



Relighting Neural Radiance Fields With Any Environment Map

Originally published at <https://www.unite.ai/relighting-neural-radiance-fields-with-any-environment-map/>




A [new paper](#) from the Max Planck Institute and MIT has proposed a technique to obtain true [disentanglement](#) of Neural Radiance Fields ([NeRF](#)) content from the lighting that was present when the data was gathered, allowing *ad hoc* environment maps to completely switch out the illumination in a NeRF scene:



The new technique applied to [real data](#). It's noteworthy that the method works even on archived data of this type, which did not take the novel pipeline into account. Source: <https://arxiv.org/pdf/2207.13607.pdf>

The new approach uses the popular open source 3D animation program Blender to create a 'virtual light stage', where numerous iterations of possible lighting scenarios are rendered out and eventually trained into a special layer in the NeRF model that can accommodate any environment map that the user wants to employ to light the scene.



A depiction of the part of the pipeline that leverages Blender to create virtual light stage views of the extracted geometry. Prior methods following similar lines have provided this data, which is a burdensome requirement for discrete objects, and an impossible one for exterior environment views. In the top-left of the right-most two environment maps that dictate the lighting of the scene. These can be arbitrarily created by the end user, bringing NeRF a stage closer to the flexibility of a modern

The approach was tested against the [Mitsuba2](#) inverse rendering framework, and also against prior works [PhySG](#), [RNR](#), [Neural-PIL](#) and [NeRFactor](#), employing only a direct illumination model, and obtained the best scores:

Method	Novel View Synthesis			Novel View Synthesis & Relighting		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
Mitsuba2 [28]	23.50	0.7567	0.0763	21.69	0.5722	0.0812
PhySG [51]	20.52	0.8563	0.2577	17.30	0.6252	0.2736
Neural-PIL [4]	17.07	0.5563	0.1159	14.76	0.4895	0.1328
NeRFactor [54]	21.97	0.6394	0.1691	15.83	0.6470	0.2033
RNR [5]	22.54	0.8122	0.0960	18.06	0.7009	0.1081
Ours	28.73	0.9151	0.0454	23.06	0.8247	0.0692

Results of the new technique, compared against comparable approaches under a variety of loss functions. The researchers claim that their approach yields the highest results, with the results evaluated through Peak Signal-to-noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and the effective if eccentric Learned Image Patch Similarity (LPIPS).

The paper states:

‘Our qualitative and quantitative results demonstrate a clear step forward in terms of the recovery of scene parameters as well as the synthesis quality of our approach under novel views and lighting conditions when comparing to the previous state of the art.’

The researchers state that they will eventually release the code for the project.

The Need for NeRF Editability

This kind of disentanglement has proved a notable challenge for researchers into Neural Radiance Fields, since NeRF is essentially a photogrammetry technique that calculates the pixel value of thousands of possible paths from a viewpoint, assigning RGBD values, and assembling a matrix of these values into a volumetric representation. At its core, NeRF is defined by lighting.

In fact, despite its impressive visuals and lavish adoption by NVIDIA, NeRF is notably ‘rigid’ – in CGI terms, ‘baked’. Therefore the research community has centered on improving its tractability and versatility in this respect over the last 12-18 months.

In terms of significance, the stakes for this kind of milestone are high, and include the possibility of transforming the visual effects industry from a creative and collaborative model centered around mesh generation, motion dynamics and texturing, to a model built around *inverse rendering*, where the VFX pipeline is fueled by real-world photos of real things (or even, conceivably, of real and synthesized models), rather than estimated, artisanal approximations.

For now, there’s relatively little cause for concern among the visual effects community, at least from Neural Radiance Fields. NeRF has only nascent abilities in terms of [rigging](#), [nesting](#), [depth control](#), [articulation](#)...and certainly also in regard to [lighting](#). The [accompanying video](#) for another [new paper](#), which offers rudimentary deformations for NeRF geometry, illustrates the enormous chasm between the current state of the art in CGI and the seminal efforts of neural rendering techniques.


Sifting the Elements

Nonetheless, since it’s necessary to start somewhere, the researchers for the new paper have adopted CGI as an intermediary controlling and production mechanism, by now a [common approach](#) towards the rigid latent spaces of GANs and the almost impenetrable and linear networks of NeRF.

Effectively, the central challenge is to compute [global illumination](#) (GI, which has no direct applicability in neural rendering) into an equivalent [Precomputed Radiance Transfer](#) (PRT, which can be adapted to neural rendering) calculation.

GI is a now-venerable CGI rendering technique that models the way light bounces off surfaces and onto other surfaces, and incorporates these areas of reflected light into a render, for added realism.


Dynamic Diffuse Global Illumination Part I Sample Image



□


PRT is used as an intermediary lighting function in the new approach, and the fact that it’s a discrete and editable component is what achieves the disentanglement. The new method models the material of the NeRF object with a learned PRT.

Local, Deformable Precomputed Radiance Transfer



□

The actual scene illumination of the original data is recovered as an environment map in the process, and scene geometry itself is extracted as a Signed Distance Field (SDF) which will eventually provide a traditional mesh for Blender to operate on in the virtual light stage.




An overview of the pipeline for the new technique.

The first stage in the process is to extract the scene geometry from the available multiple view images through implicit surface reconstruction, via techniques used in the 2021 [NeuS research](#) collaboration.

In order to develop a neural radiance transfer field (NRTF, which will accommodate the illumination data), the researchers used the Mitsuba 2 differentiable path tracer.


Mitsuba Renderer - Getting Started



□

This facilitates the joint optimization of a bidirectional scattering distribution function (BSDF), as well as the generation of an initial environment map. Once the BSDF is created, the path tracer can be used in Blender (see embedded video directly above) to create virtual one-light-at-a-time (OLAT) scene renders.

The NRTF is then trained with a combined loss between photoreal material effects and the synthetic data, which are not entangled with each other.



A comparison with predecessor NeRFactor, on the challenges of novel view synthesis and relighting.

The Road to Illumination

The training requirements for this technique, though notably lesser than the original NeRF training times, are not insignificant. On a [NVIDIA Quadro RTX 8000](#) with 48GB of VRAM, preliminary training for initial light and texture estimation takes 30 minutes; OLAT training (i.e. the training of the virtual light stage captures) takes eight hours; and the final joint optimization between the disentangled synthetic and real data takes a further 16 hours to reach optimal quality.

Further, the resulting neural representation cannot run in real time, taking, according to the researchers 'several seconds per frame'.

The researchers conclude:

'Our results demonstrate a clear improvement over the current state of the art while future work could involve further improving the runtime and a joint reasoning of geometry, material, and scene lighting.'

First published 28th July 2022.

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



Discover More