Using Photorealistic Face Synthesis and Domain Adaptation to Improve Facial Expression Analysis

Behzad Bozorgtabar[1], Mohammad Saeed Rad[1],
Hazım Kemal Ekenel[2], Jean-Philippe Thiran[1]

[1]École Polytechnique Fédérale de Lausanne, Switzerland, [2]Istanbul Technical University, Istanbul, Turkey

Abstract— **Cross-domain synthesizing realistic faces to learn deep models has attracted increasing attention for facial expression analysis as it helps to improve the performance of expression recognition accuracy despite having small number of real training images. However, learning from synthetic face images can be problematic due to the distribution discrepancy between low-quality synthetic images and real face images and may not achieve the desired performance when the learned model applies to real world scenarios. To this end, we propose a new attribute guided face image synthesis to perform a translation between multiple image domains using a single model. In addition, we adopt the proposed model to learn from synthetic faces by matching the feature distributions between different domains while preserving each domain's characteristics. We evaluate the effectiveness of the proposed approach on several face datasets on generating realistic face images. We demonstrate that the expression recognition performance can be enhanced by benefiting from our face synthesis model. Moreover, we also conduct experiments on a near-infrared dataset containing facial expression videos of drivers to assess the performance using in-the-wild data for driver emotion recognition.**

I. INTRODUCTION

Face image synthesis has received increasing attentions as it has practical applications in human-computer interactions, facial animation, face recognition, and facial expression recognition. The most notable solution in image synthesis was the incredible breakthroughs in generative models. In particular, Generative Adversarial Network (GAN) [7] variants have achieved state-of-the-art results for image-to-image translation task. These GAN models could discover relations between two visual domain using paired [12] or unpaired data [15], [41] during training process. In addition, most existing GAN models [29], [41] are proposed to synthesize images of a single attribute, which make their training inefficient in the case of having multiple attributes, since for each attribute a separate model is needed.

In this paper, we pursue several objectives; synthesizing realistic faces by controlling the facial attributes of interest (e.g. face expression), preserving the facial identity after manipulation, and to investigate learning from synthetic facial images for improving expression recognition accuracy (see Fig. 1). Our objective is to use a single model to synthesize face photos with a desired attribute and translate an input image from one domain into multiple domains without having matching image pairs. Our proposed method is based on encoder-decoder structure, using the image latent

(a) Photorealistic face synthesis
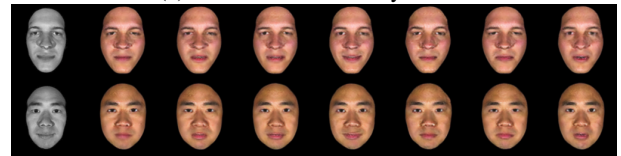
(b) Realism refinement using domain adaptation

Fig. 1: Attribute guided face image generation using our method. (a) The input *neutral* faces are fed into our model to exhibit specified attributes. **Left to right**: input *neutral* face and eight face attributes including *angry*, *contemptuous*, *disgusted*, *fearful*, *happiness*, *neutral*, *sadness* and *surprised*, respectively. (b) Our model improves the realism of synthetic face (**left**), while preserving the identity and face pose information during the realism refinement for a specified attribute.

representation, where we model the shared latent representation across image domains. Therefore, during inference step, by changing face attributes, we can generate plausible face images owing attribute of interest. We also introduce bidirectional loss for the latent representation, which can resolve generator mode collapse to ensure diverse outputs can be produced depending on input attribute. Our paper makes the following contributions:

1) We extend the previous work [4] and show that how the proposed approach can be used to generate photo-realistic facial images using synthetic face image and unlabeled real face images as the input. We compared our results with SimGAN method [31] in terms of expression recognition accuracy to see improvement in the realism of generated faces using video data recorded in real world conditions;

2) Compared to other variants of GAN models [41], [28], we demonstrate that the learnt representation achieves high-quality image synthesis results and preserves a

certain expression that contribute to improve the performance of expression recognition accuracy to focus on the data augmentation process;

3) Lastly, unlike existing methods in face expression synthesis, which have only been validated on the face datasets captured in a lab-controlled environment, we tested our approach on the videos in the wild dataset, which contains arbitrary face poses, illumination and self-occlusions.

II. RELATED WORK

GAN based models [7] have achieved impressive results in many image synthesis applications, including image super-resolution [21], image-to-image translation (pix2pix) [12] and face image editing [25], [35]. We summarize contributions of few important related works using GANs in the following subsections:

A. Image-to-Image Translation

Many of existing image-to-image translation methods e.g. [12], [31] formulated GANs in the supervised setting, where example image pairs are available. However, collecting paired training data can be difficult. On the other side, there are other GAN based methods, which do not require matching pairs of samples. For example, CycleGAN [41] is capable to learn transformations from source to target domain without one-to-one mapping between two domain's training data. However, these GAN based methods could only train one specific model for each pair of image domains. Unlike the aforementioned approaches, we use a single model to synthesize multiple photo-realistic images, each having specific attribute. Recently, to manipulate attributes of image during image synthesis, conditional information, such as image labels, is used. As examples, IcGAN [28] and StarGAN [5] proposed image editing using AC-GAN [26] with conditional information. However, we use domain adaptation by adding the realism to the simulated faces and there is no such a solution in these methods. Similar to [28], Fader Networks [19] proposed image synthesis model without needing to apply a GAN to the decoder output. However, these methods impose constraints on image latent space to enforce it to be independent from the attributes of interest, which may result in loss of information in generating attribute guided images.

