

Polarimetric Thermal to Visible Face Verification via Attribute Preserved Synthesis

Xing Di¹, He Zhang², Vishal M. Patel¹

¹Johns Hopkins University, 3400 N. Charles St, Baltimore, MD 21218, USA

²Rutgers University, 94 Brett Rd, Piscataway Township, NJ 08854, USA

xdi1@jhu.edu, he.zhang92@rutgers.edu, vpatel136@jhu.edu

Abstract

Thermal to visible face verification is a challenging problem due to the large domain discrepancy between the modalities. Existing approaches either attempt to synthesize visible faces from thermal faces or extract robust features from these modalities for cross-modal matching. In this paper, we take a different approach in which we make use of the attributes extracted from the visible image to synthesize the attribute-preserved visible image from the input thermal image for cross-modal matching. A pre-trained VGG-Face network is used to extract the attributes from the visible image. Then, a novel Attribute Preserved Generative Adversarial Network (AP-GAN) is proposed to synthesize the visible image from the thermal image guided by the extracted attributes. Finally, a deep network is used to extract features from the synthesized image and the input visible image for verification. Extensive experiments on the ARL Polarimetric face dataset show that the proposed method achieves significant improvements over the state-of-the-art methods.

1. Introduction

Face Recognition (FR) is one of the most widely studied problems in computer vision and biometrics research communities due to its applications in authentication, surveillance and security. Various methods have been developed over the last two decades that specifically attempt to address the challenges such as aging, occlusion, disguise, variations in pose, expression and illumination. In particular, convolutional neural network (CNN) based FR methods have gained a lot of traction in recent years [24, 23]. Deep CNN-based methods [19, 29, 35, 2, 20, 21] have achieved impressive performances on the current FR benchmarks.

Despite the success of CNN-based methods in addressing various challenges in FR, they are fundamentally limited to recognize face images that are collected near-visible




Figure 1. Sample Stokes as well as polarimetric and visible images corresponding to a subject in the ARL dataset [9].

spectrum. In many practical scenarios such as surveillance in low-light conditions, one has to detect and recognize faces that are captured using thermal modalities [9, 27, 31, 36, 26, 14, 18, 16, 1]. However, the performance of many deep learning-based methods degrades significantly when they are presented with thermal face images. For example, it was shown in [36, 26] that simply using deep features extracted from both raw polarimetric thermal and visible facial images are not sufficient enough for cross-domain face recognition. The performance degradation is mainly due to the significant distributional change between the thermal and visible domains as well as a lack of sufficient data for training the deep networks for cross-modal matching.

In many recent approaches, the polarization-state information of thermal emissions has been used to achieve improved cross-spectrum face recognition performance [9, 27, 31, 36, 26] since it captures geometric and textural details of faces that are not present in the conventional thermal facial images [31, 9]. A polarimetric thermal image consists

of four Stokes images: S_0 , S_1 , S_2 , and degree-of-linear-polarization (DoLP), where S_0 indicates the conventional total intensity thermal image, S_1 captures the horizontal and vertical polarization-state information, S_2 captures the diagonal polarization-state information and DoLP describes the portion of an electromagnetic wave that is linearly polarized [9]. These Stokes images along with the visible and the polarimetric images corresponding to a subject in the ARL dataset [9] are shown in Figure 1. It can be observed that S_1 , S_2 and DoLP tend to preserve more textural details compared to S_0 . Similar to [36, 26], we also refer to Polar as the three channel polarimetric image with S_0 , S_1 and S_2 as the three channels.

Several attempts have been made to address the polarimetric thermal-visible face recognition problem [26, 27, 36]. For instance, Riggan *et al.* [27] proposed a two-step procedure (visible feature estimation and visible image reconstruction) to solve this cross-modal matching problem. Zhang *et al.* [36] proposed an end-to-end generative adversarial network by fusing the different Stokes images as a multi-channel input to synthesize the visible image given the corresponding polarimetric signatures. Recently, Riggan *et al.* [26] developed a global and local region-based technique to improve the discriminative quality of the synthesized visible imagery. Though these methods are able to synthesize photo-realistic visible face images to some extent, the synthesized results in [36, 25, 26] are still far from optimal and they tend to lose some semantic attribute information such as mouth open, mustache, etc. Such reconstructions may degrade the performance of thermal to visible face verification.

