

DVG-Face: Dual Variational Generation for Heterogeneous Face Recognition

Chaoyou Fu, Xiang Wu, Yibo Hu, Huaibo Huang, and Ran He, *Senior Member, IEEE*

Abstract—Heterogeneous Face Recognition (HFR) refers to matching cross-domain faces and plays a crucial role in public security. Nevertheless, HFR is confronted with challenges from large domain discrepancy and insufficient heterogeneous data. In this paper, we formulate HFR as a dual generation problem, and tackle it via a novel Dual Variational Generation (DVG-Face) framework. Specifically, a dual variational generator is elaborately designed to learn the joint distribution of paired heterogeneous images. However, the small-scale paired heterogeneous training data may limit the identity diversity of sampling. In order to break through the limitation, we propose to integrate abundant identity information of large-scale visible data into the joint distribution. Furthermore, a pairwise identity preserving loss is imposed on the generated paired heterogeneous images to ensure their identity consistency. As a consequence, massive new diverse paired heterogeneous images with the same identity can be generated from noises. The identity consistency and identity diversity properties allow us to employ these generated images to train the HFR network via a contrastive learning mechanism, yielding both domain-invariant and discriminative embedding features. Concretely, the generated paired heterogeneous images are regarded as positive pairs, and the images obtained from different samplings are considered as negative pairs. Our method achieves superior performances over state-of-the-art methods on seven challenging databases belonging to five HFR tasks, including NIR-VIS, Sketch-Photo, Profile-Frontal Photo, Thermal-VIS, and ID-Camera. The related code will be released at <https://github.com/BradyFU>.

Index Terms—Heterogeneous face recognition, cross-domain, dual generation, contrastive learning.

1 INTRODUCTION

IN the last two decades, face recognition for visible (VIS) images has made a great breakthrough using deep convolutional neural networks [1], [2]. Nevertheless, in some practical security applications, a face recognition network often needs to match cross-domain face images rather than merely VIS, raising the problem of Heterogeneous Face Recognition (HFR). For example, the near-infrared (NIR) imaging sensor has been integrated into many mobile devices, because it offers an inexpensive and effective solution to obtain clear NIR images in extreme lightings or even dark environment. At the same time, the enrollments of face templates are usually VIS images, which requires the recognition network to match the heterogeneous NIR-VIS images. Unfortunately, due to the large domain discrepancy, the performance of the recognition network trained on VIS images often degrades dramatically in such a heterogeneous case [3]. Other HFR tasks include Sketch-Photo [4], Profile-Frontal Photo [5], Thermal-VIS [6], and ID-Camera [7]. In order to bridge the domain discrepancy between the heterogeneous data, extensive research efforts have been undertaken to match cross-domain features [8]. However, due to the difficulty of data acquisition, it is usually not feasible to collect large-scale heterogeneous databases. Over-fitting often occurs when the recognition network is trained on such insufficient data [3].

With the rapid development of deep generative models,

“recognition via generation” has been a hot research topic in the computer vision community [10]. Deep generative models are good at learning the mapping among different domains, and are therefore usually used to reduce the domain discrepancy [11]. For example, [9] proposes to transfer NIR images to VIS ones via a two-path generative model. [12] introduces a multi-stream feature fusion manner with Generative Adversarial Networks (GANs) [13] to synthesize photo-realistic VIS images from polarimetric thermal ones. However, these conditional image-to-image translation based methods face two challenging issues: (1) **Diversity**. Since these methods adopt a “one-to-one” translation manner, the generator can only synthesize one new image of the target domain for each input image of the source domain. Hence, these methods can merely generate a limited number of new data to reduce the domain discrepancy. Meanwhile, as exemplified in the left part of Fig. 1, the generated images have the same attributes (*e.g.* the pose and the expression) as the inputs except for the spectrum, which limits the intra-class diversity of the generated images. (2) **Consistency**. These methods require that the generated images belong to specific classes consistent with the inputs. Nevertheless, it is hard to meet this requirement, because the widely adopted identity preserving loss [10], [14] only constrains the feature distance between the generated images and the targets, while ignores both intra-class and inter-class distances. The above two issues make it rather difficult for deep generative models to effectively boost the performance of HFR.

In order to tackle the above challenges, we propose an unconditional Dual Variational Generation (DVG-Face) framework that can generate large-scale new diverse paired heterogeneous images. Different from the aforementioned conditional generation in HFR, unconditional generation

- C. Fu, X. Wu, Y. Hu, H. Huang, and R. He are with the National Laboratory of Pattern Recognition, CASIA, Center for Research on Intelligent Perception and Computing, CASIA, Center for Excellence in Brain Science and Intelligence Technology, CAS, and School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing 100190, China. E-mail: {chaoyou.fu, rhe}@nlpr.ia.ac.cn, {alfredxiangwu, huyibo871079699}@gmail.com, huaibo.huang@cripac.ia.ac.cn. (Corresponding author: Ran He.)

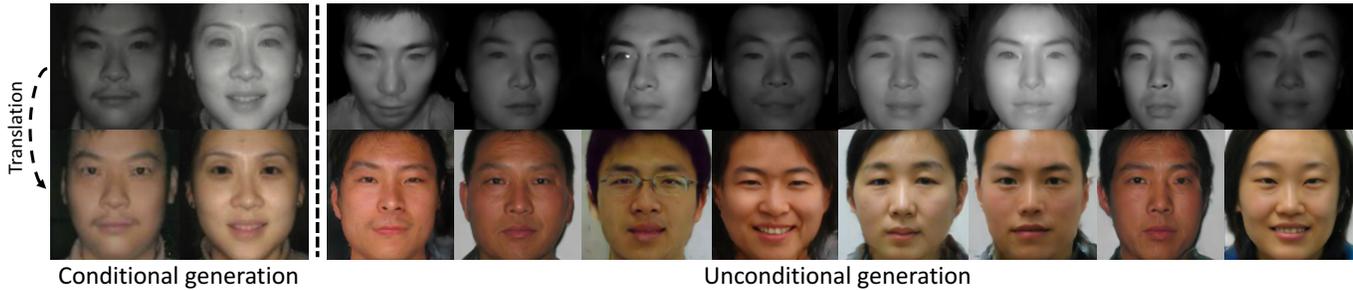


Fig. 1: Comparisons between conditional image-to-image generation [9] (on the left, the top is the input NIR and the bottom is the generated VIS) and our proposed unconditional dual variational generation (on the right, all paired NIR-VIS images are generated from noises). For the conditional generation, given one NIR image, the generator can only synthesize one VIS image with the same attributes (*e.g.* the pose and the expression) except for the spectrum. By contrast, DVG-Face has the ability to generate large-scale new paired heterogeneous images with abundant intra-class diversity from noises.

has the ability to synthesize new data by sampling from noises [15], [16]. However, the small number of the heterogeneous training data may limit the identity diversity of sampling. In order to break through the limitation, we introduce abundant identity information of large-scale VIS images into the generated images. By contrast, these VIS images are merely leveraged to pre-train the HFR network in previous works [17], [18], as discussed in Section 4.3.6. Moreover, DVG-Face focuses on the identity consistency of the generated paired heterogeneous images, rather than whom the generated images belong to. This avoids the consistency problem of the previous conditional generation based methods. Finally, a contrastive learning mechanism is employed to take advantage of these generated unlabeled images to train the HFR network. Specifically, benefiting from the identity consistency property, the generated paired heterogeneous images are regarded as positive pairs. Due to the identity diversity property, the images obtained from different samplings are considered as negative pairs, according to the instance discrimination criterion [19], [20].

The framework of our DVG-Face is depicted in Fig. 2. We first train the dual variational generator with both paired heterogeneous data and large-scale unpaired VIS data. In the latent space, the joint distribution of paired heterogeneous data is disentangled into identity representations and domain-specific attribute distributions, which is implemented by an angular orthogonal loss. Fig. 3 presents the disentanglement results. This disentanglement enables the integration of abundant identity information of the unpaired VIS images into the joint distribution. In the pixel space, we impose a pairwise identity preserving loss on the generated paired heterogeneous images to guarantee their identity consistency. After the training of the dual variational generator, we can obtain large-scale new diverse paired heterogeneous images with the same identity by sampling from noises. These generated images are organized as positive and negative pairs, and then are fed into the HFR network. A contrastive learning mechanism further facilitates the HFR network to learn both domain-invariant and discriminative embedding features.

The main contributions are summarized as follows:

- We provide a new insight into HFR, considering it as a dual generation problem. This results in a novel un-

conditional Dual Variational Generation (DVG-Face) framework, which samples large-scale new diverse paired heterogeneous data from noises to boost the performance of HFR.

- Abundant identity information is integrated into the joint distribution to enrich the identity diversity of the generated data. Meanwhile, a pairwise identity preserving loss is imposed on the generated paired images to ensure their identity consistency. These two properties allow us to make better use of the generated unlabeled data to train the HFR network.
- By regarding the generated paired images as positive pairs and the images obtained from different samplings as negative pairs, the HFR network is optimized via contrastive learning to learn both domain-invariant and discriminative embedding features.
- Extensive experiments on seven HFR databases demonstrate that our method significantly outperforms state-of-the-art methods. In particular, on the challenging low-shot Oulu-CASIA NIR-VIS database, we improve the best $\text{TPR@FPR}=10^{-5}$ by 29.2%. Besides, compared with the baseline trained on the Tufts Face database, $\text{VR@FAR}=1\%$ is increased by 26.1% after adding the generated data.

This paper is an extension of our previous conference version [21], and there are three major improvements over the preliminary one: (1) **The generated images have richer identity diversity.** For the preliminary version, the generator can only be trained with the small-scale paired heterogeneous data, thereby limiting the identity diversity of the generated images. In the current version, the architecture and the training manner of the generator are redesigned, allowing it to be trained with both paired heterogeneous data and large-scale unpaired VIS data. The introduction of the latter greatly enriches the identity diversity of the generated images. (2) **The generated images are leveraged more efficiently.** The preliminary version trains the HFR network with the generated paired data via a pairwise distance loss, resorting to the identity consistency property. On this basis, benefiting from the aforementioned identity diversity property, the current version further regards the images obtained from different samplings as negative pairs, formulating a contrastive learning mechanism. Hence,

the preliminary version can only leverage the generated images to reduce the domain discrepancy, while the current version utilizes the generated images to learn both domain-invariant and discriminative embedding features. (3) **More insightful analyses and more experiments are added.** Apart from NIR-VIS and Sketch-Photo, we further explore Profile-Frontal Photo, Thermal-VIS, and ID-Camera HFR tasks. Moreover, the current version gains significant improvements over the preliminary one on all the databases.

2 RELATED WORK

2.1 Heterogeneous Face Recognition

HFR has attracted increasing attention of researchers for its crucial practical value. In this subsection, we review the development of HFR from the perspectives of feature-level learning and image-level learning.

Feature-level learning based methods aim at seeking discriminative feature representations. Some methods try to capture a common latent space between heterogeneous data for relevance measurement. [22] employs a prototype random subspace to improve the HFR performance. [23] introduces a generalized transfer subspace with a low-rank constraint. There are also some methods that focus on learning domain-invariant features. Traditional works mainly use handcrafted features. For example, [24] extracts Histograms of Oriented Gradients (HOG) features with sparse representations to obtain cross-domain semantics. [25] proposes a common encoding model to obtain discriminant information for HFR. Recently, the thriving deep learning has powerful feature extraction capability, and thus has been widely applied in HFR. [26] explores multiple different metric learning methods to reduce the domain gap. [17] designs a residual compensation framework that learns different modalities in separate branches. However, due to the lack of large-scale heterogeneous databases, the deep learning based methods often tend to over-fit [27]. In order to alleviate the over-fitting problem, our method generates large-scale new diverse paired data as an augmentation, thus significantly boosting the recognition performance.

Image-level learning based methods mainly operate on the pixel space to improve the performance of HFR. [28], [29], [30] focus on Sketch-Photo synthesis to reduce the domain gap. [31] improves the recognition performance by using coupled dictionary learning for cross-domain image synthesis. [32] proposes to composite facial parts, *e.g.* eyes, to augment the small database to a larger one. [33] uses a CNN to perform a cross-spectral hallucination, reducing the domain gap in pixel space. Recently, the rapidly developed GANs [13] provide a solution to the problem of HFR. [34] utilizes image-to-image translation methods to synthesize new data, and then incorporates the synthesized data into the training set to augment the intra-class scale and reduce inter-class diversity. [35] proposes a cascaded face sketch synthesis model to deal with the illumination problem. As discussed in Section 1, the image-to-image translation based methods are limited in diversity and consistency. In order to break through the two limitations, DVG-Face employs a novel unconditional generation manner to generate large-scale new diverse paired images from noises, and a new

pairwise identity preserving loss to ensure the identity consistency of the generated paired heterogeneous images.

