

Wood log reconstruction from Gaussian representations

Erik Mažgon, Ciril Bohak, Matija Marolt, Žiga Lesar

University of Ljubljana, Faculty of Computer and Information Science

E-mail: em89658@student.uni-lj.si, {ciril.bohak, matija.marolt, ziga.lesar}@fri.uni-lj.si

Abstract

Accurate 3D reconstruction of natural objects such as wood logs poses significant challenges due to their complex surface geometry and textural variation. We present a mesh reconstruction pipeline that combines classical Structure-from-Motion (SfM) with 3D Gaussian Splatting and the SuGaR mesh extraction method, disregarding the use of textures. Starting from a set of 133 images of a beech log, we first generate a dense point cloud using COLMAP and compare traditional mesh extraction methods with a Gaussian-based approach. Results of qualitative evaluation show that the SuGaR method produces significantly more complete and detailed meshes, outperforming classical methods in surface continuity, feature preservation, and geometric plausibility.

1 Introduction

Reconstructing 3D models of natural objects, such as tree logs, is notoriously difficult due to their uneven surfaces, intricate textures, and frequent occlusions. These characteristics often exceed the capabilities of traditional geometry reconstruction pipelines, which rely on SfM and Multi-View Stereo (MVS) techniques to generate dense point clouds from images [8]. While tools like COLMAP [8] can produce structurally sound reconstructions, the transition from point cloud to mesh remains a weak link, commonly resulting in incomplete, over-smoothed, or artifact-laden surfaces.

To address these limitations, recent research has shifted toward neural representations that encode both geometry and appearance using learned primitives. Among these, 3D Gaussian Splatting [6] has emerged as a particularly efficient method for real-time rendering and radiance field modeling. The SuGaR framework [3] extends this approach by offering direct mesh extraction from trained Gaussian fields through geometric regularization techniques.

This paper investigates the viability of applying such neural methods to the specific task of reconstructing wood logs from photographic inputs. We assess reconstruction quality through a comparative study, using both classical mesh extraction techniques and SuGaR-based meshing, and demonstrate the practical advantages of Gaussian representations in capturing detailed surface features and producing watertight, high-fidelity meshes.

2 Related Work

Traditional mesh generation from point clouds often results in disparities such as noise and incomplete structures [7, 2]. Addressing such shortcomings, post-processing tools like MeshLab provide utility for enhancing the quality of extracted meshes; however, they do require significant manual input and can still fall short in handling complex geometries [9].

Recent advancements have sought to integrate modern computational techniques with classical approaches to bolster the accuracy and completeness of 3D reconstructions. The incorporation of Gaussian-based methods has demonstrated substantial improvements in generating high-quality mesh representations [10]. The SuGaR mesh extraction technique capitalizes on these findings, boasting superior performance metrics in detail preservation and surface continuity compared to traditional methods [3].

Furthermore, the concerted effort to combine neural network methodologies with classical reconstruction techniques aligns with broader trends in the field where automatic and intelligent systems are increasingly leveraged for enhanced efficiency and accuracy in 3D model generation. For instance, advancements in neural radiance fields (NeRFs) have underscored the significant potential of these approaches in capturing detailed surface information and facilitating novel view synthesis, particularly advantageous in scenarios involving intricate textures and geometries [11, 4].

The present study builds upon these established techniques. By systematically analyzing the efficacy of both the SuGaR mesh extraction method and traditional approaches in the context of 3D reconstruction from natural wood log imagery, we provide conclusions that affirm the superiority of the neural-based extractors in generating more complete and geometrically plausible representations. This work not only contributes to the existing literature on 3D reconstruction methodologies but also integrates complementary frameworks to address the inherent challenges associated with reconstructing complex natural surfaces.

3 Method

3.1 Image Acquisition

High-resolution images of various tree logs were provided by the Biotechnical Faculty of the University of Ljubljana. The logs were photographed from multiple viewpoints to ensure full coverage of the visible surface geometry. For this study, we selected a beech log due to its distinctive bark texture and strong visual contrast with the dark background. A total of 133 images (see Figure 1) in JPEG format were used for further processing. JPEG compression artifacts are negligible and did not influence the resulting reconstructions in any meaningful way.




Figure 1: Basic point cloud reconstruction of a log using COLMAP with highlighted camera positions in red.

3.2 Structure-from-Motion and Dense Reconstruction with COLMAP

We employed COLMAP, an established SfM and MVS tool, to reconstruct a 3D point cloud from the image set. The pipeline included feature detection, camera pose estimation, and stereo depth fusion. The resulting dense point cloud captured the overall geometry of the log and served as the foundation for both naive and advanced mesh reconstruction methods. Minor adjustments to point size and rendering parameters were made to facilitate visualization and comparative analysis (see Figure 1).

