Gaussian splatting for the reconstruction of complex and highly detailed object

Sara Gonizzi Barsanti [1]*, Dario Billi [2,3], Etienne Sommer [2], Erika Elefante [1], Arnadi Murtiyoso [2]

[1] Department of Engineering, Università degli Studi della Campania Luigi Vanvitelli 81031 Via Roma 29, Aversa (CE) Italy – (sara.gonizzibarsanti, erika.elefante)@unicampania.it
[2] Université de Strasbourg, INSA Strasbourg, CNRS, ICube Laboratory UMR 7357, Photogrammetry and Geomatics Group, 67000 Strasbourg, France – (arnadi.murtiyoso, etienne.sommer, dario.billi)@insa-strasbourg.fr
[3] Department of Civil and Industrial Engineering, ASTRO Laboratory, University of Pisa, Largo Lucio Lazzarino, 56122 Pisa, Italy – dario.billi@phd.unipi.it

Keywords: Structure from Motion, 3D Gaussian Splatting, Cultural Heritage, Complex Structure

Abstract

In recent years, Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS) have emerged as advanced methods for photogrammetry-based 3D reconstruction. Since its introduction in 2020, NeRF has gained significant attention due to its capability to generate high-fidelity reconstructions from multi-view imagery. More recently, 3D Gaussian Splatting (3DGS), introduced in 2023, has proposed an alternative explicit scene representation based on a collection of anisotropic Gaussian primitives optimized directly in 3D space. This representation allows efficient rendering and scalable modelling of complex scenes while maintaining high visual quality. This paper analyses the performance of different 3DGS methods when dealing with complex geometry and less-cooperative surfaces compared to standard SfM IM procedures. Included in the comparison is also the Mesh-In-the-Loop Gaussian Splatting for Detailed and Efficient Surface Reconstruction (MILo), a novel meshing method using Gaussian splats. Three Gaussian splatting methods as implemented in the Postshot commercial software were also tested. Our experiments show that MILo shows very promising results in terms of detail reconstruction, while standard Gaussian splatting excels in visualisation but is still plagued by a high rate of noise especially when converted into a geometric point cloud form.

1. Introduction

In the field of computer graphics, particularly in 3D rendering, Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS) have recently emerged as prominent techniques. A comparative analysis of Structure from Motion (SfM), NeRF, and 3DGS highlights significant methodological and conceptual differences in approaches to 3D reconstruction and scene representation. SfM is a geometry-oriented method that reconstructs sparse three-dimensional point clouds from overlapping two-dimensional images. SfM is often followed by a dense matching process via Multi-View Stereo (MVS) to create a dense point cloud. This process involves feature detection and matching, followed by the estimation of intrinsic and extrinsic camera parameters. SfM is widely used in cultural heritage documentation due to its robustness, transparency, and compatibility with established surveying practices. Although it can achieve high geometric accuracy under suitable acquisition conditions, SfM does not inherently model view-dependent reflectance effects or complex light transport phenomena. Radiance field representations mark a significant shift compared to traditional point-based 3D models. Instead of describing geometry solely as discrete Cartesian points with optional attributes, radiance fields capture both the structure of a scene and its view-dependent appearance. This allows for a richer and more realistic representation of how light interacts with surfaces.

NeRF (Mildenhall et al., 2021) introduces a learning-based framework for volumetric reconstruction. Rather than generating explicit geometric primitives, NeRF represents a scene as a continuous volumetric function encoded within a neural network, mapping spatial coordinates and viewing directions to colour and density values. This formulation implicitly captures geometry within a radiance field. Compared to SfM, NeRF excels in photorealistic rendering and in modelling view-dependent visual effects, such as specular reflections and complex illumination. However, it typically requires substantial computational resources, dense sampling during both training and inference, and careful optimisation. Furthermore, extracting explicit and metrically reliable geometry suitable for quantitative or structural analysis remains challenging.

3DGS can be considered an intermediate approach that combines aspects of explicit geometric representation with volumetric rendering strategies (Fei et al., 2025). In this framework, a scene is represented as a set of anisotropic three-dimensional Gaussians, each defined by spatial position, covariance, opacity, and radiance attributes. Rendering is performed through rasterization-based projection and compositing of these primitives. Compared to NeRF, 3DGS significantly reduces computational demands and enables near real-time rendering, making it particularly suitable for large-scale environments and interactive applications. It is also effective for representing diffuse or visually ambiguous objects, as well as semi-transparent materials, where volumetric density variations are relevant. Compared to SfM, it incorporates volumetric radiance modelling, although geometry remains discretely defined rather than continuously parameterised. Despite its visual fidelity and efficiency, its geometric accuracy and suitability for metric or structural analysis require further systematic evaluation. In summary, SfM emphasises explicit and metrically controlled geometric reconstruction; NeRF prioritises continuous, photorealistic volumetric representation with implicit geometry; and 3D Gaussian Splatting offers a computationally efficient, discrete volumetric approximation optimised for real-time visualisation. The choice of methodology should therefore be guided by project-specific objectives, including geometric accuracy, visual realism, computational performance, and integration with downstream analytical workflows.