B. Face Attribute Manipulation and Face Generation

Li et al. [22] proposed a Deep convolutional network model for Identity-Aware Transfer (DIAT) of the facial attributes. The work [29] and [14] proposed to edit only single facial attribute. Lu et al. [24] proposed attribute-guided face generation to translate low-resolution face images to high-resolution face images. Zhang et al. [38] introduced the spatial attention mechanism into GAN framework, to only alter the attribute-specific face region and keep the rest unchanged. Huang et al. [11] proposed a Two-Pathway Generative Adversarial Network (TP-GAN) for photorealistic face synthesis by simultaneously considering local face details and global structures.

C. Expression Transfer and Face Frontalization

Zhang et al. [37] proposed a method by disentangling the attributes (expression and pose) for simultaneous pose-invariant facial expression recognition and face images synthesis. Instead, we seek to learn attribute-invariant information in the latent space by imposing auxiliary classifier to classify the generated images. Zhou et al. [40] introduced a conditional difference adversarial autoencoder (CDAAE) to use emotion states as a conditional attribute for face generation. Lai et al. [18] proposed a multi-task GAN-based network that learns to synthesize the frontal face images from profile face images. However, they require paired training data of frontal and profile faces. Instead, we seek to add realism to the synthetic frontal face images without requiring real frontal face images during training. Our method could generate realistic frontal face images using synthetic faces and real faces with arbitrary poses as input.

III. METHODS

We first introduce our proposed attribute guided face synthesis model in Section III-A. Then, we explore learning from simulated face images through adversarial training in Section III-B. Finally, we discuss our implementation details and experimental results in Sections IV and V, respectively.

A. Attribute-guided face image synthesis

As the input of our framework, we have input face image, and the attributes to be edited (e.g. facial expression) and side image, which provides additional information to guide photo-realistic face synthesis. Let $\mathcal{X}$ and $\mathcal{S}$ denote original image and side conditional image domains, respectively and $\mathcal{Y}$ set of possible facial attributes. As the training set, we have m triple inputs $(x_i \in \mathcal{X}, s_i \in \mathcal{S}, y_i \in \mathcal{Y})$, where x_i and y_i are the i^{th} input face image and binary attribute, respectively and s_i represents the i^{th} conditional side image. Then, for any categorical attribute vector y from the set of possible facial attributes $\mathcal{Y}$, the objective is to train a model that will generate photo-realistic version (x' or s') of the inputs (x and s) from image domains $\mathcal{X}$ and $\mathcal{S}$ with desired attributes y.

Our model is based on the encoder-decoder architecture with domain adversarial training. As the input to our expression synthesis model (see Fig. 2), we propose to incorporate individual-specific facial shape model (face landmark heatmap) as the side conditional information s in addition to the original input image x. The face landmark heatmap contains 2D Gaussians centered at the landmarks locations, which are then simply concatenated channel-wise with the input image. Our goal is then to train a single generator G with the encoder G_{enc} – decoder G_{dec} networks to translate the input pair (x, s) from source domains into their corresponding output images (x', s') in the target domain conditioned on the target domain attribute y and the inputs latent representation $G_{enc}(x, s)$, $G_{dec}(G_{enc}(x, s), y) \rightarrow x', s'$. The encoder $G_{enc} : (\mathcal{X}^{source}, \mathcal{S}^{source}) \rightarrow \mathbb{R}^{n \times \frac{h}{16} \times \frac{w}{16}}$ is a fully convolutional neural network with parameters θ_{enc} that encodes the input images into a low-dimensional feature

space $G_{enc}(x, s)$, where n, h, w are the number of the feature channels and the input images dimensions, respectively. The decoder $G_{dec} : \left(\mathbb{R}^{n \times \frac{h}{16} \times \frac{w}{16}}, \mathcal{Y}\right) \to \mathcal{X}^{target}, \mathcal{S}^{target}$ is the sub-pixel [30] convolutional neural network with parameters θ_{dec} that produce realistic images with target domain attribute y and given the latent representation $G_{enc}(x, s)$. The precise architectures of the neural networks are described in Section IV-A. During training, we randomly use a set of target domain attributes y to make the generator more flexible in synthesizing images. In the following, we introduce the objectives for the proposed model optimization.

1) Adversarial Loss: We propose to learn attribute-invariant information in the latent space representing the shared features of the images sampled for different attributes. It means if the original and target domains are semantically similar (e.g. facial images of different expressions), we expect the common features across domains to be captured by the same latent representation. Then, the decoder must use the target attribute to do image-to-image translation from the original domain to the target domain. However, this learning process is unsupervised as for each training image from the source domain, its counterpart image in the target domain with attribute y is unknown. Therefore, we propose to train an additional neural network called the discriminator D (with the parameters θ_{dis}) using an adversarial formulation to not only distinguish between real and fake generated images, but also to classify the image to its corresponding attribute categories. We use Wasserstein GAN [8] objective with a gradient penalty loss $\mathcal{L}_{gp}$ [2] formulated as below:

$$
\begin{aligned}
\mathcal{L}_{adv} = & \mathbb{E}_{x,s}\left[D_{src}(x, s)\right] \\
& - \mathbb{E}_{x,s,y}\left[D_{src}(G_{dec}(G_{enc}(x, s), y))\right] - \lambda_{gp}\mathcal{L}_{gp}(D_{src}),
\end{aligned}
\tag{1}
$$

The term $D_{src}(\cdot)$ denotes a probability distribution over image sources given by D. The hyper-parameter λ_{gp} is used to balance the GAN objective with the gradient penalty. A generator (encoder-decoder networks) used in our model has to play two roles: learning the attribute invariant representation for the input images and is trained to maximally fool the discriminator in a *min-max* game. On the other hand, the discriminator simultaneously seeks to identify the fake examples for each attribute.

