Learning When to Trust Geometric Priors for 3D Surface Reconstruction

(Masters-level internship)

Supervised by Hugo Blanc and George Drettakis
GRAPHDECO, Inria Sophia Antipolis, France
hugo.blanc@inria.fr george.drettakis@inria.fr

3D Computer Graphics, Deep Learning, and Differentiable Rendering

Context and goal

Reconstructing accurate 3D surfaces from RGB images is a fundamental problem in computer graphics. Recent differentiable rendering methods such as Geo-NeuS [3] and PGSR [1] exploit geometric and photometric constraints across multiple views, but reconstruction remains difficult in poorly textured, reflective, or sparsely observed regions where the images do not sufficiently constrain the geometry.

A common remedy is to add *geometric priors* — depth maps or surface normals predicted by pre-trained models — which provide constraints when the images alone are insufficient [2, 4]. However, such priors are not always reliable: inaccurate predictions can corrupt the recovered geometry, and they add little where multi-view observations already constrain the surface well. The goal of this internship is to determine **where geometric priors can improve surface reconstruction and how to adapt their influence to both their reliability and the information available from the input images**.



Figure 1: *Surface reconstructions and corresponding normal maps produced by PGSR [1].*

Approach

We will start with an experimental study of how depth and surface-normal priors can be incorporated into an existing surface reconstruction method, to understand when they improve the result and when they degrade it, focusing on sparsely observed regions, weakly textured surfaces, and prediction errors.

Building on this, we will identify regions where priors are likely to help using two complementary cues: the reliability of the predictions, estimated from their uncertainty and cross-view consistency, and the degree to which the training images constrain the geometry (camera coverage, parallax, texture). The central objective is then a learning-based approach that automatically adapts the influence of priors during reconstruction — for example through spatially varying weights, by combining multiple prediction sources, or by selectively regularizing displacements poorly constrained by the images.

The method will be evaluated on the DTU and Tanks and Temples datasets and on Evermotion scenes, varying the number of views and the prior quality, and compared against no-prior, fixed-weight, and confidence-based baselines.

Work environment and requirements

The internship will take place within the GraphDeco team at the Inria Centre at Université Côte d'Azur in Sophia Antipolis, near Nice and Antibes.

We are looking for a Master's or final-year engineering student with a strong interest in computer graphics and machine learning. Experience with Python/PyTorch and a solid grasp of projective geometry or numerical optimization are desirable; familiarity with NeRF or Gaussian Splatting is a plus.

References

[1] D. Chen et al. *PGSR: Planar-Based Gaussian Splatting for Efficient and High-Fidelity Surface Reconstruction*. IEEE TVCG, 2024. Paper.

[2] Z. Yu et al. *MonoSDF: Exploring Monocular Geometric Cues for Neural Implicit Surface Reconstruction*. NeurIPS, 2022. Paper.

[3] Q. Fu et al. *Geo-NeuS: Geometry-Consistent Neural Implicit Surfaces Learning for Multi-view Reconstruction*. NeurIPS, 2022. Paper.

[4] H. Zhou and Z. Lähner. *Confidence Matters: Leveraging Multi-view Geometric Priors for GS-based Reconstruction*. 2026. Preprint.