A Hybrid Approach using Gaussian Splatting and Parametric Models based on 3D Renders for Real-Time Visualisation

Etienne Sommer[1], Arnadi Murtiyoso[1], Mathieu Koehl[1], Pierre Grussenmeyer[1]

[1]Université de Strasbourg, INSA Strasbourg, CNRS, ICube Laboratory UMR 7357, 67084 Strasbourg, France
(etienne.sommer, arnadi.murtiyoso, mathieu.koehl, pierre.grussenmeyer)@insa-strasbourg.fr

Keywords: Gaussian Splatting, 3D Reconstruction, Blender, Real-Time Visualisation, Cultural Heritage.

Abstract

This study investigates the use of synthetic images generated within Blender for reconstruction via 3D Gaussian Splatting (3DGS). These synthetic images are derived from a 4D parametric model of a Rhenish castle, incorporating its surroundings and distant environment. While such parametric models offer high-fidelity data, they are computationally intensive for real-time applications. 3DGS is therefore employed to produce high-quality visualisations from images with known spatial orientations. Two reconstruction methods are compared in this study: the open-source native code and the commercial Postshot solution with its Splat3 model.
The primary objective is to demonstrate the applicability of this method using synthetic imagery to create lightweight visualisations of digital twins of theoretical 4D states. The underlying parametric model, comprising numerous distinct objects and procedural textures, achieves high photorealism at the expense of substantial computational resources. Consequently, the reconstruction of this dataset via 3DGS facilitates the export and online dissemination of the complex model, decoupling visualisation quality from geometric complexity. The approach is quantitatively validated by comparing the 3DGS output against the original ground truth. Results demonstrate that the Splat3 model outperforms the native open-source approach in visual fidelity, processing speed, and geometric accuracy when handling high-resolution datasets. Both reconstruction methods achieve rendering performances well above 100 frames per second. This confirms that 3DGS can successfully be used with synthetic images to transform computationally heavy parametric models into highly optimised digital representations, ensuring near real-time visualisation suitable for immersive virtual reality and public dissemination.

1. Introduction

The importance of 3D modelling applied to built heritage is now well established. It provides an essential digital record for monitoring, conservation, and the creation of digital archives in the event of destruction (Mouaddib et al., 2024). Techniques to achieve this vary, ranging from conventional topographical surveys to HBIM modelling (Rocha et al., 2020). However, beyond conservation, the enhancement and dissemination of heritage to the public represent an equally crucial objective (Innocente et al., 2024). To achieve this objective, 3D models must not only be accurate but also interoperable and sufficiently optimised to enable easy sharing and real-time visualisation (Bolognesi and Manfredi, 2024). This is where traditional methods, which often generate dense and heavy meshes (Gonizzi Barsanti et al., 2022; Clini et al., 2025), show their limitations.

In response to this challenge, recent advances in neural rendering offer promising solutions in reducing computational time (Murtiyoso and Grussenmeyer, 2023). Following initial work on Neural Radiance Fields (NeRF) (Mildenhall et al., 2021), an explicit and more efficient approach, 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023), has emerged. This method represents the scene using 3D Gaussian primitives that are projected (or "splatted") directly onto the 2D image plane. This process enables significantly faster rendering speeds compared to Multi-View Stereo (MVS) or NeRF techniques. This efficiency makes 3DGS an ideal candidate for addressing the challenges of heritage visualisation and dissemination (Apopei, 2025). While the application of radiance fields in the heritage context has previously been explored (Croce et al., 2024), this paper evaluates the practical utility of 3DGS for creating high-performance visualisations of high-poly parametric models intended for the enhancement of cultural heritage. Synthetic images are employed rather than photographs, as the parametric model represents a 4D reconstruction of a castle that is now in ruins. A review of the literature regarding the employment of synthetic imagery reveals a scarcity of studies on this specific subject. Franczuk (2025) proposes the creation of synthetic images to establish ground truth for the validation of photogrammetric processes, whereas Frommholz (2019) utilise such data to train neural networks with perfectly labelled datasets.

Unlike these existing approaches, the images generated in this study correspond to a specific historical phase of the edifice which no longer exists (Koehl et al., 2025). This phase differs from the castle's current state, which would otherwise be documented using a traditional photographic acquisition campaign.

The original 3D Gaussian Splatting reconstruction approach was introduced by Kerbl et al. (2023). Following this highly effective initial method, numerous new models and enhancements have rapidly emerged (Kheradmand et al., 2024; LichtFeld Studio, 2025). Furthermore, Guédon et al. (2025) propose the utilisation of 3DGS specifically for surface reconstruction. In addition to these open-source, free-to-use proposals, commercial solutions have been developed in parallel, such as the Jawset Postshot software. This tool facilitates the utilisation of multiple 3DGS versions, including the "Splat3" model native to the software, for which the exact methodology is not available in a research publication.