2) Attribute Classification Loss: We use a classifier by returning additional output from the discriminator to perform an auxiliary task of classifying the synthesized and the real facial images into their respective attribute values. An attribute classification loss of real images $\mathcal{L}_{cls_r}$ to optimize the discriminator parameters θ_{dis} is defined as follow:

$$
\min_{\theta_{dis}} \mathcal{L}_{cls_r} = \mathbb{E}_{x,s,y'}\left[\ell_r(x, s, y')\right],
$$

$$
\begin{aligned}
\ell_r(x, s, y') = & \sum_{i=1}^{m} -y_i' \log D_{cls}(x, s) \\
& - (1 - y_i') \log(1 - D_{cls}(x, s)),
\end{aligned}
\tag{2}
$$

Here, y' denotes original attributes categories for the real images. ℓ_r is is the summation of binary cross-entropy losses

of all attributes. Besides, an attribute classification loss of fake images $\mathcal{L}_{cls_f}$ used to optimize the generator parameters $(\theta_{enc}, \theta_{dec})$, formulated as follow:

$$
\min_{\theta_{enc}, \theta_{dec}} \mathcal{L}_{cls_f} = \mathbb{E}_{x,s,y'}\left[\ell_f(x', s', y)\right],
$$

$$
\begin{aligned}
\ell_f(x', s', y) = & \sum_{i=1}^{m} -y_i \log D_{cls}(x', s') \\
& - (1 - y_i) \log(1 - D_{cls}(x', s')),
\end{aligned}
\tag{3}
$$

where x' and s' are the generated images and auxiliary outputs, which should correctly own the target domain attributes y. ℓ_f denotes summing up the cross-entropy losses of all fake images.

3) Identity Loss: Using the identity loss, we aim to preserve the attribute-excluding facial image details such as facial identity before and after image translation. We use a pixel-wise l_1 loss to enforce the facial details consistency and suppress the blurriness:

$$
\mathcal{L}_{id} = \mathbb{E}_{x,s,y'}\left[\|G_{dec}(G_{enc}(x, s), y') - x\|_1\right], \tag{4}
$$

4) Bidirectional Loss: Using adversarial loss alone usually leads to mode collapse, and the generator learns to ignore the attributes and changing y at the test time could not generate diverse facial images. This issue has been observed in various applications of conditional GANs [12], [6] and to our knowledge, there is still no proper approach to deal with this issue. To address this problem, we show that using the trained generator, images of different domains can be translated bidirectionally. We decompose this objective into two terms: a bidirectional loss for the image latent representation, and a bidirectional loss between synthesized images and original input images, respectively. This objective is formulated using l_1 loss as follow:

$$
\begin{aligned}
\mathcal{L}_{bi} = & \mathbb{E}_{x,s,y'}\left[\|x - \hat{x}\|_1 + \|s - \hat{s}\|_1\right] + \\
& \mathbb{E}_{x,s,y}\left[\|G_{enc}(x, s) - G_{enc}(x', s')\|_1\right], \\
& x', s' = G_{dec}(G_{enc}(x, s), y), \\
& \hat{x}, \hat{s} = G_{dec}(G_{enc}(x', s'), y'),
\end{aligned}
\tag{5}
$$

In the above equation, $\hat{x}$ and $\hat{s}$ denote the reconstructed original image and the side conditional image, respectively. Unlike [41], where only the cycle consistency losses are used at the image level, we additionally seek to minimize the reconstruction loss using latent representation.

5) Overall Objective: Finally, the generator G is trained with a linear combination of four loss terms: adversarial loss, attribute classification loss for the fake images, bidirectional loss, and identity loss. Meanwhile, the discriminator D is optimized using an adversarial loss and attribute classification loss for the real images:

$$
\begin{aligned}
\mathcal{L}_G &= \mathcal{L}_{adv} + \lambda_{bi}\mathcal{L}_{bi} + \lambda_{cls}\mathcal{L}_{cls_f} + \lambda_{id}\mathcal{L}_{id}, \\
\mathcal{L}_D &= -\mathcal{L}_{adv} + \lambda_{cls}\mathcal{L}_{cls_r},
\end{aligned}
\tag{6}
$$

where λ_{bi}, λ_{id} and λ_{cls} are hyper-parameters, which tune the importance of identity loss, bidirectional loss and attribute classification loss, respectively.
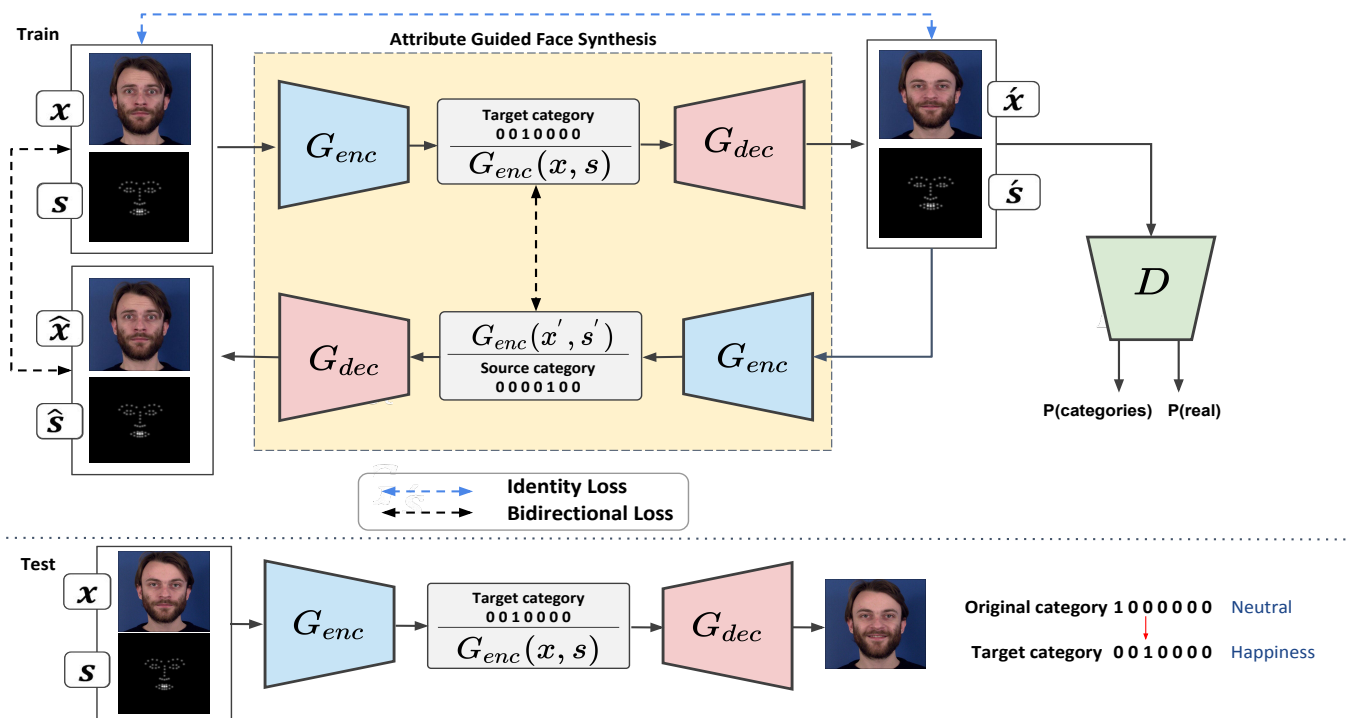
Fig. 2: Overview of our proposed method. Training process of the attribute-guided face synthesis, consisting of two components, an encoder-decoder generator G and a discriminator D (**top**). Face expression synthesis at the test time (**bottom**).