2.2 Generative Model

As a hot topic in machine learning and computer vision communities, generative model has made gratifying progress in recent years [13], [15], [37], [38]. Usually, it can be divided into two types: unconditional generative model and conditional generative model.

Unconditional generative model synthesizes data from noises. The most prominent models include Variational AutoEncoders (VAEs) [15] and Generative Adversarial Networks (GANs) [13]. VAEs adopt variational inference to learn data distribution, and consist of an encoder network $q_\phi(z|x)$ and a decoder network $p_\theta(x|z)$. The Evidence Lower Bound Objective (ELBO) is derived for an approximate optimization:

$$\log p_\theta(x) \geq \mathbb{E}_{q_\phi(z|x)} \log p_\theta(x|z) - D_{\text{KL}}(q_\phi(z|x)||p(z)), \quad (1)$$

where the first term is a reconstruction loss and the second term is a Kullback-Leibler divergence. $p(z)$ is usually set to a standard Gaussian distribution. Recently, some variants are proposed to improve synthesis quality. IntroVAE [39] generates high-resolution images via an introspective manner. VQ-VAE [40] obtains high-quality images via learning discrete representations with an autoregressive prior.

In a different way, GANs leverage an adversarial mechanism to learn data distribution implicitly. Generally, GANs consist of two networks: a generator G and a discriminator D . G generates data from a prior $p(z)$ to confuse D , while D is trained to distinguish the generated data from the real data. Such an adversarial mechanism is formulated as:

$$\min_G \max_D \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))]. \quad (2)$$

GANs are good at generating high-quality images. PGGAN [16] significantly improves the resolution of the synthesized images to 1024×1024 . SinGAN [41] generates realistic images merely using one single training image. CoGAN [42] concerns a problem similar to the one tackled by our method. A weight-sharing manner is adopted to generate paired data from two domains. However, due to the lack of explicit identity constraints, it is rather challenging for such a weight-sharing manner of CoGAN to generate paired data with the same identity, as discussed in Section 4.3.1.

Conditional generative model synthesizes data according to the given conditions. pix2pix [11] realizes photo-realistic paired image-to-image translation with a conditional generative adversarial loss. CycleGAN [43] employs a cycle-consistent network to tackle the unpaired image translation problem. GAN Compression [44] reduces the model size and the inference time of Conditional GANs by the ways of Knowledge Distillation (KD) and Neural Architecture Search (NAS). MUNIT [45] decomposes image representations into a content code and a style code, realizing exemplar based image synthesis. BigGAN [46] is the pioneer work for high-resolution natural image synthesis. StarGAN [47] is designed for unpaired multi-domain image-to-image translation. StyleGAN [48] proposes a style-based generator that can separate facial attributes automatically. SPADE

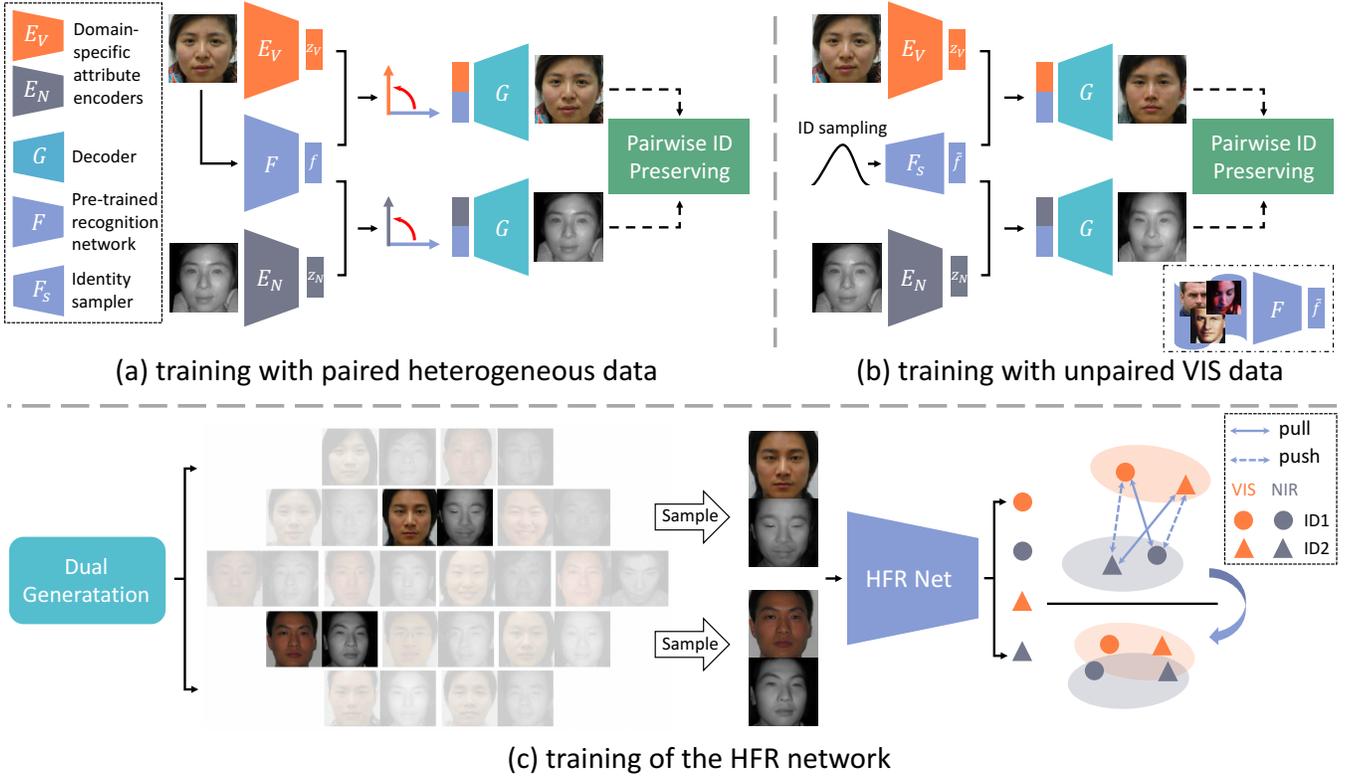


Fig. 2: Illustration of our DVG-Face. (a) and (b): The training of the dual variational generator involves both paired heterogeneous data and unpaired VIS data (MS-Celeb-1M [36]). The former is used to disentangle identity and attribute representations. The latter is introduced to enrich the identity diversity of the generated images. The identity representations of the latter can be obtained from either the identity sampler F_s or the pre-trained face recognition network F , as discussed in Section 4.3.4. In addition, a pairwise identity preserving loss is imposed on the generated paired images to guarantee their identity consistency. (c): After training the generator, we leverage it to synthesize a great deal of new paired heterogeneous images. Benefiting from the identity consistency property, the generated paired images are regarded as positive pairs. Due to the identity diversity property, the images obtained from different samplings are considered as negative pairs. A contrastive learning mechanism is imposed on the HFR network, yielding both domain-invariant and discriminative embedding features.

[49] synthesizes high-quality images with the guidance of semantic layouts. InterFaceGAN [50] edits face images according to the interpretable latent semantics learned by GANs. In-Domain GAN inversion [51] proposes to invert a pre-trained GAN model for real image editing, resorting to an image reconstruction in the pixel space and a semantic regularization in the latent space.

3 METHOD

The goal of our method is to generate a large number of paired heterogeneous data from noises to boost the performance of HFR. Correspondingly, our method is divided into two parts to tackle the problems of (1) *how to generate diverse paired heterogeneous data* and (2) *how to effectively take advantage of these generated data*. These two parts are introduced in the following two subsections, respectively. We take NIR-VIS as an example for illustration. Other heterogeneous face images are in the same way.

3.1 Dual Generation

As stated in Section 1, we expect the generated paired images to have two properties: identity diversity and identity

TABLE 1: Meaning of the symbols used in our method.

Symbol	Meaning
E_N/E_V	Attribute encoder for NIR/VIS
G	Decoder network
F	Pre-trained face recognition network
F_s	Pre-trained identity sampler
I_N/I_V	NIR/VIS image
z_N/z_V	Attribute distribution of I_N/I_V
f	Identity representation of I_V
\hat{I}_N/\hat{I}_V	Reconstructed NIR/VIS image
\hat{f}_N/\hat{f}_V	Identity representation of \hat{I}_N/\hat{I}_V
\tilde{f}	Identity representation sampled from F_s
\tilde{I}_N/\tilde{I}_V	Generated NIR/VIS image with \tilde{f}
\tilde{f}_N/\tilde{f}_V	Identity representation of \tilde{I}_N/\tilde{I}_V

consistency. On the one hand, in order to promote the identity diversity, we adopt an unconditional generation manner because it has the ability to generate new samples from noises [15], [16]. Meanwhile, considering that the small number of the paired heterogeneous training data may limit the identity diversity of the sampling, we introduce abun-

dant identity information of large-scale VIS data into the generated images. On the other hand, in order to constrain the identity consistency property, we impose a pairwise identity preserving loss on the generated paired images, minimizing their feature distances in the embedding space.

The above ideas are implemented by a dual variational generator and its framework is shown in Figs. 2 (a) and (b). The generator contains two domain-specific attribute encoders E_N and E_V , a decoder G , a pre-trained face recognition network F , and an identity sampler F_s . Among them, E_N and E_V are utilized to learn domain-specific attribute distributions of NIR and VIS data, respectively. F is used to extract identity representations. F_s can flexibly sample abundant identity representations from noises. The joint distribution of the paired heterogeneous data consists of the identity representations and the attribute distributions. G maps the joint distribution to the pixel space. The parameters of F and F_s are fixed, while those of the other networks are updated during the training with both paired heterogeneous data and unpaired VIS data. Table 1 summarizes the meaning of the symbols used in our method.

3.1.1 Training with Paired Heterogeneous Data

Given a pair of NIR-VIS data I_N and I_V with the same identity, the dual variational generator learns disentangled joint distribution in the latent space. Specifically, a face recognition network pre-trained on MS-Celeb-1M [36] is adopted as the feature extractor F . The compact embedding features extracted by F are thought to be only identity related [52]. Meanwhile, considering that F is better at extracting the features of VIS images than those of NIR images, we use $f = F(I_V)$ as the identity representation of both I_N and I_V . Then, the encoders E_N and E_V are leveraged to learn domain-specific attribute distributions $z_N = q_{\phi_N}(z_N|I_N)$ and $z_V = q_{\phi_V}(z_V|I_V)$, respectively. ϕ_N and ϕ_V are the parameters of the encoders. According to the reparameterization trick [15], $z_N = u_N + \sigma_N \odot \epsilon$ and $z_V = u_V + \sigma_V \odot \epsilon$, where u denotes mean and σ denotes standard deviation. \odot denotes Hadamard product and ϵ is a noise sampled from a multi-variate standard Gaussian distribution. Subsequently, in order to ensure that z_N and z_V are merely attribute related, we impose an angular orthogonal loss between the attribute and the identity representations. Finally, the disentangled identity and attribute representations constitute the joint distribution of paired NIR-VIS data, and then are fed to the decoder G to reconstruct the inputs I_N and I_V .

A total of four loss functions are involved in the above process, including an angular orthogonal loss, a distribution learning loss, a pairwise identity preserving loss, and an adversarial loss. These losses are introduced below.

Angular Orthogonal Loss. The angular orthogonal loss is imposed between z_N and f and between z_V and f . For the L_2 normalized z_V and f , their cosine similarity is formulated as:

$$\cos(z_V, f) = \langle z_V, f \rangle, \quad (3)$$

where $\langle \cdot, \cdot \rangle$ denotes inner product. When the two representations are orthogonal, Eq. (3) is equal to zero. Minimizing the absolute value of Eq. (3) will force z_V and f to be orthogonal, and thus they are disentangled. Ultimately, the

angular orthogonal loss that considers both normalized z_N and z_V is defined as:

$$\mathcal{L}_{\text{ort}} = |\langle z_V, f \rangle| + |\langle z_N, f \rangle|. \quad (4)$$

In the following parts, z_N , z_V , and the output of $F(\cdot)$, e.g. $f = F(I_V)$, all denote the normalized ones.

Distribution Learning Loss. Inspired by VAEs [15], the posterior distributions $q_{\phi_N}(z_N|I_N)$ and $q_{\phi_V}(z_V|I_V)$ are constrained by a Kullback-Leibler divergence:

$$\mathcal{L}_{\text{kl}} = D_{\text{KL}}(q_{\phi_N}(z_N|x_N)||p(z_N)) + D_{\text{KL}}(q_{\phi_V}(z_V|x_V)||p(z_V)), \quad (5)$$

where both the priors $p(z_N)$ and $p(z_V)$ are multi-variate standard Gaussian distributions. After obtaining the identity representation f as well as the domain-specific attribute representations z_N and z_V , the decoder network is required to reconstruct the inputs I_N and I_V :

$$\mathcal{L}_{\text{rec}} = -\mathbb{E}_{q_{\phi_N}(z_N|I_N) \cup q_{\phi_V}(z_V|I_V)} \log p_{\theta}(I_N, I_V|f, z_N, z_V). \quad (6)$$

In practice, \mathcal{L}_{rec} is implemented as:

$$\mathcal{L}_{\text{rec}} = \|\hat{I}_N - I_N\|_1 + \|\hat{I}_V - I_V\|_1, \quad (7)$$

where $\hat{I}_N = G(f, z_N)$ and $\hat{I}_V = G(f, z_V)$.