3.3 Initial Mesh Extraction via Traditional Methods

The COLMAP-generated point cloud was imported into MeshLab for baseline mesh extraction. Two common surface reconstruction techniques were applied: Ball Pivoting [1] and Screened Poisson [5] (see Figure 2). Both algorithms are widely used for generating watertight meshes from unstructured point clouds. The resulting meshes were exported and used as references to evaluate improvements achieved with neural-based reconstruction methods.

3.4 Gaussian Representation Training

To convert the point cloud into a neural 3D representation, we used the GraphDECO implementation of 3D Gaussian Splatting [6]. The training procedure was split into two stages, comprising 7,000 and 30,000 steps respectively, following the developers’ recommendations. The resulting Gaussian model encoded the radiance and geometry




Figure 2: Results of Ball Pivoting approach (top) and Screened Poisson (bottom).

of the scene using compact anisotropic primitives. Visualization was performed using tools provided within the GraphDECO framework and is shown in Figure 3.



Figure 3: Gaussian splitting reconstructed logs after 7,000 iterations (above) and 30,000 iterations (below).

3.5 Mesh Extraction using SuGaR

To obtain a high-quality mesh, we employed SuGaR [3], a recent method that extracts surface meshes directly from Gaussian splats using regularization strategies. Among the supported modes are: sdf, density, and consistency. We selected the *consistency* regularization, as it yielded the most accurate and stable surface reconstructions in our experiments, as shown in Figure 4.

Due to SuGaR’s Linux-specific dependencies, we executed this pipeline within the Windows Subsystem for Linux (WSL) environment. After successful mesh extraction, the result was exported for further analysis in Blender.

3.6 Post-processing and Visualization

The final mesh was imported into Blender for qualitative evaluation. Blender’s extensive rendering capabilities




Figure 4: Output mesh of a log generated using SuGaR method from three viewpoints.




Figure 5: Postprocessed output mesh of a log generated using SuGaR method from three viewpoints.

enabled detailed inspection of surface geometry under various lighting and camera conditions. To isolate the object of interest, background elements were removed, allowing us to assess the fidelity and continuity of the reconstructed log mesh. The final cleaned mesh was exported as the definitive output of the reconstruction pipeline. An example can be seen in Figure 5.

4 Results

To evaluate the performance of different mesh reconstruction methods, we imported the resulting models into Blender for side-by-side inspection. Specifically, we compared three meshes: one generated using the Ball Pivoting algorithm, one using the Screened Poisson surface reconstruction, and one produced using the SuGaR method.

4.1 Classical Reconstruction Results

The mesh obtained via the Ball Pivoting algorithm exhibited significant discontinuities, especially along the bottom and lateral surfaces of the log. These artifacts

were primarily caused by sparse or missing data in regions not sufficiently captured during image acquisition. As a result, the algorithm failed to form a watertight mesh, leaving gaps and isolated fragments (see fig. 2, top).

The Screened Poisson method produced a more complete and coherent mesh, preserving coarse geometric features such as the bark’s undulating patterns. However, it excessively smoothed fine-scale surface details. In regions with sparse point density, the algorithm generated spurious geometry and inflated volumes that deviated from the log’s actual structure (see fig. 2, bottom).

4.2 Gaussian-Based Reconstruction

The mesh extracted using the SuGaR method significantly outperformed the classical approaches. It successfully reconstructed detailed surface features, including the fine bark texture and characteristic notches present at the end of the log (see Figure 5). These features were missed or inadequately represented in the other two methods.

SuGaR produced a watertight mesh with high surface continuity and sharp edge preservation. Although some minor imperfections were observed—such as small surface craters not present in the original imagery—they were minimal and did not detract from the overall fidelity. Notably, the method avoided introducing artificial structures outside the original surface bounds, maintaining geometric plausibility throughout the model.

4.3 Qualitative Comparison

Figure 6 shows a visual comparison of the three reconstruction outputs from multiple viewpoints. The SuGaR reconstruction not only preserves surface realism but also offers a more faithful spatial interpretation of the log geometry. This demonstrates the superiority of Gaussian-based reconstruction in handling complex natural textures and incomplete input data.

Overall, the SuGaR method produced the most accurate and visually convincing result, highlighting the potential of neural implicit representations for high-fidelity 3D reconstruction tasks.

5 Discussion

Although the final results demonstrate the effectiveness of Gaussian-based reconstruction, the development was not without challenges.

