

Supplementary Material for
**Steered Diffusion: A Generalized Framework for
Plug-and-Play Conditional Image Synthesis**

Nithin Gopalakrishnan Nair¹ Anoop Cherian² Suhas Lohit² Ye Wang²
Toshiaki Koike-Akino² Vishal M. Patel¹ Tim K. Marks²

¹ Johns Hopkins University ² Mitsubishi Electric Research Labs (MERL)

`{ngopala2, vpatel136}@jhu.edu {acherian, slohit, yewang, koike, tmarks}@merl.com`

1. Visualization of diffusion steering process

We present a visualization of intermediate outputs of the diffusion steering process in Figure 2 and Figure 3. The left image denotes the case without the steering loss and the right denotes the case with a steering loss. As we can see, with the presence of the steering loss, the generated images are consistent with the semantic maps from an early stage and the results continue to improve. The consistency grows stronger as timesteps progress.

2. Illustrating sample diversity using our method

We present non-cherry-picked results for various conditional generation tasks to demonstrate our method’s photo-realistic image generation quality and the diversity of examples generated by our method. Here we use the same noise levels across different examples, and the generation condition is shown in the first image in each sequence of images.




Figure 1. Qualitative comparisons for segmentation labels from Figure 5 from original paper.

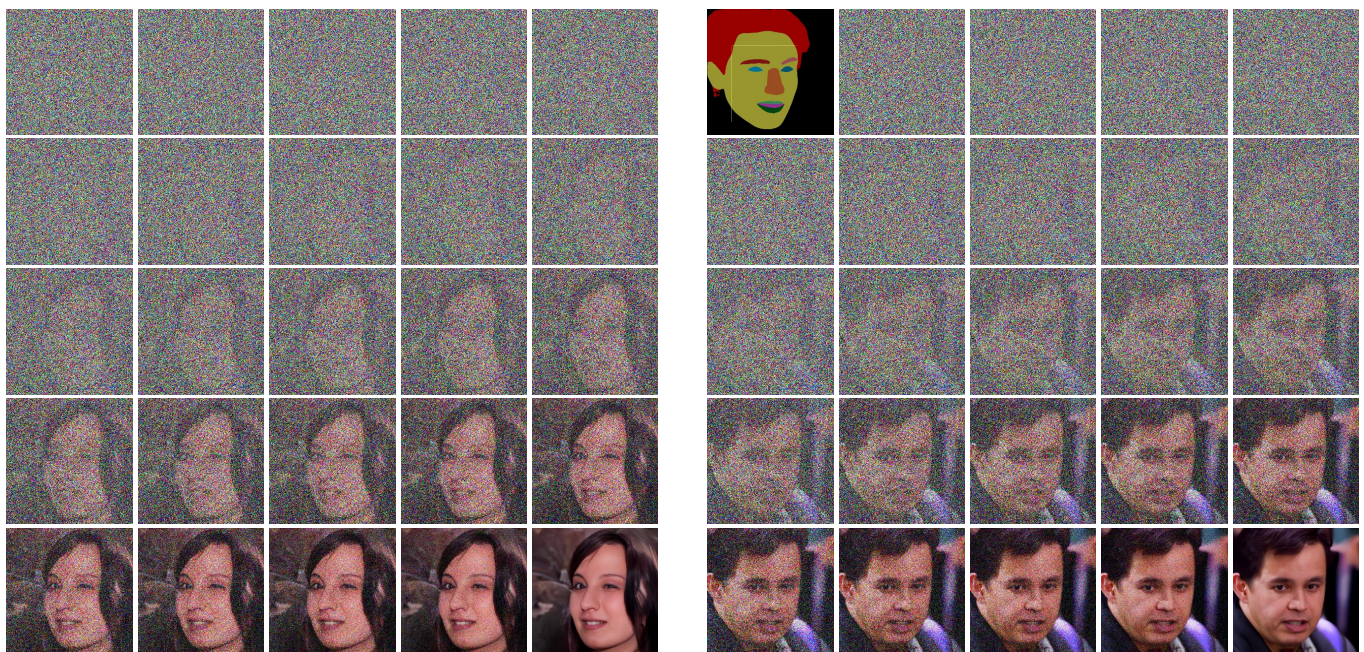


Figure 2. Evolution of sampling process without and with guidance

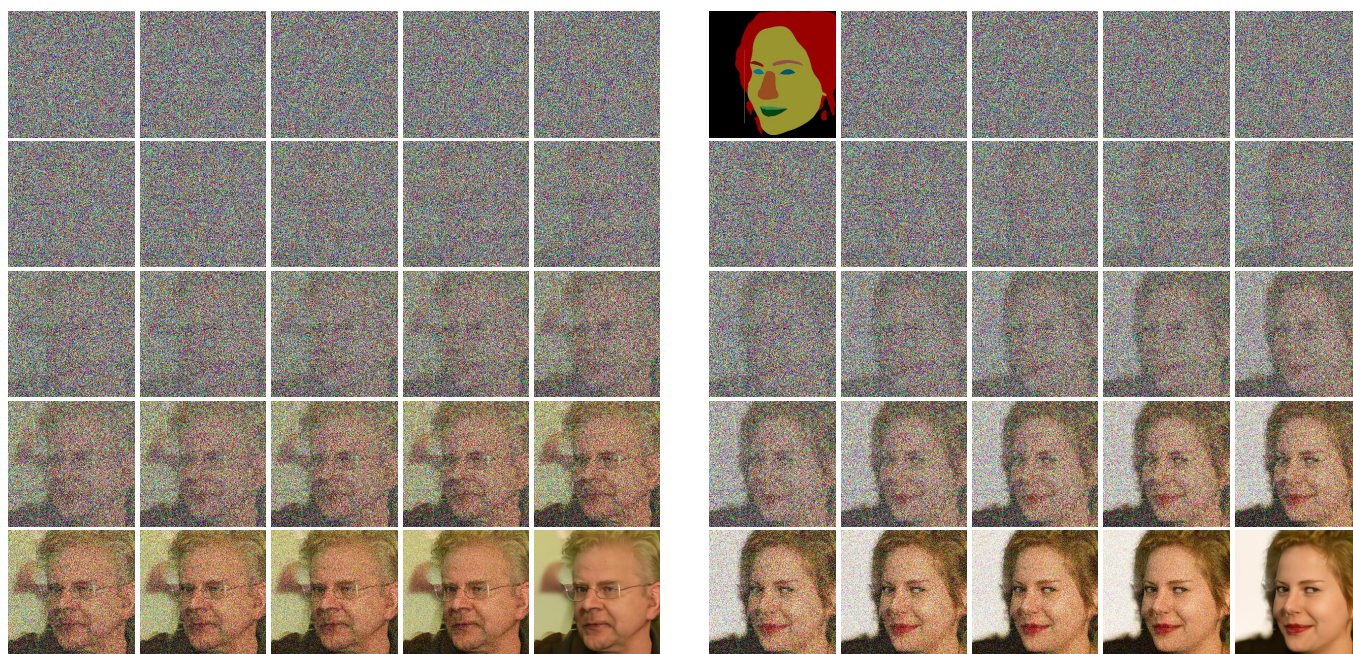


Figure 3. Evolution of sampling process without and with guidance



Figure 4. Non-cherry-picked samples from our method corresponding to the same semantic map. The upper left figure shows the corresponding semantic map. Samples across different examples with same sampling location has same random seeds




