

# Copy and Paste GAN: Face Hallucination from Shaded Thumbnails

Yang Zhang<sup>1,2,3</sup>, Ivor W. Tsang<sup>3</sup>, Yawei Luo<sup>4</sup>, Changhui Hu<sup>1,2,5</sup>, Xiaobo Lu<sup>1,2,\*</sup>, Xin Yu<sup>3,6</sup>

<sup>1</sup> School of Automation, Southeast University, China

<sup>2</sup> Key Laboratory of Measurement and Control of Complex Systems of Engineering, Ministry of Education, Southeast University, China

<sup>3</sup> Centre for Artificial Intelligence, University of Technology Sydney, Australia

<sup>4</sup> School of Computer Science and Technology, Huazhong University of Science and Technology, China

<sup>5</sup> School of Automation, Nanjing University of Posts and Telecommunications, China

<sup>6</sup> Australian Centre for Robotic Vision, Australian National University, Australia

## Abstract

Existing face hallucination methods based on convolutional neural networks (CNN) have achieved impressive performance on low-resolution (LR) faces in a normal illumination condition. However, their performance degrades dramatically when LR faces are captured in low or non-uniform illumination conditions. This paper proposes a Copy and Paste Generative Adversarial Network (CPGAN) to recover authentic high-resolution (HR) face images while compensating for low and non-uniform illumination. To this end, we develop two key components in our CPGAN: internal and external Copy and Paste nets (CPnets). Specifically, our internal CPnet exploits facial information residing in the input image to enhance facial details; while our external CPnet leverages an external HR face for illumination compensation. A new illumination compensation loss is thus developed to capture illumination from the external guided face image effectively. Furthermore, our method offsets illumination and upsamples facial details alternately in a coarse-to-fine fashion, thus alleviating the correspondence ambiguity between LR inputs and external HR inputs. Extensive experiments demonstrate that our method manifests authentic HR face images in a uniform illumination condition and outperforms state-of-the-art methods qualitatively and quantitatively.

## 1. Introduction

Human faces are important information resources since they carry information on identity and emotion

\*Corresponding author (xblu2013@126.com).

This work was done when Yang Zhang (zhangyang201703@126.com) was a visiting student at University of Technology Sydney.

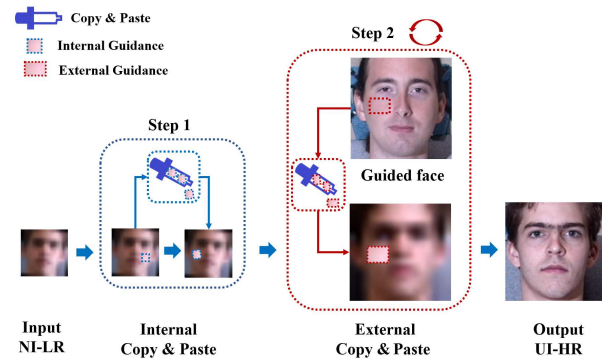


Figure 1: Motivation of CPGAN. Internal and external CPnets are introduced to mimic the Clone Stamp Tool. Internal CPnet copies good details to paste them on shadow regions. External CPnet further retouches the face using an external guided face from the UI-HR face dataset during the upsampling process to compensate for uneven illumination in the final HR face.

changes in daily activities. To acquire such information, high-resolution and high-quality face images are often desirable. Due to spacing distance and lighting conditions between cameras and humans, captured faces may be tiny or in poor illumination conditions, thus hindering human perception and computer analysis (Fig.2(a)).

Recently, many face hallucination techniques [32, 35, 41, 4, 5, 30, 37] have been proposed to visualize tiny face images by assuming uniform illuminations on face thumbnail databases, as seen in Fig.2(b) and Fig.2(c). However, facial details in shaded thumbnails become obscure in case of low/non-uniform illumination conditions, which leads to failures in hallucination due to inconsistent intensities. For instance, as Fig.2(d) shows, the hallucinated result generated by the state-of-the-art face hallucination method [33] is semantically and per-

ceptually inconsistent with the ground truth (GT), yielding blurred facial details and non-smooth appearance.

Meanwhile, various methods are proposed to tackle illumination changes on human faces. The state-of-the-art methods of face inverse lighting [6, 39] usually fit the face region to a 3D Morphable Model [29] by facial landmarks and then render illumination. However, these methods are unsuitable to face thumbnails, because facial landmark cannot be detected accurately in such low-resolution images and erroneous face alignment leads to artifacts in the illumination normalized results. This will increase difficulty to learn the mappings between LR and HR faces. Image-to-image translation methods, such as [12, 40, 22], can be an alternative to transfer illumination styles between faces without detecting facial landmarks. Due to the variety of illumination conditions, the translation method [40] fails to learn a consistent mapping for face illumination compensation, thus distorting facial structure in the output (Fig.2(e)).

As seen in Fig.2(f) and Fig.2(g), applying either face hallucination followed by illumination normalization or illumination normalization followed by hallucination produces results with severe artifacts. To tackle this problem, our work aims at hallucinating LR inputs under non-uniform low illumination (NI-LR face) while achieving HR in uniform illumination (UI-HR face) in a unified framework.