The present study aims to evaluate the performance of 3DGS in the documentation of complex and challenging artefacts. The test object is a wooden gilded table, chosen among the furniture of the Royal Palace of Caserta (Fig. 1). The choice was done

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-1355-2026 |

considering the complex geometry and non-cooperative reflective material of the object.

2. State of the art

The assessment of 3DGS for three-dimensional reconstruction has garnered significant interest within academic circles, with tests conducted on a variety of object types and environmental conditions (Wu, 2024). (Ji and Yao, 2025) established a benchmarking framework aimed at evaluating the quality of outputs from Gaussian Splatting, concentrating on laboratory datasets and performance related to compression. Challenges regarding the deformation of 3DGS outcomes have been scrutinised and addressed, leading to a novel method that incorporates a mesh-based 3DGS, which utilises both vertex positions and deformation gradients to enhance guidance in the 3DGS process (Sun et al., 2025). Furthermore, comparative studies have been initiated. Sambugaro et al. (2024) conducted an evaluation of reconstructions produced by Colmap 3.14, NeRF, and 3DGS in outdoor settings, examining differences in geometric accuracy and visual quality. Billi et al. (2025) investigated the metric accuracy of 3D reconstruction of transparent objects, comparing conventional photogrammetry, 3DGS, and 2DGS, demonstrating that visual quality does not necessarily imply geometric reliability and that 2DGS offers the best trade-off between accuracy and appearance for culturally relevant glass artefacts. Within indoor environments, where precise spatial data is crucial for various applications, including object detection, virtual reconstruction, and Building Information Modelling (BIM), (Wu, 2024) offered both qualitative and quantitative analyses of datasets generated through SfM (utilising Colmap 3.14), NeRF, and 3DGS. Additionally, Ye et al. (2024) compared the results of these methodologies with Simultaneous Localisation and Mapping (SLAM) systems, evaluating performance based on tracking accuracy, mapping quality, and view synthesis capabilities, while (Wu et al., 2025) compares and integrates 3DGS with LiDAR data.

In the realm of accurately documenting cultural heritage, numerous studies have introduced modified versions of the original approach. Fang et al. (2025) detailed a revised 3DGS technique tailored for the documentation of cultural artefacts, particularly targeting issues such as inaccurate image alignment, irrelevant background data, and gaps due to inadequate photographic coverage. A dual-prior driven 3DGS framework was introduced in (He et al., 2026), featuring a feature-aware sampling algorithm that constructs a detailed geometric prior with absolute scale from the input cloud, while an ideal visual prior supervises the generation of synthetic views. Lastly, Jia et al. (2026) proposed a comprehensive framework that merges a convolution-Transformer hybrid super-resolution network with an optimised 3D Gaussian Splatting pipeline. In architecture, (Murtiyoso and Macher, 2025) examined the geometric precision of 3DGS in orthophoto generation for cultural heritage digitisation, comparing results against Terrestrial Laser Scanning (TLS) and MVS-derived data.

Regarding 3DGS, the commercial software Postshot by Jawset displays a user-friendly approach to train, edit and render radiance fields. Postshot is capable of estimating camera poses directly from input images using camera tracking. However, importing a precomputed alignment generally improves the quality of the resulting splats and enables the creation of metrically accurate scenes with minimal computational cost. After importing both images and camera poses, a few parameters need to be configured. The following descriptions are based on the tool’s official documentation and on Fadilah et al. (2026):