For the purposes of this study, the native code provided by Kerbl et al. (2023) is utilised and compared against the "Splat3" model available with Postshot. A series of reconstructions is proposed, involving variations in processing parameters, image resolution and image count. The objective is to determine the optimal reconstruction strategy for synthetic datasets by comparing open-source and commercial models. The analysis is conducted based on qualitative and quantitative criteria, whilst also accounting for the processing time required for each reconstruction. For the subsequent stages of the study, all tests were performed on a workstation equipped with 128 GB of RAM and an NVIDIA Quadro RTX 6000 graphics card featuring 24 GB of VRAM.

2. Virtual Camera Configuration and Photogrammetric Pre-processing

3DGS requires two inputs: a set of images and a camera orientation file. The images employed in this study are generated synthetically by Blender, using a pre-existing parametric model to create synthetic renders. The Birkenfels castle, located in Ottrott around 30 kilometres from Strasbourg, Alsace, France, was used as a case study. The castle is thought to have been built in the XIII$^{\text{th}}$ century, then gradually abandoned in the early XVI$^{\text{th}}$ century. Its parametric 4D reconstruction is described in Sommer et al. (2025b). This 4D reconstruction uses an instancing system that allows an initial object to be referenced, duplicated, translated, scaled, or rotated without increasing the scene's polygon count. This instancing system is applicable to any type of object, provided that its geometry remains unmodified. The primary advantage of this system is that it enables the creation of complex and dense scenes while maintaining a low polygon count. On the other hand, a limitation of this instancing system is quickly encountered: as this system is specific to Blender, instances cannot be exported directly. Converting them into actual geometry generates an excessive number of polygons, which hinders external software compatibility and frequently causes unexpected software termination.

Furthermore, texturing presents comparable challenges. Although automated texture baking via Python is feasible, individually exporting every element with specific UV unwrapping and appropriate resolution scales creates severe file management issues and unnecessary data storage. Employing 3DGS addresses these limitations directly by translating the heavy parametric geometry and procedural textures into a volumetric Gaussian representation, facilitating efficient export and real-time rendering on external devices.

The images are rendered using Blender's native "Cycles" path-tracing engine, which is essential for supporting procedural textures. Generally, 3D renders aim for photorealism by simulating natural lighting, reflections, and cast shadows (Sommer et al., 2025a). However, as the objective here is to use these images for 3DGS training, these effects were deliberately removed. Instead, the lighting environment was configured to diffuse a uniform white light, and cast shadows originating from the scene geometry were disabled. The rendering process is automated using Blender's Python API. Each virtual camera is systematically positioned on the vertices of an icosphere generated around the object of interest. From each position, the camera is oriented towards the 3D model, and the render is generated. Figure 1 shows an example of a synthetic image generated with this workflow. The images were rendered with an iterative process, involving rapid preliminary reconstructions to identify adequately visible areas and those requiring additional acquisition angles.

Figure 1. Example of a rendering in Blender used for 3DGS training

The virtual camera can be treated as a standard physical device, defined by attributes such as sensor size and pixel count. It is therefore possible to calculate a Ground Sampling Distance (GSD) value utilising traditional photogrammetric methods (Kraus, 2007). To simulate a full-frame acquisition, the virtual camera sensor was defined with dimensions of 36×24 mm and a resolution of 7680×5120 pixels. This configuration yields a pixel pitch of approximately 4.69 µm (calculated as $36 \div 7680 = 24 \div 5120 \approx 4.69$), resulting in an average GSD of 4.5 mm.

For the orientation phase in Agisoft Metashape, the alignment accuracy was set to "High", as illustrated in figure 2. Given the exact a priori knowledge of the camera poses, the positioning accuracy was constrained to 0.00001 metres, the minimum threshold permitted by the software. Furthermore, as the images were generated using a "perfect" virtual lens, they are free from optical aberrations; consequently, no distortion correction was applied. The outcome of this orientation stage can be verified using the tie point reprojection error. In this specific case, the error is 0.38 pixels, confirming that synthetic images can be oriented using classical SfM algorithms originally designed for real-world imagery.

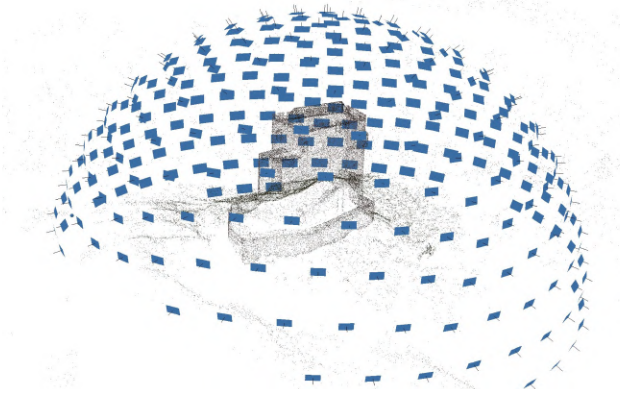

Figure 2. Orientation of the 3D renders in Metashape

3. Experimental Results on Native 3DGS

The first part of this study is based on the initial 3DGS solution published by Kerbl et al. (2023). The objective is to determine whether this original open-source solution can faithfully reconstruct a scene using synthetic renders derived from parametric modelling in Blender.

