

Deep Portrait Lighting Enhancement with 3D Guidance

Fangzhou Han[†], Can Wang[†], Hao Du and Jing Liao[‡]

City University of Hong Kong




Figure 1: Portrait lighting enhancement results of our method. With 3D guidance, our method enables producing photo-realistic lighting-enhanced images.

Abstract

Despite recent breakthroughs in deep learning methods for image lighting enhancement, they are inferior when applied to portraits because 3D facial information is ignored in their models. To address this, we present a novel deep learning framework for portrait lighting enhancement based on 3D facial guidance. Our framework consists of two stages. In the first stage, corrected lighting parameters are predicted by a network from the input bad lighting image, with the assistance of a 3D morphable model and a differentiable renderer. Given the predicted lighting parameter, the differentiable renderer renders a face image with corrected shading and texture, which serves as the 3D guidance for learning image lighting enhancement in the second stage. To better exploit the long-range correlations between the input and the guidance, in the second stage, we design an image-to-image translation network with a novel transformer architecture, which automatically produces a lighting-enhanced result. Experimental results on the FFHQ dataset and in-the-wild images show that the proposed method outperforms state-of-the-art methods in terms of both quantitative metrics and visual quality.

CCS Concepts

• Computing methodologies → Computational photography; Image processing;

1. Introduction

In portrait photography, lighting is one of the most important elements for photo quality and aesthetics. Professional portrait photographers can capture desirable illumination of their subjects by using specialized equipment, such as flashes and reflectors. How-

ever, most casual photographers are not able to create such compelling photographs and their photos taken in poor lighting conditions may suffer from overexposure, underexposure, low contrast, and high ISO noise problems. To better emulate professional-looking portraits, we aim to develop a deep learning method for enhancing the lighting of portrait photos taken by casual users.

Previous deep learning methods for lighting enhancement fall into two categories. One is to learn an image-to-image translation network from the source distribution of bad lighting images to the

[†] equal contribution
[‡] corresponding author

target distribution of enhanced lighting images. Generative adversarial networks (GANs) are often served as the backbone networks because of their excellent performance in synthesizing images [JGL*19, WZF*19]. Another category is to reformulate this task as an image-specific curve estimation problem [YS12, GLG*20]. By learning the parameters of curve functions from input images, pixel-wise adjustments can be made to obtain enhanced results. Despite their major progress on improving the overall brightness of input images, these generic methods fail to generate delicate lighting enhancement effects on portraits. This is because these methods are limited in the 2D image domain, ignoring vital 3D information, such as face geometry and lighting directions.

To fully take advantage of the 3D information, we present a 3D-guided portrait lighting enhancement framework with two stages. In the first stage, a **lighting parameter correction network** estimates the original lighting parameters from the input image and predicts the corrected parameters to adjust the original lighting as well. A differentiable renderer [LLCL19] is adopted to render the reconstructed 3D morphable model (3DMM) of the input with lighting parameters, while enabling the network to receive supervision from 2D images. As a result, a rendered portrait image with corrected lighting is obtained to serve as 3D guidance in the second stage. In the second stage, an **image enhancement network** learns a mapping from the bad lighting image distribution to the enhanced lighting image distribution, conditioned on the 3D guidance provided earlier. The guidance image, containing shading and texture information under the corrected lighting, is greatly helpful for image enhancement, and the major challenge is how to exploit semantically meaningful correlations between the input image and the guidance. Convolution Networks (CNN) models are less effective in this task as the local inductive priors of convolution operation make it difficult to exploit long-range correlations. In contrast, the transformer architecture [VSP*17] abandons the baked-in local inductive prior and thus supports global interactions via the dense attention module. Therefore, we adopt the transformer structure in our image enhancement network. Unlike many visual transformer tasks that only consider self-correlations in the input [CMS*20, DBK*20], we model the correlations between the input and the guidance by setting the guidance as query and the input as key. This helps recover missing details of the input, especially in underexposure and overexposure regions, as missing details can be borrowed from other parts of the face by leveraging long-range correlations.

Besides leveraging 3D information, another challenge of portrait lighting enhancement is how to collect bad-and-enhanced lighting image pairs to support training. To the best of our knowledge, current public image light enhancement datasets [WWYL18, CGZ18] are for scenes rather than for portraits. Though some face relighting datasets have been proposed [GMC*10, ZHSJ19], their data contains portraits with various lighting positions yet lacks explicit definition of enhanced or desirable lighting, which is inappropriate for the task. Furthermore, collecting real data is expensive and laborious as it requires a large number of identities as well as professional light stage system [MGAD19]. To this end, we propose an efficient data synthesis method. We first ask volunteers to vote for visually pleasant images taken under good lighting conditions from the FFHQ dataset [KLA19] as our target images. To create training

pairs, each target image is then degraded to a bad lighting condition by extrapolating its lighting parameters away from the center of the target dataset, followed by post-processing and manual screening. We demonstrate the model trained on our synthesized dataset generalizes well on in-the-wild images.

Our technique contributions are summarized as follows:

- We design a two-stage framework for portrait lighting enhancement with 3D guidance, which outperforms previous 2D approaches by exploiting 3D facial information.
- A transformer-based image-to-image translation network is proposed to model long-range correlations between the guidance image and the input image, which further improves the performance of our method in lighting enhancement.
- We propose an efficient data synthesis method that can produce photo-realistic training pairs for the portrait lighting enhancement task. This dataset will be released for research purposes.

2. Related Work

2.1. Image Enhancement

Image enhancement aims to beautify images based on certain standards, such as from low-light to normal-light or from low dynamic range to high dynamic range. A series of methods use reinforcement learning [HHX*18, YLZ*18] or GAN [DLT18] to learn the best adjustment parameters, including brightness, contrast, curves, etc. For example, Hu *et al.* [HHX*18] consider these basic image exposure adjustment operations as basic action units, and use reinforcement learning to solve for the optimal action sequence and parameters of each action to adjust the exposure of the input image. On the basis of [HHX*18], Yu *et al.* [YLZ*18] use semantical masks to adjust the exposure by regions. EnhanceGAN [DLT18] uses GAN for weakly supervised image enhancement. Its generator performs a global color adjustment, and the discriminator determines whether the results are enhanced images or not. This kind of methods is explainable and fast, yet they can only adjust the brightness of pixels, and cannot recover missing details in the underexposure or overexposure regions.

The second series of methods, such as [CWKC18, ZX20] directly learn the mapping between the two image domains and generate the enhanced target image by image-to-image translation networks. Chen *et al.* [CWKC18] use a two-stage GAN similar to CycleGAN [ZPIE17] to perform image enhancement. The generator uses the structure of U-Net [RFB15] and takes global features as input to reveal high-level information and to determine local adjustments for individual pixels. Comparing to learning curves or parameters, directly generating images can enhance images with greater flexibility.

