Surfel-Based Neural-Inspired Global Illumination

(Masters 2 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 is no way to control the level of the resulting error. We will develop an approximate global illumination algorithm that exploits neural rendering methodologies but is amenable to error control.

Context and goal

Neural rendering [Diolatzis22, Rainer22] has recently been demonstrated as a very competitive alternative to path tracing [Pharr16]. Using a smart training strategy where ground truth images are generated on the fly, 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 such approaches 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 (the error diminishes with further computation, even though there is no formal guarantee to correctness). In this internship we will design a new neural rendering algorithm, inspired by both older previous work (e.g., photon mapping [Hachisuka09], point-based global illumination [Christensen08]) and modern radiance field optimization techniques [Kerbl23]. The main goal of this new algorithm is that it has a parameter *k* such that for greater values of the parameter the error is guaranteed to decrease.

Approach

Photon-mapping trivially satisfies the requirement of the error parameter mentioned above [Misso22]. Inspired by this fact, we will design a new algorithm that uses primitives similar to photons (i.e., receptors of incoming radiance attached to surfaces in the scene) to represent incoming light. We will use surfel-like primitives that can be rendered extremely fast using modern GPUs [Kerbl23], allowing photon-map like splatting to be used during optimization for rendering. Such primitives may also include a directional component that will also be optimized during the training phase representing the light arriving at a given surfel. The directional component may however be built into a hierarchical structure, inspired by point-based global illumination [Christensen08] and/or include neural rendering components.

We will design an optimization strategy for learning the incoming light similar to those used in [Diolatzis22,Rainer22]. The optimization strategy will combine predetermined sampling strategies for the samples, with on-the-fly densification approaches.

Once optimized, the method will provide a parameter *k* similar to photon-mapping (i.e., the footprint of the surfels), satisfying the properties required to apply the debiasing strategy of [Misso23], thus allowing a learned representation to provide unbiased results, without the expense of full path tracing, e.g., for a moving viewpoint.

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 be passionate about computer graphics and neural rendering methods, and preferably have strong programming and mathematical skills as well as knowledge in computer graphics, geometry processing and machine learning, with experience in C++, real-time rendering techniques, path-tracing (knowledge of mitsuba1/2/3 is a plus), OpenGL and GLSL on the graphics side. Some knowledge of machine learning is a plus.

References

[Christensen08] Christensen, Per. "Point-based approximate color bleeding." *Pixar Technical Notes* 2.5 (2008): 6.

[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/

[Hachisuka09] Hachisuka, Toshiya, and Henrik Wann Jensen. "Stochastic progressive photon mapping." In *ACM SIGGRAPH Asia 2009 papers*, pp. 1-8. 2009.

[Kerbl23] Kerbl, Bernhard, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. "3d gaussian splatting for real-time radiance field rendering." *ACM Transactions on Graphics (ToG)* 42, no. 4 (2023): 1-14. https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/

[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.

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

[Philip21] Philip, Julien, Sébastien Morgenthaler, Michaël Gharbi, and George Drettakis. "Free-viewpoint indoor neural relighting from multi-view stereo." *ACM Transactions on Graphics (TOG)* 40, no. 5 (2021): 1-18.