

High Fidelity Face Manipulation with Extreme Poses and Expressions

Chaoyou Fu, Yibo Hu, Xiang Wu, Guoli Wang, Qian Zhang, and Ran He, *Senior Member, IEEE*

Abstract—Face manipulation has shown remarkable advances with the flourish of Generative Adversarial Networks. However, due to the difficulties of controlling structures and textures, it is challenging to model poses and expressions simultaneously, especially for the extreme manipulation at high-resolution. In this paper, we propose a novel framework that simplifies face manipulation into two correlated stages: a boundary prediction stage and a disentangled face synthesis stage. The first stage models poses and expressions jointly via boundary images. Specifically, a conditional encoder-decoder network is employed to predict the boundary image of the target face in a semi-supervised way. Pose and expression estimators are introduced to improve the prediction performance. In the second stage, the predicted boundary image and the input face image are encoded into the structure and the texture latent space by two encoder networks, respectively. A proxy network and a feature threshold loss are further imposed to disentangle the latent space. Furthermore, due to the lack of high-resolution face manipulation databases to verify the effectiveness of our method, we collect a new high-quality Multi-View Face (MVF-HQ) database. It contains 120,283 images at 6000×4000 resolution from 479 identities with diverse poses, expressions, and illuminations. MVF-HQ is much larger in scale and much higher in resolution than publicly available high-resolution face manipulation databases. We will release MVF-HQ soon to push forward the advance of face manipulation. Qualitative and quantitative experiments on four databases show that our method dramatically improves the synthesis quality.

Index Terms—Face Manipulation, Extreme Pose and Expression, High-Resolution, MVF-HQ.

I. INTRODUCTION

PHOTO-realistic face manipulation with poses and expressions is a meaningful task in a wide range of fields, such as movie industry, entertainment, and photography technologies. With the flourish of Generative Adversarial Networks (GANs) [1], [2], face manipulation has achieved significant advances in recent years [3], [4], [5]. However, on the one hand, existing face manipulation methods mainly focus on only one facial variation, *i.e.* only changing poses [6] or expressions [4]. On the other hand, the methods of manipulating large poses and expressions have still been limited to a low-resolution (128×128) [7]. Hence, it is challenging to model both poses

and expressions [8], especially for the extreme manipulation (facial poses beyond $\pm 60^\circ$) at high-resolution.

For face manipulation, a straightforward way is applying image-to-image translation methods [9], [10]. However, in the case of extreme manipulation, it is hard for this way to guarantee facial structures and textures. As shown in Fig. 1 (a), the facial local structures, such as the eyes, the nose, and the mouth, are unclear. Recent observations [10], [11] show that the boundary information is crucial to high fidelity image synthesis. Meanwhile, several structure-guided methods have been proposed for face manipulation [7], [12]. For example, CAPG-GAN [7] leverages facial landmarks to rotate faces. SC-FEGAN [12] realizes facial local editing with sketch guidance. Most structure-guided methods directly concatenate a face image and its structure guidance in the image space. However, such direct concatenation is infeasible for the extreme manipulation at high-resolution, due to more complex structures and textures. As shown in Fig. 1 (b), the local structures of the synthesized face are fuzzy and the textures are blurry, such as the mouth. This phenomenon may be caused by lacking disentanglement, which is important for interpretable image manipulation [13], between structures and textures. Recently, [14] proposes a simple disentanglement method. It introduces a face recognition network to extract structure invariant features, and then concatenates the structure invariant features with structure features to synthesize faces. As shown in Fig. 1 (c), the disentanglement does make the structure of the synthesized face clearer, but the textures are somewhat lost. This may be due to the fact that the features of the face recognition network are too compact, leading to severe texture loss in such a high-resolution case.

Based on the above observations, we propose a novel framework for high-resolution extreme face manipulation, as shown in Fig. 3. Our framework simplifies this challenging task into two correlated stages: a boundary prediction stage and a disentangled face synthesis stage. The first stage models poses and expressions jointly via boundary images. Different from most of the previous methods that require manually manipulation of landmarks [11], [7], we can flexibly generate the desired boundary images merely by controlling the pose and the expression vectors. Specifically, a conditional encoder-decoder network is employed to predict the boundary image of the target face in a semi-supervised way. Pose and expression estimators are introduced to improve the prediction performance. The second stage encodes the predicted boundary image and the input face image into the structure and the texture latent space by two encoder networks, respectively. Since it is hard to disentangle structures and textures [13], we

C. Fu, Y. Hu, X. Wu, 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 the School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing 100190, China. E-mail: {chaoyou.fu, rhe}@nlpr.ia.ac.cn, {huyibo871079699, alfredxiangwu}@gmail.com. Q. Zhang is with Horizon Robotics, Beijing 100190, China. E-mail: qian01.zhang@horizon.ai. G. Wang is with the Department of Automation, Tsinghua University, Beijing 100084, China. E-mail: wangguoli1990@mail.tsinghua.edu.cn. (Corresponding author: Ran He.)




Fig. 1: Visualization comparisons (512×512 resolution) of different methods. (a) Image-to-image translation [10]. The local structures, *e.g.* the mouth, are unclear; (b) Concatenating the input face and the boundary of the target face. The local structures are fuzzy and the textures are blurry; (c) Directly utilizing a face recognition network to disentangle structures and textures [14]. The local structures are clearer, but the textures are somewhat lost; (d) Our method. Both of the structures and the textures are maintained well.

introduce a face recognition network as a proxy to facilitate disentanglement. Instead of leveraging the compacted features of the proxy [14] directly, we propose a feature threshold loss to control the compactness between our learned face features and the compacted features. The result of our method in Fig. 1 (d) has not only clear structures, but also realistic textures. More synthesis results (1024×1024 resolution) are presented in Fig. 8, from which we observe that our method generates high-quality images even under the largest 90°. Note that it is rather challenging to manipulate high-resolution faces with such **extreme facial structure changes**. As far as we know, this has never been achieved except for our work.

In addition, most publicly available face manipulation databases are limited in resolution, *e.g.* the resolution of the original images in the MultiPIE database [15] is only 640×480. Although recent high-resolution databases CelebA-HQ [16] and FFHQ [17] reach a resolution of 1024×1024, the pose variants in these databases are inadequate. According to our statistics (estimating poses via the detected facial landmarks), about 92% of the images in CelebA-HQ are nearly frontal (within $\pm 30^\circ$), and only about 0.7% of the images have large poses (beyond $\pm 60^\circ$), as shown in Fig. 2. For FFHQ, the nearly frontal images and large pose images account for about 90% and 0.8%, respectively. Due to the limitation of these databases, it seems impossible for existing GANs-based methods, such as StyleGAN [17], [18], to generate high-quality faces with large poses. Therefore, we collect a new high-quality Multi-View Face (MVF-HQ) that contains 13 views from -90° to $+90^\circ$ (the interval is 15°), which fills in the blank of the high-resolution multi-view database. The comparisons between our MVF-HQ database and other




Fig. 2: Facial pose distribution of high-resolution databases, including CelebA-HQ [16], FFHQ [17], and our MVF-HQ.

publicly available high-resolution face manipulation databases are presented in Fig. 2 and Table I. It is obvious that MVF-HQ has the following three advantages: (1) **Large-Scale**. MVF-HQ consists of 120,283 images, far more than other high-resolution databases (70,000 at most [17]). (2) **High-Resolution**. The resolution of the original images in MVF-HQ is up to 6000×4000, while other databases only reach 1024×1024 resolution [16], [17]. (3) **Abundant Variants**. MVF-HQ contains 13 views from -90° to $+90^\circ$ as well as diverse expressions and illuminations. It will be released soon to push forward the advance of face manipulation.