These two categories of methods are limited in the 2D image domain without considering 3D information. So they are feasible when enhancing the overall brightness or color tone of generic images, but it would lead to unnatural effects when enhancing the lighting of faces, which are sensitive to 3D shapes. In our method, 3D facial guidance is involved and proved useful.

2.2. Face Relighting

Portrait relighting aims to change the lighting condition of face images given a group of arbitrary lighting parameters. A series of methods [SBT*19,ZTP*20] try to solve the problem directly on image level with image-to-image translation. Sun *et al.* [SBT*19] use an encoder-decoder structure based on U-Net [RFB15] to implicitly express the unlighted intermediate state and let the network learn the geometry and reflectance information. Zhang *et al.* [ZTP*20] use two GridNet [FEF*17] to remove shadows introduced by external objects, and then to soften shading as well as shadows projected by facial features. Both methods use Light Stage to collect data as it can produce image sets with single lighting sources and simulates arbitrary lighting environments, yet building such a dataset is expensive and time-consuming. There are some other work trying to construct dataset without Light Stage by reconstructing face models [ZHSJ19] or restoring images of Light Stage through color gradient images [MHP*19].

Besides relighting portraits on image level, some work first restore the albedo of the input images along with the corresponding normal or mesh, and then render corresponding images under different illuminations. Shu *et al.* [SYH*17] propose using GAN to end-to-end infer disentangled intrinsic facial attributes, such as shape, albedo and lighting. SfSNet [SKCJ18] also decomposes the image into shape, reflection and lighting, with a residual block structure to learn the relation between high-frequency variation and physical attributes. Nestmeyer *et al.* [NLML20] implement delighting and relighting by restoring albedo of the input image through the diffuse physics-based image generation model with Light Stage data.

Unlike the relighting task that requires users to input lighting parameters, our method automatically predicts the lighting parameters to correct the original lighting of input image. Meanwhile, as the reconstruction and rendering process will lead to unavoidable distortion, we do not use the rendered results with corrected lighting as our final outputs like the second category of methods. Instead, we use it to guide our image-to-image translation network for generating more photo-realistic lighting-enhanced images.

2.3. Differentiable Renderer

Differentiable renderer takes shape, texture, lighting parameters, and camera pose parameters as input and outputs rendered image and depth image. The most important feature of differentiable renderers is that it makes it possible to calculate the gradient of the rendered image to the input variables. The key challenge is how to make the rasterization process differentiable. To solve this problem, [LB14,KUH18,GCM*18] try to approximate the gradients in an inverse manner, while [RRR*15,LLCL19,CLG*19] try to simulate the forward rasterization process.

Differentiable renderers are widely used in various applications, including object reconstruction [YYY*16,TZEM17], human pose estimation [PZZD18], light source estimation [NJR18], and so on. In lighting correction stage of the proposed method, we integrate the differentiable renderer after the lighting parameter correction network to enable applying image-level loss to 3D mesh and light-

ing parameters. It serves as a bridge to connect the deep neural network and the reconstructed 3D information.

2.4. Transformer

Transformer is first proposed in natural language processing (NLP) [VSP*17], whose success makes it gradually popular in computer vision areas. For example, DETR [CMS*20] leverages the transformer as the backbone to cope with the object detection problem. ViT [DBK*20] first applies transformer for image recognition. Wan *et al.* [WZCL21] adopt transformer to address high-fidelity free-form pluralistic image completion and have achieved visually pleasing results. In these works, transformer architecture shows promising performance because of its strong capability to model global image structures by building dense correlations. However, these works only consider modeling self correlations involved in the input features. In contrast, our method aims to model dense correlations between the input and a guidance image to exploit the texture and light information involved in the guidance to help enhance the input image.

3. Method

Our method aims to enhance the lighting of input portrait images taken under some unsatisfactory lighting conditions. Considering that the shape of human faces is rich in geometric details and the shading on the face is also delicate, we do not follow previous methods which model this task as an image-to-image translation problem and directly learn a mapping from the bad lighting image distribution $P(I_s)$ to target light-enhanced image distribution $P(I_r)$, defined as $F : P(I_s) \rightarrow P(I_r)$. Instead, we achieve it in two stages to incorporate 3D information in learning. Specifically, we first learn a mapping from $P(I_s)$ to the distribution of the 3D guidance I_g , defined as $F : P(I_s) \rightarrow P(I_g)$, and then we learn a mapping from $P(I_s)$ to $P(I_r)$ conditioned on the 3D guidance, defined as $P(I_s|I_g) \rightarrow P(I_r)$.

Figure 2 shows our framework. Given an input image I_s , in the lighting correction stage (Section 3.1), its 3D mesh and albedo map are first reconstructed by a 3DMM reconstruction block **C**. And then the lighting parameter correction network **L** combined with a differentiable renderer **R**, is learned to predict the corrected lighting parameters, and yields a guidance image I_g containing a face rendered with corrected lighting, for the next stage. The image enhancement stage (Section 3.2) is based on a GAN network that consists of a transformer encoder **T**, a generator **G**, and a discriminator **D**. It takes the guidance I_g as a condition to modify the shading of the input face image I_s and eventually generate an enhanced result image, which is visually pleasing and photorealistic with good lighting.

3.1. Lighting Correction Stage

To include 3D facial information into our framework, we first reconstruct the shape of the input face and disentangle the albedo information from the texture. Following the idea of [FSGF16] which matches 3D models to portraits, we adopt a reconstruction network from [DYX*19], which takes the input image and use a neural network to regress identity coefficients α , expression coefficients β ,

