Supplementary material for Rayleigh EigenDirections (REDs): Nonlinear GAN latent space traversals for multidimensional features

Guha Balakrishnan[1,2], Raghudeep Gadde[2], Aleix Martinez[2], and Pietro Perona[2]

[1] Rice University, Houston TX 77005, USA
[2] Amazon, Seattle WA 98109, USA

A Equivalence of maximizing the Rayleigh quotient and solving a generalized eigenvalue problem

This is a classic result which we reproduce here. Consider the problem of maximizing r with respect to a vector $v \in \mathrm{R}^N$, where r is defined as follows:

$$r = \frac{v^T Av}{v^T Bv} \tag{1}$$

Set $\frac{dr}{dv} = 0$ to find v^* that maximizes r:

$$\frac{(v^{*T} Bv^*)Av^* - Bv^*(v^{*T} Av^*)}{\|v^{*T} Bv^*\|^2} = 0 \tag{2}$$

assuming $\|v^{*T} Bv^*\|^2 > 0$ we can simplify the expression to obtain:

$$(v^{*T} Bv^*)Av^* = Bv^*(v^{*T} Av^*). \tag{3}$$

Rearranging terms, we obtain the generalized eigenvalue equation:

$$Av^* = r^* Bv^*, \tag{4}$$

where r^* is the maximum value of r.

B Traversal algorithms

Algorithms 2 and 3 provide full details for our *Linear* and *Projection* traversal methods using the *getREDs* method described in Algorithm 1 of main paper.

Algorithm 2: *Linear* traversal

Input: $\mathbf{z}_0$, β_f, β_c, T (path length), s (path step)
Output: $\mathbf{z}_1, \cdots, \mathbf{z}_T$

$A_f, A_c \leftarrow$ finite difference approximation around $\mathbf{z}_0$
$R_0 \leftarrow getREDs(A_f, A_c, \beta_f, \beta_c)$
$\mathbf{v} \sim \text{Unif}(-1, 1)$
$\delta\mathbf{z}_0 \leftarrow R_0\mathbf{v}$

for $i = 0 \cdots T-1$ **do**
  $\mathbf{z}_{i+1} \leftarrow \mathbf{z}_i + s \cdot \delta\mathbf{z}_0 / \|\delta\mathbf{z}_0\|$

Algorithm 3: *Projection* traversal

Input: $\mathbf{z}_0$, β_f, β_c, T (path length), s (path step)
Output: $\mathbf{z}_1, \cdots, \mathbf{z}_T$

for $i = 0 \cdots T-1$ **do**
  $A_f, A_c \leftarrow$ finite difference approximation around $\mathbf{z}_i$
  $R_i \leftarrow getREDs(A_f, A_c, \beta_f, \beta_c)$
  $\mathbf{v} \sim \text{Unif}(-1, 1)$
  $\delta\mathbf{z}_i \leftarrow R_i\mathbf{v}$
  $\mathbf{z}_{i+1} \leftarrow \mathbf{z}_i + s \cdot \delta\mathbf{z}_i / \|\delta\mathbf{z}_i\|$

C Analysis of $\boldsymbol{\beta_f}$

Fig. A1 show the effect of varying β_f on our method (REDs-proj) when controlling eye and mouth facial regions (top two plots), and facial identity, geometry, and hairstyle (bottom two plots). See Figs. 7 and 3 in the main text for corresponding experiments and visual samples. Results vary smoothly with a change in β_f. The user may choose the appropriate value based on their application. In general, $\beta_f \in [0.95, 0.99]$ works well for our applications.

D Additional examples of region-based face traversals, and comparison to previous methods [1,2]

Fig. A2 presents additional samples of face traversals controlled by mouth and eye region bounding boxes. See Fig. 7 of the main text for more samples and description of experiments.

E Additional identity-geometry-hairstyle face traversals

Fig. A3 shows additional samples to our identity-geometry-hairstyle face traversals shown in Fig. 3 of the main text.