Towards this goal, we propose a Copy and Paste Generative Adversarial Network (CPGAN). CPGAN is designed to explore internal and external image information to normalize illumination, and to enhance facial details of input NI-LR faces. We first design an internal Copy and Paste net (internal CPnet) to approximately offset non-uniform illumination features and enhance facial details by searching for similar facial patterns within input LR faces for subsequent upsampling procedure. Our external CPnet is developed to copy illumination from a HR face template, then pass the illumination information to the input. In this way, our network learns how to compensate for illumination of inputs. To reduce the illumination transferring difficulty, we alternately upsample and transfer the illumination in a coarse-to-fine manner. Moreover, Spatial Transformer Network (STN) [13] is adopted to align input NI-LR faces, promoting more effectively feature refinement as well as facilitating illumination compensation. Furthermore, an illumination compensation loss is proposed to capture the normal illumination pattern and transfer the normal illumination to the inputs. As shown in Fig. 2(h), the upsampled HR face is not only realistic but also resembles the GT with normal illumination.

The contributions of our work are listed as follows:

- We present the first framework, dubbed CPGAN, to address face hallucination and illumination compensation together, in an end-to-end manner, which is optimized by the conventional face hallucination loss and a new illumination compensation loss.
- We introduce an internal CPnet to enhance the facial details and normalize illumination coarsely, aiding subsequent upsampling and illumination compensation.
- We present an external CPnet for illumination compensation by learning illumination from an external HR face. In this fashion, we are able to learn illumination explicitly rather than requiring a dataset with the same illumination condition.
- A novel data augmentation method, Random Adaptive Instance Normalization (RaIN), is proposed to generate sufficient NI-LR and UI-HR face image pairs. Experiments show that our method achieves photo-realistic UI-HR face images.

## 2. Related work

### 2.1. Face Hallucination

Face hallucination methods aim at establishing the intensity relationships between input LR and output HR face images. The prior works can be categorized into three mainstreams: holistic-based techniques, part-based methods, and deep learning-based models.

The basic principle of holistic-based techniques is to represent faces by parameterized models. The representative models conduct face hallucination by adopting linear mapping [27], global appearance [17], subspace learning techniques [15]. Then, part-based methods are proposed to extract facial regions and then up-sample them. SIFT flow [26] and facial landmarks [28] are introduced to locate facial components of input LR images.

Deep learning is an enabling technique for large datasets, and has been applied to face hallucination successfully. Huang *et al.* [10] introduce wavelet coefficients prediction into deep convolutional networks to super-resolve LR inputs with multiple upscaling factors. Zhu *et al.* [41] develop Cascade bi-network to hallucinate low-frequency and high-frequency parts of input LR faces, respectively. Several recent methods explore facial prior knowledge, such as facial attributes [31], parsing maps [5] and component heatmaps [30], for advanced hallucination results.

However, existing approaches mostly focus on hallucinating tiny face images with normal illumination.




Figure 2: Face hallucination and illumination normalization results of state-of-the-art methods and our proposed CPGAN. (a) Input NI-LR image ( $16 \times 16$  pixels); (b) Guided UI-HR image ( $128 \times 128$  pixels). (c) UI-HR image ( $128 \times 128$  pixels, **not available in training**). (d) Result of a popular face hallucination method, TDAE [33]; (e) Illumination normalization result on (a) by applying [40] after bicubic upsampling; (f) Face hallucination result on (e) by [34]; (g) Face hallucination and illumination normalization result on (a) by [34] and [23]; (h) Result of CPGAN ( $128 \times 128$  pixels). **Above all, our CPGAN achieves a photo-realistic visual effect when producing authentic UI-HR face images.**

Thus, in case of non-uniform illuminations, they usually generate serious blurred outputs.

## 2.2. Illumination Compensation

Face illumination compensation methods are proposed to compensate for the non-uniform illumination of human faces and reconstruct face images in a normal illumination condition.

Recent data driven approaches for illumination compensation are based on the illumination cone [2] or the Lambertian Reflectance theory [1]. These approaches learn the disentangled representations of facial appearance and mimic various illumination conditions based on the modeled illumination parameters. For instance, Zhou *et al.* [39] propose a lighting regression network to simulate various lighting scenes for face images. Shu *et al.* [24] propose a GAN framework to decompose face images into physical intrinsic components, geometry, albedo, and illumination base. Zhu *et al.* [40] propose a cycle consistent network to render a content image to an image with different styles. In this way, the illumination condition of the style image can be transferred to the content image.

However, these methods only compensate for non-uniform illumination without well retaining the accurate facial details, especially when the input face images are impaired or low-resolution. Due to the above limitations, simply cascading face hallucination and illumination compensation methods is incompetent to attain high-quality UI-HR faces.

## 3. Hallucination with “Copy” and “Paste”

To reduce the ambiguity of the mapping from NI-LR to UI-HR caused by non-uniform illumination, we present a CPGAN framework that takes a NI-LR face as the input and an external HR face with normal illumination as a guidance to hallucinate a UI-HR one. In CPGAN, we develop Copy and Paste net (CPnet) to flexibly

“copy” and “paste” the uniform illumination features according to the semantic spatial distribution of the input one, thus compensating for illumination of the input image. A discriminator is adopted to force the generated UI-HR face to lie on the manifold of real face images. The whole pipeline is shown in Fig. 3.

