Synthetic Data for Diffusion-Based Relighting

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

George Drettakis 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/

simon.lucas@inria.fr, https://simon-lucas.fr

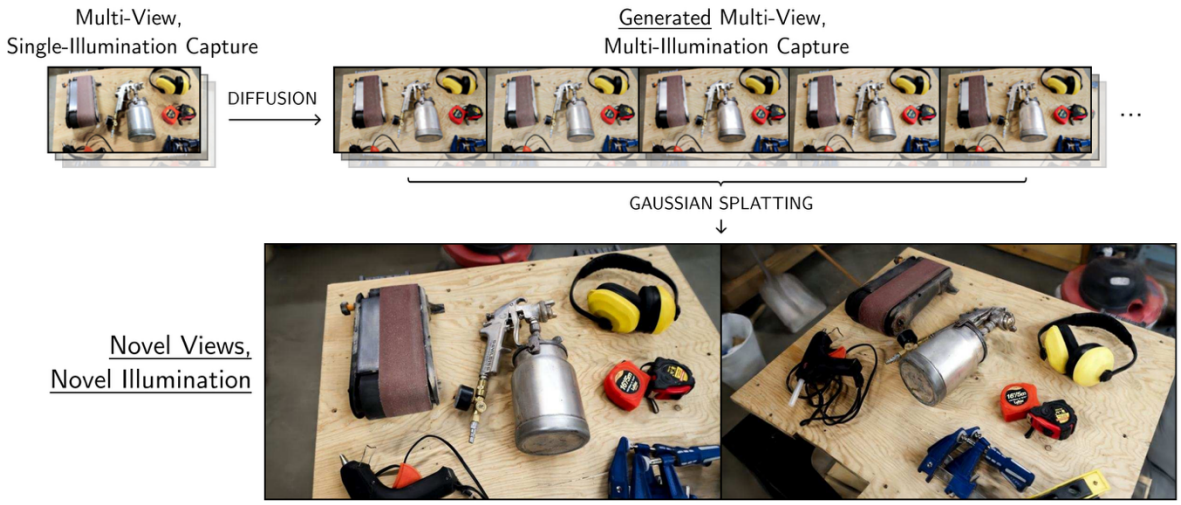

Figure 1: Using diffusion-based priors can be very effective for relighting radiance fields. However, lack of data limits the capabilities of the method. (from https://repo-sam.inria.fr/fungraph/generative-radiance-field-relighting/)

Context and goal

Capturing real scenes so they can be displayed in 3D using Computer Graphics techniques is a domain that has seen great growth recently, with Neural Radiance Fields (NeRFs [Mildenhall et al. 2020] and 3D Gaussian Splatting (3DGS) [Kerbl et al. 2023] (developed at GRAPHDECO). Most often however, it is not possible to change the lighting of the scenes. There has been extensive work on relighting *objects* that are captured with the assumption of distance lighting [REFS] but there is only little work on relighting *full scenes.* One method that addresses the latter has been developed at GRAPHDECO (in collaboration with U. Laval) [Poirier-Ginter et al. 24]. In

this method, a diffusion model was used via ControlNet to learn to relight images using a dataset of real scenes captured under different illumination [Murmann et al 19]. However, this dataset is very limited, resulting in convincing, but very constrained relighting (see Fig. 1). The success of diffusion-model based relighting using synthetic data mixed with real data for single images [Magar et al. 25] is a promising indication that such an approach could be applied to the method of Poirier-Ginter et al.

In this internship, we will investigate how to use synthetic dataset to overcome this limitation.

Approach

The project will have two main components. The first is the definition of the appropriate parameterization of lights and lighting in the scene, and the second is the actual creation and augmentation of the synthetic dataset that will be used to relight scenes.

For the first, several approaches are possible. A first possibility will be the use of point (or small spherical) lights, distributed in “reasonable places” in the scene; if this is promising, this could be extended to more complex light sources. An accompanying encoding will be developed to allow conditioning of the ControlNet for appropriate lighting.

For the second, GRAPHDECO has already developed significant infrastructure using Blender and either cycles or Mitsuba as renderers. We will build on this infrastructure and develop data augmentation techniques inspired by those used in LightLab [Magar et al. 25] to achieve the best possible results.

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). The intern will have access to local, Inria-wide and national GPU compute infrastructure.

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

[Kerbl et al. 23] 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 .

[Poirier-Ginter et al. 24] Poirier-Ginter, Y., Gauthier, A., Phillip, J., Lalonde, J.F. and Drettakis, G., 2024, July. A Diffusion Approach to Radiance Field Relighting using Multi-Illumination Synthesis. In Computer Graphics Forum (Vol. 43, No. 4, p. e15147).

[Murmann et al 19] Murmann L, Gharbi M, Aittala M, Durand F. A dataset of multi-illumination images in the wild. InProceedings of the IEEE/CVF International Conference on Computer Vision 2019 (pp. 4080-4089).

[Magar et al. 25] Magar N, Hertz A, Tabellion E, Pritch Y, Rav-Acha A, Shamir A, Hoshen Y. LightLab: Controlling Light Sources in Images with Diffusion Models. InProceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers 2025 Aug 10 (pp. 1-11).