# Learning Branched Fusion and Orthogonal Projection for Face-Voice Association

Muhammad Saad Saeed<sup>1\*</sup>, Shah Nawaz<sup>2\*†</sup>, Muhammad Haris Khan<sup>3</sup>, Sajid Javed<sup>4</sup>,

Muhammad Haroon Yousaf<sup>1</sup>, Alessio Del Bue<sup>2,5</sup>

<sup>1</sup>Swarm Robotics Lab (SRL)-NCRA, University of Engineering and Technology Taxila, <sup>2</sup>Pattern Analysis & Computer Vision (PAVIS) - Istituto Italiano di Tecnologia (IIT), <sup>3</sup>Mohamed Bin Zayed University of Artificial Intelligence, <sup>4</sup>Khalifa University of Science and Technology, <sup>5</sup>Visual Geometry & Modelling (VGM) - Istituto Italiano di Tecnologia (IIT)

Abstract— Recent years have seen an increased interest in establishing association between faces and voices of celebrities leveraging audio-visual information from YouTube. Prior works adopt metric learning methods to learn an embedding space that is amenable for associated matching and verification tasks. Albeit showing some progress, such formulations are, however, restrictive due to dependency on distance-dependent margin parameter, poor run-time training complexity, and reliance on carefully crafted negative mining procedures. In this work, we hypothesize that an enriched representation coupled with an effective yet efficient supervision is important towards realizing a discriminative joint embedding space for face-voice association tasks. To this end, we propose a light-weight, plug-and-play mechanism that exploits the complementary cues in both modalities to form enriched fused embeddings and clusters them based on their identity labels via orthogonality constraints. We coin our proposed mechanism as fusion and orthogonal projection (FOP) and instantiate in a two-stream network. The overall resulting framework is evaluated on VoxCeleb1 and MAV-Celeb datasets with a multitude of tasks, including cross-modal verification and matching. Results reveal that our method performs favourably against the current state-of-the-art methods and our proposed formulation of supervision is more effective and efficient than the ones employed by the contemporary methods. In addition, we leverage cross-modal verification and matching tasks to analyze the impact of multiple languages on face-voice association. Code is available: https://github.com/msaadsaeed/FOP

Index Terms—Multimodal, Face-voice association, Crossmodal verification and matching

## I. INTRODUCTION

T is a well-studied and understood fact that humans can associate voices and faces of people because the neurocognitive pathways for voices and faces share same structure [1], [2]. Recently, Nagrani et al. [3], [4], [5] introduced the face-voice association task to vision community with the creation of a large-scale audio-visual dataset, comprising faces and voices of 1, 251 celebrities. Since then, the face-voice association task has gained increased research interest [6], [7], [4], [3], [8], [9], [10], [11]. In addition, we observe the creation of new audio-visual datasets for studying novel tasks [7], [12], [13], [14]. For example, fundamental face-voice association tasks are studied to analyze the impact of multiple languages, see Fig. 1.

Most existing works [7], [3], [4], [12] tackle face-voice association as a cross-modal biometric task. The two prominent challenges in developing an effective method for this task are learning of a common yet discriminative embedding space, where instances from two modalities are sufficiently aligned and instances of semantically similar identities are nearby. Often separate networks for face and voice modalities are leveraged to obtain the respective feature embeddings and contrastive or triplet loss formulations are employed to construct this embedding space.

Although showing some effectiveness in this task, such loss formulations, however, are restrictive in following ways. First, they require tuning of a margin hyperparameter, which is hard as the distances between instances can alter significantly while training. Secondly, the run-time training complexity for contrastive and triplet losses are  $O(n^2)$  and  $O(n^3)$ , respectively, where *n* is the number of available instances for a modality. Finally, to mitigate the high run-time training complexity challenge, different variants of carefully crafted negative mining strategies are used, which are both timeconsuming and performance sensitive.

A few methods e.g., [8] have attempted to replace the contrastive/triplet loss formulations by utilizing auxiliary identity centroids [15]. The training process alternates between the following two steps: 1) clustering embeddings around their identity centroids and pushing embeddings away from all other identity centroids, and 2) updating these centroids using the mini-batch instances. Such centroid based losses are used with traditional classification loss (i.e. softmax crossentropy (CE)). However, their co-existence is unintuitive and ineffective because the former promotes margins in Euclidean space whereas latter implicitly achieves separability in the angular domain.

In this work, we hypothesize that an enriched unified feature representation, encompassing complementary cues from both modalities, alongside an effective yet efficient supervision formulation is crucial towards realizing a discriminative joint embedding space for improved face-voice association. To this end, we propose a light-weight, plug-and-play mechanism that exploits the best in both modalities through fusion and

<sup>\*</sup> Equal contribution

<sup>†</sup> Current Affiliation: Deutsches Elektronen-Synchrotron



Fig. 1. Face-voice association is established in cross-modal verification and matching tasks. The goal of cross-modal verification task is to verify if a clip of a voice and face image belong to the same identity. In a cross-modal matching task, Voice-Face(V-F): given a clip of a voice and two or more face images, select the face image that belongs to the voice. (Face-Voice) F-V: given an image of a face, determine the respective voice. In addition, we analyze the impact of multiple of languages on cross-modal verification and matching tasks.

semantically aligns fused embeddings with their identity labels via orthogonality constraints. We instantiate our proposed mechanism in the two-stream pipeline, which provides face and voice embeddings, and the resulting overall framework is an effective and efficient approach for face-voice matching and verification tasks.

