

Lightweight Cross-Spectral Face Recognition via Contrastive Alignment and Distillation

Anjith George, *Member, IEEE*, Sébastien Marcel, *Fellow, IEEE*

Abstract—Heterogeneous Face Recognition (HFR) aims at matching face images captured across different sensing modalities, such as thermal-to-visible or near-infrared-to-visible, enhancing the usability of face recognition systems in challenging real-world conditions. Although recent HFR methods have achieved significant improvements in performance, many rely on computationally expensive models, making them impractical for deployment on resource-limited edge devices. In this work, we introduce a lightweight yet effective HFR framework by adapting a hybrid CNN-Transformer model originally developed for RGB homogeneous face recognition. Our approach enables efficient end-to-end training with only a small amount of paired heterogeneous data, while still maintaining strong performance on standard RGB face recognition benchmarks. This makes it suitable for both homogeneous and heterogeneous settings. Comprehensive experiments on several challenging HFR and face recognition benchmarks show that our method achieves state-of-the-art or competitive performance while keeping computational requirements low.

Index Terms—Face Recognition, Heterogeneous Face Recognition, Cross-Spectral Recognition, Lightweight Models, Layer Normalization, Knowledge Distillation

1 INTRODUCTION

Face recognition (FR) has become a ubiquitous modality in biometric authentication, particularly in access control, due to its efficiency and non-intrusive nature. With advances in deep learning, especially convolutional neural networks (CNNs), face recognition has achieved near-human performance even in unconstrained settings [1]. However, most of the existing FR methods are designed for homogeneous environments, where both gallery and probe images are captured using visible-spectrum cameras.

In many practical use cases, such as surveillance, mobile authentication, or defense applications, relying only on visible-light imagery is insufficient. Images captured outside the visible spectrum, such as near-infrared (NIR) or thermal imagery, provide several advantages. For example, NIR is more resistant to lighting variations and more robust against spoofing attempts [2], [3]. Despite these advantages, developing effective FR models for these modalities is difficult, mainly due to the lack of large-scale annotated heterogeneous paired datasets. Heterogeneous Face Recognition (HFR) addresses this challenge by enabling face matching across modalities, for instance, comparing thermal or NIR images with visible-light references [4], [5]. A key aspect of HFR is Cross-Spectral Face Recognition (CFR), which focuses on handling the appearance variations caused by spectral differences between imaging domains. CFR becomes especially important in low-light or long-range scenarios where visible imaging is unreliable.

Although deep neural networks (DNNs) have greatly improved Heterogeneous Face Recognition (HFR), the task remains challenging due to the large modality gap between source and target domains. Models trained on RGB images often fail to generalize to non-RGB inputs [6]. Moreover, the collection of large-scale paired cross-modal datasets is both costly and difficult, which

requires the development of methods capable of generalizing from limited training data. Many state-of-the-art HFR systems also rely on computationally heavy architectures, which are impractical for deployment on edge or mobile devices. Consequently, there is growing interest in lightweight models that maintain competitive accuracy while reducing computational overhead. Vision Transformers (ViTs) have demonstrated a strong capability to capture global dependencies [7], making them a useful complement to CNN-based architectures. These architectures provide an opportunity to design compact yet effective HFR frameworks suitable for real-world, resource-limited environments.

In this work, we propose a parameter-efficient adaptation framework that extends pretrained RGB face recognition models to heterogeneous face recognition without increasing inference complexity. Instead of introducing modality-specific branches or additional network modules, our approach selectively adapts LayerNorm parameters and early convolutional layers while keeping the remainder of the backbone frozen. Combined with contrastive alignment and self-distillation, this strategy enables effective cross-spectral adaptation using only a small amount of paired heterogeneous data while preserving the original RGB recognition performance. We demonstrate this approach using the lightweight EdgeFace [8] architecture as a backbone.

The main contributions of this work are as follows.

- We propose a parameter-efficient adaptation framework that extends pretrained RGB face recognition models to heterogeneous face recognition without increasing inference complexity.
- We show that adapting only LayerNorm parameters and shallow layers, combined with contrastive alignment and self-distillation, enables effective cross-modal learning using limited paired data.
- Extensive experiments across six heterogeneous benchmarks demonstrate competitive or state-of-the-art performance while maintaining significantly lower computational

• All authors are with Idiap Research Institute, Martigny, Switzerland. Sébastien Marcel is also affiliated with Université de Lausanne (UNIL), Lausanne, Switzerland. E-mail: {anjith.george, sebastien.marcel}@idiap.ch
Manuscript received April 19, 2005; revised August 26, 2015.

cost. Code is available at ¹.

2 RELATED WORK

Heterogeneous Face Recognition: focuses on matching faces captured across different imaging modalities such as visible light (VIS), near-infrared (NIR), thermal cameras, or even hand-drawn sketches. The main challenge in HFR is the modality gap, meaning the substantial distribution shift between these modalities. This gap causes standard face recognition models trained solely on RGB images to perform poorly when applied to other domains. To address this issue, recent studies introduce a variety of solutions that generally fall into three categories: learning modality-invariant features, projecting data into a shared representation space, and generating source domain samples through synthesis-based methods.

Invariant feature based approaches focus on learning facial representations that remain stable across different modalities. Early studies focused on handcrafted descriptors such as Difference of Gaussian (DoG) filters, multi-scale LBP [9], SIFT, and MLBP [10] to capture local texture cues. With the rise of deep learning, CNN-based models were introduced to learn modality-invariant features [6], [11], while other works proposed improved handcrafted features like the Local Maximum Quotient (LMQ) descriptor [12] or explored composite feature fusion at the score level [13]. In [14], the authors proposed a NIR-VIS heterogeneous face recognition framework that augments lightweight face-recognition models with Gabor filter-derived invariant features by appending the filter’s imaginary response as an additional input channel, followed by PCA-based dimensionality reduction and Mahalanobis-distance matching. This strategy improves NIR-VIS recognition performance across benchmark datasets while largely preserving VIS-domain performance and incurring only minimal computational overhead.

Common-space projection methods address the domain gap by mapping features from different modalities into a unified latent space. Classical techniques include Canonical Correlation Analysis (CCA) [15], Partial Least Squares (PLS) [16], and various forms of coupled regression [17], all of which use linear or non-linear transformations to preserve discriminative properties while reducing the domain gap. More recent work leverages deep architectures with domain-specific units [18], domain-invariant modules [19], coupled attribute-aware loss functions [20], and semi-supervised collaborative representations [21], enhancing alignment even under limited annotations or unpaired data. Further progress [22], [23] shows that conditioning intermediate feature maps can effectively bridge the modality gap, later extended toward modality-agnostic learning [24].

Synthesis-based methods adopt a different strategy by generating cross-modal images, typically translating inputs into the visible domain, so that standard face recognition models can be directly applied. Early solutions performed patch-level reconstruction via Markov Random Fields [25] or employed manifold learning approaches such as LLE [26]. The advent of GAN-based frameworks, including CycleGAN [27], greatly improved this approach, enabling unpaired translation and photorealistic facial synthesis [28], [29]. More recent advances explore latent disentanglement [30], memory-enhanced transformers for unsupervised reference-guided generation [31], and plug-and-play

modules like Prepended Domain Transformers (PDT) [32], which improve cross-domain representation alignment without explicit image generation. Nonetheless, synthesis-based approaches often introduce significant computational overhead because they require both image translation module and a separate face recognition model for matching.

Recent advances in NIR-VIS heterogeneous face recognition have increasingly focused on reducing the dependence on manually annotated labels, as acquiring large-scale labeled cross-domain datasets is both costly and impractical for real-world deployment. To address this, many recent works adopt unsupervised or semi-supervised learning strategies that exploit pseudo-label generation, contrastive learning, and prototype-based representation learning. For instance, [33] formulates the task as an unsupervised domain adaptation problem and introduces the RPC network, which combines NIR cluster-based pseudo-label sharing with both domain-specific and inter-domain contrastive learning to produce compact and domain-invariant representations, achieving over 99% pseudo-label assignment accuracy and strong benchmark performance. Similarly, [34] addresses the problem in a semi-supervised setting through the LPL framework, which combines cross-domain pseudo-label association, intra-domain compact representation learning, and prototype-based inter-domain invariant learning to iteratively refine cluster structure and extract robust cross-domain identity features, reaching performance comparable to recent supervised methods. Along the same line, [35] proposes HERE (HEterogeneous learning and Residual-invariant Enhancement), an unsupervised framework that employs a homogeneous-to-heterogeneous learning strategy, combining modality-adversarial contrastive learning, cross-modal pseudo-label estimation, refined contrastive learning, and residual-invariant feature enhancement to learn robust modality-invariant representations. Together, these studies demonstrate that competitive NIR-VIS recognition performance can be achieved with limited or no explicit identity supervision.