3.1 Training Scores Analysis

The native 3DGS code cannot process the full-resolution dataset due to RAM limitations. Consequently, downsampling is applied during execution, and several factors were tested to assess their impact on performance and rendering. Analysis of figure 3 shows that processing time is proportional to image resolution. A massive gap is observed: the R2 scenario (high resolution corresponding to downsampling by 2) requires approximately 180 minutes. This is almost 4 times longer than R8 (low resolution corresponding to downsampling by 8), which completes training in under 50 minutes.

Furthermore, the evaluation of image quality metrics shows a direct correlation between the downsampling factor and the resulting quantitative scores. As shown in the graphs, the configurations with lower image resolution achieve better numerical results. The SSIM and the PSNR are highest for the R8 scenario, reaching approximately 0.85 and 29 dB respectively by the end of the training. Conversely, the higher resolution R2 configuration yields the lowest scores. This trend is corroborated by the LPIPS metric, where a lower value indicates better perceptual agreement. The R8 method achieves the most favourable LPIPS score of roughly 0.20, whereas R2 produces a significantly higher value of 0.50. Finally, all three image quality metrics demonstrate a rapid phase of improvement during the first 10,000 to 15,000 iterations, after which the learning curves gradually stabilise as they approach the final limit of 30,000 iterations.

However, these metrics should be interpreted with caution. A high score at low resolution simply means the model successfully reproduced a lower-quality image. Although R2 shows lower scores, qualitative analysis is essential for 3DGS. It may reveal that this high-resolution setting is the only one capable of providing a sharp and detailed render, unlike the low-resolution versions which are mathematically better but visually not satisfactory.

3.2 Qualitative Analysis of the Reconstruction

As the primary benefit of 3DGS lies in visualisation, qualitative analysis is crucial for validating the method’s suitability. Overall, the reconstructions are consistent. There are no holes in the reconstruction, the colours are faithful to the original images, and no major artefacts appear in the model’s immediate surroundings. However, a detailed inspection of the textures, as illustrated in figure 4, reveals disparities that are masked by the statistical metrics. The R2 scenario is the only one that preserves high-frequency information. The joints between the stones remain sharp, and the grain of the material on the wooden defensive structures is clearly distinguishable. Conversely, reducing the resolution leads to progressive degradation. A smoothing effect appears with R4, which becomes significantly more pronounced with R8. In the latter case, although the PSNR is high, the visual rendering becomes unclear and noisy: the distinction between stones fades, and fine geometric edges lose their definition. On