We summarize our key contributions as follows. 1) We propose to harness the complementary features from both modalities in forming enriched feature embeddings, that are consistent with semantics of identity, thereby allowing improved identity recognition. 2) We propose to impose orthogonality constraints on the fused embeddings. They are not only coherent with the angular characteristic of the commonly employed classification loss but are very efficient as they operate directly on mini-batches. 3) Experimental results on largescale VoxCeleb1 [5] show the effectiveness of our method on both face-voice verification and matching tasks. Further, we note that our method performs favourably against the existing state-of-the-art methods. 4) We perform a thorough ablation study to analyze the impact of different components and empirically show that the proposed supervision formulation for face-voice retrieval is more effective and efficient than the ones employed by the contemporary works.

A preliminary version of this work was published in [16]. In addition, the current manuscript makes following new contributions. We study the effectiveness of our approach in the challenging setting of establishing face-voice association across multiple languages spoken by the same set of persons. The experiments are performed on Multilingual Audio-Visual (MAV-Celeb) dataset [12]. Our results indicate that, when training on language 'A' and testing on *unheard* language 'B', the performance in face-voice association tasks deteriorates notably, and this performance is significantly poor when the same model (i.e. trained on language 'A') is tested on *heard* 

language 'A'. However, the performance in the former setting (*unheard* language 'B') is still better than random, which is of significance considering the challenging nature and the configuration of the proposed evaluation protocol. Our analysis reveal that this performance drop is due to notable mismatch between the feature distributions of the two languages, typically known as the domain shift problem. Finally, to our knowledge, we provide a first rigorous comparison of loss functions employed in existing face-voice association methods by utilizing the same underlying branch architecture, input features, and evaluation protocols on VoxCeleb1 dataset (see Fig. 3).

## II. RELATED WORK

We begin by discussing face-voice association from a cognitive neuroscience and psychological perspective, and then detail overview of related work on face-voice (F-V) association.

**Cognitive Psychological Perspective.** Studies have shown that humans are capable of associating voices of unknown identities to their faces. In [17], [18], authors performed a studies to show that humans memorize and recall voices of persons which they have seen (learned) previously. In a similar study [1], [2], authors reveal the presence of equivalent amount of information about an identity from audio and visual streams. Hence, we can say that multimodality can improve the perception capability of humans [19]. In [19], authors show that human perceptual accuracy improves when visual modality supplements the auditory modality in a noisy environment. We note that these studies have been conducted previously under the cognitive psychological perspective in detail and have recently inspired the vision community [3], [4].

Face-voice Association. The work of Nagrani et al. [3] leveraged audio and visual information to establish an association



Fig. 2. Architecture of our method. Fundamentally, it is a two-stream pipeline which generates face and voice embeddings. We believe that an enriched embedding coupled with an effective yet efficient supervision is key in learning a discriminative joint embedding space for improved face-voice association. To this end, we propose a light-weight, plug-and-play mechanism, dubbed as fusion and orthogonal projection (FOP) (shown in dotted red box). It is capable of exploiting complementary cues in both modalities and clustering the enriched fused embeddings based on their identity labels via imposing orthogonality constraints with the traditional classification loss (i.e. softmax cross-entropy).

between faces and voices in a cross-modal biometric matching task. Similarly, some recent metric learning work [7], [4], [9], [12], [20], [21] introduced joint embeddings to establish correlation between face and voice of an identity. These methods extract audio and face embeddings and then minimize the distance between embeddings of same identities while maximize the distance among embeddings from different ones. Zheng et al. [22] proposes Adversarial-Metric Learning model combing metric and adversarial learning. The aim of adversarial learning is to generate modality independent representation while the metric learning learns a metric for audio-visual cross-modal matching and retrieval tasks. Wen et al. [10] presented a disjoint mapping network to learn a shared representation for audio and visual information by mapping them individually to common covariates (gender, nationality, identity) without needing to construct pairs or triplets at the input. Similarly, Nawaz et al. [8] extracted audio and visual information with a single stream network to learn a shared deep latent representation, leveraging identity centroids to eliminate the need of pairs or triplets [4], [3]. Both Wen et al. [10] and Nawaz et al. [8] show that effective face-voice representations can be learned without pairs or triplets formation. More recently, Ning et al. [23] proposed to disentangle the alignable latent identity factors and nonalignable the modality dependent factors for cross-modal biometric matching.

In contrast to existing, our method differs in following respects. First, it proposes to construct enriched embeddings via exploiting complementary cues from the embeddings of both modalities through a attention-based fusion. Second, it clusters the embeddings of same identity and separates embeddings of different identities via orthogonality constraints. The instantiation of both proposals in a two-stream pipeline results in an effective and efficient face-voice association framework.

## III. OVERALL FRAMEWORK

Our aim is to learn a discriminative joint face-voice embedding that is amenable to a multitude of tasks such as face-voice/voice-face matching, verification and retrieval. To actualize this, we develop a new framework for cross-modal face-voice association (See Fig. 2) that is fundamentally a twostream pipeline (sec. III-A) and features a light-weight module that exploits complementary cues from both face and voice embeddings and facilitates discriminative identity mapping via orthogonality constraints (sec. III-B).

## A. Preliminaries

**Problem Settings.** Without the loss of generality, we consider cross-modal retrieval of bimodal data, i.e., for face and voice. Given that we have N instances of face-voice pairs,  $\mathcal{D} = \{(x_i^f, x_i^v)\}_{i=1}^N$ , where  $x_i^f$  and  $x_i^v$  are the face and voice examples of the  $i_{th}$  instance, respectively. Each pair of an instance  $(x_i^f, x_i^v)$  has an associated label  $y_i \in \{0, 1\}$ , where  $y_i = 1$  if  $x_i^f$  and  $x_i^v$  belong to the same identity and  $y_i = 0$  if  $x_i^f$  and  $x_i^v$  belong to a different identity. Both face and voice embeddings typically lie in different representation spaces owing to their different superficial statistics and are mostly unaligned semantically, rendering them incomparable for cross-modal tasks. Cross-modal learning aims at projecting both into a common yet discriminative representation space, where they are sufficiently aligned and instances from the same identity are nearby while from a different are far apart.

