Multi-Directional Subspace Editing in Style-Space
(Supplementary Material)

Chen Naveh Yacov Hel-Or
School of Computer Science, Reichman University
navehchen1@gmail.com toky@runi.ac.il

In this supplementary document we provide some more results of our proposed method, higher resolution images and comparisons against other methods.

1. Qualitative Comparisons

Fig. 1 shows real image editing results of our method with comparison to baseline models. When editing the pose, we can inspect evidence of glasses in Sefa [3] and StyleFlow [1]. This suggests that those attributes are somewhat entangled. We can also inspect inferior image quality when all attributes are changed at once. Since our latent directions are orthogonal, by definition, changes within a subspace do not affect other subspaces. When performing multiple edits, this results in sharper images with less artifacts.

Figure 1. Real image editing comparison between our method and InterfaceGAN [2], Sefa [3] and StyleFlow [1]. On the right-most column we apply all previous edits at once.

2. Multi-Directional Edits

To demonstrate our multi-directional edits we first inverted real images into the latent space of StyleGAN. Then we chose different vectors inside a subspace to edit each image. The results are shown in Figs. 2 to 4.

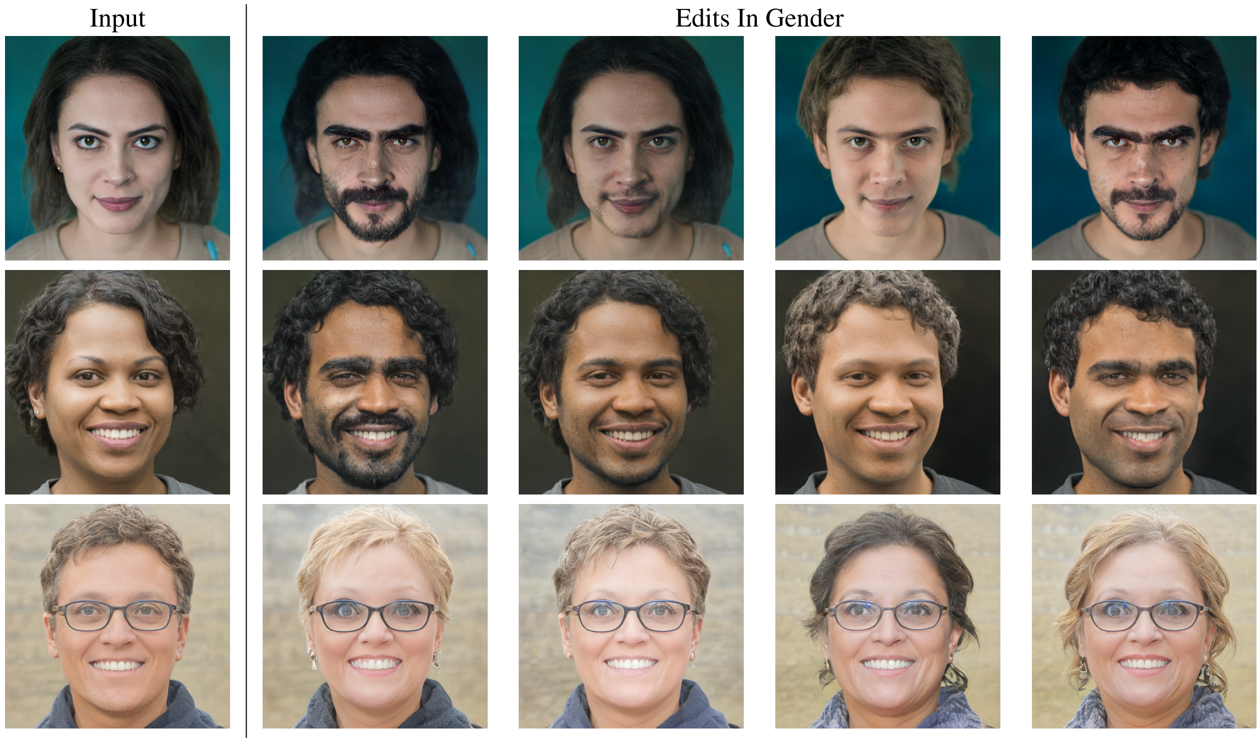

Figure 2. Multi-directional gender editing of real images. Each row displays different edits in the associated subspace.

Figure 3. Multi-directional age editing of real images. Each row displays different edits in the associated subspace.

