Efficient Neural Rendering:
Getting the Best from Points and Volumes

(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/

(a) NeRF [3] *(b) Point-based Neural Rendering[1]*

Figure 1: (a) Neural Radiance Fields – NeRF [3] achieve impressive results in neural rendering, but suffer from artifacts such as blurring in vegation (top) or background (bottom). (b) Point-based solutions [1] can resolve blurry artifacts in vegetation (top) and background, but can create other artifacts with spurious 3D reconstruction (bottom, red arrows).

Context and goal

Recently, neural rendering [4] has seen an explosion in novel research results [5,6,8,9], many of which are centered around volumetric approaches where a Multi-layer Perceptron (MLP) is used to encode the radiance field in a scene, aka NeRF [3] Fig. 1(a). These representations have many advantages, notably the fact that they are naturally differentiable, simplifying training to reconstruct 3D geometry from photos. They can also handle (semi-)transparent objects, deal relatively well with reflections and moving highlights. However, in most scenes where these methods have been demonstrated, the content is almost entirely opaque surfaces. Consequently, the expensive volumetric ray-marching approach used for view synthesis seems conceptually wasteful.

In our recent work, we have shown that point-based representations [1,2] can actually perform better in some scenes in terms of visual quality, while maintaining some, but not all of the advantages of volumetric NeRFs (Fig. 1(b)). An important benefit of point-based representations is that they are naturally suited to the display of opaque objects (as opposed to volumetric ray-marching), and amenable to fast rendering, completely compatible with the traditional graphics pipeline.

In this internship, we will develop novel solutions that exploit the advantages of both representations, leading to an algorithm that will on the one hand maintain the power and flexibility of NeRF to capture high-quality geometry and appearance representations of real objects during training, but will create a rendering-efficient, graphics-compatible point-based representation for rendering. The intern will work will work in collaboration with G. Kopanas, Ph.D. student in the group and first author of [1].

Approach

The internship will start with an analysis of the training process of NeRF, and in particular NeRF++ [5] for the kind of real-world outdoors scenes we wish to treat (see [7] and Fig. 1). The group has already performed extensive experiments for other projects, and we have identified parts of the process that we can be completely redesigned to avoid the costly step of volumetric ray-marching, replacing a significant part of the training process by operations on points. These points will be enhanced with features, in a manner analogous to [1]. We envisage a progressive algorithm that starts in a manner similar to standard NeRF to maintain all the advantages in terms of high-quality geometry reconstruction in difficult cases such as thin structures etc., and progressively transitions the representation to a point-based solution that only maintains volumetric components when necessary.

The development of this algorithm faces many challenges. The first important challenge is to define the theoretical framework that will ensure that the progressive algorithm maintains the correct solution while using the dual representation. This will involved rethinking the basic image synthesis equations of both volumetric and point-based approaches. The second challenge is the definition of principled metrics that will allow the algorithm to transition from volumetric to point-based without loss of rendering quality. Finally, the development of a mixed rendering algorithm that will seamlessly and efficient support volumetric and point-based rendering is a major challenge.

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++, OpenGL and GLSL on the graphics side, and pytorch for deep learning.

References

[1] Georgios Kopanas, Julien Philip, Thomas Leimkühler, George Drettakis, Point-Based Neural Rendering with Per-View Optimization, *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering), Volume 40, Number 4*, June 2021 https://repo-sam.inria.fr/fungraph/differentiable-multi-view/
[2] Aliev et al. Neural point-based graphics. In ECCV, 2020. https://saic-violet.github.io/npbg/
[3] Mildenhall et al. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020. https://www.matthewtancik.com/nerf
[4] Tewari et al. State of the art on neural rendering. Computer Graphics Forum, 2020.
[5] Zhang, K., Riegler, G., Snavely, N., & Koltun, V. (2020). Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492.
[6] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, Angjoo Kanazawa, "PlenOctrees: for Real-time Rendering of Neural Radiance Fields", ICCV 2021
[7] Knapitsch, Arno, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. "Tanks and temples: Benchmarking large-scale scene reconstruction." ACM Transactions on Graphics (ToG) 36, no. 4 (2017): 1-13.
[8] 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.\
[9] Reiser, C., Peng, S., Liao, Y., & Geiger, A. (2021). KiloNeRF: Speeding up Neural Radiance Fields with Thousands of Tiny MLPs. arXiv preprint arXiv:2103.13744.