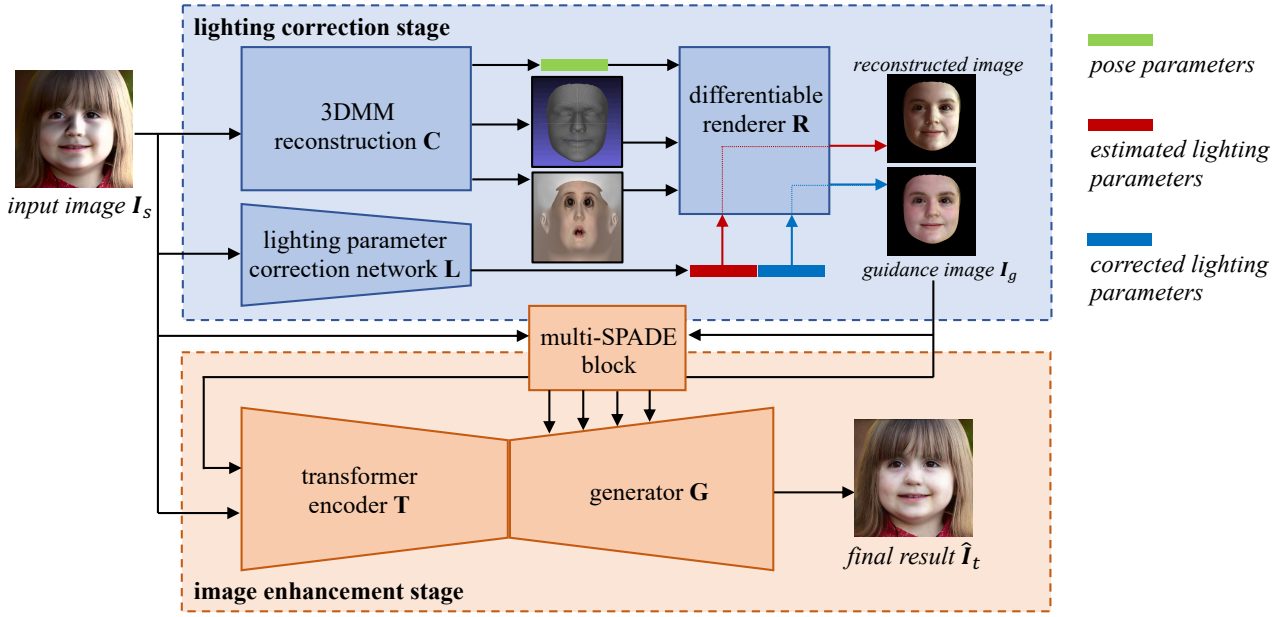


Figure 2: The schematic illustration of our method.

and texture coefficients ζ of a 3DMM face model [PKA*09] as well as the pose parameters γ corresponding to the input. The face shape \mathbf{S} and the albedo texture \mathbf{A} can be presented as:

$$\begin{aligned} \mathbf{S} &= \mathbf{S}(\alpha, \beta) = \bar{\mathbf{S}} + \mathbf{B}_{id}\alpha + \mathbf{B}_{exp}\beta \\ \mathbf{A} &= \mathbf{A}(\zeta) = \bar{\mathbf{A}} + \mathbf{B}_t\zeta \end{aligned} \quad (1)$$

where $\bar{\mathbf{S}}$ and $\bar{\mathbf{A}}$ are the average face shape and texture, \mathbf{B}_{id} , \mathbf{B}_{exp} , and \mathbf{B}_t are the PCA bases of identity, expression, and texture, respectively. As shading is not considered during the formation of the texture bases, the interpolation result of the texture can only represent albedo information regardless of the input illumination, resulting in a desired disentangle effect.

After reconstruction, we train a lighting correction network to predict lighting parameters corresponding to an enhanced lighting condition. To back-propagate gradients from loss functions defined on 2D images to the network, we use a differentiable renderer [LLCL19] to render the shape and albedo texture of the input image with lighting parameters obtained from the network. The renderer \mathbf{R} replaces the traditional rendering steps of rasterization and hidden face removal with a differentiable aggregate function. For i^{th} pixel, its color can be represented as:

$$I^i = \sum_j w_j^i C_j + w_b^i C_b,$$

where C_j is the color of j^{th} triangle, C_b is the color of the background and w^i are weights corresponding to the i^{th} pixel. The weights are negatively correlated with the distance between the pixel and the triangle, as well as the depth of the triangle.

For network \mathbf{L} , a good prediction requires accurately estimating the input lighting and then mapping the input lighting to its correction target. Experiments show that training the network to learn

the two steps simultaneously and to directly output the corrected lighting is difficult. Therefore, we design a bi-branch pipeline to explicitly learn the estimation and the mapping with supervision from input and target images, respectively. With the bi-branch design, the lighting parameter correction network encodes two sets of parameters, the estimated lighting parameters ϵ_{est} and the difference between the estimated lighting parameters and the corrected lighting parameters δ_{SH} , with which the corrected lighting parameters ϵ_{crt} can be calculated by $\epsilon_{crt} = \epsilon_{est} + \delta_{SH}$. Both ϵ_{est} and ϵ_{crt} are input to the differentiable renderer along with the reconstruction results to obtain the reconstructed image $\hat{\mathbf{I}}_s = \mathbf{R}(\mathbf{S}, \mathbf{A}, \gamma, \epsilon_{est})$ and the guidance image with corrected illumination $\mathbf{I}_g = \mathbf{R}(\mathbf{S}, \mathbf{A}, \gamma, \epsilon_{crt})$. The loss function can be written as:

$$\mathcal{L}_{light} = |\hat{\mathbf{I}}_s - \mathbf{I}_s| + \lambda_{crt} |\mathbf{I}_g - \mathbf{I}_t| \quad (2)$$

where λ_{crt} is a weight parameter.

Although the corrected lighting parameters are obtained and used in the rendering process, the rendered image cannot be directly used as the enhancement result, as there is unavoidable distortion in the reconstruction process and the lighting model in the renderer cannot perfectly simulate the real lighting effects. Therefore, we use an image-to-image translation network to guarantee the reality of the enhancement result, while using the rendered image as a guidance to provide the translation process with realistic shading and texture information.

3.2. Image Enhancement Stage

Given the input low-light face image \mathbf{I}_s and the guidance image \mathbf{I}_g generated in the lighting correction stage, the image enhancement stage outputs the desired enhanced face image. As mentioned in Section 1, to better exploit the visual relations between \mathbf{I}_s and \mathbf{I}_g

is the key to reconstruct the desired enhanced face image. We apply two modifications on the basis of pix2pix [IZZ17], the transformer encoder and the multi-SPADE block, to make better use of the guidance from the lighting correction stage and generate visually-pleasing results.

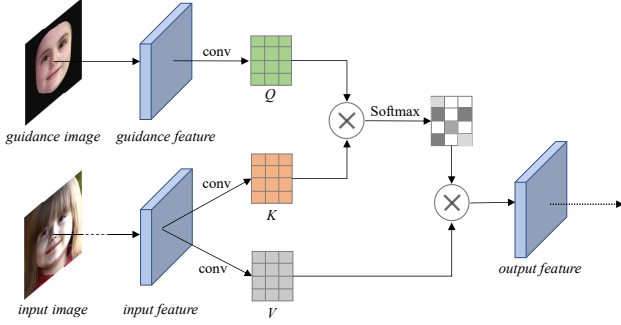


Figure 3: Transformer block.

