




NeRF for Outdoor Scene Relighting

– Supplementary Material –

Viktor Rudnev^{1,2}, Mohamed Elgharib¹, William Smith³, Lingjie Liu¹,
Vladislav Golyanik¹, and Christian Theobalt¹

¹MPI for Informatics, SIC ²Saarland University, SIC ³University of York

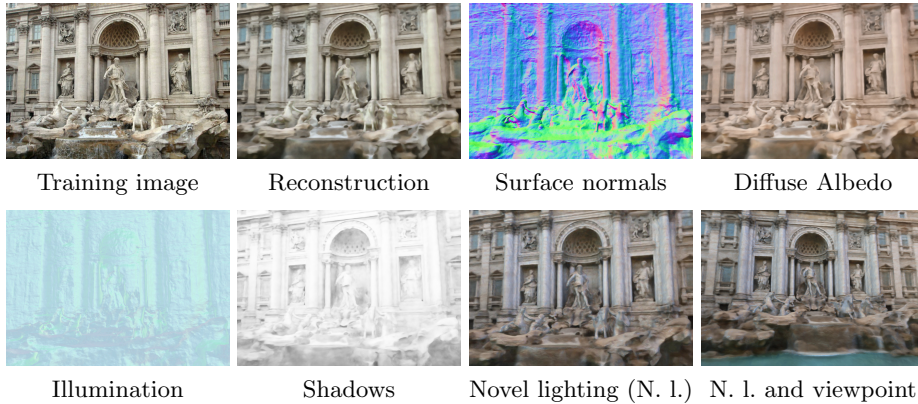


Fig. 1: **NeRF-OSR is the first neural radiance fields approach for outdoor scene relighting.** We learn a neural representation of the scene geometry, diffuse albedo and illumination-dependent shadows from a set of images capturing the same site from different viewpoints and at different times. The learnt intrinsics enable simultaneous editing of both the scene’s lighting and viewpoint.

This supplemental document provides more details on the new dataset for outdoor scene relighting and the experiments.

1 SH Environment Estimation

Recall that our method represents the target illumination through SH coefficients. For experiments in Sec. 5 we estimate the SH coefficients from a 360 environment photo using least squares, inline with Yu et al. [2]. Here, the 360 environment maps allows us to perform numerical evaluations against ground truth (see Sec. 5). We also show in Sec. 5 interactive relighting application where the SH coefficients are directly controlled by the user. Furthermore, the supplemental video (at 3:31-3:43) shows results under novel lighting using externally defined SH coefficients.

2 Statistics of the Dataset

Tab. 1 lists the statistics of our new dataset. It consists of eight sites photographed from 3240 views in 110 different sessions. Our dataset is the first to

	Sessions	Views		Sessions	Views
Site 1	18	373	Site 5	13	493
Site 2	17	423	Site 6	12	379
Site 3	16	372	Site 7	11	468
Site 4	11	401	Site 8	12	331
Total				110	3240

Table 1: Statistics of our new benchmark dataset. The dataset currently contains eight sites recorded in 110 different sessions, each with a 360° environment map captured by LG R105. The total number of views captured by a DSLR camera Canon EOS 550D is 3240.

allow numerical evaluation of relighting methods on real data against ground truth, thanks to the environment maps and the captured colour chequerboards. Please find the download link on our project page: <https://4dqv.mpi-inf.mpg.de/NeRF-OSR/>.

3 Ablative Study

Fig. 3 demonstrates the impact of the design choices in NeRF-OSR on the final novel view renderings with relighting. Not using frequency annealing leads to an evident degradation in the output (the fourth row). This includes circular-shaped artefacts on the ground (the second and the third columns) and clear artefacts on the building (the first three columns, from the left). Removing the shadow regulariser often causes the shadow layer to learn all the illumination components, except the chromaticity, leading to significant artefacts (the first and the last columns). Removing the ray jitter leads to clear artefacts, as shown in the first column. Finally, removing shadow learning and shadow jitter produces less accurate reconstruction (the third column). The strength of shadow learning is more evident during timelapse relighting (see the supplementary video). The full model produces the best results, which is also reflected numerically in Tab. 1 of the main manuscript.