Two-stream pipeline. We employ a two-stream pipeline [4] to obtain the respective feature embeddings of both face and voice inputs. The first stream corresponds to a pre-trained convolutional neural network (CNN) on image modality. We take the penultimate layer's output, denoted as  $b_i$ , of this CNN as the feature embeddings for an input face image. Likewise, the second stream is a pre-trained audio encoding network that outputs a feature embedding, denoted as  $e_i$ , for an input audio signal (typically a short-term spectrogram). Existing approaches handling face-voice retrieval [4], [8], mostly resort to triplet and contrastive objectives with carefully crafted negative mining strategies, which significantly increases computational time and are performance-sensitive, to learn a discrminative embedding space. To this end, we introduce a light-weight mechanism that exploits complementary cues from both modality embeddings to form enriched fused embeddings and imposes orthogonal constraints on them for learning discriminative joint face-voice embeddings.

#### B. Learning Discriminative Joint Embedding

In this section, we first describe extracting complementary cues, via multimodal fusion, from both face and voice embeddings obtained through their respective pre-trained networks. We then discuss clustering fused embeddings belonging to the same identity and pushing away the ones with different identity via orthogonality constraints.

Prior to multimodality fusion, we project the face embeddings  $\mathbf{b}_i \in \mathbb{R}^F$  to a new d-dimensional embedding space  $\mathbf{u}_i \in \mathbb{R}^d$  with a fully-connected layer. Similarly, we project the voice embedding  $\mathbf{e}_i \in \mathbb{R}^V$  to a similar d-dimensional embedding space  $\mathbf{v}_i \in \mathbb{R}^d$  with another fully-connected layer. We then L2 normalize both  $\mathbf{u}_i$  and  $\mathbf{v}_i$  which can now be fused to get  $\mathbf{l}_i$ , using the procedure described next.

**Multimodal fusion.** We propose to extract complementary features from both modalities, some of which could be related to age, gender and nationality, to form an enriched unified feature representation which is crucial towards learning a discriminative joint embedding space. Inspired by [24], [25], we employ an attention mechanism to first compute the attention scores (affinity) between the embeddings of two modalities and then fuse these individual modality embeddings after recalibrating them with the attention scores (see Fig. 2). We compute attention scores k between  $u_i$  and  $v_i$  as:

$$\mathbf{k} = \sigma(F_{att}([\mathbf{u}_i, \mathbf{v}_i]), \tag{1}$$

where  $\sigma$  is a sigmoid operator, and  $F_{att}$  are the attention layers. Finally, we fuse  $\mathbf{u}_i$  and  $\mathbf{v}_i$  after modulating them with the attention scores  $\mathbf{k}$  to obtain the fused embeddings  $\mathbf{l}_i$  as:

$$\mathbf{l}_{i} = \mathbf{k} \odot tanh(\mathbf{u}_{i}) + (1 - \mathbf{k}) \odot tanh(\mathbf{v}_{i}), \qquad (2)$$

where  $\odot$  is element-wise multiplication.

Supervision via orthogonality constraints. We want the fused embeddings to encapsulate the semantics of the identity. In other words, these embeddings should be able to predict the identity labels with good accuracy. This is possible if the instances belonging to the same identity are placed nearby whereas the ones with different identity labels are far away. A popular choice to achieve this is softmax cross entropy (CE) loss, which also allows stable and efficient training. Specifically, we use an identity linear classifier with weights denoted as  $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, ..., \mathbf{w}_C] \in \mathbb{R}^{d \times C}$  to compute the logits corresponding to  $\mathbf{l}_i$ . Where d is the dimensionality of embeddings and C is the number of identities. Now, identity classification loss with fused embeddings is computed as:

$$\mathcal{L}_{CE} = -log \frac{exp(\mathbf{l}_i^T \mathbf{w}_{y_i})}{\sum_{j=1}^C exp(\mathbf{l}_i^T \mathbf{w}_j)}$$
(3)

Since softmax CE loss does not enforce margins between pair of identities, it is prone to constructing differently-sized class regions which affects identity separability [26], [27]. Some works attempt to include margin between classes in the Euclidean space [15], [28], which is not well synergized with the CE loss as it achieves separation in the angular domain. Therefore, we propose to impose orthogonality constraints on the fused embeddings to explicitly minimize intra-identity separation while maximizing inter-identity separability [29]. These constraints complement better with the innate angular characteristic of CE loss. Further, since they directly operate on mini-batches, they show greater training efficiency compared to the complex negative mining procedures required in contrastive and triplet loss formulations [4], [30] (sec. IV). Formally, the constraints enforce fused embeddings of different identities to be orthogonal and the fused embeddings with same identity to be similar:

$$\mathcal{L}_{OC} = 1 - \sum_{i,j \in B, y_i = y_j} \langle \mathbf{l}_i, \mathbf{l}_j \rangle + \left| \sum_{i,j \in B, y_i \neq y_k} \langle \mathbf{l}_i, \mathbf{l}_k \rangle \right|, \quad (4)$$

where  $\langle ., . \rangle$  is the cosine similarity operator, and *B* represents the mini-batch size. The first term in Eq. 4 ensures intraidentity compactness, while the second term enforces interidentity separation. Note, the cosine similarity involves the normalization of fused embeddings, thereby projecting them to a unit hyper-sphere:

$$\langle \mathbf{l}_{i}, \mathbf{l}_{j} \rangle = \frac{\mathbf{l}_{i} \cdot \mathbf{l}_{j}}{\|\mathbf{l}_{i}\|_{2} \cdot \|\mathbf{l}_{j}\|_{2}}.$$
 (5)

**Overall Training Objective.** To train the proposed framework, we minimize the joint loss formulation, comprising of  $\mathcal{L}_{CE}$  and  $\mathcal{L}_{OC}$  as:

$$\mathcal{L} = \mathcal{L}_{CE} + \alpha \mathcal{L}_{OC},\tag{6}$$

where  $\alpha$  balances the contribution of two terms in  $\mathcal{L}$ . We empirically set  $\alpha$  to 1.0 based on validation set performance. It is important to mention that CE loss operates in logit space and orthogonal constraints are imposed in the embedding space, however, both of them synergizes well with each other owing to their common angular domain characteristic.

