



Learning Warped Guidance for Blind Face Restoration

Xiaoming Li¹, Ming Liu¹, Yuting Ye¹, Wangmeng Zuo¹,
Liang Lin², and Ruigang Yang³

¹ School of Computer Science and Technology, Harbin Institute of Technology,
Harbin, China

{csxmli,wmzuo}@hit.edu.cn, csmliu@outlook.com, yeyuting.jlu@gmail.com

² School of Data and Computer Science, Sun Yat-sen University, Guangzhou, China
linliang@ieee.org

³ Department of Computer Science, University of Kentucky, Lexington, USA
ryang@cs.uky.edu

Abstract. This paper studies the problem of blind face restoration from an unconstrained blurry, noisy, low-resolution, or compressed image (i.e., degraded observation). For better recovery of fine facial details, we modify the problem setting by taking both the degraded observation and a high-quality guided image of the same identity as input to our guided face restoration network (GFRNet). However, the degraded observation and guided image generally are different in pose, illumination and expression, thereby making plain CNNs (e.g., U-Net) fail to recover fine and identity-aware facial details. To tackle this issue, our GFRNet model includes both a warping subnetwork (WarpNet) and a reconstruction subnetwork (RecNet). The WarpNet is introduced to predict flow field for warping the guided image to correct pose and expression (i.e., warped guidance), while the RecNet takes the degraded observation and warped guidance as input to produce the restoration result. Due to that the ground-truth flow field is unavailable, landmark loss together with total variation regularization are incorporated to guide the learning of WarpNet. Furthermore, to make the model applicable to blind restoration, our GFRNet is trained on the synthetic data with versatile settings on blur kernel, noise level, downsampling scale factor, and JPEG quality factor. Experiments show that our GFRNet not only performs favorably against the state-of-the-art image and face restoration methods, but also generates visually photo-realistic results on real degraded facial images.

Keywords: Face hallucination · Blind image restoration · Flow field

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-01261-8_17) contains supplementary material, which is available to authorized users.

1 Introduction

Face restoration aims to reconstruct high quality face image from degraded observation for better display and further analyses [4, 5, 8, 9, 17, 32, 49, 52–54, 59]. In the ubiquitous imaging era, imaging sensors are embedded into many consumer products and surveillance devices, and more and more images are acquired under unconstrained scenarios. Consequently, low quality face images cannot be completely avoided during acquisition and communication due to the introduction of low-resolution, defocus, noise and compression. On the other hand, high quality face images are sorely needed for human perception, face recognition [42] and other face analysis [1] tasks. All these make face restoration a very challenging yet active research topic in computer vision.

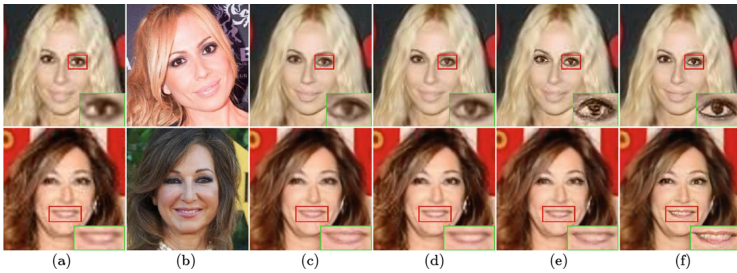


Fig. 1. Restoration results on real low quality images: (a) real low quality image, (b) guided image, and the results by (c) U-Net [43] by taking low quality image as input, (d) U-Net [43] by taking both guided image and low quality image as input, (e) our GFRNet without landmark loss, and (f) our full GFRNet model. Best viewed by zooming in the screen.

Many studies have been carried out to handle specific face restoration tasks, such as denoising [2, 3], hallucination [4, 5, 8, 17, 32, 49, 52–54, 59] and deblurring [9]. Most existing methods, however, are proposed for handling a single specific face restoration task in a non-blind manner. In practical scenario, it is more general that both the degradation types and degradation parameters are unknown in advance. Therefore, more attentions should be given to blind face restoration. Moreover, most previous works produce the restoration results purely relying on a single degraded observation. It is worth noting that the degradation process generally is highly ill-posed. By learning a direct mapping from degraded observation, the restoration result inclines to be over-smoothing and cannot faithfully retain fine and identity-aware facial details.

In this paper, we study the problem of guided blind face restoration by incorporating the degraded observation and a high-quality guided face image. Without loss of generality, the guided image is assumed to have the same identity with the degraded observation, and is frontal with eyes open. We note that such guided restoration setting is practically feasible in many real world applications.

For example, most smartphones support to recognize and group the face images according to their identities¹. In each group, the high quality face image can thus be exploited to guide the restoration of low quality images. In film restoration, it is also encouraging to use the high quality portrait of an actor to guide the restoration of low-resolution and corrupted face images of the same actor from an old film. For these tasks, further incorporation of guided image not only can ease the difficulty of blind restoration, but also is helpful in faithfully recovering fine and identity-aware facial details.

Guided blind face restoration, however, cannot be addressed well by simply taking the degraded observation and guided image as input to plain convolutional networks (CNNs), due to that the two images generally are of different poses, expressions and lighting conditions. Figure 1(c) shows the results obtained using the U-Net [43] by only taking degraded observation as input, while Fig. 1(d) shows the results by taking both two images as input. It can be seen that direct incorporation of guided image brings very limited improvement on the restoration result. To tackle this issue, we develop a guided face restoration network (GFRNet) consisting of a warping subnetwork (WarpNet) and a reconstruction subnetwork (RecNet). Here, the WarpNet is firstly deployed to predict a flow field for warping the guided image to obtain the warped guidance, which is required to have the same pose and expression with degraded observation. Then, the RecNet takes both degraded observation and warped guidance as input to produce the final restoration result. To train GFRNet, we adopt the reconstruction learning to constrain the restoration result to be close to the target image, and further employ the adversarial learning for visually realistic restoration.