In general, according to [15], the distribution learning loss is a combination of \mathcal{L}_{kl} and \mathcal{L}_{rec} :

$$\mathcal{L}_{\text{dis}} = \mathcal{L}_{\text{kl}} + \mathcal{L}_{\text{rec}}. \quad (8)$$

Pairwise Identity Preserving Loss. In order to preserve the identity of the generated data, previous conditional generation based methods usually adopt an identity preserving loss [10], [14], [53]. A pre-trained face recognition network is used to extract the embedding features of the generated data and those of the real target data respectively, and then the two features are forced to be as close as possible. However, since there are neither intra-class nor inter-class constraints, it is challenging to guarantee that the generated images belong to specific classes consistent with the targets.

As discussed in Section 1, different from the previous methods, we focus on the identity consistency of the generated paired images, rather than whom the generated images belong to. Therefore, we propose a pairwise identity preserving loss that constrains the feature distance between $\hat{f}_N = F(\hat{I}_N)$ and $\hat{f}_V = F(\hat{I}_V)$:

$$\mathcal{L}_{\text{ip-pair}} = 1 - \langle \hat{f}_N, \hat{f}_V \rangle. \quad (9)$$

Minimizing Eq. (9) will increase the cosine similarity between \hat{f}_N and \hat{f}_V . Besides, in order to stabilize the training process, we also constrain the feature distance between \hat{f}_N and f and that between \hat{f}_V and f :

$$\mathcal{L}_{\text{ip-rec}} = (1 - \langle \hat{f}_N, f \rangle) + (1 - \langle \hat{f}_V, f \rangle). \quad (10)$$

The pairwise identity preserving loss that considers both $\mathcal{L}_{\text{ip-pair}}$ and $\mathcal{L}_{\text{ip-rec}}$ is formulated as:

$$\mathcal{L}_{\text{ip}} = \mathcal{L}_{\text{ip-pair}} + \mathcal{L}_{\text{ip-rec}}. \quad (11)$$

Adversarial Loss. Same as [54], an extra adversarial loss \mathcal{L}_{adv} [13] is also introduced to increase the sharpness of the generated data, where the dual variational generator and a discriminator are optimized alternately. We first fix the

discriminator and train the generator with the adversarial loss to confuse the discriminator:

$$\mathcal{L}_{\text{adv-g}} = \log(1 - D(\hat{I}_N)) + \log(1 - D(\hat{I}_V)), \quad (12)$$

where D denotes the discriminator. Then, we fix the generator and train the discriminator to distinguish the generated data from the real data:

$$\mathcal{L}_{\text{adv-d}} = -\log(D(I_N)) - \log(1 - D(\hat{I}_N)) - \log(D(I_V)) - \log(1 - D(\hat{I}_V)). \quad (13)$$

Overall Loss. The overall loss for the training with paired heterogeneous data is a weighted sum of the above distribution learning loss \mathcal{L}_{dis} , the angular orthogonal loss \mathcal{L}_{ort} , the pairwise identity preserving loss \mathcal{L}_{ip} , and the adversarial loss \mathcal{L}_{adv} :

$$\mathcal{L}_{\text{pair}} = \mathcal{L}_{\text{dis}} + \lambda_1 \mathcal{L}_{\text{ort}} + \lambda_2 \mathcal{L}_{\text{ip}} + \lambda_3 \mathcal{L}_{\text{adv}}, \quad (14)$$

where λ_1 , λ_2 , and λ_3 are trade-off parameters. Among them, λ_3 is fixed to 0.1 according to [54].

3.1.2 Training with Unpaired VIS Data

The number of the paired heterogeneous training data is limited, which may affect the identity diversity of the generated images. Therefore, we introduce abundant identity information of large-scale VIS data into the generated images. For the acquisition of the identity information, a straightforward approach is to use a pre-trained face recognition network to extract identities from the large-scale VIS data, as shown in the lower right corner of Fig. 2 (b). However, in this situation, if we desire to generate large-scale new paired data at the testing phase, we must have the same number of VIS data with different identities. This will make our framework a conditional generative model, leading to the diversity problem as mentioned in Section 1. In order to overcome this obstacle, inspired by [55], we introduce an identity sampler to replace the recognition network. Specifically, we first adopt the recognition network to extract the embedding features of MS-Celeb-1M [36], and then leverage these embedding features to train a VAE model [15]. After the training, the decoder of VAE is used as the identity sampler, which can map the points in a standard Gaussian noise to identity representations, as shown in Fig. 2 (b). Equipped with the identity sampler, our framework becomes an unconditional generative model. The required identity representations can be flexibly sampled from noises. A detailed discussion about the face recognition network and the identity sampler is provided in Section 4.3.4.

Since these sampled identity representations do not have corresponding ground truth paired heterogeneous images, we propose to train the generator in an unpaired way. To begin with, we sample an identity representation \tilde{f} via the identity sampler F_s . Then, \tilde{f} as well as the attribute representations z_N and z_V in Eq. (4) are fed into the decoder G . Finally, a pair of new heterogeneous images, *i.e.* $\tilde{I}_N = G(\tilde{f}, z_N)$ and $\tilde{I}_V = G(\tilde{f}, z_V)$, that does not belong to the heterogeneous database is generated. We constrain \tilde{I}_N and \tilde{I}_V from the aspects of appearance and semantic. For appearance, we introduce a small-weight reconstruction

Algorithm 1 Training process of the generator

Require: Heterogeneous training set $I_i, i \in \{N, V\}$; A pre-trained face recognition network F ; A pre-trained identity sampler F_s ;

Ensure: The parameters of the dual variational generator, including E_N, E_V , and G ;

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: # *Training with paired heterogeneous data.*
 - 3: Sample a batch of paired training data $\{I_N, I_V\}$;
 - 4: Use F to extract the identity feature $f = F(I_V)$;
 - 5: Feed I_N, I_V , and f to the generator;
 - 6: Calculate the loss in Eq. (14);
 - 7: # *Training with unpaired VIS data.*
 - 8: Sample a random identity feature \tilde{f} via F_s ;
 - 9: Feed I_N, I_V , and \tilde{f} to the generator;
 - 10: Calculate the loss in Eq. (17);
 - 11: **end for**
 - 12: **return** E_N, E_V , and G .
-

loss [52] to force the appearance of the generated images to be consistent with that of the inputs:

$$\mathcal{L}_{\text{rec}}^u = \eta \left(\|\tilde{I}_N - I_N\|_1 + \|\tilde{I}_V - I_V\|_1 \right), \quad (15)$$

where η is set to 0.1 according to [52]. For semantic, same as the aforementioned pairwise identity preserving manner in Eq. (11), we constrain $\tilde{f}_N = F(\tilde{I}_N)$ and $\tilde{f}_V = F(\tilde{I}_V)$ via:

$$\mathcal{L}_{\text{ip}}^u = (1 - \langle \tilde{f}_N, \tilde{f}_V \rangle) + (1 - \langle \tilde{f}_N, \tilde{f} \rangle) + (1 - \langle \tilde{f}_V, \tilde{f} \rangle). \quad (16)$$

Consequently, the overall loss for the training with unpaired VIS data is a weighted sum of $\mathcal{L}_{\text{rec}}^u$, $\mathcal{L}_{\text{ip}}^u$, and an extra adversarial loss $\mathcal{L}_{\text{adv}}^u$ (replacing \hat{I}_N/\hat{I}_V in Eq. (12) and Eq. (13) with \tilde{I}_N/\tilde{I}_V):

$$\mathcal{L}_{\text{unpair}} = \mathcal{L}_{\text{rec}}^u + \lambda_2 \mathcal{L}_{\text{ip}}^u + \lambda_3 \mathcal{L}_{\text{adv}}^u, \quad (17)$$

where λ_2 and λ_3 are the same as those in Eq. (14). Algorithm 1 shows the training process with both paired heterogeneous data and unpaired VIS data.

3.2 Heterogeneous Face Recognition

After the training of the dual variational generator, we first employ it to generate large-scale paired heterogeneous images, and then leverage these images to facilitate the training of the HFR network. For HFR, we use a face recognition network pre-trained on MS-Celeb-1M as the backbone, and train it with both the limited number of real heterogeneous data I_i ($i \in \{N, V\}$) and the large-scale generated heterogeneous data \tilde{I}_i ($i \in \{N, V\}$).

For the real heterogeneous data, we utilize a softmax loss to optimize the HFR network:

$$\mathcal{L}_{\text{cls}} = \sum_{i \in \{N, V\}} \text{softmax}(f_i, y), \quad (18)$$

where y is the class label of the input I_i and $f_i = F(I_i)$. F denotes the pre-trained face recognition network, same as the feature extractor in Section 3.1. The difference is that the parameters of the former are updated while those of the latter are fixed.

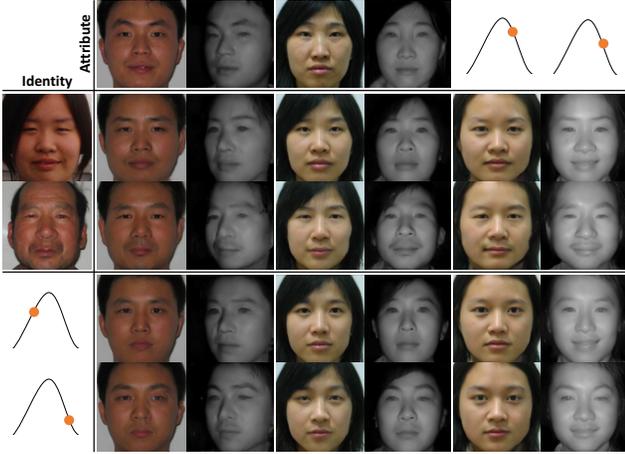


Fig. 3: Disentanglement experiments on CASIA NIR-VIS 2.0. The face images in the first row and the first column are real images, while the remaining face images are generated. After the disentanglement, our method can transfer the image attributes or the sampled attributes (the last two columns) to any identity, including the identities of the real face images and the identities sampled through the identity sampler (the last two rows).

For the generated data, since there is no specific class label, the above softmax loss is inapplicable. However, benefiting from the properties of identity consistency and identity diversity, we propose to take advantage of these generated data via a contrastive learning mechanism [56]. To be specific, as depicted in Fig. 2 (c), we first randomly sample two paired heterogeneous images $(\tilde{I}_N^1, \tilde{I}_V^1)$ and $(\tilde{I}_N^2, \tilde{I}_V^2)$ from the generated database. Based on the identity consistency property, the paired heterogeneous images $(\tilde{I}_N^1, \tilde{I}_V^1)$ and $(\tilde{I}_N^2, \tilde{I}_V^2)$ are set as positive pairs. At the same time, thanks to the identity diversity property, the images obtained from different samplings, *i.e.* $(\tilde{I}_N^1, \tilde{I}_V^2)$ and $(\tilde{I}_N^2, \tilde{I}_V^1)$, are regarded as negative pairs. Note that we do not set $(\tilde{I}_N^1, \tilde{I}_N^2)$ and $(\tilde{I}_V^1, \tilde{I}_V^2)$ as negative pairs because HFR is dedicated to cross-domain matching.

Formally, the contrastive loss for the generated data is:

$$\mathcal{L}_{\text{cont}} = \sum_{j \neq k} (1 - \langle \tilde{f}_N^j, \tilde{f}_V^j \rangle) + (1 - \langle \tilde{f}_N^k, \tilde{f}_V^k \rangle) + \max(0, \langle \tilde{f}_N^j, \tilde{f}_V^k \rangle - m) + \max(0, \langle \tilde{f}_N^k, \tilde{f}_V^j \rangle - m), \quad (19)$$

where $\tilde{f}_N^j = F(\tilde{I}_N^j)$, $\tilde{f}_V^j = F(\tilde{I}_V^j)$, and m is a margin value. By minimizing Eq. (19), the first two terms assist in reducing domain discrepancy, while the last two terms facilitate the learning of discriminative embedding features. Considering both the softmax loss \mathcal{L}_{cls} and the contrastive loss $\mathcal{L}_{\text{cont}}$, the overall loss for the HFR network is given by:

$$\mathcal{L}_{\text{hfr}} = \mathcal{L}_{\text{cls}} + \alpha \mathcal{L}_{\text{cont}}, \quad (20)$$

where α is a trade-off parameter.

