

---

# Learning to Decouple the Lights for 3D Face Texture Modeling

---

Tianxin Huang<sup>1</sup>    Zhenyu Zhang<sup>2</sup>    Ying Tai<sup>2</sup>    Gim Hee Lee<sup>1</sup>

<sup>1</sup>School of Computing, National University of Singapore

<sup>2</sup>Nanjing University

huangtx@nus.edu.sg, gimhee.lee@nus.edu.sg

## Abstract

Existing research has made impressive strides in reconstructing human facial shapes and textures from images with well-illuminated faces and minimal external occlusions. Nevertheless, it remains challenging to recover accurate facial textures from scenarios with complicated illumination affected by external occlusions, e.g. a face that is partially obscured by items such as a hat. Existing works based on the assumption of single and uniform illumination cannot correctly process these data. In this work, we introduce a novel approach to model 3D facial textures under such unnatural illumination. Instead of assuming single illumination, our framework learns to imitate the unnatural illumination as a composition of multiple separate light conditions combined with learned neural representations, named Light Decoupling. According to experiments on both single images and video sequences, we demonstrate the effectiveness of our approach in modeling facial textures under challenging illumination affected by occlusions. Our codes would be open sourced at <https://github.com/Tianxinhuang/DeFace.git>.

## 1 Introduction

Recently, 3D face reconstruction has made significant progress [4, 12, 13, 2, 30, 22] with the rapid development of digital human and meta-universe technologies. These techniques have become increasingly proficient in recovering details of face shapes and textures from single images or video sequences. Their performances have been particularly commendable on well-illuminated faces.

Despite the advancements, the real world presents more complex scenarios. As shown in the blue and red rectangled regions of input images in Fig. 1, self occlusions from facial parts such as the nose, or external occlusions such as hats or hairs introduce illumination changes and produce shadows in certain regions. Although recent works [16, 2, 21, 23, 22] replace the linear models [6, 4] with non-linear GAN-inversion-based textured models to greatly enhance the quality of textures and complete occluded facial areas, they still rely on the assumption that the environment illumination is single and uniform. As the illumination affected by self and external occlusions is unnatural, it may deviate drastically from the single and uniform illumination assumption, which would bring unexpected influence to the textures modelled with existing methods.

As shown in Fig. 1-a, methods [16, 10, 2] based on the diffuse-only texture and mesh render [10] bake both shadows caused by self occlusion of the nose and external occlusion of the hat onto the texture, where the re-rendered output also appears unrealistic and lacks authenticity. Recent works proposed by Dib *et al.* [12, 13, 11] combine local reflectance texture model incorporating diffuse, specular, and roughness albedos with ray-tracing render [26] to implicitly model influence of the self occlusion. As shown in Fig. 1-b, these methods can create more realistic rendered output and avoid the shadows of self occlusion on the recovered texture in the blue rectangle, while the influence from external occlusion in the red rectangle remains unresolved.

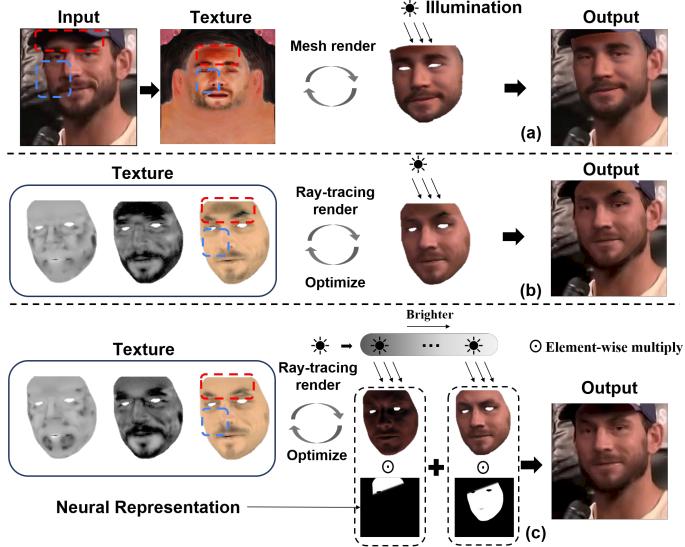


Figure 1: Blue and red rectangles mark regions affected by self and external occlusions, respectively. (a) Texture modeling with diffuse-only texture map. (b) Texture modeling based on diffuse, specular, and roughness albedos from local reflectance model [12], while optimizing with ray-tracing render. (c) Our method learns neural representations to decouple the original illumination into multiple light conditions, where the influence from external occlusions can be modeled as one of the conditions. White and black regions in the masks denote 1 and 0, respectively.

To mitigate the above-mentioned issues, we propose a face texture modeling framework for faces under complicated and unnatural illumination affected by external occlusions. Given the limited physical information about external occlusions presented in a single facial image or video sequence, accurately modeling how occlusion impacts illumination becomes nearly unattainable. Regarding this challenge, we propose to use *multiple separate light conditions* to imitate the local illumination of different facial areas after being affected by external occlusions. As shown in Fig. 1-c, our method can model the influence from external occlusion as one of the light conditions and eliminate its effect on the recovered texture as highlighted in the red rectangle, where the rendered output is also more realistic and closer to the input image especially on the forehead and eyes. Specifically, we use a spatial-temporal neural representation to predict masks for different light conditions, which are used to combine rendered results under multiple light conditions into final output. The number of light masks is adaptively adjusted during optimization. As external occlusions may directly cover parts of the face, we also introduce another neural representation to provide a continuous prediction for the available face region. Furthermore, our approach incorporates realistic constraints by introducing priors from the statistical model and pre-trained face perceptual models to ensure our extracted textures construct lifelike human faces.

Our contribution in this work can be summarized as:

- We present a new face texture modeling framework by learning to decouple the environmental illumination into multiple light conditions with neural representations;
- We propose realistic constraints to help improve the authenticity of texture by introducing human face priors;
- Extensive experiments on images and videos confirm that our method delivers more realistic renderings and improved facial textures compared to existing approaches under challenging illumination affected by occlusions.

## 2 Related Works

### 2.1 3D Morphable Model

The 3D Morphable Model (3DMM) [4, 15, 32, 6] is a widely employed linear statistical framework used for modeling the geometry and texture of faces. It is constructed based on aligned face images

using Principal Component Analysis (PCA). Within the 3DMM framework, the coordinates and colors of facial mesh vertices are computed through a linear combination of numerical parameters. These parameters can be estimated to enable face reconstruction through optimization-based fitting [4, 6] or learning-based approaches [10, 30, 43, 42], with the goal of minimizing the disparities between the rendered and original face images. Traditional 3DMMs, although effective, are limited by their linear PCA basis. Recent advancements [16, 21, 23, 44] have introduced non-linear basis by incorporating pre-trained mesh or texture decoder networks. In these methods, latent codes for these decoders are optimized to align with the original face images. This approach significantly enhances the robustness of the model when faced with issues such as occlusions and missing textures. However, such methods demand access to large, well-preprocessed datasets for effective pre-training.

## 2.2 Face Texture and Illumination Modeling

Texture and illumination modeling plays a pivotal role in 3D face reconstruction, directly impacting the color rendering in generated images. Accurate modeling for 3D face textures and environmental illumination is vital for subsequent applications such as face relighting [36, 18, 31] or animation [38, 30]. To achieve high-resolution rendering, recent approaches [9, 21, 16, 3] usually opt for UV mapping to model facial texture, as opposed to the earlier methods that focused on vertex colors [32]. The UV map can either be initialized with a PCA basis [15] or estimated through the use of a pre-trained non-linear texture decoder [22, 2, 16]. There are two commonly used models to estimate human skin reflectance for the texture modeling framework: the Lambertian model [1] and the Blinn-Phong model [5]. The Lambertian model is more computationally efficient, while the Blinn-Phong model produces more realistic rendering results by accounting for specular attributes in textures.

As capturing environmental illumination directly can be challenging, most existing methods assume it as primarily uniform Spherical Harmonics [33, 34]. This approach, however, tends to “bake” shadows cast by facial characters or external occlusions into the textures. Although method [45] introduce implicit texture modeling to achieve better face rendering performances, it requires multi-view images under uniform illumination for optimization, which is not appropriate for face images affected by occlusions. 2D shadow removal methods [41, 17, 28] try to eliminate the influence of shadows directly with networks. These methods have higher inference efficiency, while the size of their training set may greatly affect their performances. Notably, Dib *et al.* [12, 13, 11] have introduced techniques that implicitly model self-shadows through ray-tracing, demonstrating exceptional performance on faces with severe self occlusions. Despite these advancements, existing methods usually encounter challenges in handling shadows caused by external occlusions such as hats or hair, tending to take the shadows as part of textures.

To address these issues, we propose a novel framework to recover clear textures under challenging illumination affected by external occlusions. Our approach achieves this by decoupling the unnatural affected illumination into multiple light conditions using learned neural representations.

## 3 Our Methodology

### 3.1 Problem Definition

This work focus on the texture modeling problem of human faces under challenging environment illumination affected by external occlusions. Given an input image or video sequences  $I_{in}$  taken from affected illuminations, our task is to recover clear and accurate texture  $T$  from  $I_{in}$  and ensure that  $T$  can synthesize results  $I_{out}$  close to  $I_{in}$ .