### C. Existing Loss Formulations in F-V Methods

In this section, we formally overview several existing loss formulations typically employed in existing face-voice association methods with an aim to provide a fair comparison with ours, see Fig. 3.

**Center Loss.** It simultaneously learns class centers from features in a mini-batch and penalizes the distance between each class center and corresponding features [31]. Recently, Nawaz et al. [8] introduces a single stream architecture to extract audio-visual information to bridge the gap between them, leveraging center loss. The loss is formulated as follow:

$$\mathcal{L}_{C} = \frac{1}{2} \sum_{i=1}^{b} \|\mathbf{I}_{i} - \mathbf{c}_{y_{i}}\|_{2}^{2}$$
(7)

The  $c_{y_i}$  represents the  $y_i$ th class center of features. Wen et al. [15] observed that the center loss is very small which may degraded to zeros, thus, it is jointly trained with CE loss as follow:

$$\mathcal{L} = \mathcal{L}_{CE} + \alpha_c \mathcal{L}_C \tag{8}$$

A scalar value  $\alpha_c$  is used for balancing center loss and CE loss.

**Git Loss.** It improves center loss by maximizing the distance between features belonging to different classes (push) while keeping features of the same class compact (pull) [28]. The loss is formulated as follow:

ı.



Fig. 3. Multimodal frameworks based on two-stream network to embed face and voice modalities in the latent space. We compare our loss formulation with various others typically adopted in existing face-voice association methods, including Contrastive [4], Triplet [3], Center [8] and Git [28], under the same feature embeddings, network architecture, and evaluation protocol with an aim to provide a fair comparison.

$$\mathcal{L}_{G} = \sum_{i,j=1, i \neq j}^{b} \frac{1}{1 + \left\| \mathbf{I}_{i} - \mathbf{c}_{y_{j}} \right\|_{2}^{2}}$$
(9)

Git loss is trained jointly trained with center loss and CE loss as follow:

$$\mathcal{L} = \mathcal{L}_{CE} + \alpha_c \mathcal{L}_C + \alpha_q \mathcal{L}_G \tag{10}$$

Scalar values  $\alpha_c$  and  $\alpha_g$  are used for balancing center loss, git loss and CE loss.

**Contrastive Loss [4].** It construct positive and negative pairs to learn joint embedding. A set  $\rho$  is formulated comprising of training pairs  $\{(\mathbf{u}_i, \mathbf{v}_i)\}$  using d-dimensional face  $(\mathbf{u}_i)$  and voice  $(\mathbf{v}_i)$  embeddings. Each pair has an associated label  $y_i \in \{0, 1\}$ , where  $y_i = 1$  if  $\mathbf{u}_i$  and  $\mathbf{v}_i$  belong to the same class and  $y_i = 0$  if  $\mathbf{u}_i$  and  $\mathbf{v}_i$  belong to different class.

**Triplet Loss [3].** It encourages that dissimilar pairs be distant from similar pairs by a certain margin value by formulating a triplet. A triplet consists of an anchor, positive and negative sample to formulate a triplet. Given d-dimensional face embedding  $\mathbf{u}_i$  and voice embeddings  $\mathbf{v}_i$ , we take a voice anchor  $\mathbf{v}_i^a$  and choose a positive face  $\mathbf{u}_i^p$  and negative face  $\mathbf{u}_i^n$  constructing a triplet { $\mathbf{v}_i^a, \mathbf{u}_i^p, \mathbf{u}_i^n$ } such that:

$$\||\mathbf{v}_{i}^{a} - \mathbf{u}_{i}^{p}||_{2}^{2} < ||\mathbf{v}_{i}^{a} - \mathbf{u}_{i}^{n}||_{2}^{2} \quad \forall (v_{i}^{a}, u_{i}^{p}, u_{i}^{n}) \in \rho$$
(11)

where  $\rho$  is the set of all triplets. In this work, we have used negative hard mining to generate triplets [30]. The loss function can now be defined as:

$$\mathcal{L}_{triplet} = \sum_{i=1}^{N} [||\mathbf{v}_{i}^{a} - \mathbf{u}_{i}^{p}||_{2}^{2} - ||\mathbf{v}_{i}^{a} - \mathbf{v}_{i}^{n}||_{2}^{2} + m]$$
(12)

where m is margin factor which helps keeping a distance between anchor and positive and negative samples. The loss function minimizes the distance between the anchor voice and positive face i.e. push  $||\mathbf{v}_i^a - \mathbf{u}_i^p||$  and penalize  $||\mathbf{v}_i^a - \mathbf{u}_i^n||$  to be greater than  $||\mathbf{v}_i^a - \mathbf{u}_i^p|| + m$ . Intuitively, this helps in producing closely packed clusters of samples from same class while simultaneously pushing apart clusters of other classes.

## IV. EXPERIMENTS

We perform various cross-modal experiments on *verification* and *matching* tasks to evaluate the effectiveness of our method and provide ablation study with analysis to show the impact of different components in our method. In addition, we analyze the impact of multiple languages on face-voice association.