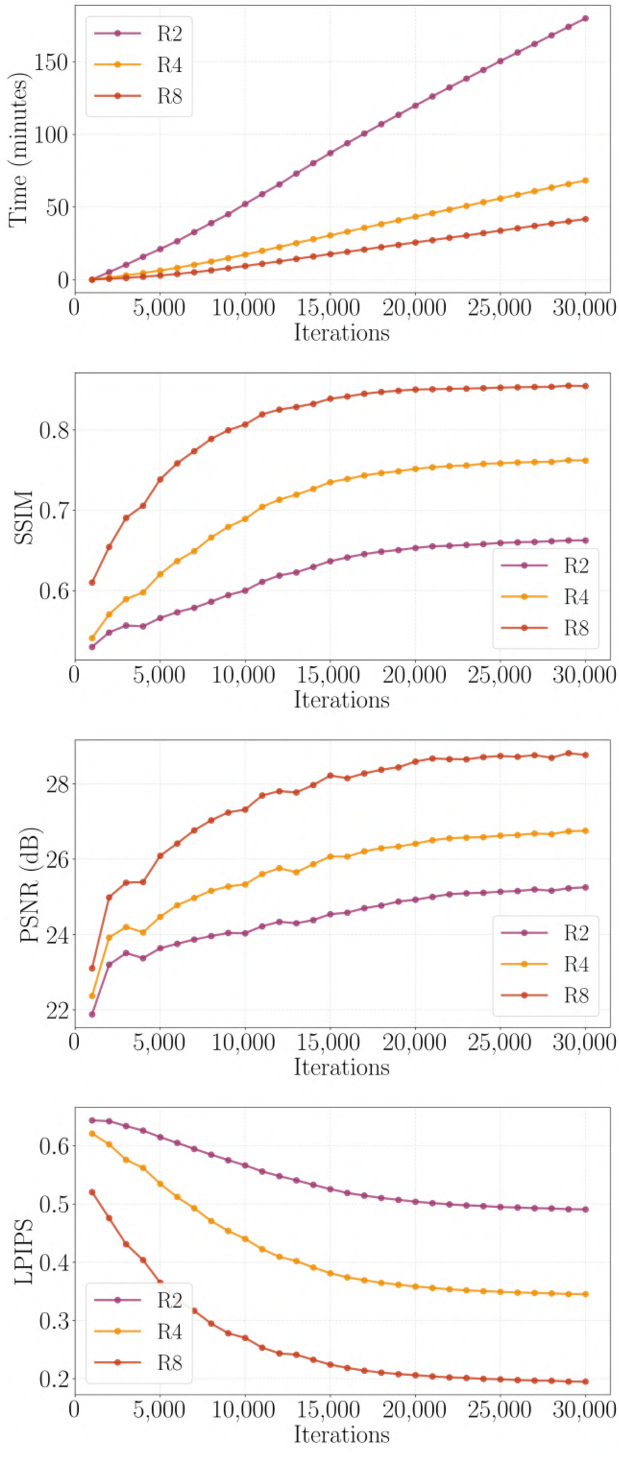

Figure 3. Native 3DGS training scores

the other hand, a singular behaviour is observed: specific areas remain blurry at high resolution, whereas they appear better resolved with smaller images. This suggests that visual convergence is more difficult with high-definition images, likely requiring more iterations.

3.3 Qualitative Analysis of the Environment

3DGS effectively handles complex geometries like vegetation, which heavily surrounds the main model in this study. Contrary to the architectural analysis, an inverse resolution trend shows that lower resolution yields better vegetation reconstruction.

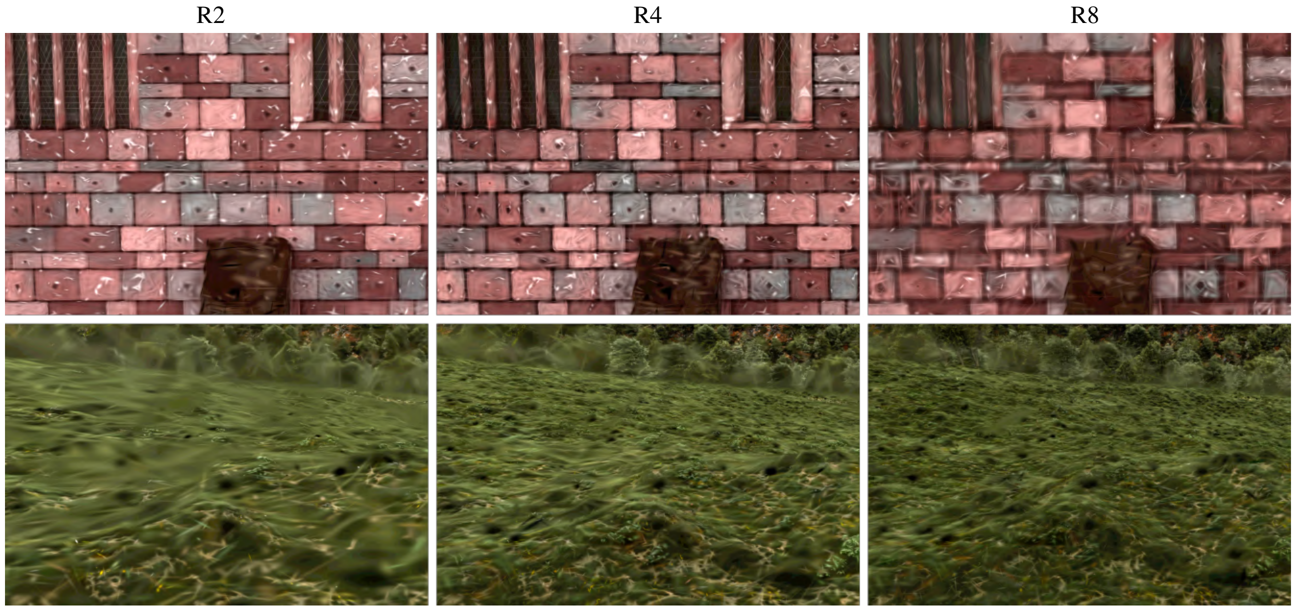

Figure 4. Qualitative analysis of rendering with native 3DGS code on the castle and its surroundings based on downsampling

As figure 4 illustrates, the high-resolution R2 produces noisy, blurred foreground grass and an unstructured background forest. Conversely, R8 generates a clearer overall environment with sharper, more coherent vegetation. This improvement likely occurs because downsampling acts as a low-pass filter, reducing the high-frequency complexity that hinders model convergence. Furthermore, decreasing the resolution expands the satisfactory reconstruction zone. This results in a better-structured distant forest unlike high-resolution models that focus on the scene’s centre. This difference in convergence capability likely explains the metric variations discussed in Section 3.1.

4. Experimental Results on Postshot using the Splat3 Model

In addition to existing open-source solutions, commercial solutions for 3DGS reconstruction have emerged. They offer the advantage of being easier to use, generally appearing as standalone software without the need to install specific dependencies or use command-line interfaces, although they come at a cost. Furthermore, the underlying reconstruction methods are not always published in scientific papers. Consequently, it is not always possible to explain why one result is better or worse than another. Nevertheless, their use is widespread, and they deserve to be evaluated alongside open-source solutions. For this study, the Jawset Postshot software is used, and the scene is reconstructed using the Splat3 model, which is specified by the developers as the highest performing option.