In summary, the main contributions are as follows:

- High-resolution extreme face manipulation is formulated as a stage-wise learning problem, which contains two correlated stages: a boundary prediction stage and a disentangled face synthesis stage.
- The joint modeling of poses and expressions is implemented by the boundary prediction in the first stage, which makes it flexible to obtain desired boundary images. Besides, a proxy network and a feature threshold loss are introduced in the second stage to disentangle structures and textures for realistic synthesis.
- We collect a new MVF-HQ database that contains 120,283 diverse face images at 6000×4000 resolution from 479 identities. It has much larger scale and much higher resolution than publicly available high-resolution face manipulation databases.
- We are the first to achieve the rather challenging high-resolution face manipulation with extreme facial structure changes, even for faces at 90° . Extensive experiments on the MultiPIE [15], the RaFD [19], the CelebA-HQ [16], and our MVF-HQ databases fully demonstrate the effectiveness of our method.

II. RELATED WORK

A. Face Manipulation

Face manipulation has attracted great attention to computer vision and graphics [20], [21], [22], [23], [24], [25], [26], [27], [28]. Pose rotation [29], [7] and expression editing [4], [30] are the two of the main research directions. TP-GAN [6] adopts a two-pathway generative network architecture, achieving photo-realistic face frontalization from a single image. DA-GAN [31] employs a dual attention mechanism into face frontalization,

including a self-attention for the generator and a face-attention for the discriminator. FaceID-GAN [5] introduces an identity classifier as a competitor to better preserve identity for face synthesis. UV-GAN [32] completes the facial UV map to improve the performance of pose-invariant face recognition. By controlling the magnitude of Action Units (AUs), GAN-imation [4] renders expressions in a continuum. MoCoGAN [30] separates the hidden features as a content subspace and a motion subspace, which makes it possible to synthesize different expressions. AF-VAE [33] introduces an additive Gaussian Mixture prior in the structure space, facilitating multi-model face synthesis.

Other face manipulation tasks also have achieved considerable development, such as facial makeup [34], face inpainting [35], [36], cross-spectral hallucination [37], [38]. PairedCycleGAN [34] introduces a new cycle generative network that transfers makeup styles and removes the makeup in an asymmetric manner. By this means, it is able to wear makeup for a target face based on a reference. StarGAN [3], [39] realizes multi-domain face attribute transfer by a single generator. [35] utilizes the surrounding background patches to facilitate image inpainting. InterFaceGAN [40] realizes face editing via analyzing the latent semantics encoded by GANs. [41] first proposes a unified framework to achieve pose-invariant facial expression recognition, face synthesis, and face alignment simultaneously. The elaborate design enables these modules to complement and enhance each other.

Besides the above deep generative model based methods, there are also many approaches based on the 3D morphable model (3DMM) [42]. As a comparatively mature face analysis technology, 3DMM has been incorporated into various face manipulation tasks. [43] suggests the best facial makeup automatically for users. [44] realizes real-time facial expression transfer via an RGB-D sensor. Face2Face [45] is able to reenact a video sequence according to a source actor. [46] makes the generated video match a given audio track.

B. Image Synthesis

As one of the primary means for face manipulation, image synthesis has made great progress in recent years [1], [47], [48], [49], [50]. Image synthesis can be done either in an unconditional manner [16], [51], [52] or in a conditional manner [9], [53]. For unconditional synthesis, images are generated from noises without any conditions. GANs [1], [54], [2] are the representative unconditional synthesis models that consist of a generator and a discriminator to play a min-max game. The generator synthesizes data from a prior to confuse the discriminator, while the discriminator tries to distinguish the generative data from the real data. PG-GAN [16] significantly improves the synthesis quality by progressively growing the generator and the discriminator. Variational AutoEncoders (VAEs) [47] are the other representative unconditional synthesis models, which optimize the evidence lower bound objective (ELBO) to learn the data distribution. IntroVAE [55] employs an introspective variational generation model to synthesize high-resolution images without discriminators. VQ-VAE [56], [57] learns discrete representations with an autoregressive prior for high-quality generation.

For conditional image synthesis, the synthesized images need to meet the given conditions. pix2pix [9] introduces a conditional generative adversarial loss for paired image-to-image translation. pix2pixHD [10] further improves pix2pix with a coarse-to-fine generator and a multi-scale discriminator, realizing higher fidelity image translation. CycleGAN [58] proposes a cycle consistent way for unpaired image-to-image translation. BigGAN [59] first achieves high-resolution (512×512) conditional image synthesis on the ImageNet database [60]. StyleGAN [17], [18] employs an alternative generator to learn attributes automatically. Benefitting from the proposed spatially-adaptive normalization, SPADE [53] synthesizes photo-realistic landscapes by the semantic layout.

C. Face manipulation databases

Publicly available face manipulation databases can be divided into two categories: low-resolution databases and high-resolution databases. Representative low-resolution databases include Celeb-A [61], CAS-PEAL-R1 [62], and MultiPIE [15]. The resolutions of the original images in these three databases are 505×606 , 640×480 , and 640×480^1 , respectively. Celeb-A is an in-the-wild face database that contains 202,599 images with 40 attribute annotations. It is widely used in face attribute editing. Both the CAS-PEAL-R1 and the MultiPIE databases, which have 30,863 and 750,000+ images respectively, are in the controlled environment. They are mainly adopted for pose invariant face recognition.

Due to the high cost of data acquisition, there are a limited number of high-resolution face manipulation databases. Among them, RaFD [19] is widely used for facial expression analysis. However, both the number of images (8,040) and the resolution (681×1024) are limited. BioHDD [63] is mainly leveraged to evaluate biometric identification in the case of heavy degradation. The resolution of the registration images in BioHDD is up to 4368×2912 , whereas the number of these images is only 606^2 . Recently released CelebA-HQ [16] and FFHQ [17] reach 1024×1024 resolution. The former contains 30k images that mainly derive from the Celeb-A database. Researchers adopt image processing techniques, such as super-resolution, to convert low-resolution images into higher-resolution ones. The latter consists of 70k images that are crawled from Flickr. However, according to our statistics, most of the images in CelebA-HQ and FFHQ are nearly frontal, making it hard to edit large poses on these databases.

III. METHOD

Given an input face I^a , the goal of our method is to synthesize the target face I^b , according to a given pose vector p^b and an expression vector e^b . In addition, we denote the boundary image of the input face and the target face as B^a and B^b , respectively. The face manipulation task is explicitly divided into two stages: a boundary prediction stage and a disentangled face synthesis stage, as shown in Fig. 3. In the rest of this section, we will present them in detail.

¹The MultiPIE database contains a small number of frontal images at 3072×2048 resolution, but most of the images are 640×480 resolution.

²BioHDD also has many low-resolution images with various attributes. Please refer to [63] for details.

A. Boundary Prediction

At this stage, we predict the target boundary image according to the given conditional vectors, including a pose vector and an expression vector. As shown in Fig. 3, we utilize an encoder network Enc and a decoder network Dec to realize this conditional boundary prediction. Specifically, through Enc , we first map the input boundary image B^a into the latent space $z^a = Enc(B^a)$. Then, the target pose vector p^b and the expression vector e^b are concatenated with the hidden variable z^a to provide conditional information. Lastly, the target boundary image is generated by the decoder network $\hat{B}^b = Dec(z^a, p^b, e^b)$.

The pose and the expression are discrete in the database, e.g. the MultiPIE database [15] only contains 15 discrete poses and 6 discrete expressions. However, we expect that this stage can generate boundary images with unseen poses and expressions. Hence, we introduce a semi-supervised training manner. For the poses and the expressions in the database, we can utilize the corresponding ground truth to constrain the generated boundary images. For the poses and the expressions that do not exist in the database, we utilize two pre-trained estimators, including a pose estimator F_p and an expression estimator F_e , to constrain the generated boundary images by conditional regression.

Two loss functions are involved in this stage, including a pixel-wise loss and a conditional regression loss.