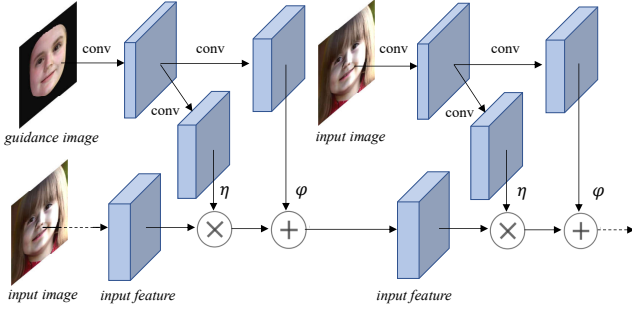


Figure 4: Multi-SPADE block. η and ϕ are modulation parameters.

Transformer Encoder: Traditional transformers capture self correlations by calculating Q (Query), K (Key), and V (Value) using the input feature, but we want to model the relations between the input and the guidance. Therefore, we propose a transformer block as shown in Figure 3. In the proposed block, the input feature is extracted from the input image and the guidance \mathbf{I}_g will first be warped into the same dimension as the input feature. Then Q is calculated from $\text{Warp}(\mathbf{I}_g)$, while K and V are calculated from the input feature. The attention map is obtained by performing a matrix multiplication between Q and K and applying a softmax operation to its result. By multiplying V and the attention map, we acquire the output feature.

Multi-SPADE Block: SPADE (Spatially-Adaptive Normalization) [PLWZ19] is an effective layer in a generator for synthesizing photo-realistic images given an input semantic layout. Unlike the original SPADE block which only injects the guidance image as condition and totally ignores the semantics involved in the original input image, we adopt a Multi-SPADE block from [MWSL20] to make the generated image better preserve semantics-consistence with the input. Specifically, we cascade two SPADE blocks, one takes \mathbf{I}_g as input and the other one takes \mathbf{I}_s as input, as shown in

Figure 4. Such a cascaded block ensures the generated image be consistent with \mathbf{I}_g at light level and simultaneously be semantics-consistent to the input \mathbf{I}_s .

There are three loss functions used in training the enhancement network. We use the adversarial loss to minimize the distribution distance between the ground-truth and output normal light distributions. For a face image, some local areas require much attention to adaptive but a global discriminator fail to provide the desired adaptivity. Similar to PatchGAN [IZZ17], we use a multi-scale patch discriminator $\mathbf{D} = \{\mathbf{D}_1, \mathbf{D}_2, \mathbf{D}_3\}$ to discriminate the real and fake images at different scales. Thus, the adversarial loss is defined as follows:

$$\min_{T,G} \max_{D_1, D_2, D_3} \sum_{k=1,2,3} \mathcal{L}_{\text{GAN}}(T, G, D_k) \quad (3)$$

$$\mathcal{L}_{\text{GAN}} = E_{\mathbf{I}_s}[\log D_k(\mathbf{I}_s)] + E_{(\mathbf{I}_s, \mathbf{I}_g)}[\log(1 - D_k(G(T(\mathbf{I}_s), \mathbf{I}_g)))] \quad (4)$$

To make the training process robust, we also adopt a feature matching loss [WLZ*18] between different layer features extracted by the discriminator of the real and fake images.

$$\mathcal{L}_{\text{FM}}(T, G, D_k) = E_{(\mathbf{I}_s, \mathbf{I}_g, \mathbf{I}_f)} \sum_{i=1}^L \frac{1}{N_i} \left[\left\| D_k^{(i)}(G(T(\mathbf{I}_s), \mathbf{I}_g)) - D_k^{(i)}(\mathbf{I}_f) \right\|_1 \right] \quad (5)$$

where i means the i^{th} layer, L represents the number of layers, and N_i is the number of parameters of the i^{th} layer.

Also, a perception loss is used to further improve the performance:

$$\mathcal{L}_{\text{percep}}(T, G) = \sum_{i=1}^K \frac{1}{M_i} \left[\left\| F^{(i)}(\mathbf{I}_s) - F^{(i)}(G(T(\mathbf{I}_s), \mathbf{I}_g)) \right\|_1 \right] \quad (6)$$

where i means the i^{th} layer, K represents the number of layers, M_i is the number of parameters of the i^{th} layer, and $F^{(i)}$ is the feature of the i^{th} layer of a VGG19 network.

In our work, these losses can be optimized jointly and the total objective can be defined as:

$$\min_{T,G} \max_{D_1, D_2, D_3} \sum_{k=1,2,3} \mathcal{L}_{\text{GAN}}(T, G, D_k) + \lambda_{\text{FM}} \sum_{k=1,2,3} \mathcal{L}_{\text{FM}}(T, G, D_k) + \lambda_{\text{percep}} \mathcal{L}_{\text{percep}}(T, G) \quad (7)$$

where λ_{FM} and λ_{percep} are weight coefficients.

4. Experiments

4.1. Dataset Construction

As mentioned in Section 1, currently there exists no high-quality dataset for portrait lighting enhancement. To build such a dataset, we first ask volunteers to manually select a group of portrait images that are believed to have good lighting from FFHQ dataset [KLA19]. This group of images forms the enhanced lighting distribution as shown by the blue oval of Figure 5. They will be used as the ground truth for training to guarantee the outputs to be more natural than using synthetic data. The corresponding input images are degraded from target images. Specifically, we

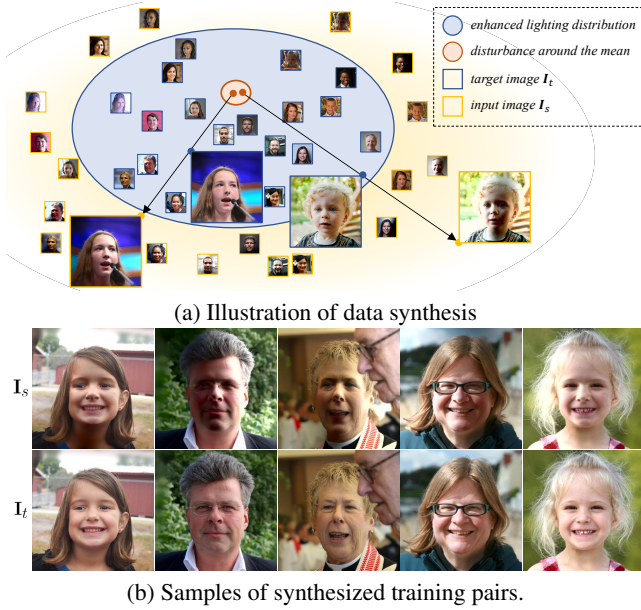


Figure 5: Dataset construction process and samples. I_s is the degraded low-light image and I_t is the pseudo ground-truth image.

first estimate SH parameters of all target images following the iterative optimization method used in [ZHSJ19] and calculate their mean value \overline{SH} . Given the SH parameters of i^{th} target image, denoted as $SH_{target,i}$, we randomly sample an ideal SH parameter within a small radius of the mean, denoted as \overline{SH}_i , and extrapolate $SH_{target,i}$ away from \overline{SH}_i to get the degraded lighting: $SH_{input,i} = SH_{target,i} + \lambda_{sh} * (SH_{target,i} - \overline{SH}_i)$ where $SH_{input,i}$ represents the SH parameters of the i^{th} input image and $\lambda_{sh} \in [1, 2]$ is a randomly sampled extrapolation parameter. The randomness in ideal point and extrapolation parameter sampling introduce lighting variations in dataset. With degraded lighting parameters $SH_{input,i}$, the i^{th} degraded image is generated by the data synthesis pipeline introduced in [ZHSJ19]. After degradation, to make the synthetic data as realistic as possible, a blurring filter is applied to the degraded image, followed by a Gaussian noise to simulate detail loss and noise artifacts in poorly-relit images. The strength of the blurring filter and the Gaussian noise is not uniform but in positive correlation with the shading darkness. And thus, the training pair of the i^{th} input and target images is formed.