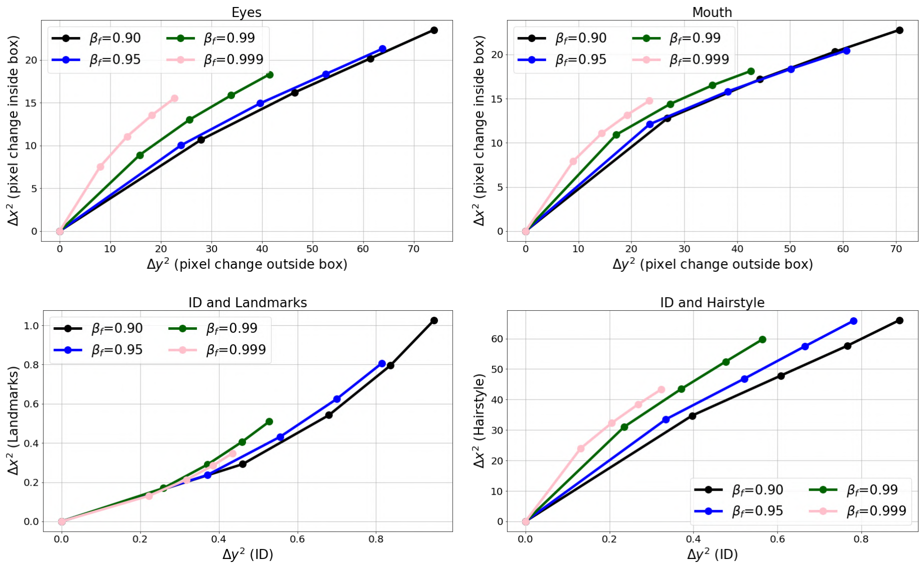

Fig. A1. Results over four β_f values when using REDs with *Projection* traversal for mouth and eye-region traversals. We use the same experimental setup as we used to produce the plots in Fig. 7 of the main text.

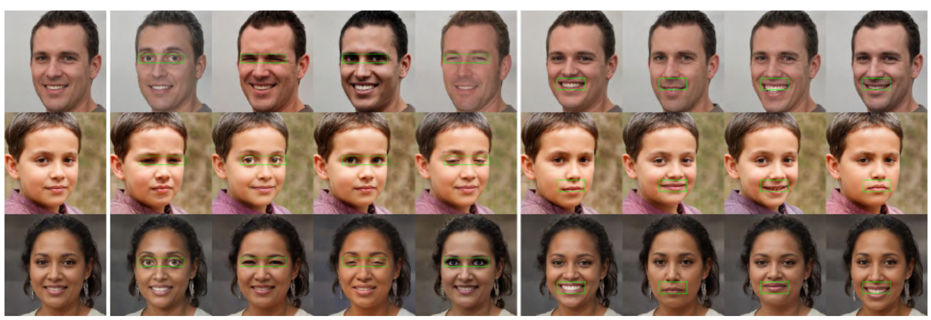

Fig. A2. Additional results of traversals controlled by spatial image regions. We let the changing features be the pixels inside a bounding box (green boxes overlaid on images for visualization), and the fixed features be pixels outside the box. See Fig. 7 of main text for more samples and description of experiments.

F Living-room traversals

Figs. A4, A5 and A6 present long versions of traversals. We also show an example traversal using each baseline method in Fig. A7.