4 EXPERIMENTS

In this section, we report insightful analyses of our proposed method and its evaluations against state-of-the-art methods

on seven databases belonging to five HFR tasks, including NIR-VIS, Sketch-Photo, Profile-Frontal Photo, Thermal-VIS, and ID-Camera. **Note that the experiments on the Oulu-CASIA NIR-VIS [57], the Multi-PIE [5], and the NJU-ID [7] databases are presented in Sections 1.2.1, 1.2.2, and 1.2.3 of the supplementary material, respectively.** In the following four subsections, we describe (1) an introduction to the databases and protocols, (2) detailed experimental settings, (3) thorough qualitative and quantitative experimental analyses, and (4) comprehensive comparisons with state-of-the-art methods, respectively.

4.1 Databases and Protocols

CASIA NIR-VIS 2.0 Face Database [58] contains a total of 725 identities, and each identity has 1-22 VIS and 5-50 NIR images. The protocol consists of ten-fold experiments, each including about 6,100 NIR images and 2,500 VIS images from 360 identities for training. Meanwhile, for each fold, 358 VIS images from 358 identities (each identity has one VIS image) constitute the gallery set. Over 6,000 NIR images from the same 358 identities compose the probe set. There is no identity overlap between the training and the testing. We conduct experiments on each fold separately, and then report the mean and standard deviation of the results. The Rank-1 accuracy, Verification Rate (VR) @ False Accept Rate (FAR)=0.1%, and VR@FAR=0.01% are reported.

BUAA-VisNir Face Database [59] is a popular heterogeneous face database that is usually utilized to evaluate domain adaption. It consists of 150 identities, and each identity includes multiple NIR and VIS images. The training set has about 1,200 images from 50 identities, and the testing set contains about 1,300 images from the remaining 100 identities. For testing, only one VIS image per identity is selected as the gallery set and the NIR images constitute the probe set. We report the Rank-1 accuracy, VR@FAR=1%, and VR@FAR=0.1% for comparisons.

IIIT-D Sketch Viewed Database [4] is developed for Sketch-Photo recognition and contains 238 image pairs. Among them, 67 pairs are from the FG-NET aging database, 99 pairs are from the Labeled Faces in Wild (LFW) database, and 72 pairs are from the IIIT-D student & staff database. Since the IIIT-D Sketch Viewed database only has a small number of images, following [18], we use it as the testing set and adopt CUHK Face Sketch FERET (CUFSF) [60] as the training set. CUFSF is also a Sketch-Photo face database that contains 1,194 image pairs. The Rank-1 accuracy is reported.

Tufts Face Database [6] is collected for the study of imaging systems and provides Thermal-VIS face images. There are a total of 113 identities with 74 females and 39 males from more than 15 countries. Each identity has multiple cross-domain face images with different poses. We randomly select 50 identities as the training set, including 1,490 VIS images and 450 thermal images. The remaining 63 identities are used as the testing set with 1,916 VIS images and 530 thermal images. We report the Rank-1 accuracy, VR@FAR=1%, and VR@FAR=0.1%.

The testing process for the HFR network is as follows: Taking the Tufts Face database as an example, the 530 thermal images and the 1,916 VIS images constitute the probe set and the gallery set, respectively. During the testing

TABLE 2: Network architectures of the encoder E_N/E_V and the decoder G . K/S/P denotes kernel size/stride/padding.

(a) Encoder architecture.			(b) Decoder architecture.		
LAYER	K/S/P	OUTPUT	LAYER	K/S/P	OUTPUT
Conv1	5/1/2	32×128^2	FC2	-	2048
Conv2	3/2/1	64×64^2	Deconv1	2/2/0	128×8^2
Conv3	3/2/1	128×32^2	Deconv2	2/2/0	128×16^2
Conv4	3/2/1	256×16^2	Deconv3	2/2/0	64×32^2
Conv5	3/2/1	256×8^2	Deconv4	2/2/0	64×64^2
Conv6	3/2/1	256×4^2	Deconv5	2/2/0	32×128^2
FC1	-	512	Out	1/1/0	3×128^2

phase, we first employ the trained HFR network to extract the embedding features of the images in both the probe and the gallery sets. Then, each thermal feature is matched against each VIS feature via calculating the cosine similarity between them, yielding a similarity matrix of size 530×1916 . Finally, we calculate the Rank-1 accuracy and the verification rate according to the similarity matrix.

4.2 Experimental Settings

For dual generation, we utilize LightCNN [61] pre-trained on MS-Celeb-1M [36] as the feature extractor F . The network architectures of the encoder E_V/E_N and the decoder G are reported in Table 2. ‘FC’ denotes a fully connected layer. ‘Conv’ contains a convolutional layer, an instance normalization (IN) layer, and a ReLU activation layer. Compared with ‘Conv’, ‘DeConv’ replaces the convolutional layer with a transposed convolutional layer and adds a resblock block [62] after ReLU. In addition, inspired by [47], [48], we introduce Adaptive Instance Normalization (AdaIN) [63] to facilitate the learning of attributes. For NIR and VIS, the affine parameters of AdaIN are learned from the attribute representations z_N and z_V via two extra multilayer perceptrons [45], respectively. Following [48], AdaIN is only applied on the low-level layer to avoid destroying the disentanglement of the attribute and the identity. As a result, IN in ‘DeConv5’ is replaced by AdaIN. Furthermore, for NIR and VIS, the weight parameters of the decoder are shared except for these of ‘Out’, which is a single convolutional layer. All of the networks are trained via an Adam optimizer with a fixed learning rate of $2e-4$. The parameters λ_1 and λ_2 in Eq. (14) are set to 50 and 0.5, respectively. We determine the values of the two parameters by balancing the magnitude of the corresponding losses \mathcal{L}_{ort} and \mathcal{L}_{ip} .

For HFR, we adopt the pre-trained LightCNN as the backbone. Face images are first aligned to a resolution of 144×144 , and then are randomly cropped to a resolution of 128×128 as inputs. One batch consists of 64 real heterogeneous images that are randomly sampled from the real training set, and 64 generated paired heterogeneous images that are randomly sampled from the generated dataset. We adopt Stochastic Gradient Descent (SGD) as the optimizer, where the learning rate is set to $1e-3$ initially and is reduced to $5e-4$ gradually. The momentum is set to 0.9 and the weight decay is set to $2e-4$. The trade-off parameter α and the margin m in Eq. (20) are set to $1e-3$ and 0.5, respectively.

TABLE 3: Experimental analyses on Tufts Face. MS (Mean Similarity, higher is better) denotes the mean cosine similarity between the generated paired images. MIS (Mean Instance Similarity, lower is better) presents the mean cosine similarity between two randomly sampled images, including VIS-VIS and NIR-VIS. FID (Fréchet Inception Distance, lower is better) measures the distribution distance between the generated data and the real data. VR@FAR=1% suggests the recognition performance after using the generated data to train the HFR network via contrastive learning.

Method	MS	MIS _{VIS-VIS}	MIS _{NIR-VIS}	FID	VR@FAR=1%
CoGAN [42]	0.23	0.32	0.16	0.77	22.1
DVG [21]	0.52	0.45	0.41	0.65	39.3
DVG-Face	0.53	0.14	0.14	0.57	68.5

4.3 Experimental Analyses

4.3.1 Identity Consistency

As discussed in Section 3, the generated paired heterogeneous images are required to have the same identity. In order to verify this property, we propose to measure MS (Mean Similarity) between the paired images. Specifically, we utilize the pre-trained LightCNN to extract the L_2 normalized embedding features of the paired images, and then calculate their cosine similarity. The MS value is a mean of the similarities of 50K randomly generated paired heterogeneous images. As reported in Table 3, the MS value of DVG-Face is 0.53, while that of the real database is 0.31. It is observed that the former is even larger than the latter, demonstrating that our method does in fact guarantee the identity consistency property. In addition, we can see that the preliminary version DVG also has this property due to the usage of the same pairwise identity preserving manner. By contrast, the MS value of CoGAN [42] is only 0.23 that is significantly smaller than that of the real database. This reveals that it is difficult for the weight-sharing manner of CoGAN to ensure the identity consistency property.

4.3.2 Identity Diversity

Identity diversity plays a crucial role in our method. The more diverse the generated data, the more valuable it will be for the HFR network. Inspired by the instance discrimination criterion [19], [20] that regards each image as an independent class, we propose a MIS (Mean Instance Similarity) metric to measure the identity diversity property. Intuitively, if the generated images have abundant identity diversity, the images obtained from different samplings can be considered to have different identities.

Specifically, by randomly sampling twice from noises, two paired images $(\tilde{I}_N^1, \tilde{I}_V^1)$ and $(\tilde{I}_N^2, \tilde{I}_V^2)$ are generated. Then, we calculate the cosine similarities of the embedding features (extracted by LightCNN) between $(\tilde{I}_V^1, \tilde{I}_V^2)$ and between $(\tilde{I}_N^1, \tilde{I}_N^2)$, yielding IS_{VIS-VIS} (Instance Similarity) and IS_{NIR-VIS}, respectively. After repeating the above sampling and calculation process 50K times, we average these IS_{VIS-VIS} and IS_{NIR-VIS} values to obtain MIS_{VIS-VIS} and MIS_{NIR-VIS} values, respectively. As listed in Table 3, both CoGAN and DVG get a much higher MIS_{VIS-VIS} value than DVG-Face, suggesting the presence of massive similar VIS images in



Fig. 4: Visualization comparisons of dual generation on Tufts Face. (a) CoGAN [42]. (b) DVG [21]. (c) DVG-Face.

TABLE 4: Recognition performances under different identity sampling approaches on Tufts Face. From top to bottom, the first one leverages LightCNN in both the training and the testing stages of dual generation. The second one uses LightCNN in the training stage and the identity sampler in the testing stage. The third one employs the identity sampler in the both stages.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
LightCNN - LightCNN	74.6	67.9	33.7
LighCNN - Identity Sampler	73.5	66.4	31.0
Identity Sampler - Identity Sampler	75.7	68.5	36.5

the sampled data of CoGAN and DVG. We also observe that CoGAN gets a low $MIS_{NIR-VIS}$ value due to the identity inconsistency between the generated paired heterogeneous images. The lowest $MIS_{VIS-VIS}$ and $MIS_{NIR-VIS}$ values of DVG-Face prove that it can really generate more diverse data. In addition, when using the generated data to train the HFR network via the contrastive loss in Eq. (19), it is observed that both CoGAN and DVG obtain a much worse recognition performance than DVG-Face. This is because that the validity of the contrastive loss is based on the identity consistency and the identity diversity properties, neither of which is simultaneously satisfied by CoGAN and DVG.

4.3.3 Distribution Consistency

The distribution of the generated data and that of the real data should be consistent, otherwise it will be difficult to boost the recognition performance with the generated data. FID (Fréchet Inception Distance) [64] is employed to measure the distance between these two distributions in the feature space. In addition, considering the advantage of the face recognition network in extracting facial features, we use LightCNN rather than the conventional Inception model as the feature extractor. Specifically, we first generate 50K paired NIR-VIS data, and then calculate the FID value between the generated NIR data and the real NIR data, as well as the FID value between the generated VIS data and the real VIS data. Subsequently, these two FID values are averaged to get the final FID value. Table 3 suggests that both DVG and DVG-Face obtain a much lower FID value than CoGAN, proving that our method can really learn the distribution of the real data.

Furthermore, Fig. 4 shows the visualization comparisons of CoGAN, DVG, and DVG-Face. We can see that the

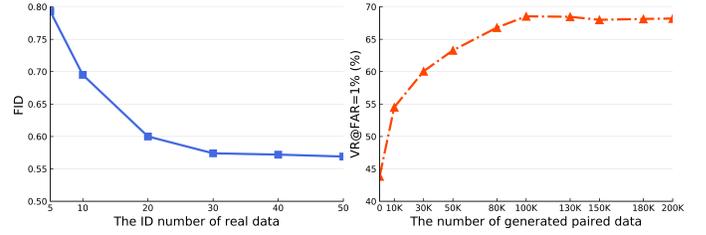


Fig. 5: Studies of the real data (left) and the generated data (right) on Tufts Face.

generated results of CoGAN have many artifacts. It seems difficult for CoGAN to generate good results under such a large domain discrepancy case. Besides, it is observed that the generated paired images of CoGAN are not particularly similar, which is consistent with the MS value in Table 3. By contrast, both DVG and DVG-Face generate photo-realistic results, further demonstrating the distribution consistency property of our method.