1. Radiance Field Profile – This defines the reconstruction approach.
    a. Splat3 is considered the optimal choice, as it provides high-quality reconstruction of fine details in both foreground and background. It also makes better use of high-resolution image data compared to other profiles. It is developed by Jawset and it is said to be the best approach for reconstruction fine details (https://www.jawset.com, accessed on: 24 March 2026).
    b. Splat3 MCMC applies a more stochastic sampling strategy. It basically replaces the ADC mechanism with a probabilistic framework that leverages Markov Chain Monte Carlo sampling, using Stochastic Gradient Langevin Dynamics (SGLD). While this can improve exploration of the scene, it may result in fewer fine details.
    c. Splat3 ADC (adaptive density control) progressively densifies the scene during training. Unlike Splat3 and Splat3 MCMC, it does not allow setting a maximum number of splats. Instead, model growth is controlled through the Splat Density parameter. It allows the creation of a dense collection of 3D Gaussians, it works by removing Gaussian with opacity below a specified threshold, and duplicating those in regions that are under-reconstructed, while dividing the ones in areas that are over-reconstructed.
2. Max Splat Count – Available only for Splat3 and Splat3 MCMC profiles, this parameter limits the total number of Gaussian splatting primitives generated during training.
3. Anti-Aliasing – Enabling this option enhances overall model quality and reduces visual artifacts, particularly when viewing the scene from perspectives distant from the original camera positions.

Furthermore, Guédon et al. (2025) introduced an alternative approach based on a differentiable mesh extraction framework integrated into the optimization of 3DGS representations. This method, called MILo (Mesh-In-the-Loop Gaussian Splatting for Detailed and Efficient Surface Reconstruction), performs mesh extraction recovering both vertex positions and connectivity at each training iteration directly from the Gaussian parameters. This design enables gradients to propagate from the mesh back to the Gaussians, fostering bidirectional consistency between the Gaussian model, aka volumetric representation and the extracted mesh (the surface representation).

This paper presents an evaluation of 3DGS, including MILo, when used for the reconstruction of challenging objects in terms of complex geometry and difficult materials.

3. Dataset

The main dataset used in this study consists of a late Baroque table from the 18th century (Figure 1), part of the historic furnishings of the Royal Palace of Caserta (Italy), located within a UNESCO World Heritage Site. The object was created by Raffaele Giovine and was specifically selected to evaluate performance on highly complex decorative geometries, characterized by intricate carved ornaments, gold-plated wood surfaces with strongly reflective material properties, porcelain medallions inlaid into the tabletop, a Sèvres corbeille, and a gilded and chiselled bronze vase featuring acanthus leaves and winged putti. This combination of highly specular materials, fine surface detail, and heterogeneous object composition makes

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-1355-2026 | © Author(s) 2026. CC BY 4.0 License.

accurate geometric and radiometric reconstruction particularly challenging.

For this dataset, RGB imaging was employed: a total of 306 RGB images were acquired using a professional Canon EOS R6 Mark II mirrorless camera. The dataset was designed to assess reconstruction performance under challenging material conditions, particularly for man-made objects exhibiting complex reflectance behaviour, extreme surface detail, and mixed material composition. All analyses and experiments were conducted on the same hardware configuration to ensure consistency across evaluations. Specifically, tests were performed using an NVIDIA GeForce RTX 4090 GPU 24 GB VRAM and an AMD Ryzen 9 7950X 16-Core CPU.

4. Methodology

The analysis encompasses a geometric comparison initiated from a shared input: camera orientation parameters (both intrinsic and extrinsic) together with the sparse point cloud derived from SfM processing in Metashape v.2.0.0. This establishes a unified reference framework for the parallel reconstruction workflows that were executed independently.

The first workflow applies conventional MVS algorithms to generate a dense point cloud, which is subsequently converted into a triangulated mesh representing the complete 3D surface geometry.

Figure 1. The Baroque table from the 18th century datasets used for GS techniques evaluation

The second workflow utilizes the MILo framework, which produces two distinct outputs from a single training process. As in standard Gaussian Splatting, MILo initializes Gaussian primitives at the sparse point positions and iteratively refines their parameters as position, scale, rotation, colour, spherical harmonics (SH) coefficients, and opacity, by minimizing a photometric loss against the input images, yielding a densified point splat representation. Simultaneously and distinctively, signed distance values are computed directly from the Gaussian parameters, and a differentiable surface mesh is extracted via Delaunay triangulation. Bidirectional consistency between the volumetric Gaussian representation and the surface mesh is enforced throughout optimization, resulting in an iteratively refined mesh that is geometrically coupled to the learned scene representation.

The third, fourth, and fifth workflows are implemented within the Postshot platform, which provides three distinct Gaussian Splatting optimization strategies: MCMC, ADC, and Splat3. Unlike MILo, Postshot does not produce a surface mesh; instead, its native output consists exclusively of point splats, i.e. the set of optimized 3D Gaussian primitives resulting from the photometric training process. These three approaches differ in the way Gaussian primitives are initialized, densified, and refined during training, allowing for a direct assessment of how different algorithmic formulations of the Gaussian representation affect the fidelity of the reconstructed geometry. Postshot offers the advantage of enabling reconstruction using images at full resolution, without the memory overload that is often associated with conventional 3DGS reconstruction methods. Consequently, for this study, the MCMC, ADC and Splat3 models were used with the images at full resolution.