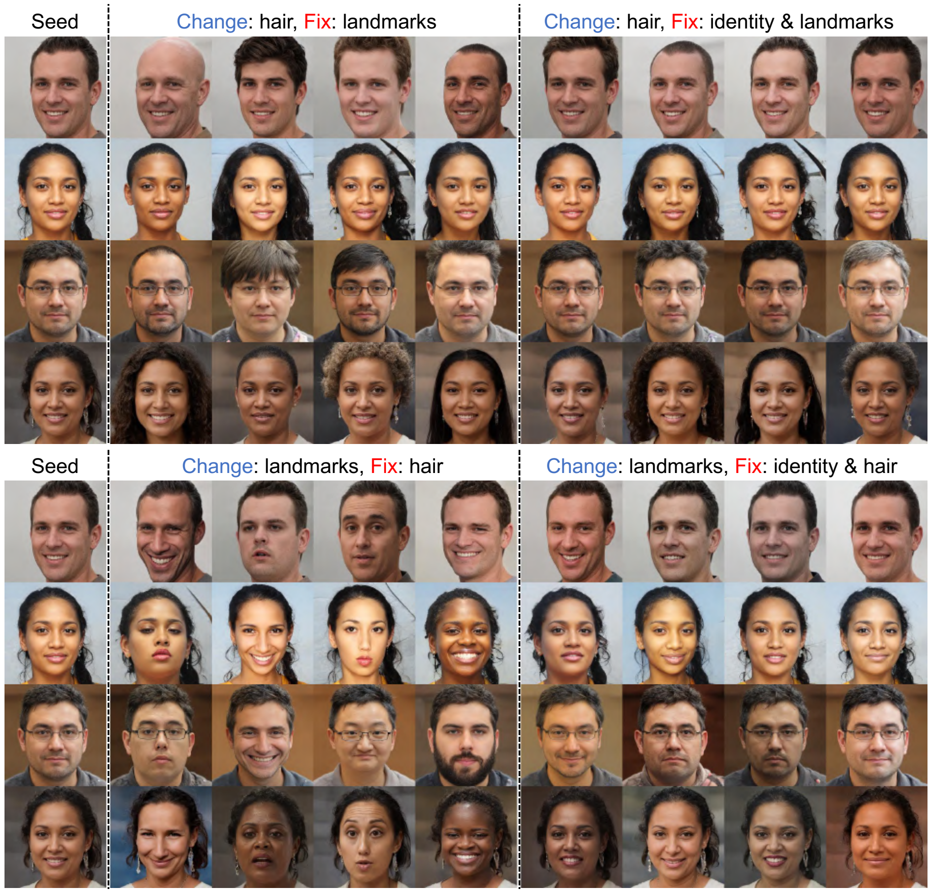

Fig. A3. Additional results for face traversals controlled by hairstyle, landmarks (geometry) and identity. See Fig. 3 of main text for more samples and description of experiment.

Table 1. Comparison of methods on object preserving traversals for livingrooms. We present the mean square difference of pixels inside (In) and outside (Out) the fixed bounding box roughly delineating the mouth for four traversal steps. We show results of our method (*REDs-proj*), along with two baselines from previous studies [2, 1]. Notice the step increase in mean square error on the pixels preserving an object as the length of the traversal increases for [2, 1]. See Figs. A4,A5 and A6 for visual samples.

	Step 1		**Step 2**		**Step 3**		**Step 4**	
Model	In	Out	In	Out	In	Out	In	Out
REDs-proj	0.010	0.038	0.015	0.044	0.017	0.047	0.020	0.068
LowRankGAN [2]	0.011	0.040	0.021	0.048	0.038	0.068	0.041	0.069
StyleSpace [1]	0.019	0.053	0.043	0.087	0.072	0.090	0.102	0.132

