Investigating Hybrid Representations for Radiance Fields

(Masters 2 or Pre-Doc internship, possible extension to a Ph.D)

George Drettakis, Linus Franke and Simon Lucas, GRAPHDECO, Inria Sophia Antipolis (France)

http://team.inria.fr/graphdeco

George.Drettakis@inria.fr, http://www-sop.inria.fr/members/George.Drettakis/
Linus.Franke@inria.fr, https://lfranke.github.io
simon.lucas@inria.fr, https://simon-lucas.fr

Figure 1: An illustration of a hybrid representation based both on a volumetric representation (left) and explicit triangles (right).

Context and goal

Recent advances in neural rendering, such as Neural Radiance Fields (NeRFs) [Mildenhall et al. 2020] and 3D Gaussian Splatting (3DGS, developed at GRAPHDECO) [Kerbl et al. 2023], have

shown the effectiveness of volumetric formulations for multi-view reconstruction and novel view synthesis. These methods benefit from differentiable volume rendering, which enables stable optimization and smooth gradient propagation across views.

However, this simplification oversmooths the result and is not physically accurate, as few objects in the real world are truly volumetric. Recent works explore bridging volumetric and surface representations: SuGaR [Guedon and Lepetit 2024] aligns Gaussians with surface geometry before extracting meshes, while MILo [Guedon et al. 2025] (developed jointly with our group) directly optimizes a mesh during volumetric reconstruction.

These approaches highlight an open challenge: how to achieve high-fidelity, explicit geometry while retaining the differentiability and rendering quality of volumetric methods.

Approach

This project proposes a hybrid reconstruction pipeline that augments Gaussian Splatting with explicit surface representations such as triangle meshes. We will investigate several different options to address the challenge above.

One possibility would begin with a standard volumetric reconstruction using 3DGS, exploiting its stable optimization for coarse geometry and appearance modeling. Once local regions reach sufficient convergence, surface alignment regularization could be introduced to encourage Gaussians to concentrate along well-defined surfaces, guided by normals, depth cues, or implicit SDF priors. From these aligned splats (or other suitable representation), meshes could be extracted and integrated back into the pipeline, yielding a hybrid representation in which explicit triangles replacing Gaussians for opaque surfaces, while Gaussians remain for semi-transparent or uncertain regions.

Differentiable rendering losses can be used to combine volumetric splatting and mesh rasterization to jointly refine both components. By fusing the strengths of volumetric optimization with explicit surface fidelity, the project aims to produce reconstructions that achieve both high-quality novel view synthesis and accurate, mesh-based scene geometry suitable for downstream applications.

Work environment and requirements

The internship will take place at Inria Sophia Antipolis in the GRAPHDECO group (http://team.inria.fr/graphdeco) (the inventors of 3D Gaussian Splatting).

Candidates should be passionate about computer graphics and neural rendering methods, and have strong programming and mathematical skills. Knowledge in one or more of computer graphics, geometry processing and machine learning, experience in python, pytorch, cuda, C++,

real-time rendering techniques, path-tracing (knowledge of mitsuba3 is a plus), OpenGL and GLSL on the graphics side are desirable.

How to apply

Applicants should either be Masters (5th year) students for an internship, or if applying for a pre-doc, they should already have a Masters degree in Computer Science, specialized in Computer Graphics and/or Computer Vision. Please email George.Drettakis@inria.fr with your CV, motivation letter and your transcripts for the last 2-3 years of study.

References

[Guédon et al. 2025] Guédon, A., Gomez, D., Maruani, N., Gong, B., Drettakis, G., & Ovsjanikov, M. (2025). MILo: Mesh-In-the-Loop Gaussian Splatting for Detailed and Efficient Surface Reconstruction. ACM Transactions on Graphics. https://anttwo.github.io/milo

[Guédon and Lepeit 2025] Guédon, A., & Lepetit, V. (2024). SuGaR: Surface-Aligned Gaussian Splatting for Efficient 3D Mesh Reconstruction and High-Quality Mesh Rendering. CVPR .

[Kerbl et al. 2023] 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/

[Mildenhall et al. 2020] Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2020). NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. ECCV .