Lightweight Face Recognition: As mobile devices and edge computing platforms have become ubiquitous, face recognition (FR) research has increasingly prioritized compact models that provide strong performance under tight computational and memory constraints. This shift has driven the development of a wide range of efficient architectures adapted for FR. MobileFaceNets [36], built on the MobileNet family [37], [38], were among the first to achieve high performance with under 1M parameters. MixFaceNets [39] further improved efficiency by integrating MixConv [40], while ShiftFaceNet [41] leveraged ShiftNet operations to achieve competitive results with just 0.78M parameters. ShuffleFaceNet [42], inspired by ShuffleNetV2 [43], introduced model variants ranging from 0.5M to 4.5M parameters without sacrificing accuracy. Neural architecture search has also advanced lightweight FR: PocketNet [44], designed via DARTS on CASIA-WebFace [45] with multi-stage knowledge distillation (KD), and VarGFaceNet [46], winner of the ICCV 2019 LFR challenge [47] used variable group convolutions to optimize efficiency. Recent work, SynthDistill [48], demonstrated that synthetic data [49] combined with online KD can effectively train TinyFaR models [50] to approximate high-performance teacher networks. GhostFaceNets [51] reduced redundancy in convolutional operations, achieving extremely compact designs with as little as 61M FLOPs using depthwise convolutions. Most recently, EdgeFace [8] combined convolutional and transformer components leveraging EdgeNeXt [52] framework and employed low-rank linear layers

¹<https://www.idiap.ch/paper/lightweighthfr>

to cut both parameters and FLOPs, achieving near–state-of-the-art results at a fraction of the complexity and securing the top position among compact models in the IJCB EFAr 2023 challenge [53].

Lightweight Heterogeneous Face Recognition: While lightweight FR architectures are ideal for edge applications, their adaptation to the more challenging Heterogeneous Face Recognition (HFR) has not been addressed in literature. Many current HFR systems rely on heavy backbones or synthesis-based pipelines, both of which introduce substantial computational costs that hinder deployment in resource-limited environments. To address this gap, we introduce a compact and efficient HFR framework tailored for edge devices, with strong cross-modal performance while keeping computational demands and data requirements minimal.

3 PROPOSED APPROACH

Heterogeneous face recognition (HFR) poses a significant challenge largely due to the limited availability of paired cross-modal training data. To address this, a widely used approach is to start with large-scale models pretrained on visible-spectrum (RGB) images and then adapt them to heterogeneous domains. However, directly fine-tuning these models on small HFR datasets often results in overfitting and catastrophic forgetting, where the model’s original RGB recognition performance deteriorates heavily. To mitigate this, prior works [18], [32] have introduced architectural changes such as modality-specific branches or asymmetric processing pathways. While these designs help retain RGB performance, they also increase model size and introduce parameter redundancy, an undesirable trade-off when aiming for lightweight models. Further, many existing approaches assume a fixed representation for the RGB modality and force the other modality (e.g., NIR, thermal, sketch) to align with the source modality in the original latent space. This rigid alignment can limit performance, especially when cross-modal differences are highly nonlinear.

In this work, our goal is to develop a unified model that effectively handles both standard (homogeneous) and heterogeneous face recognition without adding noticeable computational overhead or reducing accuracy. We introduce a lightweight yet robust adaptation strategy for pretrained FR models that prevents catastrophic forgetting while enabling strong cross-modal matching performance. Unlike prior modulation-based approaches such as PDT, CAIM, and SSMB, which introduce modality-specific branches or additional inference-time modules, our method performs parameter-efficient adaptation within the existing pretrained backbone. Specifically, we selectively update LayerNorm parameters and early convolutional layers while keeping the remaining network frozen. This design avoids additional architectural complexity and maintains the original inference cost. Furthermore, the proposed self-distillation objective preserves the original RGB performance, reducing catastrophic forgetting during heterogeneous adaptation.

Our method builds upon EdgeFace [8], a lightweight architecture that integrates convolutional layers with transformer modules. The main idea in our approach is that Layer Normalization [54] plays a useful role in modality adaptation [55]. Instead of modifying the backbone structure or adding redundant pathways, we treat LayerNorm as a modulation mechanism that adjusts modality-specific statistics, enabling the network to learn discriminative features for both RGB and other modalities within a single shared architecture.

To accomplish our goal, we adopt a contrastive self-distillation training approach. The objective consists of two key components:

- Contrastive Modality Alignment, which encourages paired samples (e.g., RGB–NIR) to produce closer embeddings in the shared latent space promoting modality-invariant representations; and
- Self-Distillation Regularization, which preserves the pre-trained model’s RGB accuracy by transferring its knowledge to the adapted model reducing catastrophic forgetting.

Together, these components allow effective fine-tuning on limited HFR data while retaining strong RGB performance and maintaining a lightweight design. The following subsections provide detailed explanations of the training objectives, backbone configuration, and implementation specifics.

LayerNorm Adaptation: Prior work has shown that the statistical characteristics of feature maps in deep neural networks (DNNs) encode key stylistic cues of images first demonstrated by Gatys et al. [55]. This insight has made normalization layers crucial for stabilizing and improving deep model training. Layer Normalization (LayerNorm), introduced by Ba et al. [54], overcomes the limitations of batch normalization by computing normalization statistics over the feature dimension of each individual sample rather than across the batch. As a result, LayerNorm maintains consistent behavior during both training and inference, making it well-suited for settings involving variable input structures or non-i.i.d. data. In large language models (LLMs) and their multimodal variants (MLLMs), recent work by Zhao et al. [56] has shown that selectively fine-tuning LayerNorm parameters inside attention blocks can yield significant efficiency gains while reducing computational cost. This approach outperforms several parameter-efficient fine-tuning techniques, such as Low-Rank Adaptation (LoRA) [57]. LayerNorm’s scale and shift parameters serve as natural modulation points: small updates to these parameters can reshape the distribution of intermediate features across layers without modifying the pretrained backbone. As observed in Xu et al. [58] and Zhao et al. [56], the expected gradient of LayerNorm diminishes with network depth, while the gradient variance remains low, properties indicative of good generalization. These low-magnitude, low-variance gradients make it possible to adapt the model to new domains by updating only the LayerNorm parameters, preserving the stability of the pretrained network. This characteristic is especially valuable in our setting, where robust cross-modal adaptation must be achieved without extensively altering the underlying model weights, preventing catastrophic forgetting.

For modalities such as NIR or Thermal, the modality gap is primarily spectral/photometric rather than geometric. The facial structure remains largely unchanged, while the image statistics seen by the network such as intensity distribution, local contrast, and channel-wise feature responses shift substantially across modalities. In pretrained face recognition models, these low-level statistics are mainly processed in the early convolutional layers and normalization layers. Adapting these components therefore provides a principled way to compensate for modality-dependent shifts at the feature-statistics level, while leaving most of the deeper identity-discriminative representation unchanged.

Problem Setting: We start with a pretrained face recognition network F , whose parameters Θ_{FR} have been learned from a large-scale RGB (visible-spectrum) dataset. Let (X_{s_i}, X_{t_i}, y_i) denote a triplet consisting of a pair of images X_{s_i} and X_{t_i} from the source (e.g., RGB) and target (e.g., NIR or thermal) modalities

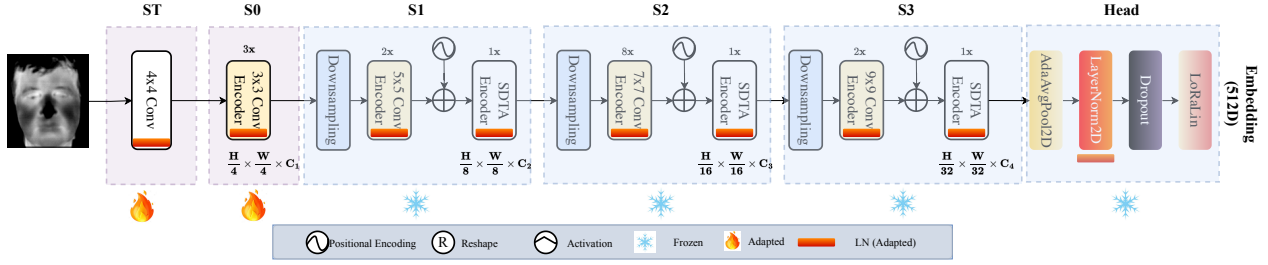


Fig. 1. Model architecture of xEdgeFace models: The highlighted modules (LN-LayerNorm, ST-Conv. Stem, Stages-S0, S1, S2) are adapted while other network components remain frozen. The two loss components ensure modality alignment while preserving source-domain FR performance. Computational complexity remains unchanged in new models.

respectively, and a binary identity label $y_i \in \{0, 1\}$, where $y_i = 1$ indicates that both images correspond to the same identity and $y_i = 0$ otherwise.

The goal is to adapt F into a heterogeneous face recognition network \hat{F} , parameterized by Θ_{HFR} , such that the resulting embeddings $e_{s_i} = \hat{F}(X_{s_i})$ and $e_{t_i} = \hat{F}(X_{t_i})$ are well-aligned in a shared embedding space if they belong to the same identity, while also preserving the discriminative ability of the original model F on the source modality.