4.1 Performance and Resource Management

The Postshot software manages resources differently from certain open-source methods, enabling the processing of all 602 images at their initial 8K resolution. Subsequently, various downsampling tests are conducted following the same procedure as in Section 3 to ensure direct comparability. The analysis of the Splat3 training curves reveals that processing time increases linearly with the number of iterations and heavily depends on image resolution. Specifically, the R1 configuration

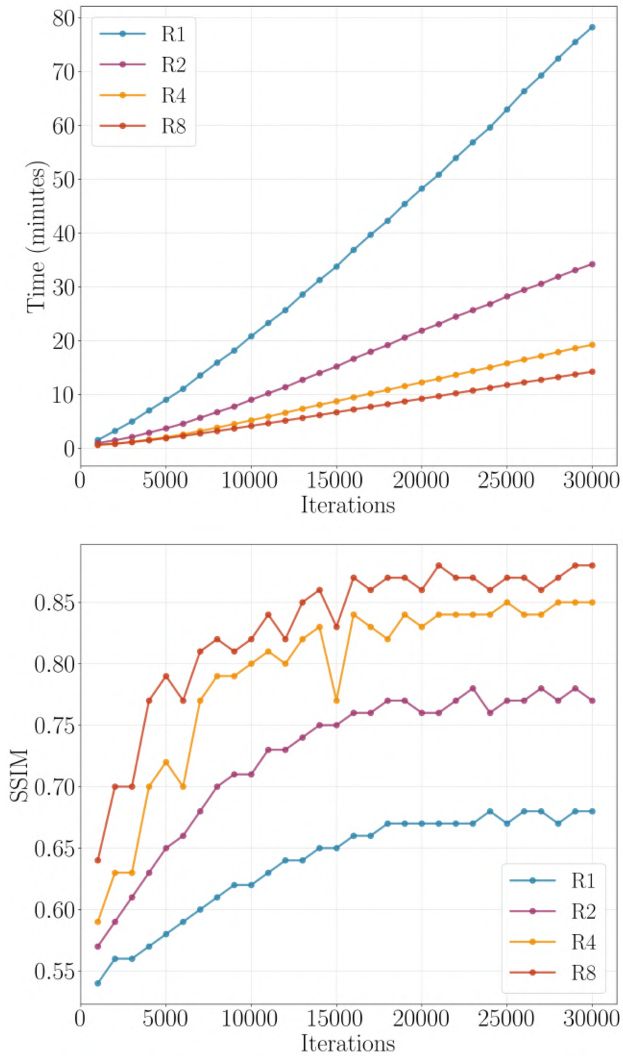

Figure 5. Splat3 training scores

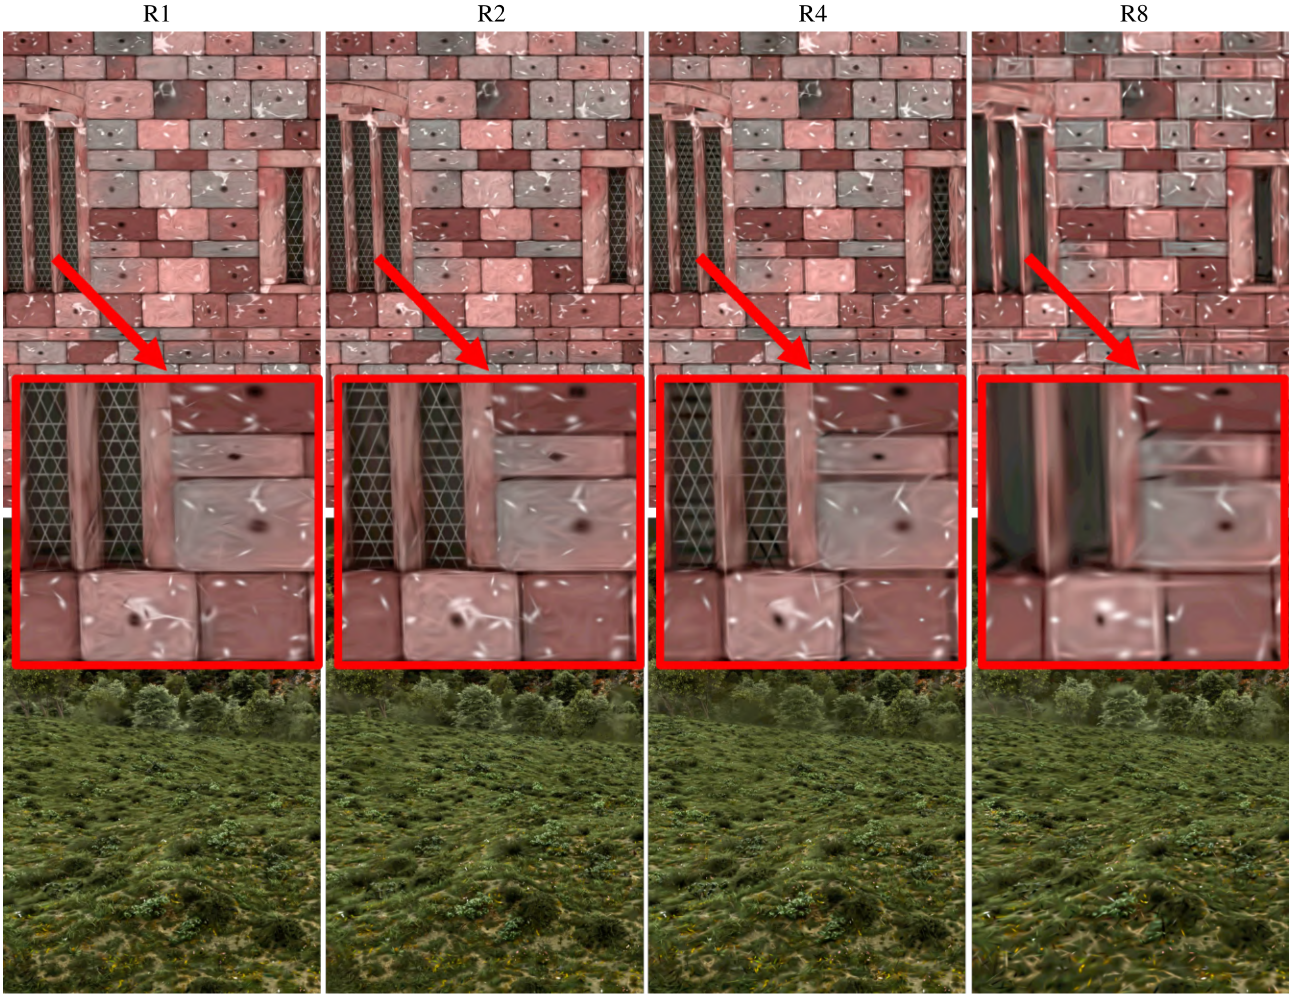

Figure 6. Qualitative analysis of rendering with Splat3 on the castle and its surroundings based on downsampling

requires nearly 80 minutes to train, whereas the R4 and R8 configurations complete in under 20 minutes. Regarding image quality, the SSIM exhibits a rapid initial increase before stabilising after 15,000 iterations. Higher downsampling factors consistently yield better metric scores, with R8 approaching an SSIM of 0.88. In contrast, the highest resolution R1 model achieves the lowest structural similarity with a score of 0.68.

4.2 Visual Fidelity Analysis

Similar to the analysis of the native code, the visual assessment of the Postshot reconstruction, as presented in figure 6, focuses first on the central main element, and subsequently on the vegetation elements constituting the surrounding environment.

Initially, it is observed that for an equivalent number of iterations, the reconstruction is superior to that of the native code. Fine details are better represented, and the overall reconstruction appears significantly sharper. This superiority of the Splat3 model is particularly evident in the management of high frequencies. When comparing the stonework reconstruction, the native code produces a smoothed rendering, erasing a little bit the roughness of the stone, whereas the Splat3 rendering restores the granularity and surface relief. However, the most discriminating element remains the reconstruction of fine structures: the bullseye glass windows (on the left of the image). With the native code, the reconstruction is a bit more blurred, indicating an inability to resolve this fine and complex feature. Conversely, Splat3 successfully represents this area. Similarly, the wooden structures benefit from this increased precision. In particular, the wood grain and the details of the small roofing are better defined in the Splat3 version. Furthermore, no blurred areas are present on the model, unlike with the native code.