B. Realism Refinement Using Domain Adaptation

In an unconstrained face expression recognition (FER), accuracy will drop significantly for large pose variations. One possible solution would be using simulated faces rendered in frontal view. In particular, we utilize a 3D Morphable Model using bilinear face model [34] to construct a simulated frontal face image. Fig. 3 shows examples of simulated faces. However, learning from synthetic face images can be problematic due to a distribution discrepancy between real and synthetic images. Using proposed attribute guided face synthesis in Section III-A, the model takes simulated frontal face image x and real face image with arbitrary pose s as inputs, and generates photo-realistic version of the synthetic face x' during the realism refinement. With a side input image as condition, the model has enough information about the appearance of the desired face in advance and we transfer a texture from a given unlabeled real face image with arbitrary pose to a synthetic frontal face (see Fig. 4). Here, the discriminator's role is to discriminate the realism of single generated image using real profile face images. In addition, using the same discriminator, we can generate images exhibiting arbitrary attributes e.g., different expressions.

We compare the pose-normalized face attribute transfer results of our proposed method with SimGAN method [31] on the BU-3DFE dataset [36]. SimGAN method [31] considers learning from simulated and unsupervised images through adversarial training. Our method differs from SimGAN in following aspects: 1) we aim to synthesize photo-realistic frontal faces by preserving the face pose to address the

Fig. 3: Examples of face frontalization process [34] using facial landmarks. A simulated frontal view (**bottom row**) is generated from two side cameras (**top and middle rows**).

challenges in unconstrained face expression recognition, whereas SimGAN is designed for simpler scenarios e.g., eye image refinement. 2) Another shortcoming of this method would be to ignore categorical information, which limits its performance. In contrast, our proposed method overcomes this issue by introducing attribute classification loss into our objective function. For a fair comparison with SimGAN method, we add the attribute classification loss by modifying the SimGAN's discriminator, while keeping the rest of network unchanged. We achieve more visually pleasing results on test data compared to the SimGAN method (see Fig. 6). Our proposed method can preserve the face image content

while modifying only the attribute-related part of the images using the latent representation.
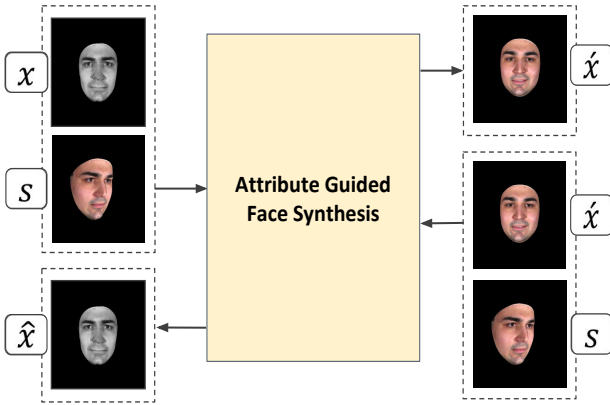

Fig. 4: Overview of our proposed realism refinement. As the inputs, takes a simulated frontal face image and unlabeled real image as input and generates photo-realistic version of the frontal face.

IV. IMPLEMENTATION DETAILS

All networks are trained using Adam optimizer [16] ($\beta_1 = 0.5, \beta_2 = 0.999$) and with a base learning rate of 0.0001. We linearly decay learning rate after the first 100 epochs. We use a simple data augmentation with only flipping the images horizontally. The input image size and the batch size are set to 128×128 and 8 for all experiments, respectively. We update the discriminator five times for each generator (encoder-decoder) update. The hyper-parameters in Eq. 6 and Eq. 1 are set as: $\lambda_{bi} = 10$ and $\lambda_{id} = 10$, $\lambda_{gp} = 10$ and $\lambda_{cls} = 1$, respectively. The whole model is implemented using PyTorch on a single NVIDIA GeForce GTX 1080.