One major obstacle was the complexity of the tools and methods employed. Many of the frameworks used (COLMAP, GraphDECO, and SuGaR) require familiarity with command-line interfaces, environment configuration, and deep learning workflows. For the primary author, who was encountering most of these tools for the first time, a substantial portion of the development time was spent resolving installation issues, dependency conflicts, and runtime errors. These difficulties were exacerbated by incomplete or outdated documentation, particularly in the case of SuGaR, which required adaptation to run within a WSL environment.

Despite these technical hurdles, our work successfully demonstrated that combining classical Structure-from-




Figure 6: Comparison of the mesh reconstruction of the same log from two viewpoints using different techniques. The best result is obtained using the SuGaR approach (the backmost log in the left image and the frontmost log in the right image). The worst result is produced using the Ball Pivoting approach (the frontmost log in the left image and the backmost log in the right image). The mid logs are generated using the Screened Poisson surface reconstruction.

Motion techniques with modern neural representations can significantly improve reconstruction quality. The pipeline, starting from traditional image-based reconstruction and culminating in mesh extraction via SuGaR, achieved detailed and continuous 3D models that were not possible with point cloud-based methods alone.

The Gaussian representation approach offers several advantages: it avoids the voxelization or surface-fitting heuristics typical of classical methods, and it enables a more accurate encoding of both geometry and appearance. While training such models remains computationally expensive and time-consuming, the improvement in surface fidelity justifies the overhead for applications requiring high-quality reconstructions.

Finally, the results underscore the importance of input data quality and coverage. Even the most advanced reconstruction techniques struggle in the absence of sufficient viewpoints or consistent lighting. Future work should consider strategies for automated data acquisition and augmentation, as well as integration of photometric priors to further enhance mesh realism and completeness.

6 Conclusion

In this paper, we presented the complete process of generating a 3D mesh of a tree log based on 3D Gaussian splatting and the SuGaR method. Qualitative evaluation shows that Gaussian-based methods surpass classical approaches in both detail and surface realism. There are many possibilities for further work, but the best would be to start with the extraction of textured meshes as an upgrade of the SuGaR method. However, texture extraction and automatic mesh unwrapping pose their own unique and difficult challenges, which for now is out of the scope of our work.

Acknowledgment

This research was conducted as part of the basic research project *Cell visualization of unified microscopic data and procedurally generated sub-cellular structures* [project number J2-50221], funded by the Slovenian Research and Innovation Agency (Javna agencija za znanstvenoraziskovalno in inovacijsko dejavnost RS) from the state budget.

References

- [1] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5(4):349–359, 1999.
- [2] Z. Chen, H. Lv, L. Lou, and J. H. Doonan. Fast and accurate 3d reconstruction of plants using mvsnets and multi-view images. *Advances in Intelligent Systems and Computing*, pages 390–399, 2021.
- [3] A. Guédon and V. Lepetit. SuGaR: Surface-Aligned Gaussian Splatting for Efficient 3D Mesh Reconstruction and High-Quality Mesh Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5354–5363, June 2024.
- [4] Z. Ibork, A. Nouri, O. Lézoray, C. Charrier, and R. Touahni. No reference 3d mesh quality assessment using deep convolutional features. *2023 International Symposium on Image and Signal Processing and Analysis (ISPA)*, pages 1–6, 2023.
- [5] M. Kazhdan and H. Hoppe. Screened poisson surface reconstruction. *ACM Trans. Graph.*, 32(3), jul 2013.
- [6] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4), July 2023.
- [7] S. Liu, W. P. Bonelli, P. Pietrzyk, and A. Bucksch. Comparison of open-source three-dimensional reconstruction pipelines for maize-root phenotyping. *The Plant Phenome Journal*, 6, 2023.
- [8] J. L. Schönberger, E. Zheng, J.-M. Frahm, and M. Pollefeys. Pixelwise View Selection for Unstructured Multi-View Stereo. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *Computer Vision – ECCV 2016*, pages 501–518, Cham, 2016. Springer International Publishing.
- [9] Y. Xu, Y. So, and S. Woo. Plane fitting in 3d reconstruction to preserve smooth homogeneous surfaces. *Sensors*, 22:9391, 2022.
- [10] J. Zhang, B. Deng, H. C. Yang, Y. Peng, W. Qin, and L. Liu. Static/dynamic filtering for mesh geometry. *IEEE Transactions on Visualization and Computer Graphics*, 25:1774–1787, 2019.
- [11] Z. Zhang, Y. Wang, P. K. Jimack, and H. Wang. Meshingnet: a new mesh generation method based on deep learning. *Lecture Notes in Computer Science*, pages 186–198, 2020.