The limit on the number of splats was set at 3 million, and the number of iterations at 30 million. Previous studies have shown that these parameters enable convergence towards the result whilst maintaining a short processing time compared to reconstruction and texturing using conventional MVS techniques (Sommer et al., 2025). Figure 2 shows an example of the level of detail that can be achieved using 3DGS reconstruction with Postshot’s Splat3 model. This method allows for the accurate reconstruction of fine details and shiny surfaces, which is typically a challenge in photogrammetric reconstructions. The resulting outputs from all five pipelines were subjected to rigorous comparative evaluation against a reference geometry derived directly from the dataset.

Figure 2. Example of the details achieved with the Splat3 Model

In the absence of acquisition, the reference was established by extracting a geometric primitive from the dense point cloud produced by the MVS pipeline, using the RANSAC Shape Detection plugin in CloudCompare v.2.13.2. The horizontal working surface of the table was identified as the most suitable reference primitive, owing to its geometric regularity and unambiguous planarity properties that ensure reliable fitting via RANSAC and meaningful cross-pipeline comparison. The choice of the MVS output as the source for primitive extraction was further motivated by its higher point density and greater geometric consistency relative to Gaussian-based representations.

To harmonize the heterogeneous nature of the pipeline outputs prior to comparison, a unified point cloud representation was derived for each method. For the MVS and MILo pipelines, vertex positions were extracted from the reconstructed meshes, effectively converting surface geometry into comparable point sets. For the three Postshot-based approaches (MCMC, ADC, and Splat 3), the optimized point splats constituting the native output were used directly, without any further conversion.
All five point clouds were then evaluated against the RANSAC-derived reference plane through Cloud-to-Primitive (C2Prim) distance computation in CloudCompare, performed independently for each pipeline. The comparative evaluation was structured across two complementary levels of analysis.
The primary quantitative and qualitative assessment was conducted using the RANSAC-derived planar reference, providing a geometrically controlled and well-defined ground truth against which all five pipelines were compared via C2Prim distance metrics. This analysis encompassed mean error, standard deviation, outlier percentage, and time processing enabling a rigorous numerical comparison of surface accuracy across methods.

In addition, a broader qualitative evaluation was performed at the scale of the entire object, with the aim of assessing global reconstruction fidelity beyond the controlled planar region. For this purpose, the complete MVS mesh generated in Metashape was adopted as the reference model, and Multiscale Model-to-Model Cloud Comparison (M3C2) analysis was performed in CloudCompare to compare it against the four Gaussian-based pipeline outputs independently. This global comparison allowed for a visual and spatial assessment of deviation patterns, surface completeness, and the capacity of each Gaussian Splatting approach to faithfully reproduce the complex geometry and fine decorative detail of the object across its full extent.
To ensure fair and meaningful comparison, all reconstruction pipelines were configured with equivalent processing parameters where applicable (e.g., medium resolution for mesh generation), and all share the same SfM-derived camera poses and sparse point cloud as their common starting point.

5. Results and Discussions

5.1 Comparison to an ideal surface

The quantitative evaluation of surface accuracy was carried out by computing C2Prim signed distances between each pipeline's output and the planar reference extracted via RANSAC from the MVS dense point cloud, corresponding to the horizontal working surface of the table. Results are summarized in Table 1 and shown in Figure 3.

	AM - MVS	MILo - GS	Postshot - ADC	Postshot - MCMC	Postshot - Splat3
St. dev. (cm)	0.60	0.60	0.70	0.70	0.70
Mean (cm)	0.20	0.30	0.10	0.00	0.00
Nr. Points	619,523	610,607	50,889	300,600	176,498
Outliers	74,838	79,104	12,683	61,278	41,985
% Outliers	12.05%	12.95%	24.92%	20.39%	23.79%

Table 1. C2Prim distance metrics computed against the RANSAC-derived planar reference for all five reconstruction pipelines.

The MVS pipeline achieved a mean deviation of 0.20 cm with a standard deviation of 0.60 cm, and a relatively low outlier rate of 12.08%, reflecting the geometric regularity and high point density characteristic of dense stereo reconstruction. MILo produced comparable results, with a slightly higher mean deviation of 0.30 cm and an identical standard deviation of 0.60 cm, while maintaining a similarly low outlier percentage of 12.95%. Both pipelines yielded point clouds more than 600,000 points over the reference region, confirming their capacity for dense surface sampling.