We initialize \hat{F} with the pretrained parameters Θ_{FR} , and decompose Θ_{HFR} into three disjoint subsets:

$$\Theta_{\text{HFR}} = \left\{ \Theta_{\text{LN}}^{(1:K)}, \Theta_{\text{Adapted}}, \Theta_{\text{Frozen}} \right\}, \quad (1)$$

where $\Theta_{\text{LN}}^{(1:K)}$ denotes the set of all LayerNorm parameters (from K layers), Θ_{Adapted} includes all trainable parameters except LayerNorms, and Θ_{Frozen} refers to the set of parameters that remain fixed during training.

To enforce alignment between embeddings from different modalities, we use a cosine-based contrastive loss defined as follows:

$$\mathcal{L}_{\text{C}}(e_{s_i}, e_{t_i}, y_i) = y_i \cdot (1 - \cos(e_{s_i}, e_{t_i})) + (1 - y_i) \cdot \max(0, \cos(e_{s_i}, e_{t_i}) - m), \quad (2)$$

where $\cos(e_{s_i}, e_{t_i}) = \frac{e_{s_i} \cdot e_{t_i}}{\|e_{s_i}\|_2 \|e_{t_i}\|_2}$ is the cosine similarity between the embeddings and $m \in [0, 1]$ is a contrastive margin.

To retain the model's original recognition ability on the source modality, we incorporate a self-distillation loss that guides the adapted model \hat{F} to remain consistent with the pretrained model F on source-domain embeddings. Formally, this loss is defined as:

$$\mathcal{L}_{\text{SDL}}(e_{F_{s_i}}, e_{\hat{F}_{s_i}}) = 1 - \cos(e_{F_{s_i}}, e_{\hat{F}_{s_i}}), \quad (3)$$

where $e_{F_{s_i}} = F(X_{s_i})$ is the frozen embedding from the original model and $e_{\hat{F}_{s_i}} = \hat{F}(X_{s_i})$ is the adapted embedding for the same image.

The overall training objective for the adapted network (\hat{F}) combines the contrastive loss for modality alignment with the self-distillation loss that preserves source-domain performance, resulting in the following formulation:

$$\mathcal{L}_{\text{total}} = (1 - \lambda) \cdot \mathcal{L}_{\text{C}}(e_{s_i}, e_{t_i}, y_i) + \lambda \cdot \mathcal{L}_{\text{SDL}}(e_{F_{s_i}}, e_{\hat{F}_{s_i}}), \quad (4)$$

where $\lambda \in [0, 1]$ is a balancing hyperparameter controlling the trade-off between cross-modal alignment and self-regularization.

In all of our experiments, we set $\lambda = 0.75$ and the margin $m = 0$ unless otherwise indicated. This configuration empirically provided the best balance between adapting to the heterogeneous domain and retaining source modality performance.

Although our implementation uses EdgeFace as the backbone, the proposed contrastive alignment and self-distillation objectives are architecture-agnostic. In principle, they can be applied to other pretrained face recognition networks, enabling similar parameter-efficient heterogeneous adaptation without modifying the underlying architecture.

Face Recognition Backbone We use the pretrained EdgeFace [8] model as our face recognition (FR) backbone. EdgeFace is a hybrid convolutional–transformer architecture that employs LayerNorm instead of the more common BatchNorm, enabling stable training across modalities. The model is trained on the large-scale WebFace12M dataset [59], which includes over 12 million RGB images from more than 600,000 identities. All input faces are resized to 112×112 and aligned via a similarity transform to standardize eye locations. For thermal images, which are single-channel, we replicate the channel three times to match the RGB input format required by the backbone.

Implementation Details Our HFR framework incorporates a frozen copy of the pretrained EdgeFace model as a regularization teacher network, guiding the fine-tuning of a shallow, trainable surrogate network through self-distillation (Figure 1). The surrogate is initialized with the pretrained weights, and only selected early modules including LayerNorm layers, are unfrozen for adaptation, while the remaining components stay frozen. To further improve cross-modal consistency, we apply a contrastive loss on the surrogate's embeddings for both RGB and thermal inputs.

The framework is implemented in PyTorch and built upon the Bob library [60], [61]². Training uses the Adam optimizer with a learning rate of 1×10^{-4} , a batch size of 256, and runs for 20 epochs. The contrastive loss margin m is set to 0, and the weighting factor λ is fixed at 0.75 for all experiments. Although both pretrained and surrogate networks are used during training, only the adapted surrogate model is required for inference.

4 EXPERIMENTS

This section reports the results of an extensive set of experiments conducted using the proposed framework. We evaluate heterogeneous face recognition (HFR) performance on established

²<https://www.idiap.ch/software/bob/>

benchmarks and compare our method against state-of-the-art approaches. To verify that the adaptation process does not introduce catastrophic forgetting, we also assess performance on standard face recognition datasets. Across all experiments, cosine distance is used as the similarity metric for evaluation.

4.1 Datasets and Protocols

For our evaluations, we used the following datasets:

Tufts Face Dataset: The Tufts Face Database [62] contains a broad collection of facial images captured across multiple modalities, making it well-suited for heterogeneous face recognition tasks. In our experiments, we followed the *VIS-Thermal* protocol and used the thermal subset of the dataset. Tufts includes 113 identities (39 males and 74 females) covering diverse demographic groups, with each subject represented across several modalities. Following the protocol in [29], we randomly selected 50 identities for training (with 45 for training and 5 as validation) and used the remaining subjects for testing.

MCXFace Dataset: The MCXFace dataset [32], [63] consists of facial images from 51 participants, collected under different illumination conditions across three sessions and multiple sensing channels. These channels include RGB, thermal, near-infrared (850 nm), short-wave infrared (1300 nm), depth, and depth estimated from RGB. The dataset provides five folds, each created by randomly splitting identities into training and development sets. Our evaluations focus on the challenging *VIS-Thermal* protocols, which serve as standard benchmarks for cross-modal face recognition on this dataset.

Polathermal Dataset: The Polathermal dataset [64], collected by the U.S. Army Research Laboratory (ARL), is an HFR dataset containing both polarimetric long-wave infrared (LWIR) imagery and visible-spectrum color images for 60 subjects. In addition to polarimetric data, the dataset provides conventional thermal images for each subject. Following the five-fold partitioning protocol introduced in [18], we use the conventional thermal images, assigning 25 identities for training and the remaining 35 identities for testing.

SCFace Dataset: The SCFace dataset [65] includes high-quality enrollment images paired with low-quality probe images captured in realistic surveillance environments using multiple cameras. It is organized into four evaluation protocols: *close*, *medium*, *combined*, and *far*, with the “far” protocol presenting the highest level of difficulty. In total, the dataset contains 4,160 static images from 130 subjects, recorded in both visible and infrared domains.

CUFSF Dataset: The CUHK Face Sketch FERET Database (CUFSF) [66] consists of 1,194 facial photographs from the FERET dataset [67], each paired with an artist-drawn sketch. Due to the stylized and exaggerated nature of the sketches, CUFSF poses a challenging HFR setting. Following the protocol in [68], we use 250 identities for training and evaluate on the remaining 944 identities.

CASIA NIR-VIS 2.0 Dataset: The CASIA NIR-VIS 2.0 Face Database [69] contains images from 725 subjects captured in both visible (VIS) and near-infrared (NIR) modalities. Each subject has around 1-22 VIS images and 5-50 NIR images. Experiments follow a fixed 10-fold cross-validation protocol with 360 identities used for training. The gallery and probe sets include 358 distinct individuals, ensuring no identity overlap between training and evaluation.

Metrics: We evaluate model performance using several standard metrics commonly adopted in the heterogeneous face recognition literature. These include the Area Under the Curve (AUC), Equal Error Rate (EER), Rank-1 identification accuracy, and Verification Rates at multiple false acceptance rates (0.01%, 0.1%, 1%, and 5%). For datasets with multiple folds, we report the mean performance along with the corresponding standard deviation across folds.

4.2 Model Complexity

The main objective of this work is the development of lightweight models for heterogeneous face recognition. Therefore, it is important to compare the computational footprint of our proposed models with those commonly used in prior literature. We assess computational efficiency using two standard metrics: the number of floating-point operations (in terms of GFLOPs) and the total number of parameters (in millions, denoted as MPARAMS). As summarized in Table 1, the xEdgeFace [70] variants achieve significantly lower computational cost and parameter count, highlighting their suitability for deployment in resource-limited environments. Throughout the remaining comparisons, it is important to note that the xEdgeFace base model has *one-third the parameters* and demands *roughly one-twentieth of the compute* of the state-of-the-art models.

TABLE 1
Comparison of computational complexity between the proposed method and state-of-the-art HFR approaches, reported in terms of floating point operations (GFLOPs) and number of parameters (MPARAMS).

	GFLOPS	MPARAMS
CAIM(1-3) [23]	26.3	65.6
DIU [19]	24.2	65.2
SSMB [24]	24.2	65.5
PDT [32]	24.2	65.2
xEdgeFace - Base	1.39	18.23
xEdgeFace - S ($\gamma = 0.5$)	0.31	3.65
xEdgeFace - XS ($\gamma = 0.6$)	0.15	1.77
xEdgeFace - XXS	0.09	1.24

4.3 Ablation Studies

Given the large set of design choices and hyperparameters involved, we begin with a comprehensive ablation study to examine the impact of each component on overall model performance. All ablation experiments are conducted on the Tufts Face Dataset following the *VIS-Thermal* protocol, which represents one of the most challenging heterogeneous face recognition (HFR) settings due to its large modality gap. In the experiments that follow, adapted HFR versions of the base models are denoted as *xEdgeFace*.

