

Unifying Neural Representations for Efficient Rendering

George Drettakis, GRAPHDECO, Inria Sophia Antipolis (France)

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

George.Drettakis@inria.fr

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



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 vegetation (top) or background (bottom). (b) Point-based solutions [1] can resolve blurry artifacts in vegetation (top) and background.

Context and goal

Recently, neural rendering [4] has seen an explosion in novel research results [5,6,8,9]. All recent methods share the same central property: the scene representation is differentiable, allowing optimization from input images to reconstruct the neural representation of the scene. Three interesting cases are Neural Radiance Fields (NeRF) [3], Signed Distance Functions (SDFs) [10] and Point-Based Representations [2] (including our recent work [1]). Each has relative advantages and disadvantages, and the choice of the best fitting representation often depends on the capture conditions and the use case. For example, the continuous nature of the NeRF representation has many advantages, but often results in blurry results (Fig. 1, left), while point-based representations are a discrete sampling, but can provide sharper renderings (Fig. 1, right, using our point-based solution [1]). Ideally, we would like to have a unified representation

that allows efficient rendering of captured scenes combining the advantages of all previous piecemeal representations. In this Ph.D. we will investigate the qualities of the different representations in the context of efficient rendering, and develop such a new, unified representation that will be powerful and general enough to achieve this goal.

Approach

We will start by developing a unified model for NeRF and point-based representation, by building on the observation that they share similarities in terms of the actual “alpha-blended” rendering. This will involve an in-depth theoretical analysis of the properties of each representation, and proposing a new model that is theoretically equivalent for both. The development of efficient rendering algorithms will follow, by exploiting the capabilities of modern GPUs that have both deep learning and ray-tracing support. We will continue with the harder case of SDFs, which have the advantage of defining an actual surface which can be particularly efficient in many cases for rendering. We first need to determine a new scene representation model that will allow the incorporation of SDF into a unified model so that the positive aspects of all three representations can be exploited while avoiding the pitfalls. Efficient rendering algorithms will then be developed, again exploiting hardware advances.

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

- [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/npg/>
- [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*.

[10] Sharp, Nicholas, and Alec Jacobson. "Spelunking the Deep: Guaranteed Queries for General Neural Implicit Surfaces." *ACM Transactions on Graphics (TOG)*, 41(4) 2022