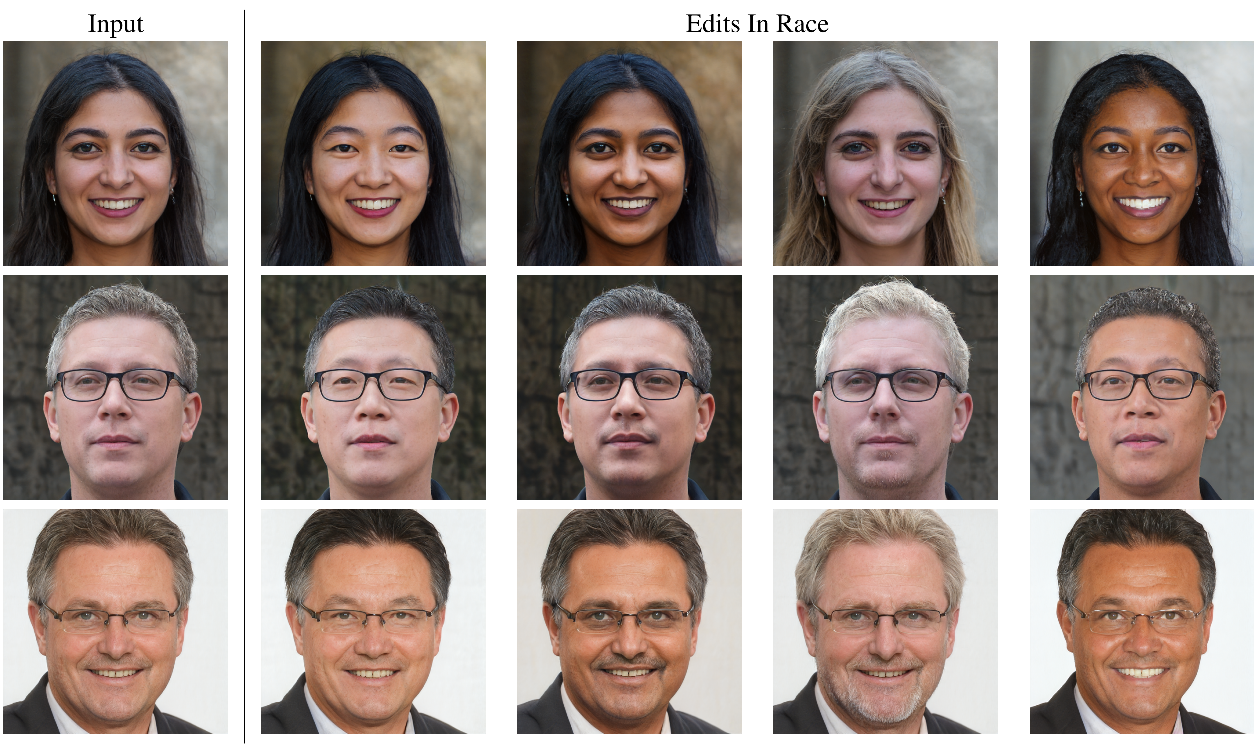

Figure 4. Multi-directional race editing of real images. Each row displays different edits in the associated subspace.

3. Comparison to StyleFlow

We compare our diverse edits results to StyleFlow as it outperforms all other baselines. The results are shown in Fig. 5. Different edits of our method are shown in Columns (a)-(c). We can observe the attributes manifest in various ways with each edit. Looking at gender, we can see that (b) and (c) better preserve the original hairstyle, while (c) also control facial hair. Changes in age may be expressed in hair color, hairstyle and facial wrinkles. Additionally, when comparing our results, we notice better preservation of unmodified attributes. For example, in the first two rows, our images preserve the original smile better than StyleFlow.