### 3.2 Overall Illustration

As illustrated in Fig. 2, instead of making the assumption that the face is exposed to a single and uniform illumination, we propose to model  $n$  possible separate light conditions  $\gamma_1 \sim \gamma_n$  in the environment, where the masks for possible regions  $M_N$  are predicted with a neural representation  $f(\cdot)$ . Effective masks  $M_L$  and rendered faces  $I_{Rs}$  are selected from original  $M_N$  and rendered  $I_{Rn}$  during optimization with the Adaptive Condition Estimation (ACE) strategy. The final rendered face under multiple light conditions is  $I_R = \sum I_{Rs} \odot M_L$ , which is merged with the input images to construct the output  $I_{out} = I_R \odot M_o + I_{in} \odot (1 - M_o)$ . The face shape is modeled with statistical

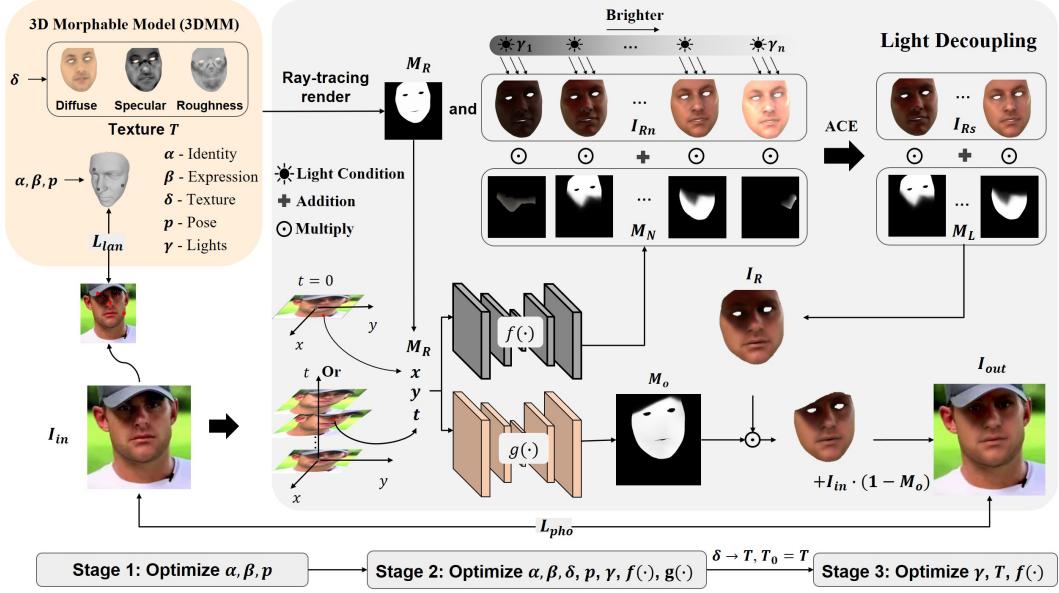


Figure 2: Illustration of our framework. The pipeline is proposed to recover texture  $T$  and 3DMM statistical coefficients  $\alpha, \beta, \delta, p, \gamma$  from the input image  $I_{in}$ . The statistical coefficient  $\delta$  is used to initialize  $T$ . Render mask  $M_R$  and Faces  $I_{Rn}$  under  $n$  light conditions  $\gamma = \{\gamma_1 \sim \gamma_n\}$  are acquired through ray-tracing rendering.  $f(\cdot)$  and  $g(\cdot)$  are neural representations predicting light masks  $M_N$  and facial region mask  $M_o$ . ACE is introduced to select effective masks  $M_L$  and rendered faces  $I_{Rs}$ .  $I_{Rs}$  are combined into  $I_R$  with  $M_L$ , where  $I_R$  is merged with surroundings in  $I_{in}$  with  $M_o$  to construct output image  $I_{out}$ .  $L_{pho}$  and  $L_{lan}$  are photometric loss and landmark loss, respectively.

coefficients  $\alpha, \beta$ , and  $p$ . Following [12], they are mainly optimized with landmark loss:

$$L_{lan} = \frac{1}{|q_{in}|} \|q_{out} - q_{in}\|_2, \quad (1)$$

where  $q_{in}$  is 2D key points detected from  $I_{in}$ .  $|q_{in}|$  is the number of points in  $q_{in}$ .  $q_{out}$  denotes projected 3D key points on 2D plane. The color-related variables such as textures and lights are mainly optimized with photometric loss:

$$L_{pho} = \frac{1}{|I_{in}|} \|I_{out} - I_{in}\|, \quad (2)$$

where  $|I_{in}|$  is defined as the number of pixels of input image  $I_{in}$ . More details are presented below.

### 3.3 Light Decoupling

**Light Condition Initialization.** Following [12, 13] using  $B = 9$  bands Spherical Harmonics (SH) and ray-tracing rendering to model the illumination under self occlusions, we use  $n$  separate SH to model  $n$  possible light conditions. The coefficients are then simply initialized as  $\gamma_i = 2 \cdot \frac{i}{n} \mathbf{I} - 1$  for  $\gamma_i \in \gamma = \{\gamma_1, \dots, \gamma_n\}$ , where  $\mathbf{I}$  is a all-one matrix with the same shape as the SH coefficient. Please note that we use each SH to imitate the local illumination after being affected by the external occlusion, instead of the global natural illumination. This means that we do not need to consider physical influence of occlusion in each SH. Instead, we optimize each SH independently to directly imitate the illumination in different face regions.

**Neural Representations for Face Segment.** To decouple the illumination into multiple light conditions, we design a pair of spatial-temporal continuous neural representations to segment the face into regions for different light conditions. As illustrated in Fig. 2, spatial positions  $x, y$  and temporal position  $t$  of pixels are normalized and embedded into a coordinates system.  $t$  is decided by the number of frames in the input image/video sequence. For the single image reconstruction,  $t$  is set as a constant 0, where it would be  $i/k$  for the  $i_{th}$  frame of a  $k$  frames video sequences. For the convenience of learning,  $(x, y)$  are normalized into  $[-1.0, 1.0]$  and  $t$  is normalized into  $[0.0, 1.0]$ .

A Multi-Layer Perceptron (MLP)  $f(\cdot)$  is then introduced to predict the probability of assignment to each light condition. Given the render mask  $M_R$  in the ray-tracing-assisted rendering, the light masks can be obtained as  $M_N = M_R \odot f(x, y, t)$ . Effective masks  $M_L$  are selected from  $M_N$  with Adaptive Condition Estimation (ACE) to provide the final segment for different light conditions.

Similarly, we use another MLP  $g(\cdot)$  to predict the probability that each pixel belongs to the face regions to avoid the influence from direct occlusion such as hat or hair. The Face mask is thus given by:  $M_o = M_R \odot g(x, y, t)$ . A pre-trained semantic segmentation network [27] is introduced for distillation to  $g(\cdot)$ , while the probability for labels of all face components such as eyes, mouth, or nose are added together to construct a classifier  $h(\cdot)$  to predict the association of a pixel at  $x, y, t$  to the face. The distillation loss  $L_{seg}$  can be simply defined as:

$$L_{seg} = \frac{1}{|I_{in}|} \|g(x, y, t) - h(x, y, t)\|_2. \quad (3)$$

Note that  $g(\cdot)$  is co-optimized together with the  $L_{pho}$  and  $L_{seg}$  to further distinguish some occlusions hard to be fitted with the 3DMM statistical model.

**Adaptive Condition Estimation (ACE).** As the complexity of illumination is mutative in different surroundings, we design a strategy to estimate the number of light conditions existing in this environment during optimization. Specifically, light masks with larger area than a pre-defined threshold  $\epsilon$  in  $M_N$  are preserved in  $M_L$ , while smaller ones are dropped with corresponding light conditions and not further optimized in later iterations. Given the number of light masks  $M_N$  same as light condition  $n$ , then  $M_N = \{M_N^i\}_{i=1}^n$ ,  $M_L = \{M_L^i\}_{i=1}^{n_L} = \{M_N^i | \frac{1}{|I_{in}|} \sum_{m \in M_N^i} m > \epsilon, i \leq n\}$ .  $n_L \leq n$  is the number of preserved light conditions.  $m$  denotes the pixel in  $M_N^i$ . We introduce an regularization  $L_{area}$  to encourage the area concentration on fewer masks:

$$L_{area} = \frac{1}{|I_{in}|} \sum \frac{1}{n} \sum_{i=1}^n e^{-\|M_N^i - \frac{\sum_{i=1}^n M_N^i}{n}\|_2} - 1, \quad (4)$$

where  $|I_{in}|$  is the number of pixels in the mask.

To ensure each light condition is mainly contained in one mask of  $M_L$ , we introduce a binary regularization  $L_{bin}$  to ensure that the predicted probability for each light condition tends to 0 or 1:

$$L_{bin} = \frac{1}{|I_{in}|} \sum \frac{1}{n_L} \sum_{i=1}^{n_L} e^{-\|M_L^i - \frac{\sum_{m \in M_L^i} m}{|I_{in}|}\|_2} - 1, \quad (5)$$

where  $m \in M_L^i$  denotes each pixel value in  $i^{\text{th}}$  light mask  $M_L$ .  $n_L$  is the number of masks in  $M_L$ .

Note that  $L_{area}$  is different from  $L_{bin}$  as it pushes the mask values  $M_N^i$  far from the mean of all masks, while  $L_{bin}$  pushes mask values away from mean of different positions in the same mask  $M_L^i$ . ACE is executed at a specific iteration  $iter_0$  to select  $M_L$  from  $M_N$ . We use  $L_{area}$  to encourage area concentration before  $iter_0$ , while using  $L_{bin}$  to push the mask to 0 or 1 after  $iter_0$ .

### 3.4 Realistic Constraints

To ensure the reconstructed texture is reasonable and lifelike, we propose global prior constraint  $L_{GP}$ , local prior constraint  $L_{LP}$ , and human prior constraint  $L_{HP}$  by introducing both priors from 3DMM statistical model and the pre-trained perceptual model.