4 Video Results

We demonstrate the ability of NeRF-OSR to edit the camera viewpoint and illumination in our supplementary video. Thus, we show timelapse relighting, where the camera viewpoint is fixed and the illumination changes by rotating the lighting 360° around the building. Our approach handles known lighting conditions and can generalise well to new ones. Even though some synthesised lightings can not occur in real life (*e.g.*, due to the sun trajectory covering only 180° of the sky at most), NeRF-OSR still produces a highly photorealistic output. We also change the viewpoint while keeping the scene illumination fixed. Moreover, we show results when both scene illumination and viewpoint were not

seen during the training. Finally, we visualise the scene intrinsics, *i.e.*, normals, albedo, shadow and shading, as recovered by our technique.

Fig. 4 provides several screenshots from the video results, including simultaneous relighting and novel view synthesis of Site 3 (Fig. 4-(a)). In the figure, we show strong hard shadow synthesis. Next, NeRF-OSR can be used for relighting using unrealistic and synthetic lighting (Fig. 4-(b)). Finally, while NeRF-W [1] can interpolate between the learnt appearances, our method decomposes the scene in its intrinsic components; it enables manipulation of the lighting, which results in interpretable editing of the novel views (Fig. 4-(c)). Also this allows direct editing of albedo, shadows, independently of other intrinsics, real-time interactive VR rendering of the extracted model, which we demonstrate in Secs. 5.4-5.5 of the main manuscript and in the supplementary video. Such shown applications are not possible with style-based techniques by any means as they do not perform intrinsics decomposition.

5 Additional Renderings of Various Sites

Similarly to Fig. 5 of the main document, we show more reconstruction, novel view and lighting synthesis for more sites of our proposed dataset in Fig. 2.

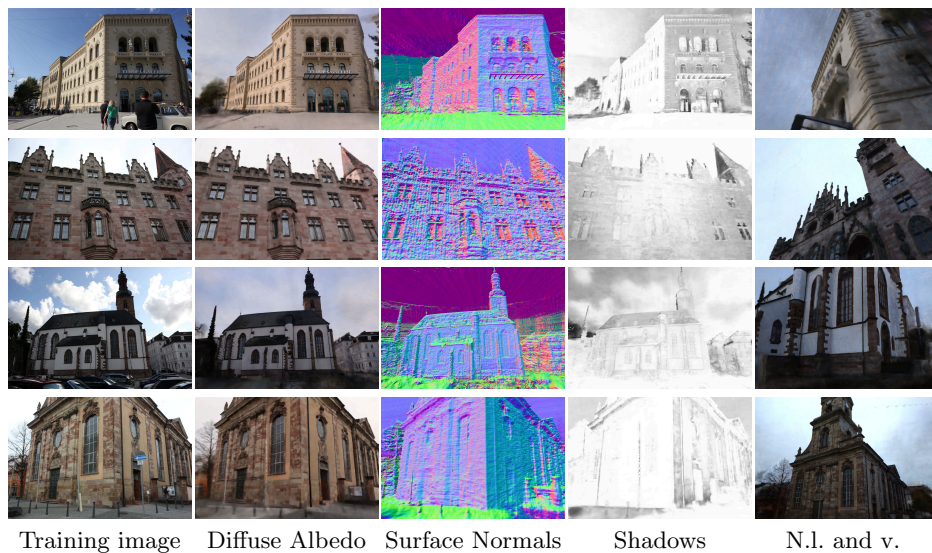


Fig. 2: Additional decompositions and novel lighting and viewpoint synthesis for various sites of our proposed dataset with NeRF-OSR. “N.I. and v.” stands for “Novel lighting and viewpoint”. We kindly ask the readers to ignore inaccuracies in the sky as sky modelling is outside the scope of the work.



Fig. 3: The impact of the various design choices in NeRF-OSR (Site 1). The four columns show the view-lighting combinations used in the quantitative evaluation against the ground truth (Tab. 1 of the main manuscript). The best result is obtained using the full model (the second row from the top). Best viewed with zoom.



Fig. 4: Additional visualisations for various experiments. (a): Relighting and novel view synthesis of Site 3; (b): Relighting of Site 2 using natural and unrealistic light sources (illuminations); (c): Qualitative comparisons to NeRF-W [1]. For the corresponding full videos, see our supplementary video.

References

1. Martin-Brualla, R., Radwan, N., Sajjadi, M.S.M., Barron, J.T., Dosovitskiy, A., Duckworth, D.: NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In: Computer Vision and Pattern Recognition (CVPR) (2021)
2. Yu, Y., Meka, A., Elgharib, M., Seidel, H.P., Theobalt, C., Smith, W.: Self-supervised outdoor scene relighting. In: European Conference on Computer Vision (ECCV) (2020)