A. Network Architecture

Tables I and II demonstrate the detailed network architectures of our proposed attribute-guided face image synthesis model. For the discriminator, we use PatchGAN [12] that penalizes structure at the scale of image patches. Regarding the generator's decoder, we use sub-pixel convolution instead of transposed convolution followed by instance normalization [3]. Our experiments verify that it works remarkably better than transposed convolution for the face image synthesis.

V. EXPERIMENTAL RESULTS

A. Datasets

Near IR Drivers' Video Dataset: We introduce the Near IR dataset that contains videos of emotion data captured from 26 subjects driving the cars in the multiple camera setup. This dataset is collected to support drivers by Advanced Driver Assistance Systems (ADAS). The drivers show six basic facial expressions including anger, disgust, fear, happiness, sadness, surprise plus neutral faces. In our experiments, we use frames (peak expressions) of 20 subjects for training and validation, and 6 subjects for the test, respectively.

TABLE I: The generator architecture. There are some notations; n_y denotes the the dimension of domain attributes. IN and RB denote instance normalization and residual block, respectively.

Part	Layers	Input Size → Output Size	Filter Size	Stride	Padding
Encoder	Conv+IN+ReLU	$(h, w, 6) \rightarrow (h, w, 64)$	7×7	1	3
	Conv+IN+ReLU	$(h, w, 64) \rightarrow \left(\frac{h}{2}, \frac{w}{2}, 128\right)$	4×4	2	1
	Conv+IN+ReLU	$\left(\frac{h}{2}, \frac{w}{2}, 128\right) \rightarrow \left(\frac{h}{4}, \frac{w}{4}, 256\right)$	4×4	2	1
	Conv+IN+ReLU	$\left(\frac{h}{4}, \frac{w}{4}, 256\right) \rightarrow \left(\frac{h}{8}, \frac{w}{8}, 512\right)$	4×4	2	1
	Conv+IN+ReLU	$\left(\frac{h}{8}, \frac{w}{8}, 512\right) \rightarrow \left(\frac{h}{16}, \frac{w}{16}, 1024\right)$	4×4	2	1
Encoder Bottleneck	RB:Conv+IN+ReLU	$\left(\frac{h}{16}, \frac{w}{16}, 1024\right) \rightarrow \left(\frac{h}{16}, \frac{w}{16}, 1024\right)$	3×3	1	1
	RB:Conv+IN+ReLU	$\left(\frac{h}{16}, \frac{w}{16}, 1024\right) \rightarrow \left(\frac{h}{16}, \frac{w}{16}, 1024\right)$	3×3	1	1
	RB:Conv+IN+ReLU	$\left(\frac{h}{16}, \frac{w}{16}, 1024\right) \rightarrow \left(\frac{h}{16}, \frac{w}{16}, 1024\right)$	3×3	1	1
	RB:Conv+IN+ReLU	$\left(\frac{h}{16}, \frac{w}{16}, 1024\right) \rightarrow \left(\frac{h}{16}, \frac{w}{16}, 1024\right)$	3×3	1	1
	RB:Conv+IN+ReLU	$\left(\frac{h}{16}, \frac{w}{16}, 1024\right) \rightarrow \left(\frac{h}{16}, \frac{w}{16}, 1024\right)$	3×3	1	1
	RB:Conv+IN+ReLU	$\left(\frac{h}{16}, \frac{w}{16}, 1024\right) \rightarrow \left(\frac{h}{16}, \frac{w}{16}, 1024\right)$	3×3	1	1
Decoder	Sub-Pixel Conv+IN+ReLU	$\left(\frac{h}{16}, \frac{w}{16}, 1024 + n_y\right) \rightarrow \left(\frac{h}{8}, \frac{w}{8}, 512\right)$	3×3	2	1
	Sub-Pixel Conv+IN+ReLU	$\left(\frac{h}{8}, \frac{w}{8}, 512\right) \rightarrow \left(\frac{h}{4}, \frac{w}{4}, 256\right)$	3×3	2	1
	Sub-Pixel Conv+IN+ReLU	$\left(\frac{h}{4}, \frac{w}{4}, 256\right) \rightarrow \left(\frac{h}{2}, \frac{w}{2}, 128\right)$	3×3	2	1
	Sub-Pixel Conv+IN+ReLU	$\left(\frac{h}{2}, \frac{w}{2}, 128\right) \rightarrow (h, w, 64)$	3×3	2	1
	Image output:Conv+Tanh	$(h, w, 64) \rightarrow (h, w, 3)$	7×7	1	3
	Side output:Conv+Tanh	$(h, w, 64) \rightarrow (h, w, 3)$	7×7	1	3

TABLE II: The discriminator architecture. FC and m denote fully connected layer and the number of target attributes, respectively

Part	Layers	Input Size → Output Size	Filter Size	Stride	Padding
Discriminator Hidden Layers	Conv+Leaky ReLU	$(h, w, 6) \rightarrow \left(\frac{h}{2}, \frac{w}{2}, 64\right)$	4×4	2	1
	Conv+Leaky ReLU	$\left(\frac{h}{2}, \frac{w}{2}, 64\right) \rightarrow \left(\frac{h}{4}, \frac{w}{4}, 128\right)$	4×4	2	1
	Conv+Leaky ReLU	$\left(\frac{h}{4}, \frac{w}{4}, 128\right) \rightarrow \left(\frac{h}{8}, \frac{w}{8}, 256\right)$	4×4	2	1
	Conv+Leaky ReLU	$\left(\frac{h}{8}, \frac{w}{8}, 256\right) \rightarrow \left(\frac{h}{16}, \frac{w}{16}, 512\right)$	4×4	2	1
	Conv+Leaky ReLU	$\left(\frac{h}{16}, \frac{w}{16}, 512\right) \rightarrow \left(\frac{h}{32}, \frac{w}{32}, 1024\right)$	4×4	2	1
	Conv+Leaky ReLU	$\left(\frac{h}{32}, \frac{w}{32}, 1024\right) \rightarrow \left(\frac{h}{64}, \frac{w}{64}, 2048\right)$	4×4	2	1
Outputs	**Output Layer**:Conv	$\left(\frac{h}{64}, \frac{w}{64}, 2048\right) \rightarrow \left(\frac{h}{64}, \frac{w}{64}, 1\right)$	3×3	1	1
	Output Layer:FC	$\left(\frac{h}{64}, \frac{w}{64}, 2048\right) \rightarrow FCm$	–	–	–

Oulu-CASIA VIS [39]: This dataset contains 480 sequences (from 80 subjects) of six basic facial expressions under the visible (VIS) normal illumination conditions. We conducted our experiments using subject-independent 10-fold cross-validation strategy.