Nonetheless, even though the WarpNet can be end-to-end trained with reconstruction and adversarial learning, we empirically find that it cannot converge to the desired solution and fails to align the guided image to the correct pose and expression. Figure 1(e) gives the results of our GFRNet trained by reconstruction and adversarial learning. One can see that its improvement over U-Net is still limited, especially when the degraded observation and guided images are distinctly different in pose. Moreover, the ground-truth flow field is unavailable, and the target and guided images may be of different lighting conditions, making it infeasible to directly use the target image to guide the WarpNet learning. Instead, we adopt the face alignment method [57] to detect the face landmarks of the target and guided images, and then introduce the landmark loss as well as the total variation (TV) regularizer to train the WarpNet. As in Fig. 1(f), our full GFRNet achieves the favorable visual quality, and is effective in recovering fine facial details. Furthermore, to make the learned GFRNet applicable to blind face restoration, our model is trained on the synthetic data generated by a general degradation model with versatile settings on blur kernel, noise level, downsampling scale factor, and JPEG quality factor.

Extensive experiments are conducted to evaluate the proposed GFRNet for guided blind face restoration. The proposed GFRNet achieves significant performance gains over the state-of-the-art restoration methods in quantitative metrics

¹ <https://support.apple.com/HT207103>.

and visually perceptual quality as well as real degraded image. Moreover, our GFRNet also performs favorably on real degraded images as shown in Fig. 1(f). To sum up, the main contribution of this work includes:

- The GFRNet architecture for guided blind face restoration, which includes a warping subnetwork (WarpNet) and a reconstruction subnetwork (RecNet).
- The incorporation of landmark loss and TV regularization for training the WarpNet.
- The promising results of GFRNet on both synthetic and real face images.

2 Related Work

Recent years have witnessed the unprecedented success of deep learning in many image restoration tasks such as super-resolution [11, 24, 28], denoising [46, 55], compression artifact removal [10, 12], compressed sensing [22, 26, 34], and deblurring [7, 27, 36, 37]. As to face images, several CNN architectures have been developed for face hallucination [5, 8, 17, 59], and the adversarial learning is also introduced to enhance the visual quality [52, 53]. Most of these methods, however, are suggested for non-blind restoration and are restricted by the specialized tasks. Benefitted from the powerful modeling capability of deep CNNs, recent studies have shown that it is feasible to train a single model for handling multiple instantiations of degradation (e.g., different noise levels) [35, 55]. As for face hallucination, Yu et al. [53, 54] suggest one kind of transformative discriminative networks to super-resolve different unaligned tiny face images. Nevertheless, blind restoration is a more challenging problem and requires to learn a single model for handling all instantiations of one or more degradation types.

Most studies on deep blind restoration are given to blind deblurring, which aims to recover the latent clean image from noisy and blurry observation with unknown degradation parameters. Early learning-based or CNN-based blind deblurring methods [7, 45, 48] usually follow traditional framework which includes a blur kernel estimation stage and a non-blind deblurring stage. With the rapid progress and powerful modeling capability of CNNs, recent studies incline to bypass blur kernel estimation by directly training a deep model to restore clean image from degraded observation [16, 27, 36–38]. As to blind face restoration, Chrysos and Zafeiriou [9] utilize a modified ResNet architecture to perform face deblurring, while Xu et al. [49] adopt the generative adversarial network (GAN) framework to super-resolve blurry face image. It is worth noting that the success of such kernel-free end-to-end approaches depends on both the modeling capability of CNN and the sufficient sampling on clean images and degradation parameters, making it difficult to design and train. Moreover, the highly ill-posed degradation further increases the difficulty of recovering the correct fine details only from degraded observation [31]. In this work, we elaborately tackle this issue by incorporating a high quality guided image and designing appropriate network architecture and learning objective.

Several learning-based and CNN-based approaches are also developed for color-guided depth image enhancement [15, 18, 29], where the structural interdependency between intensity and depth image is modeled and exploited to reconstruct high quality depth image. For guided depth image enhancement, Hui et al. [18] present a CNN model to learn multi-scale guidance, while Gu et al. [15] incorporate weighted analysis representation and truncated inference for dynamic guidance learning. For general guided filtering, Li et al. [29] construct CNN-based joint filters to transfer structural details from guided image to reconstructed image. However, these approaches assume that the guided image is spatially well aligned with the degraded observation. Due to that the guided image and degraded observation usually are different in pose and expression, such assumption generally does not hold true for guided face restoration. To address this issue, a WarpNet is introduced in our GFRNet to learn a flow field for warping the guided image to the desired pose and expression.

Recently, spatial transformer networks (STNs) are suggested to learn a spatial mapping for warping an image [21], and appearance flow networks (AFNs) are presented to predict a dense flow field to move pixels [13, 58]. Deep dense flow networks have been applied to view synthesis [40, 58], gaze manipulation [13], expression editing [50], and video frame synthesis [33]. In these approaches, the target image is required to have the similar lighting condition with the input image to be warped, and the dense flow networks can thus be trained via reconstruction learning. However, in our guided face restoration task, the guided image and the target image usually are of different lighting conditions, making it less effective to train the flow network via reconstruction learning. Moreover, the ground-truth dense flow field is not available, further increasing the difficulty to train WarpNet. To tackle this issue, we use the face alignment method [57] to extract the face landmarks of guided and target images. Then, the landmark loss and TV regularization are incorporated to facilitate the WarpNet training.

3 Proposed Method

This section presents our GFRNet to recover high quality face image from degraded observation with unknown degradation. Given a degraded observation I^d and a guided image I^g , our GFRNet model produces the restoration result $\hat{I} = \mathcal{F}(I^d, I^g)$ to approximate the ground-truth target image I . Without loss of generality, I^g and I are of the same identity and image size 256×256 . Moreover, to provide richer guidance information, I^g is assumed to be of high quality, frontal, non-occluded with eyes open. Nonetheless, we empirically find that our GFRNet is robust when the assumption is violated. For simplicity, we also assume I^d has the same size with I^g . When such assumption does not hold, e.g., in face hallucination, we simply apply the bicubic scheme to upsample I^d to the size of I^g before inputting it to the GFRNet.

In the following, we first describe the GFRNet model as well as the network architecture. Then, a general degradation model is introduced to generate synthetic training data. Finally, we present the model objective of our GFRNet.

