Towards Error Bounds for Neural Rendering

(Masters or last year Engineering internship)

George Drettakis, GRAPHDECO, Inria Sophia Antipolis (France)
http://team.inria.fr/graphdeco

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

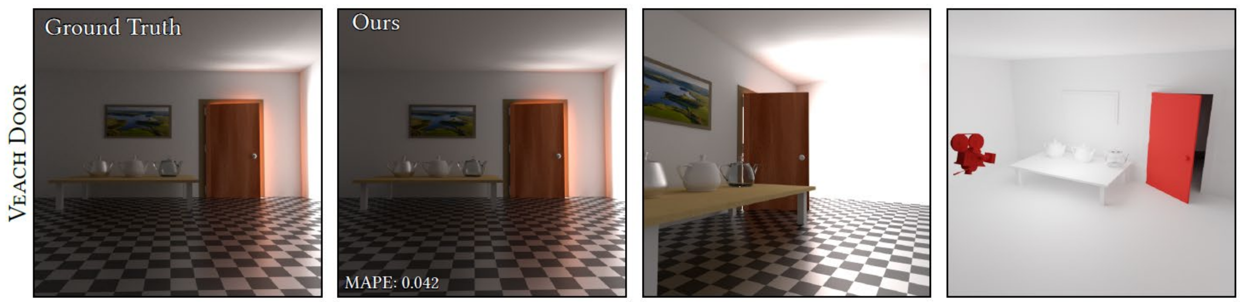

Figure 1: Our Active Exploration Neural Rendering approach [Diolatzis22] allows high quality image synthesis at interactive rates in scenes with moving camera and geometry (red on the far right); however there are no guarantees on the resulting error.

Context and goal

Neural rendering [Diolatzis22, Rainer22] has recently been demonstrated as a very competitive alternative to path tracing [Pharr16] for rendering on synthetic content (see Fig. 1). Using a smart training strategy where ground truth images are generated on the fly and interleaved with neural network training, we can train a relatively lightweight multi-layer perceptron indexed by 3D coordinates to predict global illumination for scenes with variations, e.g., moving light sources and objects, and of course moving viewpoints [Diolatzis22].

However, a major concern of this approach is that neural renderings have essentially no guarantees on error, and in particular -- in their current form -- are neither unbiased (i.e., are guaranteed to converge to the correct solution) or even consistent (i.e., the error diminishes in a predictable way with further computation, even though there is no formal guarantee to correctness). In this internship we will investigate solutions that move towards the directions of converting neural renderings to unbiased or consistent images, with a certain additional computational cost.

Approach

As a first step, we will investigate two variance/bias reduction strategies. The first will involve accumulating bias and variance data during the training phase of the neural network in a 3D spatial data structure. This is of course possible, since we constantly compute ground truth values at different points in space. The second will involve training the neural network to predict bias together with radiance in 2D -- again this achievable for the same reason. During inference, the neural rendering will produce a biased image; if an unbiased solution is required, an additional path-tracing step will be applied, but restricted to regions in image and 3D space that are identified as having high variance or bias.

Evidently, this approach still does not provide guarantees on error, although it does reduce bias in the result. In a second step, we will study the structure of the network weights, and the nature of the gradients during training to determine if there is a way to bound bias for a given trained scene. Given such bounds, it will also be possible to at least develop a consistent solution, i.e., a solution in which error decreases with more samples. We also expect that these bounds will improve the training strategy further, compared to the approach of Diolatzis et al [Diolatzis22]. We will also study more traditional techniques for "unbiasing" (e.g., [Misso22]) and determine if and how they can be applied in the context of training using stochastic gradient descent, as well as more learning-related methods that automatically invert neural networks that seek inverse solutions in the vicinity of reliable samples from path tracing that have been used for training the neural renderer [Ansari22].

Work environment and requirement

The internship will take place at Inria Sophia Antipolis in the GRAPHDECO group (http://team.inria.fr/graphdeco). Inria will provide a monthly stipend of around 1100 euros for EU citizens in their final year of masters, and 400 euros for other candidates.

Candidates should have strong programming and mathematical skills as well as knowledge in computer graphics, geometry processing and machine learning, with experience in C++, patht-tracing (knowledge of mitsuba1/2/3 is a plus), OpenGL and GLSL on the graphics side, and pytorch for deep learning.

References

[Pharr16] Pharr, M., Jakob, W. and Humphreys, G., 2016. *Physically based rendering: From theory to implementation*. Morgan Kaufmann. http://www.pbrt.org

[Diolatzis22] 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/

[Misso22] Misso, Zackary, Benedikt Bitterli, Iliyan Georgiev, and Wojciech Jarosz. "Unbiased and consistent rendering using biased estimators." *ACM Transactions on Graphics (TOG)* 41, no. 4 (2022): 1-13.

[Rainer22] Rainer, Gilles, Adrien Bousseau, Tobias Ritschel, and George Drettakis. "Neural Precomputed Radiance Transfer." In Computer Graphics Forum, vol. 41, no. 2, pp. 365-378. 2022. https://repo-sam.inria.fr/fungraph/neural-prt/

[Ansari22] Ansari, N., Seidel, H.P., Ferdowsi, N.V. and Babaei, V., 2022. Autoinverse: Uncertainty Aware Inversion of Neural Networks. arXiv preprint arXiv:2208.13780. https://arxiv.org/pdf/2208.13780