Pixel-Wise Loss. For the poses and the expressions that belong to the database, a pixel-wise L_1 loss is utilized to constrain the predicted boundary image $\hat{B}^b = Dec(Enc(B^a), p^b, e^b)$:

$$\mathcal{L}_{\text{pix-bound}} = |Dec(Enc(B^a), p^b, e^b) - B^b|, \quad (1)$$

where B^b is the ground truth of the target boundary image.

Conditional Regression Loss. For the poses and the expressions that do not exist in the database, we first randomly produce p^r and e^r to generate a boundary image $B^r = Dec(z^a, p^r, e^r)$. Then, we utilize a pose estimator F_p and an expression estimator F_e , which are pre-trained on the boundary images of the training databases, to estimate the pose $\hat{p}^r = F_p(B^r)$ and the expression $\hat{e}^r = F_e(B^r)$, respectively. The estimated \hat{p}^r and \hat{e}^r are used to constrain the generated boundary image. The intuition is that the estimated \hat{p}^r and \hat{e}^r of B^r should be equal to the conditional vectors p^r and e^r , respectively. Hence, the conditional regression loss, including a pose regression term and an expression regression term, is formulated as:

$$\begin{aligned} \mathcal{L}_{\text{reg}} = & \|F_p(Dec(z^a, p^r, e^r)) - p^r\|_2^2 \\ & + \|F_e(Dec(z^a, p^r, e^r)) - e^r\|_2^2. \end{aligned} \quad (2)$$

The parameters of the pre-trained F_p and F_e are fixed during the training procedure.

B. Disentangled Face Synthesis

This stage leverages the predicted boundary image to perform realistic face synthesis. As shown in Fig. 3, we first utilize two encoders G_{enc}^B and G_{enc}^I to map the predicted

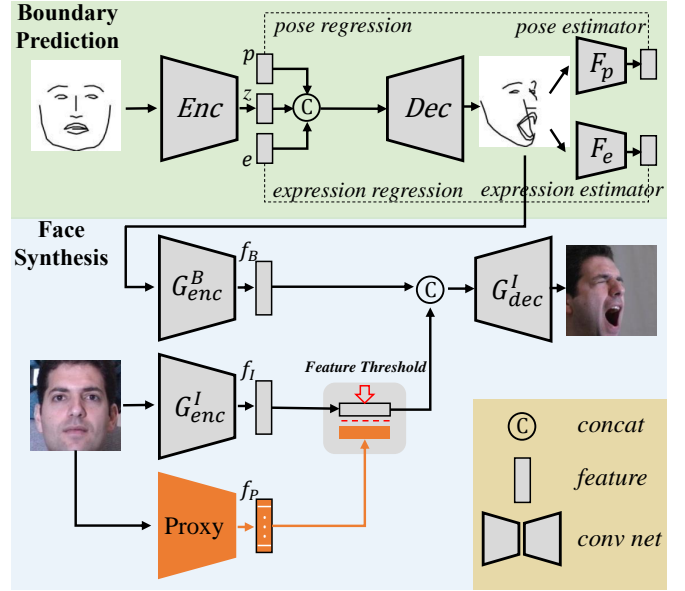


Fig. 3: The framework of our method, which consists of a boundary prediction stage and a disentangled face synthesis stage. The first stage predicts the boundary image of the target face in a semi-supervised way. Pose and expression estimators are employed to improve the prediction performance. The second stage leverages the predicted boundary image to synthesize the target face. A proxy network and a feature threshold loss are introduced to disentangle structures and textures in the latent space.

boundary image \hat{B}^b and the input face I^a to $f_{B^b} = G_{enc}^B(\hat{B}^b)$ and $f_{I^a} = G_{enc}^I(I^a)$, respectively. Then, we disentangle structures and textures in the latent space, by a proxy network $Proxy$ and a feature threshold loss. After the disentanglement, the boundary features f_{B^b} and the image feature f_{I^a} are concatenated to feed into the decoder G_{dec}^I , synthesizing the final target face $\hat{I}^b = G_{dec}^I(f_{B^b}, f_{I^a})$.

The loss functions in this stage are presented as below, including a feature threshold loss, a multi-scale pixel-wise loss, a multi-scale conditional adversarial loss, and an identity preserving loss.

Feature Threshold Loss. The feature threshold loss is designed to assist in disentangling structures and textures in the latent space. Given that directly disentangling is difficult, we employ a face recognition network LightCNN [64] pre-trained on MS-Celeb-1M [65] as a proxy network $Proxy$, whose features $f_{P^a} = Proxy(I^a)$ are thought to be structure invariant [14]. In addition, instead of directly utilizing the compact features f_{P^a} that will result in texture loss, as shown in Fig. 1 (c), we introduce a feature threshold loss for disentanglement. Specifically, the feature threshold loss controls the feature distance between the learned face features $f_{I^a} = G_{enc}^I(I^a)$ and the compact features $f_{P^a} = Proxy(I^a)$ with a threshold margin m :

$$\mathcal{L}_{\text{thr}} = [\|G_{enc}^I(I^a) - Proxy(I^a)\|_2^2 - m]^+, \quad (3)$$

where $[\cdot]^+ = \max(0, \cdot)$. As the loss \mathcal{L}_{thr} decreases, the face

features f_{I^a} are closer to the compact features f_{P^a} , which means structures and textures are better disentangled. Meanwhile, the threshold margin m controls the compact degree of the face features f_{I^a} , which is employed to maintain textures. The parameter analysis of m is presented in Section V-D.

Multi-Scale Pixel-Wise Loss. We introduce a multi-scale pixel-wise loss to constrain the synthesized face on different scales. Specifically, with the downsampling operation on factors of 2 and 4, we first obtain an image pyramid of 3 scales of the synthesized faces and the ground truth faces, respectively. Then, we calculate the pixel-wise loss on these 3 scales faces:

$$\mathcal{L}_{\text{pix-mul}} = \sum_{s=1,2,3} |G_{dec}^I(f_{B^b}, f_{I^a})_s - I_s^b|, \quad (4)$$

where s denotes the scales. The pixel-wise loss at the top of the image pyramid pays more attention to the global information, because it has a larger receptive field. On the contrary, the pixel-wise loss in the bottom of the image pyramid is more concerned with the recovery of details.

Multi-Scale Conditional Adversarial Loss. To improve the sharpness of the synthesized face image, we also introduce a conditional adversarial loss. The discriminator tries to distinguish the fake image pair $\{\hat{I}^b, B^b\}$ from the real image pair $\{I^b, B^b\}$, while the generator tries to fool the discriminator:

$$\begin{aligned} \mathcal{L}_{\text{adv}} = & \mathbb{E}_{I^b \sim P(I^b)} [\log D(I^b, B^b)] \\ & + \mathbb{E}_{\hat{I}^b \sim P(\hat{I}^b)} [\log(1 - D(\hat{I}^b, B^b))]. \end{aligned} \quad (5)$$

In order to improve the ability of the discriminator, we adopt the strategy of [10] that discriminates the synthesized images on three different scales.

Identity Preserving Loss. In order to further preserve the identity information of the synthesized faces, we adopt an identity preserving loss as [7]. Specifically, a pre-trained face recognition network [64] is introduced as a feature extractor D_{ip} . It forces the identity features of the synthesized face \hat{I}^b to be as close to the identity features of the real face I^b as possible. The identity preserving loss is formulated as:

$$\mathcal{L}_{\text{ip}} = \|D_{ip}^p(\hat{I}^b) - D_{ip}^p(I^b)\|_2^2 + \|D_{ip}^{fc}(\hat{I}^b) - D_{ip}^{fc}(I^b)\|_2^2. \quad (6)$$

where D_{ip}^p and D_{ip}^{fc} denote the output of the last pooling layer and the fully connected layer, respectively.

