Modelling Transparent Surfaces in Heritage Artefacts with Gaussian Splatting

Marco Medici[1], Andrea Sterpin[1,2], Stefano Settimo[1], Matteo Bevilacqua[1],
Gianluca Bertolasi[3], Simone Rigon[3], Elisa Mariarosaria Farella [3], Fabio Remondino [3]

[1] INCEPTION, Spin-off of the University of Ferrara, Italy
Email: (marco.medici, andrea.sterpin, stefano.settimo, matteo.bevilacqua)@unife.it

[2] Department of Architecture, University of Ferrara, Italy

[3] 3D Optical Metrology (3DOM) unit, Bruno Kessler Foundation (FBK), Trento, Italy
Email: (gbertolasi, srigon, elifarella, remondino)@fbk.eu

KEYWORDS: Gaussian Splatting; non-collaborative surfaces digitisation; glass objects; cultural heritage

Abstract

The 3D reconstruction of cultural heritage artefacts plays a crucial role in documentation, conservation and dissemination. While recent advances in photogrammetry, laser scanning and neural rendering techniques have significantly improved the geometric accuracy and visual realism of digitised assets, the reconstruction of transparent and reflective materials - typical in museal collections - remains a major challenge. Materials such as glass, glazes and varnishes exhibit complex optical behaviours, leading to incomplete or inaccurate 3D models. Recent developments in Gaussian Splatting (GS) offer a potential alternative by enabling efficient, high-fidelity scene representation without explicit surface modelling. However, their application to non-Lambertian and transparent heritage objects remains largely unexplored. This paper presents a study on GS methods for the 3D digitisation of transparent cultural heritage artefacts. Through a series of experimental reconstructions, the work investigates the potential and limitations of GS, highlight the opportunities of hybrid pipelines for addressing long-standing challenges in the digitisation of non-collaborative materials.

1. Introduction

3D reconstruction of cultural heritage artefacts has become an indispensable and requested task for documentation, analysis, conservation, restoration and dissemination purposes. Advances in photogrammetry, laser scanning and, more recently, neural scene representations (Mildenhall et al., 2020; Tewari et al., 2022; Mazzacca et al., 2023) have supported these purposes and intensely improved the realism and geometric fidelity of 3D digitised heritage assets.

However, reflective, transparent and semi-transparent materials, such as glass, glazes, varnishes and crystal, remain among the most challenging elements for 3D digitisation. Their optical behaviour, characterised by refraction, reflection and transmission, violates the basic assumptions of most optical and vision-based methods, which typically rely on Lambertian reflectance models. As a result, most of the artefacts featuring glass inlays, polished stones or glazed ceramics are only partially reconstructed, creating significant gaps in both geometry and appearance. Optical methods based on shape from heating (Pelletier and Maldague, 1997; Aubreton et al., 2009), shape from amplitude (Liu et al., 2006) and light polarisation (Morel et al., 2006; Menna et al., 2016) during image acquisition demonstrated to aid 3D reconstruction processes. Transparent surfaces present a significant challenge for conventional digitisation techniques (Karami et al., 2022): structured-light 3D scanning and photogrammetry typically fail on non-collaborative surfaces, transparent or highly reflective materials, as the light passes through or scatters unpredictably. Traditionally, conservators resort to spraying transparent heritage artefacts with a temporary opaque coating to capture their geometry with laser or structured-light scanners (Yang et al., 2019). This process adds time and cost, notwithstanding the loss of the objects' transparent appearance (Stavroulakis et al., 2023).

Recent developments in Gaussian Splatting (Kerbl et al., 2023; Chen et al., 2025) have introduced a new paradigm for novel view synthesis and neural radiance field (NeRF)-like rendering that combines efficiency, high visual fidelity and almost real-time capabilities. Gaussian Splatting (GS) results represent a scene as a collection of 3D Gaussians with learned opacity, colour, and spatial extent, allowing continuous volumetric rendering without explicit meshing. While current implementations have achieved remarkable results in opaque Lambertian scenarios (Dalal et al., 2024; Jamil and Brennan, 2025), their application to non-Lambertian and transparent surfaces in heritage contexts remains quite unexplored (Jiang et al., 2024). High-fidelity transparent surface reconstruction from multi-views poses still significant challenges (Li et al., 2025; Tian et al., 2026).

The aim of this work is to report developments and experiences related to the 3D digitisation of many transparent museum artefacts using GS methods.

2. The XRculture project

The XRculture project[1], co-funded by the EC within the Digital Europe Programme, focuses on enriching Europe's common data space for cultural heritage with high-quality 3D models using innovative digitisation methods.

Among its activities, the project explores how modern image-based 3D reconstruction workflows based on GS methods could support the 3D digitisation of ancient Roman glass artefacts located in six Italian archaeological museums. These collections comprise over 400 ancient Roman glass artefacts distributed across several locations in Italy, including the museums of Adria, Verona, Aquileia, Marche region, Lomellina and Ravenna. Each museum contributes a representative selection of finely crafted vessels and glassware, offering a diverse corpus of shapes, colours, and manufacturing techniques that together form an exceptional dataset for testing advanced 3D digitisation methods. Our activities aim to overcome actual 3D digitisation limitations by leveraging 3D GS, a recent explicit point-based rendering and reconstruction technique, to directly digitise glass artefacts

[1] https://xrculture.eu/

without altering them, preserving their translucency and reflective qualities (Fei et al., 2024). By doing so, we seek to achieve a level of visual fidelity - including transparency and volumetric appearance - that is unprecedented in cultural heritage digitisation and valuable primarily for visualisation and consultation purposes. Indeed, a GS algorithm, starting from the photogrammetric sparse point cloud, iteratively refines the position, colour, opacity and scale of many 3D Gaussians to match the input photographs. This representation captures view-dependent effects such as reflections, refractions and colour attenuation - directly learned from the imagery. GS uses full pixel information rather than keypoints, making it well-suited for shiny or translucent surfaces.

