Mature male voice, (d) Female voice, (e) Female voice. Deviate positives prone to hold lower  $\rho(i)$ .

cost, especially in large-scale datasets. On the contrary, our method leverages an online strategy which achieves representation consistency without apparent computation overhead. (2) Our method is designed for cross-modal contrastive learning, while these methods mainly focus on one modality.

### 3.3 Instance Recalibration

To address the deviate positive problem, we propose a *Instance Recalibration* mechanism to reduce the importance of deviate positives on representation learning.

For a video instance, it is more likely to be deviate if (1) the similarity of its voice representation  $\mathbf{v}_i$  and face representation  $\mathbf{f}_i$  is small, and (2) the similarity of its corresponding voice prototype  $\mathbf{c}_{s_1}^V$  and face prototype  $\mathbf{c}_{s_2}^F$  is large. Base on the above observations, we design a deviation function  $\rho(\cdot)$  which measures the possibility of an instance being deviate:

$$\rho(i) = \mathbf{v}_i^T \mathbf{f}_i - \left( \frac{1}{R} \sum_{r=1}^R (\mathbf{c}_{s_1}^V)^T \mathbf{c}_{s_2}^F \right) \quad (8)$$

The deviation function describes how far the instance representation similarity deviates from prototype similarity. The smaller  $\rho(i)$  is, the more likely the  $i$ -th instance is a deviate positive. We then compute a recalibration factor  $w_i$  by:

$$w_i = F_{\mathcal{N}}(\rho(i)). \quad (9)$$

Inspired by [Morgado *et al.*, 2021a], we use the cumulative distribution function (CDF) to control the shape of recalibration function.  $F_{\mathcal{N}}$  is the CDF of a Gaussian distribution  $\mathcal{N}(\mu + \delta\sigma, \kappa\sigma^2)$ , where  $\mu$  and  $\sigma$  are the mean and the variance of the deviation function  $\rho(i)$  on all instances, and  $\delta$  and  $\kappa$  are hyper-parameters to control the shape of the recalibration function. The recalibration factor  $w_i$  is then use to rescale the CMPC loss of each instance. The final loss function are defined as:

$$\mathcal{L} = \frac{1}{\sum_n w_n} \sum_i w_i \mathcal{L}_{\text{CMPC}}(\mathbf{v}_i, \mathbf{f}_i) \quad (10)$$

| Method                                       | Kind of Sup.    | V-F Matching (Acc %) |             |             |             | F-V Matching (Acc %) |             |             |             | Mean        |
|--|-----------------|----------------------|-------------|-------------|-------------|----------------------|-------------|-------------|-------------|-------------|
|  |                 | U                    | G           | N           | GN          | U                    | G           | N           | GN          |             |
| SVHF [Nagrani <i>et al.</i> , 2018b]         | Supervised      | 81.0                 | 63.9        | -           | -           | 79.5                 | 63.4        | -           | -           | -           |
| DIMNet-I [Wen <i>et al.</i> , 2019a]         | Supervised      | 83.5                 | 70.9        | 82.0        | 69.9        | 83.5                 | 71.8        | 82.4        | 70.9        | 76.9        |
| DIMNet-G [Wen <i>et al.</i> , 2019a]         | Supervised      | 72.9                 | 50.3        | 71.9        | 50.2        | 72.5                 | 50.5        | 72.2        | 50.6        | 61.4        |
| Horiguchi’s [Horiguchi <i>et al.</i> , 2018] | Supervised      | 78.1                 | 61.7        | -           | -           | 77.8                 | 60.8        | -           | -           | -           |
| Wen’s †[Wen <i>et al.</i> , 2021]            | Supervised      | 81.4                 | 67.8        | 79.6        | 66.7        | 79.9                 | 66.1        | 79.3        | 66.9        | 73.5        |
| Kim’s [Kim <i>et al.</i> , 2018]             | Self-supervised | 78.2                 | 62.9        | 76.4        | 61.6        | 78.6                 | 61.6        | 76.7        | 61.2        | 69.7        |
| CID [Morgado <i>et al.</i> , 2021b]          | Self-supervised | 78.3                 | 64.7        | 77.8        | 64.2        | 77.6                 | 64.7        | 77.2        | 64.1        | 71.1        |
| CMPC (Ours)                                  | Self-supervised | <b>82.2</b>          | <b>67.4</b> | <b>81.1</b> | <b>67.1</b> | <b>81.7</b>          | <b>67.4</b> | <b>81.1</b> | <b>66.8</b> | <b>74.4</b> |
| CID [Morgado <i>et al.</i> , 2021b]          | Low-shot        | 82.9                 | 68.6        | 82.6        | 66.3        | 79.9                 | 68.7        | 82.1        | 68.7        | 75.0        |
| CMPC (Ours)                                  | Low-shot        | <b>85.3</b>          | <b>70.6</b> | <b>84.3</b> | <b>70.2</b> | <b>84.9</b>          | <b>70.6</b> | <b>84.4</b> | <b>71.1</b> | <b>77.7</b> |