### 3.1. Overview of CPGAN

CPGAN is composed of the following components: internal CPnet, external CPnets, spatial transformer networks (STNs) [13], deconvolutional layers, stacked hourglass module [21] and discriminator network. Unlike previous works [5, 30] which only take the LR images as inputs and then super-resolve them with the facial prior knowledge, we incorporate not only input facial information but also an external guided UI-HR face for hallucination. An encoder module is adopted to extract the features of the guided UI-HR image. Note that, our guided face is different from the GT of the NI-LR input.

As shown in Fig. 3, the input NI-LR image is first passed through the internal CPnet to enhance facial details and normalize illumination coarsely by exploiting the shaded facial information. Then the external CPnet resorts to an external guided UI-HR face for further illumination compensation during the upsampling process. Because input images may undergo misalignment, such as in-plane rotations, translations and scale changes, we employ STNs to compensate for misalignment [37, 36], as shown in the yellow blocks in Fig. 3. Meanwhile, inspired by [3], we adopt the stacked hourglass network [21] to estimate vital facial landmark heatmaps for preserving face structure.

#### 3.1.1 Internal CPnet

Due to the shading artifacts, the facial details (high-frequency features) in the input NI-LR face image be-

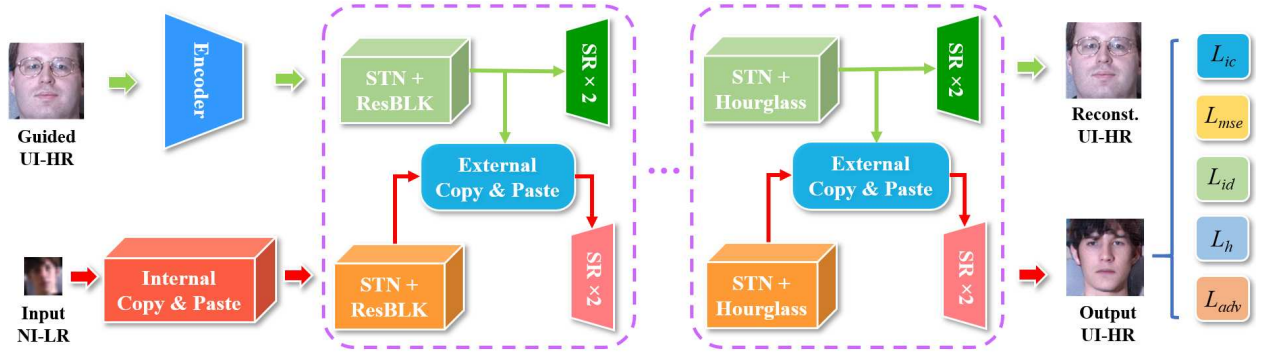


Figure 3: The pipeline of the proposed CPGAN framework. The upper and bottom symmetrical layers in the purple blocks share the same weights.

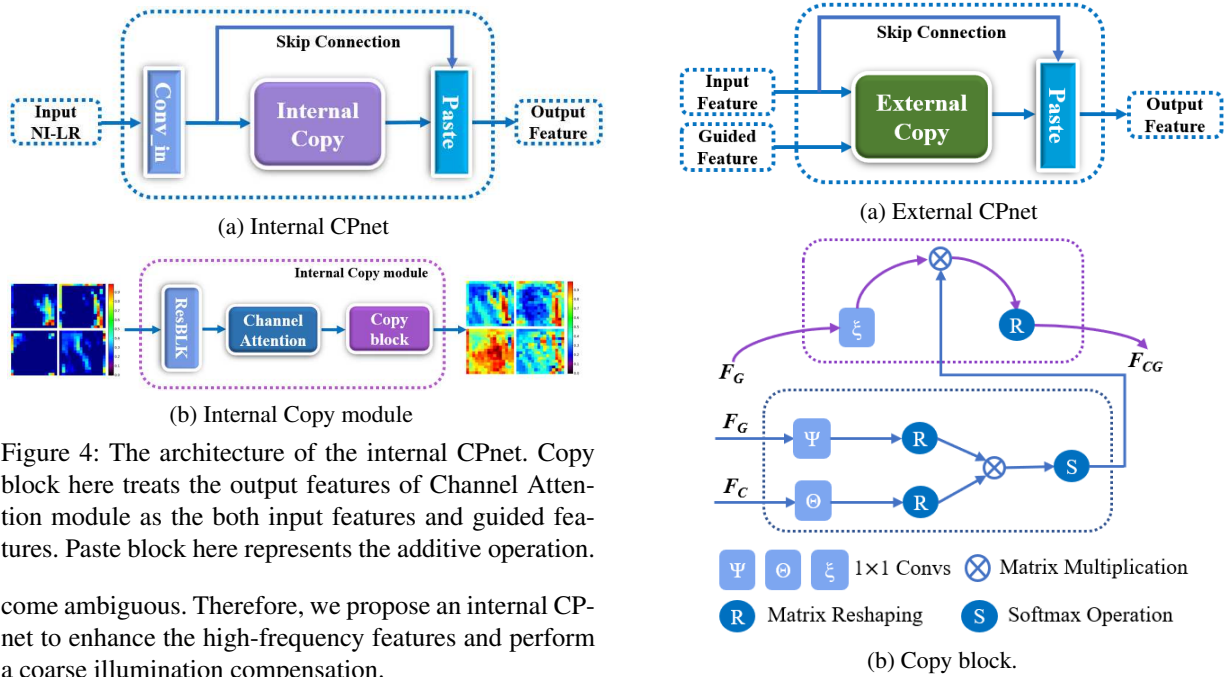


Figure 4: The architecture of the internal CPnet. Copy block here treats the output features of Channel Attention module as the both input features and guided features. Paste block here represents the additive operation.

come ambiguous. Therefore, we propose an internal CPnet to enhance the high-frequency features and perform a coarse illumination compensation.

