Supplementary Materials to
"GAN Prior Embedded Network for Blind Face Restoration in the Wild"

Tao Yang[1], Peiran Ren[1], Xuansong Xie[1], and Lei Zhang[1,2]
[1]DAMO Academy, Alibaba Group
[2]Department of Computing, The Hong Kong Polytechnic University
yangtao9009@gmail.com, peiran_r@sohu.com, xingtong.xxs@taobao.com, cslzhang@comp.polyu.edu.hk

In this supplementary file, we provide the following materials:

- More visual comparisons of different methods on synthetic face images (referring to Section 4.4);
- More visual comparisons of different methods on real face images in the wild (referring to Section 4.5);
- Some preliminary results on face inpainting and face colorization.

1. Experiments on Synthetic Images

This section shows more visual results of competing methods on the blind face restoration (BFR) and face super-resolution (FSR) tasks. The experimental settings can be found in Section 4.4 in the main paper. As in the main paper, we compare our GPEN method with Pix2PixHD [8], Super-FAN [1], GFRNet [5], GWAInet [2], DFDNet [4] and HiFaceGAN [10] on the task of BFR. As for the task of FSR, we compare with Super-FAN [1], GFRNet [5], GWAInet [2], DFDNet [4], HiFaceGAN [10], mGANprior [3], PULSE [6] and pSp [7]. The visual comparisons on the BFR and FSR tasks are presented in Figure 1 and Figure 2, respectively.

2. Experiments on Images in the Wild

The $1,000$ real-world low quality face images we collected from internet and our BFR results will be made publically available. As in Section 4.5 in the main manuscript, the methods Pix2PixHD [8], Super-FAN [1], GFRNet [5], GWAInet [2], DFDNet [4] and HiFaceGAN [10] are used in the comparison. Figure 3 shows the visual comparisons, demonstrating the superior performance of our method on restoring photo-realistic facial details.

3. Face Inpainting and Face Colorization

Though our method is designed for BFR, it can serve as a generic solution for other image-to-image tasks, such as face inpainting and face colorization, in which GAN prior plays a critical role.

Face Inpainting. Face inpainting aims to recover the missing pixels indicated by a binary mask in a face image. In this experiment, we treat the task as a blind face inpainting problem without using the binary mask. During training, we generate random holes with arbitrary shape in the high-quality face images on-the-fly as inputs. The model is updated following the same strategies and settings as in our main paper.

Figure 4 shows the qualitative comparisons of our method with the state-of-the-art face inpainting methods Deepfill v2 [11] and GMCNN [9], both of which require an extra binary mask to indicate the location of missing pixels. Our model demonstrates much better performance and it reproduces high-quality faces in a resolution of 1024^2.

Face Colorization. Given a grayscale face as input, our model can also hallucinate a plausible color version of it. We update our model by taking a colored face image and its grayscale counterpart as a training pair. The training strategies and settings are inherited from our main paper.

We compare our GPEN with mGANprior [3], which uses the multi-code GAN prior, and the methods in [12, 13], which are specially designed for colorization task. Figure 5 presents the qualitative comparisons. It can be seen that our model can achieve favorably better face colorization results.

(a) (b) (c) (d) (e) (f) (g) (h) (i)

Figure 1: Blind face restoration results on synthetic face images. (a) Degraded faces; (b) Super-FAN [1]; (c) GFRNet [5]; (d) GWAInet [2]; (e) Pix2PixHD [8]; (f) DFDNet [4]; (g) HiFaceGAN [10]; (h) GPEN; (i) Ground truth.