C. Overall Loss

The boundary prediction stage and the disentangled face synthesis stage are trained separately. That is, the boundary prediction stage is first trained, and then the face synthesis stage is trained based on the predicted boundary. For the boundary prediction stage, the overall loss is:

$$\mathcal{L}_{\text{bp}} = \mathcal{L}_{\text{pix-bound}} + \lambda_1 \mathcal{L}_{\text{reg}}. \quad (7)$$

For the face synthesis stage, the overall loss is:

$$\mathcal{L}_{\text{fs}} = \mathcal{L}_{\text{adv}} + \alpha_1 \mathcal{L}_{\text{thr}} + \alpha_2 \mathcal{L}_{\text{pix-mul}} + \alpha_3 \mathcal{L}_{\text{ip}}, \quad (8)$$

where λ_1 , α_1 , α_2 , and α_3 are trade-off parameters.




Fig. 4: Examples of the expressions.




Fig. 5: Examples of the illuminations.

IV. MULTI-VIEW FACE (MVF-HQ) DATABASE

Due to the lack of high-resolution face databases to verify the effectiveness of our method, we collect the MVF-HQ database. This section introduces the details of MVF-HQ in terms of technical setup, data acquisition, data processing, and comparisons.

A. Technical Setup

13 Canon EOS digital SLR cameras (EOS 1300D/1500D with 55mm prime lens) are equipped to take photos. The resolution of these photos is up to 6000×4000 . In order to ensure the accuracy of the poses of the collected photos, we elaborately design and build a horizontal semicircular bracket with a radius of 1.5m, located at the same height as the head. All cameras are placed on the bracket with 15° interval and point to the center of the semicircular bracket, as shown in Fig. 7. Meanwhile, all cameras are connected to one computer through USB interfaces. We design a software that can control all cameras to take photos simultaneously. The taken photos are automatically stored on the hard drive.

We also use 7 flashes for illuminations. These flashes are placed on the above, front, front-above, front-below, behind, left and right, respectively. By turning on one flash and turning off the others, we can simulate different weak lighting conditions. Besides, a chair is placed in the center of the semicircular bracket to fix the pose of the participants. Furthermore, we set a uniform white background for the acquisition environment.

B. Data Acquisition

We invite a total of 500 participants and all of them have signed data acquisition licenses before taking photos. Each participant is asked to sit down in the chair, and then fine-tune the height of the chair to ensure the head is at the same height as cameras. During the process of data acquisition, the participant is asked to look into the direction of the camera on 0° (see Fig. 7) and show three facial expressions, including neutral, smile and surprise, respectively. Each expression is photographed under all illuminations. The flashes are switched automatically and quickly to guarantee the consistency of



Fig. 6: Examples (6000×4000 resolution) of the thirteen views from -90° to $+90^\circ$.

TABLE I: Comparisons of publicly available low-resolution (the first four databases) and high-resolution (the last five databases) face manipulation databases. ‘-’ means no label.

Database	Images	Resolution	ID	Poses	Expressions
PIE [66]	41,000+	640×486	68	13	4
MultiPIE [15]	750,000+	640×480	337	15	6
CelebA [61]	202,599	505×606	10,177	-	-
CAS-PEAL-R1 [62]	30,863	640×480	1,040	21	5
RaFD [19]	8,040	681×1024	73	5	8
CelebA-HQ [16]	30,000	1024×1024	-	-	-
FFHQ [17]	70,000	1024×1024	-	-	-
BioHDD [63]	606	4368×2912	101	3	1
MVF-HQ	120,283	6000×4000	479	13	3

poses and expressions under different illuminations. The examples of different views, expressions, and illuminations are displayed in Fig. 6, Fig. 4, and Fig. 5, respectively.

C. Data Processing

After data acquisition, we carefully check each original image to clean the database. The blurred images and the images with nonstandard poses are removed from the database. Ultimately, we select 120,283 images from 479 identities. Since it is hard for landmark detection algorithms to accurately detect landmarks under large poses, we manually mark five landmarks for the faces with poses of $\pm 60^\circ$, $\pm 75^\circ$, and $\pm 90^\circ$. The landmarks of other poses are automatically detected by the algorithm and checked by human. All facial landmarks will be released along with the MVF-HQ database.

D. Comparisons

Table I presents the comparisons between our MVF-HQ database and publicly available high-resolution databases. We observe that MVF-HQ has the following three advantages: (1) Large-Scale. MVF-HQ consists of 120,283 images, far more than other databases (70,000 at most [17]). (2) High-Resolution. The high-performance digital SLR camera enables the obtained images to have a resolution of 6000×4000, while




Fig. 7: Technical setup.

other high-resolution databases only reach 1024×1024 resolution [16], [17]. (3) Abundant Variants. MVF-HQ contains diverse poses, expressions, and illuminations.

V. EXPERIMENTS

In this section, we evaluate our method on the MultiPIE [15], the RaFD [19], the CelebA-HQ [16], and our newly built MVF-HQ databases. The introductions of these databases, such as the number of the images, the resolution, and other attributes, are listed in Table I. The details of these databases and experimental settings are reported in Section V-A. In Section V-B and Section V-C, extensive qualitative and quantitative results are provided, respectively. In Section V-D, detailed experimental analyses are described.

A. Databases and Settings

MultiPIE is a low-resolution multi-view database for face recognition and synthesis. In our experiments, we adopt two different settings for quantitative and qualitative experiments, respectively. For the quantitative recognition experiments, following the Setting 2 protocol of [67], [7], we only use the faces with the natural expression. 200 subjects are used for training and the remaining 137 subjects are used for testing. For the testing set, the first face of each subject is used as the gallery and the other faces are used as probes. The setting



Fig. 8: Synthesis results (1024×1024 resolution) on the MVF-HQ database. The lower right corner is the input face.

protocol of our qualitative experiments is mainly based on the above Setting 2 protocol. The difference is that, besides the natural expression, we also use the other 5 expressions for expression editing. All images are aligned and cropped to 128×128 resolution.

RaFD is a high-resolution face database for expression analysis. It consists of 8 expressions (see Fig. 12) and 5 camera angles ($\pm 90^\circ$, $\pm 45^\circ$, and 0°). Furthermore, each identity also contains three different gazed directions (left, frontal, and right). We randomly select 10 identities as the testing set and use the remaining identities as the training set. Each image is aligned and cropped to 512×512 resolution.

CelebA-HQ is a high-quality version of the CelebA database [61], containing abundant attributes but with low image quality. As shown in Fig. 2, most of the faces in CelebA-HQ are nearly frontal, making it hard to edit large pose faces. In order to enrich the poses of CelebA-HQ, we utilize a 3D model [68] to create profiles. In addition, given that the created profiles have many artifacts, as shown in Fig. 15, we only conduct face frontalization experiments, *i.e.* rotating profiles

to the frontal view. We randomly choose 3,000 images as the testing set and use the remaining images as the training set. Each image is resized to 512×512 resolution. Furthermore, considering the abundant pose variants in the CelebA database, we also verify the effectiveness of our method on it. 2000 images are selected as the testing set. Each image is aligned and cropped to 128×128 resolution.

MVF-HQ is the newly built high-resolution multi-view face database, the details of which are reported in Section IV. In experiments, we randomly select 336 identities as the training set and the remaining 143 identities as the testing set. There are no identity overlaps between training and testing. In addition, due to the limited GPU memory, we only conduct experiments at 512×512 and 1024×1024 resolutions.

Experimental Settings. For the boundary image, we use an open-source toolkit [69] to automatically detect 68 landmarks, and manually check and revise the landmarks with extreme poses ($\pm 60^\circ$, $\pm 75^\circ$, and $\pm 90^\circ$) to ensure the accuracy. The adjacent landmarks are connected to get the boundary image that is in RGB format, as shown in Fig. 1. The boundary




Fig. 9: Continuous pose (the first row) and expression (the second row) synthesis on the MultiPIE database. Only the poses and expressions with red boxes are in the database. Please zoom in for details.