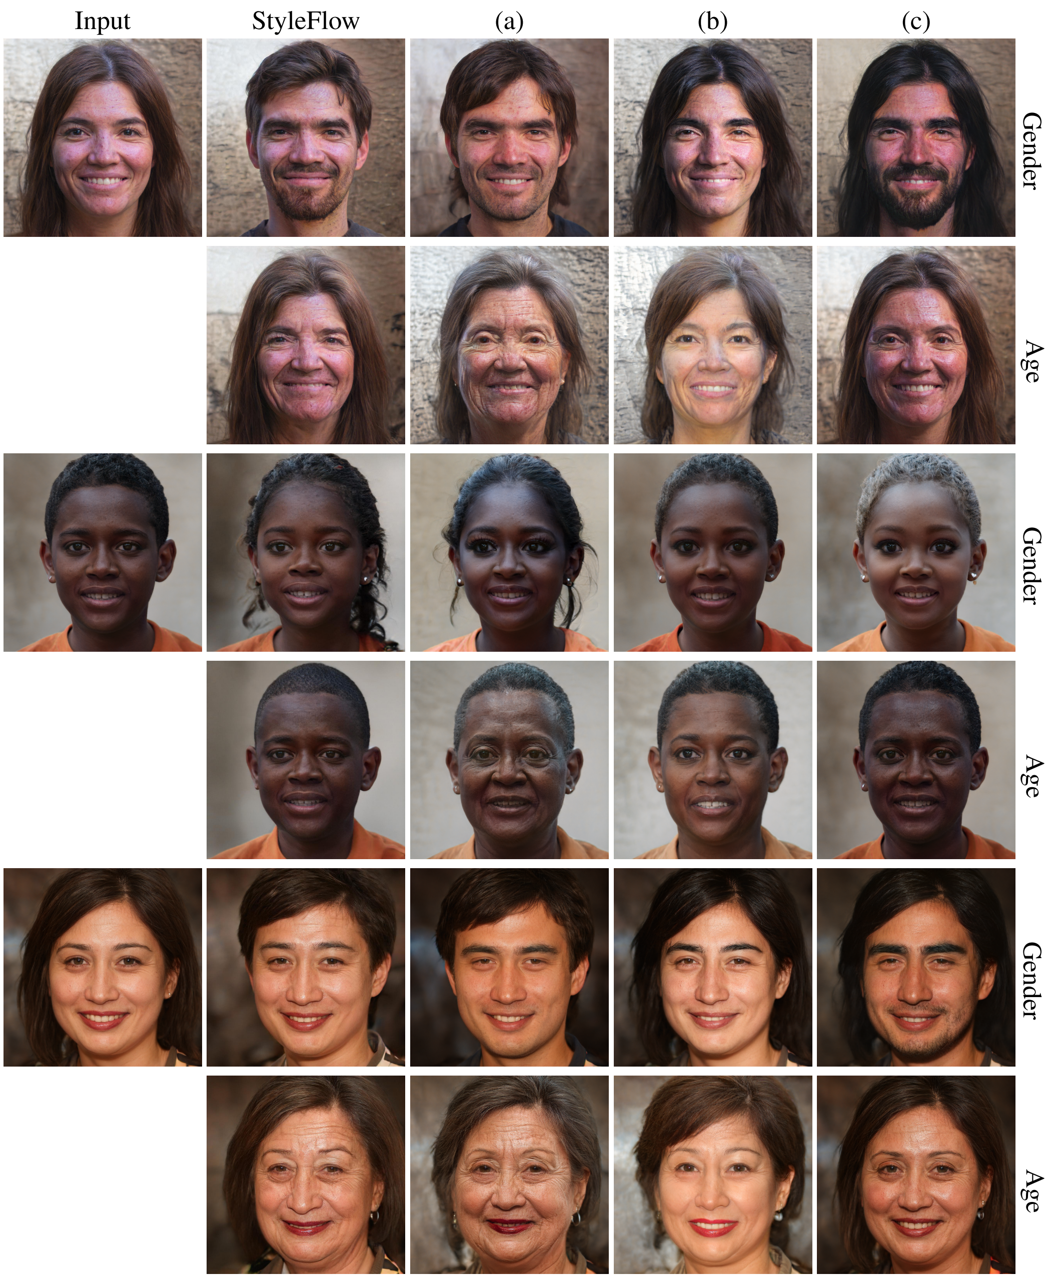

Figure 5. Multi-directional editing of real images in full resolution (1024×1024) and comparison to StyleFlow. Columns (a)-(c) represent edits in different directions inside the corresponding subspace (gender and age), allowing for rich and diverse edits.

4. Outside the Domain Editing

Here we display additional editing results of images outside the FFHQ dataset. These include paintings and old photography. Our model can generalize and perform diverse editing on images outside the training set. The results can be seen in the figure below.