In this paper, we take a different approach to the problem of thermal to visible matching. Figure 2 compares the traditional cross-modal verification problem with that of the proposed attribute-preserved cross-modal verification approach. Given a visible and thermal pair, the traditional approach first extracts some features from these images and then verifies the identity based on the extracted features [14] (see Figure 2(b)). In contrast, we propose a novel framework in which we make use of the attributes extracted from the visible image to synthesize the attribute-preserved visible image from the input thermal image for matching (see Figure 2(b)). In particular, a pre-trained VGG-Face model [19] is used to extract the attributes from the visible image. Then, a novel Attribute Preserved Generative Adversarial Network (AP-GAN) is proposed to synthesize the visible image from the thermal image guided by the extracted attributes. Finally, a deep network is used to extract features from the synthesized and the input visible images for verification.

The proposed AP-GAN model is inspired by the recent image generation from attributes/text works [25, 36, 3]. The AP-GAN consists of two parts: (i) a multimodal compact

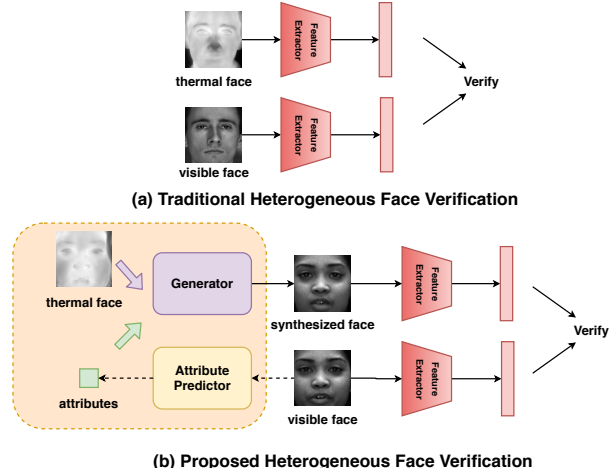
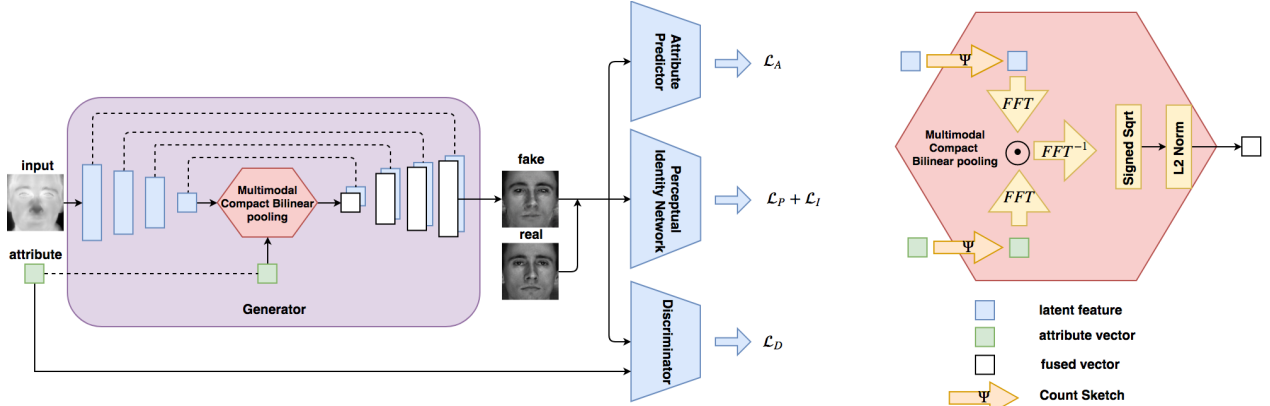


Figure 2. (a) Traditional heterogeneous face verification approaches use the features directly extracted from different modalities for verification [8, 14, 11, 34]. (b) The proposed heterogeneous face verification approach uses a thermal face and semantic attributes to synthesize a visible face. Finally, deep features extracted from the synthesized and visible faces are used for verification.

bilinear (MCB) pooling-based generator [4, 5], and (ii) a triplet-pair discriminator. The generator fuses the extracted attribute vector with the image feature vector in the latent space. On the other hand, the discriminator uses triplet pairs (real image/true attributes, fake image/true attributes, real image/wrong attributes) to not only discriminate between real and fake images but also to discriminate between the image and the attributes. In order to generate high-quality and attribute-preserved images, the generator is optimized by a multi-purpose objective function consisting of adversarial loss [6], L_1 loss, perceptual loss [12], identity loss [36] and attribute preserving loss. The entire AP-GAN framework is shown in Figure 3.