3.1 Guided Face Restoration Network

The degraded observation I^d and guided image I^g usually vary in pose and expression. Directly taking I^d and I^g as input to plain CNNs generally cannot achieve much performance gains over taking only I^d as input (See Fig. 1(c) and (d)). To address this issue, the proposed GFRNet consists of two subnetworks: (i) the warping subnetwork (WarpNet) and (ii) reconstruction subnetwork (RecNet).

Figure 2 illustrates the overall architecture of our GFRNet. The WarpNet takes I^d and I^g as input to predict the flow field for warping guided image,

$$\Phi = \mathcal{F}_w(I^d, I^g; \Theta_w), \quad (1)$$

where Θ_w denotes the WarpNet model parameters. With Φ , the output pixel value of the warped guidance I^w at location (i, j) is given by

$$I_{i,j}^w = \sum_{(h,w) \in \mathcal{N}} I_{h,w}^g \max(0, 1 - |\Phi_{i,j}^y - h|) \max(0, 1 - |\Phi_{i,j}^x - w|), \quad (2)$$

where $\Phi_{i,j}^x$ and $\Phi_{i,j}^y$ denote the predicted x and y coordinates for the pixel $I_{i,j}^w$, respectively. \mathcal{N} stands for the 4-pixel neighbors of $(\Phi_{i,j}^x, \Phi_{i,j}^y)$. From Eq. (2), we note that I^w is subdifferentiable to Φ [21]. Thus, the WarpNet can be end-to-end trained by minimizing the losses defined either on I^w or on Φ .

The predicted warping guidance I^w is expected to have the same pose and expression with the ground-truth I . Thus, the RecNet takes I^d and I^w as input to produce the final restoration result,

$$\hat{I} = \mathcal{F}_r(I^d, I^w; \Theta_r), \quad (3)$$

where Θ_r denotes the RecNet model parameters.




Fig. 2. Overview of our GFRNet. The WarpNet takes the degraded observation I^d and guided image I^g as input to predict the dense flow field Φ , which is adopted to deform I^g to the warped guidance I^w . I^w is expected to be spatially well aligned with I . Thus the RecNet takes I^w and I^d as input to produce the restoration result \hat{I} .




Fig. 3. The network architectures of (a) WarpNet and (b) RecNet.

Warping Subnetwork (WarpNet). The WarpNet adopts the encoder-decoder structure (see Fig. 3(a)) and is comprised of two major components:

- The input encoder extracts feature representation from I^d and I^g , consisting of eight convolution layers and each one with size 4×4 and stride 2.
- The flow decoder predicts the dense flow field for warping I^g to the desired pose and expression, consisting of eight deconvolution layers.

Except the first layer in encoder and the last layer in decoder, all the other layers adopt the convolution-BatchNorm-ReLU form. The detailed structure of WarpNet is given in the supplementary material.

Reconstruction Subnetwork (RecNet). For the RecNet, the input (I^d and I^w) are of the same pose and expression with the output (\hat{I}), and thus the U-Net can be adopted to produce the final restoration result \hat{I} . The RecNet also includes two components, i.e., an encoder and a decoder (see Fig. 3(b)). The encoder and decoder of RecNet are of the same structure with those adopted in WarpNet. To circumvent the information loss, the i -th layer is concatenated to the $(L - i)$ -th layer via skip connections (L is the depth of the U-Net), which has been demonstrated to benefit the rich and fine details of the generated image [20]. The detailed structure of RecNet is given in the supplementary material.

3.2 Degradation Model and Synthetic Training Data

To train our GFRNet, a degradation model is required to generate realistic degraded images. We note that real low quality images can be the results of either defocus, long-distance sensing, noise, compression, or their combinations. Thus, we adopt a general degradation model to generate degraded image $I^{d,s}$,

$$I^{d,s} = ((I \otimes \mathbf{k}_\varrho) \downarrow_s + \mathbf{n}_\sigma)_{JPEG_q}, \tag{4}$$

where \otimes denotes the convolution operator. \mathbf{k}_ϱ stands for the Gaussian blur kernel with the standard deviation ϱ . \downarrow_s denotes the downsampling operator with scale factor s . \mathbf{n}_σ denotes the additive white Gaussian noise (AWGN) with the noise level σ . $(\cdot)_{JPEG_q}$ denotes the JPEG compression operator with quality factor q .

In our general degradation model, $(I \otimes \mathbf{k}_\varrho) \downarrow_s + \mathbf{n}_\sigma$ characterizes the degradation caused by long-distance acquisition, while $(\cdot)_{JPEG_q}$ depicts the degradation caused by JPEG compression. We also note that Xu et al. [49] adopt the degradation model $(I \otimes \mathbf{k}_\varrho + \mathbf{n}_\sigma) \downarrow_s$. However, to better simulate the long-distance image acquisition, it is more appropriate to add the AWGN on the downsampled image. When $s \neq 1$, the degraded image $I^{d,s}$ is of different size with the ground-truth I . So we use bicubic interpolation to upsample $I^{d,s}$ with scale factor s , and then take $I^d = (I^{d,s}) \uparrow_s$ and I^g as input to our GFRNet.

In the following, we explain the parameter settings for these operations:

- **Blur kernel.** In this work, only the isotropic Gaussian blur kernel \mathbf{k}_ϱ is considered to model the defocus effect. We sample the standard deviation of Gaussian blur kernel from the set $\varrho \in \{0, 1 : 0.1 : 3\}$, where 0 indicates no blurring.
- **Downsampler.** We adopt the bicubic downsampler as [5, 8, 17, 49, 59]. The scale factor s is sampled from the set $s \in \{1 : 0.1 : 8\}$.
- **Noise.** As for the noise level σ , we adopt the set $\sigma \in \{0 : 1 : 7\}$ [49].
- **JPEG compression.** For economic storage and communication, JPEG compression with quality factor q is further operated on the degraded image, and we sample q from the set $q \in \{0, 10 : 1 : 40\}$. When $q = 0$, the image is only losslessly compressed.

By including $\varrho = 0$, $s = 1$, $\sigma = 0$ and $q = 0$ in the set of degradation parameters, the general degradation model can simulate the effect of either the defocus, long-distance acquisition, noising, compression or their versatile combinations.