#### A. Experimental Setup

**Training Details and Datasets.** We train our method on Quadro P5000 GPU for 50 epochs using a batch-size of 128 using Adam optimizer with exponentially decaying learning rate (initialised to  $10^{-5}$ ). We extract face and voice embeddings from VGGFace [32] and Utterance Level Aggregation [33]. Note that, we only backprop. through FOP module while the weights of face and voice subnetworks remains unchanged.

Nagrani et. al [5] introduced a large-scale dataset of audiovisual human speech videos extracted 'in the wild' from YouTube. Learnable Pins [4] introduced two train/test splits out of this dataset to perform various cross-modal verification and matching tasks. The first split consists of disjoint videos from the same set of speakers while the second split contains disjoint identities. We train the model using these two training sets, allowing us to evaluate on both test sets, the first one for *seen-heard* identities, and the second for *unseen-unheard* identities. Note that we followed the same train, validation and test split configurations as used in [4] for fair comparisons.

Recently Nawaz et al. [12] curated MAV-Celeb dataset to study the impact of language on face-voice association

TABLE I CROSS-MODAL VERIFICATION RESULTS FOR OUR (JOINT) LOSS AND OTHER LOSSES UNDER TWO CONFIGURATIONS AND TWO ERROR METRICS.

Method	EER	AUC	EER	AUC
	Seen-Heard		Unseer	-Unheard
CE Loss	21.8	86.6	26.8	81.7
Center Loss [15], [8]	19.8	88.6	29.7	77.5
Git Loss [28]	19.6	88.9	29.5	77.8
Contrastive Loss [4]	23.4	84.7	29.1	79.5
Triplet Loss [30]	20.7	88.0	27.1	81.4
Ours	19.3	89.3	24.9	83.5

task; it consists of video and audio recordings of celebrities speaking more than one language. The dataset contains two splits English–Urdu (EU) and English–Hindi (EH) to analyze performance measure across multiple languages, leveraging *unseen-unheard* and completely *unheard* evaluation protocol. In this work, we only used EU split in our experiments.

**Experimental Setup.** We perform our experiments on following two tasks: The first task is to perform cross-modal verification where the goal is to verify if an audio segment and a face image belong to the same identity. Two inputs are considered i.e. face and voice and verification between the two depends upon a threshold on the similarity value. The threshold can be adjusted in accordance to wrong rejections of true match and/or wrong acceptance of false match. We report results on standard verification metrics i.e. ROC curve (AUC) and equal error rate (EER). The second task consists of *cross-modal matching* where the goal is to match the input modality to the varying gallery size  $n_c$  which consists of the other modality. In our experiments, we increase  $n_c$  to analyze its impact on the performance. For example, in 1:2 matching task, we are given a modality at input, e.g. face, and the gallery consists of two inputs from other modality, e.g. audio. One of them contains a true match while the other serves as an imposter. We employ matching metric i.e. accuracy to report performance. We perform this task in five settings where in each setting the  $n_c$  is increased as 2, 4, 6, 8, 10.

### B. Results

**Comparison with other F-V losses.** We compare our proposed FOP against various losses typically employed in F-V association methods. The first is *center loss* [15], and is adapted by [8] in a single stream network to learn F-V association for cross-modal verification and matching tasks. The second is *Git loss* [28], which in addition to center loss also maximizes the intra-identity distances of embeddings with other centroids and has shown effectiveness for verification and matching tasks. In implementation, we simply replaced our orthogonal constraints loss formulation ( $\mathcal{L}_{OC}$ ) with center loss and then Git loss while keeping all other settings unaltered. The third one is *Contrastive Loss* that is adapted by [4] to associate face and voice of a person in a cross-modal verification task. The last is *Triplet Loss*, that was leveraged by [3] in their face-voice association framework.

6

TABLE II Theoretical/empirical training complexity of our (joint) loss and others. n: # training instances in a modality, and B is mini-batch size.

Method	Empirical	Theoretical
	Time (s)	Worst Case
CE Loss	.02	$\mathcal{O}(n)$
Center Loss [15], [8]	6.8	$\mathcal{O}(n+\frac{n^2}{B_0})$
Git Loss [28]	6.2	$\mathcal{O}(n + \frac{n^2}{B})$
Contrastive Loss [4]	568.2	$\mathcal{O}(n^2)$
Triplet Loss [30]	619.7	$\mathcal{O}(n^3)$
Ours	0.7	$\mathcal{O}(n)$

In implementation, we remove our FOP module in the overall architecture and plug first contrastive and then triplet loss with hard negative mining strategy [30] while keeping the rest of settings fixed.

Table I shows the comparison of our supervision formulation with center and Git losses, and our FOP mechanism with the contrastive and triplet losses on cross-modal verification task with seen-heard and unseen-unheard configurations. We observe that our proposal performs better than other loss formulations across all configurations and both error metrics. We can attribute this improved performance across full spectrum to two possible reasons. First, our orthogonality constraints don't depend on some margin hyperparameter which is certainly difficult to tune with large number of instances and identities (as in VoxCeleb1). Second, as opposed to centroid based losses e.g., center or Git, these orthogonality constraints complement well with the innate angular characteristic of softmax CE loss. We compare both the theoretical and empirical runtime training complexity of our (joint) loss formulation and various others commonly employed in existing F-V methods (Table II). The empirical run-time is the time to complete one epoch. Our loss formulation is superior than all others listed in terms of both theoretical and empirical training efficiency. To further validate the effectiveness of our proposed FOP mechanism, we examine the effect of Gender (G), Nationality (N), Age (A) and its combination (GNA) separately, which influence both face and voice verification (Table III). FOP achieves consistently better performance on G, N, A and the combination (GNA) in both seen-heard and unseen-unheard configurations than other loss formulations.