Table 1: **Matching Results.** Bidirectional (V-F and F-V) matching performance in accuracy. The U/G/N/GN represent different demographic attributes to construct testing groups. All the test data are from the unseen and unheard identities. †: It uses different face data from previous and ours, we reproduce it on VoxCeleb dataset.

## 4 Experiments

### 4.1 Dataset

We use the VoxCeleb dataset [Nagraniy *et al.*, 2017] in our experiments, which contains 21,063 video clips of 1251 celebrities, extracted from videos uploaded to YouTube. The dataset is gender-balanced and speakers span a range of different ethnicity, accent, profession, and age. Following the previous works [Nagrani *et al.*, 2018a; Wen *et al.*, 2021], random 1001 identities are randomly sampled as training data and the rest 250 identities as the test set. The face images are generated by facial bounding boxes cropping the video frames. The voice clips are extracted from the soundtrack of the videos with a 16kHz sampling rate at 16-bit PCM.

The test set is annotated with different demographic attributes to construct 5 testing groups, including the unstratified group (U), stratified by gender (G), stratified by nationality (N), stratified by age (A), stratified by gender and nationality (GN), and stratified by (G,N,A). The demographic groups allow us to control the important attributes in the experiments. For example, in the G-group, all the alternative instances have the same gender. The matching and verification tasks are conducted under different demographic partitions. All the evaluation lists are released.

## 5 Evaluation Protocols

To appraise our learned representations, following protocols are conducted for evaluation:

**Matching.** We are given a probe from one modality, and two alternatives from another modality, including one that belongs to the same person as the probe, and another imposter that comes from another person. The task is to affirm which alternative matches the probe. We report the bidirectional matching performance, namely voice to face (V-F) matching accuracy and face to voice (F-V) matching accuracy.

**Verification.** We are given a voice clip and face image. The task is to judge whether they belong to the same identity. The

basis of judgment is to compare the similarity of two input features with a given threshold. With different thresholds, we could get various false positive rates and true positive rates. The final performance is reported by the *Area Under the ROC curve* (AUC).

**Retrieval.** Given a probe from one modality, and a large size gallery which contains candidates from another modality. One or more candidates might match the probe. The task is to order the gallery such that the matched candidates lie at the top of the ordering. We report the bidirectional performance in terms of *Mean Average Precision* (mAP).

### 5.1 Implementation Details

**Voice and face preprocessing.** We randomly crop a 5 seconds length speech and extract the commonly used *logmel* feature which is the mel-spectrogram at log-scale during training. The *logmel* features are transformed with 100ms windows width, 10ms frame shift and 64 mel-bins, yielding an input at size of  $64 \times 500$ . No further speech-specific preprocessing (e.g. noise suppression, voice activity detection) is applied. The face frames are extracted at 1fps from the video stream and cropped surrounding a human face. The input is a series of RGB images resized at  $224 \times 224$  pixels. The only augmentation used on the face frames is a horizontal flip.