3. Methodology

To investigate the benefits of Gaussian Splatting for challenging transparent museal objects, we designed and tested a specific image-capture and processing workflow, as afterwards reported.

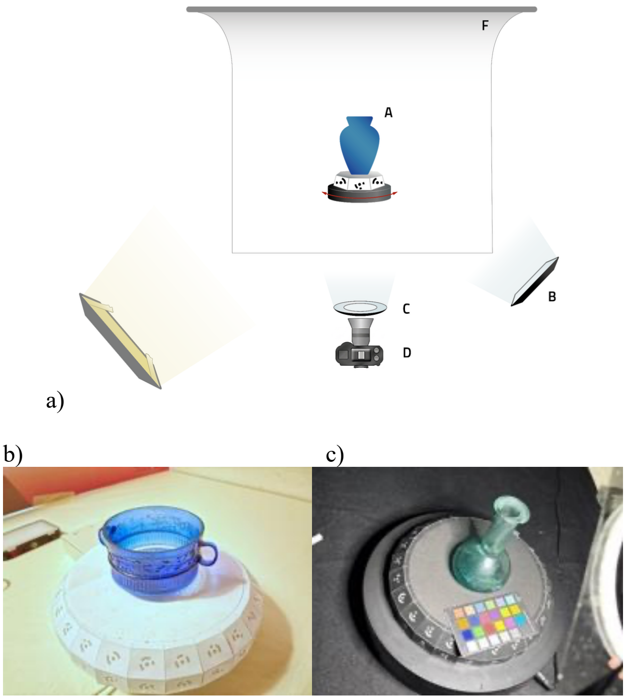

Figure 1. Rotating table setup with coded targets, a colour checker and white (left) or black (centre) background to support automated orientation processes.

3.1 The capturing setup

Figure 1 illustrates the components of the operational setup used when the archaeological artefact could be freely reoriented on the supporting platform, including inverted (top-to-bottom) orientations. In front of the camera (D), a motorised rotating table was positioned to ensure controlled and repeatable imaging conditions. To facilitate accurate photogrammetric reconstruction of transparent glass artefacts (A), a custom 3D-printed support integrated with photogrammetric markers was used. Since glass surfaces generally lack reliable geometric features for camera orientation, these markers served as high-contrast reference points for image matching. To prevent visual interference due to the object's translucency, the markers were strategically positioned on the vertical and inclined surfaces of the support rather than directly on the turntable base. Behind the system, a background panel (F), either white or black depending on the colour of the artefact, was arranged to ensure sufficient chromatic contrast. The illumination configuration included a polarised LED ring light (C) and/or a lateral LED light source (B). In some cases, faint ambient illumination was present when uncontrolled light sources could not be completely eliminated. The polarised ring light, combined with a CPL filter mounted on the lens (D), enabled a cross-polarisation configuration that provided uniform scene illumination while suppressing glare, allowing reliable acquisition of geometric and colour information useful for photometric consistency and feature extraction. The non-polarised lateral LED source illuminated the scene from a specific direction, preserving reflective and refractive phenomena such as highlights, internal reflections, and interference patterns within the transparent objects. Information captured under polarised illumination provides a critical input for appearance-based reconstruction approaches, enabling methods such as Gaussian Splatting to better model the material's optical behaviour and volumetric depth (Ferraton et al., 2009). The lateral positioning of the LED source also minimised frontal reflections that could otherwise compromise the quality of the reconstructed models.
Among the tested configurations, the combination of a polarised ring light and lateral LED illumination provided the most satisfactory results: the ring light reduced the shadows produced by the lateral source, while the latter enhanced the characteristic optical effects of the glass material, which were less pronounced in images acquired exclusively under cross-polarised conditions. This configuration allowed a clearer separation between the object's physical geometry and its appearance features (e.g., highlights and specular reflections), in line with recent approaches to transparent surface reconstruction (Li et al., 2025).
Figures 2a and 2b illustrate the operational setup used for artefacts that could not be flipped top-to-bottom or rotated laterally due to their fragility, a critical precaution when handling fragile archaeological materials.

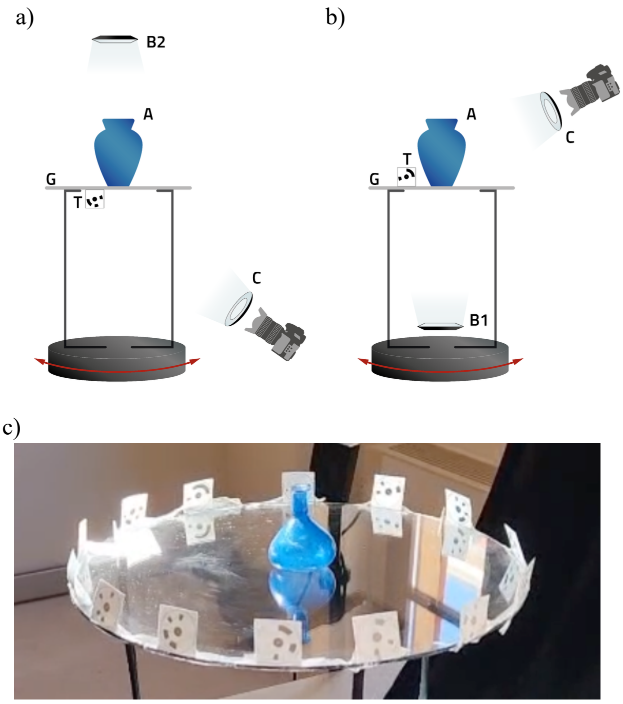

Figure 2. Transparent acrylic platform for image acquisition from beneath and avoiding an object's rotation.