MUG [1]: The MUG dataset contains image sequences of seven different facial expressions belonging to 86 subjects comprising 51 men and 35 women. The image sequences were captured with a resolution of 896×896. We used image sequences of 52 subjects and the corresponding annotation, which are available publicly via the internet.

BU-3DFE [36]: The Binghamton University 3D Facial Expression Database (BU-3DFE) [36] contains 3D models from 100 subjects, 56 females and 44 males. The subjects show a neutral face as well as six basic facial expressions and at four different intensity levels. Following the setting in [33] and [37], we used an openGL based tool from the database creators to render multiple views from 3D models in seven pan angles ($0°, \pm15°, \pm30°, \pm45°$).

RaFD [20] : The Radboud Faces Database (RaFD) contains 4,824 images belonging to 67 participants. Each subject makes eight facial expressions.

1) Qualitative evaluation: From qualitative results in Fig. 5, it is obvious that our facial attribute transfer test results (unseen images during the training step) are more visually pleasing compared to other baselines including IcGAN [28] and CycleGAN [41]. We believe our proposed identity loss

helps to preserve the face image details and identity. IcGAN even fails to generate subjects with desired attributes, while our proposed method could learn attribute invariant features applicable to synthesize multiple images with desired attributes. In addition, to evaluate the proposed realism
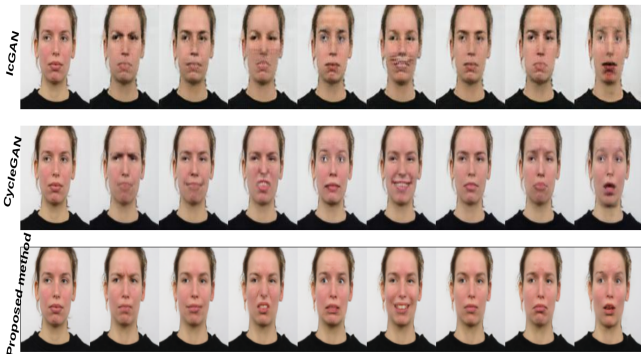

Fig. 5: Facial attribute transfer results of our method compared with IcGAN [28] and CycleGAN [41], respectively.

refinement, the face attribute transfer results of our proposed method have been compared with the SimGAN method [31] on the BU-3DFE dataset [36] (see Fig. 6).
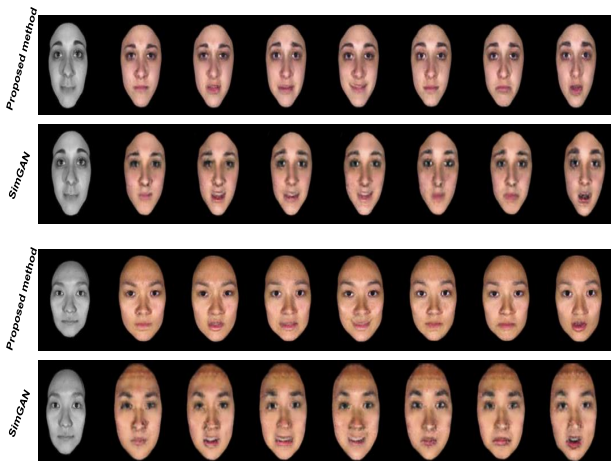

Fig. 6: Pose-normalized face attribute transfer results of our proposed method compared with SimGAN method [31] on the BU-3DFE dataset [36]. **Left to right**: input simulated frontal face and seven different exhibited attributes including *angry*, *disgusted*, *fearful*, *happiness*, *neutral*, *sadness* and *surprised*, respectively.

2) User Study: We also evaluate the realism of our results with a user study to compare our model with CycleGAN [41]. We asked 15 subjects to select results that are more realistic and facial expression is well distinguishable through pairwise comparisons. In addition, a third choice as "None" was also introduced in the case if none of them could generate realistic result. 16 random images with the corresponding emotion transfer results from RaFD [20] dataset

were presented in a randomized fashion to each person. The Pie chart shown in Fig. 7 illustrates that the results reconstructed by our approach are more appealing to the users.
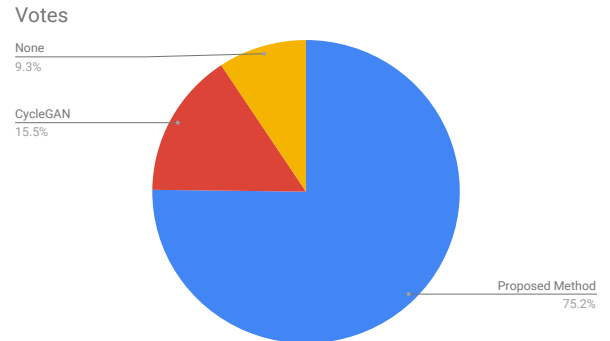

Fig. 7: The evaluation results of our user studies to compare our method with CycleGAN [41].