Adapting Different Sets of Layers: To identify the most effective subset of layers for adaptation, we perform a series of controlled ablation experiments by selectively unfreezing the LayerNorm (LN) layers, the initial convolutional stem (ST), and successive backbone stages: Stage 0 (S0), Stage 1 (S1), and Stage 2 (S2). In addition to cumulative configurations, we also evaluate individual stages and deeper-stage adaptation settings to better isolate the contribution of each component. The results, presented in Table 2, show that adapting only the LayerNorm layers already provides a significant improvement over the pre-trained baseline, while incorporating the stem and early stages

further enhances performance. Among single-stage adaptations, S0 performs best (51.76%), outperforming S1 and S2, and the same trend is observed when combined with LayerNorm, where (LN, S0) (54.92%) clearly surpasses (LN, S1) and (LN, S2). The best overall result is obtained with (LN, ST, S0) (56.03%), whereas adapting deeper stages provides limited additional benefit and slightly degrades performance. These results indicate that the main gains come from LayerNorm and early-stage adaptation, which offer the best balance between heterogeneous face recognition performance and parameter efficiency in training. At the same time, the average face recognition accuracy remains nearly unchanged across all configurations (96.6–96.8%), confirming that the proposed adaptation improves heterogeneous verification without sacrificing standard RGB recognition.

TABLE 2

Ablation study on the Tufts Face Dataset using different configurations of adapted layers. The verification rate on the Tufts Face dataset and the average accuracy on the face recognition benchmarks are shown.

Adapted Layers	VR@FAR=0.01%	Avg. FR Acc. (%)
LN	32.47	96.62 ± 0.93
ST	13.73	96.83 ± 0.90
LN,ST	48.79	96.64 ± 0.95
LN,ST,S0	56.03	96.72 ± 0.92
LN,ST,S0,S1	53.62	96.71 ± 0.82
LN,ST,S0,S1,S2	52.88	96.71 ± 0.85
S0	51.76	96.56 ± 0.88
S1	48.05	96.78 ± 0.85
S2	34.69	96.66 ± 0.90
LN,S0	54.92	96.74 ± 0.85
LN,S1	44.53	96.75 ± 0.85
LN,S2	40.45	96.65 ± 0.89
S1,S2	44.16	96.65 ± 0.87
S0,S1,S2	49.17	96.68 ± 0.87
ST,S0,S1,S2	44.90	96.69 ± 0.86
LN,S0,S1,S2	53.99	96.69 ± 0.90

Effect of Varying λ : The hyperparameter λ controls the balance between the supervision signal from the pretrained model and the modality alignment objective, both essential components for effective heterogeneous face recognition (HFR) task. Table 3 summarizes the results of this ablation experiment. Setting $\lambda = 0$ emphasizes only the modality alignment term, removing all guidance from the pretrained network. Although this encourages alignment to the new modality, it quickly leads to overfitting due to the limited size of the training set. In contrast, using $\lambda = 1$ depends entirely on pretrained supervision, resulting in inadequate cross-modal alignment and poor HFR performance. The values of $\lambda = 0.50$ and $\lambda = 0.75$ both provide strong results; however, $\lambda = 0.75$ achieves the most favorable trade-off by placing slightly more weight on the pretrained guidance, which is important given the small fine-tuning dataset. This choice not only improves HFR accuracy but also reduces catastrophic forgetting, preserving performance on the original RGB domain (Table 6).

Effect of Training Set Size: In Table 4, we investigate how the number of subjects influences performance. In particular, our training protocol already operates in a highly data-efficient regime: using 100% of the available training data corresponds to only 45 subjects with paired RGB–Thermal samples: orders of magnitude

TABLE 3
Ablation study with varying values of hyperparameter λ .

λ	AUC	EER	Rank-1	VR@FAR=1%
0.00	94.52	12.23	52.96	46.57
0.25	97.87	7.24	79.53	81.82
0.50	98.46	5.95	83.3	84.79
0.75	97.91	6.68	82.59	86.83
1.00	87.47	20.22	42.19	42.86




Fig. 2. Performance evolution using different fractions of the training data

fewer identities than typical RGB face recognition models, which are often trained on more than 100K identities. Despite this limited supervision, our method achieves strong results, showing that the method can still learn useful representations from limited paired data.

TABLE 4

Experimental results using different fractions of the training set on the Tufts Face Dataset.

% of Training Data	Number of Subjects	AUC	EER	Rank-1	VR@FAR=1%
100%	45	97.91	6.68	82.59	86.83
50%	22	97.32	7.65	79.17	81.82
20%	9	95.71	11.13	70.38	74.40
10%	4	92.70	15.62	58.89	61.04
5%	2	90.15	18.37	48.11	46.94

To further evaluate the model, we progressively reduce the training set and show the performance trend in Fig. 2. As expected, performance decreases as less training data is used; however, the method remains relatively stable at moderate fractions and only shows a sharper decline below 20% of the data (less than 10 subjects in the training set). This indicates that our approach is resilient under constrained data availability while also suggesting clear room for further improvement with more paired data.

Experiments with Other EdgeFace Variants: Although our primary experiments use the EdgeFace-Base model, we additionally evaluated the proposed adaptation strategy on smaller variants of the architecture. Table 5 compares the performance of the original pretrained models with their adapted versions, denoted as *xEdgeFace*. The results show that absolute heterogeneous face recognition (HFR) performance scales with the quality of the pretrained weights. Nonetheless, our adaptation method delivers substantial improvements across all model sizes, achieving relative

gains of 103%, 194%, 257%, and 361% from the largest to the most compact variant. These findings demonstrate both the scalability and robustness of the proposed approach, showcasing its effectiveness even when applied to extremely lightweight model variants.

TABLE 5
Comparison with different variants of EdgeFace

Model	AUC	EER	Rank-1	VR@FAR=1%
EdgeFace - Base	87.44	20.37	42.73	42.86
xEdgeFace - Base	97.91	6.68	82.59	86.83 (↑ 103%)
EdgeFace - S ($\gamma = 0.5$)	80.93	27.30	24.78	25.05
xEdgeFace - S ($\gamma = 0.5$)	96.93	8.89	71.10	73.65 (↑ 194%)
EdgeFace - XS ($\gamma = 0.6$)	77.76	30.24	18.67	19.11
xEdgeFace - XS ($\gamma = 0.6$)	96.28	10.02	68.22	68.27 (↑ 257%)
EdgeFace - XXS	75.72	31.73	17.41	12.80
xEdgeFace - XXS	95.14	11.35	60.50	59.00 (↑ 361%)

Face Recognition Performance of Adapted HFR Models:

To assess the face recognition (FR) capability of the adapted models beyond the heterogeneous setting, we evaluate the **xEdgeFace** variants on standard FR benchmarks. Specifically, we report accuracies on LFW [71], CA-LFW [72], CP-LFW [73], CFP-FP [74], and AgeDB-30 [75]. We compare xEdgeFace models adapted under both the VIS-NIR and VIS-Thermal HFR settings, with the latter representing the most challenging domain shift in our experiments. As shown in Table 6, the adapted model maintains FR performance that is nearly identical to the original EdgeFace backbone, even after adaptation to the severe VIS-Thermal scenario. At the same time, xEdgeFace achieves substantial improvements in HFR accuracy, demonstrating strong performance in both homogeneous and cross-modal recognition tasks. This behavior highlights the effectiveness of the self-distillation component, which acts as a regularizer that mitigates catastrophic forgetting and preserves the discriminative power of the pretrained RGB model. Overall, the proposed training strategy successfully extends the model’s capabilities to heterogeneous face recognition without degrading its original FR performance, thereby showcasing reliable performance in both homogeneous and heterogeneous settings.

TABLE 6
Face recognition performance of the pretrained and adapted model.

Model	LFW [71]	CALFW [72]	CPLFW [73]	CFP-FP [74]	AGEDB-30 [75]
EdgeFace - Base	99.83 ± 0.24	96.07 ± 1.03	93.75 ± 1.16	97.01 ± 0.94	97.60 ± 0.70
xEdgeFace - Base (VIS-Thermal)	99.78 ± 0.27	95.85 ± 1.16	93.62 ± 1.31	96.83 ± 0.99	97.50 ± 0.88
xEdgeFace - Base (VIS-NIR)	99.82 ± 0.26	96.07 ± 0.99	93.78 ± 1.24	96.94 ± 0.97	97.28 ± 0.83

Visualizations In this section, we compare the score distributions of the models before and after the adaptation process. Figure 3 shows the score distributions on the Tufts dataset under the VIS-Thermal protocol, illustrating the baseline performance where the genuine and impostor VIS-Thermal pairs are plotted. Prior to adaptation, the genuine score distribution heavily overlaps with the impostor distribution. After adaptation, however, the genuine distribution shifts to the right, indicating improved separability. This clearly demonstrates the improvement achieved by our proposed pipeline in the challenging VIS-Thermal matching scenario.

