RB Journal of Lib & Information Science ISSN: 0972-2750
(UGC Care Group I Listed Journal) Vol-12 Issue-06 No.01: 2022

TOWARDS REAL-WORLD BLIND FACE RESTORATION USING GFP

AVS RADHIKA[1]

ASSISTANT PROFESSOR, DEPARTMENT OF CSE, BHOJ REDDY ENGINEERING COLLEGE FOR WOMEN, VINAY NAGAR, HYDERABAD-59

BOLUGAM AKSHAYA GOUD[2]

UG SCHOLAR, DEPARTMENT OF CSE, BHOJ REDDY ENGINEERING COLLEGE FOR WOMEN, VINAY NAGAR, HYDERABAD-59

THALLAPALLI HARIKA[3]

UG SCHOLAR, DEPARTMENT OF CSE, BHOJ REDDY ENGINEERING COLLEGE FOR WOMEN, VINAY NAGAR, HYDERABAD-59

YANAMALA HRUSHITHA[4]

UG SCHOLAR, DEPARTMENT OF CSE, BHOJ REDDY ENGINEERING COLLEGE FOR WOMEN, VINAY NAGAR, HYDERABAD-59

ABSTRACT

Blind face restoration usually relies on facial priors, such as facial geometry prior or reference prior, to restore realistic and faithful details. However, very low-quality inputs cannot offer accurate geometric prior while highquality references are inaccessible, limiting the applicability in real-world scenarios. In this work, we propose GFP-GAN that leverages rich and diverse priors encapsulated in a pretrained face GAN for blind face restoration. This Generative Facial Prior (GFP) is incorporated into the face restoration process via spatial feature transform layers, which allow our method to achieve a good balance of realness and fidelity. Thanks to the powerful generative facial prior and delicate designs, our GFP-GAN could jointly restore facial details and enhance colors with just a single forward pass, while GAN inversion methods require image-specific optimization at inference. Extensive experiments show that our method achieves superior performance to prior art on both synthetic and real-world datasets.

INTRODUCTION

Blind face restoration aims at recovering high-quality faces from the low-quality counterparts suffering from unknown degradation, such as low-resolution [13, 48, 9], noise [71], blur [39, 58], compression artifacts [12], etc. When applied to real-world scenarios, it becomes more challenging, due to more complicated degradation, diverse poses and expressions. Previous works [9, 69, 6] typically exploit face-specific priors in face restoration, such as facial landmarks [9], parsing maps [6, 9], facial component heatmaps [69], and show that those geometry facial priors are pivotal to recover accurate face shape and details. However, those priors are usually estimated from input images and inevitably degrades with very low-quality inputs in the real world. In addition, despite their semantic guidance, the above priors contain limited texture information for restoring facial details (e.g., eye pupil). Another category of approaches investigates reference priors, i.e., high-quality guided faces [46, 45, 11] or facial component dictionaries [44], to generate realistic results and alleviate the dependency on degraded inputs. However, the inaccessibility of high-resolution references limits its practical applicability, while the limited capacity of dictionaries restricts its diversity and richness of facial details. In this study, we leverage Generative Facial Prior (GFP) for real-world blind face restoration, i.e., the prior implicitly encapsulated in pretrained face Generative Adversarial Network (GAN) [18] models such as StyleGAN [35, 36]. These face GANs are capable of generating faithful faces with a high degree of variability, and thereby providing rich and diverse priors such as geometry, facial textures and colors, making it possible to jointly restore facial details and enhance colors (Fig. 1). However, it is challenging to incorporate such generative priors into the restoration process. Previous attempts typically use GAN inversion [19, 54, 52]. They first 'invert' the degraded image back to a latent code of the pretrained GAN, and then conduct expensive imagespecific optimization to reconstruct images. Despite visually realistic outputs, they usually produce images with low fidelity, as the low-dimension latent codes are insufficient to guide accurate restoration. To address these challenges, we propose the GFP-GAN with delicate designs to achieve a good balance of realness and fidelity in a single forward pass. Specifically, GFPGAN consists of a degradation removal module and a pretrained face GAN as facial prior. They are connected by a direct latent code mapping, and several Channel-Split Spatial Feature Transform (CS-SFT) layers in a coarse-to-fine manner. The proposed CS-SFT layers perform spatial modulation on a split of features and leave the left features to directly pass through for better information preservation, allowing our method to effectively incorporate generative prior while retraining high fidelity. Besides, we introduce facial component loss with local discriminators to further enhance perceptual facial details, while employing identity preserving loss to further improve fidelity. We summarize the contributions as follows. (1) We leverage rich and diverse generative facial priors for blind face restoration. Those priors contain sufficient facial textures and color information, allowing us to jointly perform face restoration and color enhancement. (2) We propose the GFP-GAN framework with delicate designs of architectures and losses to incorporate generative facial prior. Our GFP-GAN with CS-SFT layers achieves a good balance of fidelity and texture faithfulness in a single forward pass. (3) Extensive experiments show that our method achieves superior performance to prior art on both synthetic and realworld datasets.

EXISITNG SYSTEM

Despite visually realistic outputs, they usually produce images with low fidelity, as the low-dimension latent codes are insufficient to guide accurate restoration.