**Network architecture.** The ResNet-34 [He *et al.*, 2016] is adopted as the voice and face encoder network. Our network is trained from scratch without any pretrained data. The original fully-connected layer of each encoder is removed. The weights of the two encoders are not shared, while they are trained jointly by the approach introduced above. We get the features after the last convolutional layer as the representations, and they are average pooled to the 512-dimensions. For both the voice and face representations, they are projected normalized into the unit sphere by L2 norm. In the evaluation phase, the projection network is removed, and only calculates the similarity between learned representations.



| Method                                       | Kind of Sup.    | Verification (AUC) |             |             |             |             | Retrieval (mAP %) |             |
|--|-----------------|--------------------|-------------|-------------|-------------|-------------|-------------------|-------------|
|  |                 | U                  | G           | N           | A           | GNA         | V-F               | F-V         |
| DIMNet-I [Wen <i>et al.</i> , 2019a]         | Supervised      | 82.5               | 71.0        | 81.1        | 77.7        | 62.8        | 4.25              | 4.17        |
| DIMNet-IG [Wen <i>et al.</i> , 2019a]        | Supervised      | 83.2               | 71.2        | 81.9        | 78.0        | 62.8        | 4.42              | 4.23        |
| Horiguchi’s [Horiguchi <i>et al.</i> , 2018] | Supervised      | -                  | -           | -           | -           | -           | 2.18              | 1.96        |
| Wen’s [Wen <i>et al.</i> , 2021]             | Supervised      | 80.3               | 67.6        | 79.3        | 76.2        | 62.3        | 4.31              | 3.91        |
| PINs [Nagrani <i>et al.</i> , 2018a]         | Self-supervised | 78.5               | 61.1        | 77.2        | 74.9        | 58.8        | -                 | -           |
| CID [Morgado <i>et al.</i> , 2021b]          | Self-supervised | 78.2               | 65.2        | 78.4        | 75.1        | 62.9        | 3.69              | 3.55        |
| CMPC (Ours)                                  | Self-supervised | <b>82.6</b>        | <b>68.6</b> | <b>81.3</b> | <b>77.8</b> | <b>63.7</b> | <b>3.80</b>       | <b>3.70</b> |
| CID [Morgado <i>et al.</i> , 2021b]          | Low-shot        | 80.8               | 67.4        | 80.7        | 77.7        | 65.3        | 4.17              | 4.10        |
| CMPC (Ours)                                  | Low-shot        | <b>84.6</b>        | <b>71.6</b> | <b>82.5</b> | <b>79.8</b> | <b>65.1</b> | <b>4.31</b>       | <b>4.27</b> |

Table 2: Verification and Retrieval Results.

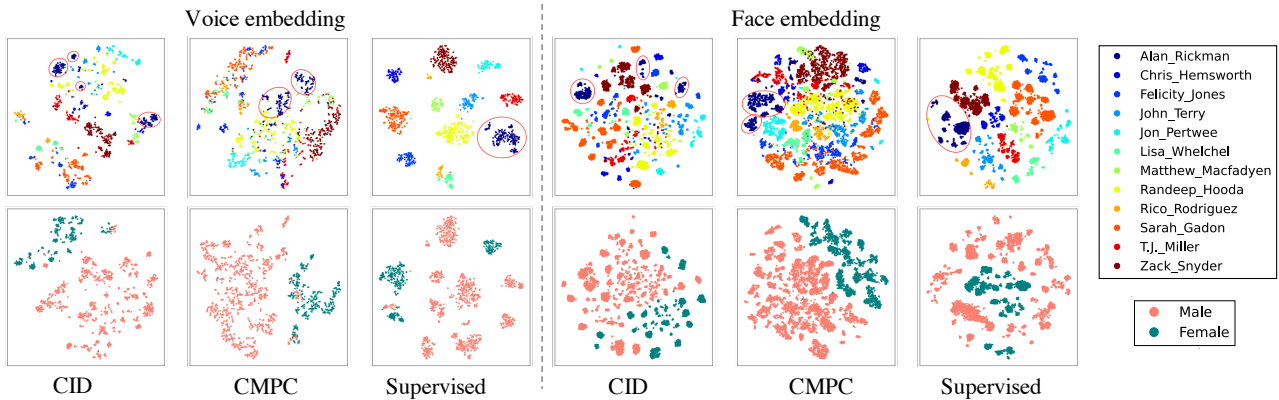


Figure 4: T-SNE visualization of voice and face representations. We compare our approach with CID and the fully supervised approach. We randomly choose 12 identities’ videos. Points in the figure correspond to a face frame or voice clip representation. Best view in color.