Fig. 10: Visualization comparisons between our method and CAPG-GAN [7] on the MultiPIE database. For fair comparison, the two methods adopt same boundary images as geometry guidance, the same discriminator, and the same losses except for the feature threshold loss.

image mainly consists of five facial components, including eyebrows, eyes, nose, mouth, and jaw. These components can clearly present the pose, the expression, and the shape of faces. The pose vectors are directly calculated according to the detected facial landmarks. The expression vectors are denoted by Action Units (AUs) [70], which are collected by an open-source toolkit [71]. The pose estimator F_p and the expression estimator F_e in Section III-A are pre-trained on the above four databases as well as the large-scale in-the-wild database CelebA [61]. The dimensionalities of the latent vector z , the pose vector p , and the expression vector e in Eq. (2) are 128, 3, and 17, respectively. Our discriminator is the same as [10]. The parameters λ_1 , α_1 , α_2 , and α_3 in Section III-C are set to 0.1, 0.01, 50, and 0.02, respectively. The parameter m in Eq. (3) is set to 7. Adam [72] ($\beta_1 = 0.9$, $\beta_2 = 0.999$) is adopted as the optimizer with a fixed learning rate 0.0002. The high-resolution experiments on MVF-HQ are conducted on 8 NVIDIA Titan RTX GPUs with 24GB memory. Training takes about 12 days for 1024×1024 resolution and about 7 days for 512×512 resolution.

B. Qualitative Experiments

Experimental Results on the MultiPIE database. According to the given conditional vectors, our method can render an input face to the corresponding poses and expressions. Another state-of-the-art work to tackle the similar task is CAPG-GAN [7], which rotates a face to the target pose by controlling 5 facial landmarks. For fair comparisons, the 5 facial landmarks are replaced with the same boundary images as our method. Moreover, the discriminator and the losses



Fig. 11: Synthesis results on the MultiPIE database. The boundary is generated in our boundary prediction stage.

(except for the feature threshold loss) are also unified. The comparison results between our method and CAPG-GAN on the MultiPIE database are shown in Fig. 10. CAPG-GAN concatenates the input faces and the target facial landmarks in the image space, and then feeds them into the generator. As mentioned in Section I, such a concatenation manner cannot hold the textures well, which is also embodied in Fig. 10. For example, the details of eyes are somewhat blurry. On the contrary, the synthesized images of our method are closer to the ground truth in terms of both structures and textures. We owe the superiority of our method over CAPG-GAN to the effective disentanglement in the latent space. In addition, for CAPG-GAN, the target landmarks are manually given if they want to manipulate the input faces. Such a manual manner is difficult and time-consuming. Differently, our method can flexibly generate desired boundary images with the specified pose and expression vectors, as shown in Fig. 11. Another advantage of our method over CAPG-GAN is that, besides rendering poses, we can also edit facial expressions (see Fig. 11). We observe that the structures of the synthesized faces are consistent with the boundary images. At the same time, the textures of the synthesized images are preserved well, even under the extreme case.

Fig. 9 further presents the continuous pose and expression synthesis. All synthesized poses and expressions except for the three images with red boxes are unseen in the training stage. The first row displays the results of the continuous rotation from 18.75° to 56.25° . The angle interval of the synthesized images is 3.75° , while the angle interval in the MultiPIE database is 15° . In addition, the second row shows the continuous expression variation from neutral to screaming. The synthesized expressions are quite vivid. The second person with neutral expression begins with opening the mouth

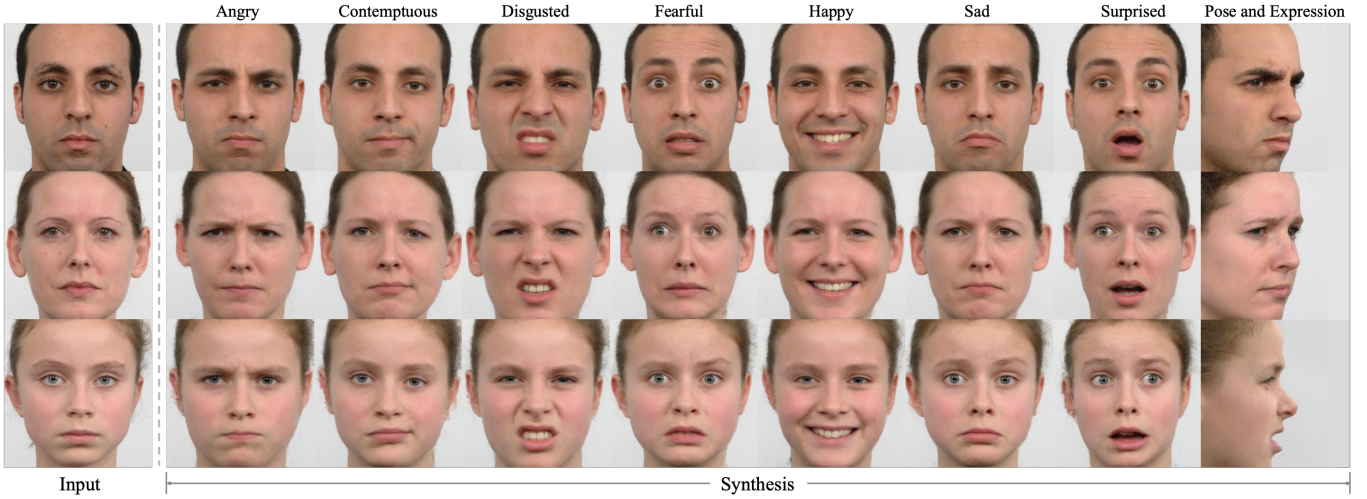


Fig. 12: Facial expression and pose synthesis (512×512 resolution) on the RaFD database. The first column is the input, and the remaining columns are synthesis results with different expressions and poses.




Fig. 13: Visualization comparisons (512×512 resolution) with pix2pixHD [10] on the RaFD database (the first row) and the MVF-HQ database (the second row).

gradually. Then, the eyebrows are raising and the eyes begin to close. Ultimately, he makes a scream expression. The results of the continuous variation demonstrate the generalization ability of our method.

Experimental Results on the RaFD database. We compare our method with pix2pixHD [10], which is a state-of-the-art high-resolution conditional image-to-image translation method. The first row of Fig. 13 plots the expression manipulation results. We observe that, although this task only needs to make slight facial changes, pix2pixHD fails to maintain structures and textures well. The reason may be that pix2pixHD does not disentangle structures and textures. Contrastively, our method gets much better synthesis results. Moreover, in Fig. 12, we also show the synthesis results of angry, contemptuous, disgusted, fearful, happy, sad, and surprised expressions, respectively. Each synthesized expression is vivid and matches the target expression label. Besides, in the last column of Fig. 12, we also present the high-quality synthesis results of the joint pose and expression variation. The high quality of the synthesized images indicates the good performance of our method.

Experimental Results on the MVF-HQ database. The




Fig. 14: Visualization comparisons (512×512 resolution) between our method and CAPG-GAN on the MVF-HQ database. Please zoom in for details.

second row of Fig. 13 plots the visualization comparisons between our method and pix2pixHD. For the synthesized image of pix2pixHD, the local structures, such as the outline of eyes, are unclear, and the facial textures are blurry. Conversely, our synthesized image has clear structures and realistic textures. Furthermore, we discover that the result of our method on the MVF-HQ database is much better than that on the RaFD database. This is because that the MVF-HQ database contains more data (about 120k images) than the RaFD database (about 8k images). It is rather challenging to synthesize high-resolution results with limited training images.

Fig. 14 shows the comparison results between our method and CAPG-GAN on the MVF-HQ database. Same as the experiments on the MultiPIE database, we also unify the the geometry guidance, the discriminator, and the losses. We can see that our method displays significant advantages than CAPG-GAN in such high-resolution cases. Compared with our results, the synthesis results of CAPG-GAN have neither clear local structures (e.g. the structure of the eyes) nor realistic textures, especially at large poses. These results demonstrate that it is really infeasible to directly apply the existing low-resolution method to the high-resolution extreme face manipulation.