This analysis applies to the structures linked to the main object, but also to the environment across all resolutions. The environmental reconstruction is also significantly better and extends further. It is also worth noting that, unlike the native code, better visual results are obtained here on vegetation when the image resolution is higher. Where the native code produced a blurred foreground and an indiscernible background forest at high resolution, Splat3 manages to structure the scene with high fidelity. Individual ground vegetation elements are clearly distinguished, and the forest edge in the background is well-defined, revealing the distinct geometry of each tree rather than a uniform textured mass.

Next, it is interesting to note that very few differences are observed between the full image resolution and R2. Visually, the details and reconstructions do not change. It is even observed that certain fine details are more visible in the R2 image, suggesting that the result may have converged more quickly given the same number of iterations. This nuance is perceptible on very high-frequency elements, such as the bullseye glass windows. On the full-resolution version, although definition is maximal, slight visual noise or excessive granularity is perceived on fine details, indicating that the model is still attempting to optimise micro-details pixel by pixel. Conversely, the R2 rendering appears smoother and more legible: the bullseye glass windows are sharper and more continuous, and the wood grain on the right-hand structure is equally precise. This suggests that the full resolution has not yet reached its convergence optimum after 30,000 iterations, whereas the R2 factor represents an ideal compromise, eliminating superfluous noise while preserving the essential structure.

5. Quantitative Metric Analysis of the 3DGS Reconstructions

The previous sections described reconstruction methods using the native code of 3DGS and the Splat3 model from the Postshot software. These two methods are compared in this section, focusing on the training score, on the geometric quality of the reconstructions, and on the graphical rendering performance of *.ply* files.