As in the previous configuration, a rotating table was positioned in front of the camera. Mounted on the table, a 3D-printed support held a thin circular plexiglass plate (3 mm) on which the artefact (A) was placed while allowing images to be captured from below. Since the artefact could not be physically reoriented due to its fragility, this configuration allowed the lower portion to be captured without manipulating the object. The acrylic platform was required to be of very high optical quality, clear and free of surface defects, to avoid unwanted refractions or distortions in the captured images that could interfere with the reconstruction process. Behind the system, as before, a black or white background panel (F) ensured chromatic contrast. Ground control points (T) were placed on both sides of the plate and arranged to minimise potential occlusions of the glass artefact as well as visual disturbances caused by the visibility of the markers through the transparent glass surface. To compensate for the thickness of the plexiglass plate and avoid reconstruction issues, the Z coordinates of the upper and lower markers were corrected during the camera alignment stage of the photogrammetric processing. Given the minimal thickness and high optical quality of the acrylic plate, potential refractive effects were considered negligible for the purposes of camera pose estimation and image-based reconstruction.

As in the previous setup, the polarised ring light combined with the CPL filter on the lens (D) enabled a cross-polarisation configuration, providing diffuse illumination of the scene and attenuating the shadows generated by the non-polarised LED light (B). The latter illuminated the scene along the vertical axis (nadir in Figure. 2a; zenith in Figure. 2b), generating controlled reflections intended to highlight details of the artefact through the interaction between light and the glass material. As in the previous configuration, faint ambient illumination was sometimes present when uncontrolled light sources could not be completely eliminated.

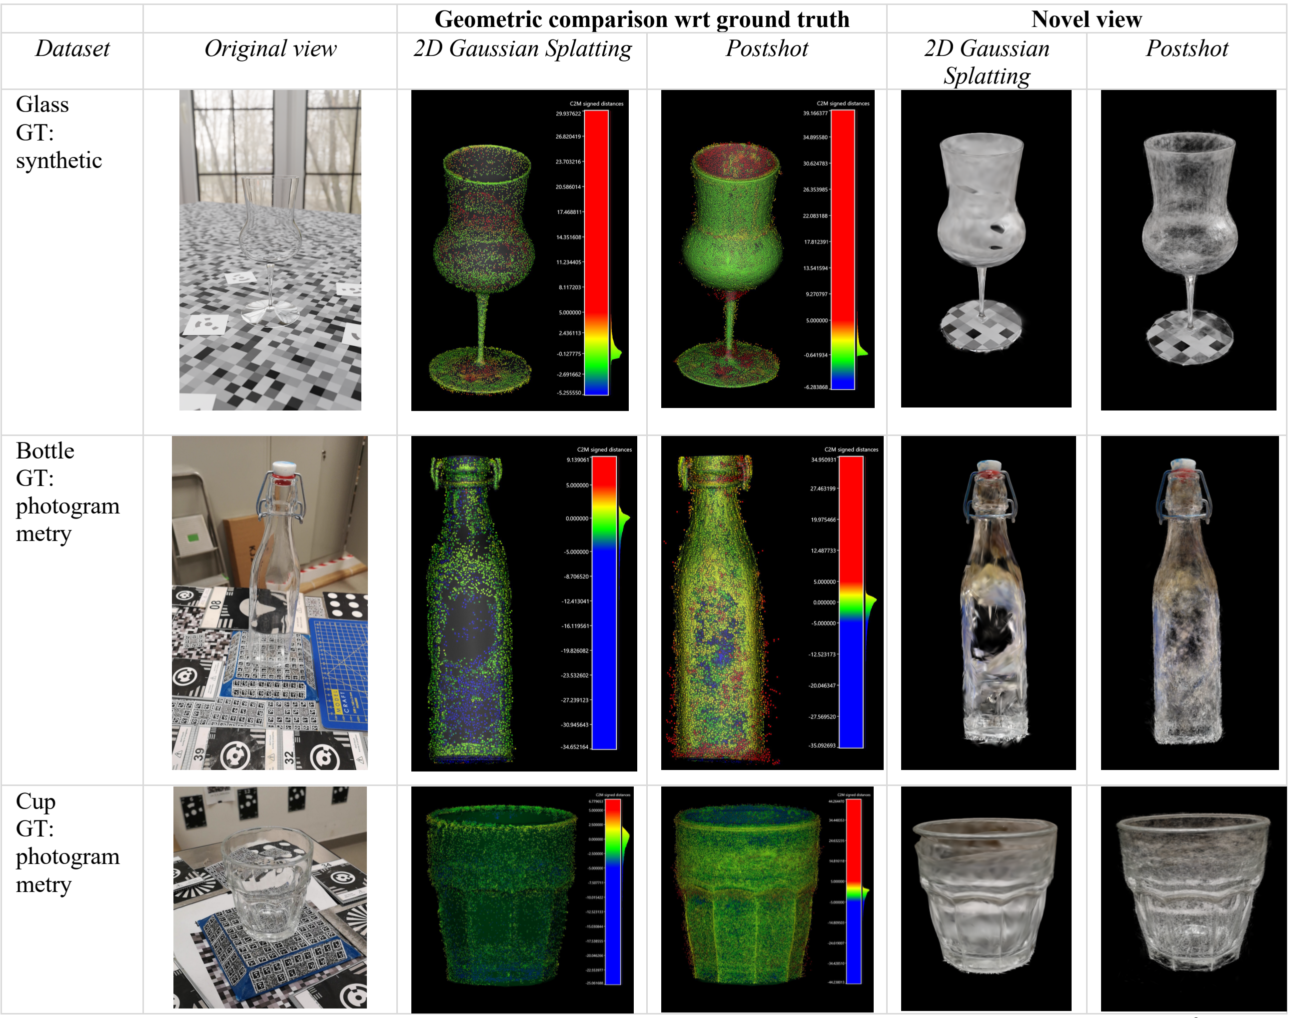

Figure 3. Evaluation (mm) of GS results on one synthetic and two real transparent objects from the NerFBK dataset[2].