In Fig. 8, we present more synthesis results (1024×1024 resolution) of different poses and expressions on the MVF-



Fig. 15: Synthesis results (512×512 resolution) on the CelebA-HQ database. The lower right corner shows the created profile.




Fig. 16: Synthesis results on the CelebA database. (a) Diverse face manipulation results. (b) Multiple manipulations of one input image.

HQ database. We observe that our method obtains satisfactory results, such as recovering the unseen ears of the third face in the first row. More synthesized details, such as the double eyelids of the first face image in the third row, demonstrate the superiority of our method.

Experimental Results on the CelebA-HQ database. All the above face manipulation databases are in the controlled environment. In order to explore the expansibility of our method under the in-the-wild situation, we conduct experiments on the CelebA-HQ database. As stated in Section V-A, most of the images in CelebA-HQ are nearly frontal. The profiles are created by a 3D model and have too many artifacts. In this case, we only conduct face frontalization experiments. Due to the effects of uncontrolled variants, such as diverse illuminations and backgrounds, it is challenging to perform high-resolution face frontalization under the in-the-wild setting. Fig. 15 shows the results of the synthesized frontal faces. Although the created profiles have massive artifacts and lose many facial textures, our method successfully eliminates the artifacts and completes the lost textures.

We further perform experiments on the original version of the CelebA-HQ database, *i.e.* the CelebA database [61] that has abundant poses but with lower image quality. Fig. 16 (a) plots the manipulation results with diverse poses, from which we observe that our method has the ability to rotate extreme poses, such as the first set of images. Nevertheless, since the training set of the CelebA database has massive low-quality images, the synthesized images inevitably appear some artifacts. For example, there are some artifacts in the outlines of the two profile faces in the bottom left. Fig. 16 (b) presents the results of multiple manipulations of one input face. The above visualization results demonstrate the superior performance of our method in the uncontrolled environment.

C. Quantitative Experiments

In this subsection, we quantitatively evaluate the identity preserving property and the synthesis quality of our method. As shown in Fig. 11, our method effectively recovers structures and textures from the profiles, which can be used to improve the performance of face recognition under large poses [7]. Therefore, we compare the recognition accuracy of our method with that of the state-of-the-art face frontalization methods, including DA-GAN [31], 3D-PIM [73], CAPG-GAN [7], PIM [74], TP-GAN [6], FF-GAN [75], and DR-GAN [29] on the MultiPIE Setting 2 protocol. The probe set consists of profiles with various views and the gallery set only contains one frontal face per subject. The profiles in the probe set are frontalized by the above methods, and the pre-trained LightCNN [64] is used to extract features. Cosine distances are calculated as the similarities to obtain the Rank-1 accuracies, the comparisons of which are tabulated in Table II. ‘LightCNN’ means evaluation on the original profiles via the pre-trained LightCNN model. ‘Ours’ means calculating the Rank-1 accuracy on the synthesized frontal faces by the same LightCNN model. We observe that as the face angle increases, the accuracies of all the methods drop gradually. The degradation may be caused by the loss of facial appearance of profiles. Furthermore, the Rank-1 accuracies of all the methods are comparable under small poses ($\pm 15^\circ$, $\pm 30^\circ$, and $\pm 45^\circ$). But at the larger poses, the superiority of our method




Fig. 17: Visualization results (512×512 resolution) of the ablation study: comparisons between our method and its five variants. Please zoom in for details.

TABLE II: Comparisons of Rank-1 recognition rates (%) across views under the MultiPIE Setting 2.

Method	$\pm 15^\circ$	$\pm 30^\circ$	$\pm 45^\circ$	$\pm 60^\circ$	$\pm 75^\circ$	$\pm 90^\circ$
DR-GAN [29]	94.0	90.1	86.2	83.2	-	-
FF-GAN [75]	94.6	92.5	89.7	85.2	77.2	61.2
TP-GAN [6]	98.6	98.0	95.3	87.7	77.4	64.6
PIM [74]	99.3	99.0	98.5	98.1	95.0	86.5
CAPG-GAN [7]	99.8	99.5	97.3	90.6	83.0	66.0
3D-PIM [73]	99.6	99.4	98.8	98.3	95.2	86.7
DA-GAN [31]	99.9	99.8	99.1	97.2	93.2	81.5
LightCNN [64]	98.5	97.3	92.1	62.0	24.1	5.5
Ours	99.9	99.9	99.4	98.7	96.3	87.4

TABLE III: Comparisons of FID (lower is better) and Rank-1 recognition rates (%) on the MVF-HQ database.

Method	FID	$\pm 15^\circ$	$\pm 30^\circ$	$\pm 45^\circ$	$\pm 60^\circ$	$\pm 75^\circ$	$\pm 90^\circ$
CAPG-GAN [7]	36.68	99.9	99.9	99.4	95.2	82.1	53.7
pix2pixHD [10]	45.62	98.5	97.5	93.3	86.1	67.4	39.0
pix2pixHD+boundary	43.37	99.2	98.1	94.5	87.3	68.7	40.3
LightCNN [64]	-	100	100	99.6	95.7	65.2	23.9
Ours	12.94	100	100	99.6	96.5	84.6	60.4

is obvious. Particularly, our method significantly improves the accuracy under the challenging $\pm 90^\circ$, and obtains the best performance compared with other state-of-the-art methods.

The settings of the probe and the gallery of our MVF-HQ database are analogous to the MultiPIE database. Table III shows the comparison results of our method against other state-of-the-art methods, including pix2pixHD (‘pix2pixHD+boundary’ denotes concatenating the face image with a boundary image, which is generated by the boundary prediction stage, as the input of pix2pixHD) and CAPG-GAN. It is obvious that our method outperforms its competitors by a large margin under the extreme poses ($\pm 75^\circ$ and $\pm 90^\circ$). We owe the significant improvements to the introduced proxy network and the feature threshold loss, which disentangle structures and textures in the latent space. The quantitative recognition results are consistent with the qualitative visualization results in Fig. 10 and Fig. 13. The recognition results on MultiPIE and MVF-HQ prove that our method can effectively improve the recognition performance under large poses.

Furthermore, in order to evaluate the quality of the syn-

TABLE IV: Quantitative results of the ablation study on the MVF-HQ database.

Method	FID	$\pm 15^\circ$	$\pm 30^\circ$	$\pm 45^\circ$	$\pm 60^\circ$	$\pm 75^\circ$	$\pm 90^\circ$
w/o \mathcal{L}_{thr}	36.71	96.2	90.7	77.3	67.1	55.6	47.6
w/o \mathcal{L}_{adv}	53.92	98.5	96.1	87.9	80.3	68.4	50.3
w/o \mathcal{L}_{ip}	21.34	88.7	81.5	75.3	65.4	50.7	45.5
w/o $\mathcal{L}_{pix-mul}$	34.34	98.3	94.3	82.4	81.6	71.3	55.7
Ours	12.94	100	100	99.6	96.5	84.6	60.4

thesized images, we compare the Fréchet Inception Distance (FID) results [76] with CAPG-GAN and pix2pixHD. FID is used to measure the distance between the real faces and the synthesized faces. The results in Table III qualitatively reveal the high-quality synthesis character of our method.

D. Experimental Analysis

Ablation Study. In this part, we investigate the roles of the five loss functions in our method, including the conditional regression loss \mathcal{L}_{reg} in Eq. (2), the feature threshold loss \mathcal{L}_{thr} in Eq. (3), the multi-scale pixel-wise $\mathcal{L}_{pix-mul}$ loss in Eq. (4), the conditional adversarial loss \mathcal{L}_{adv} in Eq. (5), and the identity preserving loss \mathcal{L}_{ip} in Eq. (6). Both qualitative and quantitative experimental results are reported for better comparisons.