The three Postshot-based approaches, ADC, MCMC, and Splat 3, exhibited a standard deviation of 0.70 cm across all variants, marginally higher than the mesh-based methods. Mean deviations were generally lower (0.00 cm for MCMC and Splat 3; −0.10 cm for ADC, indicating a slight systematic negative offset), suggesting good planar alignment on average, though at the cost of significantly higher outlier rates ranging from 20.39% to 24.92%. This is consistent with the inherently noisy and less geometrically constrained nature of point splat representations, which tend to produce sparser and less regularly distributed point sets as reflected in the markedly lower point counts, particularly for ADC (50,889 points).

Qualitatively, the C2Prim deviation analysis (Figure 3), reveals distinctive spatial patterns that complement and enrich the numerical metrics. The MVS and MILo outputs display a broadly coherent deviation distribution across the tabletop surface, with large areas rendered in green tones (deviations within approximately ±0.004 m), confirming close adherence to the reference plane. However, both maps exhibit a pronounced ring of high positive deviations (yellow to red, reaching +0.015 m) along the central oval opening of the tabletop and along the outer perimeter.
This pattern is consistent with the known convexity of the actual table surface, the planimetric reference extracted via RANSAC represents an ideal flat plane, and the real surface curves slightly upward towards its edges and around the central void, producing systematic positive offsets in these regions rather than reconstruction errors per se. A small, localized anomaly visible as a blue spot (strong negative deviation) in the MVS map near the lower-left corner of the tabletop likely corresponds to a surface artifact or a minor misalignment in the mesh at that location.

The three Postshot outputs present a markedly different visual character. While MCMC and Splat 3 maintain a spatial deviation structure broadly similar to the mesh-based methods, with the central oval void and perimeter edges displaying the same positive deviation ring, their maps appear considerably noisier, with a more scattered distribution of coloured points reflecting the irregular and less geometrically constrained nature of the point splat representation.

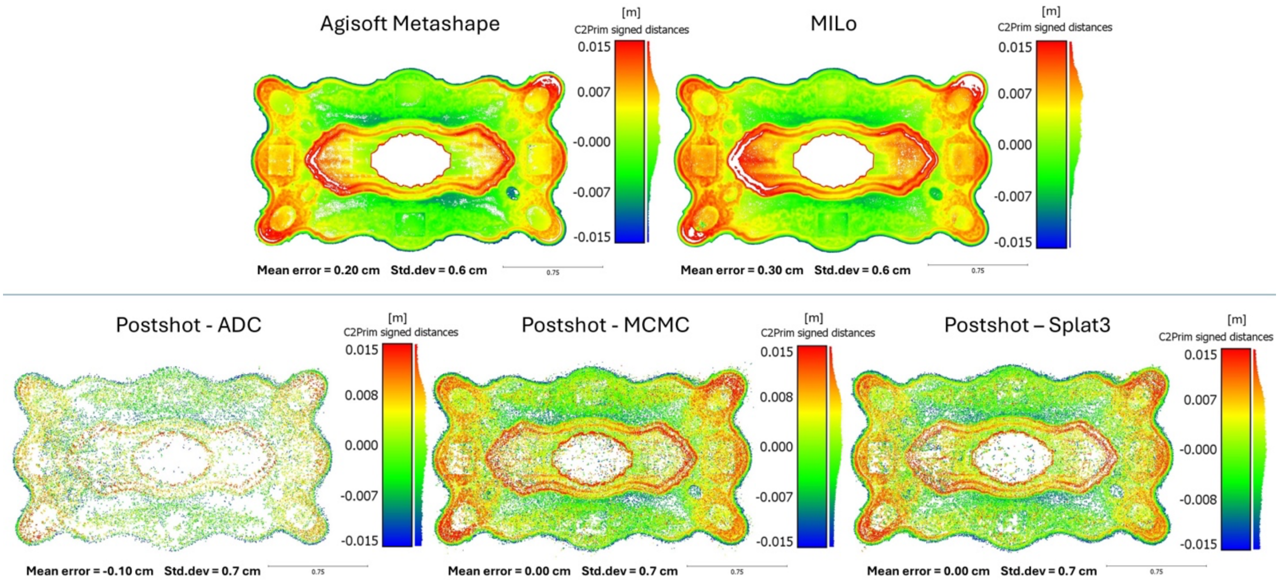

Figure 3. C2Prim distance comparison between Metashape, GS-MILo and the three Postshot workflows.

The ADC output is the most visually degraded of the three: the extremely low point density results in a sparse, fragmented map with large void areas across the tabletop interior, severely limiting the spatial completeness of the deviation analysis and making direct comparison with the denser outputs less straightforward.