Furthermore, we compare our FOP against aforementioned loss functions on a cross-modal matching task,  $1 : n_c$  with  $n_c = 2, 4, 6, 8, 10$  in Fig. 4. In this setting, the input is voice while the matching gallery consists of faces. For example, consider the case where the input is voice and the task is 1 : 2matching, we find that out the entry in gallery which best matches the input. We see that our proposed FOP outperforms the counterpart loss formulations for all values of  $n_c$ . Note, this is a relatively challenging task as the common trend for all competitors is that upon increasing  $n_c$  the matching performance deteriorates. In addition to the numerical results, we present qualitative results in Fig. 5.

Comparison with state-of-the-art. Table IV compares our

 TABLE III

 CROSS-MODAL BIOMETRICS RESULTS UNDER VARYING DEMOGRAPHICS FOR seen-heard AND unseen-unheard CONFIGURATIONS.

Demographic	Random	G	Ν	А	GNA	Random	G	Ν	А	GNA
Seen-Heard			Unseen-Unheard							
CE	86.6	78.0	85.0	86.3	77.3	81.7	65.9	53.6	76.0	52.8
Center	88.6	79.2	87.0	88.2	78.1	77.5	62.4	51.7	72.5	54.2
Git	88.9	79.7	87.4	88.6	78.5	77.9	62.6	51.8	72.8	54.2
Contrastive	84.7	69.7	83.7	84.5	69.2	79.5	61.0	53.5	74.7	51.8
Triplet	88.0	76.3	86.7	87.6	75.6	81.7	65.5	53.4	76.3	52.2
Ours	89.3	76.7	87.9	88.6	76.6	83.5	68.8	54.9	78.1	54.2



Fig. 4. Cross-modal matching results with varying gallery size for our proposed FOP and different loss formulations employed by the existing face-voice association methods.

method against existing state-of-the-art works (DIMNet [10], Learnable Pins [4], MAV-Celeb [12], Single Stream Network [8], Multi-view Approach [34], Adversarial-Metric Learning [22], Disentangled Representation Learning [23], Voice-face Discriminative Network [20]). Our method shows the best performance on unseen-unheard protocol, however, it achieves the second best performance on seen-heard protocol. Adversarial-Metric Learning [8] records top performance on seen-heard configuration. We further show a comparison of our method with existing state-of-the-art methods on crossmodal matching (Fig. 6). In particular, we perform the  $1: n_c$ matching tasks, where  $n_c = 2, 4, 6, 8, 10$  and report the results. Our method outperforms [4] while achieves competitive performance against DIMNet [10]. Note that, in addition to identity, DIMNet leverages two additional source of supervisions from gender and nationality, which are not always readily available.

Ablation study and analysis. Table V reports the results when varying the hyperparameter  $\alpha$  that is used to balance the contribution of softmax CE and orthogonal constraints based loss in our joint formulation (Eq. 6). Our method performance is mostly robust to the choice of  $\alpha$ . There is only a 0.8% drop in EER when varying  $0.5 \le \alpha \le 2.0$  around the best reported value of 1.0. However, the performance starts to deteriorate beyond 0.5 and above 2.0, where the contribution of softmax CE, which synergizes with orthogonality constraints, is almost negligible. We also experiment with replacing gated multimodal fusion with a much simpler linear fusion (Table VI). Linear fusion naively sums elementwise the embeddings from two modalities. Results reveal the gated fusion significantly works better than the linear fusion. This highlights the importance of first finding the correlating features between the two modality embeddings, modulating them accordingly, and finally fusing them.

**Decomposing Orthogonal Constraints (OC).** Taking a step further, we break down OC into its sub-components as defined in Eq. 4. These components are employed to quantify the intra-class compactness and inter-class separation in any given embedding space. Therefore, it will be interesting to compare the contribution of OC towards inter-class maximisation and intra-class minimization of an embedding space in contrast to the traditional CE based training scenario. In Fig. 7, the effect of OC on learning face-voice association is visualized. It is evident that OC has a strong impact on overall feature orthogonality in fused embedding space. Moreover, it helps to increase the similarity index between features of similar classes as compared to only CE while simultaneously enforcing higher intra-class dissimilarity.

Impact of multiple languages on face-voice association. MAV-Celeb dataset provides language annotation of celebrities to analyze the impact of multiple languages on cross-modal verification and matching tasks. For example, a celebrity named 'Imran Khan' has audio information in Urdu and English languages. Thus, we can a train a model on one language and tested on an unheard one. Table VII shows the performance of face-voice association on cross-modal verification task across multiple languages using unseen-unheard and completely unheard configuration. The performance is dropped when our face-voice association model is trained on language 'A' (English) and tested on a completely unheard 'B' (Urdu) language. Similarly, Fig. 8 and 9 show the performance of face-voice association on cross-modal matching task with varying gallery size  $(n_c)$ . Interestingly, the performance is better than random verification and matching, which is not trivial considering the challenging configuration unseen-unheard and completely unheard evaluation protocol. In order to find the reason of performance drop, we use feature distributions of three classes in test sets belonging to Urdu and English languages, see Fig. 10. We observed that the performance degradation is due to different feature distributions of the two languages. In addition, both languages share common characteristics.



Fig. 5. Qualitative evaluation of V-F matching task: We evaluate the performance of our method by examining the failure cases of CE Loss. A voice query is shown on the left, and a gallery of 10 template images are shown on the right. There are three successful predictions by our proposed method (highlighted in green). On the last row the wrong prediction is highlighted in red whereas ground truth is highlighted in green. It is clear that the embedding produced with FOP is discriminative in contrast to the embedding produced by the traditional CE loss. Best viewed zoomed in and in colour.