Fig. 17 shows the qualitative visualization results of our method and its five variants. We discover that without \mathcal{L}_{reg} , the generated boundary images are unsatisfactory, when the given conditional vectors are not completely consistent with those in the databases. The outlines of many facial components, such as the nose and jaw, are unclear, resulting in incomplete synthesized faces. It demonstrates the effectiveness of the conditional regression loss \mathcal{L}_{reg} . Without the feature threshold loss \mathcal{L}_{thr} , the local structures, e.g. the eyes and nose, are unclear and the textures are blurry, indicating the effectiveness of the disentanglement. Moreover, the parameter m in \mathcal{L}_{thr} has non-negligible impacts on the synthesis results, which will be discussed in the following part. Without the multi-scale pixel-wise loss $\mathcal{L}_{pix-mul}$ (only utilizing one scale pixel-wise loss), the global structure is clear but the local textures, e.g. the teeth, are blurry. Hence, the multi-scale pixel-wise loss $\mathcal{L}_{pix-mul}$ contributes to the recovery of texture details. Without \mathcal{L}_{adv} , there are many artifacts in the synthesized images, revealing the validity of the multi-scale conditional




Fig. 18: Visualization results (512×512 resolution) of the parameter analysis of m in Eq. (3). Please zoom in for details.



Fig. 19: Synthesis results of the cross-database experiment (trained on MultiPIE and tested on MVF-HQ).

TABLE V: FID and Rank-1 accuracies ($\pm 90^\circ$) on the MVF-HQ database under different values of m in Eq. (3).

m	10	9	8	7	6	5
FID	38.65	23.62	17.63	12.94	20.35	35.43
Rank-1	48.2	54.7	57.3	60.4	56.6	50.7

adversarial loss. Without \mathcal{L}_{ip} , the local textures, such as the beard, are somewhat light. Thus, the identity preserving loss may be beneficial to the enhancement of local textures.

Table IV further tabulates the FID and Rank-1 results of different variants of our method. We observe that the FID will increase and the Rank-1 will decrease if one loss is not adopted, which are consistent with the qualitative visualization results in Fig. 17. These qualitative and quantitative results verify that each component in our method is essential for the high-resolution extreme face manipulation.

Cross-Database Experiments. Fig. 19 plots the results of cross-database experiments. That is, the model is trained on the MultiPIE database and tested on the MVF-HQ database. There is a large domain gap between the two databases, because of the differences in the acquisition equipment, participants, backgrounds, etc. Although the synthesized images on the MVF-HQ database inevitably bring some domain information of the MultiPIE database, such as the backgrounds, our method successfully manipulates the input faces. The results of cross-database experiments further demonstrate the generalization ability of our method.

Parameter Analysis. As mentioned in Section III-B, the value of the parameter m in the feature threshold loss \mathcal{L}_{thr} (Eq. (3)) has non-negligible effects on the disentanglement. In Fig. 18, we plot the visualization results under different values. We observe that when the value of m is too large, the synthesized faces are blurry due to the weak disentanglement.

On the contrary, when the value of m is too small, the textures of the synthesized faces will be somewhat lost because of the too compact texture features. The best results are obtained when $m = 7$. In addition, the quantitative FID and Rank-1 results are listed in Table V. The results of these quantitative indicators are consistent with the visualization results in Fig. 18. When m equals 7, we obtain the minimum FID and the highest accuracy.

VI. CONCLUSION

This paper has developed a stage-wise framework for high-resolution extreme face manipulation. It simplifies the face manipulation into two correlated stages: a boundary prediction stage and a disentangled face synthesis stage. The first stage predicts the boundary image of the target face in a semi-supervised way. The second stage utilizes the predicted boundary to perform realistic face synthesis. A proxy network and a feature threshold loss are introduced to disentangle structures and textures in the latent space. Furthermore, a new high-resolution MVF-HQ database has been created, which consists of 120,283 images at 6000×4000 resolution from 479 identities. It is much larger in scale and much higher in resolution than publicly available high-resolution face manipulation databases. In the future, we will continue to collect more data to further enrich MVF-HQ. Extensive experiments on four databases show that our method significantly pushes forward the advance of extreme face manipulation.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NeurIPS*, 2014.
- [2] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *ICML*, 2017.
- [3] 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 *CVPR*, 2018.