Overall, the C2Prim analysis confirms that MVS and MILo provide the most geometrically dense and spatially coherent reconstructions of the reference planar surface, while Postshot-based approaches, despite achieving competitive mean deviation values, suffer from higher noise levels and substantially greater outlier rates, reflecting the intrinsic limitations of unstructured Gaussian splat representations when evaluated against a precise geometric primitive.

5.2 Qualitative global analysis

The global evaluation was conducted using M3C2 in CloudCompare, with the point cloud derived from the vertices of the complete MVS mesh produced in Metashape serving as the reference model. The four Gaussian-based pipelines, MILo, Postshot ADC, MCMC, and Splat 3, were independently compared against this reference. Results are summarized in Table 2 and illustrated in Figure 4.

	MILo - GS	Postshot - ADC	Postshot - MCMC	Postshot - Splat3
St. dev. (cm)	0.40	0.50	0.50	0.50
Mean (cm)	0.00	0.00	0.00	0.00
Nr. Points	5,922,233	359,849	2,397,053	1,319,976
Outliers	498,364	62,330	318,853	230,601
% Outliers	8.42%	17.32%	13.30%	17.47%
Time processing	203.27 min	45.8 min	96.1 min	72.6 min

Table 2. M3C2 distance metrics computed against MVS reference for all five reconstruction pipelines.

All four Gaussian-based pipelines returned a mean error of 0.00 cm against the MVS reference, indicating no systematic global offset. MILo achieved the lowest standard deviation (0.40 cm) and the lowest outlier rate (8.42%), while also producing by far the densest point representation with over 5.9 million points. This suggests that MILo's coupled mesh-Gaussian optimization yields a more geometrically faithful and complete reconstruction at the global scale.

The three Postshot pipelines all exhibited a standard deviation of 0.50 cm and higher outlier percentages, ranging from 13.30% (MCMC) to 17.47% (Splat 3). Point cloud density varied considerably across Postshot variants: MCMC produced the densest output among the three (approximately 2.4 million points), followed by Splat 3 (approximately 1.3 million) and ADC (approximately 360,000), the latter exhibiting the sparsest coverage of all evaluated methods.

Qualitatively, as visible in Figure 4, the M3C2 deviation maps reveal that all four Gaussian-based methods maintain good agreement with the MVS reference across the main body of the table, with deviations predominantly within the ±0.007 m range (green tones). Larger deviations (yellow to red) are systematically concentrated along the outer silhouette edges and in the most intricate carved decorative elements, areas where the fineness of the geometry and the complexity of material reflectance challenge all photogrammetric and Gaussian-based reconstruction approaches alike. MILo's deviation map shows the most spatially coherent low-error distribution, while the Postshot outputs, particularly ADC, display a more scattered and noisier pattern across the tabletop surface, consistent with the sparser and less geometrically constrained nature of their point splat representations.