5.1 Comparison of the Training Scores

A comparative analysis of the training metrics, as illustrated for the native 3DGS code in figure 3 and for the Splat3 approach in figure 5, reveals distinct performance characteristics. Regarding processing speed, both methods exhibit a linear increase in training time relative to the number of iterations. However, the Splat3 approach demonstrates faster processing speeds than the native code across all comparable resolutions. For example, training the native model at a downsampling factor of 2 requires nearly 180 minutes to reach 30,000 iterations, whereas the Splat3 method completes the exact same configuration in approximately 35 minutes. Furthermore, Splat3 is capable of processing the full-resolution R1 dataset in under 80 minutes, which highlights superior computational efficiency. In terms of image quality, a comparison of the SSIM curves shows that the Splat3 method achieves higher final scores than the native code across all shared resolutions. For instance, at a downsampling factor of 2, Splat3 reaches an SSIM of approximately 0.77, compared to 0.66 for the native model. Nevertheless, the native method exhibits highly stable and smooth learning curves throughout the entire training process. Conversely, the Splat3 progression is characterised by significant fluctuations, indicating a less stable convergence.

5.2 Geometric Analysis of the Reconstruction

The geometric evaluation of the 3DGS reconstruction is performed using Gaussian centroids, which are transformed into a point cloud. This assessment is conducted against a portion of the parametric mesh converted into geometry for export, serving as ground truth. The entire model, including both the castle elements and the surrounding vegetation, cannot be directly exported as a mesh due to the instancing system employed. Instead, only a specific wall section is exported to facilitate geometric evaluation on a complex surface with relief. It has been demonstrated that 3DGS reconstructions are often noisy (Sommer et al., 2025c; Fadilah et al., 2026), although methods exist to mitigate this via Gaussian filtering to preserve the geometric structure (Guo et al., 2026). In this study, no noise removal filters were applied to evaluate the reconstruction.

To ensure a reliable assessment of the reconstructed area, it is necessary to isolate the element of interest along with its inherent noise while excluding points belonging to neighbouring objects. This selection is automated via a Geometry Nodes process within Blender, which filters the point cloud representing the Gaussians in two successive steps:

- Proximity filtering: The system first compares the position of each point relative to the reference mesh. A buffer zone is defined such that all Gaussians located at a distance greater than one metre from the mesh surface are automatically removed. This ensures that only data close to the object of interest is retained. This one-metre value was selected because the number of points situated beyond this distance is negligible and does not represent the actual reconstruction. Furthermore, these isolated points could be easily removed and therefore do not cause confusion.
- Vertical alignment: To refine the selection, the algorithm verifies altitude consistency along the Z-axis. The height of each remaining point is compared with that of the reference mesh. If a point is not vertically aligned with the model surface, effectively floating above or below it, it is removed. This ensures the elimination of points that do not correspond to the segmented mesh, which would otherwise bias the quantitative evaluation.

The quantitative analysis is firstly performed on the number of generated points relative to the image downsampling factor, as illustrated in figure 7. At high resolutions, specifically downsampling factors 1 and 2, the Splat3 approach produces a significantly higher number of points compared to the native method at downsampling factor 2. As the image resolution decreases at downsampling factors 4 and 8, the number of points generated by Splat3 drops sharply. In contrast, the native method maintains a higher number of points at these lower resolutions.

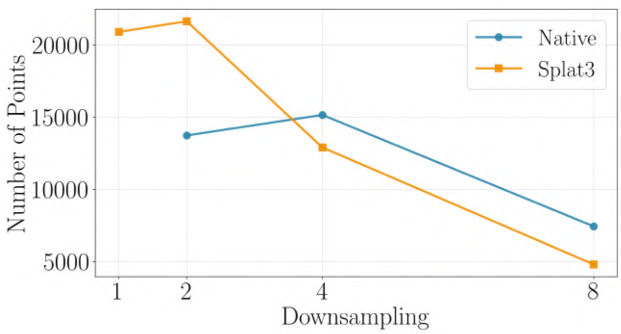

Figure 7. Density of Gaussians for native and Splat3 methods according to downsampling for the analysed sample

The implications of these point densities are clarified by assessing geometric accuracy and precision presented in figure 8. Regarding the mean deviation, the Splat3 model maintains an error close to zero across all resolutions. This indicates a reconstruction that is consistently faithful to the true reference geometry. Conversely, the native method exhibits a systematic negative bias that progressively gets worse as resolution decreases, dropping to approximately −0.06 metres at downsampling factor 8. This confirms the previously noted tendency of the native method to place points behind the actual surface due to the creation of internal Gaussians. Furthermore, an analysis of precision highlights the better stability of the Splat3 model. The standard deviation of the distances for Splat3 remains less than ±0.05 metres regardless of the downsampling factor. On the other hand, the native method displays a wider dispersion, with standard deviation values from ±0.10 metres at downsampling factor 2 to ±0.17 metres at downsampling factor 8. Consequently, the point cloud generated by Splat3 is less noisy and less dispersed.