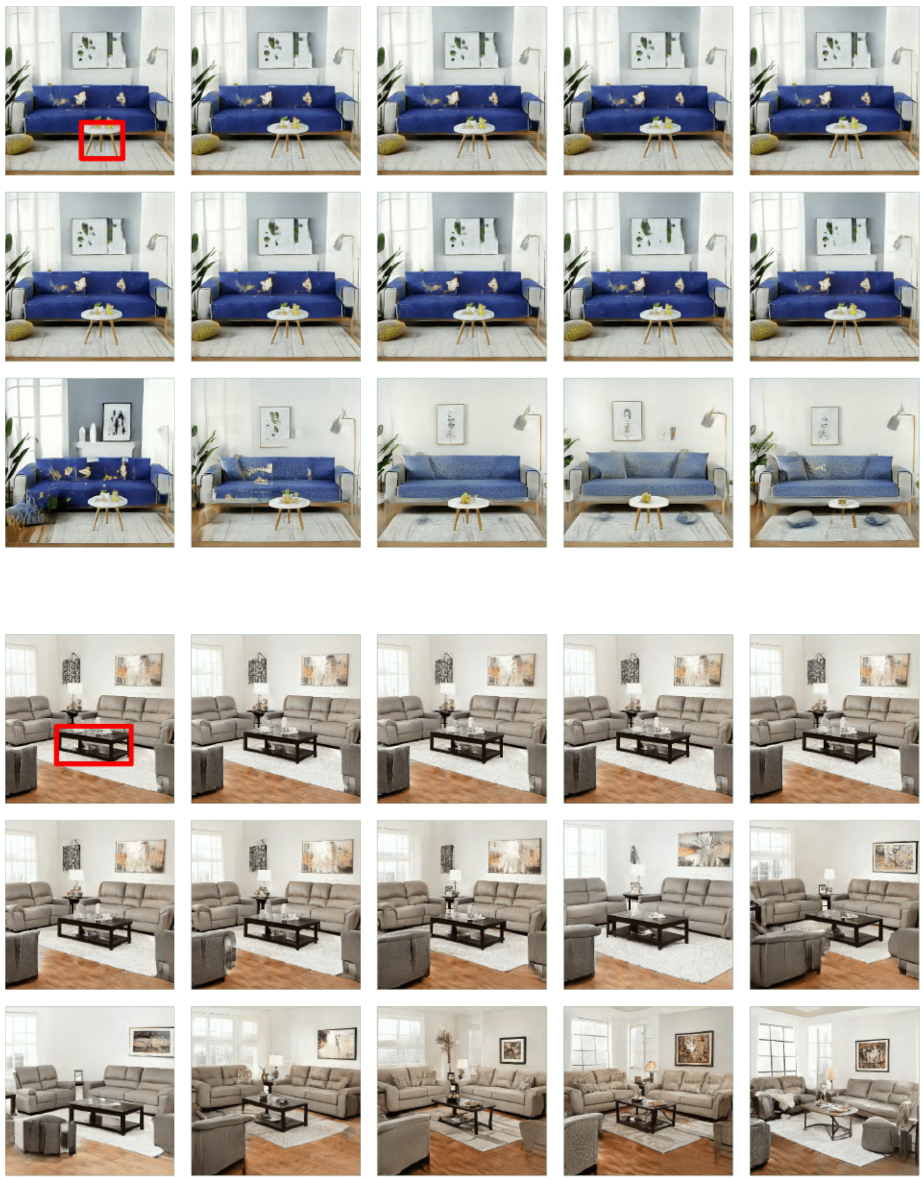

Fig. A4. Two object-preserving living room traversals using REDs and *Projection* traversal. Sequences are in raster-scan order.