3) Quantitative Evaluation: We quantitatively demonstrate the usefulness our proposed model in synthesizing photo-realistic facial images controlled by the expression category. Doing so, we augment real images from the Oulu-CASIA VIS dataset with the synthetic expression images generated by our model as well as its variants and then compare with other methods to train an expression classifier. The purpose of this experiment is to introduce more variability and enrich the dataset further to improve the expression recognition performance. In particular, from each of the six expression category, we generate 0.5K, 1K, 2K, 5K and 10K images, respectively. The classifier has an identical network architecture used in synthesizing (RaFD) [20] images except the number of neurons used in the discriminator's fully connected layer. The accuracy results for the expression recognition are shown in Fig. 8. We can observe that when the number of synthetic images is increased to 30K, the accuracy is improved drastically, reaching to 86.95%. The performance starts to become saturated when more images (60K) are used. We achieved a higher recognition accuracy value using the images generated from our method than the state-of-the-arts including CNN-based methods e.g., DTAGN [13]. This suggests that our model has learned to generate more diverse realistic images. In addition, we evaluate the sensitivity of the results for different components of our proposed method (bidirectional loss and side conditional image, respectively). Moreover, we evaluate the performance of our proposed method on the MUG facial expression dataset, [1] using the video frames of the peak expressions. In Table III, we report the results of average accuracy of a facial expression on synthesized images. We trained a facial expression classifier with $(90\%/10\%)$ splitting for training and test sets using a ResNet-50 [10], resulting in a high accuracy of 90.42%. We then trained each of baseline models including CycleGAN and IcGAN using the same training set and performed image-to-image translation on the same test set. Finally, we classified the expression of these generated
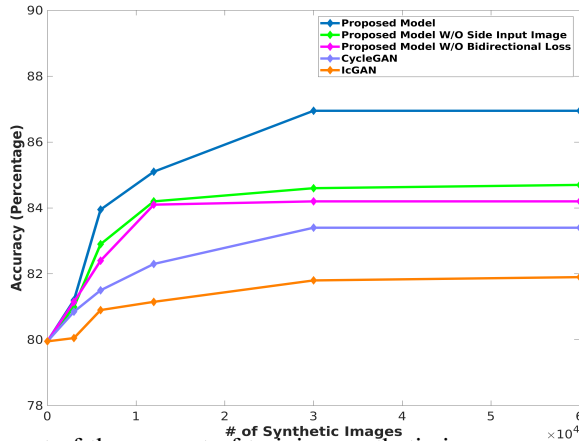
Method	Accuracy
HOG 3D [17]	70.63%
AdaLBP [39]	73.54%
Atlases [9]	75.52%
STM-ExpLet [23]	74.59%
DTAGN [13]	81.46%
Proposed Method	**86.95%**

Fig. 8: Impact of the amount of training synthetic images on performance in terms of expression recognition accuracy (**left**). Performance comparison of expression recognition accuracy between the proposed method and other state-of-the-art methods (**right**).

images using the above-mentioned classifier. As can be seen in Table III, our model achieves the highest classification accuracy (close to real image), demonstrating that our model could generate the most realistic expressions among all the methods compared.

For the near IR drivers' dataset, we conducted two set of experiments. In the first experiment, we trained facial expression classifiers with subject-independent subsets (20 subjects for training and validation and 6 subjects for the test). We used multi-view convolutional neural network (MVCNN) [32] as our baseline. The VGG-Face model [27] is used as the bottleneck network. In the second experiment setup, we utilized face frontalization scheme and added realisms to the simulated faces using our proposed approach and [31], respectively. As can be seen in Table IV, our model achieves the highest classification accuracy, demonstrating that our realism refinement facilitates the synthesized images to preserve much detail of face expression. Finally, as our

TABLE III: Performance comparison on the MUG dataset.

Method	Accuracy
Real Test Images	90.42%
CycleGAN [41]	84.40%
IcGAN [28]	80.32%
Proposed Method	**89.91%**

TABLE IV: Performance comparison on the near IR drivers' dataset.

Method	Accuracy
MVCNN [32]	65.35%
SimGAN [31]	69.21%
Proposed Method	**78.85%**

last experiment, we performed 5-fold cross validation using 100 subjects for the BU-3DFE dataset [36]. Training data includes images of 80 (non-frontal face) subjects, while test data includes images of 20 subjects with varying poses. We use the VGG-Face model [27], which is pretrained on the

(RaFD) [20] and then we further fine-tune it on the frontal face images from BU-3DFE dataset. It can be observed from Table V that face frontalization contributes to the expression recognition performance of the profile faces (ranging from 15 to 45 degrees in 15 degrees steps). Having said that, adding realism to the synthetic images (simulated frontal face) helps to bring additional gains in terms of expression recognition accuracy.

TABLE V: Recognition accuracies on face images at different pose yaw angles.

Method	±15	±30	±45
Real Profile Images	70.15%	66.50%	58.90%
Synthetic Frontal Face Images	70.91%	65.90%	59.30%
Proposed Method	**72.10%**	**68.52%**	**63.35%**

VI. CONCLUSION AND FUTURE WORK

In this paper we propose attribute guided face image synthesis method, which is capable to synthesize photo-realistic face images conditioned on desired attributes. Using our proposed attribute classification objective and incorporating bidirectional learning, we demonstrate a proper way to model latent representation among different domains leading realistic face images as the result. More importantly, we seek to reduce the domain distribution mismatch between synthetic and real faces. In addition, we demonstrate that the synthetic images generated by our method can be used for data augmentation to train face expression classier. We achieve significantly higher average accuracy compared with the state-of-the-art result. In particular, the proposed method surpasses previous approaches by a significant margin of 5.5% on Oulu-CASIA VIS dataset. For the future work, we plan to apply our model to translate dynamic textures of a face from a single image in the context of video domain.