**Global Prior Constraint.** The global prior constraint is used to ensure the consistency of overall hue between optimized texture  $T$  and initialized texture  $T^0$  from the 3DMM statistical model calculated with  $\delta$ . As the colors mainly come from diffuse albedo in  $T$ , we estimate the overall hue of  $T^0$  with a K-means algorithm on its diffuse albedo  $T_D^0$  and get a  $4 \times 4$  color matrices  $C$ . Given the diffuse albedo of Texture  $T$  as  $T_D$ , this term is given by:

$$L_{GP} = \frac{1}{|I_{in}|} \sum \min_{c \in C} \|T_D - c\|_2. \quad (6)$$

**Local Prior Constraint.** The local prior constraint is used to ensure the local smoothness of optimized texture  $T$  by constraining its local variation to be similar to  $T^0$ . Let us define the diffuse, specular, and roughness albedos of texture  $T$  as  $T_D, T_S, T_R$ , respectively. We set the local variations

---

**Algorithm 1** Training Process

---

```
1: Set the number of  $iter_0 \sim iter_3$ 
2: Stage 1:
3: for  $n = 1$  to  $iter_1$  do
4:   Calculate the Loss of Stage 1 :  $L_1 = L_{lan}$ 
5:   Update coefficients with  $L_1$ :  $\nabla_{\alpha, \beta, p} L_1$ .
6: end for
7: Stage 2:
8: for  $n = 1$  to  $iter_2$  do
9:   if  $n < iter_0$  then
10:    Calculate the Loss of Stage 2 on  $M_N$  and  $I_{Rn}$ :
11:     $L_2 = \omega_0 L_{pho} + \omega_1 L_{lan} + \omega_2 L_{seg} + \omega_3 L_{area}$ 
12:   else if  $n = iter_0$  then
13:    Select  $M_L$  and  $I_{Rs}$  from  $M_N$  and  $I_{Rn}$  with ACE;
14:   else
15:    Calculate the Loss of Stage 2 on  $M_L$  and  $I_{Rs}$ :
16:     $L_2 = \omega_0 L_{pho} + \omega_1 L_{lan} + \omega_2 L_{seg} + \omega_4 L_{bin}$ 
17:   end if
18:   Let  $\theta_f, \theta_g$  be the parameters of  $f(\cdot)$  and  $g(\cdot)$ .
19:   Update coefficients with  $L_2$ :  $\nabla_{\alpha, \beta, p, \delta, \gamma, \theta_f, \theta_g} L_2$ .
20: end for
21: Stage 3:
22: Extract texture  $T$  from  $\delta$  as a separate variable.
23: for  $n = 1$  to  $iter_3$  do
24:   Calculate the Loss of Stage 3 :
25:    $L_3 = \omega_0 L_{pho} + \omega_5 L_{GP} + \omega_6 L_{LP} + \omega_7 L_{HP}$ 
26:   Update coefficients with  $L_3$ :  $\nabla_{T, \gamma, \theta_f} L_3$ .
27: end for
```

---

of  $T_D, T_S, T_R$  as  $N_D, N_S, N_R$ , which are computed with the  $5 \times 5$  neighbors around each pixel,  $N_D = T_D - \text{Neighbor}(T_D)$ .  $N_D^0, N_S^0, N_R^0$  are local variations calculated from  $T^0$ . The local prior constraint  $L_{LP}$  is then defined as:

$$L_{LP} = \frac{1}{|I_{in}|} (\|N_D - N_D^0\| + \|N_S - N_S^0\| + \|N_R - N_R^0\|). \quad (7)$$

**Human Prior Constraint.** Human prior constraint is introduced to enhance the optimization of texture  $T$  in Stage 3 with a face recognition network FaceNet [35] pre-trained on large dataset such as VGGFace2 [7] or Casia-Webface [39]. The recognition model pre-trained on VGGFace2 and Casia-webface tends to classify the input images to 8,631 and 10,575 identities with different genders, ethnicity, etc. In this work, we propose a human prior constraint by maximizing the probability that the rendered face is recognized on one of the identities by FaceNet. Defining the FaceNet as  $f_r(\cdot)$ , the human prior constraint can then be written as:

$$L_{HP} = \sum \min_{s \in f_r(I_{Rs})} \|1 - s\|_2. \quad (8)$$

$I_{Rs}$  are rendered under multiple light conditions.  $L_{HP}$  constrains the textures to be more reasonable.

### 3.5 Training Pipeline

The pipeline of the entire training process of our proposed framework is presented in Alg. 1. As illustrated in Fig. 2, the training of our proposed method consists of 3 stages, which is similar to NextFace [12]. In Stage 1, the expression  $\beta$ , shape  $\alpha$ , and pose  $p$  coefficients are optimized to construct the basic face shape. In Stage 2, all coefficients, and the neural representations  $f(\cdot)$  and  $g(\cdot)$  are optimized together to reconstruct the face with statistical model. In Stage 3, the texture  $T$  is directly optimized without coefficient  $\delta$ , while  $\gamma$  and  $f(\cdot)$  are optimized with small learning rates.  $\omega_0 \sim \omega_7$  are pre-defined weights. We use Adam optimizer [19] for optimization.

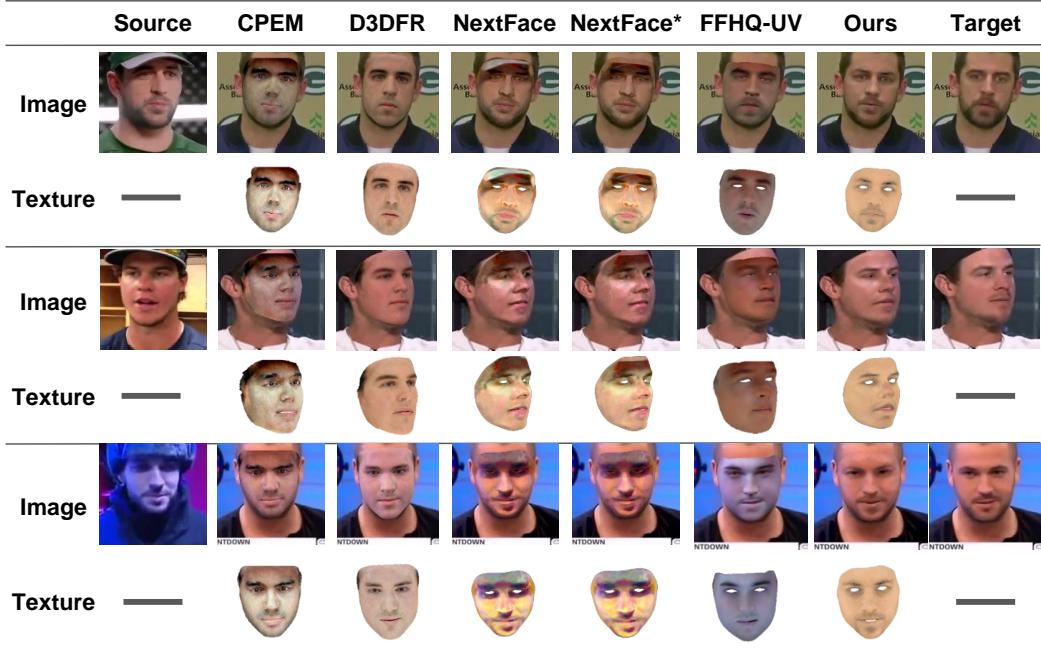


Figure 3: Comparison on Voxceleb2 images. The diffuse albedo is visualized as the texture because it contains most of the color information. Textures from source images are used to synthesize the target images. NextFace\* denotes results optimized within regions selected with face parsing [27]. We do not have textures for CPEM [30] or D3DFR [10] as they predict vertex colors instead of uv textures.

## 4 Experiments

### 4.1 Dataset and Implementation Details

We follow NextFace [12] for the ray-tracing rendering and multi-stage organization of optimization pipeline. Detailed hyper-parameter settings can be found in our supplementary. For our experiments, we utilize two datasets: Voxceleb2 [8] and CelebAMask-HQ [29, 24]. VoxCeleb2 [8] is a diverse dataset encompassing numerous videos collected from interviews, movies and videos, where the same person may have multiple separate videos. CelebAMask-HQ [29, 24] is a large scale face image dataset with fine attributes annotation and high resolution, widely used in face editing and generation.

**Evaluation.** We construct a collection of evaluation data with challenging illumination affected by external occlusions for comparative analysis. This collection includes 38 pairs of single images, 24 pairs of videos from Voxceleb2 [8] with  $256 \times 256$  resolution, and 62 single images from CelebAMask-HQ [29] with  $512 \times 512$  resolution. Each pair consists of source images affected by external occlusions and target images without occlusion, both from the same person, where each video is sampled to 8 frames for sequence-based comparisons.

To quantitatively assess the quality of recovered textures, face textures extracted from the occluded source images are leveraged to synthesize the unoccluded target images. Specifically, we optimize source and target images separately following Sec. 3. Then, keeping the face shape, pose, and illumination invariant, we replace the target image’s texture with the source’s and re-render it to the synthetic result. This allows us to measure texture quality by quantifying the differences between the synthesized target images and the original target images. We also introduce differences between original and reconstructed source images within facial regions acquired by face parsing [27] as an assistant metric to see if textures restore source images well.

Our quantitative comparison employs three metrics: PSNR, SSIM, and perceptual error LPIPS [40] of AlexNet [20]. State-of-the-art methods D3DFR [10], CPEM [30], NextFace [12], and RGB Fitting methods proposed in FFHQ-UV [2] are introduced for comparison. For D3DFR, we adopt its enhanced PyTorch version. As NextFace [12] does not distinguish the face and external occlusion when optimizing the texture, we introduce **NextFace\*** in subsequent comparison by adding face parsing [27] to select the face region during optimization.

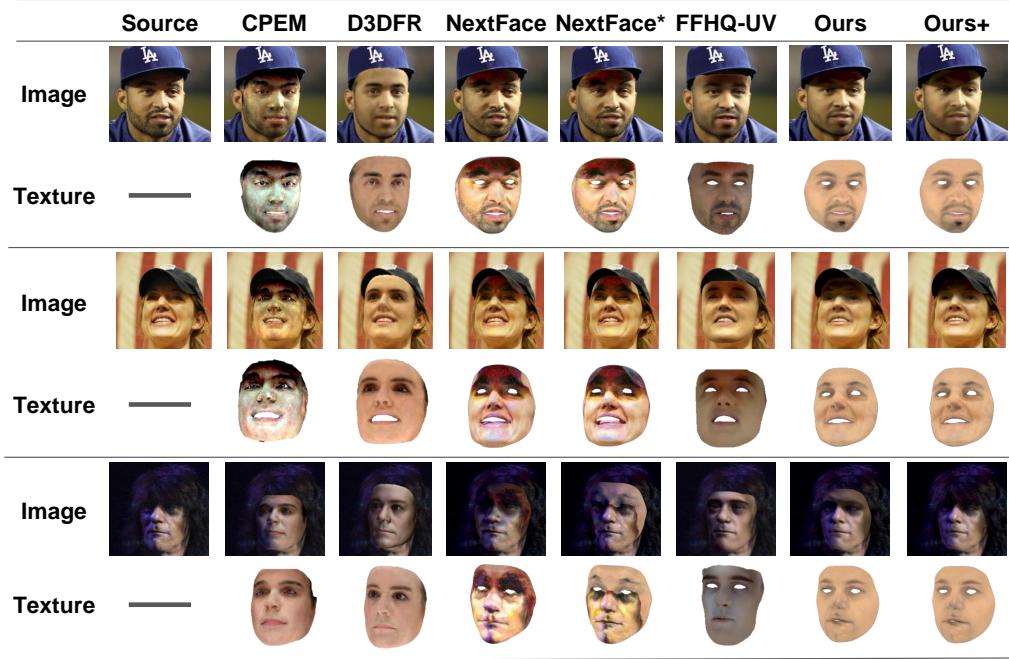


Figure 4: Comparison results on the CelebAMask-HQ dataset. Ours and Ours+ denote our rendered results  $I_R$  directly overlapped onto original images, and results combined with environments:  $I_{out} = M_o \odot I_R + (1 - M_o) \odot I_{in}$ , respectively.