TABLE IV CROSS-MODAL VERIFICATION RESULTS ON *seen-heard* AND *unseen-unheard* CONFIGURATIONS OF OUR METHOD AND EXISTING STATE-OF-THE-ART METHODS.

Methods	Dataset	EER	AUC	EER	AUC
		Seen	Heard	Unseen	-Unheard
Voice-face Discriminative Network [20]	VoxCeleb2	-	-	22.5	85.4
DIMNet [10]	VoxCeleb1	-	-	24.9	-
Learnable Pins [4]	VoxCeleb1	21.4	87.0	29.6	78.5
MAV-Celeb [12]	VoxCeleb1	-	-	29.0	78.9
Single Stream Network [8]	VoxCeleb1	17.2	91.1	29.5	78.8
Multi-view Approach [34]	VoxCeleb1	-	-	28.0	-
Adversarial-Metric Learning [22]	VoxCeleb1	-	92.3	-	80.6
Disentangled Representation Learning [23]	VoxCeleb1	-	-	25.0	84.6
Ours	VoxCeleb1	19.3	89.3	24.9	83.5



Fig. 6. Cross-modal matching results of our method and existing state-of-the-art methods with varying gallery size.

TABLE V COSS-MODAL VERIFICATION RESULTS ON *unseen-unheard* CONFIGURATIONS TO ILLUSTRATE THE EFFECT OF VARYING  $\alpha$  VALUE.

$\mathcal{L}_{CE} + \alpha \mathcal{L}_{OC}$						
$\alpha$	0.0	0.1	0.5	1.0	2.0	5.0
EER	26.8	26.1	25.8	24.9	25.9	26.0
AUC	81.7	82.4	82.8	83.5	82.7	82.6

TABLE VI COSS-MODAL VERIFICATION RESULTS ON *unseen-unheard* CONFIGURATIONS WITH LINEAR AND GATED FUSION STRATEGIES.

Fusion Strategy	EER	AUC
Linear Fusion	25.6	82.7
Gated Fusion	24.9	83.5

#### TABLE VII

CROSS-MODAL VERIFICATION BETWEEN FACE AND VOICE ACROSS MULTIPLE LANGUAGE ON VARIOUS TEST CONFIGURATIONS OF *MAV-Celeb* DATASET. THE DOWN ARROW( $\downarrow$ ) REPRESENTS PERCENTAGE DECREASE IN PERFORMANCE. (**EER: LOWER IS BETTER**)

		EU			
Method	Configuration	Eng. test (EER)	Urdu test (EER)		
MAV-Celeb [12]	Eng. train	45.1	48.3↓7.1		
	Urdu train	47.0↓6.3	44.3		
Ours	Eng. train	31.0	36.6↓17.4		
	Urdu train	44.4↓41.7	31.2		

#### V. CONCLUSION

We proposed a light-weight module, namely FOP, for F-V association task. It attempts to harness the best in both face and voice modalities through attention-based fusion and clusters the fused embeddings based on their identity-labels via orthogonality constraints. We integrated this module in a two-stream pipeline, used for extracting face and voice embeddings, and the resulting overall framework is evaluated on a large-scale VoxCeleb1 dataset for F-V matching and verification tasks. Our method performs favourably against the existing state-of-the-art methods and proposed FOP outperforms competitors both in accuracy and efficiency. Moreover, we study the performance of our method in the challenging setting of F-V association across multiple languages. Results indicate the opportunity of developing more effective F-V methods that are robust to varying feature distributions in



Fig. 7. (a) Feature Orthogonality ( $\downarrow$ ) (b) Similarity of same class features ( $\uparrow$ ) (c) Similarity of different class features ( $\downarrow$ ). FOP simultaneously achieves higher inter-class similarity and intra-class dissimilarity in comparison with the standard CE.



Fig. 8. Cross-modal matching results of our method on *heard* language (Urdu) and completely *unheard* language (English).



Fig. 9. Cross-modal matching results of our method on *heard* language (English) and completely *unheard* language (Urdu).

cross-modal verification and matching tasks.

#### REFERENCES

- Miyuki Kamachi, Harold Hill, Karen Lander, and Eric Vatikiotis-Bateson, "Putting the face to the voice': Matching identity across modality," *Current Biology*, vol. 13, no. 19, pp. 1709–1714, 2003.
- [2] Pascal Belin, Shirley Fecteau, and Catherine Bedard, "Thinking the voice: neural correlates of voice perception," *Trends in cognitive sciences*, vol. 8, no. 3, pp. 129–135, 2004.
- [3] Arsha Nagrani, Samuel Albanie, and Andrew Zisserman, "Seeing voices and hearing faces: Cross-modal biometric matching," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8427–8436.
- [4] Arsha Nagrani, Samuel Albanie, and Andrew Zisserman, "Learnable pins: Cross-modal embeddings for person identity," in *Proceedings of* the European Conference on Computer Vision (ECCV), 2018, pp. 71–88.
- [5] Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, "Voxceleb: a large-scale speaker identification dataset," *arXiv preprint arXiv*:1706.08612, 2017.
- [6] Shota Horiguchi, Naoyuki Kanda, and Kenji Nagamatsu, "Face-voice matching using cross-modal embeddings," in *Proceedings of the 26th* ACM international conference on Multimedia, 2018, pp. 1011–1019.
- [7] Changil Kim, Hijung Valentina Shin, Tae-Hyun Oh, Alexandre Kaspar, Mohamed Elgharib, and Wojciech Matusik, "On learning associations of faces and voices," in *Asian Conference on Computer Vision*. Springer, 2018, pp. 276–292.
- [8] Shah Nawaz, Muhammad Kamran Janjua, Ignazio Gallo, Arif Mahmood, and Alessandro Calefati, "Deep latent space learning for cross-modal mapping of audio and visual signals," in 2019 Digital Image Computing: Techniques and Applications (DICTA). IEEE, 2019, pp. 1–7.
- [9] 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," arXiv preprint arXiv:2103.07293, 2021.
- [10] Yandong Wen, Mahmoud Al Ismail, Weiyang Liu, Bhiksha Raj, and Rita Singh, "Disjoint mapping network for cross-modal matching of voices and faces," in 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019, 2019.
- [11] Boqing Zhu, Kele Xu, Changjian Wang, Zheng Qin, Tao Sun, Huaimin Wang, and Yuxing Peng, "Unsupervised voice-face representation learning by cross-modal prototype contrast," *arXiv preprint* arXiv:2204.14057, 2022.
- [12] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1682–1691.
- [13] Cho-Ying Wu, Chin-Cheng Hsu, and Ülrich Neumann, "Cross-modal perceptionist: Can face geometry be gleaned from voices?," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 10452–10461.