To summarize, the following are our main contributions:

- A novel thermal-visible face verification framework is proposed in which AP-GAN is developed for synthesizing visible faces from thermal (conventional or polarimetric) images using facial attributes.
- A novel MCB pooling [4, 5] based generator is proposed to fuse the given attributes with the image features.
- A novel triplet-pair discriminator is proposed, where the discriminator [25] not only learns to discriminate between real/fake images but also to discriminate between the image and the corresponding semantic attributes.
- Extensive experiments are conducted on the ARL Facial Database [9] and comparisons are performed



(a) AP-GAN Framework

(b) Multimodal Compact Bilinear (MCB) pooling

Figure 3. (a) A U-net based generator with MCB pooling is proposed to fuse the semantic attribute information with the image feature in the latent space. A triplet-pair is adopted for the discriminator in order to discriminate fake/real images as well as the corresponding semantic attributes. In order to generate high-quality and attribute-preserving images, a multi-purpose loss is optimized for training the network. (b) The architecture of MCB. Here, FFT indicates the Fast Fourier Transform and FFT^{-1} indicates the inverse FFT.

against several recent state-of-the-art approaches. Furthermore, an ablation study is conducted to demonstrate the improvements obtained by including semantic attribute information for synthesis.

2. Related Work

In this section, we review some related works on thermal to visible face synthesis and recognition.

2.1. Traditional Thermal-Visible Face Recognition

As described in Figure 2, traditional thermal to visible face verification methods first extract features from the visible and thermal images and then verify the identity based on the extracted features. Both hand-crafted and learned features have been investigated in the literature. Hu *et al.* [8] proposed a partial least squares (PLS) regression-based approach for cross-modal matching. Klare *et al.* [15] developed a generic framework for heterogeneous face recognition based on kernel prototype nonlinear similarities. Another multiple texture descriptor fusion-based method was proposed by Bourlai *et al.* in [34] for cross-modal face recognition. In [11] PLS-based discriminant analysis approaches were used to correlate the thermal face signatures to the visible face signatures. Some of the other visible to thermal cross-modal matching methods include [7, 30, 32].

2.2. Synthesis-based Thermal-Visible Face Verification

Unlike the above mentioned traditional methods, synthesis-based thermal to visible face verification algorithms leverage the synthesized visible faces for verification. Due to the success of CNNs and recently introduced generative adversarial networks (GANs) in synthesizing re-

alistic images, various deep learning-based approaches have been proposed in the literature for thermal to visible face synthesis [26, 36, 39, 27]. For example, Riggan *et al.* [27] proposed a two-step procedure (visible feature estimation and visible image reconstruction) to solve the thermal-visible verification problem. Zhang *et al.* [36] proposed an end-to-end GAN-based approach for synthesizing photo-realistic visible face images from their corresponding polarimetric images. Recently Riggan *et al.* [26] proposed a new synthesis method to enhance the discriminative quality of generated visible face images by leveraging both global and local facial regions.

3. Proposed Method

In this section, we discuss details of the proposed AP-GAN method. In particular, we discuss the proposed attribute predictor, generator and discriminator networks as well as the loss function used to train the network.

3.1. Attribute Predictor

To efficiently extract attributes from a given visible face, an attribute predictor is fine-tuned based on the VGG-Face network [19] using the ten annotated attributes. This network is trained separately from AP-GAN. The fine-tuned network is used in both obtaining the visible face attributes and for capturing the attribute loss.

3.2. Generator

A U-net structure [28] is used as the building block for the generator since it is able to better capture large receptive field and also able to efficiently address the vanishing gradient problem. In addition, to effectively combine the extra facial attribute information into the building block, we fuse

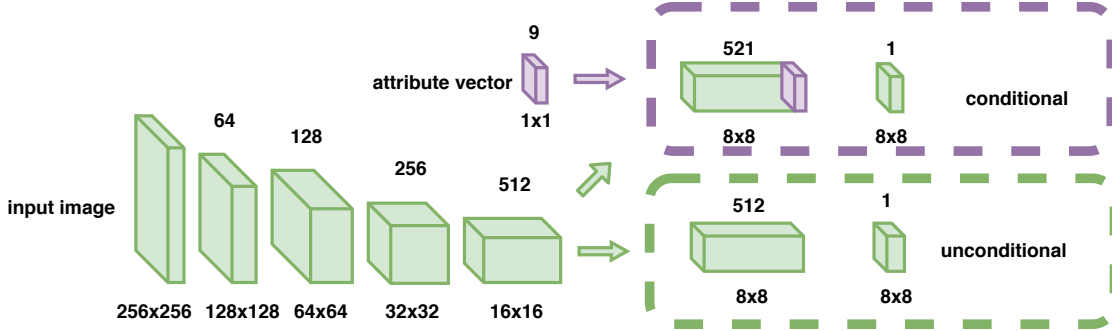


Figure 4. An overview of the triplet-pair discriminator. The triplet-pair discriminator is composed of a conditional and an unconditional streams. The unconditional stream aims to discriminate the fake and real images. The conditional stream aims to discriminate between the image and the corresponding attributes.