To further prevent unrealistic synthetic images, volunteers are asked to perform manual filtering, after which 6828 out of 20000 image pairs are selected to form the final dataset. Among the 6828 pairs, 6371 pairs are splitted as the training set, while 457 pairs are splitted as the testing set. Some samples from the dataset are shown in Figure 5.

This dataset construction method is peculiarly efficient and the dataset is easy to extend. Training with the dataset, networks will learn to adjust the lighting of the input images towards the elected standard, yet the results will still preserve the lighting characteristic of the input images, instead of undesirably mapping all lighting conditions to a normalized one. Although the domain gap

between synthetic data and real data is inevitable, even with the above-mentioned delicate designs, the testing results on real images demonstrate that our synthetic dataset enables networks to generalize well on real poorly-relit images.

4.2. Implement Details

In the lighting correction stage, L is modified from the structure of ResNet50 [HZRS16], which concatenates the $I_s \in R^{512 \times 512 \times 3}$ and $I_g \in R^{512 \times 512 \times 3}$ as input and outputs a matrix of size $B \times 54$, where $B = 4$ is the batch size and 54 is the vector length representing 2 groups of 3×9 parameters, ϵ_{est} and δ_{SH} . The numbers of SH we used are 27 (9 for each channel of RGB). During the training of the lighting correction stage, we optimize the network parameters using Adam with $\beta_1 = 0.95$ and $\beta_2 = 0.90$. The initial learning rate is 0.001 and linearly decays for every 15 epochs by a factor of 0.7, with total epochs being 120. The value of λ_{crt} is set to 1.

In the image enhancement stage, E has 7 convolution layers and 4 transformer blocks in total. It takes $I_s \in R^{512 \times 512 \times 3}$ and $I_g \in R^{512 \times 512 \times 3}$ as input and outputs a feature map of shape $B \times 1024 \times 8 \times 8$, where $B = 16$ is the batch size. G has 7 upsampling layers with scale factor 2, and each is accompanied by a Multi-SPADE block. It reconstructs an image of the same size as the input. D has a similar architecture as PatchGAN [Izze17] and downsampling factors for D_1, D_2, D_3 are 1, 2, and 4 respectively. We optimize the network parameters using Adam with $\beta_1 = 0.95$ and $\beta_2 = 0.90$. The initial learning rate is 0.0002 and linearly decays to 0 after 50 epochs. Values of λ_{FM} and λ_{percep} are all set as 10. Besides, in Equation 5, L equals 5 including 1 input layer, 2 intermediate layers, and 1 output layer. And in Equation 6, K equals 5 which means we apply the perceptual loss to the output of the first 5 layers of the VGG19 network.

4.3. Ablation Study

	Albedo	w/o BB	w/o SH	Ours
PSNR	9.868	9.994	10.17	10.25
SSIM	0.2774	0.2798	0.2813	0.2842

Table 1: Quantitative evaluation results of ablation study on the lighting correction stage.

Lighting correction stage

To demonstrate the effectiveness of the bi-branch design and the SH lighting model of the lighting correction stage, we conduct an ablation experiment including two comparisons with three results:

- Est. lighting w/o **SH**: the estimated lighting of the method using the default ambient and directional lighting model instead of SH lighting model (**SH**) in the differentiable renderer.
- Crt. lighting w/o **BB**: the corrected lighting of the method not using the bi-branch design (**BB**) but use a single-branch to directly estimate the corrected lighting.
- Crt. lighting w/o **SH**: the corrected lighting of the method using the default ambient and directional lighting model in the differentiable renderer.

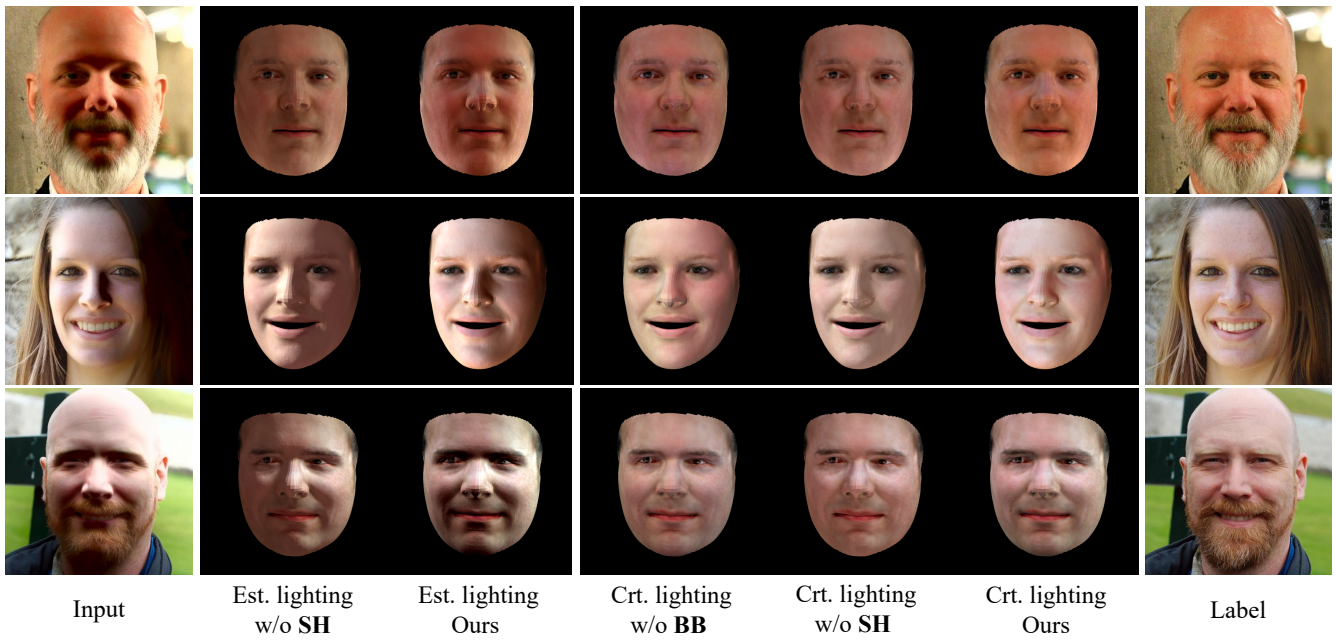


Figure 6: Example results of ablation study on the lighting correction stage.