DISADVANTAGES

- When the degradation of real images is severe, the restored facial details by GFP-GAN are twisted with artifacts.
- Our method also produces unnatural results for very large poses

PROPOSED SYSTEM

- In this we propose the GFP-GAN with delicate designs to achieve a good balance of realness and fidelity in a single forward pass.
- Specifically, GFP-GAN consists of a degradation removal module and a pretrained face GAN as facial prior.

ADVANTAGES

- Pretrained GAN provides rich and diverse features for restoration.
- Our method performs well on most dark skinned faces and various population groups.

RELATED WORK

Image Restoration typically includes super-resolution [13, 48, 60, 49, 74, 68, 22, 50], denoising [71, 42, 26], deblurring [65, 39, 58] and compression removal [12, 21]. To achieve visually-pleasing results, generative adversarial network [18] is usually employed as loss supervisions to push the solutions closer to the natural manifold [41, 57, 64, 7, 14], while our work attempts to leverage the pretrained face GANs as generative facial priors (GFP). Face Restoration. Based on general face hallucination [5, 30, 66, 70], two typical face-specific priors: geometry priors and reference priors, are incorporated to further improve the performance. The geometry priors include facial landmarks [9, 37, 77], face parsing maps [58, 6, 9] and facial component heatmaps [69]. However, 1) those priors require estimations from low-quality inputs and inevitably degrades in real-world scenarios. 2) They mainly focus on geometry constraints and may not contain adequate details for restoration. Instead, our employed GFP does not involve an explicit geometry estimation from degraded images, and contains adequate textures inside its pretrained network. Reference priors [46, 45, 11] usually rely on reference images of the same identity. To overcome this issue, DFDNet [44] suggests to construct a face dictionary of each component (e.g., eyes, mouth) with CNN features to guide the restoration. However, DFDNet mainly focuses on components in the dictionary and thus degrades in the regions beyond its dictionary scope (e.g., hair, ears and face contour), instead, our GFP-GAN could treat faces as a whole to restore. Moreover, the limited size of dictionary restricts its diversity and richness, while the GFP could provide rich and diverse priors including geometry, textures and colors. Generative Priors of pretrained GANs [34, 35, 36, 3] is previously exploited by GAN inversion [1, 76, 54, 19], whose primary aim is to find the closest latent codes given an input image. PULSE [52] iteratively optimizes the latent code of StyleGAN [35] until the distance between outputs and inputs is below a threshold. mGANprior [19] attempts to optimize multiple codes to improve the reconstruction quality. However, these methods usually produce images with low fidelity, as the low-dimension latent codes are insufficient to guide the restoration. In contrast, our proposed CS-SFT modulation layers enable prior incorporation on multi-resolution spatial features to achieve high fidelity. Besides, expensive iterative optimization is not required in our GFP-GAN during inference. Channel Split Operation is usually explored to design compact models and improve model representation ability. MobileNet [28] proposes depthwise convolutions and GhostNet [23] splits the convolutional layer into two parts and uses fewer filters to generate intrinsic feature maps. Dual path architecture in DPN [8] enables feature re-usage and new feature exploration for each path, thus improving its representation ability. A similar idea is also employed in super-resolution [75]. Our CS-SFT layers share the similar spirits, but with different operations and purposes. We adopt spatial feature transform [63, 55] on one split and leave the left split as identity to achieve a good balance of realness and fidelity. Local Component Discriminators. Local discriminator is proposed to focus on local patch distributions [32, 47, 62]. When applied to faces, those discriminative losses are imposed on separate semantic facial regions [43, 20]. Our introduced facial component loss also adopts such designs but with a further style supervision based on the learned discriminative features

CONCLUSION

We have proposed the GFP-GAN framework that leverages the rich and diverse generative facial prior for the challenging blind face restoration task. This prior is incorporated into the restoration process with channel-split spatial feature transform layers, allowing us to achieve a good balance of realness and fidelity. Extensive comparisons demonstrate the superior capability of GFP-GAN in joint face restoration and color enhancement for real-world images, outperforming prior art.

REFERENCES

[1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In ICCV, 2019. 2, 4
[2] Yochai Blau, Roey Mechrez, Radu Timofte, Tomer Michaeli, and Lihi Zelnik-Manor. The 2018 pirm challenge on perceptual image super-resolution. In ECCVW, 2018. 6
[3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096, 2018. 2
[4] 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. 6, 7
[5] Qingxing Cao, Liang Lin, Yukai Shi, Xiaodan Liang, and Guanbin Li. Attention-aware face hallucination via deep reinforcement learning. In CVPR, 2017. 2
[6] Chaofeng Chen, Xiaoming Li, Lingbo Yang, Xianhui Lin, Lei Zhang, and Kwan-Yee K. Wong. Progressive semantic-aware style transformation for blind face restoration. arXiv:2009.08709, 2020. 1, 2, 6, 7, 8
[7] Jingwen Chen, Jiawei Chen, Hongyang Chao, and Ming Yang. Image blind denoising with generative adversarial network based noise modeling. In CVPR, 2018. 2
[8] Yunpeng Chen, Jianan Li, Huaxin Xiao, Xiaojie Jin, Shuicheng Yan, and Jiashi Feng. Dual path networks. In NeurIPS, 2017. 2
[9] Yu Chen, Ying Tai, Xiaoming Liu, Chunhua Shen, and Jian Yang. Fsrnet: End-to-end learning face super-resolution with facial priors. In CVPR, 2018. 1, 2, 4
[10] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In CVPR, 2019. 5, 6