- [4] A. Pumarola, A. Agudo, A. M. Martinez, A. Sanfeliu, and F. Moreno-Noguer, “Ganimation: Anatomically-aware facial animation from a single image,” in *ECCV*, 2018.
- [5] Y. Shen, P. Luo, J. Yan, X. Wang, and X. Tang, “Faceid-gan: Learning a symmetry three-player gan for identity-preserving face synthesis,” in *CVPR*, 2018.
- [6] 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 *ICCV*, 2017.
- [7] Y. Hu, X. Wu, B. Yu, R. He, and Z. Sun, “Pose-guided photorealistic face rotation,” in *CVPR*, 2018.
- [8] F. Zhang, T. Zhang, Q. Mao, and C. Xu, “Joint pose and expression modeling for facial expression recognition,” in *CVPR*, 2018.
- [9] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” in *CVPR*, 2017.
- [10] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, “High-resolution image synthesis and semantic manipulation with conditional gans,” in *CVPR*, 2018.
- [11] L. Ma, X. Jia, Q. Sun, B. Schiele, T. Tuytelaars, and L. Van Gool, “Pose guided person image generation,” in *NeurIPS*, 2017.
- [12] Y. Jo and J. Park, “Sc-fegan: Face editing generative adversarial network with user’s sketch and color,” in *ICCV*, 2019.
- [13] Z. Shu, M. Sahasrabudhe, R. Alp Guler, D. Samaras, N. Paragios, and I. Kokkinos, “Deforming autoencoders: Unsupervised disentangling of shape and appearance,” in *ECCV*, 2018.
- [14] J. Bao, D. Chen, F. Wen, H. Li, and G. Hua, “Towards open-set identity preserving face synthesis,” in *CVPR*, 2018.
- [15] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, “Multi-pie,” *Image and Vision Computing*, vol. 28, no. 5, pp. 807–813, 2010.
- [16] T. Karras, T. Aila, S. Laine, and J. Lehtinen, “Progressive growing of gans for improved quality, stability, and variation,” in *ICLR*, 2018.
- [17] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” in *CVPR*, 2019.
- [18] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, “Analyzing and improving the image quality of stylegan,” in *CVPR*, 2020.
- [19] 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*, vol. 24, no. 8, pp. 1377–1388, 2010.
- [20] V. Blanz, C. Basso, T. Poggio, and T. Vetter, “Reanimating faces in images and video,” in *Computer Graphics Forum*, 2003.
- [21] Y. Wang, L. Zhang, Z. Liu, G. Hua, Z. Wen, Z. Zhang, and D. Samaras, “Face relighting from a single image under arbitrary unknown lighting conditions,” *TPAMI*, vol. 31, no. 11, pp. 1968–1984, 2008.
- [22] F. Yang, J. Wang, E. Shechtman, L. Bourdev, and D. Metaxas, “Expression flow for 3d-aware face component transfer,” in *ACM SIGGRAPH*, 2011.
- [23] C. Cao, Y. Weng, S. Zhou, Y. Tong, and K. Zhou, “Facewarehouse: A 3d facial expression database for visual computing,” *TVCG*, vol. 20, no. 3, pp. 413–425, 2013.
- [24] I. Kemelmacher-Shlizerman, S. Suwajanakorn, and S. M. Seitz, “Illumination-aware age progression,” in *CVPR*, 2014.
- [25] Y.-C. Chen, X. Shen, Z. Lin, X. Lu, I. Pao, and J. Jia, “Semantic component decomposition for face attribute manipulation,” in *CVPR*, 2019.
- [26] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, G. Liu, A. Tao, J. Kautz, and B. Catanzaro, “Video-to-video synthesis,” in *NeurIPS*, 2018.
- [27] E. Zakhharov, A. Shysheya, E. Burkov, and V. Lempitsky, “Few-shot adversarial learning of realistic neural talking head models,” in *ICCV*, 2019.
- [28] K. Songsri-in and S. Zafeiriou, “Face video generation from a single image and landmarks,” *arXiv:1904.11521*, 2019.
- [29] L. Tran, X. Yin, and X. Liu, “Disentangled representation learning gan for pose-invariant face recognition,” in *CVPR*, 2017.
- [30] S. Tulyakov, M.-Y. Liu, X. Yang, and J. Kautz, “Mocogan: Decomposing motion and content for video generation,” in *CVPR*, 2018.
- [31] Y. Yin, S. Jiang, J. P. Robinson, and Y. Fu, “Dual-attention gan for large-pose face frontalization,” in *FG*, 2020.
- [32] J. Deng, S. Cheng, N. Xue, Y. Zhou, and S. Zafeiriou, “Uv-gan: Adversarial facial uv map completion for pose-invariant face recognition,” in *CVPR*, 2018.
- [33] S. Qian, K.-Y. Lin, W. Wu, Y. Liu, Q. Wang, F. Shen, C. Qian, and R. He, “Make a face: Towards arbitrary high fidelity face manipulation,” in *ICCV*, 2019.
- [34] H. Chang, J. Lu, F. Yu, and A. Finkelstein, “Pairedcyclegan: Asymmetric style transfer for applying and removing makeup,” in *CVPR*, 2018.
- [35] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, “Generative image inpainting with contextual attention,” in *CVPR*, 2018.
- [36] K. Nazeri, E. Ng, T. Joseph, F. Qureshi, and M. Ebrahimi, “Edge-connect: Generative image inpainting with adversarial edge learning,” *arXiv:1901.00212*, 2019.
- [37] 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 *CVPR*, 2017.
- [38] B. Duan, C. Fu, Y. Li, X. Song, and R. He, “Pose agnostic cross-spectral hallucination via disentangling independent factors,” in *CVPR*, 2020.
- [39] Y. Choi, Y. Uh, J. Yoo, and J.-W. Ha, “Stargan v2: Diverse image synthesis for multiple domains,” in *CVPR*, 2020.
- [40] Y. Shen, J. Gu, X. Tang, and B. Zhou, “Interpreting the latent space of gans for semantic face editing,” in *CVPR*, 2020.
- [41] F. Zhang, T. Zhang, Q. Mao, and C. Xu, “A unified deep model for joint facial expression recognition, face synthesis, and face alignment,” *TIP*, vol. 29, pp. 6574–6589, 2020.
- [42] V. Blanz and T. Vetter, “A morphable model for the synthesis of 3d faces,” in *SIGGRAPH*, 1999.
- [43] K. Scherbaum, T. Ritschel, M. Hullin, T. Thormählen, V. Blanz, and H.-P. Seidel, “Computer-suggested facial makeup,” in *Computer Graphics Forum*, 2011.
- [44] J. Thies, M. Zollhöfer, M. Nießner, L. Valgaerts, M. Stamminger, and C. Theobalt, “Real-time expression transfer for facial reenactment,” *TOG*, vol. 34, no. 6, pp. 183–1, 2015.
- [45] J. Thies, M. Zollhofer, M. Stamminger, C. Theobalt, and M. Nießner, “Face2face: Real-time face capture and reenactment of rgb videos,” in *CVPR*, 2016.
- [46] S. Suwajanakorn, S. M. Seitz, and I. Kemelmacher-Shlizerman, “Synthesizing obama: learning lip sync from audio,” *TOG*, vol. 36, no. 4, pp. 1–13, 2017.
- [47] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *ICLR*, 2014.
- [48] A. Van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves *et al.*, “Conditional image generation with pixelcnn decoders,” in *NeurIPS*, 2016.
- [49] Y. Li, K. Swersky, and R. Zemel, “Generative moment matching networks,” in *ICML*, 2015.
- [50] L. Dinh, J. Sohl-Dickstein, and S. Bengio, “Density estimation using real nvp,” in *ICLR*, 2017.
- [51] C. Fu, X. Wu, Y. Hu, H. Huang, and R. He, “Dual variational generation for low-shot heterogeneous face recognition,” in *NeurIPS*, 2019.
- [52] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, “Spectral normalization for generative adversarial networks,” in *ICLR*, 2018.
- [53] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, “Semantic image synthesis with spatially-adaptive normalization,” in *CVPR*, 2019.
- [54] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” in *ICLR*, 2016.
- [55] H. Huang, Z. Li, R. He, Z. Sun, and T. Tan, “Introvae: Introspective variational autoencoders for photographic image synthesis,” in *NeurIPS*, 2018.
- [56] A. van den Oord, O. Vinyals *et al.*, “Neural discrete representation learning,” in *NeurIPS*, 2017.
- [57] A. Razavi, A. van den Oord, and O. Vinyals, “Generating diverse high-fidelity images with vq-vae-2,” in *NeurIPS*, 2019.
- [58] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *ICCV*, 2017.
- [59] A. Brock, J. Donahue, and K. Simonyan, “Large scale gan training for high fidelity natural image synthesis,” in *ICLR*, 2019.
- [60] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *CVPR*, 2009.
- [61] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep learning face attributes in the wild,” in *ICCV*, 2015.
- [62] W. Gao, B. Cao, S. Shan, X. Chen, D. Zhou, X. Zhang, and D. Zhao, “The cas-peal large-scale chinese face database and baseline evaluations,” *IEEE SMC*, vol. 38, no. 1, pp. 149–161, 2007.
- [63] G. Santos, P. T. Fiadeiro, and H. Proença, “Biohdd: a dataset for studying biometric identification on heavily degraded data,” *IET Biometrics*, vol. 4, no. 1, pp. 1–9, 2015.
- [64] X. Wu, R. He, Z. Sun, and T. Tan, “A light cnn for deep face representation with noisy labels,” *TIFS*, vol. 13, no. 11, pp. 2884–2896, 2018.
- [65] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, “Ms-celeb-1m: A dataset and benchmark for large-scale face recognition,” in *ECCV*, 2016.
- [66] T. Sim, S. Baker, and M. Bsat, “The cmu pose, illumination, and expression (pie) database,” in *FG*, 2002.

- [67] J. Yim, H. Jung, B. Yoo, C. Choi, D. Park, and J. Kim, "Rotating your face using multi-task deep neural network," in *CVPR*, 2015.
- [68] X. Zhu, Z. Lei, X. Liu, H. Shi, and S. Z. Li, "Face alignment across large poses: A 3d solution," in *CVPR*, 2016.
- [69] A. Bulat and G. Tzimiropoulos, "How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks)," in *ICCV*, 2017.
- [70] E. Friesen and P. Ekman, "Facial action coding system: a technique for the measurement of facial movement," *Palo Alto*, vol. 3, 1978.
- [71] T. Baltusaitis, A. Zadeh, Y. C. Lim, and L.-P. Morency, "Openface 2.0: Facial behavior analysis toolkit," in *FG*, 2018.
- [72] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2014.
- [73] J. Zhao, L. Xiong, Y. Cheng, Y. Cheng, J. Li, L. Zhou, Y. Xu, J. Karlekar, S. Pranata, S. Shen, J. Xing, S. Yan, and J. Feng, "3d-aided deep pose-invariant face recognition," in *IJCAI*, 2018.
- [74] 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 *CVPR*, 2018.
- [75] X. Yin, X. Yu, K. Sohn, X. Liu, and M. Chandraker, "Towards large-pose face frontalization in the wild," in *ICCV*, 2017.
- [76] 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 *NeurIPS*, 2017.