Fig. 4(b) shows the architecture of our internal CPnet, which consists of an input convolution layer, an Internal Copy module, a Paste block as well as a skip connection. Our Internal Copy module adopts the residual block and Channel-Attention (CA) module in [38] to enhance high-frequency features firstly. Then, our Copy block (Fig. 5(b)) is introduced to “copy” the desired internal uniform illumination features for coarsely compensation. Note that, Copy block here treats the output features of CA module as the both input features ( $F_C$ ) and guided features ( $F_G$ ) in Fig. 5(b). Meanwhile, the skip connection in internal CPnet bypasses the LR input features to the Paste block. In this way, the input NI-LR face is initially refined by the internal CPnet.

To analyze the role of our proposed Internal Copy module, we can exploit the changes of the input and out-

Figure 5: The architecture of the external CPnet. External Copy module here is composed of one Copy block. Paste block represents the additive operation.

put feature maps. The input feature maps estimate the frequency band of the input NI-LR face, which consists of low-frequency facial components. In this way, they are mainly distributed over the low-frequency band (in blue color). After our Internal Copy module, the output features spread in the direction of high-frequency band (in red color), and literally spans the whole band.

Thus, we use the name “Internal Copy module” because its functionality resembles an operation that “copies” the high-frequency features to the low-frequency parts. Above all, the internal CPnet achieves effective feature enhancement, which benefits subse-

quent facial detail upsampling and illumination compensation processes.

### 3.1.2 External CPnet

CPGAN adopts multiple external CPnets and deconvolutional layers to offset the non-uniform illumination and upsample facial details alternately, in a coarse-to-fine fashion. This distinctive design alleviates the ambiguity of correspondences between NI-LR inputs and external UI-HR ones. The network of the external CPnet is shown in Fig. 5(a), and its core components are the Copy and Paste blocks.

Fig. 5(b) performs the “copy” procedure of the Copy block. The guided features  $F_G$  and input features  $F_C$  are extracted from the external guided UI-HR image and the input NI-LR image, respectively. First, the guided features  $F_G$  and input features  $F_C$  are normalized and transformed into two feature space  $\theta$  and  $\psi$  to calculate their similarity. Then, the “copied” features  $F_{CG}$  can be formulated as a weighted sum of the guided features  $F_G$  that are similar to the corresponding positions on the input features  $F_C$ . For the  $i$  th output response:

$$F_{CG}^i = \frac{1}{M(F)} \sum_{\forall j} \left\{ \exp \left( W_{\theta}^T (\bar{F}_C^i)^T \bar{F}_G^j W_{\psi} \right) F_G^j W_{\zeta} \right\} \quad (1)$$

where  $M(F) = \sum_{\forall j} \exp \left( W_{\theta}^T (\bar{F}_C^i)^T \bar{F}_G^j W_{\psi} \right)$  is the sum of all output responses over all positions.  $\bar{F}$  is a transform on  $F$  based on the mean-variance channel-wise normalization. Here, the embedding transformations  $W_{\theta}$ ,  $W_{\psi}$  and  $W_{\zeta}$  are learnt during the training process.

As a result, the Copy block can flexibly integrate the illumination pattern of the guided features into the input features. Based on the Copy and Paste blocks, our proposed external CPnet learns the illumination pattern from the external UI-HR face explicitly.

### 3.2. Loss Function

To train our CPGAN framework, we propose an illumination compensation loss ( $L_{ic}$ ) together with an intensity similarity loss ( $L_{mse}$ ), an identity similarity loss ( $L_{id}$ ) [22], a structure similarity loss ( $L_h$ ) [3] and an adversarial loss ( $L_{adv}$ ) [8]. We will detail the illumination loss shortly. For the rest, please refer to the supplementary material.

The overall loss function  $L_G$  is a weighted summation of the above terms.

$$L_G = \alpha L_{mse} + \beta L_{id} + \gamma L_h + \chi L_{ic} + \kappa L_{adv} \quad (2)$$

**Illumination Compensation Loss:** CPGAN not only recovers UI-HR face images but also compensates for the non-uniform illumination. Inspired by the style loss

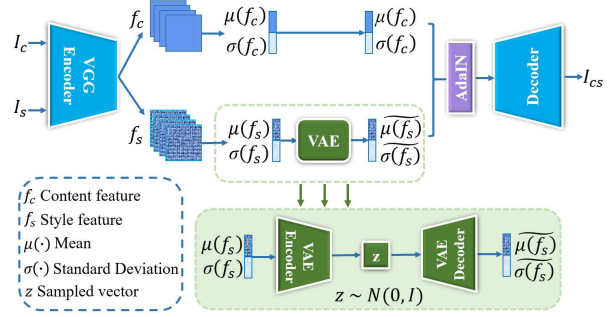


Figure 6: The training process of RaIN model.

in AdaIN [11], we propose the illumination compensation loss  $L_{ic}$ . The basic idea is to constrain the illumination characteristics of the reconstructed UI-HR face is close to the guided UI-HR one in the latent subspace.