Table 1: Quantitative comparison on single images from Voxceleb2. Source and Target denote differences evaluation on reconstructed source images and synthetic target images. LPIPS is multiplied with  $10^2$ . Underline and **bold** mark the suboptimal and optimal results, respectively.

	CPEM	D3DFR	NextFace	NextFace*	FFHQ-UV	Ours
Source	PSNR $\uparrow$ 27.06	24.26	<b>33.41</b>	32.76	28.72	<u>32.85</u>
	SSIM $\uparrow$ 0.87	0.87	<b>0.95</b>	0.94	0.92	<u>0.94</u>
	LPIPS $\downarrow$ 10.22	11.25	<u>7.37</u>	7.88	8.32	<b>6.40</b>
Target	PSNR $\uparrow$ 24.70	<b>27.23</b>	23.74	24.21	25.03	<b>29.22</b>
	SSIM $\uparrow$ 0.87	0.91	0.85	0.86	<u>0.91</u>	<b>0.91</b>
	LPIPS $\downarrow$ 9.43	7.26	10.52	10.02	<u>7.19</u>	<b>6.36</b>

## 4.2 Evaluation on Voxceleb2

**Evaluation on Single Images.** We conduct an evaluation about the performance of our method in collected single images sourced from the Voxceleb2 dataset [8]. As shown in Fig. 3, our approach demonstrates a marked improvement in recovering clearer textures from the original images taken under challenging illumination affected by external occlusions. Our method also surprisingly recovers clear textures under strong colorful lights, which may benefit from the realistic constraints to keep textures reasonable. In contrast, existing methods tend to incorporate these occlusions, shadows, or colorful lights directly into the texture, resulting in less satisfactory outcomes.

Furthermore, our methods achieves superior results on the synthesis of target images and comparable performances on the reconstruction of source images as shown in Table 1. It confirms our method consistently recovers both accurate and clear textures. Although NextFace [12] performs a little better on PSNR and SSIM of source images, it performs the worst on synthetic target images as it actually severely over-fits source images. As shown in Fig. 3, it bakes shadows and occlusions to the textures, which confirms it is not appropriate for faces affected by external occlusions.

**Evaluation on Video Sequences.** We conduct an evaluation on video sequences collected from Voxceleb2 dataset [8] to further substantiate the efficacy of our approach. To adapt our framework

Table 2: Quantitative comparison on video sequences from Voxceleb2. Source and Target denote differences evaluation on reconstructed source sequences and synthetic target sequences.

	CPEM	D3DFR	NextFace	NextFace*	FFHQ-UV	Ours
Source	PSNR $\uparrow$ 27.69	24.55	30.60	<u>30.65</u>	29.70	<b>30.67</b>
	SSIM $\uparrow$ 0.87	0.87	0.92	0.92	<u>0.92</u>	<b>0.92</b>
	LPIPS $\downarrow$ 9.21	9.92	8.17	8.45	<u>7.64</u>	<b>7.62</b>
Target	PSNR $\uparrow$ 24.33	<b>26.20</b>	24.15	24.32	24.35	<b>29.15</b>
	SSIM $\uparrow$ 0.87	0.90	0.87	0.87	<u>0.91</u>	<b>0.91</b>
	LPIPS $\downarrow$ 9.58	8.13	10.27	9.91	<u>7.66</u>	<b>6.92</b>

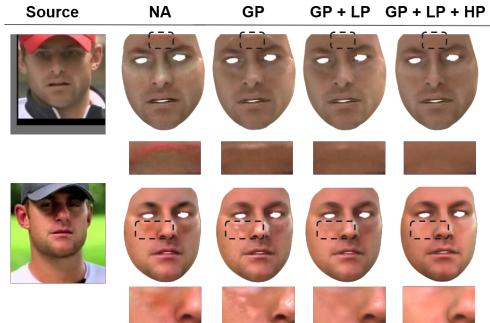


Figure 5: Ablation study for losses. GP, LP, and HP denote  $L_{GP}$ ,  $L_{LP}$ ,  $L_{HP}$ , respectively, while NA means to remove all of them.

Table 3: Quantitative ablation study for proposed Losses GP, LP, and HP.

Metrics	NA	GP	GP+LP	GP+LP+HP (ours)
PSNR $\uparrow$	27.20	28.81	28.82	<b>29.22</b>
SSIM $\uparrow$	0.88	0.89	0.90	<b>0.91</b>
LPIPS $\downarrow$	8.34	6.78	6.49	<b>6.36</b>

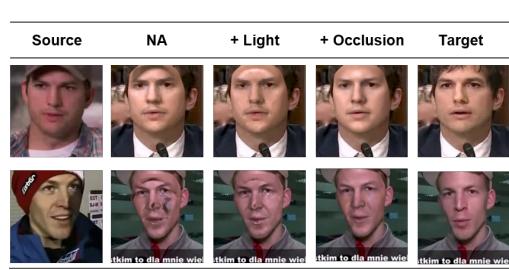


Figure 6: Ablation study for the neural representations. NA means to remove both  $f(\cdot)$  and  $g(\cdot)$ , while **+ Light** and **+ Occlusion** denote adding  $f(\cdot)$  and  $g(\cdot)$ , respectively.

Table 4: Quantitative ablation study for  $f(\cdot)$ ,  $g(\cdot)$ .

Metrics	NA	+ Light ( $f(\cdot)$ ) + Occlusion ( $g(\cdot)$ )
PSNR $\uparrow$	25.19	27.52
SSIM $\uparrow$	0.87	0.89
LPIPS $\downarrow$	9.16	7.75
		<b>29.22</b>
		<b>0.91</b>
		<b>6.36</b>

to video sequences, we share texture, illumination and shape coefficients across all frames during optimization. As D3DFR [10] and FFHQ-UV [2] do not provide support for sequences, we recurrently apply them to each single image. Quantitative comparison in Table 2 demonstrates that our method constantly performs superior to existing methods on texture modeling from continuous video sequences. Please check our Supplementary for corresponding qualitative results.

### 4.3 Evaluation on CelebAMask-HQ

We extend our comparisons on CelebAMask-HQ [29]. As this dataset lacks multiple images from the same identity, we focus on evaluating the performance based on the recovered textures and the reconstructed results of source images. The outcomes of our assessment are visually depicted in Fig. 4. We observe that our method continues to perform well in the task of recovering clear textures from challenging input. Additionally, our reconstructed results exhibit a higher degree of realism on the reconstructed results when compared to the outcomes produced by other methods.

### 4.4 Evaluation on Images with Diverse Shadows [41]

To validate the performances of our method more sufficiently, we further evaluate our method on the single image dataset proposed by [41], which includes 100 images affected by manually created external shadows, as well as corresponding ground truths. The quantitative and qualitative comparisons are presented in Table 5 and Fig. 7. We can observe that our method still outperforms other methods under faces with diverse shadows.

Table 5: Quantitative comparisons on images with diverse shadows.

Metrics	CPEM	D3DFR	NextFace	NextFace*	FFHQ-UV	Ours
PSNR $\uparrow$	24.75	23.02	32.10	31.92	26.93	<b>32.13</b>
SSIM $\uparrow$	0.83	0.83	0.94	0.94	0.89	<b>0.94</b>
LPIPS $\downarrow$	10.99	12.76	<b>5.26</b>	5.34	8.71	6.29
PSNR $\uparrow$	23.34	25.02	21.83	21.93	24.00	<b>28.97</b>
SSIM $\uparrow$	0.84	0.87	0.84	0.84	0.88	<b>0.92</b>
LPIPS $\downarrow$	9.31	9.11	9.48	9.69	7.82	<b>7.00</b>

	Source	CPEM	D3DFR	NextFace	NextFace*	FFHQ-UV	Ours	Target
<b>Image</b>								
<b>Texture</b>	—							—
<b>Image</b>								
<b>Texture</b>	—							—

Figure 7: Comparison on images with diverse shadows [41]. As [41] provides corresponding ground truths of shadow affected images, we directly use such ground truths as the target images.

#### 4.5 Ablation Study

**Ablation Study on Losses.** In this section, we analyze the impact of the proposed losses:  $L_{GP}$ ,  $L_{LP}$ , and  $L_{HP}$  by comparing the rendered results under relative brighter light conditions. As illustrated in Fig. 5, it becomes evident that the inclusion of the  $L_{GP}$  loss plays a significant role in eliminating the abnormal colors on the texture, where  $L_{LP}$  loss effectively remove large artifacts obviously different from human faces.  $L_{HP}$  can further reduce slight unreasonable defects with priors from the pre-trained recognition model [35]. Quantitative comparisons in Table 3 also confirm that each proposed loss contributes to the final performances.

**Ablation Study on Neural Representations.** To ascertain the significance of the neural representations employed, we perform an analysis where we eliminate both the neural representations:  $f(\cdot)$  and  $g(\cdot)$  from the light decoupling pipeline. As depicted in Fig. 6, the neural representation for the decoupling of light conditions  $f(\cdot)$  notably contributes in the removal of shadows artifacts from the results. Furthermore, the neural representation  $g(\cdot)$  displays its efficacy in minimizing existing irregularities in the synthetic results, which achieves smoother and cleaner results. We also present corresponding quantitative comparisons in Table 4, which demonstrates that both  $f(\cdot)$  and  $g(\cdot)$  have obvious influence on the final performances.

## 5 Conclusion

In this paper, we present a novel framework dedicated to recover clear facial textures from images taken under challenging illumination affected by external occlusion. Our approach makes use of neural representations to decouple the original illumination into multiple separate light conditions across various facial regions. Then the affected complicated illumination can be modelled with the combination of different light conditions. Furthermore, we introduce well-established human face priors through the realistic constraints to enhance the realism of our results. According to experiments on single images and video sequences, our method consistently surpasses existing techniques to recover clearer and more accurate textures from faces under affected unnatural illumination.

## Acknowledgement

This research / project is supported by the National Research Foundation (NRF) Singapore, under its NRF-Investigatorship Programme (Award ID. NRF-NRFI09-0008).