Figure 5. Non-cherry-picked samples from our method corresponding to the same semantic map. The upper left figure shows the corresponding semantic map. Samples across different examples with same sampling location has same random seeds



Figure 6. Non-cherry-picked samples from our method corresponding to the same identity image. The identity image projects out from the rest of the images. Samples across different examples with same sampling location has same random seeds

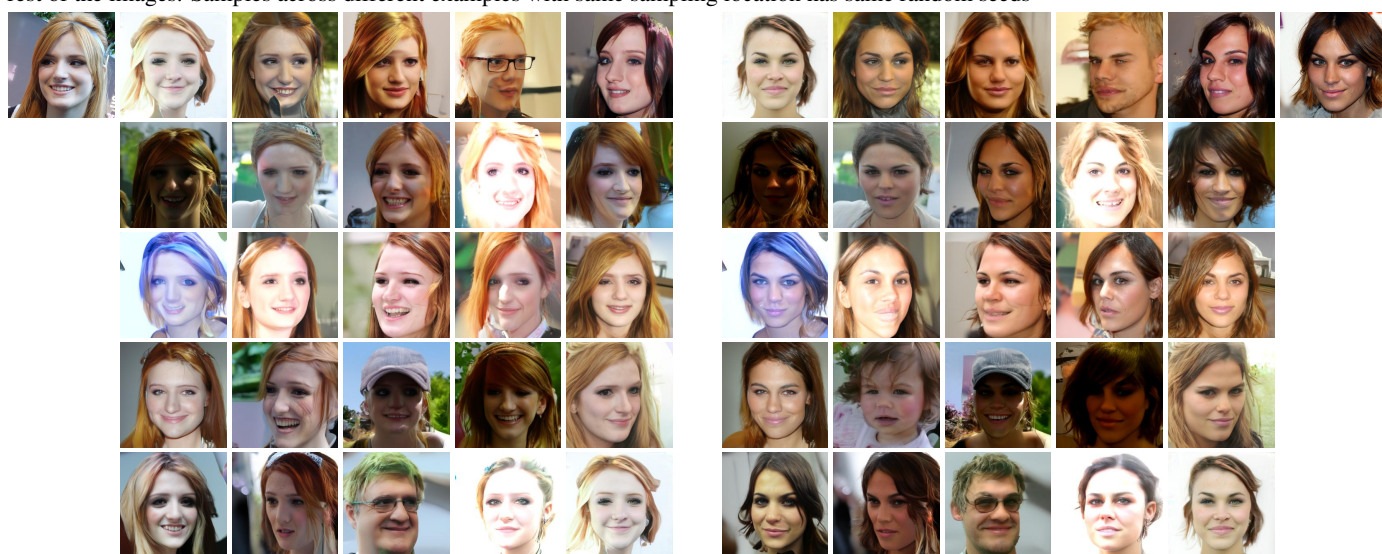


Figure 7. Non-cherry-picked samples from our method corresponding to the same identity image. The identity image projects out from the rest of the images. Samples across different examples with same sampling location has same random seeds




Figure 8. Non-cherry-picked samples from our method corresponding to grayscale to RGB. The grayscale projects out from the rest of the images. Samples across different examples with same sampling location has same random seeds




Figure 9. Non-cherry-picked samples from our method corresponding to grayscale to RGB. The grayscale projects out from the rest of the images. Samples across different examples with same sampling location has same random seeds



Figure 10. Non-cherry-picked samples from our method corresponding to LR to SR. The grayscale projects out from the rest of the images. Samples across different examples with same sampling location has same random seeds



Figure 11. Non-cherry-picked samples from our method corresponding to LR to SR. The grayscale projects out from the rest of the images. Samples across different examples with same sampling location has same random seeds