Given a ground-truth image I_i together with the guided image I_i^g , we can first sample ϱ_i , s_i , σ_i and q_i from the parameter set, and then use the degradation model to generate a degraded observation I_i^d . Furthermore, the face alignment method [57] is adopted to extract the landmarks $\{(x_j^{I_i}, y_j^{I_i})\}_{j=1}^{68}$ for I_i and $\{(x_j^{I_i^g}, y_j^{I_i^g})\}_{j=1}^{68}$ for I_i^g . Therefore, we define the synthetic training set as $\mathcal{X} = \{(I_i, I_i^g, I_i^d, \{(x_j^{I_i}, y_j^{I_i})\}_{j=1}^{68}), \{(x_j^{I_i^g}, y_j^{I_i^g})\}_{j=1}^{68})\}_{i=1}^N$, where N denotes the number of samples.

3.3 Model Objective

Losses on Restoration Result \hat{I} . To train our GFDNet, we define the reconstruction loss on the restoration result \hat{I} , and the adversarial loss is further incorporated on \hat{I} to improve the visual perception quality.

Reconstruction Loss. The reconstruct loss is used to constrain the restoration result \hat{I} to be close to the ground-truth I , which includes two terms. The first term is the ℓ_2 loss defined as the squared Euclidean distance between \hat{I} and I , i.e., $\ell_r^0(I, \hat{I}) = \|I - \hat{I}\|^2$. Due to the inherent irreversibility of image restoration, only the ℓ_2 loss inclines to cause over-smoothing result. Following [23], we define the second term as the perceptual loss on the pre-trained VGG-Face [41]. Denote

by ψ the VGG-Face network, $\psi_l(I)$ the feature map of the l -th convolution layer. The perceptual loss on the l -th layer (i.e., Conv-4 in this work) is defined as

$$\ell_p^{\psi,l}(I, \hat{I}) = \frac{1}{C_l H_l W_l} \left\| \psi_l(\hat{I}) - \psi_l(I) \right\|_2^2 \quad (5)$$

where C_l , H_l and W_l denote the channel numbers, height and width of the feature map, respectively. Finally, we define the reconstruction loss as

$$\mathcal{L}_r(I, \hat{I}) = \lambda_{r,0} \ell_r^0(I, \hat{I}) + \lambda_{r,l} \ell_p^{\psi,l}(I, \hat{I}), \quad (6)$$

where $\lambda_{r,0}$ and $\lambda_{r,l}$ are the tradeoff parameters for the ℓ_2 and the perceptual losses, respectively.

Adversarial Loss. Following [19, 30], both global and local adversarial losses are deployed to further improve the perceptual quality of the restoration result. Let $p_{data}(I)$ be the distribution of ground-truth image, $p_d(I^d)$ be the distribution of degraded observation. Using the global adversarial loss [14] as an example, the adversarial loss can be formulated as,

$$\ell_{a,g} = \min_{\Theta} \max_D \mathbb{E}_{I \sim p_{data}(I)} [\log D(I)] + \mathbb{E}_{I^d \sim p_d(I^d)} [\log(1 - D(\mathcal{F}(I^d, I^g; \Theta)))] \quad (7)$$

where $D(I)$ denotes the global discriminator which predicts the possibility that I is from the distribution $p_{data}(I)$. $\mathcal{F}(I^d, I^g; \Theta)$ denotes the restoration result by our GFRNet with the model parameters $\Theta = (\Theta_w, \Theta_r)$.

Following the conditional GAN [20], the discriminator has the same architecture with pix2pix [20], and takes the degraded observation, guided image and restoration result as the input. The network is trained in an adversarial manner, where our GFRNet is updated by minimizing the loss $\ell_{a,g}$ while the discriminator is updated by maximizing $\ell_{a,g}$. To improve the training stability, we adopt the improved GAN [44], and replace the labels 0/1 with the smoothed 0/0.9 to reduce the vulnerability to adversarial examples. The local adversarial loss $\ell_{a,l}$ adopts the same settings with the global one but its discriminator is defined only on the minimal bounding box enclosing all facial landmarks. To sum up, the overall adversarial loss is defined as

$$\mathcal{L}_a = \lambda_{a,g} \ell_{a,g} + \lambda_{a,l} \ell_{a,l}. \quad (8)$$

where $\lambda_{a,g}$ and $\lambda_{a,l}$ are the tradeoff parameters for the global and local adversarial losses, respectively.

Losses on Flow Field Φ . Although the WarpNet can be end-to-end trained based on the reconstruction and adversarial losses, it cannot be learned to correctly align I^w with I in terms of pose and expression (see Fig. 7). In [13, 50], the appearance flow network is trained by minimizing the MSE loss between the output and the ground-truth of the warped image. But for guided face restoration, I generally has different illumination with I^g , and cannot serve as the ground-truth of the warped image. To circumvent this issue, we present the landmark loss as well as the TV regularization to facilitate the learning of WarpNet.

Landmark Loss. Using the face alignment method TCDCN [57], we detect the 68 landmarks $\{(x_j^{I^g}, y_j^{I^g})|_{j=1}^{68}\}$ for I^g and $\{(x_j^I, y_j^I)|_{j=1}^{68}\}$ for I . In order to align I^w and I , it is natural to require that the landmarks of I^w are close to those of I , i.e., $\Phi^x(x_j^I, y_j^I) \approx x_j^{I^g}$ and $\Phi^y(x_j^I, y_j^I) \approx y_j^{I^g}$. Thus, the landmark loss is defined as

$$\ell_{lm} = \sum_i (\Phi_x(x_i^I, y_i^I) - x_i^{I^g})^2 + (\Phi_y(x_i^I, y_i^I) - y_i^{I^g})^2. \quad (9)$$

All the coordinates (including x , y , Φ_x and Φ_y) are normalized to range $[-1, 1]$.

