

Neural Global Illumination

George Drettakis, GRAPHDECO, Inria Sophia Antipolis (France)

<http://team.inria.fr/graphdeco>

George.Drettakis@inria.fr

<http://www-sop.inria.fr/members/George.Drettakis>




Figure 1: (1st row) Our Neural Pre-computed Radiance method [Rainer 22] encodes light transport in an MLP, allowing interactive rendering (right). (2nd row) Our Active Exploration [Diolatzis 22] approach allows fast training of an MLP that can represent light transport in scenes with variations such as moving objects, lights or changing materials; here the door can open interactively

Context and goal

Neural rendering [Tewari 20] is an emerging topic in Computer Graphics, where neural networks are used to represent scene geometry (e.g., via opacity or a signed distance function [Park 19]) and the outgoing radiance in a given viewing direction [Mildenhall 20, Barron 21]. Many of these methods focus on rendering data captured from photographs, and rely on coordinate based

multi-layer perceptrons (CB-MLPs) for compact and expressive representation. Recently, similar techniques have been proposed to compactly precompute expensive global illumination in *synthetic* scenes [Hadadan 22], including our work on Neural Precomputed Radiance Transfer [Rainer 22] (Fig 1 a) and an Active Exploration [Diolatzis 22] method to accelerate training of such networks (Fig 2 b). In parallel, real-time ray-tracing hardware has become commonplace [Benty 20], and neural networks have been used to approximate indirect illumination with on-the-fly training [Müller 21]. These advances show an extremely promising future for neural rendering of synthetic scenes, in combination with real-time ray-tracing hardware; however several significant challenges exist. First, for pre-computation methods, training timing is still high (in the tens of hours for complex light paths and scene variations); second, training is still per-scene for these methods, imposing a significant overhead for the use of such approaches. In this thesis we will develop solutions to both of these problems, and investigate ways to develop neural rendering solutions that work closely with real-time ray-tracing hardware to provide the best of both worlds.

Approach

To address the first challenge, we will first investigate hybrid solutions that integrate real-time ray-tracing (RTRT) hardware renderers in the training pipeline, automatically identifying hard global illumination paths allowing the neural network to focus on learning only these, while RTRT can be used to handle the rest. We will also investigate how to develop methods that allow fast initial training and can then be fine-tuned on-the-fly during rendering, exploiting low-capacity MLPs and appropriately developed encodings in the spirit of [Müller 22]. For the second challenge, there has been exciting recent work that shows how to separate CB-MLP methods into a scene-dependent and a generalizable component [Kania 21] as well as attempts to generalize across scenes with Transformers [Wang 21]. We will investigate ways to exploit fast training methodologies for CB-MLPs to adapt these methods to the learning needs of Neural Rendering for global illumination.

Work environment and requirement

Candidates should have a Masters in Computer Graphics and/or Computer Vision, but with knowledge in both. Strong programming and mathematical skills are required as well as knowledge in computer graphics, geometry processing and machine learning, with experience in C++, OpenGL and GLSL on the graphics side, and pytorch for deep learning.

References

[Mildenhall 20] Mildenhall et al. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020. <https://www.matthewtancik.com/nerf>

[Tewari 20] Tewari et al. State of the art on neural rendering. Computer Graphics Forum, 2020.

[Barron 21] Barron, J. T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., & Srinivasan, P. P. (2021). Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. arXiv preprint arXiv:2103.13415.

[Diolatzis 22] Diolatzis, Stavros, Julien Philip, and George Drettakis. "Active Exploration for Neural Global Illumination of Variable Scenes." *ACM Transactions on Graphics (TOG)* (2022).

<https://repo-sam.inria.fr/fungraph/active-exploration/>

[Rainer 22] Rainer, G., Bousseau, A., Ritschel, T., & Drettakis, G. (2022, April). Neural Precomputed Radiance Transfer. In *Computer Graphics Forum*. <http://www.sop.inria.fr/revs/Basilic/2022/RBRD22/>

[Hadadan 21] Hadadan, Saeed, Shuhong Chen, and Matthias Zwicker. "Neural radiosity." *ACM Transactions on Graphics (TOG)* 40, no. 6 (2021): 1-11.

[Benty 20] Nir Benty, Kai-Hwa Yao, Petrik Clarberg, Lucy Chen, Simon Kallweit, Tim Foley, Matthew Oakes, Conor Lavelle, and Chris Wyman. 2020. The Falcor rendering framework.

<https://github.com/NVIDIAGameWorks/Falcor>

[Müller 21] T. Müller, T., Rousselle, F., Novák, J., & Keller, A. (2021). Real-time neural radiance caching for path tracing. *ACM Transactions on Graphics (TOG)*, 40(4), 1-16

[Müller 22] Müller, T., Evans, A., Schied, C. and Keller, A., 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Transactions on Graphics (TOG)*, 41(4) 2022

[Wang 21] Wang, Q., Wang, Z., Genova, K., Srinivasan, P. P., Zhou, H., Barron, J. T., ... & Funkhouser, T. (2021). Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 4690-4699).

[Kania 21] Kania, K., Yi, K.M., Kowalski, M., Trzeciński, T. and Tagliasacchi, A., 2021. CoNeRF: Controllable Neural Radiance Fields. *arXiv preprint arXiv:2112.01983*.

[Park 19] Park, J. J., Florence, P., Straub, J., Newcombe, R., & Lovegrove, S. (2019). DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 165-174).