REFERENCES

[1] N. Aifanti, C. Papachristou, and A. Delopoulos. The mug facial expression database. In *Image analysis for multimedia interactive services (WIAMIS), 2010 11th international workshop on*, pages 1–4. IEEE, 2010.

[2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017.

[3] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

[4] B. Bozorgtabar, K. E. H. Rad, Mohammad Saeed, and J.-P. Thiran. Learn to synthesize and synthesize to learn. *Submitted to Journal of Computer Vision and Image Understanding (CVIU). Special Issue on Adversarial Learning in Computer Vision*, 2018.

[5] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8789–8797, 2018.

[6] A. Dosovitskiy and T. Brox. Generating images with perceptual similarity metrics based on deep networks. In *Advances in Neural Information Processing Systems*, pages 658–666, 2016.

[7] 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.

[8] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5769–5779, 2017.

[9] Y. Guo, G. Zhao, and M. Pietikäinen. Dynamic facial expression recognition using longitudinal facial expression atlases. In *Computer Vision–ECCV 2012*, pages 631–644. Springer, 2012.

[10] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[11] R. Huang, S. Zhang, T. Li, R. He, et al. Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis. *arXiv preprint arXiv:1704.04086*, 2017.

[12] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1125–1134, 2017.

[13] H. Jung, S. Lee, S. Park, I. Lee, C. Ahn, and J. Kim. Deep temporal appearance-geometry network for facial expression recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV 2015)*. IEEE, 2015.

[14] T. Kaneko, K. Hiramatsu, and K. Kashino. Generative attribute controller with conditional filtered generative adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7006–7015. IEEE, 2017.

[15] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim. Learning to discover cross-domain relations with generative adversarial networks. In *International Conference on Machine Learning*, pages 1857–1865, 2017.

[16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[17] A. Klaser, M. Marszałek, and C. Schmid. A spatio-temporal descriptor based on 3d-gradients. In *BMVC 2008-19th British Machine Vision Conference*, pages 275–1. British Machine Vision Association, 2008.

[18] Y.-H. Lai and S.-H. Lai. Emotion-preserving representation learning via generative adversarial network for multi-view facial expression recognition. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 263–270. IEEE, 2018.

[19] G. Lample, N. Zeghidour, N. Usunier, A. Bordes, L. Denoyer, et al. Fader networks: Manipulating images by sliding attributes. In *Advances in Neural Information Processing Systems*, pages 5969–5978, 2017.

[20] O. Langner, R. Dotsch, G. Bijlstra, D. H. Wigboldus, S. T. Hawk, and A. Van Knippenberg. Presentation and validation of the radboud faces database. *Cognition and emotion*, 24(8):1377–1388, 2010.

[21] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4681–4690, 2017.

[22] M. Li, W. Zuo, and D. Zhang. Deep identity-aware transfer of facial attributes. *arXiv preprint arXiv:1610.05586*, 2016.

[23] M. Liu, S. Shan, R. Wang, and X. Chen. Learning expressionlets on spatio-temporal manifold for dynamic facial expression recognition. In *2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1749–1756. IEEE, 2014.

[24] Y. Lu, Y.-W. Tai, and C.-K. Tang. Attribute-guided face generation using conditional cyclegan. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 282–297, 2018.

[25] R. Natsume, T. Yatagawa, and S. Morishima. Rsgan: Face swapping and editing using face and hair representation in latent spaces. *arXiv preprint arXiv:1804.03447*, 2018.

[26] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier gans. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2642–2651. JMLR. org, 2017.

[27] O. M. Parkhi, A. Vedaldi, A. Zisserman, et al. Deep face recognition. In *BMVC*, volume 1, page 6, 2015.

[28] G. Perarnau, J. van de Weijer, B. Raducanu, and J. M. Álvarez. Invertible conditional gans for image editing. *arXiv preprint arXiv:1611.06355*, 2016.

[29] W. Shen and R. Liu. Learning residual images for face attribute manipulation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1225–1233. IEEE, 2017.

[30] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1874–1883, 2016.

[31] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from simulated and unsupervised images through adversarial training. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 3, page 6, 2017.

[32] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 945–953, 2015.

[33] U. Tariq, J. Yang, and T. S. Huang. Maximum margin gmm learning for facial expression recognition. In *Automatic Face and Gesture Recognition (FG), 2013 10th IEEE International Conference and Workshops on*, pages 1–6. IEEE, 2013.

[34] D. Vlasic, M. Brand, H. Pfister, and J. Popović. Face transfer with multilinear models. *ACM transactions on graphics (TOG)*, 24(3):426–433, 2005.

[35] T. Xiao, J. Hong, and J. Ma. Elegant: Exchanging latent encodings with gan for transferring multiple face attributes. *arXiv preprint arXiv:1803.10562*, 2018.

[36] L. Yin, X. Wei, Y. Sun, J. Wang, and M. J. Rosato. A 3d facial expression database for facial behavior research. In *Automatic face and gesture recognition, 2006. FGR 2006. 7th international conference on*, pages 211–216. IEEE, 2006.

[37] F. Zhang, T. Zhang, Q. Mao, and C. Xu. Joint pose and expression modeling for facial expression recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3359–3368, 2018.

[38] G. Zhang, M. Kan, S. Shan, and X. Chen. Generative adversarial network with spatial attention for face attribute editing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 417–432, 2018.

[39] G. Zhao, X. Huang, M. Taini, S. Z. Li, and M. PietikäInen. Facial expression recognition from near-infrared videos. *Image and Vision Computing*, 29(9):607–619, 2011.

[40] Y. Zhou and B. E. Shi. Photorealistic facial expression synthesis by the conditional difference adversarial autoencoder. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 370–376. IEEE, 2017.

[41] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2223–2232, 2017.