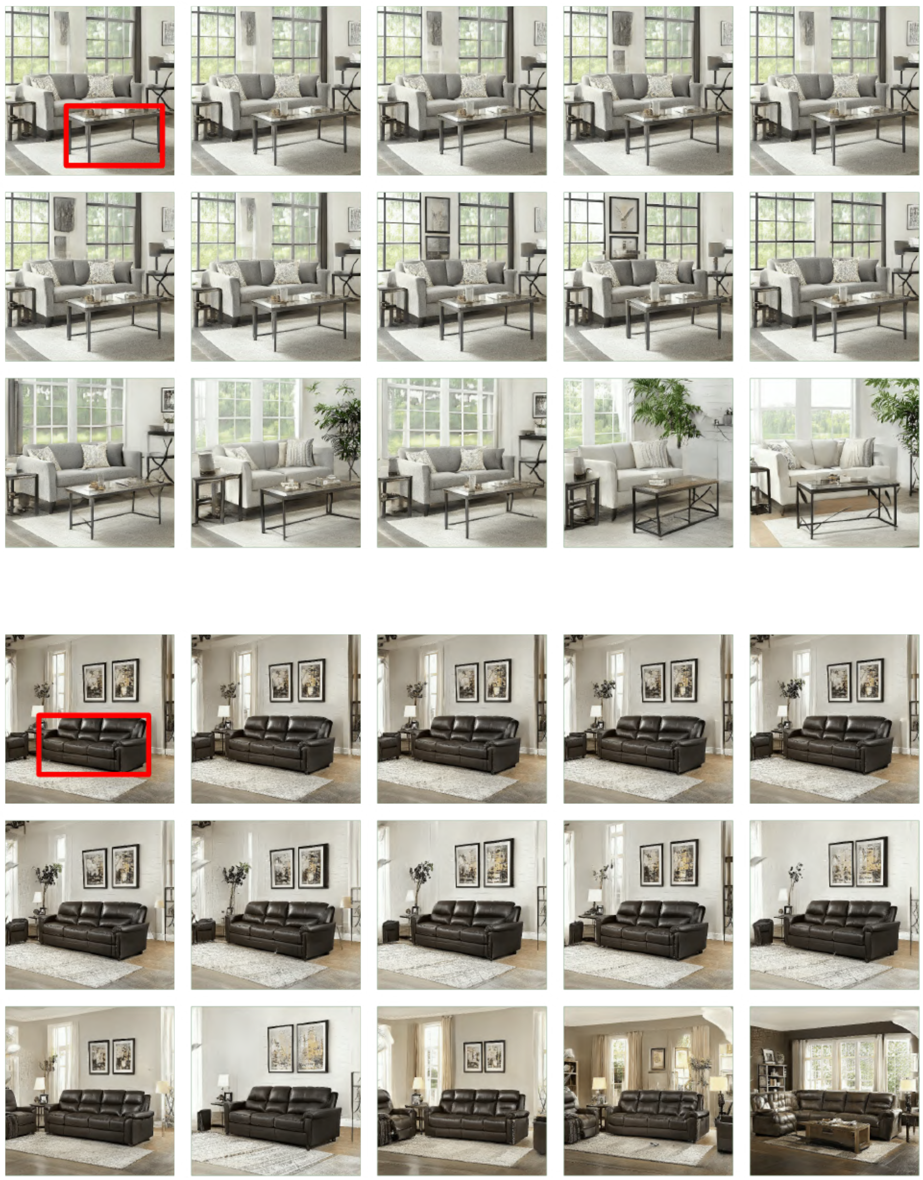

Fig. A5. Object-preserving living room traversals using REDs and *Projection* traversal. Sequences are in raster-scan order.

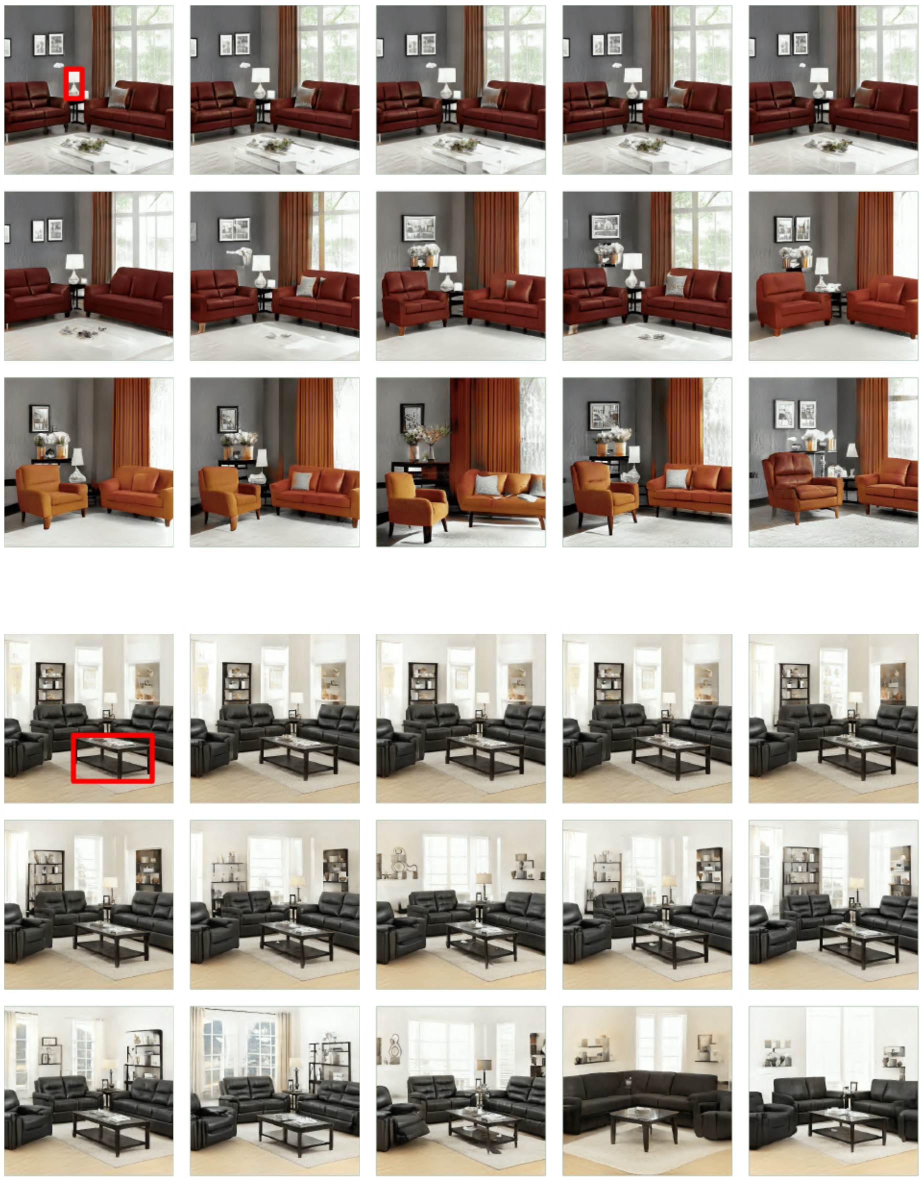

Fig. A6. Object-preserving living room traversals using REDs and *Projection* traversal. Sequences are in raster-scan order.