Some example results are given in Figure 6. To only output the corrected lighting parameters under the supervision of the ground truth is the most straightforward design of the lighting parameter correction network. However, in such a single-branch design without explicitly estimating the original lighting of the input image, it will cause ambiguities and lead to results sometimes being too close to the input (2nd example) but sometimes being too flat and losing the original lighting patterns (1st and 3rd examples). In our bi-branch design, the network can learn to estimate and adjust the lighting from separate supervision, which leads to better performance.

Another issue that influences the lighting correction performance is that the original differentiable renderer supports only ambient lighting model and direction lighting model, which sometimes cannot perfectly simulate the lighting effect of the image. In our implementation, we modify the lighting model to SH lighting, which is more flexible in simulating lighting directions and colors. In Figure 6 we can see that the SH lighting model is able to simulate more complex lighting condition and better fit the skin color of the input image, therefore leads to better corrected lighting results.

We also conduct quantitative evaluation by calculating Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity (SSIM) against target images. The results in Table 1 shows that our final method obtains the highest score on both metrics, which numerically proves the effectiveness of the bi-branch design with SH lighting model.

Image enhancement stage

To demonstrate the effectiveness of each part of the image enhancement stage, we conduct an ablation experiment to compare our method with four methods:

	Input	w/o I_g	w/o SH	w/o T	w/o MS	Ours
PSNR	18.25	24.21	26.13	26.26	26.93	27.15
SSIM	0.8816	0.9030	0.9115	0.9163	0.9188	0.9211

Table 2: Ablation study on the image enhancement stage.

- w/o I_g : the method that removes the 3D guidance injection I_g in the retouching branch. The input of this method is the image I_s only. Specifically, we set I_s as the input to the encoder T without using Multi-SPADE blocks and the guidance I_g .
- w/o SH: the method that uses guidance if point lighting rather than SH lighting in the relighting branch.
- w/o T: the method that adopts traditional convolution layers as encoder rather than the transformer encoder.
- w/o MS: the method that uses single SPADE block which only injects I_g without considering semantics consistency using Multi-SPADE blocks.

First, qualitative experiments are conducted, as shown in Figure 7. We observe that method w/o I_g has obtained the worst results of all ablations, which evaluates the necessity of using 3D guidance for image enhancement. Sometimes method w/o SH shows enhancements compared to the input, but it is still not comparable to ours. For example, the lower jaw of the 3rd example cannot be enhanced. This is because the point lighting model fails to estimate the lighting accurately and an inaccurate guidance cannot provide a right training direction of the model. Method w/o T shows severe artifacts, for example the nose of the girl in the second row. This shows that the low-light regions cannot be well-enhanced without the transformer's ability of modeling correlations between the input



Figure 7: Visual results of ablation study on the image enhancement stage.

and the guidance. Method w/o MS outperforms other ablations, but the absence of texture details makes its results non-photorealistic. For example, we cannot see enough details of the 3rd and 4th examples. Depending only on single SPADE block fails to preserve enough texture details. In contrast, our method gives superior results, which demonstrates the superior performance and generalization ability of each component of our method.

We also conduct quantitative experiments in Table 2. In contrast, the PSNR and SSIM metrics of our final method are all better than other settings, which validates the effectiveness of our method by imposing 3D guidance, SH lighting model, transformer encoder, and Multi-SPADE block. We also find more serious degradation without using our proposed 3D guidance than others. This reveals the importance of providing 3D information to the enhancement process. This 3D guidance not only guides the training process in the right direction, but also provides missing texture details of the input, especially in the underexposure and overexposure regions. Though the 3D guidance contributes more to the improvement of our method, other components are also crucial to the success of the final pleasing results.

4.4. Comparisons on Synthetic Data

To evaluate the performance of the proposed method, we compare it with state-of-the-art image lighting enhancement methods, EnlightenGAN [JGL*19] and Zero-DCE [GLG*20], and portrait relighting methods, DSIPR (Deep Single-Image Portrait Relighting) [ZHSJ19] and SMFR [HZS*21]. We train the two image lighting enhancement methods on our dataset with author-released code,

as they were previously trained for general images, and use the pre-trained model for the portrait relighting methods. As the portrait relighting methods require lighting parameters as input, we feed them with the lighting parameters predicted by our first stage.

We first compare these methods on synthetic data, both qualitatively and quantitatively, results of which are shown in Figure 8 and Table 3. In terms of very dark areas in Figure 8, all of the comparison methods cannot produce desired results. EnlightenGAN and Zero-DCE achieve better results than the relighting methods DSIPR and SMFR, but they still perform badly in generating natural face textures and pleasing lighting conditions compared to ours. Numerical results in Table 3 also support this observation. DSIPR and SMFR report the worst two results on both PSNR and SSIM metrics, while EnlightenGAN and Zero-DCE show relatively better results. However, it is obvious that the proposed method has achieved the best in terms of PSNR and SSIM metrics. This proves the superiority of the proposed method in making use of 3D information and correlations between the 3D guidance and the input.

	Input	DSIPR	SMFR	EnlightenGAN	Zero-DCE	Ours
PSNR	18.25	23.68	22.77	26.15	26.32	27.15
SSIM	0.8816	0.8970	0.8893	0.9088	0.9102	0.9211

Table 3: Comparison with state-of-the-art models.



Figure 8: Qualitative comparison with state-of-the-art methods on the FFHQ dataset on the image enhancement stage.

4.5. Comparisons on In-the-Wild Data

To evaluate the generalization ability of our method, we compare to related works on the in-the-wild images, as shown in Figures 9. All methods can enhance the lighting condition overall, but these related works still show obvious worse results than ours. Both EnlightenGAN and Zero-DCE show unrealistic skin colors after enhancing the input. For example, EnlightenGAN makes the skin color more yellow for the face of the 3rd example. Besides, Zero-DCE tends to generate over-exposure and unnatural images. In contrast, our method can successfully enhance the input and synthesize photo-realistic light-enhanced image without destroying the face color. This is because the 3D information such as face geometry, face texture and lighting directions and intensities involved in the guidance can guide our networks to generate desired enhanced images, while EnlightenGAN and Zero-DCE do not use any guidance. As for relighting-based methods DSIPR and SMFR, they cannot reconstruct face details. Though they use SH parameters as guidance to train their models, their performances are still worse than ours. This is because these SH parameters only capture lighting information, while the 3D guidance we used cannot only model target lighting but also capture face geometry, shading, and face texture. To conclude, all these methods generate unsatisfactory visual results when it comes to both brightness and naturalness. Compared to these methods, our method successfully enhances the low-right areas and also preserves facial texture details without inducing artifacts. More results please refer to our supplemental materials.