the attribute vector and the image feature in the latent space [25, 36, 3]. Note that the attributes are extracted from the given visible face using the fine-tuned model as discussed above. The architecture corresponding to the generator is shown in Figure 3(a).

In our experiments, we observe that simple concatenation of the two vectors (encoded image vector and attribute vector) does not work well. One possible reason is that both vectors are significantly different in terms of their dimensionality. Thus, we adopt the well-known MCB pooling method [4, 5] to overcome this issue. Instead of simple concatenation, MCB leverages the following two techniques: bilinear pooling and sketch count. Bilinear pooling is the outer-product and linearization of two vectors, where all elements of both vectors are interacting with each other in a multiplicative way. In order to overcome the high-dimension computation of bilinear pooling, Pham *et al.* [22] implemented the count sketch of the outer product of two vectors, which involves the Fast Fourier Transform (FFT) and inverse Fast Fourier Transform (FFT^{-1}). The architecture of MCB module is shown in Figure 3(b). The generator network we use in this paper can be described as follows:

CL(64)-CBL(128)-CBL(256)-CBL(512)-CBL(512)-CBL(512)-CBL(512)-CBL(512)-MCB(512)-DBR(512)-DBR(512)-DBR(512)-DBR(256)-DBR(128)-DBR(64)-DT(3),

where C stands for the convolutional layer (stride 2, kernel-size 4, and padding-size 1), L stands for Leaky Relu layer (negative_slope=0.02), B stands for the batch-normalization layer, MCB indicates the Multimodal Compact Bilinear module [4, 5], D stands for the deconvolutional layer (stride 2, kernel-size 4 and padding-size 1), R is the RuLU layer, and T is the Tanh function layer. All the numbers in parenthesis indicate the channel number of the output feature maps.

3.3. Discriminator

Motivated by the work [10], a patch-based discriminator D is leveraged in the proposed method and it is trained iteratively with G . As discussed above, the discriminator D not only aims to discriminate between real/fake images but also to discriminate between the image and the corresponding attributes. Similar to the discriminator in [25, 37], a triplet pair is given to the discriminator: real-image/true-attributes (Real), synthesized-image/true-attributes (Fake), real-image/wrong-attributes (Fake). Given an input image X , and attribute vector Y , the overall objective function for training D is as follows:

$$\begin{aligned} \mathcal{L}_{D_{uncond}} &= -\frac{1}{2}E_{X \sim P_{data}(X)}[\log D(X)] + \\ &\quad -\frac{1}{2}E_{G(X) \sim P_{G(X)}}[\log(1 - D(G(X)))], \\ \mathcal{L}_{D_{cond}} &= -\frac{1}{2}E_{X \sim P_{data}(X,A)}[\log D(X)] + \\ &\quad -\frac{1}{2}E_{G(X) \sim P_{G(X)}}[\log(1 - D(G(X), A))], \\ \mathcal{L}_D &= \mathcal{L}_{D_{uncond}} + \mathcal{L}_{D_{cond}}, \end{aligned} \quad (1)$$

where the unconditional loss $\mathcal{L}_{D_{uncond}}$ is to discriminate between real and synthesized samples. This information is back-propagated to G to make sure the generated samples are as realistic as possible. In addition, the conditional loss $\mathcal{L}_{D_{cond}}$ is added to discriminate whether the given image matches the attributes. This information is back-propagated to G so that it generates samples that are attribute preserving.

The architecture corresponding to the discriminator is shown in Figure 4. It consists of 6 convolutional blocks for both conditional and unconditional streams. Details of these convolutional blocks are as follows:

NCL(64)-NCBL(128)-NCBL(256)-NCBL(512)-CBL(512)-CS(1),

where N stands for the Gaussian noise layer used to improve the training stability, with zero-mean and standard

derivation of 0.01. S stands for the sigmoid activation layer. Note that the only difference between the unconditional and conditional stream is the concatenation of the attribute vector at the fifth convolutional block.