## References

[1] Edward Angel. Interactive computer graphics: a top-down approach with shader-based opengl. 2011.

[2] Haoran Bai, Di Kang, Haoxian Zhang, Jinshan Pan, and Linchao Bao. Ffhq-uv: Normalized facial uv-texture dataset for 3d face reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 362–371, 2023.

[3] Linchao Bao, Xiangkai Lin, Yajing Chen, Haoxian Zhang, Sheng Wang, Xuefei Zhe, Di Kang, Haozhi Huang, Xinwei Jiang, Jue Wang, et al. High-fidelity 3d digital human head creation from rgb-d selfies. *ACM Transactions on Graphics (TOG)*, 41(1):1–21, 2021.

[4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 157–164. 2023.

[5] James F Blinn. Models of light reflection for computer synthesized pictures. In *Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, pages 192–198, 1977.

[6] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3d morphable model learnt from 10,000 faces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5543–5552, 2016.

[7] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 67–74. IEEE, 2018.

[8] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. Voxceleb2: Deep speaker recognition. *arXiv preprint arXiv:1806.05622*, 2018.

[9] Wei-Chieh Chung, Jian-Kai Zhu, I-Chao Shen, Yu-Ting Wu, and Yung-Yu Chuang. Stylefaceuv: A 3d face uv map generator for view-consistent face image synthesis. 2022.

[10] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019.

[11] Abdallah Dib, Junghyun Ahn, Cedric Thebault, Philippe-Henri Gosselin, and Louis Chevallier. S2f2: Self-supervised high fidelity face reconstruction from monocular image. *arXiv preprint arXiv:2203.07732*, 2022.

[12] Abdallah Dib, Gaurav Bharaj, Junghyun Ahn, Cédric Thébault, Philippe Gosselin, Marco Romeo, and Louis Chevallier. Practical face reconstruction via differentiable ray tracing. In *Computer Graphics Forum*, volume 40, pages 153–164. Wiley Online Library, 2021.

[13] Abdallah Dib, Cedric Thebault, Junghyun Ahn, Philippe-Henri Gosselin, Christian Theobalt, and Louis Chevallier. Towards high fidelity monocular face reconstruction with rich reflectance using self-supervised learning and ray tracing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12819–12829, 2021.

[14] Bernhard Egger, Sandro Schönborn, Andreas Schneider, Adam Kortylewski, Andreas Morel-Forster, Clemens Blumer, and Thomas Vetter. Occlusion-aware 3d morphable models and an illumination prior for face image analysis. *International Journal of Computer Vision*, 126:1269–1287, 2018.

[15] Bernhard Egger, William AP Smith, Ayush Tewari, Stefanie Wuhrer, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, et al. 3d morphable face models—past, present, and future. *ACM Transactions on Graphics (ToG)*, 39(5):1–38, 2020.

[16] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1155–1164, 2019.

[17] Yingqing He, Yazhou Xing, Tianjia Zhang, and Qifeng Chen. Unsupervised portrait shadow removal via generative priors. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 236–244, 2021.

[18] Andrew Hou, Michel Sarkis, Ning Bi, Yiying Tong, and Xiaoming Liu. Face relighting with geometrically consistent shadows. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4217–4226, 2022.

[19] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.

[21] Alexandros Lattas, Stylianos Moschoglou, Baris Gecer, Stylianos Ploumpis, Vasileios Trianafyllou, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme: Realistically renderable 3d facial reconstruction" in-the-wild". In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 760–769, 2020.

[22] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Jiankang Deng, and Stefanos Zafeiriou. Fitme: Deep photorealistic 3d morphable model avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8629–8640, 2023.

[23] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme++: Facial shape and brdf inference with photorealistic rendering-aware gans. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9269–9284, 2021.

[24] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.

[25] Chunlu Li, Andreas Morel-Forster, Thomas Vetter, Bernhard Egger, and Adam Kortylewski. Robust model-based face reconstruction through weakly-supervised outlier segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 372–381, 2023.

[26] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 37(6):222:1–222:11, 2018.

[27] Jinpeng Lin, Hao Yang, Dong Chen, Ming Zeng, Fang Wen, and Lu Yuan. Face parsing with roi tanh-warping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5654–5663, 2019.

[28] Yaojie Liu, Andrew Z Hou, Xinyu Huang, Liu Ren, and Xiaoming Liu. Blind removal of facial foreign shadows. In *BMVC*, page 88, 2022.

[29] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaou Tang. Large-scale celebfaces attributes (celeba) dataset. *Retrieved August*, 15(2018):11, 2018.

[30] Langyuan Mo, Haokun Li, Chaoyang Zou, Yubing Zhang, Ming Yang, Yihong Yang, and Mingkui Tan. Towards accurate facial motion retargeting with identity-consistent and expression-exclusive constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1981–1989, 2022.

[31] Thomas Nestmeyer, Jean-François Lalonde, Iain Matthews, and Andreas Lehrmann. Learning physics-guided face relighting under directional light. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5124–5133, 2020.

[32] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *2009 sixth IEEE international conference on advanced video and signal based surveillance*, pages 296–301. Ieee, 2009.

[33] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 497–500, 2001.

[34] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 117–128, 2001.

[35] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.

[36] Soumyadip Sengupta, Angjoo Kanazawa, Carlos D Castillo, and David W Jacobs. Sfsnet: Learning shape, reflectance and illuminance of faces in the wild'. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6296–6305, 2018.

[37] William AP Smith, Alassane Seck, Hannah Dee, Bernard Tiddeman, Joshua B Tenenbaum, and Bernhard Egger. A morphable face albedo model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5011–5020, 2020.

- [38] Chao Xu, Yang Liu, Jiazheng Xing, Weida Wang, Mingze Sun, Jun Dan, Tianxin Huang, Siyuan Li, Zhi-Qi Cheng, Ying Tai, et al. Facechain-imagineid: Freely crafting high-fidelity diverse talking faces from disentangled audio. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1292–1302, 2024.
- [39] Dong Yi, Zhen Lei, Shengcui Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014.
- [40] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [41] Xuaner Zhang, Jonathan T Barron, Yun-Ta Tsai, Rohit Pandey, Xiuming Zhang, Ren Ng, and David E Jacobs. Portrait shadow manipulation. *ACM Transactions on Graphics (TOG)*, 39(4):78–1, 2020.
- [42] Zhenyu Zhang, Renwang Chen, Weijian Cao, Ying Tai, and Chengjie Wang. Learning neural proto-face field for disentangled 3d face modeling in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 382–393, June 2023.
- [43] Zhenyu Zhang, Yanhao Ge, Ying Tai, Xiaoming Huang, Chengjie Wang, Hao Tang, Dongjin Huang, and Zhifeng Xie. Learning to restore 3d face from in-the-wild degraded images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4237–4247, June 2022.
- [44] Mingwu Zheng, Hongyu Yang, Di Huang, and Liming Chen. Imface: A nonlinear 3d morphable face model with implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20343–20352, 2022.
- [45] Mingwu Zheng, Haiyu Zhang, Hongyu Yang, and Di Huang. Neuface: Realistic 3d neural face rendering from multi-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16868–16877, 2023.

## A Appendix / Supplementary Material

### A.1 Limitation

We use AlbedoMM [37] to initialize the face texture and optimize neural representations. Limited by the capacity of AlbedoMM, the initialized face texture may not be accurate enough for subsequent optimization in some occasions. We can explore to resolve this problem by combining our method with non-linear texture modeling methods such as [22, 2] in our future work. Furthermore, the optimization of multiple light conditions makes our method slower than other single illumination methods [12, 2]. It costs 340s for a  $256 \times 256$  image on 2080ti, which is slower than 240s of FFHQ-UV [2] but still affordable.

### A.2 Details of Parameter Settings

In Table 6, we present details about the hyper-parameters mentioned in Sec. 3. Although it seems that there are multiple parameters, *the setting is robust for both images and video sequences from adopted datasets*. We conduct all experiments on a Nvidia 2080ti GPU with a 2.9Ghz i5-9400 CPU. It costs 340s for a  $256 \times 256$  image.

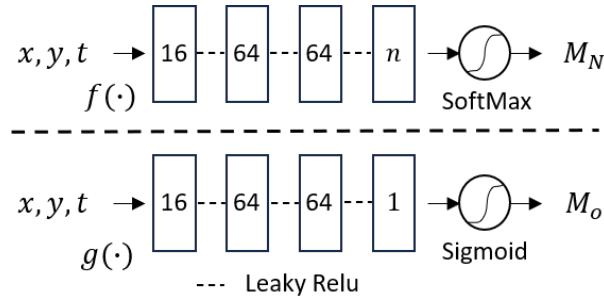


Figure 8: Detailed design of  $f(\cdot)$  and  $g(\cdot)$ .

Table 6: Hyper-parameter settings.  $n$  is the number of initial light conditions presented in Fig. 2.

Parameter	
$w_1 \sim w_7$	2e3, 1e-3, 1.5e2, 0.5, 25, 2e3, 2.0, 1.0
Landmark	Mediapipe
$iter_0 \sim iter_3$	100, 2000, 400, 200
$\epsilon$	0.17
$n$	5

### A.3 Comparisons against the combination of 2D Shadow Removal and 3D Texture Modeling

Although the former mentioned existing 3D face texture modeling methods cannot process facial shadows from external occlusions directly, there are some 2D shadow removal methods [41, 17, 28] trying to eliminate the shadows directly from the images. Therefore, another alternative simple baseline is to pre-process the image with 2D shadow-removal networks before texture modeling. In this section, we introduce the most recent method [17] with a pre-trained model to pre-process the images before feeding them to compared methods mentioned in Sec. 4.1. The quantitative and qualitative comparison are presented in Table 7, Table 8 and Fig. 11, respectively.

Table 7: Comparisons against baselines with 2D shadow-removal pre-processing on [41].

