
FDNeRF: Semantics-Driven Face Reconstruction, Prompt Editing and Relighting with Diffusion Models

Anonymous Author(s)

Affiliation

Address

email

1 *Please also watch the supplementary video for visualization of results at various viewpoints.*

2 **1 Implementation Details**

3 Our reconstruction loss and identity loss are applied to the ground truth image camera poses. Due to
4 our limited GPU memory, we only render one side view to calculate the diffusion loss and illumination
5 loss at each iteration. The camera rotation angles θ and ϕ are randomly sampled from $[\frac{\pi}{2} - \frac{\pi}{12}, \frac{\pi}{2} + \frac{\pi}{12}]$
6 and $[\frac{\pi}{2} - \frac{\pi}{12}, \frac{\pi}{2} + \frac{\pi}{12}]$, where θ and ϕ are the angles of spherical coordinate. We set the optimization
7 iterations for our editing to 500, which takes approximately 10 minutes on a 3090 GPU. We set the
8 weighting $\mathcal{L}_D, \mathcal{L}_{ID}, L_R$ to be 0.2, 0.2 and 2×10^{-5} for most editing cases, which can be finetuned
9 for each editing.

10 **Dataset and Generative Bias** We utilize trained checkpoints of EG3D on FFHQ [3], AFHQv2 [2]
11 and ShapeNet [1] for the data domain of face, cat and car respectively. For some editing and
12 generation, we notice the existence of biases in generated results caused by the bias of the training
13 dataset, especially for race, gender, etc.

14 **2 More Results**

15 As a supplement to Figure 1 and Figure 3 in our main paper, fig. 1 shows more results of our FDNeRF.
16 All three figures illustrate that given a single face image, our FDNeRF can perform semantics-driven
17 NeRF editing on various features, such as expressions, emotions, glasses, hairstyles, races, genders,
18 ages, makeup, beard, mustache, and goatee. Notably, both individual and joint editing of these
19 features can be achieved in high fidelity.

20 **Text-conditioned Generation Comparison** Our attempts to compare with Dreamfusion[4] (Im-
21 plemented on a StableDiffusion) failed since it cannot generate faithful models for human heads.
22 This problem might be caused by various reasons. First, the ambiguity of text prompts (human
23 identity, rendering directions) might result in an inconsistent denoising behavior at each iteration.
24 Also, the randomized and unconstrained NeRF optimization process might collapse. On the contrary,
25 our utilization of latent code and trained 3D generative models ensures successful and high-quality
26 generation.

27 **3 Ablation Study on Editing Prompts**

28 We investigate the influence of the input text prompt in this section. In fig. 1, and Figure 1 and 3 in
29 our main paper, we mainly show NeRF editing results when the input text prompt is a short sentence
30 or a full description of the expected NeRF, such as “A woman wearing dark red lipstick”. While
31 this helps a valid generation that avoids ambiguity in semantics, our FDNeRF can also take in a text
32 prompt that only expresses the difference between the input face and the output face, such as “Dark
33 red lipstick”, if the input face has no lipstick at all. Here, fig. 2 shows the ablation results. The right

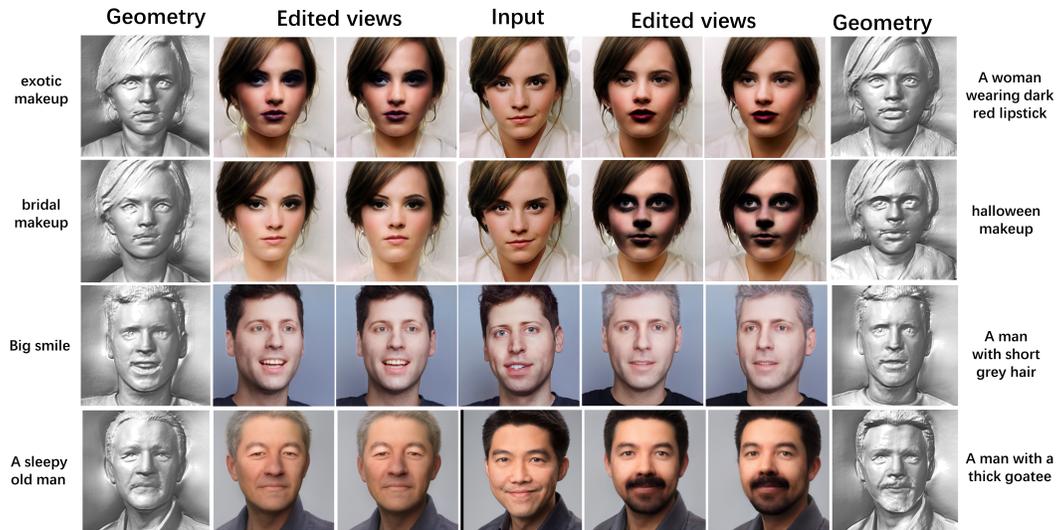


Figure 1: **More FDNerF results.** The middle column shows the input images, while the left and right halves of the images in the other two columns show the results of text prompts editing on the left and right, respectively. Additionally, the first and last columns showcase the corresponding geometries.