3.4. Object Function

The generator is optimized by minimizing the following loss

$$\mathcal{L}_{AP-GAN} = \mathcal{L}_G + \mathcal{L}_A + \lambda_P \mathcal{L}_P + \lambda_I \mathcal{L}_I + \lambda_1 \mathcal{L}_1, \quad (2)$$

where \mathcal{L}_G is the adversarial loss for generator G , \mathcal{L}_P is the perceptual loss, \mathcal{L}_I is the identity loss, \mathcal{L}_A is the attribute loss, \mathcal{L}_1 is the loss based on the L_1 -norm between the target and the reconstructed image, $\lambda_P, \lambda_I, \lambda_A, \lambda_1$ are weights respectively for perceptual loss, identity loss, attribute loss and L_1 loss.

3.4.1 Adversarial Loss

Similar to the discriminator D , the adversarial loss for the generator G consists of both conditional and unconditional parts as defined below

$$\mathcal{L}_G = \frac{1}{2} E_{G(X) \sim P_{G(X)}} [\log(1 - D(G(X)))] + \frac{1}{2} E_{G(X) \sim P_{G(X)}} [\log(1 - D(G(X), A))]. \quad (3)$$

The generator G therefore jointly approximates the image distribution conditioned (or unconditioned) on the attributes A .

3.4.2 Perceptual and Identity Loss

Perceptual loss was introduced by Johnson *et al.* [12] for style transfer and super-resolution. It has been observed that the perceptual loss produces visually pleasing results than L_1 or L_2 loss. The perceptual and identity losses are defined as follows

$$\mathcal{L}_{P,I} = \sum_{c=1}^C \sum_{w=1}^W \sum_{h=1}^H \|V(G(X))^{c,w,h} - V(Y)^{c,w,h}\|_2, \quad (4)$$

where V represents a non-linear CNN feature. VGG-16 [33] is used to extract features in this work. C, W, H are the dimensions of features from a certain level of the VGG-16, which are different for perceptual and identity losses.

In addition, L_1 loss between the synthesized image $G(X)$ and the real image Y is used to capture the low-frequency information, which is defined as follows

$$\mathcal{L}_1 = \|G(X) - Y\|_1. \quad (5)$$

Table 1. The facial attributes used in this work.

attributes	Arched_Eyebrows, Big_Lips, Big_Nose, Bushy_Eyebrows, Male, Mustache, Narrow_Eyes, No_Beard, Mouth_Slightly_Open, Young
------------	--

3.4.3 Attribute Loss

Inspired by the perceptual loss, we define an attribute preserving loss, which measures the error between the attributes of the synthesized image and the real image. To make sure the pre-trained model captures the facial attribute information, we fine-tune the pretrained VGG-Face network on the attribute dataset and regard the fine-tuned attribute classifier as the pre-trained model for the attribute preserving loss. Similar to the perceptual loss, the \mathcal{L}_A is defined as follows

$$\mathcal{L}_A = \|Q(G(X)) - Q(Y)\|_2, \quad (6)$$

where Q is the fine-tuned attribute predictor network and N is the total number of output neurons. By feeding such an attribute information into the generator during training, the generator G is able to learn semantic information corresponding to the face.

3.5. Implementation

The entire network is trained in Pytorch on a single Nvidia Titan-X GPU. During the AP-GAN training, the L_1 , perceptual and identity loss parameters are chosen as $\lambda_1 = 10$, $\lambda_P = 2.5$, $\lambda_I = 0.5$, respectively. The ADAM [13] is implemented as the optimization algorithm with parameter *betas* = (0.5, 0.999) and batch size is chosen as 3. The total epochs are 200. For the first 100 epochs, we fix the learning rate as 0.0002 and for the remaining 100 epochs, the learning rate was decreased by 1/100 after each epoch. The feature maps for the perceptual and the identity loss are from the relu1-1 and the relu2-2 layers, respectively. In order to fine-tune the attribute predictor network, we manually annotate images with the attributes tabulated in Table 1.

4. Experimental Results

The proposed method is evaluated on the ARL Multimodal Face Database [9] which consists of polarimetric (i.e. Stokes image) and visible images from 60 subjects. Similar to the protocol discussed in [27], we only use the images from Range 1 and their corresponding attributes are obtained from fine-tuned attribute predictor network. In particular, Range 1 images from 30 subjects and the corresponding attributes are used for training. The remaining 30 subjects' data are used for evaluation. We repeat this process 5 times and report the average results.

We evaluate the face verification performance of proposed method compared with several recent works [36, 26,

Table 2. Verification performance comparisons among the baseline methods, state-of-the-art methods, and the proposed AP-GAN method for both polarimetric thermal (Polar) and conventional thermal (S0) cases.