In the end, it is obvious that the real training data are crucial for the dual variational generator to learn domain distributions. This motivates us to further investigate the effect of the number of the real data on domain learning. Specifically, apart from the whole training set of Tufts Face that has 1,940 images belonging to 50 identities, we also train the generator on 5 subsets that include 5, 10, 20, 30, and 40 identities with 198, 378, 734, 1,130, and 1,499 images, respectively. Subsequently, we measure the domain distances between the generated data of the above settings and the real data of the whole training set via FID. The lower the FID value, the closer the domain of the generated data is to that of the real data. The experimental results in Fig. 5 (left) suggest that as the number of the trained identities increases, the FID value decreases gradually before 30 identities and then tends to be stable. Thus, 30 identities containing 1,130 real training images seem enough for the generator to learn the real domain.

4.3.4 Identity Sampling

As stated in Section 3.1.2, the identity representation can be obtained from the pre-trained face recognition network (LightCNN [61]) or from the identity sampler. Correspondingly, there are three potential identity sampling approaches: (1) Using LightCNN in both the training and the testing stages of the dual generation. In this case, DVG-Face becomes a conditional generative model. At the testing stage, if we desire to generate large-scale paired data, we must have the same number of VIS data with different identities to provide the required identity representations. Therefore, the sampling diversity is limited by the size of the available VIS data. (2) Using LightCNN in the training stage and the identity sampler in the testing stage. By this means, we no longer need large-scale available VIS data in the testing stage. The required rich identity representations can be easily sampled from noises. However, the identity representations generated by the identity sampler conform to the Gaussian distribution, while those of LightCNN do not satisfy the Gaussian distribution hypothesis. Thus, there may be a distribution gap between the two types of identity representations. (3) Using the identity sampler in both the

TABLE 5: Results under different usages of large-scale VIS data on Tufts Face. (A) trains LightCNN on the large-scale VIS database MS-Celeb-1M. (B) fine-tunes the pre-trained LightCNN on Tufts Face. (C) fine-tunes on both Tufts Face and MS-Celeb-1M. (D) is our method. It introduces the identities of MS-Celeb-1M to enrich the identity diversity of the generated data, and fine-tunes the pre-trained LightCNN on both Tufts Face and the generated data.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
(A)	29.4	23.0	5.3
(B)	54.5	42.4	15.6
(C)	33.9	24.7	7.9
(D)	75.7	68.5	36.5

TABLE 6: Quantitative ablation studies of the angular orthogonal loss \mathcal{L}_{ort} and the pairwise identity preserving loss \mathcal{L}_{ip} on Tufts Face.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
w/o \mathcal{L}_{ort}	55.9	43.6	16.7
w/o \mathcal{L}_{ip}	62.4	57.6	22.1
Ours	75.7	68.5	36.5

training and the testing stages. In this way, it is flexible to obtain massive identity representations in the both stages.

The performances of the above three approaches are shown in Table 4. We observe that the second approach obtains the worst results because of the aforementioned distribution gap. In addition, the third approach gets better results than the first approach, even though the identity sampler is trained on the extracted identity representations of LightCNN. This indicates that the identity sampler may produce new diverse identity representations based on the property of VAEs.

4.3.5 Number of Generated Data

We also investigate the effect of the number of the generated data on the HFR performance. In particular, we use 0, 10K, 30K, 50K, 80K, 100K, 130K, 150K, 180K, and 200K generated paired heterogeneous data to train the HFR network, respectively. Fig. 5 (right) depicts the results under different numbers of generated data on the Tufts Face database. It is clear that from 0 to 100K, the recognition performance is improved with increasing the number of the generated data. Compared with only using the real data (0 generated data), VR@FAR=1% is significantly improved by 26.1% after adding 100K generated paired data. These phenomena demonstrate that the generated data really contain abundant useful information that is helpful for recognition. Besides, as the number of the generated paired data exceeds 100K, no obvious improvement is observed, indicating that the generated data have been saturated.

4.3.6 Usage of Large-Scale VIS Data

Generally, the large-scale VIS database, *i.e.* MS-Celeb-1M [36], is only used to pre-train the HFR network [17], [18]. The intuition behind this is that training the large-scale VIS data together with the small-scale heterogeneous data will lead to severe domain biases that may degrade the performance

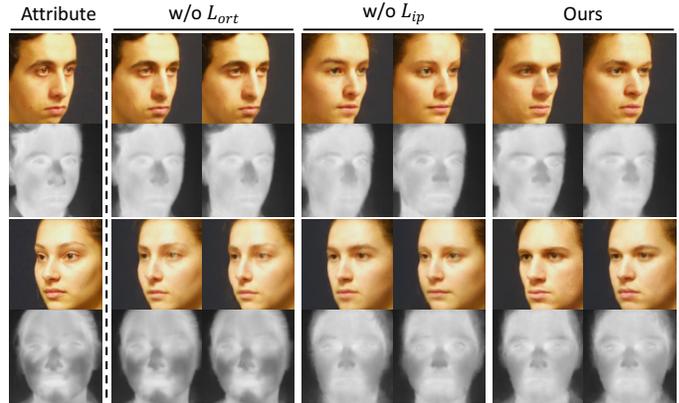


Fig. 6: Qualitative ablation studies of the angular orthogonal loss \mathcal{L}_{ort} and the pairwise identity preserving loss \mathcal{L}_{ip} on Tufts Face. Given the attribute images, these methods generate two pairs of images with two different identity representations.

TABLE 7: Results under different usages of the generated data on Tufts Face. (I) only uses the real data to train the HFR network, and (II) merely uses the generated unlabeled data. (III) leverages both the real data and the generated data, and the latter is utilized via the pairwise distance loss. (IV) is our method that replaces the pairwise distance loss with the contrastive loss.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
(I)	54.5	42.4	15.6
(II)	36.6	24.7	7.2
(III)	70.3	60.8	29.6
(IV)	75.7	68.5	36.5

of HFR. In order to verify this point, four experiments are elaborately designed: (A) directly trains LightCNN on the large-scale VIS database MS-Celeb-1M. (B) fine-tunes the pre-trained LightCNN on the heterogeneous database, and (C) fine-tunes on both the heterogeneous database and MS-Celeb-1M. (D) is our method that fine-tunes on both the heterogeneous database and the generated data.

Table 5 shows that the results of (C) are merely slightly better than those of (A), and are much worse than those of (B), revealing that the domain biases actually affect the performance of HFR. Moreover, (D) outperforms the others by a large margin. This demonstrates the effectiveness of our method that introduces the identity representations of MS-Celeb-1M to enrich the identity diversity of the generated data. In this case, the domain biases are mitigated because the generated data belong to the same domain of the real heterogeneous database.

4.3.7 Ablation Study

For Dual Generation. The angular orthogonal loss \mathcal{L}_{ort} in Eq. (4) and the pairwise identity preserving loss \mathcal{L}_{ip} in Eq. (11) (including \mathcal{L}_{ip}^u in Eq. (16)) are studied to analyze their respective roles. Both quantitative and qualitative results are presented for a comprehensive comparison.

Fig. 6 depicts the qualitative comparisons between our method and its two variants. Without \mathcal{L}_{ort} , the generated im-

TABLE 8: Verification rates on Tufts Face under different parameter values.

(a) λ_1 in Eq. (14)		(b) λ_2 in Eq. (14)		(c) η in Eq. (15)		(d) α in Eq. (20)	
λ_1	FAR=1%	λ_2	FAR=1%	η	FAR=1%	α	FAR=1%
0	43.6	0	57.6	0	46.5	0	42.4
10	68.1	0.10	66.9	0.05	68.1	2e-4	67.9
25	68.5	0.25	67.8	0.10	68.5	5e-4	68.3
50	68.5	0.50	68.5	0.20	67.6	1e-3	68.5
75	67.9	0.75	68.4	0.30	65.7	2e-3	67.6
100	67.4	1.00	68.3	0.40	64.8	3e-3	65.7

TABLE 9: Studies of the margin m in Eq. (19) on Tufts Face.

margin m	0.2	0.3	0.4	0.5	0.6	0.7
VR@FAR=1%	67.0	67.5	68.1	68.5	68.2	67.4

ages are almost unchanged when the identity representation is varied. This means that we cannot enrich the identity diversity of the generated data via injecting more identity information. When omitting \mathcal{L}_{ip} , the generated paired images look inconsistent in identities, suggesting the effectiveness of \mathcal{L}_{ip} . In addition, Table 6 displays the quantitative comparison results. It is observed that the recognition performance drops significantly if one of the losses is discarded, which further proves the importance of each loss. For example, VR@FAR=0.1% decreases by 19.8% when \mathcal{L}_{ort} is omitted.

For HFR. We design four experiments to study the usage of the generated data: (I) only uses the real data to train the HFR network. (II) discards the real data and merely uses the generated data. (III) leverages both the real data and the generated data, and the latter is utilized via the pairwise distance loss as our preliminary version DVG. (IV) is our method that uses the real data and the generated data with the softmax loss and the contrastive loss, respectively.

Table 7 shows that (II) gets the worst results. However, compared with the backbone LightCNN pre-trained on MS-Celeb-1M (Rank-1 = 29.4%, listed in Table 5), (II) still gains an improvement of 7.2% in terms of the Rank-1 accuracy, suggesting the value of the generated data. Compared with (I), the Rank-1 accuracy is boosted at least by 15.8% after introducing the generated data, including both (III) and (IV). This further demonstrates the importance of the generated data. In addition, we find that compared with our preliminary version DVG (Rank-1 = 56.1%, listed in Table 13), (III) improves the Rank-1 accuracy by 14.2%. Since the difference between the two methods lies in the generated data, we attribute the improvement to the richer identity diversity of the generated data of DVG-Face, as discussed in Section 4.3.2. (IV) gains better results than (III), which verifies the effectiveness of the introduced contrastive loss. It optimizes both the intra-class and the inter-class distances, facilitating to learn domain-invariant and discriminative embedding features. By contrast, the pairwise distance loss of DVG can only assist in reducing domain discrepancy. Note that the effectiveness of the contrastive loss is based on the identity consistency and the identity diversity properties of the generated data, as stated in Section 4.3.2.

Table 8 lists the results of sensitivity studies of the parameters λ_1 and λ_2 in Eq. (14), η in Eq. (15), and α in Eq. (20). The parameters λ_1 , λ_2 , and α are set to balance the magnitude of the corresponding loss functions. η is set

TABLE 10: Experimental results on CASIA NIR-VIS 2.0.

Method	Rank-1	VR@FAR=0.1%	VR@FAR=0.01%
H2(LBP3) [65]	43.8	10.1	-
DSIFT [66]	73.3±1.1	-	-
CDFL [67]	71.5±1.4	55.1	-
Gabor+RBM [68]	86.2±1.0	81.3±1.8	-
Recon.+UDP [69]	78.5±1.7	85.8	-
CEFD [25]	85.6	-	-
IDNet [70]	87.1±0.9	74.5	-
HFR-CNN [26]	85.9±0.9	78.0	-
Hallucination [33]	89.6±0.9	-	-
DLFace [71]	98.68	-	-
TRIVET [72]	95.7±0.5	91.0±1.3	74.5±0.7
W-CNN [3]	98.7±0.3	98.4±0.4	94.3±0.4
PACH [14]	98.9±0.2	98.3±0.2	-
RCN [17]	99.3±0.2	98.7±0.2	-
MC-CNN [73]	99.4±0.1	99.3±0.1	-
DVR [27]	99.7±0.1	99.6±0.3	98.6±0.3
DVG [21]	99.8±0.1	99.8±0.1	98.8±0.2
DVG-Face	99.9±0.1	99.9±0.0	99.2±0.1

according to [52]. It is observed that our method is not sensitive to these trade-off parameters in a large range. For instance, when λ_1 is varied from 10 to 100, the verification rate only changes from 67.4% to 68.5%. Besides, Table 8 suggests that as the parameter becomes larger gradually, the recognition performance first increases and then decreases. The optimal setting is $\lambda_1=50$, $\lambda_2=0.5$, $\eta=0.1$, and $\alpha=1e-3$. More sensitivity studies of the parameters λ_1 , λ_2 , η , and α on the BUAA-Vis-Nir, the MultiPIE, and the NJU-ID databases are presented in Tables 4, 5, and 6 of the supplementary material, respectively. Finally, we further explore the setting of the margin parameter m in Eq. (19). Table 9 reveals that the best margin value is 0.5.

Finally, it is worth noting that our method is computationally efficient. When using one TITAN Xp GPU, it takes about 20 minutes to train the generator on the Tufts Face database. In the inference stage of the generator, it merely takes about 4.7 ms to generate a pair of heterogeneous face images. Furthermore, training the HFR network with 100K generated data and about 2K real data costs about 5 hours.