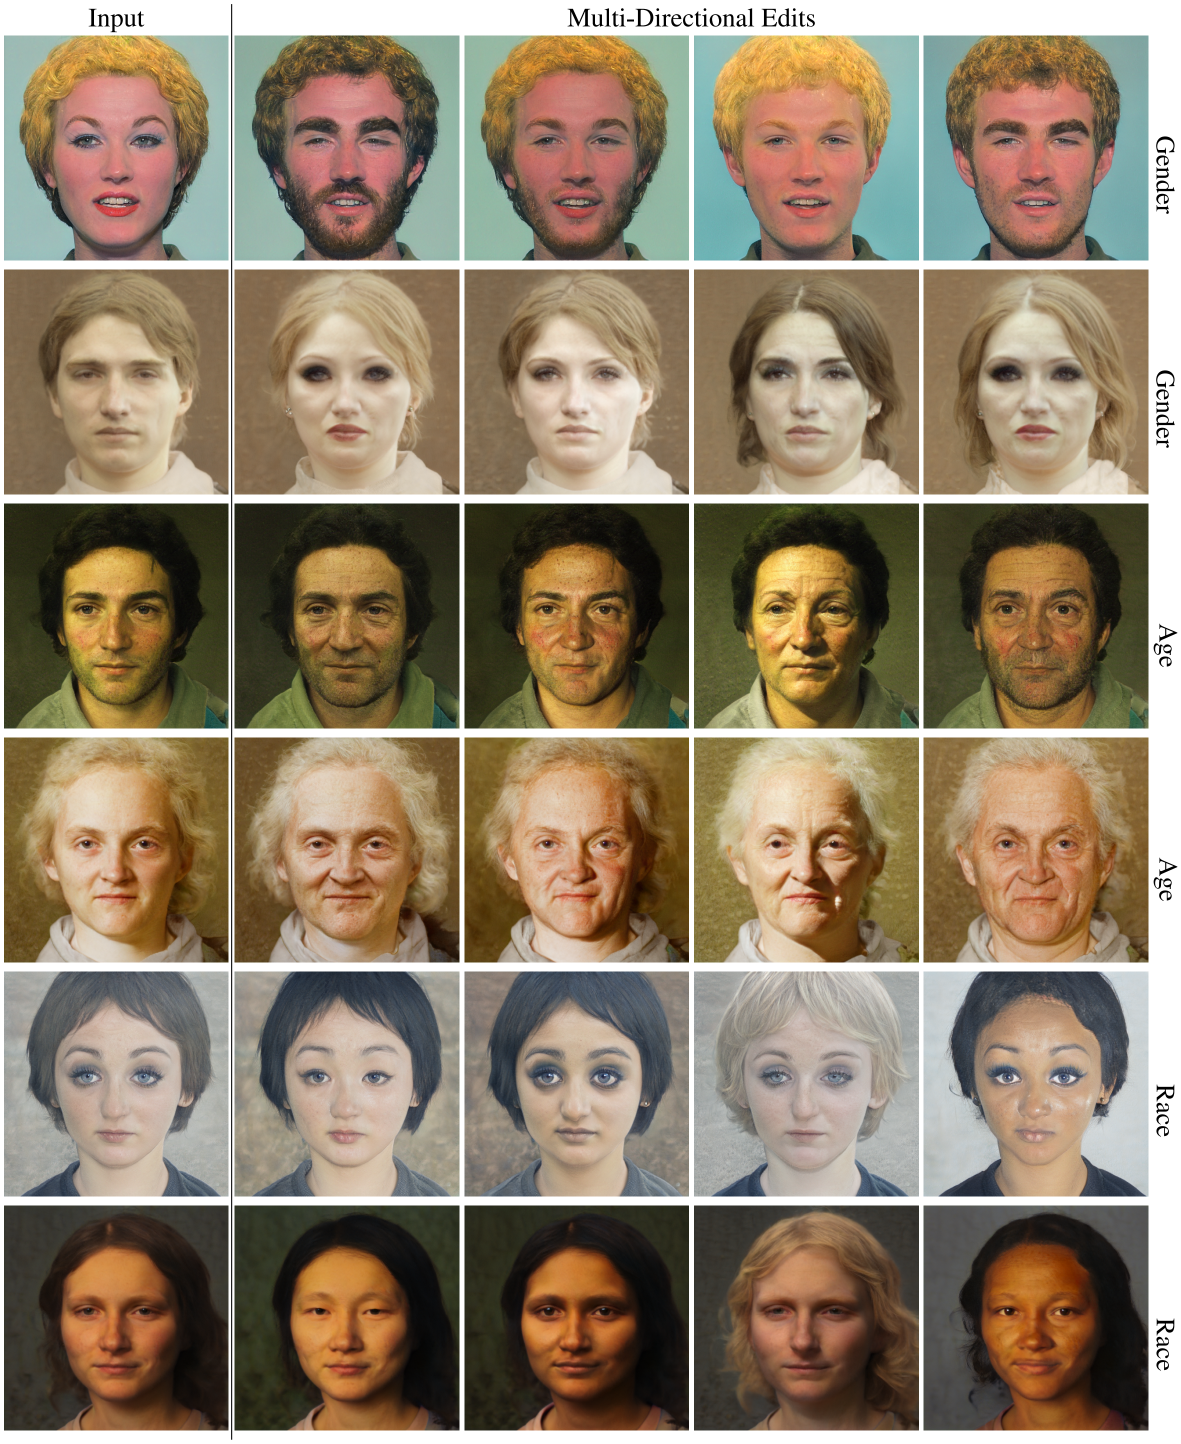

Figure 6. Additional human face editing results. Editing of images outside of the domain of StyleGAN.

References

[1] Rameen Abdal, Peihao Zhu, Niloy J Mitra, and Peter Wonka. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics (TOG)*, 40(3):1–21, 2021. 1

[2] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 1

[3] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1532–1540, 2021. 1