Method	AUC(Polar)	AUC(S0)	EER(Polar)	EER(S0)
Raw	50.35%	58.64%	48.96%	43.96%
Mahendran <i>et al.</i> [17]	58.38%	59.25%	44.56%	43.56%
Riggan <i>et al.</i> [27]	75.83%	68.52%	33.20%	34.36%
Zhang <i>et al.</i> [36]	79.90%	79.30%	25.17%	27.34%
Riggan <i>et al.</i> [26]	85.43%	82.49%	21.46%	26.25%
AP-GAN (our)	88.93% ± 1.54%	84.16% ± 1.54%	19.02% ± 1.69%	23.90% ± 1.52%
AP-GAN (GT)	91.28% ± 1.68%	86.08% ± 2.68%	17.58% ± 2.36%	23.13% ± 3.02%

10]. Moreover, the performance is evaluated on the FC-7 layer of the pre-trained VGG-Face model [19] using the receiver operating characteristic (ROC) curve, Area Under the Curve (AUC) and Equal Error Rate (EER) measures. To summarize, the proposed method is evaluated on the following two protocols:

- (a) Conventional thermal (S0) to Visible (Vis).
- (b) Polarimetric thermal (Polar) to Visible (Vis).

4.1. Preprocessing

In addition to the standard preprocessing in [9], two more pre-processing steps are used for the proposed method. First, the faces in visible images are detected by MTCNN [38]. Then, a standard central crop method is used to crop the detected faces. Since the MTCNN is implementable on the visible images only, we use the same detected rectangle coordinations to crop the S0, S1, S2 images. After preprocessing, all the images are scaled to be 256×256 and saved as 16-bit PNG files.

4.2. Comparison with state-of-the-art Methods

We evaluate and compare the performance of the proposed method with that of recent state-of-the-art methods [36, 17, 27, 26]. In addition to our method, we also conduct experiments with a baseline method 'AP-GAN(GT)' where we use the ground truth attributes in our method rather than automatically predicting them using the proposed attribute predictor. This baseline will clearly determine how effective the proposed attribute predictor is in determining the attributes from unconstrained visible faces.

Figure 5 shows the evaluation performance for two different experimental settings, S0 and Polar separately. Compared with other state-of-the-art methods in Figure 5, the proposed method performs better with a larger AUC and lower EER scores. In addition, it can be observed that the performance corresponding to the Polar modality is better than the S0 modality, which also demonstrates the advantage of using the polarimetric thermal images than the conventional thermal images. The quantitative comparisons, as shown in the Table 2, also demonstrate the effectiveness of proposed method.

In addition to the quantitative results, we also show some visual comparisons in Figure 6. The first row in Figure 6 shows one synthesized sample using S0. The second row shows the same synthesized sample using Polar. It can be observed that results of Riggan *et al.* [27] do capture the overall face structure but it tends to lose some details on the skin. Results of Mahendran *et al.* [17] poor compared to [27]. Results of Zhang *et al.* [36] are more photo-realistic but tend to lose some attribute information. The proposed AP-GAN not only generates photo-realistic images but also preserves attributes on the reconstructed images.

4.3. Ablation Study

In order to demonstrate the effectiveness of different modules in the proposed method, we conduct the following ablation studies: (1) Polar to Visible estimation with only \mathcal{L}_1 loss, (2) Polar to Visible estimation with L_1 and adversarial loss \mathcal{L}_G , (3) Polar to Visible estimation with \mathcal{L}_1 , \mathcal{L}_G , and perceptual and identity loss $\mathcal{L}_P, \mathcal{L}_I$, (4) Polar to Visible estimation with all the losses \mathcal{L}_{AP-GAN} as defined in Eq. (2). Figure 8 shows the ROC curves corresponding to each experimental setting. All the experiments in the ablation study are evaluated from one experimental split of the Polar modality. From this figure, we can observe that using all the losses together as \mathcal{L}_{AP-GAN} , we obtain the best performance. Compared to the results of \mathcal{L}_{AP-GAN} and $\mathcal{L}_1 + \mathcal{L}_G + \mathcal{L}_P + \mathcal{L}_I$, we can clearly see the improvements obtained by fusing the semantic attribute information with the image feature in the latent space.