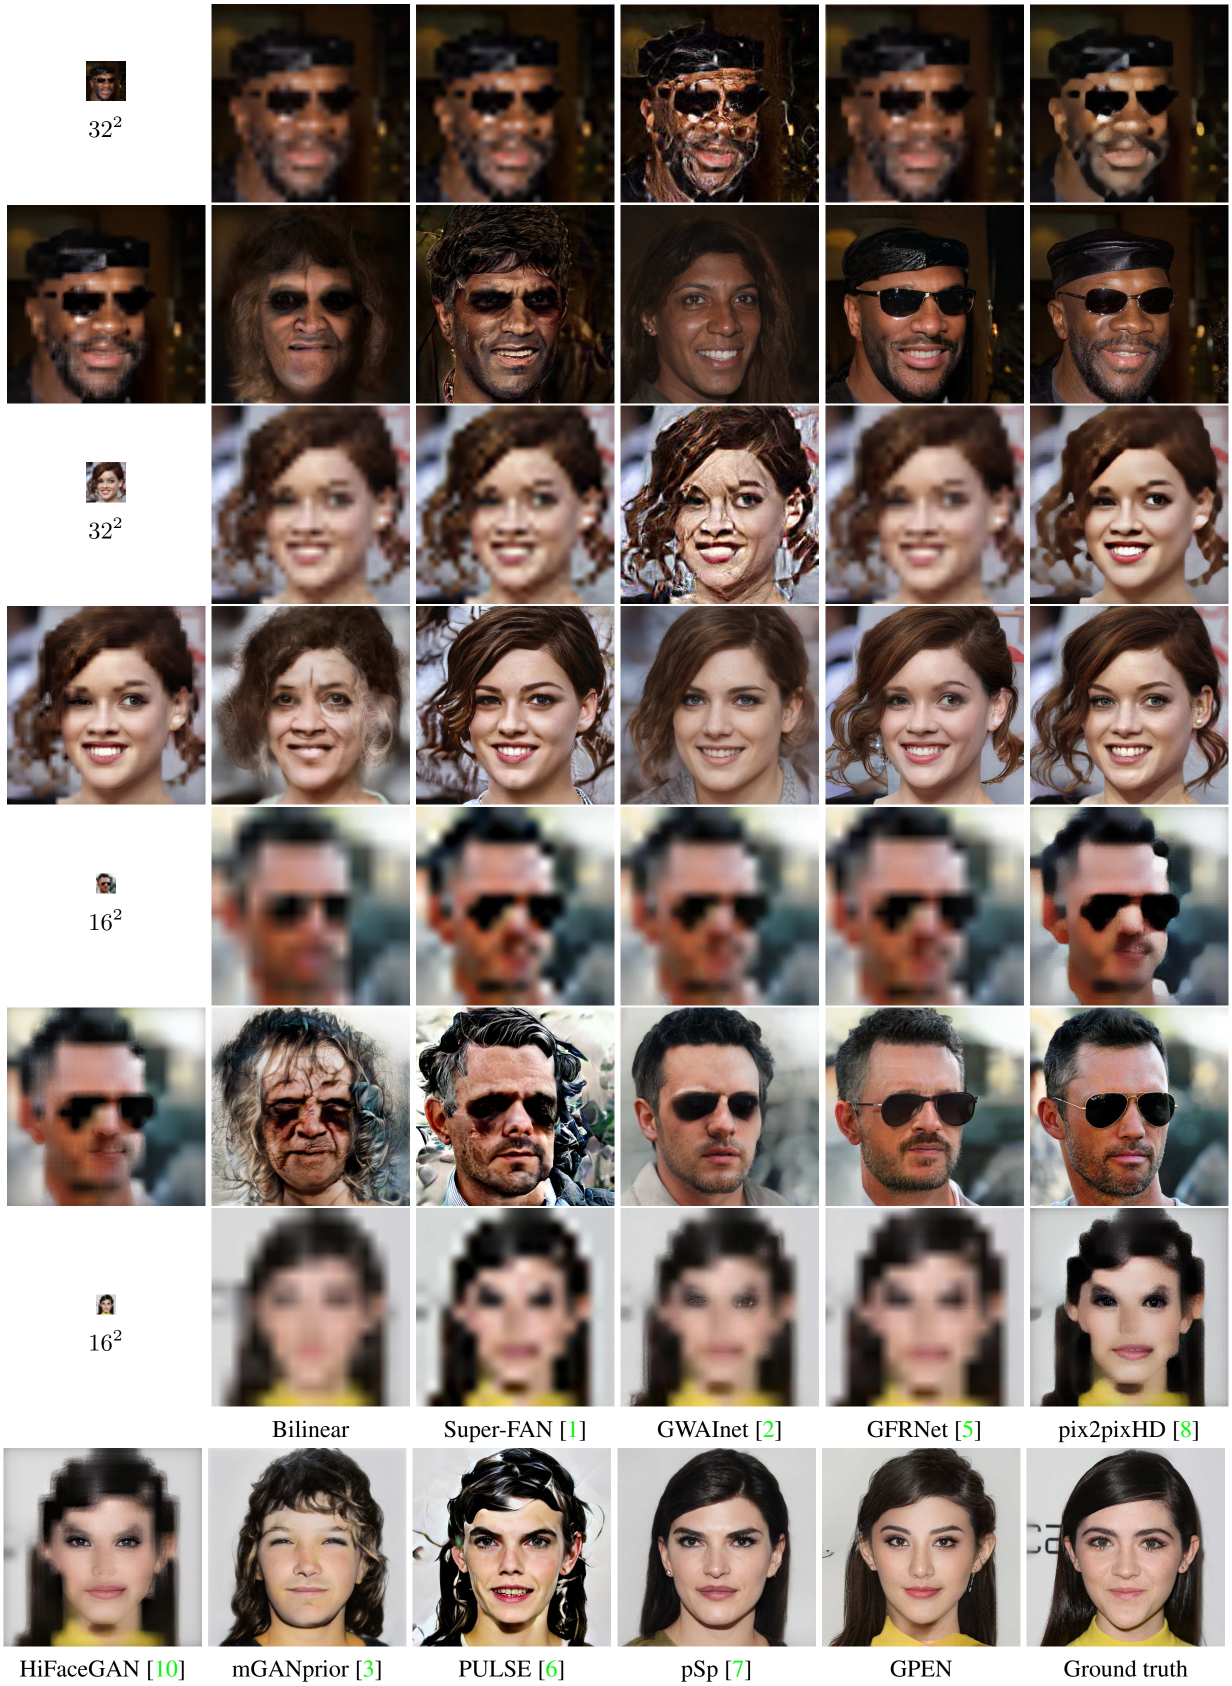

Figure 2: Face super-resolution results by state-of-the-art methods.