Dataset	*Method*	*N. points*	*Mean*	*St. dev.*	*Min*	*Max*
Glass	Postshot	246189	2.09	4.08	-6.28	39.17
	2D Gaussian Splatting	28623	0.93	2.08	-5.25	29.34
Bottle	Postshot	285985	-3.48	5.67	-35.09	34.95
	2D Gaussian Splatting	22693	-3.32	5.22	-34.65	9.14
Cup	Postshot	343044	-3.51	5.53	-44.24	44.26
	2D Gaussian Splatting	25381	-1.69	3.17	-25.06	6.78

Table 1: Metrics (mm) for the two evaluated GS methods.

[2] https://github.com/3DOM-FBK/NeRFBK

3.2 Data processing and 3D model reconstruction

Photogrammetry is used to derive the camera poses and an initial sparse point cloud of the photographed object (Figure 4). Scaling is ensured by the markers placed around the rotating table, whereas rotated images support the self-calibration process.

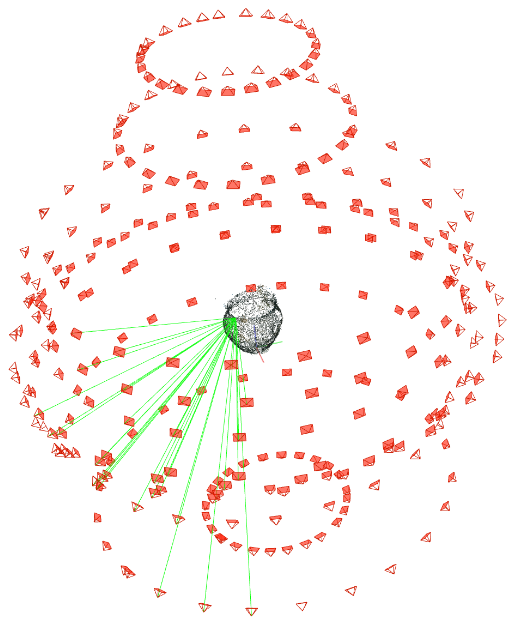

Figure 4. Camera network for an artefact photographed using the transparent platform shown in Figure 2.

Camera parameters are then used to optimise a differentiable 3D point cloud of Gaussian primitives, each defined by position, orientation, colour, opacity and scale, to reproduce the artefact’s appearance directly from the captured imagery, thereby eliminating the need for an explicit surface mesh. Among the various GS pipelines, based on efficiency and replicability, 2D Gaussian Splat (Huang et al., 2024) and Jawset PostShot[3] are evaluated. After initial laboratory testing on synthetic data to facilitate metric evaluations, PostShot proved to be the most capable solution, offering both robust handling of large multi-view image datasets, superior geometric quality in terms of more homogeneous point distribution and completeness as well as visual results on transparent and reflective materials. Table 1 and Figure 3 show quantitative and visual evaluations on a transparent object from the NerFBK dataset (Yan et al., 2023), which supported our decision. The output for each artefact is a dense set of elliptical 3D Gaussians representing position, colour, opacity, scale and orientation, normally stored in a dedicated splat format, which retains the complete Gaussian parameters required for real-time rendering and appearance-faithful visualisation. Therefore, only dedicated viewers can fully represent such 3D results.

4. Comparative analysis and results

The analysis was carried out on two digitised cultural heritage objects dating to the Modern Ages and provided by the Arheološki muzej Istre (Croatia), namely NV-1405 and S-7571 (Figure 5 and 6). These two artefacts were selected because, although both feature semi-transparent and reflective glass surfaces, their visual appearance and geometric characteristics still allowed the generation of a reference mesh through photogrammetric processing, thus enabling a direct geometric comparison between Gaussian Splatting outputs and mesh-based models. At the same time, the two objects represent complementary test cases. NV-1405 is a cup characterised by an empty inner volume, a condition in which Gaussian Splatting artefacts more frequently occur, as reflections, refractions and view-dependent highlights may be interpreted by the optimisation process as occupied 3D space, generating spurious points in regions where no physical surface is actually present. In contrast, S-7571 does not contain empty parts and therefore provides a more stable and controlled case for comparison. Their combined use made it possible to evaluate the behaviour of the method under two different yet representative conditions for the digitisation of transparent heritage artefacts.

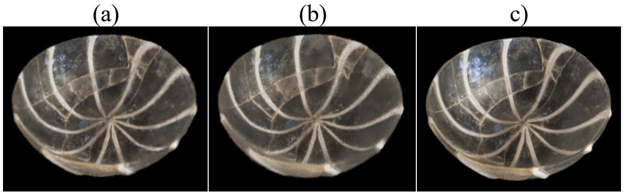

Figure 5. GS render of the object NV-1405 in three configurations of training and max splat count. a) low training and low splat count limit (V4_4), b) low training and high splat count limit (V4_1), c) high training and high splat count limit (V1_1).

For the Gaussian Splatting processing, the original full-resolution images were used in Postshot. S-7571 was reconstructed from 213 images acquired with a Sony Alpha 7R III and a 24 mm lens, with an image resolution of 7952 × 5304 pixels. The photogrammetric processing provided 72,773 tie points as input for Postshot and a reference mesh composed of 281,724 faces. NV-1405 was reconstructed from 229 images acquired with the same camera and lens, also at a resolution of 7952 × 5304 pixels. In this case, the photogrammetric alignment provided 55,970 tie points as input for Postshot and a reference mesh composed of 375,198 faces. The evaluation was carried out, in both cases, using 16 Gaussian Splatting configurations per model, obtained by combining four training lengths (3, 10, 30 and 100 ksteps) with four maximum Gaussian counts (300, 1000, 3000 and 10000 ksplats). Geometric accuracy was assessed through cloud-to-mesh distances between the GS point clouds and the reference photogrammetric meshes, while visual appearance was assessed through image-based metrics (MSE, PSNR, SSIM, MS-SSIM, LPIPS and ΔE), computed on rendered image comparisons to evaluate pixel-wise error, structural similarity, perceptual fidelity and colour consistency with respect to the reference.