$$L_{ic} = \mathbb{E}_{(\hat{h}_i, g_i) \sim p(\hat{h}, g)} \left\{ \sum_{j=1}^L \left\| \mu(\varphi_j(\hat{h}_i)) - \mu(\varphi_j(g_i)) \right\|_2 + \sum_{j=1}^L \left\| \sigma(\varphi_j(\hat{h}_i)) - \sigma(\varphi_j(g_i)) \right\|_2 \right\} \quad (3)$$

where  $g_i$  represents the guided UI-HR image,  $\hat{h}_i$  represents the generated UI-HR image,  $p(\hat{h}, g)$  represents their joint distribution. Each  $\varphi_j(\cdot)$  denotes the output of relu1-1, relu2-1, relu3-1, relu4-1 layer in a pre-trained VGG-19 model [25], respectively. Here,  $\mu$  and  $\sigma$  are the mean and variance for each feature channel.

## 4. Data augmentation

Training a deep neural network requires a lot of samples to prevent overfitting. Since none or limited NI/UI face pairs are available in public face datasets [9, 18], we propose a tailor-made Random Adaptive Instance Normalization (RaIN) model to achieve arbitrary illumination style transfer in real-time, generating sufficient samples for data augmentation (Fig. 6).

RaIN adopts the encoder-decoder architecture, in which the encoder is fixed to the first few layers (up to relu4-1) of a pre-trained VGG-19 [25]. The Adaptive Instance Normalization (AdaIN) [11] layer is embedded to align the feature statistics of the UI face image with those of the NI face image. Specially, we embed Variational Auto-Encoder (VAE) [14] before the AdaIN layer. In this way, we can efficiently produce an unlimited plausible hypotheses for the feature statistics of the NI face image (limited NI face images are provided in public datasets). As a result, sufficient face samples with arbitrary illumination conditions are generated.

Fig. 6 shows the training process of RaIN model. First, given an input content image  $I_c$  (UI face) and a






Figure 7: The generating stage of the RaIN model. A random noise in normal distribution can generate a new style sample.

style image  $I_s$  (NI face), the VAE in RaIN encodes the distributions of feature statistics (mean ( $\mu$ ) and variance ( $\sigma$ )) of the encoded style features  $f_s$ . In this way, a low-dimensional latent space encodes all possible variants for style feature statistics. Then, based on the intermediate AdaIN layer, the feature statistics ( $\mu$  and  $\sigma$ ) of the  $f_c$  are aligned with a randomly chosen style feature statistics, forming the transferred feature  $t$  via

$$t = \text{AdaIN}(I_c, I_s) = \widetilde{\sigma}(f_s) \left( \frac{f_c - \mu(f_c)}{\sigma(f_c)} \right) + \widetilde{\mu}(f_s) \quad (4)$$

where  $\mu$  and  $\sigma$  are computed across each spatial dimension independently for each channel and each sample.

Then, a randomly initialized decoder  $g$  is trained to map  $t$  back to the image space, generating the stylized image  $I_{cs}$ :

$$I_{cs} = g[t] \quad (5)$$

At training stage, RaIN is trained based on the setting of [11, 20, 19], followed by a fine-tuning procedure on the face dataset to encode the photo-realistic facial details. To generate the stylized image with a different illumination condition, we can just feed the content image along with a random noise, as shown in Fig. 7.

## 5. Experiments

In this section, we provide both qualitative and quantitative evaluations on the proposed framework. We conduct the comparisons in the following three scenarios:

- FH: Face hallucination methods (SRGAN [16], TDAE [33], FHC [30]);
- IN+FH: Illumination compensation technique (CycleGAN [40]) + Face hallucination methods (SRGAN [16], TDAE [33], FHC [30]) (use bicubic interpolation procedure to adjust input size);
- FH+IN: Face hallucination methods (SRGAN [16], TDAE [33], FHC [30]) + Illumination compensation technique (CycleGAN [40]).

For a fair comparison, we retrain all these methods using our training datasets.

### 5.1. Datasets

CPGAN is trained and tested on the Multi-PIE dataset [9] (indoor) and the CelebFaces Attributes dataset (CelebA) [18] (in the wild).

The Multi-PIE dataset [9] is a large face dataset with 750K+ images for 337 subjects under various poses, illumination and expression conditions. We choose the NI/UI face pairs of the 249 identities under 10 illumination conditions in Session 1 among Multi-PIE dataset, i.e.,  $249 \times 10 = 2.49\text{K}$  face pairs in total. To enrich the limited illumination conditions, RaIN is adopted to perform data augmentation 10 times on the training set.

Note that CelebA [18] dataset only provides faces in the wild without NI/UI face pairs. For the training purpose, we opt to generate synthesized NI faces from the UI face images. Similar to [7], Adobe Photoshop Lightroom is adopted to render various illumination conditions. We randomly select 18K NI faces from CelebA dataset to perform rendering. Then, 18K NI/UI face pairs are generated.

For the GT images, we crop the aligned UI-HR faces and then resize them to  $128 \times 128$  pixels. For the NI-LR faces, we resize the UI-HR face images to  $128 \times 128$  pixels and then apply a sequence of manual transformations, including rotations, translations, scaling and downsampling processes to obtain images of  $16 \times 16$  pixels. We choose 80 percent of the face pairs for training and 20 percent of the face pairs for testing. Specially, during the training process of CPGAN, we randomly select UI-HR images in the training set to serve as the external guided UI-HR images.