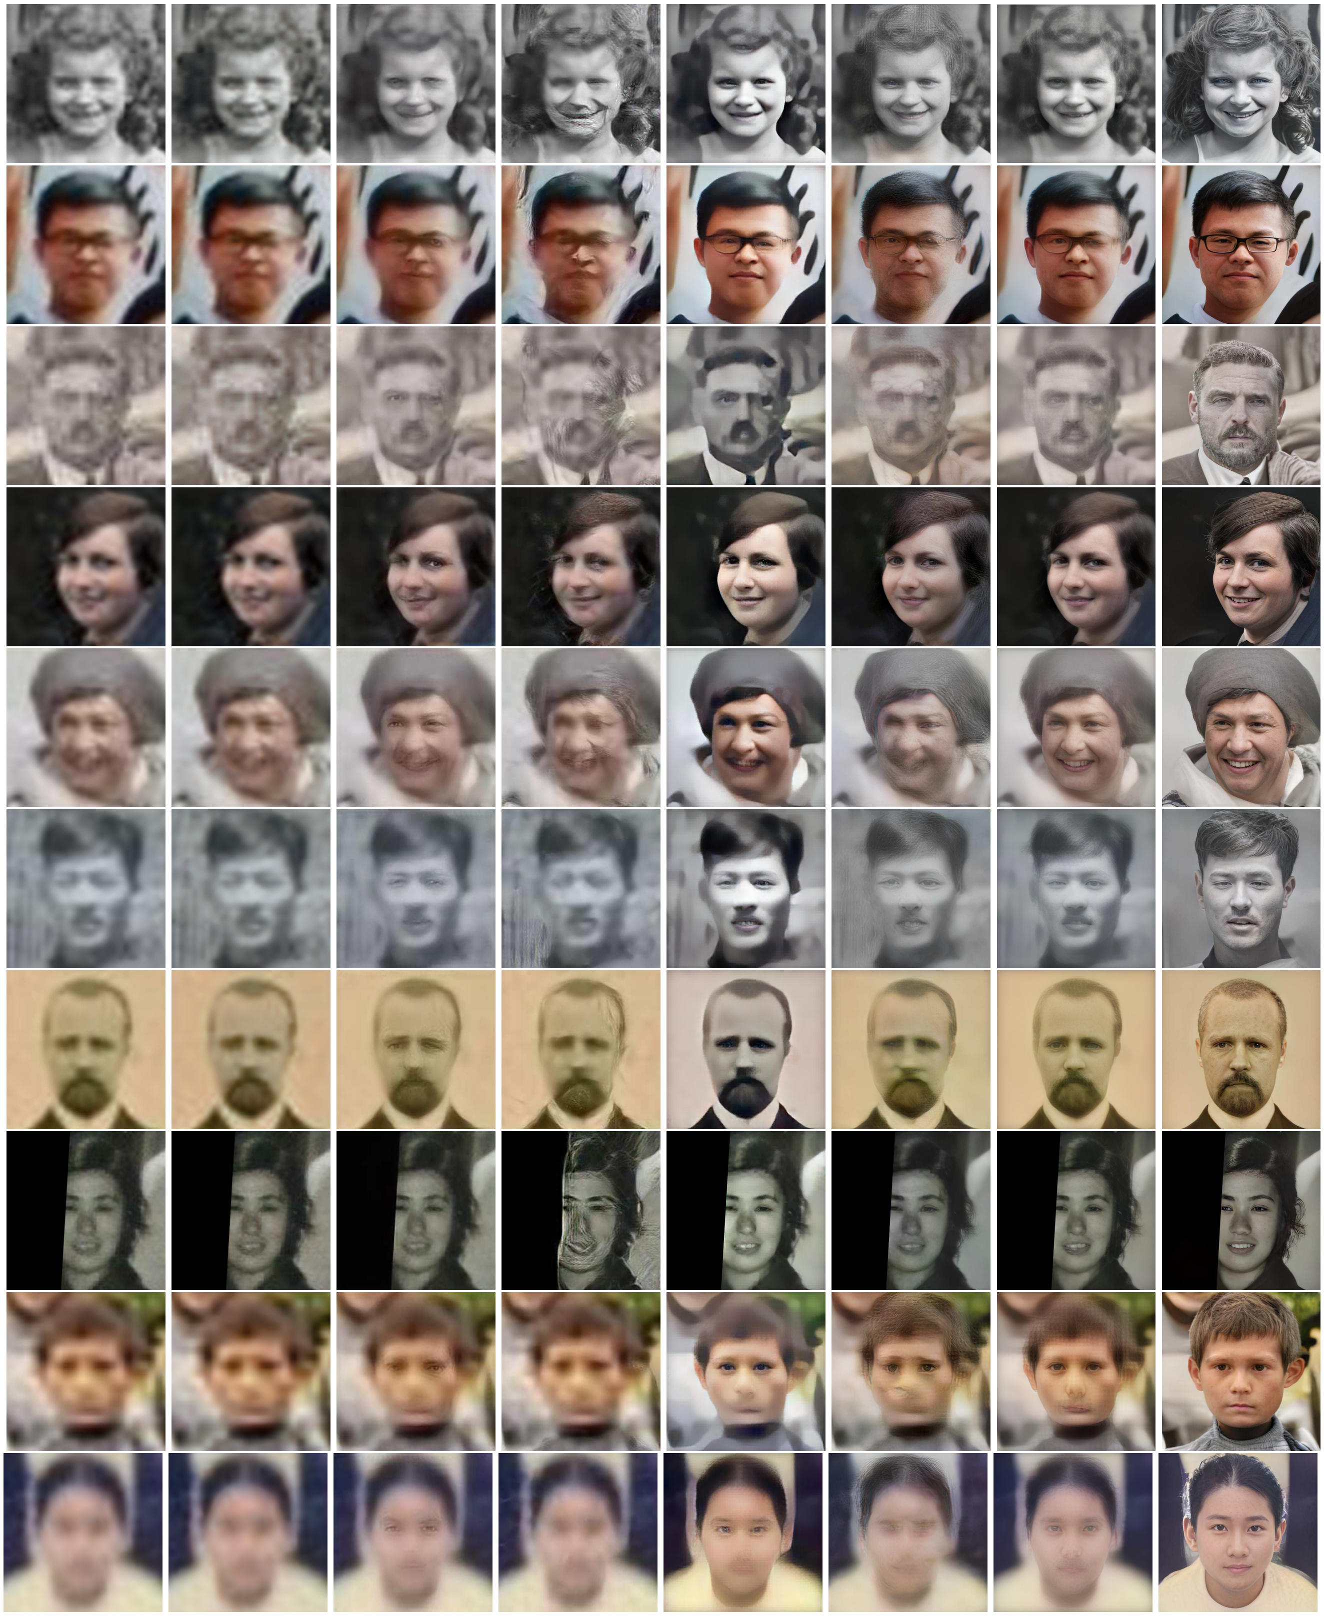

(a) (b) (c) (d) (e) (f) (g) (h)

Figure 3: Blind face restoration results on real-world degraded faces in the wild. (a) Real degraded faces; (b) Super-FAN [1]; (c) GFRNet [5]; (d) GWAInet [2]; (e) Pix2PixHD [8]; (f) DFDNet [4]; (g) HiFaceGAN [10]; (h) GPEN.