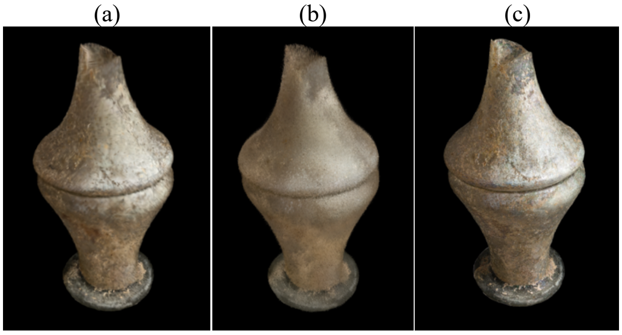

Figure 6. GS render of the object NV-1405 in three configurations of training and max splat count. a) low training and low splat count limit (V4_4), b) low training and high splat count limit (V4_1), c) high training and high splat count limit (V1_1).

[3] https://www.jawset.com/

	Input parameters		Generated splat count and geometric comparison wrt mesh			Splat count after filtering (range 0-2 mm) and geometric comparison wrt mesh		
NV-1405	**Max splat count (ksplat)**	**Training (kstep)**	**Splat/point count**	**Mean (mm)**	**St.Dev. (mm)**	**Splat/point count**	**Mean (mm)**	**St.Dev. (mm)**
V1_1	10000	100	2981439	0,527	1,803	2904474	0,361	0,293
V1_2	3000	100	2524310	0,567	2,138	2465857	0,368	0,289
V1_3	1000	100	1000000	**0,417**	**0,992**	989238	**0,356**	**0,269**
V1_4	300	100	300000	0,443	1,129	294860	**0,348**	**0,279**
V2_1	10000	30	3243091	0,652	2,479	3140569	0,366	0,29
V2_2	3000	30	2583544	0,771	3,032	2497196	0,372	0,286
V2_3	1000	30	1000000	0,504	1,645	983758	0,365	0,273
V2_4	300	30	300000	0,525	1,653	293203	0,358	0,283
V3_1	10000	10	**5189830**	0,659	2,205	**5020186**	0,404	0,295
V3_2	3000	10	2089743	0,587	1,971	2042334	0,394	0,281
V3_3	1000	10	797292	0,586	1,799	776628	0,389	0,285
V3_4	300	10	272969	0,617	1,885	264023	0,379	0,293
V4_1	10000	3	3346527	0,916	1,221	2963795	0,574	0,455
V4_2	3000	3	1103519	0,777	1,126	1011680	0,517	0,417
V4_3	1000	3	416422	0,724	1,116	386332	0,483	0,392
V4_4	300	3	160113	0,688	1,113	150146	0,46	0,372

Table 2: Metrics for the geometric evaluation of the generated GS for the object NV-1405.

	Input parameters		Generated splat count and geometric comparison wrt mesh			Splat count after filtering (range 0-2 mm) and geometric comparison wrt mesh		
S-7571	**Max splat count (ksplat)**	**Training (kstep)**	**Splat/point count**	**Mean (mm)**	**St.Dev. (mm)**	**Splat/point count**	**Mean (mm)**	**St.Dev. (mm)**
V1_1	10000	100	2491641	**1,054**	**2,339**	2177456	**0,39**	**0,442**
V1_2	3000	100	2063662	1,192	2,565	1767924	0,401	0,449
V1_3	1000	100	1000000	1,188	2,649	861542	0,392	0,45
V1_4	300	100	300000	1,261	2,706	253646	0,412	0,481
V2_1	10000	30	2859458	1,285	2,749	2418753	0,413	0,457
V2_2	3000	30	2254271	1,424	2,95	1870431	0,419	0,46
V2_3	1000	30	1000000	1,599	3,274	814041	0,42	0,468
V2_4	300	30	300000	1,666	3,263	238994	0,448	0,499
V3_1	10000	10	**5150085**	1,659	3,012	**4103939**	0,547	0,49
V3_2	3000	10	2084983	2,068	3,567	1563799	0,547	0,497
V3_3	1000	10	756277	2,327	3,898	547203	0,546	0,505
V3_4	300	10	259419	2,101	3,611	192336	0,541	0,519
V4_1	10000	3	3432350	2,692	3,487	2107426	0,83	0,548
V4_2	3000	3	1107372	2,755	3,702	690016	0,798	0,543
V4_3	1000	3	402376	2,503	3,58	267353	0,756	0,539
V4_4	300	3	154988	1,987	3,087	113854	0,709	0,536

Table 3: Metrics for the geometric evaluation of the generated GS for the object S-7571.

The geometric comparison (Table 2 and 3) shows that, in both analysed objects, increasing the number of training steps generally improves geometric agreement with the reference mesh (Figure 7), as reflected by lower cloud-to-mesh distances for configurations with the same maximum Gaussian count, with mean distances spanning from 1.054 to 2.755 mm for S-7571 and from 0.417 to 0.916 mm for NV-1405. However, the standard deviation remained relatively high in several configurations, reaching up to 3.898 mm in V3_3 for S-7571. This suggested that part of the measured discrepancy could be influenced by Gaussian Splatting artefacts rather than by deviations on the actual reconstructed surface. For this reason, a second filtered evaluation was introduced, retaining only GS points lying within the 0–2 mm distance range from the reference photogrammetric mesh and discarding all others. The aim was to reduce the influence of typical GS artefacts, such as floating splats or clusters of points generated in regions where no real surface exists. These artefacts are especially common in transparent and reflective objects, where reflections, refractions and view-dependent highlights may be interpreted by the optimisation process as occupied 3D space, producing spurious points, for example inside a cup or vessel, where apparent internal structures may be reconstructed despite being physically absent.