4.4 Comparisons with State-of-the-Art Methods

4.4.1 Results on the CASIA NIR-VIS 2.0 Database

In order to verify the effectiveness of DVG-Face, we compare it with 16 state-of-the-art methods, including 6 traditional methods and 10 deep learning based methods. The traditional methods include H2(LBP3) [65], DSIFT+PCA+LDA [66], CDFL [67], Gabor+RBM [68], Recon.+UDP [69], and CEFD [25]. The deep learning based methods consist of IDNet [70], HFR-CNN [26], Hallucination [33], DLFace [71], TRIVET [72], W-CNN [3], PACH [14], RCN [17], MC-CNN [73], and DVR [27].

Table 10 shows the results of the Rank-1 accuracy and verification rates of the above methods. We can see that most of the deep learning based methods outperform the traditional methods. Nevertheless, the deep learning based method HFR-CNN performs worse than the traditional method Gabor+RBM. This indicates that it is challenging for the deep learning based methods to train on the small-scale heterogeneous data. The methods, such as W-CNN,

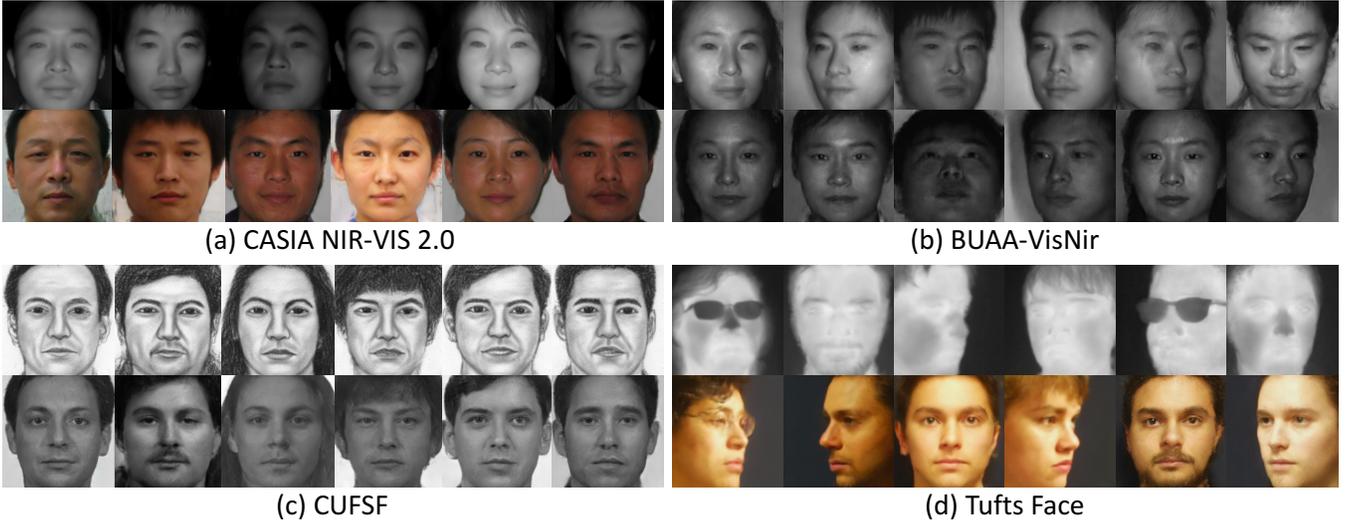


Fig. 7: Dual generation results on (a) CASIA NIR-VIS 2.0 [58], (b) BUAA-VisNir [59], (c) CUFSF [60], and (d) Tufts Face [6]. More generation results, including those on Oulu-CASIA NIR-VIS [57] and Multi-PIE [5], are shown in Fig. 2 of the supplementary material.

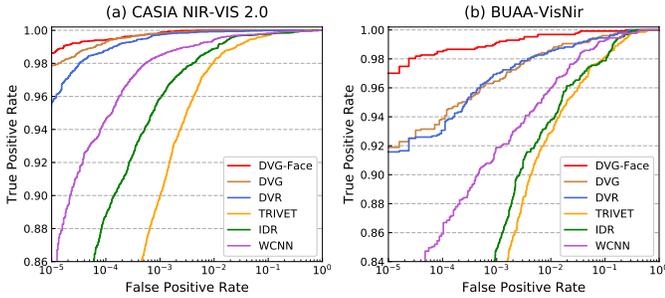


Fig. 8: ROC curves of different methods on (a) CASIA NIR-VIS 2.0 and (b) BUAA-VisNir.

MC-CNN, and DVR, which are devoted to tackling the over-fitting caused by the limited training data, get higher performances. Moreover, our method obtains the best results than all the competitors. In particular, regarding the Rank-1 accuracy and VR@FAR=0.1%, DVG-Face exceeds the conditional generation based method PACH by 1.0% and 1.6% respectively, suggesting the superiority of our unconditional generation approach. Compared with the state-of-the-art method DVR, we improve VR@FAR=0.01% from 98.6% to 99.2%. More importantly, the ROC curves in Fig. 8 (a) show that DVG-Face improves the harder indicator True Positive Rate (TPR)@False Positive Rate (FPR)= 10^{-5} by 3.0% over DVR. In addition, compared with the preliminary version DVG [21], DVG-Face further improves TPR@FPR= 10^{-5} from 97.8% to 98.6%. The improvement highlights the importance of the abundant identity diversity and the consequent contrastive loss. The generated paired NIR-VIS images of DVG-Face are exemplified in Fig. 7 (a).

4.4.2 Results on the BUAA-VisNir Database

A total of 5 traditional methods and 4 deep learning based methods are compared on the BUAA-VisNir database. The former includes MPL3 [57], KCSR [74], KPS [22], KDSR [75],

TABLE 11: Experimental results on BUAA-VisNir.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
MPL3 [57]	53.2	58.1	33.3
KCSR [74]	81.4	83.8	66.7
KPS [22]	66.6	60.2	41.7
KDSR [75]	83.0	86.8	69.5
H2(LBP3) [65]	88.8	88.8	73.4
TRIVET [72]	93.9	93.0	80.9
W-CNN [3]	97.4	96.0	91.9
PACH [14]	98.6	98.0	93.5
DVR [27]	99.2	98.5	96.9
DVG [21]	99.3	98.5	97.3
DVG-Face	99.9	99.7	99.1

and H2(LBP3) [65], and the latter includes TRIVET [72], W-CNN [3], PACH [14], and DVR [27].

Here as well, we observe from Table 11 that the deep learning based methods are always superior to the traditional methods. Moreover, the improvements of DVG-Face over the conditional generation based method PACH are significant. VR@FAR=0.1% is increased from 93.5% to 99.1%. In addition, although the state-of-the-art DVR has achieved very high results, DVG-Face still exceeds it by 2.2% in terms of VR@FAR=0.1%. The ROC curves in Fig. 8 (b) further demonstrate the advantages of our method. Regarding TPR@FPR= 10^{-5} , DVG-Face leads DVR by 5.4% and the preliminary version DVG by 5.1%. The generated paired images of DVG-Face on the BUAA-VisNir database are presented in Fig. 7 (b).

4.4.3 Results on the IIIT-D Sketch Viewed Database

Considering the crucial role of the Sketch-Photo recognition in criminal investigation, we further evaluate our method on the IIIT-D Sketch Viewed database. The compared methods include 5 traditional methods: Original WLD [76], SIFT [77], EUCLBP [78], LFDA [24], MCWLD [4], as well as 4 deep

TABLE 12: Experimental results on IIIT-D Sketch Viewed.

Method	Rank-1
Original WLD [76]	74.34
SIFT [77]	76.28
EUCLBP [78]	79.36
LFDA [24]	81.43
MCWLD [4]	84.24
LightCNN [61]	83.24
CDL [18]	85.35
MC-CNN [73]	87.40
RCN [17]	90.34
DVG [21]	96.99
DVG-Face	97.21

TABLE 13: Experimental results on Tufts Face.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
LightCNN [61]	29.4	23.0	5.3
DVG [21]	56.1	44.3	17.1
DVG-Face	75.7	68.5	36.5

learning based methods: LightCNN [61], CDL [18], RCN [17], and MC-CNN [73].

Table 12 suggests that on such a Sketch-Photo database, the overall performance discrepancy between the traditional methods and the deep learning based methods is not particularly large. The Rank-1 accuracy of MCWLD is even better than that of LightCNN, and is comparable with that of CDL. In addition, our method exhibits superior performances over all the competitors. The best Rank-1 accuracy is improved from 90.34% [17] to 97.21%. Fig. 7 (c) depicts the dual generation results of Sketch-Photo.

4.4.4 Results on the Tufts Face Database

As mentioned in Section 4.3.6, the face recognition network pre-trained on MS-Celeb-1M merely obtains a Rank-1 accuracy of 29.4% on the Tufts Face database, indicating the large domain discrepancy between the thermal data and the VIS data. At the same time, it is observed that after fine-tuning on the training set of Tufts Face, the Rank-1 accuracy only increases to 54.5% that is still unsatisfactory. This phenomenon is caused by the limited number of the Thermal-VIS training data. Obviously, data augmentation is a straightforward approach to alleviate this problem. Both the preliminary version DVG and the current DVG-Face aim to generate more data to facilitate the training of the HFR network. The comparisons of DVG, DVG-Face, and the baseline LightCNN are reported in Table 13. We can see that compared with DVG, DVG-Face boosts the Rank-1 accuracy from 56.1% to 75.7%, and VR@FAR=0.1% from 17.1% to 36.5%. The dual generation results on Tufts Face are shown in Fig. 7 (d).

5 CONCLUSION

This paper has proposed a novel DVG-Face framework that generates large-scale new paired heterogeneous images from noises to boost the performance of HFR. To begin

with, a dual variational generator is elaborately designed to train with both paired heterogeneous data and unpaired VIS data. The introduction of the latter greatly promotes the identity diversity of the generated images. Subsequently, a pairwise identity preserving loss is imposed on the generated paired images to guarantee their identity consistency. Finally, benefiting from the identity consistency and identity diversity properties, the generated unlabeled images can be employed to train the HFR network via contrastive learning. Our method obtains the best results on seven databases, exposing a new way to the HFR problem.

ACKNOWLEDGMENTS

The authors would like to greatly thank the associate editor and the reviewers for their valuable comments and advice. This work is funded by Beijing Natural Science Foundation (Grant No. JQ18017), National Natural Science Foundation of China (Grant No. 61721004, U20A20223) and Youth Innovation Promotion Association CAS (Grant No. Y201929).

REFERENCES

- [1] Y. Wen, K. Zhang, Z. Li, and Y. Qiao, "A discriminative feature learning approach for deep face recognition," in *European Conference on Computer Vision*, 2016.
- [2] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [3] 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.
- [4] H. S. Bhatt, S. Bharadwaj, R. Singh, and M. Vatsa, "Memetic approach for matching sketches with digital face images," *Tech. Rep.*, 2012.
- [5] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multiple," *Image and Vision Computing*, vol. 28, no. 5, pp. 807–813, 2010.
- [6] K. Panetta, Q. Wan, S. Agaian, S. Rajeev, S. Kamath, R. Rajendran, S. Rao, A. Kaszowska, H. Taylor, A. Samani, and Y. Xin, "A comprehensive database for benchmarking imaging systems," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 3, pp. 509–520, 2018.
- [7] J. Huo, Y. Gao, Y. Shi, W. Yang, and H. Yin, "Heterogeneous face recognition by margin-based cross-modality metric learning," *IEEE Transactions on Cybernetics*, vol. 48, no. 6, pp. 1814–1826, 2017.
- [8] S. Ouyang, T. Hospedales, Y.-Z. Song, X. Li, C. C. Loy, and X. Wang, "A survey on heterogeneous face recognition: Sketch, infra-red, 3d and low-resolution," *Image and Vision Computing*, vol. 56, pp. 28–48, 2016.
- [9] L. Song, M. Zhang, X. Wu, and R. He, "Adversarial discriminative heterogeneous face recognition," in *AAAI Conference on Artificial Intelligence*, 2018.
- [10] R. Huang, S. Zhang, T. Li, and R. He, "Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis," in *IEEE International Conference on Computer Vision*, 2017.
- [11] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [12] H. Zhang, B. S. Riggan, S. Hu, N. J. Short, and V. M. Patel, "Synthesis of high-quality visible faces from polarimetric thermal faces using generative adversarial networks," *International Journal of Computer Vision*, vol. 127, no. 6-7, pp. 845–862, 2019.
- [13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, 2014.
- [14] 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.
- [15] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *International Conference on Learning Representations*, 2014.