As there is no ground truth for the enhancement of in-the-wild images, we also conduct a user study to perform numerical eval-

uation among our method and the related works. Specifically, we randomly choose 25 bad-lighting real images as test cases and apply different methods to enhance these images. Then for each case, we show the input image and the enhanced images generated by different methods and ask 22 participants to rank the five results from the highest quality (score 4) to the lowest quality (score 0). During ranking, the participants are instructed to consider whether the lighting condition of the images are properly and naturally enhanced without any artifacts or noise. The statistics results are shown in Figure 10. Figure 10 (a) demonstrate the average ranking scores and standard deviations of each method, where our method is ranked in the first place (3.54 ± 0.79). Figure 10 (b) shows the average ranking scores for each methods by every test case, and Figure 10 (c) shows the average ranking scores for each methods by every participant, from both of which it can be seen that our method has a stable and outstanding performance over different test cases judged by different individuals.

4.6. Image Harmonization

To increase the stability for training the networks, we detect and crop the face region when pre-processing the dataset. To extend our method to handle images with random face size and location, we train an image harmonization network [KS20] to adjust the brightness of the background according to the enhanced face region. In Figure 11 we compare our harmonized enhancement results (Ours) with the inputs and the enhancement results directly padded to the original background (w/o harmonization). It can be seen that the discontinuity on the background or the hair region of the images




Figure 9: Qualitative comparison with state-of-the-art methods on in-the-wild images on the image enhancement stage.

after the enhancement is smoothed by the harmonization network, which effectively enriches the working scenarios of our method.

5. Conclusions

In this paper, we present a deep face lighting enhancement method with 3D guidance. To support face light enhancement, we propose a method to synthesize more photo-realistic training pairs and use a 3DMM and a differentiable renderer to generate 3D guidance. We show that a transformer encoder can better model the relation between the guidance and the input low-light face image. Qualitative and quantitative experimental results prove the effectiveness of the proposed method in both poorly-lit synthetic and in-the-wild images.

Although it is proved that the lighting guidance is useful and more accurate and realistic guidance will lead to better enhancement performance, a limitation of current method is that it is unable to alter the lighting attributes, such as direction or strength, of the enhancement results with modifying the lighting guidance. This is because, unlike the relighting tasks where the lighting inputs (parameters or color maps) are only entangled with the training targets, the lighting guidance in our method also has correlation with the input image. Figure 12 shows some failure cases when providing the guidance with arbitrary lighting outside our enhanced lighting distribution. In future work, the framework can be modify to perform relighting function with generating multiple guidance-target pairs


for one input image. Another potential limitation is, because a limited number of SH coefficients are used, the synthesized training pairs cannot represent all kinds of illuminations, especially those with high-frequency components. Thus, our network may not be able to handle these challenging illuminations perfectly. This is a common problem for all lighting enhancement and relighting methods with synthetic training data. While real datasets captured in light stages or real environments would be helpful to avoid such problems, they are expensive and time-consuming to collect.

6. Acknowledgements


This work is supported by Huawei Ascend.

References


- [CGZ18] CAI J., GU S., ZHANG L.: Learning a deep single image contrast enhancer from multi-exposure images. *IEEE Transactions on Image Processing* 27, 4 (2018), 2049–2062. 2
- [CLG*19] CHEN W., LING H., GAO J., SMITH E., LEHTINEN J., JACOBSON A., FIDLER S.: Learning to predict 3d objects with an interpolation-based differentiable renderer. In *Advances in Neural Information Processing Systems* (2019), pp. 9609–9619. 3
- [CMS*20] CARION N., MASSA F., SYNNAEVE G., USUNIER N., KIRILLOV A., ZAGORUYKO S.: End-to-end object detection with transformers. In *European Conference on Computer Vision* (2020), Springer, pp. 213–229. 2, 3




(a) Average ranking score and standard deviation



(b) Ranking score per test case



(c) Ranking score per participant.

Figure 10: User study result analysis.**Figure 11:** Image harmonization results. Three images of each sample are the input, the light-enhanced image without harmonization, and our result using harmonization respectively.**Figure 12:** Failure cases when sampling the guidance with arbitrary lighting outside our enhanced lighting distribution.

[CWKC18] CHEN Y.-S., WANG Y.-C., KAO M.-H., CHUANG Y.-Y.: Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 6306–6314. 2

[DBK*20] DOSOVITSKIY A., BEYER L., KOLESNIKOV A., WEISENBORN D., ZHAI X., UNTERTHINER T., DEGHANI M., MINDERER M., HEIGOLD G., GELLY S., ET AL.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020). 2, 3

[DLT18] DENG Y., LOY C. C., TANG X.: Aesthetic-driven image enhancement by adversarial learning. In *Proceedings of the 26th ACM international conference on Multimedia* (2018), pp. 870–878. 2

[DYX*19] DENG Y., YANG J., XU S., CHEN D., JIA Y., TONG X.: 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* (2019), pp. 0–0. 3

[FEF*17] FOURURE D., EMONET R., FROMONT E., MUSELET D., TREMEAU A., WOLF C.: Residual conv-deconv grid network for semantic segmentation. *arXiv preprint arXiv:1707.07958* (2017). 3

[FSGF16] FRIED O., SHECHTMAN E., GOLDMAN D. B., FINKELSTEIN A.: Perspective-aware manipulation of portrait photos. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–10. 3

[GCM*18] GENOVA K., COLE F., MASCHINOT A., SARNA A., VLASIC D., FREEMAN W. T.: Unsupervised training for 3d morphable

model regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 8377–8386. 3

[GLG*20] GUO C., LI C., GUO J., LOY C. C., HOU J., KWONG S., CONG R.: Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 1780–1789. 2, 8

[GMC*10] GROSS R., MATTHEWS I., COHN J., KANADE T., BAKER S.: Multi-pie. *Image and vision computing* 28, 5 (2010), 807–813. 2

[HHX*18] HU Y., HE H., XU C., WANG B., LIN S.: Exposure: A white-box photo post-processing framework. *ACM Transactions on Graphics (TOG)* 37, 2 (2018), 1–17. 2