	CPEM	D3DFR	NextFace	NextFace*	FFHQ-UV	Ours
PSNR $\uparrow$	25.56	25.60	25.24	25.47	26.19	<b>28.97</b>
SSIM $\uparrow$	0.86	0.88	0.87	0.87	0.90	<b>0.92</b>
LPIPS $\downarrow$	8.56	8.68	<b>6.37</b>	6.41	7.02	7.00

Pre-processing with 2D shadow removal methods indeed improves the performances of baselines, while our method still outperforms them. From qualitative results, we observe that the shadow

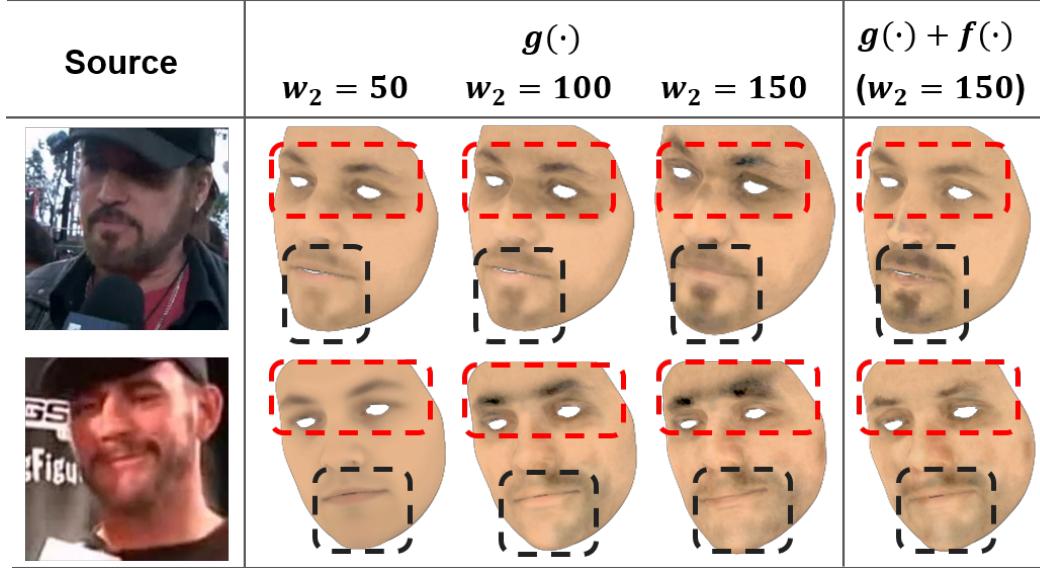


Figure 9: Discussion about the effect of  $g(\cdot)$ .  $w_2$  is the weight to constrain  $g(\cdot)$ , defined in Alg. 1. The **red** and **black** rectangles mark shadow-affected regions and detailed textures, respectively.  $g(\cdot)$  will weaken both shadows and details from textures when reducing  $w_2$  to loose its constraint.

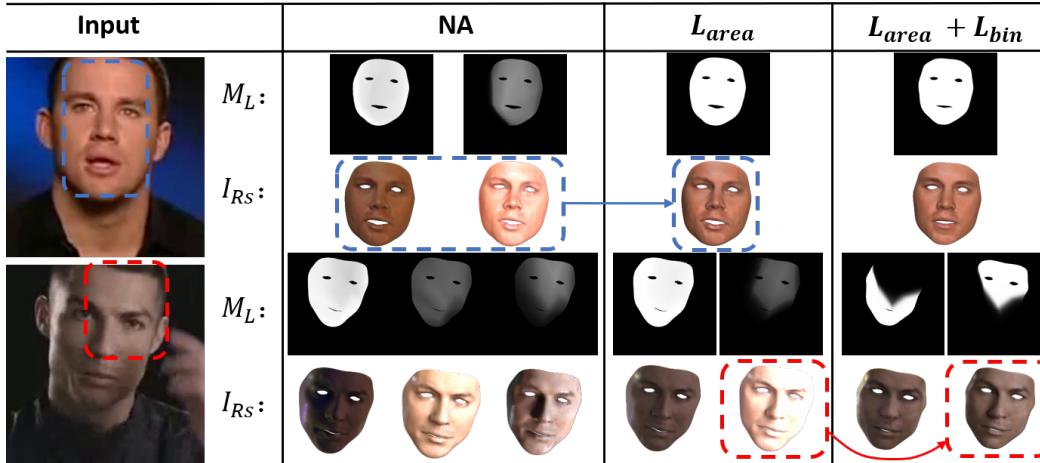


Figure 10: Ablation study for  $L_{area}$  and  $L_{bin}$  in ACE. NA denotes removing both  $L_{area}$  and  $L_{bin}$ .  $L_{area}$  can remove redundant light conditions as shown by the **blue** rectangle, while  $L_{bin}$  ensures the light condition shown in the **red** rectangle region is consistent as our observation of the input.  $M_L$  and  $I_{RS}$  are predicted masks and rendered faces defined in Fig. 2, respectively.

removal model cannot fully remove the external shadows in cases where the external shadows cover relatively large regions. Although modifying the shadow removal model to be more powerful may further improve performances, it goes beyond the range of this work. We can explore it in the future.

#### A.4 Comparisons with Deocclusion methods

Besides the former mentioned shadow removal baselines, learning-based deocclusion methods [25, 14] can remove external occlusions from faces by predicting the occluded regions. Such operations also have the potential to deal with external shadows by directly treating the shadow regions as occlusions. In this section, we conduct a brief comparison with the most recent open sourced deocclusion method [25]. The quantitative and qualitative results are presented in Table 9 and Fig. 12, respectively. The metrics are evaluated on the target images mentioned in Sec. 4.1. We can observe that the method [25] lose some facial details such as beards. The reason is that the deocclusion method [25] actually divides occlusion regions by the distances between input images

Table 8: Comparisons against baselines with 2D shadow-removal pre-processing on images from Voxceleb2.

	CPEM	D3DFR	NextFace	NextFace*	FFHQ-UV	Ours
PSNR $\uparrow$	24.84	26.78	23.77	24.51	25.35	29.22
SSIM $\uparrow$	0.87	0.90	0.85	0.86	0.91	0.91
LPIPS $\downarrow$	10.20	7.93	10.47	9.64	7.62	6.36

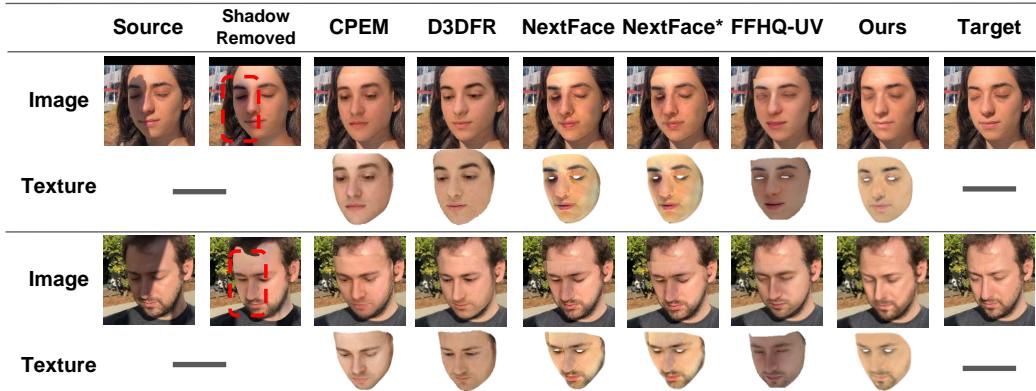


Figure 11: Qualitative Comparisons against baselines with 2D shadow-removal pre-processing.

and the reconstructed results from 3D Morphable Model (3DMM), which will also regard details hard to describe with 3DMM as occlusions. Our methods treat the shadows as different illumination instead of occlusions, which can preserve facial details in these regions in subsequent optimization.

#### A.5 Discussion about the number of Sphere Harmonics (SH) bands.

In this work, we follow NextFace [1] to use 9-bands SH to model the local illumination, which can actually capture quite fine details. However, SH with more bands would have stronger ability for the modeling of illumination. To verify if the external shadows can be directly modelled with more bands SH, we also provide the quantitative results of our method modelled with only one single global SH in 9, 12, 15, 18 bands in Table 10, where we remove our light decoupling framework. We observe that increasing the number of bands in a single global SH yields quite limited improvements. A possible reason is that the external occluded shadows on human faces represent drastic illumination changes in relatively small areas, which may not be appropriately modeled as a single global SH during optimization.

#### A.6 Discussion about the number of lighting conditions $n$ .

In this work, we initialize  $n$  different lighting conditions to imitate the illumination affected by external occlusion at the beginning, where some of them would be removed by ACE in subsequent processing. To explore the influence of different numbers of  $n$ , we conduct an ablation study for the number of  $n$  used for initialization of lighting conditions in Table 11. Single images from VoxCeleb2 are used for evaluation. We observe that  $n = 3$  produces sub-optimal results, likely because the number of lighting condition candidates is insufficient to model images with complex illuminations. In contrast,  $n = 5, 7, 9$  yield good and similar results as there are enough initial lighting conditions, and any redundant ones are removed. The result for  $n = 5$  is slightly better. While introducing larger  $n$  and further adjusting the hyper-parameters might improve performance, it would also increase the optimization burden.

#### A.7 Differences between Human Prior Constraint and Perceptual Loss

Please note that the rendered faces used to calculate Human Prior Constraint (HP) are calculated with  $I_{Rs}$ . As illustrated in Fig.2,  $I_{Rs}$  includes rendered faces under multiple different lighting conditions.

Table 9: Comparisons against baselines with Deocclusion methods on Voxceleb2.

	Single Images		Sequences	
Methods	Deocclu	Ours	Deocclu	Ours
PSNR $\uparrow$	25.13	<b>29.22</b>	25.03	<b>29.15</b>
SSIM $\uparrow$	0.88	<b>0.91</b>	0.88	<b>0.91</b>
LPIPS $\downarrow$	14.09	<b>6.36</b>	14.18	<b>6.92</b>

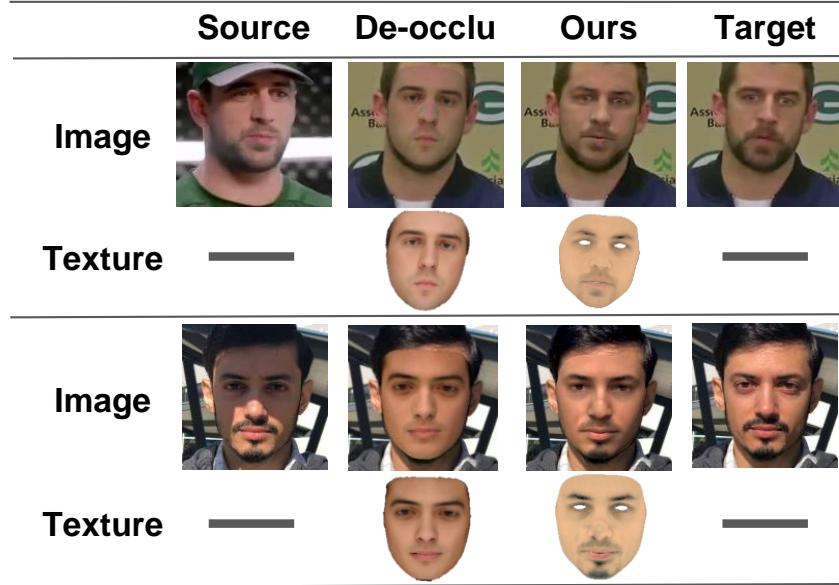


Figure 12: Qualitative Comparisons with the Deocclusion method [25].