Besides the ROC curves, we also show the visual results for each experimental setting in Figure 7. Given the input Polar image, the synthesized results from different experimental setting are shown in Figure 7. It can be observed that \mathcal{L}_1 captures the low-frequency features of images very well. $\mathcal{L}_1 + \mathcal{L}_G$ can capture both low-frequency and high-frequency features in the image. However, it adversely introduced distortions and artifacts in the synthesized image. In addition, optimizing $\mathcal{L}_P + \mathcal{L}_I$ suppresses these distortions to some extent. Finally, fusing attributes into the previous loss can not only improving the performance but also preserves facial attributes, like the mustache as shown in the




Figure 5. The ROC curve comparison with several state-of-the-art methods: (a) The performance on Polar-Visible verification. (b) The performance on S0-Visible verification.




Figure 6. The synthesized samples from different methods: Riggan et al. [27], Mahendran et al. [17], Zhang et al. [36], AP-GAN (ours), ground truth. The first row results correspond to the S0 image, and the second row results correspond to the Polar image.

red circle.

4.4. Attribute Manipulation Result

Instead of visually and quantitatively showing the performance of AP-GAN on face verification, we also show results when the attributes are manipulated.

Given a certain thermal image, by manipulating its corresponding attributes, we obtain some interesting synthesis results as shown in Figure 9. In the first row of Figure 9, the *mouth_open* attribute value was changed from -1 to 1 while the other attribute values were fixed. As can be seen from the generated figure, the synthesized image shows a slightly open mouth. In the second row, we show the results for changing the attribute value corresponding to mustache from 1 to -1 . The generated results clearly capture the attribute change as shown with a red circle.

5. Conclusion

We propose a novel Attribute Preserving Generative Adversarial Network (AP-GAN) structure for polarimetric-visible face verification via synthesizing photo realistic visible face images from the corresponding thermal (polarimetric or conventional) images with extracted attributes. Rather than use only image-level information for synthesis and verification, we take a different approach in which semantic facial attribute information is also fused during training and testing. Quantitative and visual experiments evaluated on a real thermal-visible dataset demonstrate that the proposed method achieves state-of-the-art performance compared with other existing methods. In addition, an ablation study is developed to demonstrate the improvements obtained by different combination of loss functions.



Input \mathcal{L}_1 $\mathcal{L}_1 + \mathcal{L}_G$ $\mathcal{L}_1 + \mathcal{L}_G + \mathcal{L}_P + \mathcal{L}_I$ \mathcal{L}_{AP-GAN} Ground Truth

Figure 7. The visual results of the ablation study for different experimental settings. Given input Polar image, synthesized results using different combination of losses are shown successively from left to right columns as \mathcal{L}_1 ; $\mathcal{L}_1 + \mathcal{L}_G$; $\mathcal{L}_1 + \mathcal{L}_G + \mathcal{L}_P + \mathcal{L}_I$; \mathcal{L}_{AP-GAN} . The ground truth images are on the most right side.




Figure 8. The ROC curves corresponding to the ablation study.

Acknowledgement

This work was supported by an ARO grant W911NF-16-1-0126.

References

- [1] T. Bourlai, N. Kalka, A. Ross, B. Cukic, and L. Hornak. Cross-spectral face verification in the short wave infrared (swir) band. In *Pattern Recognition (ICPR), 2010 20th International Conference on*, pages 1343–1347. IEEE, 2010.
- [2] J. C. Chen, V. M. Patel, and R. Chellappa. Unconstrained face verification using deep cnn features. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9, March 2016.
- [3] X. Di and V. M. Patel. Face synthesis from visual attributes via sketch using conditional vaes and gans. *arXiv preprint arXiv:1801.00077*, 2017.
- [4] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *Proceed-*




Figure 9. Attribute manipulation results. Given the input thermal image and manipulated attributes, the synthesized visible image shows the changes caused by these attributes. The first row shows the visual difference on mouth_open attribute changing from -1 to 1 (the other attributes are fixed). The second row shows the visual difference on mustache attribute changing from 1 to -1 (the other attributes are fixed). The attribute-based manipulation and original-synthesis column show the comparison with and without attribute change, respectively.

ings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016, 2016.