34 two columns show NeRF editing results when input text prompts are full descriptions of the expected
 35 editing results, while the left two columns are results with text prompts only describing what should
 36 be different.

37 In most cases, our FDNerF can generate similar editing results given either type of text prompt with
 38 subtle or even no tuning of weights of losses λ_{ID} , λ_R , and λ_D , as shown in fig. 2. However, sometimes
 39 ambiguity in semantics interferes editing when the input text prompt is not a full description of the
 40 expected output. As shown in fig. 3, even though our FDNerF can generate an old Elon Musk given
 41 “An old man” as the input text prompt, it fails when the input becomes just “Old”. We believe this is
 42 because “Old” has various meanings depending on its context. Therefore, a single “Old” without
 43 any context results in useless guidance from our diffusion model, thus producing poor editing results.
 44 “Elderly” and “senior” are two synonyms of “Old” in this context. As shown in fig. 3. “senior” also
 45 fails due to its semantic ambiguity with no context, while “elderly” succeeds because of its specific
 46 meaning.

47 Thus, users of our FDNerF are advised to give a better and complete text prompt to avoid semantics
 48 ambiguity.



Figure 2: **Ablation study on the text prompt.** On the right are results generated with full text prompts. The left ones are produced using simplified prompts. Given text prompt describing the wanted change, FDNeRF generates results with desired editing.

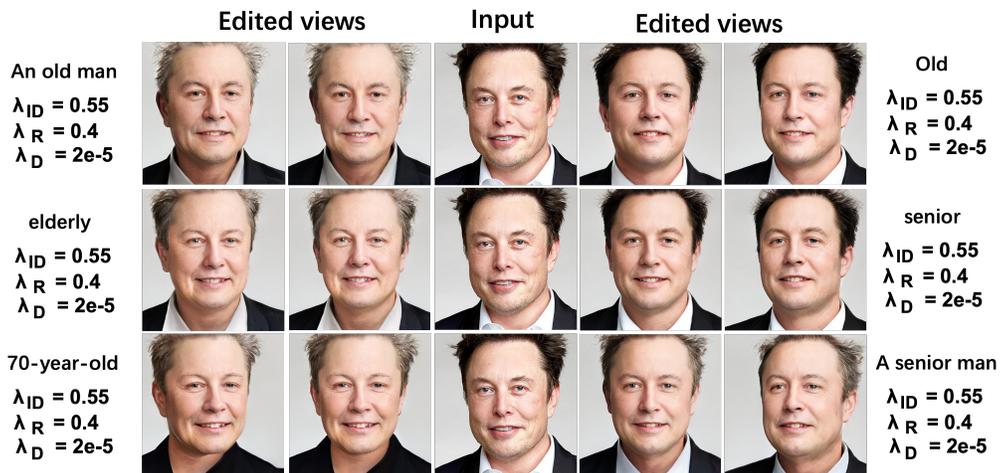


Figure 3: **Ablation study on synonyms.** Here are results generated with various Synonyms, which could affect the results. The context is important for synonyms to eliminate ambiguity, as the case of "senior" and "A senior man".

49 **References**

- 50 [1] Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva,
51 M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F., 2015. Shapenet: An information-rich 3d model
52 repository. arXiv:1512.03012.
- 53 [2] Choi, Y., Uh, Y., Yoo, J., Ha, J.W., 2020. Stargan v2: Diverse image synthesis for multiple
54 domains, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
- 55 [3] Karras, T., Laine, S., Aila, T., 2018. Flickr faces hq (ffhq) 70k from
56 stylegan. CoRR URL: [https://github.com/NVLabs/ffhq-dataset/blob/
57 93955b7cd435b7b1c724f8ca6a0e0c391300fe83/README.md](https://github.com/NVLabs/ffhq-dataset/blob/93955b7cd435b7b1c724f8ca6a0e0c391300fe83/README.md).
- 58 [4] Poole, B., Jain, A., Barron, J.T., Mildenhall, B., 2022. Dreamfusion: Text-to-3d using 2d
59 diffusion. arXiv .