TV Regularization. The landmark loss can only be imposed on the locations of the 68 landmarks. For better learning WarpNet, we further take the TV regularization into account to require that the flow field should be spatially smooth. Given the 2D dense flow field (f_x, f_y) , the TV regularizer is defined as

$$\ell_{TV} = \|\nabla_x \Phi_x\|^2 + \|\nabla_y \Phi_x\|^2 + \|\nabla_x \Phi_y\|^2 + \|\nabla_y \Phi_y\|^2, \quad (10)$$

where ∇_x (∇_y) denotes the gradient operator along the x (y) coordinate.

Combining landmark loss with TV regularizer, we define the flow loss as

$$\mathcal{L}_{flow} = \lambda_{lm} \ell_{lm} + \lambda_{TV} \ell_{TV}, \quad (11)$$

where λ_{lm} and λ_{TV} denote the tradeoff parameters for landmark loss and TV regularizer, respectively.

Overall Objective. Finally, we combine the reconstruction loss, adversarial loss, and flow loss to give the overall objective,

$$\mathcal{L} = \mathcal{L}_r + \mathcal{L}_a + \mathcal{L}_{flow}. \quad (12)$$

4 Experimental Results

Extensive experiments are conducted to assess our GFRNet for guided blind face restoration. Peak Signal-to-Noise Ratio (PSNR) and structural similarity (SSIM) indices are adopted for quantitative evaluation with the related state-of-the-arts (including image super-resolution, deblurring, denoising, compression artifact removal and face hallucination). As for qualitative evaluation, we illustrate the results by our GFRNet and the competing methods. Results on real low quality images are also given to evaluate the generalization ability of our GFRNet. More results can be found in the supplementary material. More results and the code are available at: <https://github.com/csxmli2016/GFRNet>.

4.1 Dataset

We adopt the CASIA-WebFace [51] and VggFace2 [6] datasets to constitute our training and test sets. The WebFace contains 10,575 identities and each has about 46 images with the size 256×256 . The VggFace2 contains 9,131

identities and each has an average of 362 images with different sizes. The images in the two datasets are collected in the wild and cover a large range of pose, age, illumination and expression. For each identity, at most five high quality images are selected, in which a frontal image with eyes open is chosen as the guided image and the others are used as the ground-truth to generate degraded observations. By this way, we build our training set of 20,273 pairs of ground-truth and guided images from the VggFace2 training set. Our test set includes two subsets: (i) 1,005 pairs from the VggFace2 test set, and (ii) 1,455 pairs from WebFace. In addition, 200 pairs from Web-face are chosen as a validation set, which are not included in training and testing. The images whose identities have appeared in our training set are excluded from the test set. Furthermore, low quality images are also excluded in training and testing, which include: (i) low-resolution images, (ii) images with large occlusion, (iii) cartoon images, and (iv) images with obvious artifacts. The face region of each image in VGGFace2 is cropped and resized to 256×256 based on the bounding box detected by MTCNN [56]. All training and test images are not aligned to keep their original pose and expression. Facial landmarks of the ground-truth and guided images are detected by TCDCN [57] and are only used in training.

4.2 Training Details and Parameter Setting

Our model is trained using the Adam algorithm [25] with the learning rate of 2×10^{-4} , 2×10^{-5} , 2×10^{-6} and $\beta_1 = 0.5$. In each learning rate, the model is trained until the reconstruction loss on validation set becomes non-decreasing. Then a smaller learning rate is adopted to further fine-tune the model. The tradeoff parameters are set as $\lambda_{r,0} = 100$, $\lambda_{r,l} = 0.001$, $\lambda_{a,g} = 1$, $\lambda_{a,l} = 0.5$, $\lambda_{lm} = 10$, and $\lambda_{TV} = 1$. We first pre-train the WarpNet for 5 epochs by minimizing the flow loss \mathcal{L}_{flow} , and then both WarpNet and RecNet are end-to-end trained by using the objective \mathcal{L} . The batch size is 1 and the training is stopped after 100 epochs. Data augmentation such as flipping is also adopted during training.

4.3 Results on Synthetic Images

Table 1 lists the PSNR and SSIM results on the two test subsets, where our GFRNet achieves significant performance gains over all the competing methods. Using the $4\times$ SR on WebFace as an example, in terms of PSNR, our GFRNet outperforms other methods by at least 3.5 dB. Since guided blind face restoration remains an uninvestigated issue in literature, we compare our GFRNet with several relevant state-of-the-arts, including three non-blind image super-resolution (SR) methods (SRCNN [11], VDSR [24], SRGAN [28]), three blind deblurring methods (DCP [39], DeepDeblur [36], DeblurGAN [27]), two denoising methods (DnCNN [55], MemNet [46]), one compression artifact removal method (ARCNN [10]), three non-blind face hallucination (FH) methods (CBN [59], WaveletSRNet [17], TDAE [54]), and two blind FH methods (SCGAN [49], MCGAN [49]). To keep consistent with the SR and FH methods, only two scale factors, i.e., 4 and 8, are considered for the test images. As for non-SR methods,

Table 1. Quantitative results on two test subsets. Numbers in the parentheses indicate SSIM and the remaining represents PSNR (dB). The best results are highlighted in red and second best ones except our GFRNet variants are highlighted in blue.