Fig. 10. Features distribution of English and Urdu test sets to show domain gap. First row represents feature distribution of English and Urdu test sets with a model trained on English language while the second row represents distribution of English and Urdu test sets with a model trained on Urdu language.

- [14] Hao Zhu, Wayne Wu, Wentao Zhu, Liming Jiang, Siwei Tang, Li Zhang, Ziwei Liu, and Chen Change Loy, "CelebV-HQ: A large-scale video facial attributes dataset," in *ECCV*, 2022.
- [15] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao, "A discriminative feature learning approach for deep face recognition," in *European conference on computer vision*. Springer, 2016, pp. 499–515.
- [16] Muhammad Saad Saeed, Muhammad Haris Khan, Shah Nawaz, Muhammad Haroon Yousaf, and Alessio Del Bue, "Fusion and orthogonal projection for improved face-voice association," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7057–7061.
- [17] Hunter A McAllister, Robert HI Dale, Norman J Bregman, Allyssa Mc-Cabe, and C Randy Cotton, "When eyewitnesses are also earwitnesses: Effects on visual and voice identifications," *Basic and Applied Social Psychology*, vol. 14, no. 2, pp. 161–170, 1993.
- [18] Stefan R Schweinberger, David Robertson, and Jürgen M Kaufmann, "Hearing facial identities," *Quarterly Journal of Experimental Psychology*, vol. 60, no. 10, pp. 1446–1456, 2007.
- [19] Kevin G Munhall and Eric Vatikiotis-Bateson, "The moving face during speech communication," *Hearing by eye II: Advances in the psychology* of speechreading and auditory-visual speech, pp. 123–139, 1998.
- [20] Ruijie Tao, Rohan Kumar Das, and Haizhou Li, "Audio-Visual Speaker Recognition with a Cross-Modal Discriminative Network," in *Proc. Interspeech 2020*, 2020, pp. 2242–2246.
- [21] Guangyu Chen, Deyuan Zhang, Tao Liu, and Xiaoyong Du, "Self-lifting: A novel framework for unsupervised voice-face association learning," in *Proceedings of the 2022 International Conference on Multimedia Retrieval*, 2022, pp. 527–535.
- [22] Aihua Zheng, Menglan Hu, Bo Jiang, Yan Huang, Yan Yan, and Bin Luo, "Adversarial-metric learning for audio-visual cross-modal matching," *IEEE Transactions on Multimedia*, vol. 24, pp. 338–351, 2021.
- [23] Hailong Ning, Xiangtao Zheng, Xiaoqiang Lu, and Yuan Yuan, "Disentangled representation learning for cross-modal biometric matching," *IEEE Transactions on Multimedia*, vol. 24, pp. 1763–1774, 2021.
- [24] John Arevalo, Thamar Solorio, Manuel Montes-y Gómez, and Fabio A González, "Gated multimodal units for information fusion," arXiv preprint arXiv:1702.01992, 2017.
- [25] Zhengyang Chen, Shuai Wang, and Yanmin Qian, "Multi-modality matters: A performance leap on voxceleb," *Proc. Interspeech 2020*, pp. 2252–2256, 2020.
- [26] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 4690–4699.

- [27] Munawar Hayat, Salman Khan, Syed Waqas Zamir, Jianbing Shen, and Ling Shao, "Gaussian affinity for max-margin class imbalanced learning," in *Proceedings of the IEEE/CVF International Conference* on Computer Vision, 2019, pp. 6469–6479.
- [28] Alessandro Calefati, Muhammad Kamran Janjua, Shah Nawaz, and Ignazio Gallo, "Git loss for deep face recognition," in *Proceedings* of the British Machine Vision Conference (BMVC), 2018.
- [29] Kanchana Ranasinghe, Muzammal Naseer, Munawar Hayat, Salman Khan, and Fahad Shahbaz Khan, "Orthogonal projection loss," arXiv preprint arXiv:2103.14021, 2021.
- [30] Florian Schroff, Dmitry Kalenichenko, and James Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings* of the IEEE conference on computer vision and pattern recognition, 2015, pp. 815–823.
- [31] Liwei Wang, Yin Li, and Svetlana Lazebnik, "Learning deep structurepreserving image-text embeddings," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5005–5013.
- [32] O Parkhi, A Vedaldi, and A Zisserman, "Deep face recognition," 2019, pp. 1–12, British Machine Vision Association.
- [33] Weidi Xie, Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, "Utterance-level aggregation for speaker recognition in the wild," in ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2019, pp. 5791–5795.
- [34] Leda Sari, Kritika Singh, Jiatong Zhou, Lorenzo Torresani, Nayan Singhal, and Yatharth Saraf, "A multi-view approach to audio-visual speaker verification," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP).* IEEE, 2021, pp. 6194–6198.