The filtering affected the tested configurations differently, with the strongest impact observed in models trained for fewer steps and with a high maximum splat count. The most evident case was configuration 4_1 (3 ksteps training, 10000 max ksplats), in which 39% of the points were discarded for S-7571 and 11% for NV-1405. This indicates that short training combined with a large Gaussian budget is more likely to generate unstable splats and reconstruction artefacts. The difference between the two objects is also consistent with their morphology and visibility in the image set. In S-7571, the inner part of the object is not visible in the photographs, yet GS still generated points in the interior, reconstructing non-visible areas that do not correspond to actual observed surfaces (Figure 8); these spurious internal points were removed only through the filtering process. In contrast, NV-1405 has a visible empty inner volume in the images, which better constrains the optimisation and results in a smaller amount of spurious points.

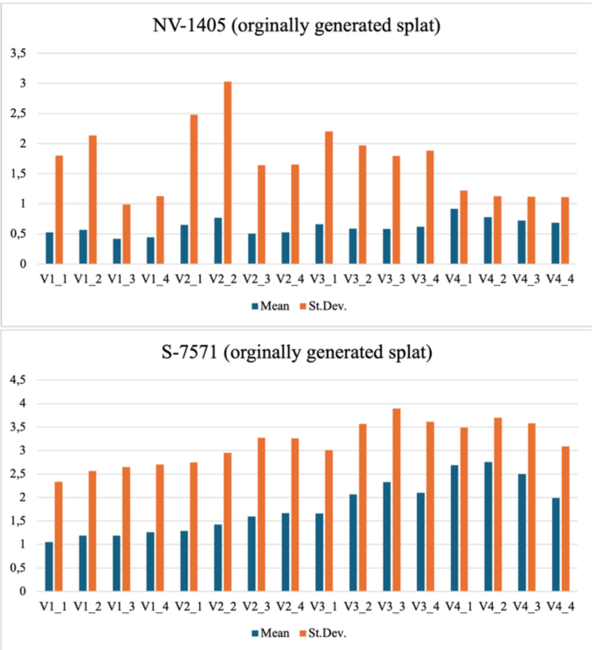

Figure 7. Comparison of the 16 originally generated GS configurations for NV-1405 (top) and S-7571 (bottom), showing mean cloud-to-mesh distances and corresponding standard deviations before filtering.

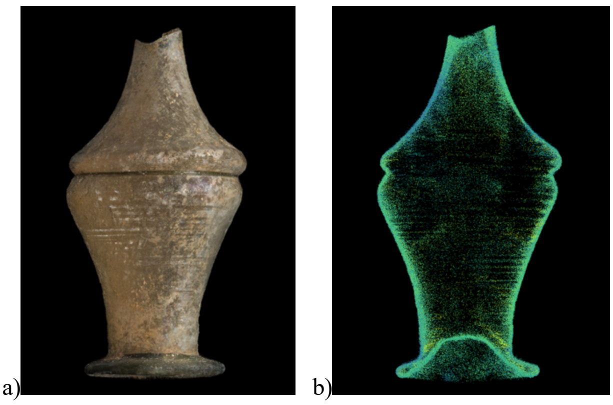
a) b)
Figure 8. (a) Front view of the GS model generated for S-7571 using the configuration with the highest training level and splat-count limit (V1_1). (b) Cross-section of the same model (5 mm depth), showing the presence of spurious internal points in non-observed regions.

As shown in the charts (Figure 9), both mean distances and standard deviations benefited from the filtering, indicating that a relevant portion of the cloud-to-mesh discrepancy was caused by GS artefacts rather than by errors on the actual reconstructed surface. Consistent with the considerations above, this effect was particularly strong for S-7571: the worst configuration (V4_2) improved from a mean distance of 2.755 mm (σ = 3.702 mm) to 0.798 mm (σ = 0.543 mm), while the best unfiltered case (V1_1) improved from 1.054 mm (σ = 2.339 mm) to 0.390 mm (σ = 0.442 mm). For NV-1405, the same behaviour was observed but to a lesser extent: V4_1 improved from 0.916 mm (σ = 1.221 mm) to 0.574 mm (σ = 0.455 mm), and V1_3 from 0.417 mm (σ = 0.992 mm) to 0.356 mm (σ = 0.269 mm).

The filtering had an even stronger effect on standard deviation than on mean distance, showing that its main contribution was the reduction of dispersed outliers and unstable GS artefacts. Before filtering, the average standard deviation was 3.15 mm for S-7571 and 1.70 mm for NV-1405, whereas after filtering it decreased to 0.49 mm and 0.32 mm, respectively. This indicates that a large part of the original variability was caused by spurious points located far from the reference surface. The effect is also evident in the reduced variability across configurations. For S-7571, the difference between the maximum and minimum standard deviation decreased from 1.56 mm before filtering to only 0.11 mm after filtering; for NV-1405, it decreased from 1.49 mm to 0.19 mm. This shows that the filtering not only improved individual results, but also made the behaviour of the different GS configurations much more consistent and comparable.

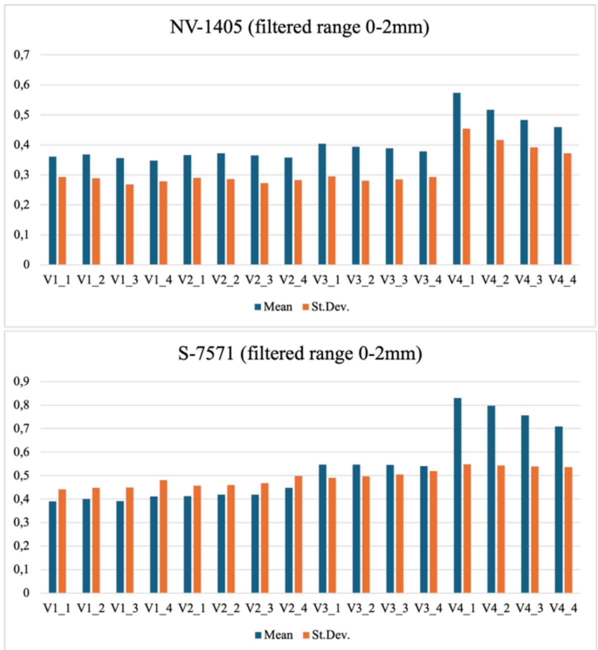

Figure 9. Comparison of the 16 originally generated GS configurations for NV-1405 (top) and S-7571 (bottom), showing mean cloud-to-mesh distances and corresponding standard deviations after filtering.