[HZRS16] HE K., ZHANG X., REN S., SUN J.: Deep residual learn-

- ing for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 770–778. 6
- [HZS*21] HOU A., ZHANG Z., SARKIS M., BI N., TONG Y., LIU X.: Towards high fidelity face relighting with realistic shadows. *arXiv preprint arXiv:2104.00825* (2021). 8
- [IZZE17] ISOLA P., ZHU J.-Y., ZHOU T., EFROS A. A.: Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017), pp. 1125–1134. 5, 6
- [JGL*19] JIANG Y., GONG X., LIU D., CHENG Y., FANG C., SHEN X., YANG J., ZHOU P., WANG Z.: Enlightengan: Deep light enhancement without paired supervision. *arXiv preprint arXiv:1906.06972* (2019). 2, 8
- [KLA19] KARRAS T., LAINE S., AILA T.: A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 4401–4410. 2, 5
- [KS20] KONSTANTIN SOFIUK POLINA POPENOVA A. K.: Foreground-aware semantic representations for image harmonization. *arXiv preprint arXiv:2006.00809* (2020). 9
- [KUH18] KATO H., USHIKU Y., HARADA T.: Neural 3d mesh renderer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 3907–3916. 3
- [LB14] LOPER M. M., BLACK M. J.: Opendr: An approximate differentiable renderer. In *European Conference on Computer Vision* (2014), Springer, pp. 154–169. 3
- [LLCL19] LIU S., LI T., CHEN W., LI H.: Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE International Conference on Computer Vision* (2019), pp. 7708–7717. 2, 3, 4
- [MGAD19] MURMANN L., GHARBI M., AITTALA M., DURAND F.: A dataset of multi-illumination images in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019), pp. 4080–4089. 2
- [MHP*19] MEKA A., HAENE C., PANDEY R., ZOLLHÖFER M., FANELLO S., FYFFE G., KOWDLE A., YU X., BUSCH J., DOURGARIAN J., ET AL.: Deep reflectance fields: high-quality facial reflectance field inference from color gradient illumination. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12. 3
- [MWSL20] MALLYA A., WANG T.-C., SAPRA K., LIU M.-Y.: World-consistent video-to-video synthesis. *arXiv preprint arXiv:2007.08509* (2020). 5
- [NJR18] NIETO G., JIDDI S., ROBERT P.: Robust point light source estimation using differentiable rendering. *arXiv preprint arXiv:1812.04857* (2018). 3
- [NLML20] NESTMEYER T., LALONDE J.-F., MATTHEWS I., LEHRMANN A.: Learning physics-guided face relighting under directional light. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 5124–5133. 3
- [PKA*09] PAYSAN P., KNOTHE R., AMBERG B., ROMDHANI S., VETTER T.: A 3d face model for pose and illumination invariant face recognition. In *2009 sixth IEEE international conference on advanced video and signal based surveillance* (2009), Ieee, pp. 296–301. 4
- [PLWZ19] PARK T., LIU M.-Y., WANG T.-C., ZHU J.-Y.: Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), pp. 2337–2346. 5
- [PZZD18] PAVLAKOS G., ZHU L., ZHOU X., DANIILIDIS K.: Learning to estimate 3d human pose and shape from a single color image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 459–468. 3
- [RFB15] RONNEBERGER O., FISCHER P., BROX T.: U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* (2015), Springer, pp. 234–241. 2, 3
- [RRR*15] RHODIN H., ROBERTINI N., RICHARDT C., SEIDEL H.-P., THEOBALT C.: A versatile scene model with differentiable visibility applied to generative pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision* (2015), pp. 765–773. 3
- [SBT*19] SUN T., BARRON J. T., TSAI Y.-T., XU Z., YU X., FYFFE G., RHEMANN C., BUSCH J., DEBEVEC P. E., RAMAMOORTHI R.: Single image portrait relighting. *ACM Trans. Graph.* 38, 4 (2019), 79–1. 3
- [SKCJ18] SENGUPTA S., KANAZAWA A., CASTILLO C. D., JACOBS D. W.: Sfsnet: Learning shape, reflectance and illuminance of faces in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 6296–6305. 3
- [SYH*17] SHU Z., YUMER E., HADAP S., SUNKAVALLI K., SHECHTMAN E., SAMARAS D.: Neural face editing with intrinsic image disentangling. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017), pp. 5541–5550. 3
- [TZEM17] TULSIANI S., ZHOU T., EFROS A. A., MALIK J.: Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017), pp. 2626–2634. 3
- [VSP*17] VASWANI A., SHAZEER N., PARMAR N., USZKOREIT J., JONES L., GOMEZ A. N., KAISER L., POLOSUKHIN I.: Attention is all you need. *arXiv preprint arXiv:1706.03762* (2017). 2, 3
- [WLZ*18] WANG T.-C., LIU M.-Y., ZHU J.-Y., TAO A., KAUTZ J., CATANZARO B.: High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2018), pp. 8798–8807. 5
- [WWYL18] WEI C., WANG W., YANG W., LIU J.: Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560* (2018). 2
- [WZCL21] WAN Z., ZHANG J., CHEN D., LIAO J.: High-fidelity pluralistic image completion with transformers. *arXiv preprint arXiv:2103.14031* (2021). 3
- [WZF*19] WANG R., ZHANG Q., FU C.-W., SHEN X., ZHENG W.-S., JIA J.: Underexposed photo enhancement using deep illumination estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 6849–6857. 2
- [YLZ*18] YU R., LIU W., ZHANG Y., QU Z., ZHAO D., ZHANG B.: Deepexposure: Learning to expose photos with asynchronously reinforced adversarial learning. In *Advances in Neural Information Processing Systems* (2018), pp. 2149–2159. 2
- [YS12] YUAN L., SUN J.: Automatic exposure correction of consumer photographs. In *European Conference on Computer Vision* (2012), Springer, pp. 771–785. 2
- [YY*16] YAN X., YANG J., YUMER E., GUO Y., LEE H.: Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In *Advances in neural information processing systems* (2016), pp. 1696–1704. 3
- [ZHSJ19] ZHOU H., HADAP S., SUNKAVALLI K., JACOBS D. W.: Deep single-image portrait relighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019), pp. 7194–7202. 2, 3, 6, 8
- [ZPIE17] ZHU J.-Y., PARK T., ISOLA P., EFROS A. A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision* (2017), pp. 2223–2232. 2
- [ZTP*20] ZHANG X. C., TSAI Y.-T., PANDEY R., ZHANG X., NG R., JACOBS D. E., ET AL.: Portrait shadow manipulation. *arXiv preprint arXiv:2005.08925* (2020). 3
- [ZX20] ZHOU Y., XIAO Q.: Gan-based facial attractiveness enhancement. *arXiv preprint arXiv:2006.02766* (2020). 2