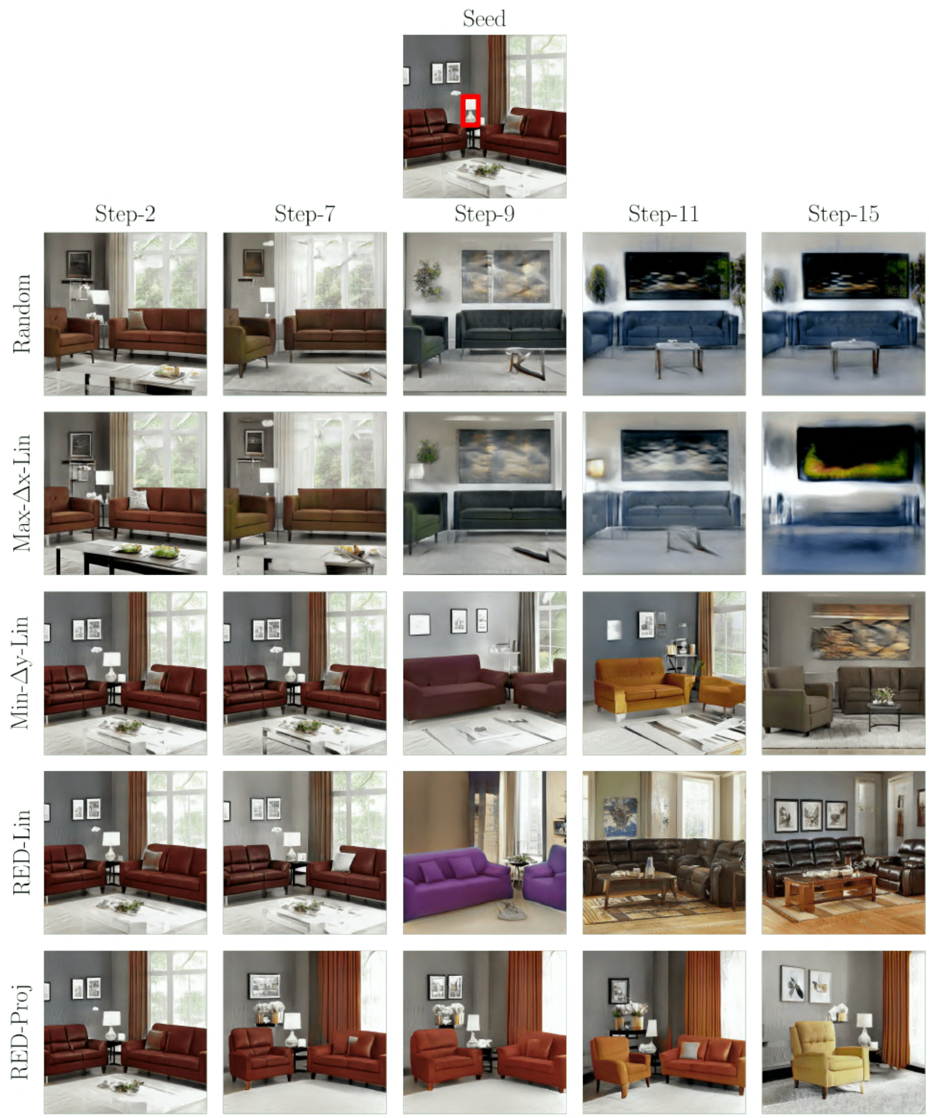

Fig. A7. Comparison to baseline approaches for one seed point. *Linear* traversals (top four rows) deviate significantly from the initial seed image within few steps while the traversal using REDs and *Projection* is consistent with the original lamp (red bounding box) image for a longer distance.

References

1. Wu, Z., Lischinski, D., Shechtman, E.: Stylespace analysis: Disentangled controls for stylegan image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12863–12872 (2021)
2. Zhu, J., Feng, R., Shen, Y., Zhao, D., Zha, Z.J., Zhou, J., Chen, Q.: Low-rank subspaces in gans. Advances in Neural Information Processing Systems **34** (2021)