We further examine the t-SNE distribution of the embeddings before and after adaptation. The t-SNE plots in Fig. 4 illustrate how the embedding space evolves after the training process. After adaptation, the identities form tighter and more coherent clusters and align more closely with the source-domain embeddings.

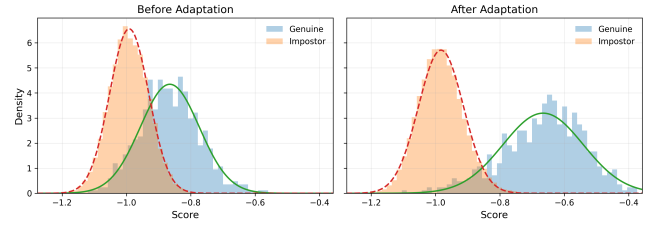


Fig. 3. Genuine and Impostor Score Distribution Before and After Adaptation

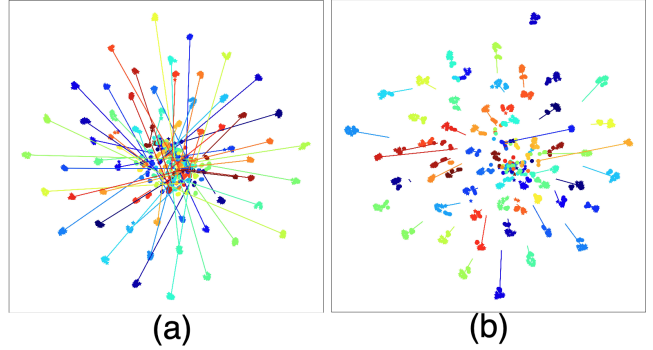


Fig. 4. t-SNE plots of visible and thermal images at different stages of the pipeline, where each color represents a distinct identity. Lines connect the cluster centers of the visible and thermal images for each identity. (a) shows the embedding space before adaptation, while (b) shows the final embedding space. As observed, the identity clusters align much more closely in the final embedding space.

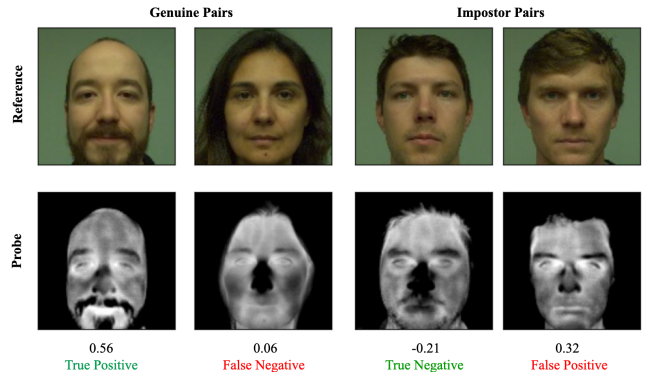


Fig. 5. Illustration of successful and failure cases in the MCXFace VIS-Thermal face-matching scenario. The reported scores correspond to the cosine similarity computed between each reference-probe pair.

We also present success and failure cases from the VIS-Thermal protocol, one of the most extreme and challenging recognition scenarios in the MCXFace dataset (Fig. 5). These examples are particularly useful for understanding the failure cases. As shown, the cold temperature in the nose region distorts key discriminative cues, leading to false negatives. In the false-positive examples, although the match scores remain relatively low, the similar facial structure, combined with the absence of rich textural information can still result in incorrect high match scores.

4.4 Comparison with State-of-the-art

In this section, we provide a comparative evaluation of the proposed *xEdgeFace* model against state-of-the-art heterogeneous

face recognition (HFR) approaches reported in the literature. However, it is important to remember that *xEdgeFace* is substantially more lightweight than the competing models, highlighting the efficiency and practicality of our method. For all experiments, we use the *xEdgeFace*-Base variant with two adaptation configurations: (LN, ST) and $(LN, ST, S0)$. The adaptation loss weight λ is fixed at 0.75 across all evaluations to maintain consistency, although this value can be further tuned for specific datasets depending on their size and the desired balance between adaptation strength and face recognition performance.

Experiments with the Tufts Face Dataset: Table 7 reports the performance of *xEdgeFace* alongside state-of-the-art methods on the VIS-Thermal protocol of the Tufts Face Dataset. This dataset is particularly challenging due to substantial pose variations, especially extreme yaw angles that negatively impact both visible-spectrum and heterogeneous face recognition systems. Despite these challenges, the *xEdgeFace* model using the $(LN, ST, S0)$ configuration achieves the highest verification rate (69.02%) and Rank-1 accuracy (82.59%), while remaining extremely lightweight compared to competing approaches.

TABLE 7

Experimental results on VIS-Thermal protocol of the Tufts Face dataset.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
LightCNN [76]	29.4	23.0	5.3
DVG [77]	56.1	44.3	17.1
DVG-Face [29]	75.7	68.5	36.5
DSU-Iresnet100 [32]	49.7	49.8	28.3
PDT [32]	65.71	69.39	45.45
CAIM [23]	73.07	76.81	46.94
SSMB [24] (N=1)	75.04	78.29	53.99
SSMB [24] (N=2)	78.46	80.33	54.55
<i>xEdgeFace</i>-Base (LN, ST)	75.76	79.59	62.89
<i>xEdgeFace</i>-Base (LN, ST, S0)	82.59	86.83	69.02

Experiments with the MCXFace Dataset: Table 8 presents the average performance over five folds for the VIS-Thermal protocol on the MCXFace dataset, with results reported as mean and standard deviation across all folds. The baseline shows the performance of the pretrained *IresNet100* face recognition model when evaluated directly on thermal images. As shown in the table, the proposed *xEdgeFace* method surpasses all competing approaches, achieving the highest average Rank-1 accuracy of 91.68%.

TABLE 8

Performance of the proposed approach in the VIS-Thermal protocol of MCXFace dataset, the Baseline is a pretrained *Iresnet100* model.

Method	AUC	EER	Rank-1
Baseline	84.45 ± 3.70	22.07 ± 2.81	47.23 ± 3.93
DSU-Iresnet100 [32]	98.12 ± 0.75	6.58 ± 1.35	83.43 ± 5.47
PDT [32]	98.43 ± 0.78	6.52 ± 1.45	84.52 ± 5.36
CAIM [23]	98.97 ± 0.24	5.05 ± 0.91	87.24 ± 2.75
<i>xEdgeFace</i>-Base (LN, ST)	99.12 ± 0.14	4.71 ± 0.31	89.98 ± 2.47
<i>xEdgeFace</i>-Base (LN, ST, S0)	99.50 ± 0.21	3.42 ± 0.78	91.68 ± 2.67

Modality agnostic performance on MCXFace dataset: We further evaluated the baseline methods and the proposed approach on the MCXFace dataset under the “VIS-UNIVERSAL” protocols (using protocols from [24]), where enrollment samples come from the VIS domain and probe samples may belong to any of the other modalities (Thermal, Near-Infrared, or Shortwave Infrared). Each protocol includes a training set that contains identities in both

the source and target modalities, as well as a development (dev) set in which the source images are used for enrollment and the target images are used for probing. Model training and selection are performed exclusively on the training set, while the dev set is used only for performance comparison and not for additional tuning. For evaluation, experiments are conducted across all five protocol splits, and we report the mean and standard deviation of the results (we follow the same evaluation scheme reported in [24]).

Table 9 presents the aggregated results across all probe modalities. The proposed approach achieves performance comparable to existing methods while operating with a substantially smaller computational footprint. Across the three probe modalities: Near-Infrared, Thermal, and Shortwave Infrared, our method attains an average verification rate of 94.11%. These results demonstrate that our approach enables modality-agnostic heterogeneous matching with a single unified model, greatly enhancing flexibility by supporting both homogeneous and heterogeneous recognition, as well as potential cross-modal matching across different spectral domains.

Experiments with the Polathermal Dataset: We evaluate our method on the thermal-to-visible face recognition protocols of the Polathermal dataset, with results summarized in Table 10. The reported values correspond to the average Rank-1 identification accuracy over the five protocols defined in [18]. As shown, the proposed *xEdgeFace* model achieves the highest performance, reaching an average Rank-1 accuracy of 97.31%.

Experiments with the SCFace Dataset: The SCFace dataset introduces a challenging form of heterogeneity due to the quality differences between the gallery images (high-resolution mugshots) and the probe images (low-resolution surveillance camera images). Table 11 reports the results for the “far” protocol, which is the most difficult among the dataset’s evaluation protocols. As shown, our method achieves the best Rank-1 accuracy of 96.36%. This demonstrates that our framework can effectively handle not only cross-modal variation but also substantial image quality degradation.