We cannot use perceptual loss here because we do not have \*\*corresponding ground truth images under the multiple decoupled lighting conditions\*\*, where HP does not need such ground truths. Nonetheless, the perceptual loss can be indeed applied between the final rendered result  $I_{out}$  and input image  $I_{in}$ . We present a comparison between such perceptual loss implementation and the HP in Table 12. We can see that HP still has better performances, which can confirm it provides more effective constraints for the textures through rendered faces under multiple lighting conditions.

#### A.8 Discussion about $g(\cdot)$ .

Both  $f(\cdot)$  and  $g(\cdot)$  have impacts to remove artifacts on the recovered textures. To confirm that they have principled differences, we implement a discussion about the effect of  $g(\cdot)$ . As presented in the third column of Fig. 9,  $g(\cdot)$  cannot fully avoid the influence of external shadows on the faces. As presented in Sec. 3.3,  $L_{seg}$  is introduced to constrain  $g(\cdot)$  to ensure that  $g(\cdot)$  predicts relatively accurate face regions.  $L_{seg}$  is weighted by  $w_2$  in optimization as defined in Alg. 3. Although the shadow effect can decrease when we reduce  $w_2$  to allow  $g(\cdot)$  to filter out more face regions, the detailed textures are equally weaken and may be fully removed as shown in the last row of Fig. 9. It confirms that  $f(\cdot)$  is still essential for this framework.

#### A.9 Analysis about $L_{bin}$ and $L_{area}$ in ACE.

In ACE mentioned in Sec 3.3, two regularization constraints  $L_{area}$  and  $L_{bin}$  are introduced. As the blue rectangle input face of Fig. 10 is not affected by any external occlusion, its illumination can be modelled with only one light condition.  $L_{area}$  helps remove redundant light conditions. However, only using  $L_{area}$  may create decoupled light conditions obviously different from the observation of input image, as illustrated in the red rectangle regions of Fig. 10. Adding  $L_{bin}$  can ensure the

Table 10: Ablation study for the number of SH bands.

B	9	12	15	18	Ours
PSNR $\uparrow$	25.87	25.26	25.27	25.34	<b>29.22</b>
SSIM $\uparrow$	0.87	0.87	0.87	0.87	<b>0.91</b>
LPIPS $\downarrow$	9.16	9.23	9.22	9.10	<b>6.36</b>

Table 11: Ablation study for the number of lighting conditions  $n$ .

n	3	5 (we use)	7	9
PSNR $\uparrow$	28.49	<b>29.22</b>	29.14	29.16
SSIM $\uparrow$	0.90	<b>0.91</b>	0.91	0.91
LPIPS $\downarrow$	6.37	<b>6.36</b>	6.37	6.39

predicted masks  $M_L$  to be nearly binarized, which makes the decoupled light conditions consistent as the human observation of the input image.

#### A.10 Differences between Stage 2 and Stage 3.

Illustrated in Fig. 2, our approach composes of both Stage 2 and Stage 3 to optimize the face texture. In this section, we conduct comparisons to demonstrate differences between textures from Stage 2 and Stage 3. The results are presented in Fig. 13. We can observe that textures from AlbedoMM [37]

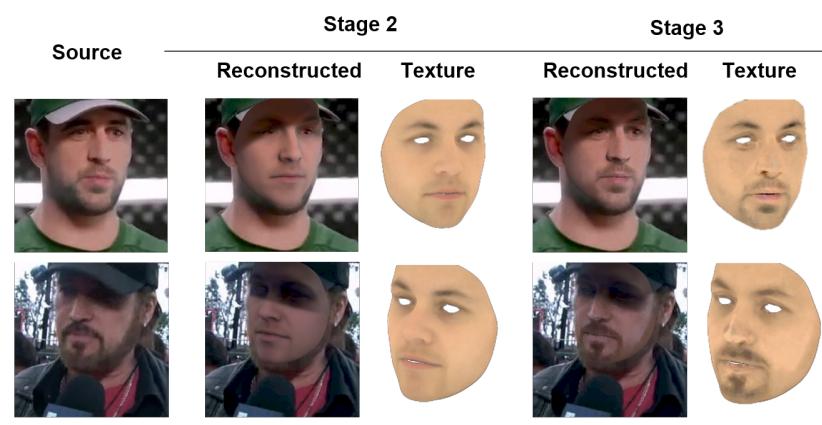


Figure 13: Differences between Stage 2 and Stage 3. In Stage 3, the texture is refined with details from the source image, such as the beard, to render a more realistic reconstructed image.

in Stage 2 are quite over-smoothed and lack of details. With Stage 3, details in source images are added to textures, which can reconstruct more realistic results than Stage 2.

#### A.11 Visualization about $M_o$

We present an ablation study to confirm the effect of  $g(\cdot)$  against direct segmented mask from face parsing [27] to predict  $M_o$ . The circled parts of Fig. 14 show that the parsed masks may be inappropriate due to the limitation of generalizability. We can see in the first row of Fig. 14 that the rendered result  $I_R$  may have rough edges and artifact colors from the occlusion, while our method can avoid this problem by refining the parsed mask with  $g(\cdot)$ . Moreover, some parts may be missing on the parsed mask. As shown in the second and third rows of Fig. 14, artifacts show up in these unconstrained regions. Our method can complete these missing regions for more reasonable results.

Table 12: Comparison between HP and Perceptual loss.

n	Perceptual Loss [10]	HP (ours)
PSNR $\uparrow$	28.72	<b>29.22</b>
SSIM $\uparrow$	0.90	<b>0.91</b>
LPIPS $\downarrow$	6.46	<b>6.36</b>

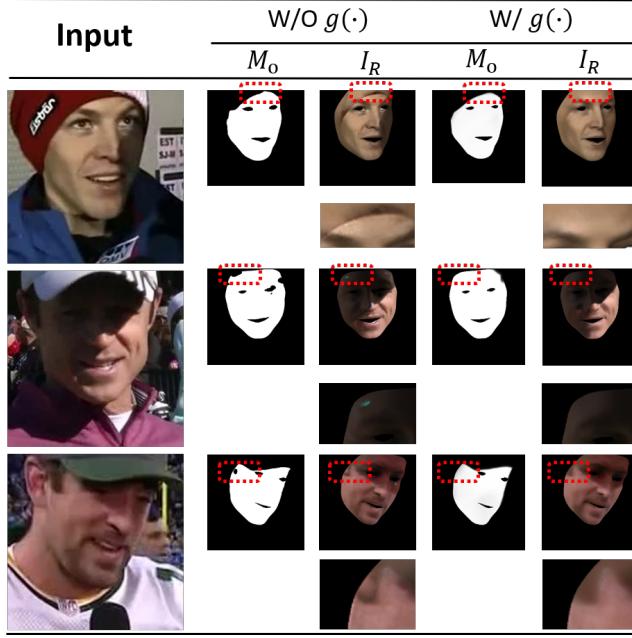


Figure 14: Ablation study for the usage of neural representation  $g(\cdot)$ .  $W/O g(\cdot)$  and  $W/ g(\cdot)$  denote using parsed mask [27] and using  $g(\cdot)$ , respectively.  $M_o$  and  $I_R$  are the mask and rendered results, as mentioned in Sec. 3.

### A.12 Discussion about the failure cases.

Except the efficiency problem mentioned in Sec. A.1, our primary limitation is the initialization with AlbedoMM. As shown in Fig. 15, for faces with many high-frequency details, such as wrinkles, our method may lose these details during reconstruction. Replacing AlbedoMM with more powerful face representations could address this issue. We plan to explore this further in future work.

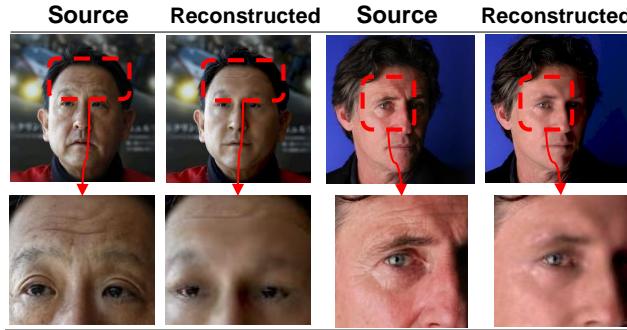


Figure 15: Some failure cases.

### A.13 Effect of Adaptive Condition Estimation

As described in Sec. 3.3, the Adaptive Condition Estimation (ACE) is proposed to select effective  $M_L$  and  $I_{Rs}$  from the initialized  $M_N$  and  $I_{Rn}$ . To remove ACE, we use initialized  $M_N$  and  $I_{Rn}$  as  $M_L$  and  $I_{Rs}$  directly. As shown in Fig. 16, we can see that the rendered results  $I_{out}$  are almost

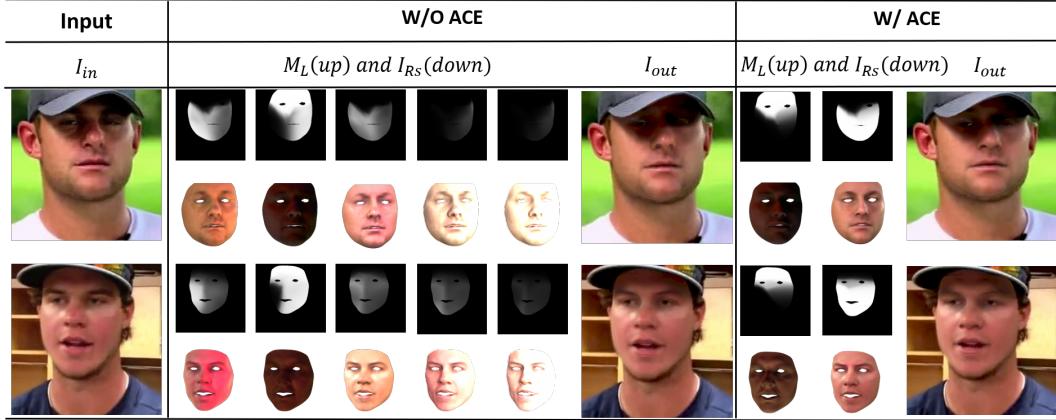


Figure 16: Ablation study for ACE. *W/O ACE* and *W/ ACE* denote removing ACE by using  $M_N$  and  $I_{Rn}$  as  $M_L$  and  $I_{Rs}$ , and using ACE to select  $M_L$  and  $I_{Rs}$ , respectively.