### 5.2. Qualitative Comparison with the SoA

A qualitative comparison with the benchmark methods is presented in Fig.8, which justifies the superior performance of CPGAN over the competing methods. Obviously, the image obtained by CPGAN is more authentic, identity-preserving, and having richer facial details.

As illustrated in Fig.8(b), the combination of bicubic interpolation and CycleGAN cannot produce authentic face images. Due to the incapability of bicubic interpolation to generate necessary high-frequency facial details and the lack of GT, CycleGAN achieves image-to-image translation with deformed face shapes.

SRGAN [16] provides an upscaling factor of  $8\times$ , but it is only trained on general patches. Thus, we retrain SRGAN on face images as well. As shown in Fig.8(c), the results of SRGAN are still blurred.

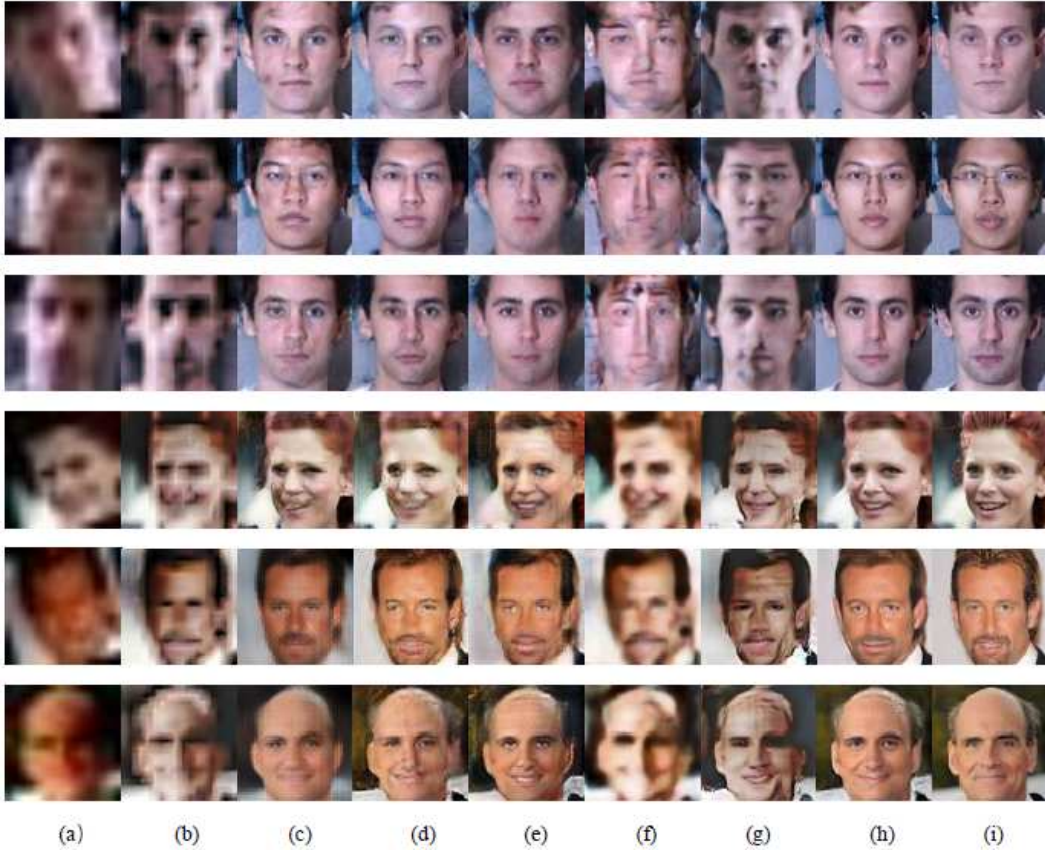


Figure 8: Comparison with state-of-the-art methods. Columns: (a) Unaligned NI-LR inputs. (b) Bicubic interpolation + CycleGAN [40]. (c) SRGAN [16]. (d) TDAE [33]. (e) FHC [30]. (f) CycleGAN [40] + SRGAN [16]. (g) TDAE [33] + CycleGAN [40]. (h) Ours. (i) GT. The first three columns: testing samples from **Multi-PIE** dataset (indoor). The last three columns: testing samples from **CelebA** dataset (in-the-wild).

Table 1: Average PSNR [dB] and SSIM values of compared methods on the testing datasets.

Method	Multi-PIE		CelebA		Multi-PIE		CelebA		Multi-PIE		CelebA	
	FH		PSNR	SSIM	IN+FH		PSNR	SSIM	FH+IN		PSNR	SSIM
	PSNR	SSIM			PSNR	SSIM			PSNR	SSIM		
Bicubic	12.838	0.385	12.79	0.480	13.960	0.386	13.018	0.494	13.315	0.399	12.945	0.352
SRGAN	16.769	0.396	17.951	0.536	14.252	0.408	16.506	0.488	15.044	0.432	15.315	0.398
TDAE	19.342	0.411	19.854	0.530	15.449	0.445	18.031	0.540	15.319	0.427	15.727	0.403
FHC	20.680	0.467	21.130	0.552	17.554	0.512	19.508	0.499	18.263	0.545	16.527	0.426
CPGAN	<b>24.639</b>	<b>0.778</b>	<b>23.972</b>	<b>0.723</b>	<b>24.639</b>	<b>0.778</b>	<b>23.972</b>	<b>0.723</b>	<b>24.639</b>	<b>0.778</b>	<b>23.972</b>	<b>0.723</b>

TDAE [33] super-resolves very low resolution and unaligned face images. It employs deconvolutional layers to upsample LR faces and a discriminative network to promote the generation of sharper results. However, it does not take illumination into account, thus, the final outputs suffer from severe artifacts (Fig.8(d)).