Experiments with the CUFSF Dataset: We next evaluate our method on the challenging sketch-to-photo face recognition task. Table 12 reports the Rank-1 accuracies of other approaches under the evaluation protocol defined in [68]. Our method achieves a Rank-1 accuracy of 80.30%, ranking second only to SSMB [24]. Despite this strong relative performance, the absolute accuracy remains lower than in other modalities such as thermal or near-infrared, showing the difficulty of this VIS-Sketch matching task. The CUFSF dataset contains viewed hand-drawn sketches [80], which, although recognizable to humans, often omit or exaggerate key discriminative facial cues used by recognition models. Unlike other sensing modalities, sketches are heavily influenced by artistic style and interpretation, introducing a substantial domain gap. Nevertheless, our approach demonstrates competitive performance under these extreme cross-modal conditions, highlighting its robustness and adaptability.

Experiments with the CASIA NIR-VIS 2.0 Dataset: We further evaluate the proposed method on the CASIA NIR-VIS 2.0 dataset to examine its performance on the VIS-NIR matching task. Due to the relatively small domain gap, VIS-pretrained models already achieve strong baseline performance. To ensure more rigorous comparison, we adopt stricter evaluation criteria, reporting VR@FAR=0.1% and VR@FAR=0.01%. Following the standard protocol, results are presented as the average and standard

TABLE 9
Experimental results on VIS-UNIVERSAL protocol of the MCXFace dataset– aggregated performance.

Modality	AUC	EER	Rank-1	VR@FAR=0.1%	VR@FAR=1%
DSU [32]	95.57±0.80	10.24±0.88	84.21±0.94	67.89±1.05	78.13±1.20
PDT [32]	96.16±1.60	9.60±2.07	80.90±2.49	64.63±5.87	76.30±2.49
CAIM [23]	99.45±0.12	3.67±0.33	90.92±1.30	79.64±2.46	91.58±0.68
SSMB [24]	99.70±0.08	2.59±0.28	92.80±0.71	84.04±1.71	94.50±1.44
xEdgeFace-Base (LN, ST, S0)	99.72±0.07	2.76±0.42	94.42±1.42	85.59±2.10	94.11±1.24

TABLE 10
Pola Thermal - Average Rank-1 recognition rate.

Method	Mean (Std. Dev.)
DPM in [64]	75.31 % (-)
CpNN in [64]	78.72 % (-)
PLS in [64]	53.05% (-)
LBP _s + DoG in [9]	36.8% (3.5)
ISV in [78]	23.5% (1.1)
GFK in [79]	34.1% (2.9)
DSU(Best Result) [18]	76.3% (2.1)
DSU-Iresnet100 [32]	88.2% (5.8)
PDT [32]	97.1% (1.3)
CAIM [23]	95.00% (1.63)
xEdgeFace-Base (LN, ST)	95.98% (1.90)
xEdgeFace-Base (LN, ST, S0)	97.31% (1.96)

TABLE 11
Performance of the proposed approach in the SCFace dataset, performance reported on the *far* protocol.

Method	AUC	EER	Rank-1
DSU-Iresnet100 [32]	97.18	8.37	79.53
PDT [32]	98.31	6.98	84.19
CAIM [23]	98.81	5.09	86.05
SSMB [24] (N=1)	98.77	5.91	87.73
SSMB [24] (N=2)	98.67	6.36	86.82
xEdgeFace-Base (LN, ST)	99.86	1.82	96.36
xEdgeFace-Base (LN, ST, S0)	99.71	2.27	93.18

deviation over 10 folds. As shown in Table 13, our approach consistently surpasses existing state-of-the-art methods, demonstrating excellent performance in the VIS-NIR setting.

4.5 Discussions

The comprehensive experimental results across six heterogeneous face recognition (HFR) benchmarks highlight the effectiveness, generalizability, and efficiency of the proposed *xEdgeFace* framework. Despite its lightweight design, *xEdgeFace* consistently outperforms or closely matches state-of-the-art methods across a wide range of modalities, including thermal, NIR, sketch, and low-resolution surveillance imagery. Notably, our approach achieves top Rank-1 accuracy on several challenging datasets with minimal degradation in performance on standard face recognition benchmarks. The ablation studies provide valuable insights into the most impactful components of the adaptation strategy and demonstrate the crucial role of self-distillation in balancing cross-modal alignment while preventing catastrophic forgetting. The framework also scales effectively to highly compact model variants, achieving large relative improvements showcasing its suitability for edge and resource-constrained scenarios. Moreover, strong performance under large domain shifts (e.g., visible to thermal) further validates

TABLE 12
CUFSF: Rank-1 recognition rate in sketch to photo recognition.

Method	Rank-1
IACycleGAN [68]	64.94
DSU-Iresnet100 [32]	67.06
PDT [32]	71.08
CAIM [23]	76.38
SSMB [24] (N=1)	81.14
SSMB [24] (N=2)	81.67
xEdgeFace-Base (LN, ST)	80.30
xEdgeFace-Base (LN, ST, S0)	78.81

TABLE 13
Experimental results on CASIA NIR-VIS 2.0.

Method	Rank-1	VR@FAR=0.1%	VR@FAR=0.01%
IDNet [81]	87.1±0.9	74.5	-
HFR-CNN [82]	85.9±0.9	78.0	-
Hallucination [83]	89.6±0.9	-	-
TRIVET [84]	95.7±0.5	91.0±1.3	74.5±0.7
W-CNN [85]	98.7±0.3	98.4±0.4	94.3±0.4
PACH [86]	98.9±0.2	98.3±0.2	-
RCN [87]	99.3±0.2	98.7±0.2	-
MC-CNN [88]	99.4±0.1	99.3±0.1	-
DVR [89]	99.7±0.1	99.6±0.3	98.6±0.3
DVG [77]	99.8±0.1	99.8±0.1	98.8±0.2
DVG-Face [29]	99.9±0.1	99.9±0.0	99.2±0.1
PDT [32]	99.95±0.04	99.94±0.03	99.77±0.09
CAIM [23]	99.96±0.02	99.95±0.02	99.79±0.11
xEdgeFace-Base (LN, ST)	99.96±0.02	99.91±0.02	99.83±0.04
xEdgeFace-Base (LN, ST, S0)	99.99±0.01	99.93±0.02	99.86±0.04

the robustness and adaptability of the proposed method across challenging heterogeneous settings. These results suggest that heterogeneous face recognition can be effectively addressed through parameter-efficient adaptation of pretrained RGB models, without requiring heavy architectures or modality-specific branches.

4.6 Limitations

The proposed framework is most effective for modality gaps dominated by spectral or photometric differences, such as VIS–NIR and VIS–Thermal, where adapting LayerNorm parameters and shallow layers can reduce domain gap while preserving RGB recognition performance. Its effectiveness is more limited for modalities with stronger structural discrepancies, such as sketch-to-photo matching, since these differences cannot be fully addressed through statistical alignment alone. This is reflected in the lower performance on CUFSF compared to the thermal and NIR benchmarks. More generally, our method should be viewed as a lightweight and parameter-efficient adaptation strategy that is particularly well suited to cross-spectral face recognition, rather than a universal solution for all heterogeneous settings. For modalities with larger structural gaps, more expressive adaptation mechanisms may be required.

5 CONCLUSIONS

In this work, we presented *xEdgeFace*, an efficient framework for lightweight heterogeneous face recognition (HFR) that extends existing face recognition models to cross-modal scenarios. By selectively adapting early convolutional layers and LayerNorm (LN) modules within a contrastive self-distillation framework, our approach achieves strong cross-modal generalization while preserving the model's original performance in the visible spectrum. This design enables the adapted model to perform robustly across challenging HFR tasks, including VIS-Thermal, VIS-NIR, and sketch-to-photo recognition while maintaining competitive accuracy on standard FR benchmarks, effectively mitigating catastrophic forgetting. Extensive experiments demonstrate that *xEdgeFace* consistently outperforms or matches state-of-the-art methods, even when applied to highly compact model variants, making it particularly suitable for edge deployment. The proposed model achieved performance comparable to, or better than, state-of-the-art models while using roughly one-twentieth of the compute. Furthermore, the results show that our adaptation strategy supports modality-agnostic heterogeneous matching within a single unified architecture, greatly enhancing flexibility by enabling both homogeneous and heterogeneous recognition, as well as potential cross-modal matching across different spectral domains. The source code and pretrained models are made publicly available to support reproducibility and facilitate further extensions of this work.

ACKNOWLEDGMENTS

This research was funded by the European Union project CarMen (Grant Agreement No. 101168325).