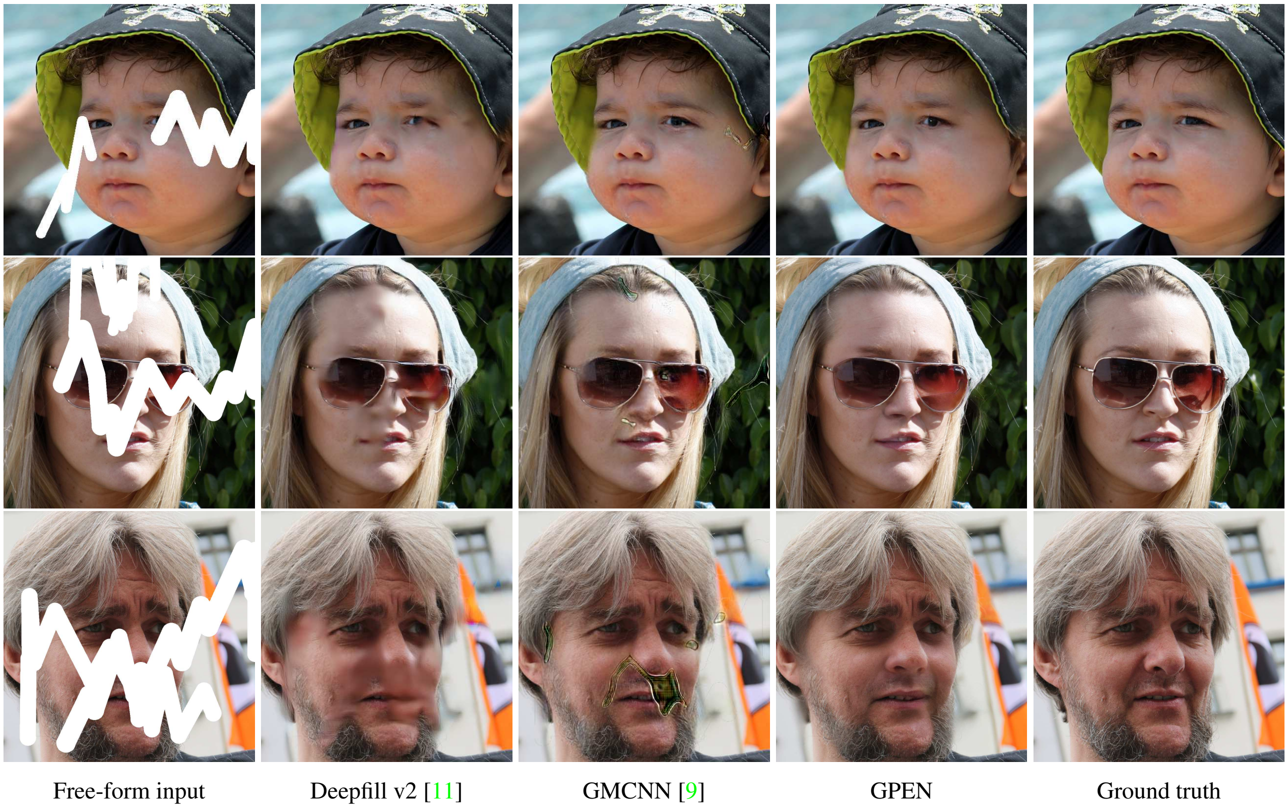

Figure 4: Qualitative comparison of different inpainting methods on high-quality faces (in 1024^2).