| Cross-modal Prototype | Instance Recalibration | Matching |      | Verification |
|-----------------------|------------------------|----------|------|--------------|
|                       |                        | V-F      | F-V  |              |
| ✗                     | ✗                      | 78.3     | 77.6 | 78.2         |
| ✓                     | ✗                      | 81.9     | 80.1 | 80.4         |
| ✓                     | ✓                      | 82.2     | 81.7 | 82.6         |

Table 3: Ablation study. We quantify the individual contributions of our proposed approach: cross-modal prototype and instance recalibration.

**Optimization.** We use the Adam optimizer, with a momentum of 0.9, a weight decay of 0.002 and a batch size of 128. We train for 3.2M iterations. We use a cosine annealing schedule with warm-up to tune our learning rate. The initial learning rate is  $1e-4$  and increases linearly to  $5e-3$  in the first 0.3M iterations, then gradually decreases to  $1e-4$ . We also warmup our network in these 0.3M iterations using CID loss. The momentum coefficient in computing instance feature memory is 0.5. We set temperature  $\tau$  to 0.03, number of clusters  $K = \{500, 1000, 1500\}$ . The  $\sigma$  and  $\kappa$  in the recalibration are set to -1 and 0.1 respectively. We use KeOps

[Charlier *et al.*, 2021] for the efficient  $k$ -means clustering. Our code, model and testing resources are released<sup>1</sup>.

## 5.2 Quantitative Results

We now compare our CMPC to prior works, including supervised and unsupervised methods on various voice-face association evaluation protocols. We can outperform previous state-of-the-art unsupervised methods on various voice-face association evaluation protocols.

Table 1 summarizes the cross-modal matching performance, and Table 2 presents the cross-modal verification and retrieval results. We could see that our CMPC outperforms the previous unsupervised approaches by a large margin, bridging the gap between unsupervised and supervised methods. As this task is challenging, even human performance is around 80% [Nagrani *et al.*, 2018b] on the matching protocol. Our results indicate that it is feasible to learn robust voice-face representation through unlabeled videos.

**Compared methods.** Among the compared methods, DIMNet [Wen *et al.*, 2019a] uses one or more covariates, such as gender and identity, by a multi-task object function. Due to

<sup>1</sup><https://github.com/Cocoxili/CMPC>

the use of more supervised information, it has achieved good results. Wen’s [Wen *et al.*, 2021] adaptively re-weights identity samples, while its re-weighting method is based on the identity label which is not used in our method. Meanwhile, their face data are obtained from the VGGFace dataset, which is different from the previous and our method. The results we report are reproducing<sup>2</sup> their method on the VoxCeleb dataset. PINs [Nagrani *et al.*, 2018a] also employs a contrastive loss to learn the representation under an unsupervised setting, while our contrastive loss is distinct from theirs in which a positive pair only compares with one negative pair at once. It can be seen as a variant of the triplet loss.

**Low-shot learning.** To further evaluate the learned representations, we conduct our experiments on the low-shot training, which only uses a few instances per-identity. These instances are labeled by the identity information. We follow the DIMNet [Wen *et al.*, 2019a] setup to finetune our network, while we do not use other covariates such as gender, nationality, etc. We representatively select 5 instances per-identity, which is about one-third of the whole dataset. We report the evaluation results in Table 1 and Table 2. Due to the better voice and face representations we learned, the results even outperform those with entire dataset supervision.

**Ablation Study.** To shed light on the contributions of our two components, cross-modal prototype contrast and instance recalibration, we report the matching and verification results on the unseen-unheard unstratified group (U) of VoxCeleb dataset as in Table 3. The other experimental settings are kept identical. Both two components yield a positive impact on the results, individually and jointly.

### 5.3 Qualitative Analyses

In Figure 4, we visualize the learned representations of voice clips and face images using t-SNE. We color the sample points by identity labels and gender labels respectively. We compare our approach with the CID method [Morgado *et al.*, 2021b] and fully supervised approach [Wen *et al.*, 2019a].