On the perceptual side, the comparison was performed by selecting one GS configuration for each training length, always preferring the version with the highest splat limit. For each of the two objects, four configurations were therefore considered, corresponding to 3, 10, 30 and 100 ksteps. The visual assessment was based on a direct comparison with original photographs from the acquisition campaign. In each case, the reference image was masked in order to isolate only the object and exclude the background. Then, using the estimated camera pose, a rendered image was generated from the corresponding GS model for the same viewpoint. This procedure was repeated for all four selected configurations and for both objects (Table 4 and 5). The similarity between each rendered image and its masked reference photograph was then evaluated through image-based metrics. In particular, MSE and PSNR were used to quantify pixel-level radiometric differences, SSIM and MS-SSIM to assess structural similarity, LPIPS to estimate perceptual similarity in a way closer to human visual judgement, and ΔE to measure colour difference. These texture-comparison metrics confirm the strong visual potential of GS for transparent artefacts, but also show that the visually best solution does not necessarily coincide with the metrically best one. In contrast with the geometric analysis, where longer training generally improved cloud-to-mesh agreement, the perceptual evaluation indicates that the best visual balance was reached at intermediate optimisation levels. Across the rendered comparisons, the 30-kstep configuration provided the most convincing overall results for both objects, combining lower MSE and LPIPS, higher PSNR, and consistently strong SSIM and MS-SSIM values. For example, it reached PSNR values of 6.77 dB and 5.45 dB, LPIPS values of 0.423 and 0.424, and ΔE values of 27.596 and 31.960 for NV-1405 and S-7571, respectively. Very short training (3 ksteps) consistently produced the weakest perceptual quality, with higher error, lower structural similarity and higher LPIPS and colour differences.

	Texture assessment method					
NV-1405	MSE↓	PSNR (dB)↑	SSIM↑	MS-SSIM↑	LPIPS↓	ΔE↓
v1_1	0,215	6,67	0,5	0,486	0,452	28,557
v2_1	**0,21**	**6,77**	0,557	0,535	**0,423**	**27,596**
v3_1	0,215	6,66	**0,588**	**0,568**	0,432	27,635
v4_1	0,242	6,16	0,522	0,521	0,518	29,418

Table 4: Metrics for the visual evaluation of the generated GS for the object NV-1405.

	Texture assessment method					
NV-1405	MSE↓	PSNR (dB)↑	SSIM↑	MS-SSIM↑	LPIPS↓	ΔE↓
v1_1	**0,285**	5,44	0,538	0,496	0,443	32,262
v2_1	**0,285**	**5,45**	**0,564**	**0,518**	**0,424**	**31,96**
v3_1	0,298	5,25	0,546	0,5	0,452	33,433
v4_1	0,347	4,59	0,469	0,419	0,526	39,265

Table 5: Metrics for the visual evaluation of the generated GS for the object S-7571.

5. Discussion and conclusions

The comparative analysis confirmed the strong potential of Gaussian Splatting for the digitisation and visualisation of transparent heritage artefacts, while also highlighting a clear trade-off between perceptual quality and geometric reliability. The results showed that the configurations producing the most convincing visual appearance do not necessarily correspond to those with the highest geometric accuracy. In particular, the texture-comparison metrics identified V2_1 as the most convincing configuration from a perceptual point of view (Figure 10).

(a) (b)

Figure 10. Original image from the digitisation campaign (a). Rendered view of the GS model generated with configuration V2_1, identified as the most convincing configuration on the basis of the texture assessment (b).

This analysis was intentionally carried out on the configurations with the highest splat budget, as these were expected to provide the greatest added value for visualisation purposes.

On the geometric side, the most reliable results were generally obtained with longer optimisation. The best geometric agreement with the reference photogrammetric meshes was reached at 100 ksteps, specifically in V1_3 for S-7571 and V1_4 for NV-1405. This behaviour is consistent with another trend observed in the tests: configurations with lower splat limits (300 and 1000 ksplats) more frequently reached the imposed maximum number of splats, suggesting that the training process was approaching its effective representational limit. At the same time, configurations with fewer training steps and higher splat limits often generated more splats than corresponding configurations trained for longer. This suggests that, in shorter training regimes, the pruning process may be less effective in removing floating and unnecessary splats, which can increase the number of artefacts without improving the actual reconstruction quality (Figure 11). The filtered geometric evaluation further supported this interpretation, showing that a relevant part of the discrepancy was caused by unstable or spurious splats, especially in poorly constrained or non-visible regions. Therefore, high splat counts alone should not be interpreted as an indicator of better reconstruction quality in transparent and reflective objects.

(a) (b)
(c) (d)

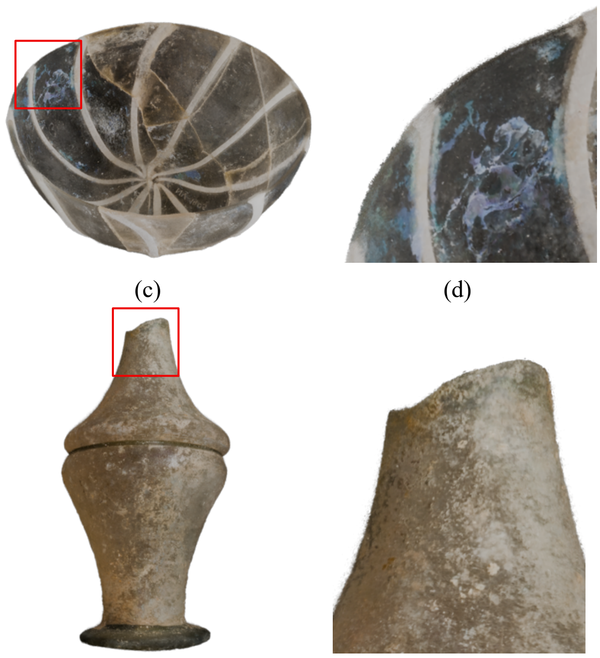

Figure 11. (a, c) Renderings of the GS models generated with configuration V3_1, which produced the highest splat count. Although the overall visual evaluation of these models is relatively good, the enlarged details (b, d) reveal the presence of spurious points and edge artefacts along the object edges.