Yu and Porikli [30] exploit the FHC to hallucinate unaligned LR face images. As visualized in Fig.8(e), FHC fails to construct realistic facial details due to inaccurate facial prior prediction caused by shading artifacts.

We provide results on the IN+FH and FH+IN com-

bination as well for completeness. However, due to the deteriorated super-resolved facial patterns caused by the misleading mappings, their generated results are contaminated with ghosting artifacts (see Figs.8(f) and (g)).

In contrast, our method is able to reconstruct authentic facial details as shown in Fig.8(h). Albeit there exists non-uniform illumination in the input NI-LR faces, our method still produces visually pleasing UI-HR faces. For instance, we hallucinate the shading artifacts covered facial parts preciously, such as the jaw and mouth, as illustrated in the third and fourth rows in Fig.8(h).



Figure 9: Ablation Analysis. (a)  $16 \times 16$  NI-LR image. (b)  $128 \times 128$  GT image. (c) Result without using the internal CPnet, using a simple input convolution layer instead. (d) Result without using the external CPnets, connecting input and output features directly. (e) Result without using the identity similarity loss ( $L_{id}$ ). (f) Result without using structure similarity loss ( $L_h$ ). (g) Result without using adversarial loss ( $L_{adv}$ ). (h) CPGAN result.



NI-LR      guided      [33]      ours  
Figure 10: Results on real NI-LR face images.

### 5.3. Quantitative Comparison with the SoA

The above qualitative performances are verified by the quantitative evaluations. From Table 1, the proposed CPGAN performs unarguably better than the rest in both indoor and in the wild datasets. For instance, in FH scenario, CPGAN outperforms the second best technique with a large margin of approximate 4 dB in PSNR on Multi-PIE dataset. Thanks to the deliberate design of the internal and external CPnets, the realistic facial details are well recovered from the shading artifacts.

### 5.4. Performance on Real NI-LR faces

Our method can also effectively hallucinate the real NI-LR faces beyond the synthetically generated face pairs. To demonstrate this, we randomly choose face images with non-uniform illumination from the CelebA dataset. As shown in Fig.10, our CPGAN can hallucinate such randomly chosen shaded thumbnails, demonstrating its robustness in various circumstances.

### 5.5. Ablation Analysis

**Effectiveness of internal CPnet:** As shown in Fig.9(c), it can be seen that the result without internal CPnet suffers from severe distortion and blur artifacts. This is because the distinctive designed internal CPnet enhances the facial details in the input NI-LR image, and it aids subsequent upsampling and illumination compensation. A quantitative study is reported in Table 2.

**Effectiveness of external CPnet:** In our method, ex-

Table 2: Ablation study of CPnet

w/o CPnet	Multi-PIE		CelebA	
	PSNR	SSIM	PSNR	SSIM
internal CPnet	22.164	0.693	21.441	0.578
external CPnet	21.925	0.680	21.032	0.536
CPGAN	<b>24.639</b>	<b>0.778</b>	<b>23.972</b>	<b>0.723</b>

ternal CPnet introduces a guided UI-HR image for illumination compensation. We demonstrate the effectiveness in Fig.9(d) and Table 2. Without external CPnet, the reconstructed face deviates from the GT appearance seriously, especially for the input part affected by shading artifacts. It implies that external CPnet learns the illumination pattern from the external UI-HR face explicitly.

**Loss Function Specifications:** Fig.9 also illustrates the perceptual performance of different training loss variants. The effect of our illumination compensation loss ( $L_{ic}$ ) is provided in the supplementary material.

## 6. Conclusion

This paper presents our CPGAN framework to jointly hallucinate the NI-LR face images and compensate for the non-uniform illumination seamlessly. With the internal and external CPnets, our method enables coarse-to-fine feature refinement based on the semantic spatial distribution of the facial features. In this spirit, we offset shading artifacts and upsample facial details alternately. Meanwhile, the RaIN model mimics sufficient face pairs under diverse illumination conditions to facilitate practical face hallucination. Experimental results validate the effectiveness of CPGAN, which yields photo-realistic visual quality and promising quantitative performance compared with the state-of-the-art.

**Acknowledgment.** This work is supported by the National Natural Science Foundation of China (No.61871123 & No.61802203), Key Research and Development Program in Jiangsu Province (No.BE2016739) and NSF (No.BK20180761). This work is partially supported by ARC under Grant (DP180100106 & DP200101328) and ARC of Excellence for Robotic Vision (CE140100016).