REFERENCES

- [1] E. Learned-Miller, G. B. Huang, A. RoyChowdhury, H. Li, and G. Hua, "Labeled faces in the wild: A survey," *Advances in face detection and facial image analysis*, vol. 1, pp. 189–248, 2016.
- [2] S. Z. Li, R. Chu, S. Liao, and L. Zhang, "Illumination invariant face recognition using near-infrared images," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 29, no. 4, pp. 627–639, 2007.
- [3] A. George, D. Geissbuhler, and S. Marcel, "A comprehensive evaluation on multi-channel biometric face presentation attack detection," *arXiv preprint arXiv:2202.10286*, 2022.
- [4] B. F. Klare and A. K. Jain, "Heterogeneous face recognition using kernel prototype similarities," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 6, pp. 1410–1422, 2012.
- [5] D. Anghelone, C. Chen, A. Ross, and A. Dantcheva, "Beyond the visible: A survey on cross-spectral face recognition," *Neurocomputing*, vol. 611, p. 128626, 2025.
- [6] R. He, X. Wu, Z. Sun, and T. Tan, "Wasserstein CNN: Learning invariant features for Nir-Vis face recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 7, pp. 1761–1773, 2018.
- [7] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah, "Transformers in vision: A survey," *ACM computing surveys (CSUR)*, vol. 54, no. 10s, pp. 1–41, 2022.
- [8] A. George, C. Ecabert, H. O. Shahreza, K. Kotwal, and S. Marcel, "Edgeface: Efficient face recognition model for edge devices," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 6, no. 2, pp. 158–168, 2024.
- [9] S. Liao, D. Yi, Z. Lei, R. Qin, and S. Z. Li, "Heterogeneous face recognition from local structures of normalized appearance," in *International Conference on Biometrics*. Springer, 2009, pp. 209–218.
- [10] B. Klare, Z. Li, and A. K. Jain, "Matching forensic sketches to mug shot photos," *IEEE transactions on pattern analysis and machine intelligence*, vol. 33, no. 3, pp. 639–646, 2010.
- [11] R. He, X. Wu, Z. Sun, and T. Tan, "Learning invariant deep representation for Nir-Vis face recognition," in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [12] H. Roy and D. Bhattacharjee, "A novel quaternary pattern of local minimum quotient for heterogeneous face recognition," *Pattern Recognition Letters*, vol. 113, pp. 19–28, 2018.
- [13] D. Liu, J. Li, N. Wang, C. Peng, and X. Gao, "Composite components-based face sketch recognition," *Neurocomputing*, vol. 302, pp. 46–54, 2018.
- [14] J. V. de Andrade, A. Freire, G. C. Pereira, C. Millan-Arias, B. Fernandes, C. Bastos-Filho, J. Tortato, L. Da Rocha, and A. M. Maciel, "Towards robust facial recognition: Gabor filter-based feature extraction for nir-vis heterogeneous face recognition," in *Proceedings of the 40th ACM/SI-GAPP Symposium on Applied Computing*, 2025, pp. 1275–1281.
- [15] D. Yi, R. Liu, R. Chu, Z. Lei, and S. Z. Li, "Face matching between near infrared and visible light images," in *International Conference on Biometrics*. Springer, 2007, pp. 523–530.
- [16] A. Sharma and D. W. Jacobs, "Bypassing synthesis: PLS for face recognition with pose, low-resolution and sketch," in *CVPR 2011*. IEEE, 2011, pp. 593–600.
- [17] Z. Lei and S. Z. Li, "Coupled spectral regression for matching heterogeneous faces," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2009, pp. 1123–1128.
- [18] T. de Freitas Pereira, A. Anjos, and S. Marcel, "Heterogeneous face recognition using domain specific units," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 7, pp. 1803–1816, 2018.
- [19] A. George and S. Marcel, "Heterogeneous face recognition using domain invariant units," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 4780–4784.
- [20] D. Liu, X. Gao, N. Wang, J. Li, and C. Peng, "Coupled attribute learning for heterogeneous face recognition," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 11, pp. 4699–4712, 2020.
- [21] D. Liu, W. Yang, C. Peng, N. Wang, R. Hu, and X. Gao, "Modality-agnostic augmented multi-collaboration representation for semi-supervised heterogeneous face recognition," in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 4647–4656.
- [22] A. George and S. Marcel, "From modalities to styles: Rethinking the domain gap in heterogeneous face recognition," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 6, no. 4, pp. 475–485, 2024.
- [23] —, "Bridging the gap: Heterogeneous face recognition with conditional adaptive instance modulation," in *2023 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2023, pp. 1–10.
- [24] —, "Modality agnostic heterogeneous face recognition with switch style modulators," in *2024 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2024, pp. 1–10.
- [25] X. Wang and X. Tang, "Face photo-sketch synthesis and recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 31, no. 11, pp. 1955–1967, 2008.
- [26] Q. Liu, X. Tang, H. Jin, H. Lu, and S. Ma, "A nonlinear approach for face sketch synthesis and recognition," in *2005 IEEE Computer Society conference on computer vision and pattern recognition (CVPR'05)*, vol. 1. IEEE, 2005, pp. 1005–1010.
- [27] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks," *arXiv:1703.10593 [cs]*, Mar. 2017.
- [28] H. Zhang, V. M. Patel, B. S. Riggan, and S. Hu, "Generative adversarial network-based synthesis of visible faces from polarimetric thermal faces," in *2017 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2017, pp. 100–107.
- [29] C. Fu, X. Wu, Y. Hu, H. Huang, and R. He, "DVG-face: Dual variational generation for heterogeneous face recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [30] D. Liu, X. Gao, C. Peng, N. Wang, and J. Li, "Heterogeneous face interpretable disentangled representation for joint face recognition and synthesis," *IEEE transactions on neural networks and learning systems*, vol. 33, no. 10, pp. 5611–5625, 2021.
- [31] M. Luo, H. Wu, H. Huang, W. He, and R. He, "Memory-modulated transformer network for heterogeneous face recognition," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 2095–2109, 2022.
- [32] A. George, A. Mohammadi, and S. Marcel, "Prepended domain transformer: Heterogeneous face recognition without bells and whistles," *IEEE Transactions on Information Forensics and Security*, 2022.
- [33] Y. Yang, W. Hu, H. Lin, and H. Hu, "Robust cross-domain pseudo-labeling and contrastive learning for unsupervised domain adaptation nir-vis face recognition," *IEEE Transactions on Image Processing*, vol. 32, pp. 5231–5244, 2023.