- [16] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of gans for improved quality, stability, and variation," in *International Conference on Learning Representations*, 2018.
- [17] Z. Deng, X. Peng, and Y. Qiao, "Residual compensation networks for heterogeneous face recognition," in *AAAI Conference on Artificial Intelligence*, 2019.
- [18] X. Wu, L. Song, R. He, and T. Tan, "Coupled deep learning for heterogeneous face recognition," in *AAAI Conference on Artificial Intelligence*, 2018.
- [19] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, "Unsupervised feature learning via non-parametric instance discrimination," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [20] P. Bachman, R. D. Hjelm, and W. Buchwalter, "Learning representations by maximizing mutual information across views," in *Advances in Neural Information Processing Systems*, 2019.
- [21] 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.
- [22] 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.
- [23] M. Shao, D. Kit, and Y. Fu, "Generalized transfer subspace learning through low-rank constraint," *International Journal of Computer Vision*, vol. 109, no. 1-2, pp. 74–93, 2014.
- [24] 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.
- [25] D. Gong, Z. Li, W. Huang, X. Li, and D. Tao, "Heterogeneous face recognition: A common encoding feature discriminant approach," *IEEE Transactions on Image Processing*, vol. 26, no. 5, pp. 2079–2089, 2017.
- [26] S. Saxena and J. Verbeek, "Heterogeneous face recognition with cnns," in *European Conference on Computer Vision*, 2016.
- [27] 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.
- [28] X. Tang and X. Wang, "Face sketch synthesis and recognition," in *IEEE International Conference on Computer Vision*, 2003.
- [29] C. Peng, X. Gao, N. Wang, D. Tao, X. Li, and J. Li, "Multiple representations-based face sketch-photo synthesis," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 27, no. 11, pp. 2201–2215, 2015.
- [30] C. Peng, X. Gao, N. Wang, and J. Li, "Superpixel-based face sketch-photo synthesis," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 27, no. 2, pp. 288–299, 2015.
- [31] D.-A. Huang and Y.-C. Frank Wang, "Coupled dictionary and feature space learning with applications to cross-domain image synthesis and recognition," in *IEEE International Conference on Computer Vision*, 2013.
- [32] G. Hu, X. Peng, Y. Yang, T. M. Hospedales, and J. Verbeek, "Frankenstein: Learning deep face representations using small data," *IEEE Transactions on Image Processing*, vol. 27, no. 1, pp. 293–303, 2017.
- [33] 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.
- [34] B. Cao, N. Wang, J. Li, and X. Gao, "Data augmentation-based joint learning for heterogeneous face recognition," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 6, pp. 1731–1743, 2018.
- [35] M. Zhang, Y. Li, N. Wang, Y. Chi, and X. Gao, "Cascaded face sketch synthesis under various illuminations," *IEEE Transactions on Image Processing*, vol. 29, pp. 1507–1521, 2019.
- [36] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, "Ms-celeb-1m: A dataset and benchmark for large-scale face recognition," in *European Conference on Computer Vision*, 2016.
- [37] A. Van den Oord, N. Kalchbrenner, L. Espeholt, k. kavukcuoglu, O. Vinyals, and A. Graves, "Conditional image generation with pixelcnn decoders," in *Advances in Neural Information Processing Systems*, 2016.
- [38] D. P. Kingma and P. Dhariwal, "Glow: Generative flow with invertible 1x1 convolutions," in *Advances in Neural Information Processing Systems*, 2018.
- [39] H. Huang, Z. Li, R. He, Z. Sun, and T. Tan, "Introvae: Introspective variational autoencoders for photographic image synthesis," in *Advances in Neural Information Processing Systems*, 2018.
- [40] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," in *Advances in Neural Information Processing Systems*, 2017.
- [41] T. R. Shaham, T. Dekel, and T. Michaeli, "Singan: Learning a generative model from a single natural image," in *IEEE International Conference on Computer Vision*, 2019.
- [42] M.-Y. Liu and O. Tuzel, "Coupled generative adversarial networks," in *Advances in Neural Information Processing Systems*, 2016.
- [43] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *IEEE International Conference on Computer Vision*, 2017.
- [44] M. Li, J. Lin, Y. Ding, Z. Liu, J.-Y. Zhu, and S. Han, "Gan compression: Efficient architectures for interactive conditional gans," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [45] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz, "Multimodal unsupervised image-to-image translation," in *European Conference on Computer Vision*, 2018.
- [46] A. Brock, J. Donahue, and K. Simonyan, "Large scale gan training for high fidelity natural image synthesis," in *International Conference on Learning Representations*, 2019.
- [47] Y. Choi, Y. Uh, J. Yoo, and J.-W. Ha, "Stargan v2: Diverse image synthesis for multiple domains," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [48] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [49] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [50] Y. Shen, J. Gu, X. Tang, and B. Zhou, "Interpreting the latent space of gans for semantic face editing," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [51] J. Zhu, Y. Shen, D. Zhao, and B. Zhou, "In-domain gan inversion for real image editing," in *European Conference on Computer Vision*, 2020.
- [52] J. Bao, D. Chen, F. Wen, H. Li, and G. Hua, "Towards open-set identity preserving face synthesis," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [53] Y. Hu, X. Wu, B. Yu, R. He, and Z. Sun, "Pose-guided photorealistic face rotation," in *IEEE International Conference on Computer Vision*, 2018.
- [54] Z. Shu, M. Sahasrabudhe, R. A. Güler, D. Samaras, N. Paragios, and I. Kokkinos, "Deforming autoencoders: Unsupervised disentangling of shape and appearance," in *European Conference on Computer Vision*, 2018.
- [55] Y. Deng, J. Yang, D. Chen, F. Wen, and X. Tong, "Disentangled and controllable face image generation via 3d imitative-contrastive learning," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [56] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2006.
- [57] J. Chen, D. Yi, J. Yang, G. Zhao, S. Z. Li, and M. Pietikainen, "Learning mappings for face synthesis from near infrared to visual light images," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [58] S. Li, D. Yi, Z. Lei, and S. Liao, "The casia nir-vis 2.0 face database," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2013.
- [59] D. Huang, J. Sun, and Y. Wang, "The buaa-visnir face database instructions," Tech. Rep., 2012.
- [60] W. Zhang, X. Wang, and X. Tang, "Coupled information-theoretic encoding for face photo-sketch recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2011.
- [61] 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.
- [62] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [63] X. Huang and S. Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization," in *IEEE International Conference on Computer Vision*, 2017.
- [64] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, G. Klambauer, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a nash equilibrium," in *Advances in Neural Information Processing Systems*, 2017.

- [65] M. Shao and Y. Fu, "Cross-modality feature learning through generic hierarchical hyperlingual-words," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, no. 2, pp. 451–463, 2016.
- [66] T. I. Dhamecha, P. Sharma, R. Singh, and M. Vatsa, "On effectiveness of histogram of oriented gradient features for visible to near infrared face matching," in *International Conference on Pattern Recognition*, 2014.
- [67] Y. Jin, J. Lu, and Q. Ruan, "Coupled discriminative feature learning for heterogeneous face recognition," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 3, pp. 640–652, 2015.
- [68] D. Yi, Z. Lei, and S. Z. Li, "Shared representation learning for heterogeneous face recognition," in *IEEE International Conference and Workshops on Automatic Face and Gesture Recognition*, 2015.
- [69] F. Juefei-Xu, D. K. Pal, and M. Savvides, "Nir-vis heterogeneous face recognition via cross-spectral joint dictionary learning and reconstruction," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2015.
- [70] 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.
- [71] C. Peng, N. Wang, J. Li, and X. Gao, "Dlface: Deep local descriptor for cross-modality face recognition," *Pattern Recognition*, vol. 90, pp. 161–171, 2019.
- [72] X. Liu, L. Song, X. Wu, and T. Tan, "Transferring deep representation for nir-vis heterogeneous face recognition," in *International Conference on Biometrics*, 2016.
- [73] 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.
- [74] Z. Lei and S. Z. Li, "Coupled spectral regression for matching heterogeneous faces," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [75] X. Huang, Z. Lei, M. Fan, X. Wang, and S. Z. Li, "Regularized discriminative spectral regression method for heterogeneous face matching," *IEEE Transactions on Image Processing*, vol. 22, no. 1, pp. 353–362, 2012.
- [76] J. Chen, S. Shan, C. He, G. Zhao, M. Pietikainen, X. Chen, and W. Gao, "Wld: A robust local image descriptor," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 9, pp. 1705–1720, 2009.
- [77] B. Klare and A. K. Jain, "Sketch-to-photo matching: a feature-based approach," in *Biometric Technology for Human Identification VII*, 2010.
- [78] H. S. Bhatt, S. Bharadwaj, R. Singh, and M. Vatsa, "On matching sketches with digital face images," in *IEEE International Conference on Biometrics: Theory, Applications and Systems*, 2010.

Supplementary Material

TABLE 1: Experimental results on Oulu-CASIA NIR-VIS.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
MPL3	48.9	41.9	11.4
KCSR	66.0	49.7	26.1
KPS	62.2	48.3	22.2
KDSR	66.9	56.1	31.9
H2(LBP3)	70.8	62.0	33.6
TRIVET	92.2	67.9	33.6
W-CNN	98.0	81.5	54.6
PACH	100.0	97.9	88.2
DVR	100.0	97.2	84.9
DVG	100.0	98.5	92.9
DVG-Face	100.0	99.2	97.3

1 EXPERIMENTS

1.1 Databases and Protocols

Oulu-CASIA NIR-VIS Database is a widely utilized NIR-VIS database that contains 80 identities with various poses and expressions. 50 identities are from the University of Oulu and the other 30 identities are from CASIA. Following the protocol of [1], only 40 identities are used in experiments, with 20 used for training and 20 used for testing. There are 48 NIR images and 48 VIS images for each identity. For testing, all of the VIS images from the 20 identities are adopted as the gallery set and the corresponding NIR images are used as the probe set. Evaluation indicators include the Rank-1 accuracy, VR@FAR=1%, and VR@FAR=0.1%.

Multi-PIE Database is often employed for multi-pose evaluations. There are 4 sessions with a total of 337 identities, and each identity has abundant illuminations, poses, and expressions. According to Setting 2 of [2], we only leverage the face images with the neutral expression under 20 illuminations. The first 200 identities are used for training and the remaining 137 identities are used for testing. For testing, the gallery set has one frontal face image per identity, and the probe set consists of $\pm 90^\circ$ face images. The Rank-1 accuracy is reported.

NJU-ID Database is built to study ID-Camera face verification. This database has a total of 256 identities, and each identity contains one ID card image with 102×126 resolution and one camera image with 640×480 resolution. Since the NJU-ID database only has a limited number of images, we merely use it as the testing set, and collect a new ID-Camera database as the training set that has 1,190 image pairs. For testing, the camera images are used as the probe set, and the ID card images are used as the gallery set. The Rank-1 accuracy, VR@FAR=1%, and VR@FAR=0.1% are reported.

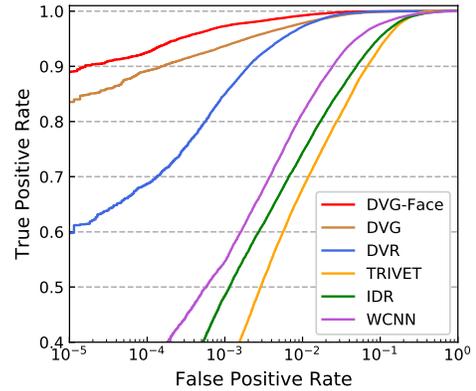


Fig. 1: ROC curves on Oulu-CASIA NIR-VIS.

TABLE 2: Experimental results under $\pm 90^\circ$ on MultiPIE.

Method	Rank-1
GridFace [3]	75.4
FF-GAN [4]	61.2
TPGAN [5]	64.6
CAPG-GAN [6]	66.1
PIM [7]	86.5
Rotate-and-Render [8]	94.4
DVG	83.9
DVG-Face	97.6

1.2 Comparisons with State-of-the-Art Methods

1.2.1 Results on the Oulu-CASIA NIR-VIS Database

Different from the CASIA NIR-VIS 2.0 database, the training set of the Oulu-CASIA NIR-VIS database only has 20 identities. Such a low-shot database poses huge challenges for the training of the HFR network. We evaluate our method against 5 traditional methods and 4 deep learning based methods. The compared traditional methods include MPL3, KCSR, KPS, KDSR, and H2(LBP3). The compared deep learning based methods include TRIVET, W-CNN, PACH, and DVR.

In Table 1, it is observed that the deep learning based methods significantly outperform the traditional methods in terms of the Rank-1 accuracy. At the same time, we also notice that compared with the Rank-1 accuracy, verification rates of all of the methods are not particularly high. This may be caused by the testing protocol, in which there are merely 20 identities in the testing set, and each identity contains up to 48 VIS images as the galleries and 48 NIR images as the probes. As a consequence, it is relatively

TABLE 3: Experimental results on NJU-ID.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
LightCNN	91.3	90.5	80.4
DVG	96.8	96.7	83.0
DVG-Face	98.4	98.1	91.2

TABLE 4: Verification rates on BUAA-VisNir under different parameter values.