Further investigations are needed to better understand the relationship between training length, splat budget, pruning behaviour and the characteristics of the input imagery. In particular, the influence of image size and resolution on the effective number of generated splats and on the final balance between visual fidelity and geometric robustness deserves deeper analysis. In addition, future developments should explore methods which derive an explicit mesh geometry from GS representations, so as to bridge appearance-oriented rendering and more conventional mesh-based representations suitable for geometric analysis and interoperability.

Acknowledgments

The presented research activities are developed as part of the EU project "XRCulture - eXplore & Reuse 3D cultural heritage within the Data Space" (GA 101174317), co-funded by the European Union under the Digital Europe Programme.

References

Aubreton, O., Erçil, A., Eren, G., Fofi, D., Mériaudeau, F., Sanchez Secades, L.A., Truchetet, F., 2009. 3D Scanner Based On Scanning From Heating. PCT application number: PCT/IB08/055328.

Chen, W., Li, Z., Guo, J., Zheng, C., & Tian, S., 2025. Trends and techniques in 3D reconstruction and rendering: A survey with emphasis on Gaussian Splatting. *Sensors*, *25*(12), 3626.

Dalal, A., Hagen, D., Robbersmyr, K.G., Knausgård, K.M., 2024. Gaussian splatting: 3D reconstruction and novel view synthesis: a review. *IEEE Access*, 12 (2024), pp. 96797-96820.

Fei, B., Xu, J., Zhang, R., Zhou, Q., Yang, W., & He, Y. (2024). 3d gaussian splatting as new era: A survey. *IEEE Transactions on Visualization and Computer Graphics*.

Ferraton, M., Stolz, C., & Mériaudeau, F., 2009. Optimization of a polarization imaging system for 3D measurements of transparent objects. *Optics Express*, 17(23), 21077-21082.

Huang, B., Yu, Z., Chen, A., Geiger, A., & Gao, S., 2024. 2D gaussian splatting for geometrically accurate radiance fields. Proc. *ACM SIGGRAPH*, pp. 1-11.

Jamil, O., Brennan, A., 2025. Immersive heritage through Gaussian splatting: a new visual aesthetic for reality capture. *Front. Comp. Sci.,* 7, p. 2025

Jiang, Y., Tu, J., Liu, Y., Gao, X., Long, X., Wang, W., & Ma, Y., 2024. Gaussianshader: 3D gaussian splatting with shading functions for reflective surfaces. Proc. *CVPR*, pp. 5322-5332.

Karami, A., Battisti, R., Menna, F., & Remondino, F., 2022. 3D digitization of transparent and glass surfaces: State of the art and analysis of some methods. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, *43*, 695-702.

Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph*. 42 (4), 139:1-139:14. doi.org/10.1145/3592433.

Li, M., Pang, P., Fan, H., Huang, H., & Yang, Y., 2025. TSGS: Improving Gaussian Splatting for Transparent Surface Reconstruction via Normal and De-lighting Priors. *Proc. ACM MM*.

Liu, C., Czuban, L., Bison, P., Grinzato, E., Marinetti, S., Maldague, X., 2006. Complex Surfaced objects: effects on phase and amplitude images in pulsed phase thermography. Proceedings *12th A-PCNDT.*

Mazzacca, G., Karami, A., Rigon, S., Farella, E. M., Trybala, P., and Remondino, F., 2023. NeRF for heritage 3D reconstruction. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.* XLVIII-M-2-2023, 1051-1058.

Menna, F., Nocerino, E., Remondino, F., Dellepiane, M., Callieri, M. and Scopigno, R., 2016. 3D Digitization of an Heritage Masterpiece - a Critical Analysis on Quality Assessment. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, Vol. XLI-B5, pp. 675-683.

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R., 2020. NeRF: representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 99-106.

Morel, O., Stolz, C., Meriaudeau, F., Gorria, P., 2006. Active Lighting Applied to 3D Reconstruction of Specular Metallic Surfaces by Polarization Imaging. *Appl. Opt.,* 45(17), 4062-4068.

Pelletier, J. F., Maldague, X., 1997. Shape from Heating: A Two-Dimensional Approach for Shape Extraction in Infrared Images. *Opt. Eng.,* 36(2), 370-375.

Stavroulakis, P. I., Ganetsos, T., & Zabulis, X., 2023. Large scale optical projection tomography without the use of refractive-index-matching liquid. *Sensors*, *23*(24), 9814.

Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P., Tretschk, E., et al., 2022. Advances in neural rendering. *Computer Graphics Forum*, SIGGRAPH courses, Wiley Online Library.

Tian, C., Hu, W., Zhang, Z., Wen, L., Ding, H., 2026. Geometry-aware Gaussian splatting of transparent object reconstruction. *Engineering Applications of Artificial Intelligence*. Vol. 167(1), article 113787.

Yan, Z., Mazzacca, G., Rigon, S., Farella, E. M., Trybala, P., Remondino, F., 2023. NERFBK: A holistic dataset for benchmarking NeRF-based 3D reconstruction. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.,* XLVIII-1/W3-2023, 219–226

Yang, Y., Chen, S., Wang, L., He, J., Wang, S. M., Sun, L., & Shao, C., 2019. Influence of coating spray on surface measurement using 3D optical scanning systems. In *International manufacturing science and engineering conference*, Vol. 58745, p. V001T02A009.

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-869-2026 |