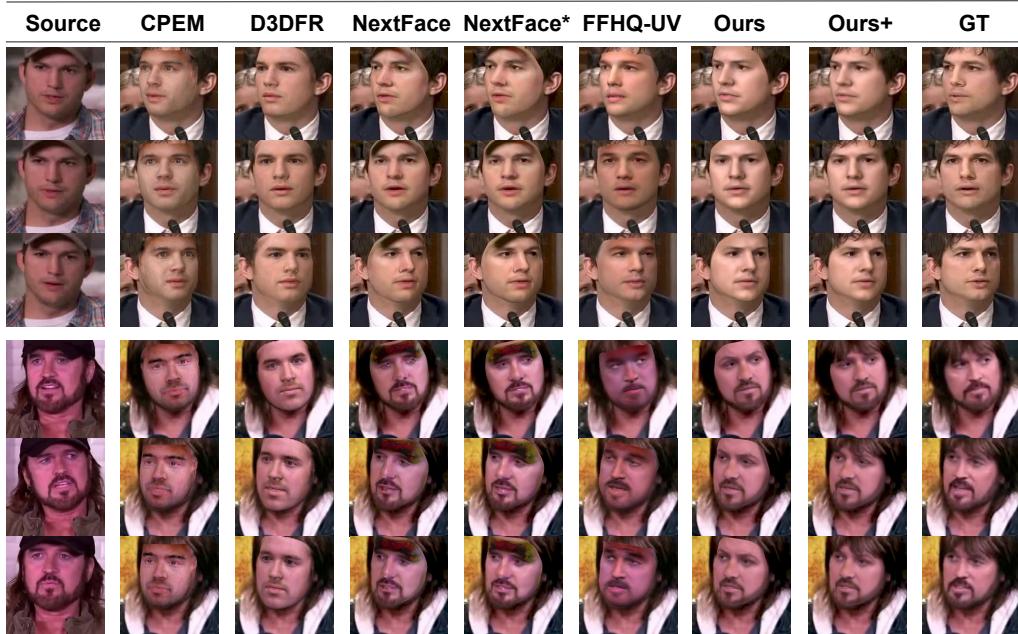


Figure 17: Comparison results on the video sequences from Voxceleb2. *Ours* and *Ours+* denote our rendered results  $I_R$  directly overlapped onto original images, and results combined with environments:  $I_{out} = M_o \odot I_R + (1 - M_o) \odot I_{in}$ , respectively. The symbols are defined following Sec. 3.

the same, while the results optimized without ACE contain multiple redundant and inaccurate light conditions in  $M_L$ . It confirms that ACE can help remove these unnecessary light conditions and keep effective ones. With ACE, our method can decouple the original illumination affected by occlusions into light conditions more consistent with actual observations of input images.

#### A.14 More Visualized Results

In this section, we present three representative examples from each sequence in Fig. 17. It is evident that our method performs the best in generating synthetic results close to the target sequences, exhibiting a high degree of realism. In contrast, other methods still produce less convincing outcomes due to negative effects from external occlusions. More results on images sourced from Voxceleb2 [8] and CelebAMask-HQ [24] are presented in Fig. 18 and Fig. 19. Our method still performs better. Please refer to the attached video for more results on sequences from Voxceleb2 [8].

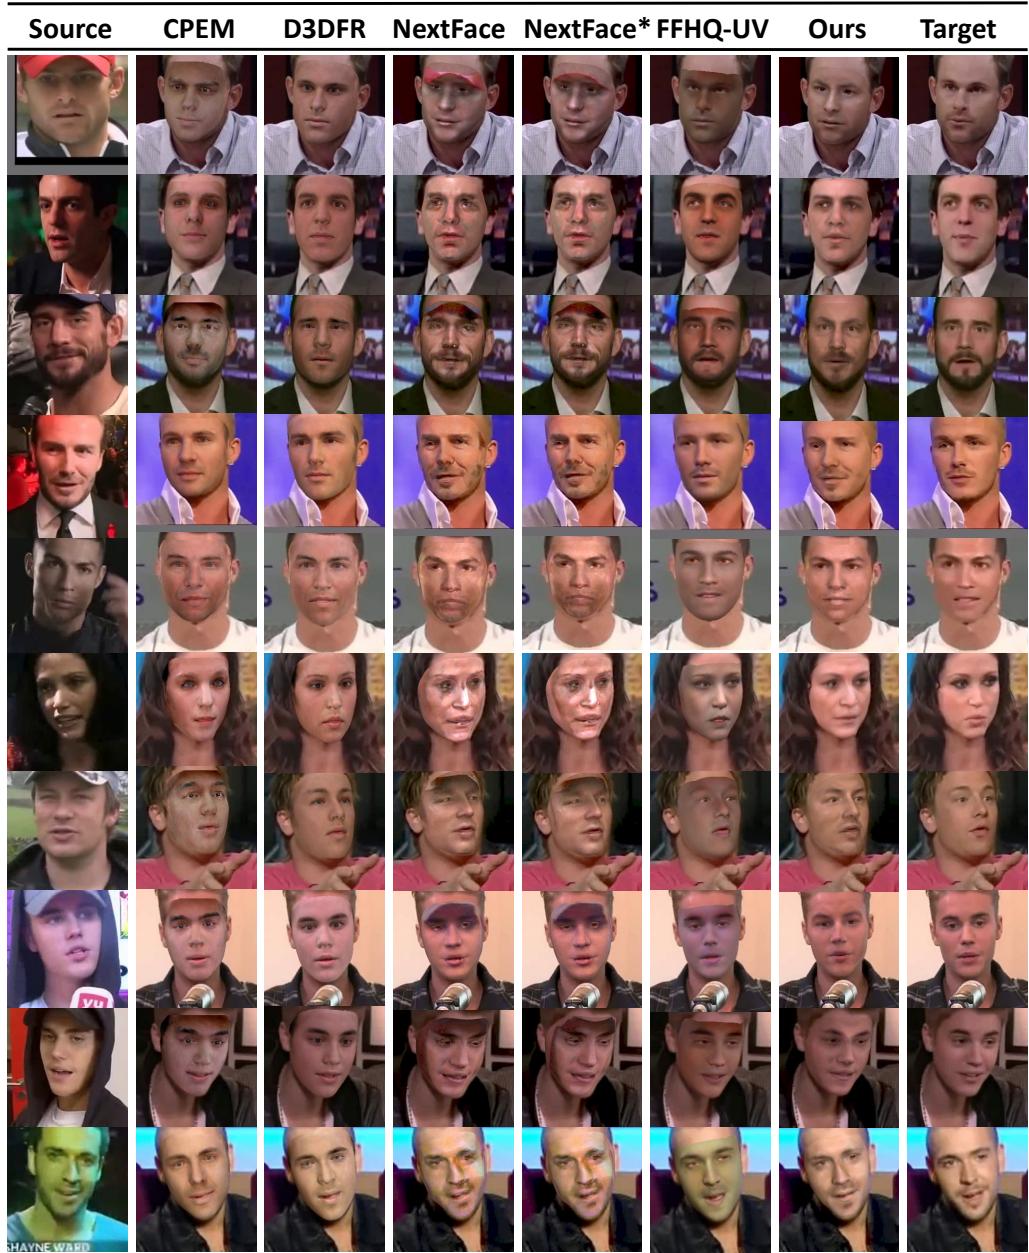


Figure 18: More results on single images from Voxceleb2 [8]. We can see that our method can synthesize more accurate target images based on textures from the source images, which validate the quality of our acquired textures.

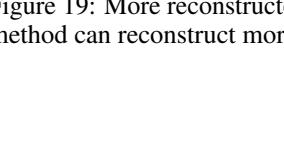
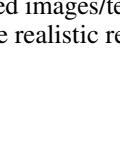
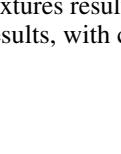
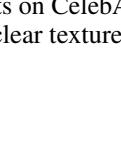
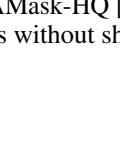
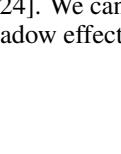
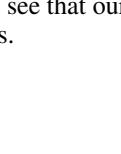
Input	CPEM	D3DFR	NextFace	NextFace*	FFHQ-UV	Ours
						
						
						
						
						

Figure 19: More reconstructed images/textures results on CelebAMask-HQ [24]. We can see that our method can reconstruct more realistic results, with clear textures without shadow effects.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: **[Yes]**

Justification: We have discussed them in the page 2.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: **[Yes]**

Justification: We have discussed about that in Sec. A.1.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: **[NA]**

Justification: We do not have theoretical contributions in this work, where our contributions are validated with experiments.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All the hyper-parameters and network organizations are provided in the supplementary. We would release our codes in <https://github.com/Tianxinhuang/DeFace.git>.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: Our codes and data would be released in <https://github.com/Tianxinhuang/DeFace.git>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: Training, test details are described in Sec. 4.1, and Sec. A.2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[No\]](#)

Justification: We follow existing related works for the setting of error bar.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We have provided that in the supplementary.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We have reviewed that and claim we conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[Yes\]](#)

Justification: Similar as existing human face reconstruction works, our method, focusing on recovering more realistic face textures under shadows from self and external occlusions, may also lead to the ethical concerns about privacy, consent, and the spread of misinformation. Strengthening the management of the actual facial data collected may alleviate this situation.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper does not have such risks as we only use public datasets.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have cited them in the references.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We use and cite existing datasets in this work. Other assets including codes/models would be released in <https://github.com/Tianxinhuang/DeFace.git>.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not include such experiments.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: All our experiments about human faces are based on existing public datasets.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.