## References

- [1] Ronen Basri and David W Jacobs. Lambertian reflectance and linear subspaces. *IEEE transactions on pattern analysis and machine intelligence*, 25(2):218–233, 2003.
- [2] Peter N Belhumeur and David J Kriegman. What is the set of images of an object under all possible illumination conditions? *International Journal of Computer Vision*, 28(3):245–260, 1998.
- [3] Adrian Bulat and Georgios Tzimiropoulos. Superfan: Integrated facial landmark localization and super-resolution of real-world low resolution faces in arbitrary poses with gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 109–117, 2018.
- [4] Qingxing Cao, Liang Lin, Yukai Shi, Xiaodan Liang, and Guanbin Li. Attention-aware face hallucination via deep reinforcement learning. In *CVPR*, pages 690–698, 2017.
- [5] Yu Chen, Ying Tai, Xiaoming Liu, Chunhua Shen, and Jian Yang. Fsrnet: End-to-end learning face super-resolution with facial priors. In *CVPR*, pages 2492–2501, 2018.
- [6] Miguel Heredia Conde, Davoud Shahlaci, Volker Blanz, and Otmar Loffeld. Efficient and robust inverse lighting of a single face image using compressive sensing. In *ICCV Workshops*, pages 226–234, 2015.
- [7] Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Style aggregated network for facial landmark detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 379–388, 2018.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [9] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker. Multi-pie. *Image and Vision Computing*, 28(5):807–813, 2010.
- [10] Huaibo Huang, Ran He, Zhenan Sun, and Tieniu Tan. Wavelet-srnet: A wavelet-based cnn for multi-scale face super resolution. In *ICCV*, pages 1689–1697, 2017.
- [11] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017.
- [12] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, pages 1125–1134, 2017.
- [13] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *NIPS*, pages 2017–2025, 2015.
- [14] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [15] Soheil Kolouri and Gustavo K Rohde. Transport-based single frame super resolution of very low resolution face images. In *CVPR*, pages 4876–4884, 2015.
- [16] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, pages 4681–4690, 2017.
- [17] Ce Liu, Heung-Yeung Shum, and William T Freeman. Face hallucination: Theory and practice. *IJCV*, 75(1):115–134, 2007.
- [18] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [19] Yawei Luo, Ping Liu, Tao Guan, Junqing Yu, and Yi Yang. Significance-aware information bottleneck for domain adaptive semantic segmentation. In *ICCV*, 2019.
- [20] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *CVPR*, 2019.
- [21] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, pages 483–499, 2016.
- [22] Fatemeh Shiri, Xin Yu, Fatih Porikli, Richard Hartley, and Piotr Koniusz. Identity-preserving face recovery from stylized portraits. *International Journal of Computer Vision*, 127(6-7):863–883, 2019.
- [23] Zhixin Shu, Sunil Hadap, Eli Shechtman, Kalyan Sunkavalli, Sylvain Paris, and Dimitris Samaras. Portrait lighting transfer using a mass transport approach. *ACM Transactions on Graphics (TOG)*, 37(1):2, 2018.
- [24] Zhixin Shu, Ersin Yumer, Sunil Hadap, Kalyan Sunkavalli, Eli Shechtman, and Dimitris Samaras. Neural face editing with intrinsic image disentangling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5541–5550, 2017.
- [25] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [26] Marshall F Tappen and Ce Liu. A bayesian approach to alignment-based image hallucination. In *ECCV*, pages 236–249, 2012.
- [27] Xiaogang Wang and Xiaoou Tang. Hallucinating face by eigentransformation. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 35(3):425–434, 2005.
- [28] Chih-Yuan Yang, Sifei Liu, and Ming-Hsuan Yang. Hallucinating compressed face images. *IJCV*, 126(6):597–614, 2018.
- [29] Fei Yang, Jue Wang, Eli Shechtman, Lubomir Bourdev, and Dimitri Metaxas. Expression flow for 3d-aware face component transfer. *ACM transactions on graphics (TOG)*, 30(4):60, 2011.

- [30] Xin Yu, Basura Fernando, Bernard Ghanem, Fatih Porikli, and Richard Hartley. Face super-resolution guided by facial component heatmaps. In *ECCV*, pages 217–233, 2018.
- [31] Xin Yu, Basura Fernando, Richard Hartley, and Fatih Porikli. Super-resolving very low-resolution face images with supplementary attributes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 908–917, 2018.
- [32] Xin Yu and Fatih Porikli. Ultra-resolving face images by discriminative generative networks. In *ECCV*, pages 318–333, 2016.
- [33] Xin Yu and Fatih Porikli. Face hallucination with tiny unaligned images by transformative discriminative neural networks. In *AAAI*, 2017.
- [34] Xin Yu and Fatih Porikli. Hallucinating very low-resolution unaligned and noisy face images by transformative discriminative autoencoders. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3760–3768, 2017.
- [35] Xin Yu and Fatih Porikli. Imagining the unimaginable faces by deconvolutional networks. *IEEE Transactions on Image Processing*, 27(6):2747–2761, 2018.
- [36] Xin Yu, Fatih Porikli, Basura Fernando, and Richard Hartley. Hallucinating unaligned face images by multiscale transformative discriminative networks. *International Journal of Computer Vision*, 128(2):500–526, 2020.
- [37] Xin Yu, Fatemeh Shiri, Bernard Ghanem, and Fatih Porikli. Can we see more? joint frontalization and hallucination of unaligned tiny faces. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [38] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 286–301, 2018.
- [39] Hao Zhou, Jin Sun, Yaser Yacoob, and David W Jacobs. Label denoising adversarial network (ldan) for inverse lighting of face images. *arXiv preprint arXiv:1709.01993*, 2017.
- [40] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
- [41] Shizhan Zhu, Sifei Liu, Chen Change Loy, and Xiaoou Tang. Deep cascaded bi-network for face hallucination. In *ECCV*, pages 614–630, 2016.