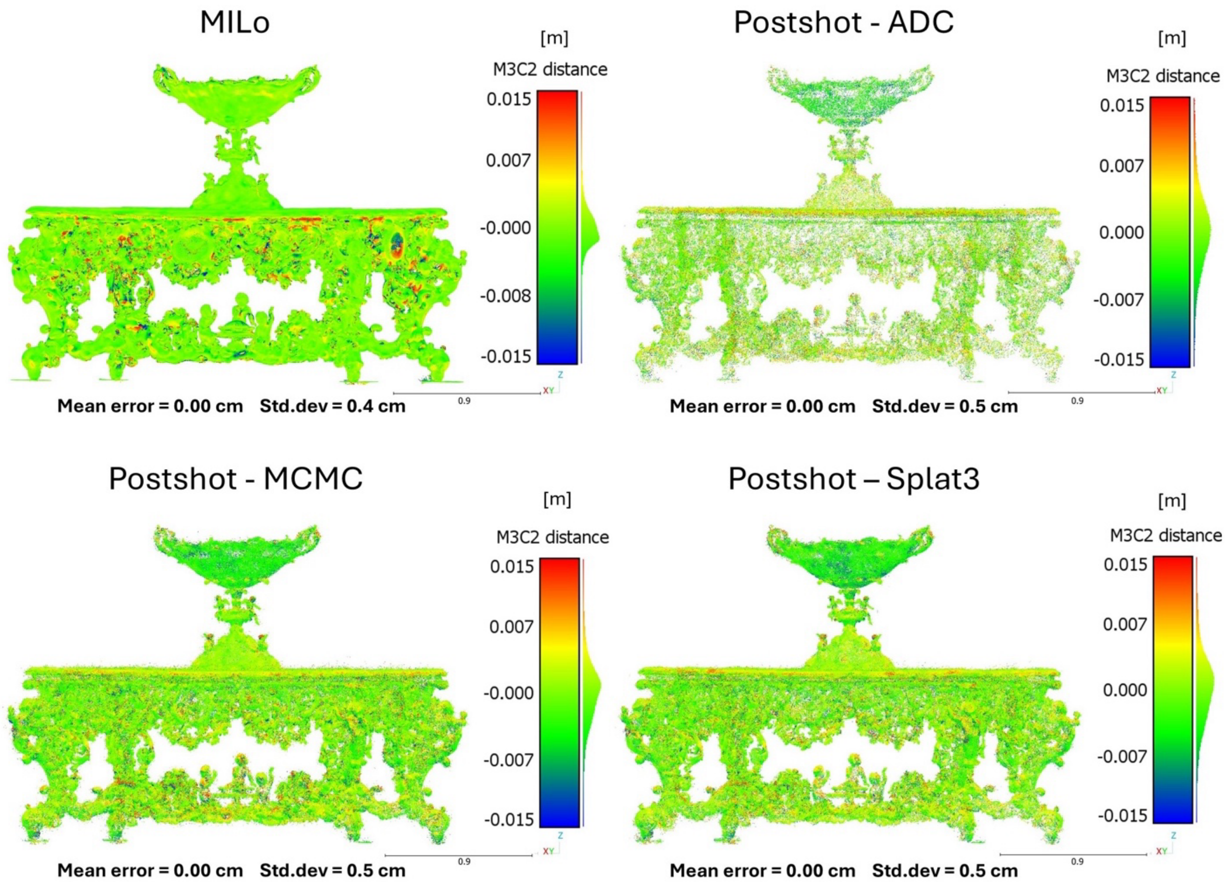

Figure 4. M3C2 distance comparison between GS – MILo and the three Postshot workflows using as reference the MVS point cloud derived from mesh vertices.

6. Conclusions

This paper evaluates four different 3DGS training strategies compared to traditional MVS procedures when applied to the 3D reconstruction of an object with complex geometry, tiny details, and made of different materials with reflective or shiny parts. The survey methodology was straightforward; the only limitation was encountered due to the usual constraints present when dealing with old, valuable assets. For example, it was not possible to stay too close to the object and targets had to be placed only on the floor and the use of a tripod was not allowed. As for the processing, the standard SfM procedure showed a strong consistency and reliability as expected since the procedure is standard and the software used is nowadays the most consistent.

The different GS approaches have shown different results considering the different comparison procedures used. When providing a quantitative and qualitative analysis using a RANSAC-derived reference plane, MILo and MVS showed comparable results which differ only in the value of mean deviation of which MILo is 1 mm higher. The standard deviation is identical, and the outlier profiles are comparable (12.08% and 12.95%). The standard deviation of the three approaches in the Postshot platform tested showed a slightly worse standard deviation (1 mm higher) Postshot also yielded notably higher levels of outliers, explained by the higher presence of noise in the GS data. Considering also a visual comparison, the ADC setting gave the worst result, consisting in a very sparse cloud.

The results of the comparison highlighted the similar behaviour of MVS and MILo, with high density data and comparable deviations, while the noise affecting the Postshot results strongly influenced the completeness and the accuracy of the data.
The comparison of all the four datasets with the MVS obtained with Metashape resulted in the same conclusions: MILo showed a strong consistency. The three Postshot procedures, even if showing a mean error of 0.00, gave a higher deviation and a density of point variable from MCMC (denser) to ADC (lower).
To summarise, MILo appeared to provide the best results both in terms of qualitative and quantitative comparison, while the three GS approaches are strongly influenced by the parameters and training strategies.

Future work will include the use of ground truth data (e.g. lidar) to provide a common comparison benchmark. Indeed, one of the weak points of this paper is that the experiment setup does not allow a proper comparison between MVS and GS methods; although we argue that it already shows a good overview of how each GS method compares to each other. Using high precision lidar data as ground truth would fortify our conclusions by including MVS as an element in the comparison. Furthermore, more case studies with different complexities both in terms of objects and architectural examples are also planned in the near future to extend this paper's observations into a larger choice of heterogeneous datasets.