The CID only focuses on the local association between instances and ignores the semantic structure. It undesirable pushes false-negative instances apart. Therefore, different instances representations of the same person will appear far away, such as the samples marked in the red circles. While, in the supervised method, every identity’s representations are brought together tightly. In our method, we keep both voice and face representations of the same identity closer together even without using any labels, due to the cross-modal prototype contrast. Another interesting thing, the representations are not well separated by gender even in the supervised method, as we do not use gender supervision. While in our approaches, they are separated naturally.

## 6 Conclusion

In this paper, we propose a novel unsupervised approach to learn voice-face association from the talking-face videos. We use the natural relevance between audio and visual frames to establish a self-supervised approach. Our method is based

on the framework of contrastive learning, where the selection of positive and negative samples is crucial. We reveal their adverse impact by analyzing their gradient flow in the training phase. To tackle the false negative problem, we propose a cross-modal prototype contrast algorithm by estimating the semantic prototypes of each instance, then we replace the instance-to-instance contrast with instance-to-prototype contrast in a cross-modal way. To tackle the deviate positive problem, we dynamically recalibrate the instances by a deviation function. Experiments on cross-modal matching, verification, and retrieval have shown the superiority of our learn representations.

## Acknowledgments

This work is supported by the major Science and Technology Innovation 2030 “New Generation Artificial Intelligence” project 2020AAA0104803.

## References

- [Alayrac *et al.*, 2020] Jean-Baptiste Alayrac, Adria Recasens, Rosalia Schneider, Relja Arandjelovic, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. *NeurIPS*, 2(6):7, 2020.
- [Alwassel *et al.*, 2020] Humam Alwassel, Dhruv Mahajan, Bruno Korbar, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. In *NerulPS*, 2020.
- [Bahrick *et al.*, 2005] Lorraine E Bahrick, Maria Hernandez-Reif, and Ross Flom. The development of infant learning about specific face-voice relations. *Developmental psychology*, 41(3):541, 2005.
- [Caron *et al.*, 2018] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, pages 132–149, 2018.
- [Caron *et al.*, 2020] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NerulPS*, 2020.
- [Charlier *et al.*, 2021] Benjamin Charlier, Jean Feydy, Joan Alexis Glaunès, François-David Collin, and Ghislain Durif. Kernel operations on the gpu, with autodiff, without memory overflows. *Journal of Machine Learning Research*, 22(74):1–6, 2021.
- [Chen *et al.*, 2020] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607. PMLR, 2020.
- [Choi *et al.*, 2020] Hyeong-Seok Choi, Changdae Park, and Kyogu Lee. From inference to generation: End-to-end fully self-supervised generation of human face from speech. In *ICLR*, 2020.
- [Chuang *et al.*, 2020] Ching-Yao Chuang, Joshua Robinson, Lin Yen-Chen, Antonio Torralba, and Stefanie Jegelka.