Methods		VggFace2 [6]		WebFace [51]	
		4×	8×	4×	8×
SR	SRCNN [11]	24.57 (.842)	22.30 (.802)	26.11 (.872)	23.50 (.842)
	VDSR [24]	25.36 (.858)	22.50 (.807)	26.60 (.884)	23.65 (.847)
	SRGAN [28]	25.85 (.911)	23.01 (.874)	27.65 (.941)	24.49 (.913)
	MSRGAN	26.55 (.906)	23.45 (.862)	28.10 (.934)	24.92 (.908)
Deblur	DCP [39]	24.42 (.894)	21.54 (.848)	24.97 (.895)	23.05 (.887)
	DeepDeblur [36]	26.31 (.917)	22.97 (.873)	28.13 (.934)	24.63 (.910)
	DeblurGAN [27]	24.65 (.889)	22.06 (.846)	24.63 (.910)	23.38 (.896)
	MDeblurGAN	25.32 (.918)	22.46 (.867)	29.41 (.952)	23.49 (.900)
Denoise	DnCNN [55]	26.73 (.920)	23.29 (.877)	28.35 (.933)	24.75 (.912)
	MemNet [46]	26.85 (.923)	23.31 (.877)	28.57 (.934)	24.77 (.909)
	MDnCNN	27.05 (.925)	23.33 (.879)	29.40 (.942)	24.84 (.912)
AR	ARCNN [10]	22.05 (.863)	20.84 (.827)	23.39 (.876)	20.47 (.858)
	MARCNN	25.43 (.923)	23.16 (.876)	28.40 (.938)	25.15 (.914)
Non-blind FH	CBN [59]	24.52 (.867)	21.84 (.817)	25.43 (.899)	23.10 (.852)
	WaveletSR [17]	25.66 (.909)	20.87 (.831)	27.10 (.937)	21.63 (.869)
	TDAE [54]	- (-)	20.19 (.729)	- (-)	20.24 (.741)
Blind FH	SCGAN [49]	25.16 (.905)	- -	26.37 (.923)	- -
	MCGAN [49]	25.26 (.912)	- -	26.35 (.931)	- -
Ours	Ours(-WG)	25.97 (.915)	22.91 (.838)	28.73 (.928)	24.76 (.884)
	Ours(-WG2)	27.20 (.932)	23.22 (.863)	29.45 (.945)	25.93 (.914)
	Ours(-W)	26.03 (.923)	23.29 (.843)	29.66 (.934)	25.20 (.897)
	Ours(-W2)	27.25 (.933)	23.24 (.864)	29.73 (.948)	25.95 (.917)
	Ours(-F)	26.61 (.927)	23.17 (.863)	31.43 (.920)	26.00 (.922)
	Ours(R)	27.90 (.943)	24.05 (.890)	31.46 (.962)	26.88 (.922)
	Ours(Full)	28.55 (.947)	24.10 (.898)	32.31 (.973)	27.21 (.935)

we take the bicubic upsampling result as the input to the model. To handle 8× SR for SRCNN [11] and VDSR [24], we adopt the strategy in [47] by applying the 2× model to the result produced by the 4× model. For SCGAN [49] and MCGAN [49], only the 4× models are available. For TDAE [54], only the 8× model is available.

Quantitative Evaluation. It is worth noting that the promising performance of our GFRNet cannot be solely attributed to the use of our training data and the simple incorporation of guided image. To illustrate this point, we retrain four competing methods (i.e., SRGAN, DeblurGAN, DnCNN, and ARCNN) by using our training data and taking both degraded observation and guided image as input. For the sake of distinction, the retrained models are represented as MSRGAN, MDeblurGAN, MDnCNN, MARCNN. From Table 1, the retrained models do achieve better PSNR and SSIM results than the original ones, but still perform inferior to our GFRNet with a large margin, especially on WebFace.

Therefore, the performance gains over the retrained models should be explained by the network architecture and model objective of our GFRNet.

Qualitative Evaluation. In Fig. 4, we select three competing methods with top quantitative performance, and compare their results with those by our GFRNet. It is obvious that our GFRNet is more effective in restoring fine details while suppressing visual artifacts. In comparison with the competing methods, the results by GFRNet are visually photo-realistic and can correctly recover more fine and identity-aware details especially in eyes, nose, and mouth regions. More results of all competing methods are included in the supplementary material.




Fig. 4. The $4\times$ SR results: (a) synthetic low quality image (Close-up in right bottom is the guided image), (b) MDnCNN [55], (c) MARCNN [10], (d) MDeblurGAN [27], (e) Ours, and (f) ground-truth.



Fig. 5. Restoration results on real low quality images with different poses: (a) real low quality images (Close-up in right bottom is the guided image), (b) MDnCNN [55], (c) MARCNN [10], (d) MDeblurGAN [27], and (e) Ours.

4.4 Results on Real Low Quality Images

Figure 5 further shows the results on real low quality images by MDnCNN [55], MARCNN [10], MDeblurGAN [27], and our GFRNet. The real images are selected from VGGFace2 with the resolution lower than 60×60 . Even the degradation is unknown, our method yields visually realistic and pleasing results in face region with more fine details, while the competing methods can only achieve moderate improvement on visual quality.

4.5 Ablative Analysis

Two groups of ablative experiments are conducted to assess the components of our GFRNet. First, we consider five variants of our GFRNet: (i) Ours(*Full*): the full GFRNet, (ii) Ours(*-F*): GFRNet by removing the flow loss \mathcal{L}_{flow} , (iii) Ours(*-W*): GFRNet by removing WarpNet (RecNet takes both I^d and I^g as input), (iv) Ours(*-WG*): GFRNet by removing WarpNet and guided image (RecNet only takes I^d as input), and (v) Ours(*R*): GFRNet by using a random I^g with different identity to I^d . Table 1 also lists the PSNR and SSIM results of these variants, and we have the following observations. (i) All the three components, i.e., guided image, WarpNet and flow loss, contribute to the performance improvement. (ii) GFRNet cannot be well trained without the help of flow loss. As a result, although Ours(*-F*) outperforms Ours(*-W*) in most cases, sometimes Ours(*-W*) can perform slightly better than Ours(*-F*) by average PSNR, e.g., for $8 \times$ SR on VggFace2. (iii) It is worth noting that GFRNet with random guidance (i.e., Ours(*R*)) achieves the second best results, indicating that GFRNet is robust to the misuse of identity. Figures 1 and 6 give the restoration results by GFRNet variants. Ours(*Full*) can generate much sharper and richer details, and achieves better perceptual quality than its variants. Moreover, Ours(*R*) also achieves the second best performance in qualitative results, but it may introduce the fine details of the other identity to the result (e.g., eye regions in Fig. 6(h)). Furthermore, to illustrate the effectiveness of flow loss, Fig. 7 shows the warped guidance by Ours(*Full*) and Ours(*-F*). Without the help of flow loss, Ours(*-F*) cannot converge to stable solution and results in unreasonable warped guidance. In contrast, Ours(*Full*) can correctly align guided image to the desired pose and expression, indicating the necessity and effectiveness of flow loss.