Acknowledgements

The authors want to thank the director of the Royal Palace of Caserta, Dr Tiziana Maffei for allowing the survey of the test object, and Dr Anna Manzone, conservationist restorer responsible for the maintenance and restoration of works of art, museum deposits and condition reports, for the help in the organisation of the survey. The research work is also supported by the Digital Twins for Cultural Heritage (Twin4CH - https://twin4ch.github.io/) Junior Professorship Chair financed by the French Agence Nationale de la Recherche (ANR).

References

Billi, D., Caroti, G., Piemonte, A., 2025. Metric Error Assessment Regarding Geometric 3D Reconstruction of Transparent Surfaces via SfM Enhanced by 2D and 3D Gaussian Splatting. *Sensors* 25(14), 4410. https://doi.org/10.3390/s25144410

Fadilah, W. A., Murtiyoso, A., Landes, T., and Grussenmeyer, P., 2026: Metric Assessment of 3D Gaussian Splatting for UAV-Based Urban Heritage Reconstruction, International Archives of Photogrammetry and Remote Sensing Spatial Information Sciences, XLVIII-2/W12-2026, 143–150, https://doi.org/10.5194/isprs-archives-XLVIII-2-W12-2026-143-2026

Fang, X., Zhang, Y., Tan, H., Liu, C., Yang, X., 2025: Performance Evaluation and Optimization of 3D Gaussian Splatting in Indoor Scene Generation and Rendering, ISPRS International Journal of Geo-Information, 14, no. 1: 21. https://doi.org/10.3390/ijgi14010021

Fei, B., Xu, J., Zhang, R., Zhou, Q., Yang, W., He, Y., 2025: 3D Gaussian Splatting as a New Era: A Survey, IEEE Transactions on Visualization and Computer Graphics, vol. 31, no. 8, pp. 4429-4449, doi:10.1109/TVCG.2024.3397828

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. arXiv preprint arXiv:2506.24096. https://doi.org/10.48550/arXiv.2506.24096

He, Y., Zhang, X.., Xie, Z., Li, D., He, J., Chen, R., 2026: A dual-prior driven Gaussian splatting framework for high-fidelity reconstruction of museum artifacts. NPJ Herit. Sci., 14, 69. https://doi.org/10.1038/s40494-026-02330-z

Ji, B., Yao, A., 2025: SfM-Free 3D Gaussian Splatting via Hierarchical Training. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 21654-21663

Jia, Q., He, J., 2026: High-fidelity 3D reconstruction of cultural heritage via super-resolution and progressive Gaussian splatting. NPJ Herit. Sci., 14, 84. https://doi.org/10.1038/s40494-026-02355-4

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

Murtiyoso, A., Macher, H., 2025: Gaussian Splatting for Facade Orthophoto Generation – Comparison with MVS and TLS, International Archives of Photogrammetry and Remote Sensing Spatial Information Sciences, XLVIII-M-9-2025, 1059–1064, https://doi.org/10.5194/isprs-archives-XLVIII-M-9-2025-1059-2025.

Sambugaro, Z., Orlandi, L., Conci, N., 2024: 3D reconstruction methods in industrial settings: a comparative study for COLMAP, NeRF and 3D Gaussian Splatting, In CEUR WORKSHOP PROCEEDINGS, volume 3762, pages 212--217.

Sommer, E., Murtiyoso, A., Koehl, M., Grussenmeyer, P., 2025: Radiance fields for archaeological documentation of medieval rhine valley castle ruin. In Proceedings of the IEEE International Conference on Imaging Systems and Techniques (IST) 1–6.

Sun, G., Liu, D., Yang, Z., An, S., Zhang, Y., 2025: 3D reconstruction and precision evaluation of industrial components via Gaussian Splatting, Computers & Graphics, Volume 132, 104422, ISSN 0097-8493, doi.org/10.1016/j.cag.2025.104422.

Wu, T., 2024: Recent Advances in 3D Gaussian Splatting, <i>arXiv e-prints</i>, Art. no. arXiv:2403.11134, doi:10.48550/arXiv.2403.11134

Wu, Y., Liu, K., Zhang, L., Chen, M., Luo, Y., Li, H., Wang, L., Cai, Y., 2025: Improving 3D Gaussian splatting for radiance field rendering using deep structure-from-motion, in Proc. SPIE 13511, Tenth Symposium on Novel Optoelectronic Detection Technology and Applications, 1351106 https://doi.org/10.1117/12.3054275

Ye, Z., Li, W., Liu, S., Qiao, P., Dou, Y. 2024: Abs3DGS: Recovering Fine Details in 3D Gaussian Splatting. In Proceedings of the 32nd ACM International Conference on Multimedia (MM '24). Association for Computing Machinery, New York, NY, USA, 1053–1061. https://doi.org/10.1145/3664647.3681361