

Rotate-and-Render: Unsupervised Photorealistic Face Rotation from Single-View Images (Supplementary Materials)

Hang Zhou^{1*} Jihao Liu^{2*} Ziwei Liu¹ Yu Liu^{1†} Xiaogang Wang¹
¹The Chinese University of Hong Kong ²SenseTime Research

{zhouhang@link, yuliu@ee, xgwang@ee}.cuhk.edu.hk liujihao@sensetime.com zwliu.hust@gmail.com



Figure 1: Additional frontalization result. The first rows are the inputs and the second rows are the rotated results.

A. Possibility for Improvements

In our paper [5], we describe the least building blocks for our proposed method to work. However, there is actually great space for improvements. For example, the 3DDFA [6] model used for 3D face fitting is quite inaccurate, thus can be replaced by any state-of-the-art model. Moreover, the Render-to-Image network could be replaced with the one in Pix2pix HD [4] or modified to a progressive growing style [1]. Upgraded training losses are also applicable.

Here we provide a simple modification to our Render-to-Image module. The idea is to use facial semantic maps to guide the generation procedure, which shares substantially similar idea with SPADE [3] and MaskGAN [2]. However, semantic maps are normally acquired by an additional face

parsing model. Differently, we leverage facial landmarks that can be directly derived from 3D models. Then we connect and expend the landmarks to get a rough prediction of the key components' semantic maps. The predicted maps are converted to weights and biases of the batch normalization parameters inside the ResNet blocks of our original generator G. Please refer to our released code ¹ for our implementation and more details.

B. Additional Results

We show more frontalization results in the figure 1 to validate that our method can work robustly with different

¹<https://github.com/Hangz-nju-cuhk/Rotate-and-Render>.

poses, illuminating conditions and qualities.

References

- [1] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *ICLR*, 2018. 1
- [2] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: towards diverse and interactive facial image manipulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 1
- [3] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019. 1
- [4] 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
- [5] Hang Zhou, Jihao Liu, Ziwei Liu, Yu Liu, and Xiaogang Wang. Rotate-and-render: Unsupervised photorealistic face rotation from single-view images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [6] Xiangyu Zhu, Xiaoming Liu, Zhen Lei, and Stan Z Li. Face alignment in full pose range: A 3d total solution. *IEEE transactions on pattern analysis and machine intelligence*, 41(1):78–92, 2017. 1