- [5] Y. Gao, O. Beijbom, N. Zhang, and T. Darrell. Compact bilinear pooling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 317–326, 2016.
- [6] 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*, pages 2672–2680, 2014.
- [7] K. P. Gurton, A. J. Yuffa, and G. W. Videen. Enhanced facial recognition for thermal imagery using polarimetric imaging. *Opt. Lett.*, 39(13):3857–3859, Jul 2014.
- [8] S. Hu, J. Choi, A. L. Chan, and W. R. Schwartz. Thermal-

- to-visible face recognition using partial least squares. *JOSA A*, 32(3):431–442, 2015.
- [9] S. Hu, N. J. Short, B. S. Riggan, C. Gordon, K. P. Gurton, M. Thielke, P. Gurrarn, and A. L. Chan. A polarimetric thermal database for face recognition research. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 119–126, 2016.
- [10] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *CVPR*, 2017.
- [11] J. Choi, S. Hu, S. S. Young, L. S. Davis. Thermal to visible face recognition. In *Proc.SPIE*, pages 8371 – 8371 – 10, 2012.
- [12] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711. Springer, 2016.
- [13] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [14] B. Klare and A. K. Jain. Heterogeneous face recognition: Matching nir to visible light images. In *Pattern Recognition (ICPR), 2010 20th International Conference on*, pages 1513–1516. IEEE, 2010.
- [15] B. F. Klare and A. K. Jain. Heterogeneous face recognition using kernel prototype similarities. *IEEE transactions on pattern analysis and machine intelligence*, 35(6):1410–1422, 2013.
- [16] 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 *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6807–6816. IEEE, 2017.
- [17] A. Mahendran and A. Vedaldi. Understanding deep image representations by inverting them. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 5188–5196, 2015.
- [18] F. Nicolo and N. A. Schmid. Long range cross-spectral face recognition: matching swir against visible light images. *IEEE Transactions on Information Forensics and Security*, 7(6):1717–1726, 2012.
- [19] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2015.
- [20] X. Peng, R. S. Feris, X. Wang, and D. N. Metaxas. A recurrent encoder-decoder network for sequential face alignment. In *European conference on computer vision*, pages 38–56. Springer, Cham, 2016.
- [21] X. Peng, R. S. Feris, X. Wang, and D. N. Metaxas. Red-net: A recurrent encoder–decoder network for video-based face alignment. *International Journal of Computer Vision*, pages 1–17, 2018.
- [22] N. Pham and R. Pagh. Fast and scalable polynomial kernels via explicit feature maps. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 239–247. ACM, 2013.
- [23] R. Ranjan, V. M. Patel, and R. Chellappa. Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2017.
- [24] R. Ranjan, S. Sankaranarayanan, A. Bansal, N. Bodla, J. C. Chen, V. M. Patel, C. D. Castillo, and R. Chellappa. Deep learning for understanding faces: Machines may be just as good, or better, than humans. *IEEE Signal Processing Magazine*, 35(1):66–83, Jan 2018.
- [25] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text to image synthesis. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML’16*, pages 1060–1069. JMLR.org, 2016.
- [26] B. S. Riggan, N. J. Short, and S. Hu. Thermal to visible synthesis of face images using multiple regions. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2018.
- [27] B. S. Riggan, N. J. Short, S. Hu, and H. Kwon. Estimation of visible spectrum faces from polarimetric thermal faces. In *Biometrics Theory, Applications and Systems (BTAS), 2016 IEEE 8th International Conference on*, pages 1–7. IEEE, 2016.
- [28] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [29] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [30] N. Short, S. Hu, P. Gurrarn, and K. Gurton. Exploiting polarization-state information for cross-spectrum face recognition. In *Biometrics Theory, Applications and Systems (BTAS), 2015 IEEE 7th International Conference on*, pages 1–6. IEEE, 2015.
- [31] N. Short, S. Hu, P. Gurrarn, K. Gurton, and A. Chan. Improving cross-modal face recognition using polarimetric imaging. *Optics letters*, 40(6):882–885, 2015.
- [32] N. Short, S. Hu, P. Gurrarn, K. Gurton, and A. Chan. Improving cross-modal face recognition using polarimetric imaging. *Opt. Lett.*, 40(6):882–885, Mar 2015.
- [33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [34] T. Bourlai, A. Ross, C. Chen, L. Hornak. A study on using mid-wave infrared images for face recognition. volume 8371, pages 8371 – 8371 – 13, 2012.
- [35] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision*, pages 499–515. Springer, 2016.
- [36] H. Zhang, V. M. Patel, B. S. Riggan, and S. Hu. Generative adversarial network-based synthesis of visible faces from polarimetric thermal faces. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 100–107, Oct 2017.

- [37] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. N. Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *CoRR*, abs/1710.10916, 2017.
- [38] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, Oct 2016.
- [39] T. Zhang, A. Wiliem, S. Yang, and B. C. Lovell. Tv-gan: Generative adversarial network based thermal to visible face recognition. *arXiv preprint arXiv:1712.02514*, 2017.