Unsupervised High-Fidelity Facial Texture Generation and Reconstruction - Supplementary Material

Ron Slossberg[1*], Ibrahim Jubran[2*], and Ron Kimmel[1]
ronslos@gmail.com, ibrahim.jub@gmail.com, ron@cs.technion.ac.il

[1] Technion Institute, Haifa, Israel
[2] University of Haifa, Haifa, Israel

1 Additional Results

First, we prepared an additional demonstration video that presents various angles of our output results; please see the following anonymous link: https://youtu.be/sm8xWxvAL9s. The video is best viewed in the highest possible resolution.

Furthermore, in what follows we include additional results obtained by our method. In Fig. 1 we demonstrated several generated results which include glasses. The glasses are generated within the facial texture causing a somewhat unrealistic result. This can be mitigated by latent feature manipulation or by simply eliminating samples with glasses from the training data. Figure 3 demonstrates numerous reconstruction results obtained from our reconstruction pipeline. We observe highly detailed fully textured faces reconstructed from any input pose. In Fig. 4 we demonstrate more generation results, further demonstrating our high-fidelity texture generation capability coupled with matching realistic geometry. Results are best viewed when zoomed-in. In Fig. 5 we demonstrate the performance of our reconstruction pipeline under side-view conditions of varying rotation angles. Finally, in Fig. 6 we demonstrate the effect of varying the regularization weight parameter k on the reconstruction outcome. As discussed previously, the proposed method decouples Lambertian lighting effects from the generated textures, allowing for realistic relighting of the produced models. We demonstrate this within the enclosed video file.

[*] These authors contributed equally to this work

Fig. 1. Generation Results that Include Glasses. Due to the presence of subjects wearing glasses within the FFHQ dataset used for training, in some cases, our output texture might contain glasses; see Discussion and Future Work Section.

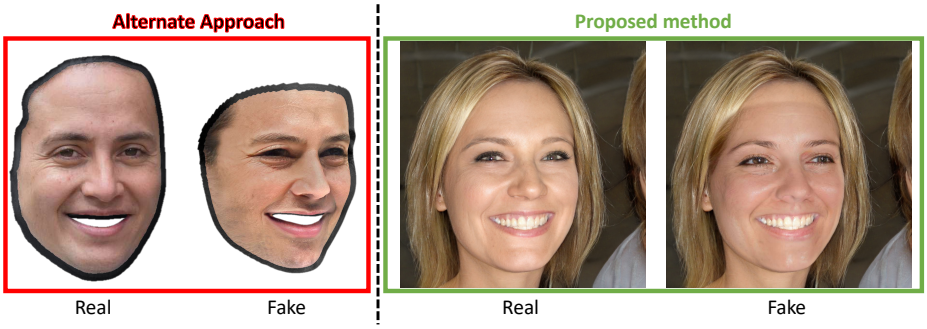

Fig. 2. Ablation study - background removal. Left: when removing background the outline of the face (bold black line) discloses the sample origin. Right: Rendering over a suitable background results in more realistic samples which are harder to distinguish from the real images.

Fig. 3. Additional Qualitative Reconstruction Results. We applied our reconstruction pipeline to numerous facial images. The 2D input images and the output textured models are presented side by side. The figure is best viewed when zoomed in.

Fig. 4. Additional Qualitative Generation Results. We used our generation pipeline to generate various random textures and geometries. The figure is best viewed when zoomed in.

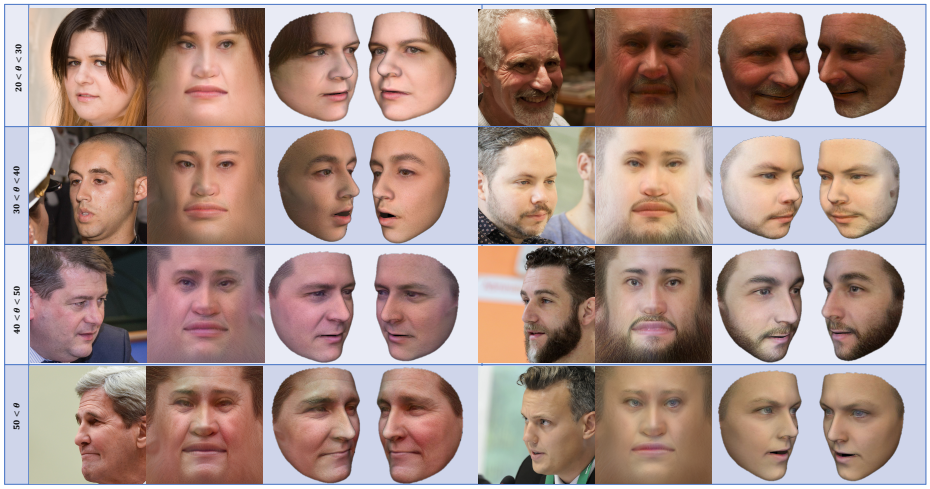

Fig. 5. Reconstruction of side views Here we demonstrate reconstruction results under various side-view conditions. We first estimate the head pose for samples of the FFHQ dataset and reconstruct several samples with head rotation from 20 up to above 50 degrees of rotation. We note that upwards of 99% of the dataset is within the range of $[-50, 50]$ degrees. We find that even when the face is highly rotated, the unobserved side of the face is reasonable completed in a realistic manner due to the learned facial model.

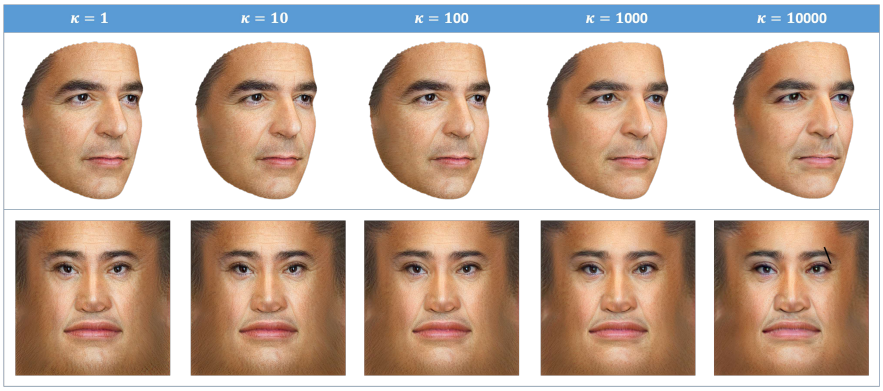

Fig. 6. Reconstruction Regularization Parameter. We demonstrate the effect of the regularization parameter κ on the reconstruction result. Stronger regularization leads to smoother texture with less details while weak regularization leads to high-details but also introduces some noise.