<sup>2</sup><https://github.com/KID-7391/seeking-the-shape-of-sound>

- Debiased contrastive learning. In *NeurIPS*, Virtual, Online, 2020.
- [Chung *et al.*, 2020] Soo-Whan Chung, Soyeon Choe, Joon Son Chung, and Hong-Goo Kang. Facefilter: Audio-visual speech separation using still images. In *INTERSPEECH*, volume 2020-October, pages 3481 – 3485, 2020.
- [Han *et al.*, 2020] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. In *ECCV*, pages 312–329. Springer, 2020.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [He *et al.*, 2020] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020.
- [Horiguchi *et al.*, 2018] Shota Horiguchi, Naoyuki Kanda, and Kenji Nagamatsu. Face-voice matching using cross-modal embeddings. In *ACM MM*, pages 1011–1019, 2018.
- [Joassin *et al.*, 2011] Frédéric Joassin, Mauro Pesenti, Pierre Maurage, Emilie Verreckt, Raymond Bruyer, and Salvatore Campanella. Cross-modal interactions between human faces and voices involved in person recognition. *Cortex*, 47(3):367–376, 2011.
- [Kalantidis *et al.*, 2020] Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. In *NeurIPS*, 2020.
- [Khosla *et al.*, 2020] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020.
- [Kim *et al.*, 2018] Changil Kim, Hijung Valentina Shin, Tae-Hyun Oh, Alexandre Kaspar, Mohamed Elgharib, and Wojciech Matusik. On learning associations of faces and voices. In *ACCV*, pages 276–292. Springer, 2018.
- [Li *et al.*, 2021] Junnan Li, Pan Zhou, Caiming Xiong, and Steven C.H. Hoi. Prototypical contrastive learning of unsupervised representations. In *ICLR*, 2021.
- [Morgado *et al.*, 2021a] Pedro Morgado, Ishan Misra, and Nuno Vasconcelos. Robust audio-visual instance discrimination. In *CVPR*, pages 12934–12945, 2021.
- [Morgado *et al.*, 2021b] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. In *CVPR*, pages 12475–12486, 2021.
- [Nagrani *et al.*, 2018a] Arsha Nagrani, Samuel Albanie, and Andrew Zisserman. Learnable pins: Cross-modal embeddings for person identity. In *ECCV*, pages 71–88, 2018.
- [Nagrani *et al.*, 2018b] Arsha Nagrani, Samuel Albanie, and Andrew Zisserman. Seeing voices and hearing faces: Cross-modal biometric matching. In *CVPR*, pages 8427–8436, 2018.
- [Nagraniy *et al.*, 2017] Arsha Nagraniy, Joon Son Chungy, and Andrew Zisserman. Voxceleb: A large-scale speaker identification dataset. In *INTERSPEECH*, volume 2017-August, pages 2616 – 2620, 2017.
- [Nawaz *et al.*, 2021] Shah Nawaz, Muhammad Saad Saeed, Pietro Morerio, Arif Mahmood, Ignazio Gallo, Muhammad Haroon Yousaf, and Alessio Del Bue. Cross-modal speaker verification and recognition: A multilingual perspective. In *CVPR (Workshop)*, pages 1682–1691, 2021.
- [Oord *et al.*, 2018] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [Pan *et al.*, 2021] Tian Pan, Yibing Song, Tianyu Yang, Wenhao Jiang, and Wei Liu. Videomoco: Contrastive video representation learning with temporally adversarial examples. In *CVPR*, pages 11205–11214, 2021.
- [Patrick *et al.*, 2021] Mandela Patrick, Yuki M Asano, Polina Kuznetsova, Ruth Fong, Joao F Henriques, Geoffrey Zweig, and Andrea Vedaldi. Multi-modal self-supervision from generalized data transformations. In *ICCV*, 2021.
- [Wen *et al.*, 2019a] Yandong Wen, Mahmoud Al Ismail, Weiyang Liu, Bhiksha Raj, and Rita Singh. Disjoint mapping network for cross-modal matching of voices and faces. In *ICLR*, 2019.
- [Wen *et al.*, 2019b] Yandong Wen, Bhiksha Raj, and Rita Singh. Face reconstruction from voice using generative adversarial networks. In *NeurIPS*, volume 32, 2019.
- [Wen *et al.*, 2021] Peisong Wen, Qianqian Xu, Yangbangyan Jiang, Zhiyong Yang, Yuan He, and Qingming Huang. Seeking the shape of sound: An adaptive framework for learning voice-face association. In *CVPR*, pages 16347–16356, 2021.
- [Wu *et al.*, 2021] Mike Wu, Milan Mosse, Chengxu Zhuang, Daniel Yamins, and Noah Goodman. Conditional negative sampling for contrastive learning of visual representations. In *ICLR*, 2021.
- [Zhan *et al.*, 2020] Xiaohang Zhan, Jiahao Xie, Ziwei Liu, Yew-Soon Ong, and Chen Change Loy. Online deep clustering for unsupervised representation learning. In *CVPR*, pages 6688–6697, 2020.
- [Zolfaghari *et al.*, 2021] Mohammadreza Zolfaghari, Yi Zhu, Peter Gehler, and Thomas Brox. Crossclr: Cross-modal contrastive learning for multi-modal video representations. In *ICCV*, pages 1450–1459, 2021.