- [34] W. Hu, Y. Yang, and H. Hu, "Pseudo label association and prototype-based invariant learning for semi-supervised nir-vis face recognition," *IEEE Transactions on Image Processing*, vol. 33, pp. 1448–1463, 2024.
- [35] Y. Yang, W. Hu, and H. Hu, "Unsupervised nir-vis face recognition via homogeneous-to-heterogeneous learning and residual-invariant enhancement," *IEEE Transactions on Information Forensics and Security*, vol. 19, pp. 2112–2126, 2023.
- [36] S. Chen, Y. Liu, X. Gao, and Z. Han, "Mobilefacenets: Efficient cnns for accurate real-time face verification on mobile devices," in *Biometric Recognition: 13th Chinese Conference, CCBR 2018, Urumqi, China, August 11-12, 2018, Proceedings 13*. Springer, 2018, pp. 428–438.
- [37] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.
- [38] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4510–4520.
- [39] F. Boutros, N. Damer, M. Fang, F. Kirchbuchner, and A. Kuijper, "Mixfacenets: Extremely efficient face recognition networks," in *2021 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2021, pp. 1–8.
- [40] M. Tan and Q. V. Le, "Mixconv: Mixed depthwise convolutional kernels," *arXiv preprint arXiv:1907.09595*, 2019.
- [41] B. Wu, A. Wan, X. Yue, P. Jin, S. Zhao, N. Golmant, A. Gholaminejad, J. Gonzalez, and K. Keutzer, "Shift: A zero flop, zero parameter alternative to spatial convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9127–9135.
- [42] Y. Martindiez-Diaz, L. S. Luevano, H. Mendez-Vazquez, M. Nicolas-Diaz, L. Chang, and M. Gonzalez-Mendoza, "Shufflefacenet: A lightweight face architecture for efficient and highly-accurate face recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019, pp. 0–0.
- [43] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, "Shuffletnet v2: Practical guidelines for efficient cnn architecture design," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 116–131.
- [44] F. Boutros, P. Siebke, M. Klemm, N. Damer, F. Kirchbuchner, and A. Kuijper, "Pocketnet: Extreme lightweight face recognition network using neural architecture search and multistep knowledge distillation," *IEEE Access*, vol. 10, pp. 46 823–46 833, 2022.
- [45] D. Yi, Z. Lei, S. Liao, and S. Z. Li, "Learning face representation from scratch," *arXiv preprint arXiv:1411.7923*, 2014.
- [46] M. Yan, M. Zhao, Z. Xu, Q. Zhang, G. Wang, and Z. Su, "Vargfacenet: An efficient variable group convolutional neural network for lightweight face recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019, pp. 0–0.
- [47] J. Deng, J. Guo, D. Zhang, Y. Deng, X. Lu, and S. Shi, "Lightweight face recognition challenge," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019, pp. 0–0.
- [48] H. O. Shahreza, A. George, and S. Marcel, "Knowledge distillation for face recognition using synthetic data with dynamic latent sampling," *IEEE Access*, 2024.
- [49] A. George and S. Marcel, "Digi2real: Bridging the realism gap in synthetic data face recognition via foundation models," in *Proceedings of the Winter Conference on Applications of Computer Vision*, 2025, pp. 1469–1478.
- [50] K. Han, Y. Wang, Q. Zhang, W. Zhang, C. Xu, and T. Zhang, "Model rubik's cube: Twisting resolution, depth and width for tinynets," *Advances in Neural Information Processing Systems*, vol. 33, pp. 19 353–19 364, 2020.
- [51] M. Alansari, O. A. Hay, S. Javed, A. Shoufan, Y. Zweiri, and N. Werghi, "Ghostfacenets: Lightweight face recognition model from cheap operations," *IEEE Access*, 2023.
- [52] M. Maaz, A. Shaker, H. Cholakkal, S. Khan, S. W. Zamir, R. M. Anwer, and F. Shahbaz Khan, "Edgenext: efficiently amalgamated cnn-transformer architecture for mobile vision applications," in *Computer Vision—ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*. Springer, 2023, pp. 3–20.
- [53] J. N. Kolf, F. Boutros, J. Elliesen, M. Theuerkauf, N. Damer, M. Alansari, O. A. Hay, S. Alansari, S. Javed, N. Werghi *et al.*, "Efar 2023: Efficient face recognition competition," in *2023 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2023, pp. 1–12.
- [54] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [55] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2414–2423.
- [56] B. Zhao, H. Tu, C. Wei, J. Mei, and C. Xie, "Tuning layernorm in attention: Towards efficient multi-modal LLM finetuning," in *The Twelfth International Conference on Learning Representations*, 2024. [Online]. Available: <https://openreview.net/forum?id=YR3ETaEINK>
- [57] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen *et al.*, "Lora: Low-rank adaptation of large language models." *ICLR*, vol. 1, no. 2, p. 3, 2022.
- [58] J. Xu, X. Sun, Z. Zhang, G. Zhao, and J. Lin, "Understanding and improving layer normalization," *Advances in neural information processing systems*, vol. 32, 2019.
- [59] Z. Zhu, G. Huang, J. Deng, Y. Ye, J. Huang, X. Chen, J. Zhu, T. Yang, J. Lu, D. Du *et al.*, "Webface260m: A benchmark unveiling the power of million-scale deep face recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10 492–10 502.
- [60] A. Anjos, L. E. Shafey, R. Wallace, M. Günther, C. McCool, and S. Marcel, "Bob: a free signal processing and machine learning toolbox for researchers," in *20th ACM Conference on Multimedia Systems (ACMMM)*, Nara, Japan, Oct. 2012. [Online]. Available: https://publications.idiap.ch/downloads/papers/2012/Anjos_Bob_ACMMM12.pdf
- [61] A. Anjos, M. Günther, T. de Freitas Pereira, P. Korshunov, A. Mohammadi, and S. Marcel, "Continuously reproducing toolchains in pattern recognition and machine learning experiments," in *International Conference on Machine Learning (ICML)*, Aug. 2017. [Online]. Available: http://publications.idiap.ch/downloads/papers/2017/Anjos_ICML2017-2_2017.pdf
- [62] K. Panetta, Q. Wan, S. Agaian, S. Rajeev, S. Kamath, R. Rajendran, S. P. Rao, A. Kaszowska, H. A. Taylor, A. Samani *et al.*, "A comprehensive database for benchmarking imaging systems," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 3, pp. 509–520, 2018.
- [63] Z. Mostaani, A. George, G. Heusch, D. Geissbuhler, and S. Marcel, "The high-quality wide multi-channel attack (hq-wmca) database," *arXiv preprint arXiv:2009.09703*, 2020.
- [64] S. Hu, N. J. Short, B. S. Riggan, C. Gordon, K. P. Gurton, M. Thielke, P. Gurrum, and A. L. Chan, "A polarimetric thermal database for face recognition research," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2016, pp. 119–126.
- [65] M. Grgic, K. Delac, and S. Grgic, "SCface—surveillance cameras face database," *Multimedia tools and applications*, vol. 51, no. 3, pp. 863–879, 2011.
- [66] W. Zhang, X. Wang, and X. Tang, "Coupled information-theoretic encoding for face photo-sketch recognition," in *CVPR 2011*. IEEE, 2011, pp. 513–520.
- [67] P. J. Phillips, H. Wechsler, J. Huang, and P. J. Rauss, "The FERET database and evaluation procedure for face-recognition algorithms," *Image and vision computing*, vol. 16, no. 5, pp. 295–306, 1998.
- [68] Y. Fang, W. Deng, J. Du, and J. Hu, "Identity-aware CycleGAN for face photo-sketch synthesis and recognition," *Pattern Recognition*, vol. 102, p. 107249, 2020.
- [69] S. Li, D. Yi, Z. Lei, and S. Liao, "The CASIA Nir-Vis 2.0 face database," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2013, pp. 348–353.
- [70] A. George and S. Marcel, "xEdgeFace: Efficient cross-spectral face recognition for edge devices," pp. 1–10, 2025.
- [71] G. B. Huang, M. Mattar, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," in *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*, 2008.
- [72] T. Zheng, W. Deng, and J. Hu, "Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments," *arXiv preprint arXiv:1708.08197*, 2017.
- [73] T. Zheng and W. Deng, "Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments," *Beijing University of Posts and Telecommunications, Tech. Rep*, vol. 5, no. 7, 2018.
- [74] S. Sengupta, J.-C. Chen, C. Castillo, V. M. Patel, R. Chellappa, and D. W. Jacobs, "Frontal to profile face verification in the wild," in *2016 IEEE winter conference on applications of computer vision (WACV)*. IEEE, 2016, pp. 1–9.
- [75] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotsia, and S. Zafeiriou, "Agedb: the first manually collected, in-the-wild age database," in *proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 51–59.

- [76] X. Wu, R. He, Z. Sun, and T. Tan, "A light CNN for deep face representation with noisy labels," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 11, pp. 2884–2896, 2018.
- [77] C. Fu, X. Wu, Y. Hu, H. Huang, and R. He, "Dual variational generation for low shot heterogeneous face recognition," in *Advances in Neural Information Processing Systems*, 2019.
- [78] T. de Freitas Pereira and S. Marcel, "Heterogeneous face recognition using inter-session variability modelling," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2016, pp. 111–118.
- [79] A. F. Sequeira, L. Chen, J. Ferryman, P. Wild, F. Alonso-Fernandez, J. Bigun, K. B. Raja, R. Raghavendra, C. Busch, T. de Freitas Pereira *et al.*, "Cross-eyed 2017: Cross-spectral iris/periocular recognition competition," in *2017 IEEE International Joint Conference on Biometrics (IJCB)*. IEEE, 2017, pp. 725–732.
- [80] S. J. Klum, H. Han, B. F. Klare, and A. K. Jain, "The facesketchid system: Matching facial composites to mugshots," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 12, pp. 2248–2263, 2014.
- [81] C. Reale, N. M. Nasrabadi, H. Kwon, and R. Chellappa, "Seeing the forest from the trees: A holistic approach to near-infrared heterogeneous face recognition," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2016.
- [82] S. Saxena and J. Verbeek, "Heterogeneous face recognition with CNNs," in *European Conference on Computer Vision*, 2016.
- [83] J. Lezama, Q. Qiu, and G. Sapiro, "Not afraid of the dark: Nir-Vis face recognition via cross-spectral hallucination and low-rank embedding," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [84] X. Liu, L. Song, X. Wu, and T. Tan, "Transferring deep representation for Nir-Vis heterogeneous face recognition," in *International Conference on Biometrics*, 2016.
- [85] R. He, X. Wu, Z. Sun, and T. Tan, "Wasserstein CNN: Learning invariant features for Nir-Vis face recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 7, pp. 1761–1773, 2018.
- [86] B. Duan, C. Fu, Y. Li, X. Song, and R. He, "Pose agnostic cross-spectral hallucination via disentangling independent factors," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [87] Z. Deng, X. Peng, and Y. Qiao, "Residual compensation networks for heterogeneous face recognition," in *AAAI Conference on Artificial Intelligence*, 2019.
- [88] Z. Deng, X. Peng, Z. Li, and Y. Qiao, "Mutual component convolutional neural networks for heterogeneous face recognition," *IEEE Transactions on Image Processing*, vol. 28, no. 6, pp. 3102–3114, 2019.
- [89] X. Wu, H. Huang, V. M. Patel, R. He, and Z. Sun, "Disentangled variational representation for heterogeneous face recognition," in *AAAI Conference on Artificial Intelligence*, 2019.



Sébastien Marcel heads the Biometrics Security and Privacy group at Idiap Research Institute (Switzerland) and conducts research on face recognition, speaker recognition, vein recognition, attack detection (presentation attacks, morphing attacks, deepfakes) and template protection. He received his Ph.D. degree in signal processing from Université de Rennes I in France (2000) at CNET, the research center of France Telecom (now Orange Labs). He is Professor at the University of Lausanne (School of Criminal Justice) and a lecturer at the École Polytechnique Fédérale de Lausanne. He is also the Director of the Swiss Center for Biometrics Research and Testing, which conducts certifications of biometric products.



Anjith George has received his Ph.D. and M-Tech degree from the Department of Electrical Engineering, Indian Institute of Technology (IIT) Kharagpur, India in 2012 and 2018 respectively. After Ph.D, he worked in Samsung Research Institute as a machine learning researcher. Currently, he is a research associate in the biometric security and privacy group at Idiap Research Institute, focusing on developing face recognition and presentation attack detection algorithms.

His research interests are real-time signal and image processing, embedded systems, computer vision, machine learning with a special focus on Biometrics.