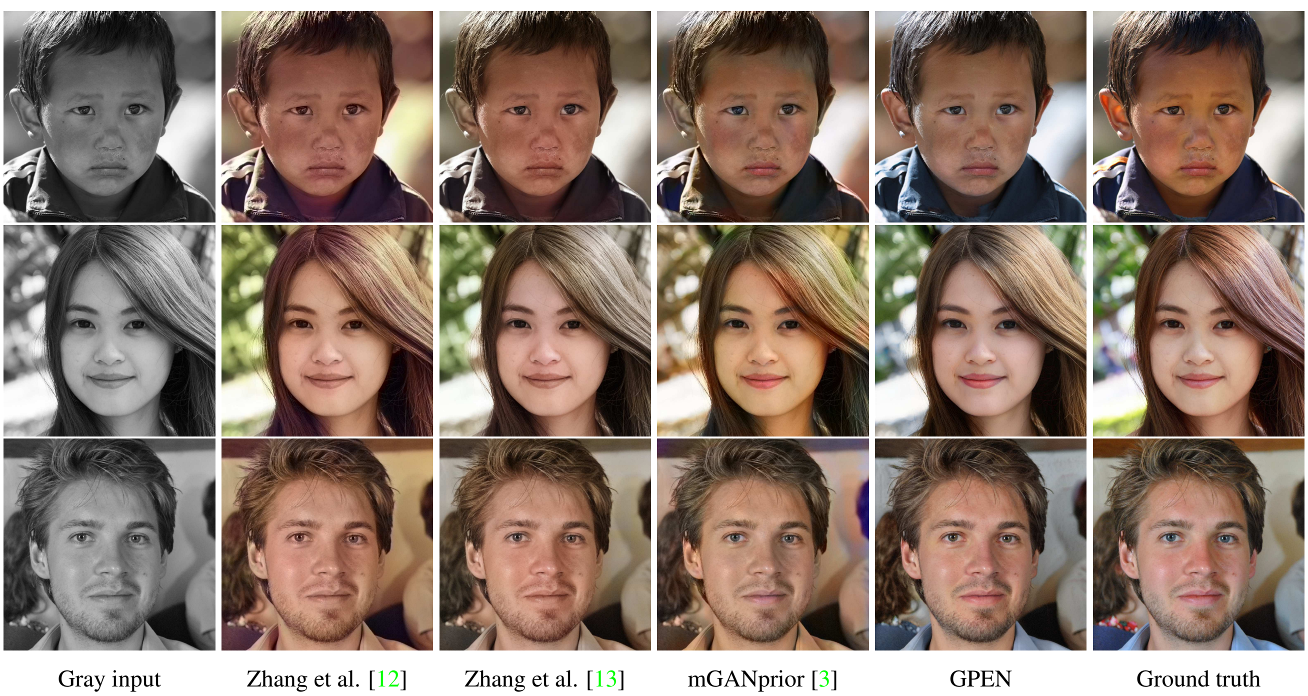

Figure 5: Qualitative comparison of different colorization methods on high-quality faces (in 1024^2).

References

[1] Adrian Bulat and Georgios Tzimiropoulos. Super-fan: Integrated facial landmark localization and super-resolution of real-world low resolution faces in arbitrary poses with gans. In *CVPR*, 2018. 1, 2, 3, 4

[2] Berk Dogan, Shuhang Gu, and Radu Timofte. Exemplar guided face image super-resolution without facial landmarks. In *CVPRW*, 2019. 1, 2, 3, 4

[3] Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code gan prior. *ArXiv*, 2019. 1, 3, 5

[4] Xiaoming Li, Chaofeng Chen, Shangchen Zhou, Xianhui Lin, Wangmeng Zuo, and Lei Zhang. Blind face restoration via deep multi-scale component dictionaries. In *ECCV*, 2020. 1, 2, 4

[5] Xiaoming Li, Ming Liu, Yuting Ye, Wangmeng Zuo, Liang Lin, and Ruigang Yang. Learning warped guidance for blind face restoration. In *ECCV*, 2018. 1, 2, 3, 4

[6] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *CVPR*, 2020. 1, 3

[7] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. *Arxiv*, 2020. 1, 3

[8] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *CVPR*, 2018. 1, 2, 3, 4

[9] Yi Wang, Xin Tao, Xiaojuan Qi, Xiaoyong Shen, and Jiaya Jia. Image inpainting via generative multi-column convolutional neural networks. *ArXiv*, 2018. 1, 5

[10] Lingbo Yang, Chang Liu, Pan Wang, Shanshe Wang, Peiran Ren, Siweia Ma, and Gao Wen. Hifacegan: Face renovation via collaborative suppression and replenishment. *Arxiv*, 2020. 1, 2, 3, 4

[11] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. *ArXiv*, 2018. 1, 5

[12] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016. 1, 5

[13] Richard Zhang, Jun-Yan Zhu, Phillip Isola, Xinyang Geng, Angela S Lin, Tianhe Yu, and Alexei A Efros. Real-time user-guided image colorization with learned deep priors. *ACM Transactions on Graphics (TOG)*, 9(4), 2017. 1, 5