Studying and Improving Neural Radiance Field 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/

(a) NeRF [Mildenhall2] *(b) Point-based [Kopanas21]* *(c) Signed Dist. Field [Vicini22]*

Figure 1: Neural Radiance Fields have seen immense growth in recent in the last few years. However, there are several different ways to render: volumetric rendering with ray-marchine (a), point-based rendering (b) but also signed distance fields (c) (typically restricted to objects) and the different tradeoffs have not been study systematically to determine the best solution for different contexts.

Context and goal

Neural Rendering research has seen an explosion in recent years, especially for Neural Radiance Fields (NeRFs [Mildenhall21]), based on volumetric rendering using ray-marching of a density function. Nonetheless, there are several other possibilities for rendering Neural Radiance Fields [Xie22] namely Signed Distance Fields [Jian20, Vicini22], and Point-Based solutions [Kopanas21,Kopanas22]. The rendering quality achievable by each representation has advanced enormously in an impressively small period of time. However, there is no clear methodology to predict which method is most appropriate for different use cases, depending on the type of scene or the speed/quality tradeoff for rendering.

The goal of this internship is to develop such a methodology, and in addition use the resulting methodology to develop new solutions for the different kinds of neural radiance field rendering approaches, by combining the strengths of each approach in novel ways while avoiding their pitfalls.

Approach

The are several fundamental differences and limitations in the aforementioned rendering approaches: the distinction between continuous (SDF, NeRF) and sampled (Points) representations, the definition of an explicit surface position (SDF, points) vs a "soft density" approach to geometry (NeRF), and the difference between Langrangian (Points) and Eulerian (NeRF) representations of view-dependent effects. We will investigate these differences to develop a well-founded methodology for choosing the best rendering algorithm for a given use case. In the process of this investigation, we will use the methodology to improve the individual algorithm building on the understanding achieved.

We will first investigate the tradeoffs of different rendering methods depending on scene content, and in particular for the typically hard-to-reconstruct cases of thin structures (vegetation) and reflective/transparent surfaces. We will determine whether specific characteristics of scene content (level of shininess or transparency) affect the choices made and develop a method to identify and categorize these tradeoffs. Using these solutions, we will investigate ways to improve the quality of rendering, e.g., for transparent surfaces and flat surfaces that are suboptimal for the state-of-the-art point-based rendering method [Kopanas22].

Second, we will investigate rendering speed vs quality tradeoffs for these algorithms. Volumetric ray-marching, sphere tracing and alpha-blended point splatting have many similarities, but also several fundamental differences in terms of the ability to fully exploit the capabilities of GPUs. We will systematically categorize these, starting with controlled environments (e.g., scenes with reflections, thin structures etc.) and investigate ways to develop algorithms that can seamlessly and efficiently switch between modes, or even better use the best solution for different parts of a scene if appropriate.

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

[Mildenhall21] Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2021). Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM, 65(1), 99-106.

[Xie22] Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V. and Sridhar, S., 2022, May. Neural fields in visual computing and beyond. In Computer Graphics Forum (Vol. 41, No. 2, pp. 641-676).

[Jian20] Jiang, Yue, Dantong Ji, Zhizhong Han, and Matthias Zwicker. "Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization." In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 1251-1261. 2020.

[Kopanas21] Kopanas, G., Philip, J., Leimkühler, T., & Drettakis, G. (2021, July). Point‐Based Neural Rendering with Per‐View Optimization. In Computer Graphics Forum (Vol. 40, No. 4, pp. 29-43).

[Kopanas22] Kopanas, G., Leimkühler, T., Rainer G., Jambon C., & Drettakis, G. (2021, July). Neural Point Catacaustics for Novel-View Synthesis of Reflections. In ACM Transactions on Graphics (SIGGRAPH Asia 2022)

[Vicini22] Vicini, Delio, Sébastien Speierer, and Wenzel Jakob. "Differentiable signed distance function rendering." *ACM Transactions on Graphics (TOG)* 41, no. 4 (2022): 1-18.