In addition, it is noted that the parameters of Ours(*Full*) are nearly two times of Ours(*-W*) and Ours(*-WG*). To show that the gain of Ours(*Full*) does not come from the increase of parameter number, we include two other variants of GFRNet, i.e., Ours(*-W2*) and Ours(*-WG2*), by increasing the channels of Ours(*-W*) and Ours(*-WG*) to 2 times, respectively. From Table 1, in terms of PSNR, Ours(*Full*) also outperforms Ours(*-W2*) and Ours(*-WG2*). Instead of the increase of model parameters, the performance improvement of Ours(*Full*) should be mainly attributed to the incorporation of both WarpNet and flow loss.




Fig. 6. Restoration results of our GFRNet variants: (a) input, (b) guided image, (c) Ours(-WG), (d) Ours(-WG2), (e) Ours(-W), (f) Ours(-W2), (g) Ours(-F), (h) Ours(R) (Close-up in right bottom is the random guided image), (i) Ours(Full), and (j) ground-truth. Best viewed by zooming in the screen.




Fig. 7. Warped guidance by Ours(Full) and Ours(-F): (a) input, (b) guided image, (c) Ours(-F), (d) Ours(Full), and (e) ground-truth.

5 Conclusion

In this paper, we present a guided blind face restoration model, i.e., GFRNet, by taking both the degraded observation and a high-quality guided image from the same identity as input. Besides the reconstruction subnetwork, our GFRNet also includes a warping subnetwork (WarpNet), and incorporates the landmark loss as well as TV regularizer to align the guided image to the desired pose and expression. To make our GFRNet applicable to blind restoration, we further introduce a general image degradation model to synthesize realistic low quality face image. Quantitative and qualitative results show that our GFRNet not only performs favorably against the relevant state-of-the-arts but also generates visually pleasing results on real low quality face images.

Acknowledgments. This work was supported in part by the National Natural Science Foundation of China under grant Nos. 61671182 and 61471146.

References

1. Andreu, Y., López-Centelles, J., Mollineda, R.A., García-Sevilla, P.: Analysis of the effect of image resolution on automatic face gender classification. In: IEEE International Conference Pattern Recognition, pp. 273–278. IEEE (2014)
2. Anwar, S., Huynh, C., Porikli, F.: Combined internal and external category-specific image denoising. In: British Machine Vision Conference (2017)
3. Anwar, S., Porikli, F., Huynh, C.P.: Category-specific object image denoising. IEEE Trans. Image Process. **26**(11), 5506–5518 (2017)
4. Baker, S., Kanade, T.: Hallucinating faces. In: IEEE International Conference on Automatic Face and Gesture Recognition, pp. 83–88. IEEE (2000)
5. Cao, Q., Lin, L., Shi, Y., Liang, X., Li, G.: Attention-aware face hallucination via deep reinforcement learning. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 690–698. IEEE (2017)
6. Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: a dataset for recognising faces across pose and age (2017). arXiv preprint: [arXiv:1710.08092](https://arxiv.org/abs/1710.08092)
7. Chakrabarti, A.: A neural approach to blind motion deblurring. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016, Part III. LNCS, vol. 9907, pp. 221–235. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_14
8. Chen, Y., Tai, Y., Liu, X., Shen, C., Yang, J.: FSRNet: end-to-end learning face super-resolution with facial priors (2017). arXiv preprint: [arXiv:1711.10703](https://arxiv.org/abs/1711.10703)
9. Chrysos, G.G., Zafeiriou, S.: Deep face deblurring. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 2015–2024. IEEE (2017)
10. Dong, C., Deng, Y., Change Loy, C., Tang, X.: Compression artifacts reduction by a deep convolutional network. In: IEEE International Conference on Computer Vision, pp. 576–584. IEEE (2015)
11. Dong, C., Loy, C.C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014, Part IV. LNCS, vol. 8692, pp. 184–199. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10593-2_13
12. Galteri, L., Seidenari, L., Marco, B., Alberto, B.D.: Deep generative adversarial compression artifact removal. In: IEEE International Conference on Computer Vision, pp. 4826–4835. IEEE (2017)
13. Ganin, Y., Kononenko, D., Sungatullina, D., Lempitsky, V.: DeepWarp: photorealistic image resynthesis for gaze manipulation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016, Part II. LNCS, vol. 9906, pp. 311–326. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_20
14. Goodfellow, I.J., et al.: Generative adversarial networks. In: Advances in Neural Information Processing Systems, vol. 3, pp. 2672–2680 (2014)
15. Gu, S., Zuo, W., Guo, S., Chen, Y., Chen, C., Zhang, L.: Learning dynamic guidance for depth image enhancement. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 3769–3778. IEEE (2017)
16. Hradiš, M., Kotera, J., Zemčík, P., Šroubek, F.: Convolutional neural networks for direct text deblurring. In: British Machine Vision Conference (2015)
17. Huang, H., He, R., Sun, Z., Tan, T.: Wavelet-SRNet: a wavelet-based CNN for multi-scale face super resolution. In: IEEE International Conference on Computer Vision, pp. 1689–1697. IEEE (2017)
18. Hui, T.-W., Loy, C.C., Tang, X.: Depth map super-resolution by deep multi-scale guidance. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016, Part III. LNCS, vol. 9907, pp. 353–369. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_22

19. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Globally and locally consistent image completion. *ACM Trans. Graph.* **36**(4), 107:1–107:14 (2017)
20. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1125–1134. IEEE (2016)
21. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: *Advances in Neural Information Processing Systems*, pp. 2017–2025 (2015)
22. Jin, K.H., McCann, M.T., Froustey, E., Unser, M.: Deep convolutional neural network for inverse problems in imaging. *IEEE Trans. Image Process.* **26**(9), 4509–4522 (2017)
23. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016, Part II*. LNCS, vol. 9906, pp. 694–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_43
24. Kim, J., Lee, J.K., Lee, K.M.: Accurate image super-resolution using very deep convolutional networks. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1646–1654. IEEE (2016)
25. Kingma, D., Ba, J.: Adam: a method for stochastic optimization (2014). arXiv preprint: [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
26. Kulkarni, K., Lohit, S., Turaga, P., Kerviche, R., Ashok, A.: Reconnet: non-iterative reconstruction of images from compressively sensed measurements. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 449–458. IEEE (2016)
27. Kupyn, O., Budzan, V., Mykhailych, M., Mishkin, D., Matas, J.: DeblurGAN: blind motion deblurring using conditional adversarial networks (2017). arXiv preprint: [arXiv:1711.07064](https://arxiv.org/abs/1711.07064)
28. Ledig, C., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4681–4690. IEEE (2017)
29. Li, Y., Huang, J.-B., Ahuja, N., Yang, M.-H.: Deep joint image filtering. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016, Part IV*. LNCS, vol. 9908, pp. 154–169. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_10
30. Li, Y., Liu, S., Yang, J., Yang, M.H.: Generative face completion. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3911–3919. IEEE (2017)
31. Lin, Z., He, J., Tang, X., Tang, C.K.: Limits of learning-based superresolution algorithms. *Int. J. Comput. Vis.* **80**(3), 406–420 (2008)
32. Liu, C., Shum, H.Y., Freeman, W.T.: Face hallucination: theory and practice. *Int. J. Comput. Vis.* **75**(1), 115–134 (2007)
33. Liu, Z., Yeh, R.A., Tang, X., Liu, Y., Agarwala, A.: Video frame synthesis using deep voxel flow. In: *IEEE International Conference on Computer Vision*, pp. 4463–4471. IEEE (2017)
34. Lucas, A., Iliadis, M., Molina, R., Katsaggelos, A.K.: Using deep neural networks for inverse problems in imaging: beyond analytical methods. *IEEE Signal Process. Mag.* **35**(1), 20–36 (2018)
35. Mao, X.J., Shen, C., Yang, Y.B.: Image denoising using very deep fully convolutional encoder-decoder networks with symmetric skip connections (2016). arXiv preprint: [arXiv:1603.09056](https://arxiv.org/abs/1603.09056)

36. Nah, S., Kim, T.H., Lee, K.M.: Deep multi-scale convolutional neural network for dynamic scene deblurring. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 3883–3891. IEEE (2017)
37. Nimisha, T., Singh, A.K., Rajagopalan, A.: Blur-invariant deep learning for blind-deblurring. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 4752–4760. IEEE (2017)
38. Noroozi, M., Chandramouli, P., Favaro, P.: Motion deblurring in the wild. In: Roth, V., Vetter, T. (eds.) GCPR 2017. LNCS, vol. 10496, pp. 65–77. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66709-6_6
39. Pan, J., Sun, D., Pfister, H., Yang, M.H.: Blind image deblurring using dark channel prior. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 1628–1636. IEEE (2016)
40. Park, E., Yang, J., Yumer, E., Ceylan, D., Berg, A.C.: Transformation-grounded image generation network for novel 3d view synthesis. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 3500–3509. IEEE (2017)
41. Parkhi, O.M., Vedaldi, A., Zisserman, A.: Deep face recognition. In: British Machine Vision Conference, pp. 41.1–41.12 (2015)
42. Phillips, P.J., et al.: Overview of the face recognition grand challenge. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 947–954. IEEE (2005)
43. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015, Part III. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
44. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. In: Advances in Neural Information Processing Systems, pp. 2234–2242 (2016)
45. Schuler, C.J., Hirsch, M., Harmeling, S., Schölkopf, B.: Learning to deblur. IEEE Trans. Pattern Anal. Mach. Intell. **38**(7), 1439–1451 (2016)
46. Tai, Y., Yang, J., Liu, X., Xu, C.: MemNet: a persistent memory network for image restoration. In: International Conference on Computer Vision, pp. 4549–4557. IEEE (2017)
47. Tuzel, O., Taguchi, Y., Hershey, J.R.: Global-local face upsampling network (2016). arXiv preprint: [arXiv:1603.07235](https://arxiv.org/abs/1603.07235)
48. Xiao, L., Wang, J., Heidrich, W., Hirsch, M.: Learning high-order filters for efficient blind deconvolution of document photographs. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016, Part III. LNCS, vol. 9907, pp. 734–749. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_45
49. Xu, X., Sun, D., Pan, J., Zhang, Y., Pfister, H., Yang, M.H.: Learning to super-resolve blurry face and text images. In: IEEE International Conference on Computer Vision, pp. 251–260. IEEE (2017)
50. Yeh, R., Liu, Z., Goldman, D.B., Agarwala, A.: Semantic facial expression editing using autoencoded flow (2016). arXiv preprint: [arXiv:1611.09961](https://arxiv.org/abs/1611.09961)
51. Yi, D., Lei, Z., Liao, S., Li, S.Z.: Learning face representation from scratch (2014). arXiv preprint: [arXiv:1411.7923](https://arxiv.org/abs/1411.7923)
52. Yu, X., Porikli, F.: Ultra-resolving face images by discriminative generative networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016, Part V. LNCS, vol. 9909, pp. 318–333. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46454-1_20

53. Yu, X., Porikli, F.: Face hallucination with tiny unaligned images by transformative discriminative neural networks. In: AAAI Conference on Artificial Intelligence, pp. 4327–4333 (2017)
54. Yu, X., Porikli, F.: Hallucinating very low-resolution unaligned and noisy face images by transformative discriminative autoencoders. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 3760–3768. IEEE (2017)
55. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a Gaussian denoiser: residual learning of deep CNN for image denoising. *IEEE Trans. Image Process.* **26**(7), 3142–3155 (2017)
56. Zhang, K., Zhang, Z., Li, Z., Qiao, Y.: Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Process. Lett.* **23**(10), 1499–1503 (2016)
57. Zhang, Z., Luo, P., Loy, C.C., Tang, X.: Learning deep representation for face alignment with auxiliary attributes. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(5), 918–930 (2016)
58. Zhou, T., Tulsiani, S., Sun, W., Malik, J., Efros, A.A.: View synthesis by appearance flow. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016, Part IV. LNCS, vol. 9908, pp. 286–301. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46493-0_18
59. Zhu, S., Liu, S., Loy, C.C., Tang, X.: Deep cascaded Bi-network for face hallucination. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016, Part V. LNCS, vol. 9909, pp. 614–630. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46454-1_37