(a) λ_1 in Eq. (14)		(b) λ_2 in Eq. (14)		(c) η in Eq. (15)		(d) α in Eq. (20)	
λ_1	FAR=0.1%	λ_2	FAR=0.1%	η	FAR=0.1%	α	FAR=0.1%
0	97.0	0	95.3	0	96.6	0	95.1
10	98.6	0.10	97.6	0.05	98.6	2e-4	98.3
25	98.9	0.25	98.9	0.10	99.1	5e-4	98.9
50	99.1	0.50	99.1	0.20	98.7	1e-3	99.1
75	98.8	0.75	98.7	0.30	97.6	2e-3	98.4
100	98.4	1.00	98.4	0.40	97.0	3e-3	97.2

easy to retrieve the correct image according to the highest similarity score, leading to the generally high Rank-1 results. However, due to the large domain discrepancy and the small number of the training data, the discriminability of the learned features is limited. This results in the relatively low verification rates. In such a challenging case, our method still gets the highest performances. In particular, the best VR@FAR=0.1% is improved from 88.2% to 97.3%. The ROC curves in Fig. 1 further demonstrate the significant superiority of our method. Compared with the state-of-the-art DVR, TPR@FPR@ 10^{-5} is improved by 29.2%. Moreover, regarding TPR@FPR@ 10^{-5} , we beat the preliminary version DVG by 5.5%. The dual generation results on Oulu-CASIA NIR-VIS are displayed in Fig. 2 (b).

1.2.2 Results on the MultiPIE Database

Since the recognition under extreme poses is still a challenging problem, we further evaluate our method on the representative MultiPIE database. The compared methods include 6 state-of-the-art methods: GridFace [3], FF-GAN [4], TPGAN [5], CAPG-GAN [6], PIM [7], and Rotate-and-Render [8]. All of these methods except for GridFace are based on conditional generation.

Table 2 suggests that our method outperforms all of the compared methods. The best Rank-1 accuracy under the most challenging $\pm 90^\circ$ is improved from 94.4% to 97.6%. Compared with the preliminary version DVG, we obtain an improvement of 13.7%. The generated paired images on the MultiPIE database are shown in Fig. 2 (e).

1.2.3 Results on the NJU-ID Database

With the popularization of face authentication systems, the ID-Camera recognition has become increasingly important. Therefore, we also evaluate our method on the representative ID-Camera database NJU-ID. The results of LightCNN, DVG, and DVG-Face are reported in Table 3. It is obvious that DVG-Face shows clear superiority over the preliminary version DVG in terms of the Rank-1 accuracy, VR@FAR=1%, and VR@FAR=0.1%. In particular, VR@FAR=0.1% is improved by 8.2%. We do not display the generated data because it may violate privacy.

TABLE 5: Rank-1 accuracies on MultiPIE under different parameter values.

(a) λ_1 in Eq. (14)		(b) λ_2 in Eq. (14)		(c) η in Eq. (15)		(d) α in Eq. (20)	
λ_1	Rank-1	λ_2	Rank-1	η	Rank-1	α	Rank-1
0	84.3	0	83.5	0	83.3	0	84.6
10	95.8	0.10	96.7	0.05	96.8	2e-4	97.1
25	96.9	0.25	97.6	0.10	97.6	5e-4	97.3
50	97.6	0.50	97.6	0.20	96.5	1e-3	97.6
75	96.3	0.75	96.7	0.30	95.9	2e-3	96.7
100	96.1	1.00	96.3	0.40	94.8	3e-3	95.8

TABLE 6: Verification rates on NJU-ID under different parameter values.

(a) λ_1 in Eq. (14)		(b) λ_2 in Eq. (14)		(c) η in Eq. (15)		(d) α in Eq. (20)	
λ_1	FAR=1%	λ_2	FAR=1%	η	FAR=1%	α	FAR=1%
0	96.5	0	95.3	0	96.9	0	91.3
10	97.7	0.10	96.7	0.05	97.8	2e-4	96.6
25	98.1	0.25	97.6	0.10	98.1	5e-4	97.4
50	98.1	0.50	98.1	0.20	97.1	1e-3	98.1
75	97.9	0.75	97.8	0.30	96.5	2e-3	97.3
100	97.3	1.00	97.6	0.40	95.7	3e-3	96.2

REFERENCES

- [1] M. Shao and Y. Fu, "Cross-modality feature learning through generic hierarchical hyperlingual-words," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, no. 2, pp. 451–463, 2016.
- [2] J. Yim, H. Jung, B. Yoo, C. Choi, D. Park, and J. Kim, "Rotating your face using multi-task deep neural network," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [3] E. Zhou, Z. Cao, and J. Sun, "Gridface: Face rectification via learning local homography transformations," in *European Conference on Computer Vision*, 2018.
- [4] X. Yin, X. Yu, K. Sohn, X. Liu, and M. Chandraker, "Towards large-pose face frontalization in the wild," in *IEEE International Conference on Computer Vision*, 2017.
- [5] R. Huang, S. Zhang, T. Li, and R. He, "Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis," in *IEEE International Conference on Computer Vision*, 2017.
- [6] Y. Hu, X. Wu, B. Yu, R. He, and Z. Sun, "Pose-guided photorealistic face rotation," in *IEEE International Conference on Computer Vision*, 2018.
- [7] J. Zhao, Y. Cheng, Y. Xu, L. Xiong, J. Li, F. Zhao, K. Jayashree, S. Pranata, S. Shen, J. Xing, S. Yan, and J. Feng, "Towards pose invariant face recognition in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [8] H. Zhou, J. Liu, Z. Liu, Y. Liu, and X. Wang, "Rotate-and-render: unsupervised photorealistic face rotation from single-view images," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.

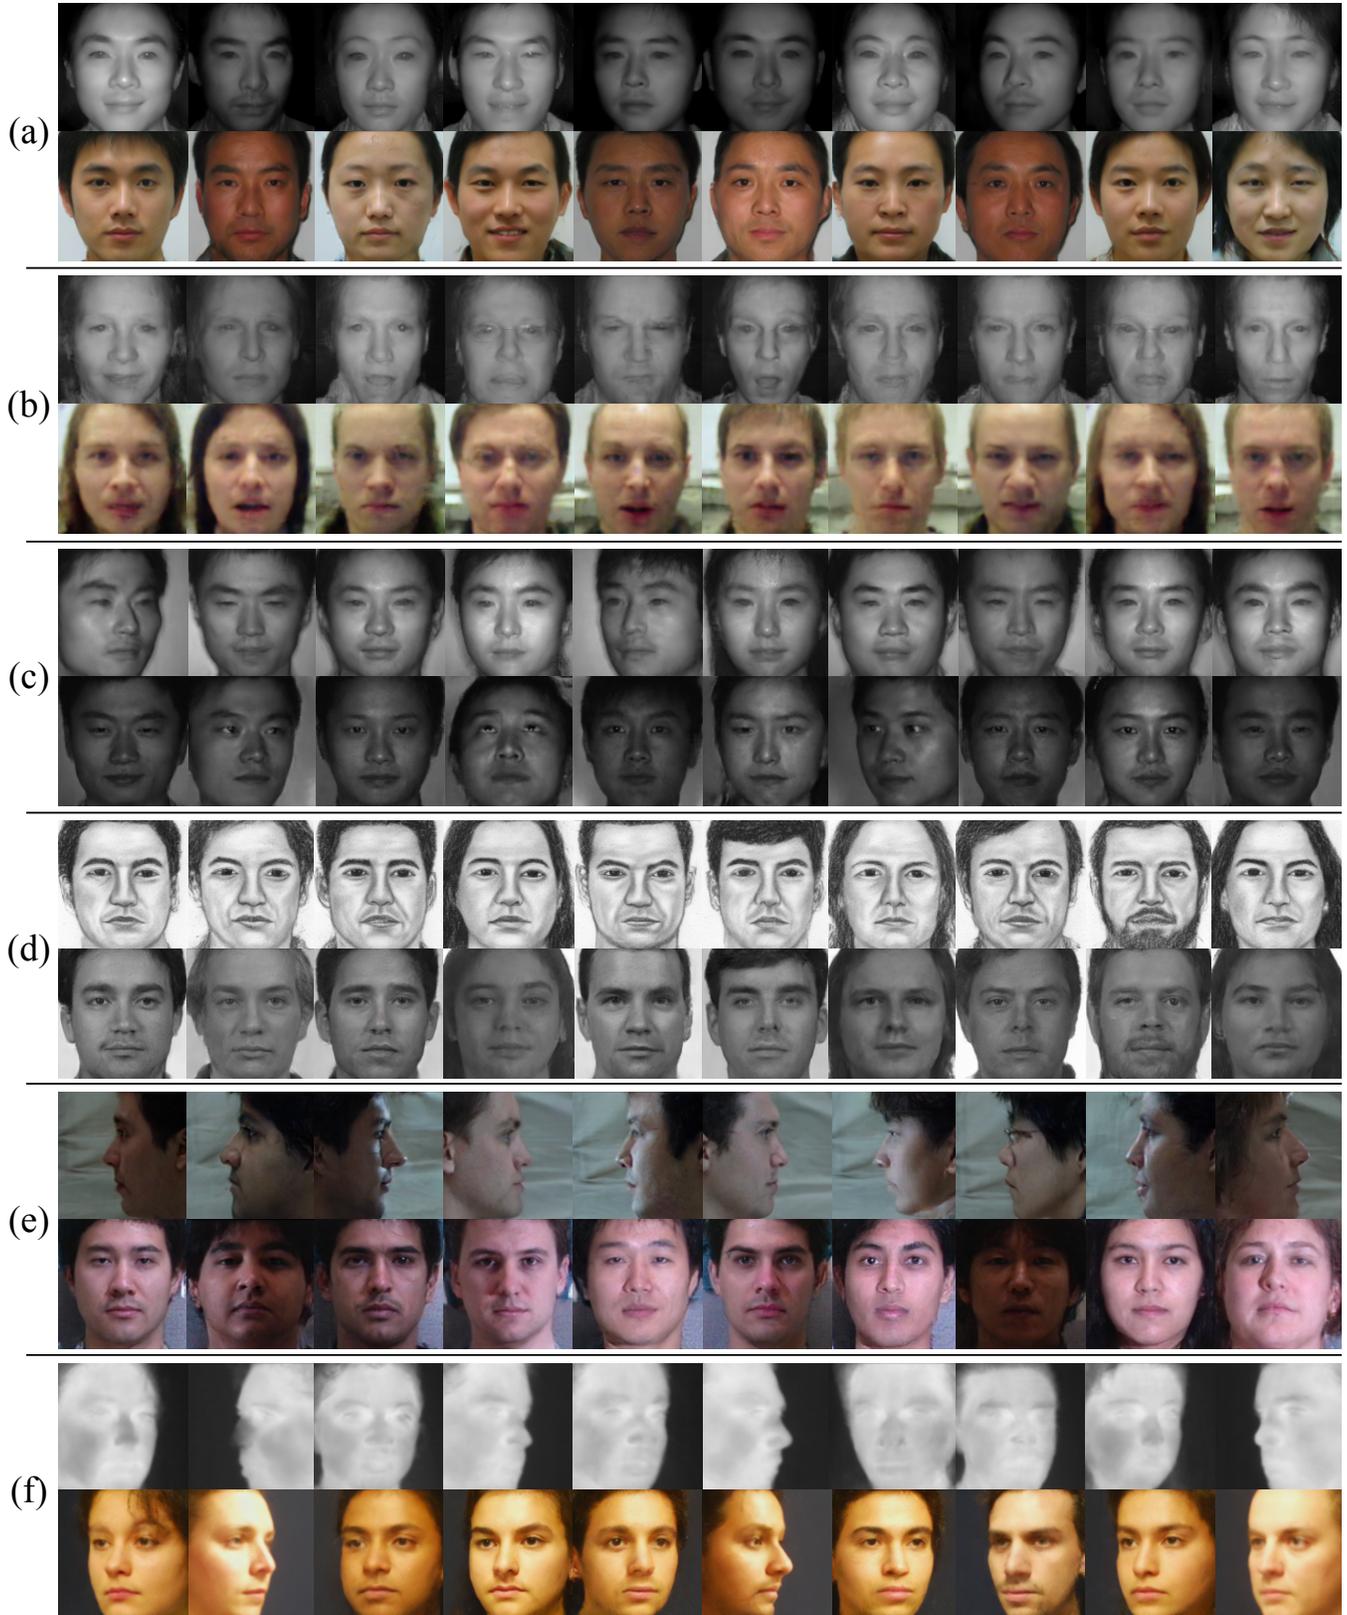


Fig. 2: Dual generation results on (a) CASIA NIR-VIS 2.0, (b) Oulu-CASIA NIR-VIS, (c) BUAA-VisNir, (d) CUFSE, (e) Multi-PIE, and (f) Tufts Face.