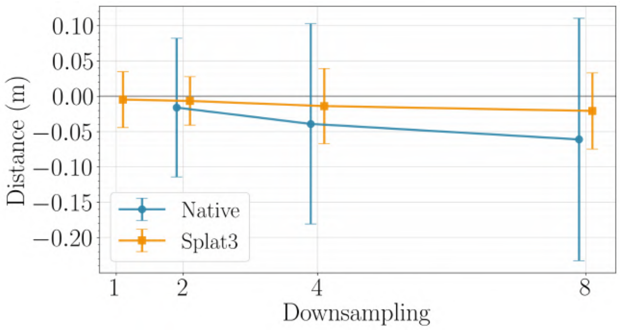

Figure 8. Geometric accuracy comparison of native and Splat3 models across downsampling factors for the analysed sample

5.3 Performance Analysis for Real-Time Rendering

A primary advantage of 3DGS is its real-time visualisation capability, which is essential for fluid viewing on platforms like SuperSplat and in virtual reality (Suwardhi et al., 2026). High frame rates are crucial in these contexts to ensure smooth navigation, user comfort, and to minimise motion sickness. To determine achievable frame rates, reconstructions are tested using an NVIDIA Quadro RTX 6000 GPU (24 GB VRAM) within the SIBR viewer (Bonopera et al., 2020). Performance is evaluated at a specific pixel resolution using interpolation mode, simulating user navigation via dynamic camera movements between training positions. Finally, fast culling and anti-aliasing are enabled, while V-Sync is disabled to capture true hardware performance limits without monitor refresh rate constraints.

Performance data acquisition is carried out using NVIDIA FrameView software over a duration of 120 seconds, ensuring that no other applications or background programs slow down performance. While FrameView captures a significant amount of information, two specific metrics are of particular interest for evaluating rendering performance: *MsBetweenDisplayChange* and *MsBetweenPresents*. The former, *MsBetweenDisplayChange*, measures the interval between displayed frames at the end of the graphics pipeline and indicates what the user actually sees on screen. However, the objective here is to evaluate performance directly related to the GPU, independent of the display device. Therefore, the analysis focuses on *MsBetweenPresents*, which captures data at the beginning of the graphics pipeline and indicates the smoothness of the animation delivered to the GPU. Figure 9 presents the varying rendering time performance across the different reconstructions. To ensure compatibility between the proprietary output of Postshot and the open-source SIBR viewer, a specific data formatting step is required. The *.ply* files generated by Postshot lack certain attributes strictly expected by the original Inria viewer implementation. Consequently, a Python script designed to bridge this format gap is employed to inject zero-value normal vectors (nx, ny, nz) into the *.ply* file. While these normals are not utilized for the splatting rendering itself, their presence in the file structure is mandatory for the viewer to load the scene without errors.

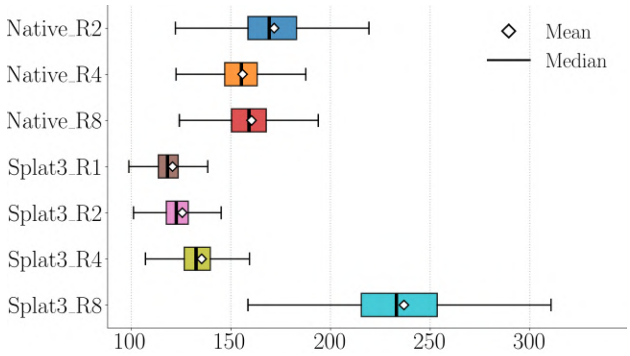

Figure 9. Rendering performance comparison in frames per second (FPS) for the native and Splat3 models across various downsampling factors

Rendering performance reveals distinct behaviours between the methods. The native model maintains stable speeds of 155 to 175 FPS across all downsampling factors, meaning reduced training resolution does not significantly alter the computational load. Conversely, Splat3 depends heavily on training resolution. At higher resolutions (R1, R2, R4), its frame rates average 120 to 135 FPS, slightly lower than the native method. This is due to higher point cloud density: Splat3 generates around 3 million splats (706 MB files) compared to the native R2 model's 1 million splats (243 MB files), requiring more GPU power. However, for the downsampling factor 8, Splat3 reaches 240 FPS. This increase correlates with a drop in generated Gaussians, which lightens the workload. Regardless of the configuration, all tests consistently exceed 100 FPS. Both models therefore satisfy the requirements for smooth, real-time visualisation in demanding applications such as virtual reality.

6. Conclusion

This study demonstrates the effectiveness of 3D Gaussian Splatting for the promotion of cultural heritage. By utilising synthetic renders generated from a 4D parametric model, the limitations of exporting dense meshes and procedural textures are overcome. The translation of this heavy geometry into a 3DGS model facilitates highly detailed and easily exportable 3D reconstructions. A comparative analysis between the native open-source code and the commercial Splat3 model reveals distinct advantages for the latter. Splat3 provides superior visual fidelity, faster processing speeds, better geometric accuracy, and the ability to effectively manage full-resolution 8K datasets. Conversely, the native method struggles with high-frequency details and internal geometric noise. Both reconstruction approaches deliver rendering performances significantly exceeding 100 frames per second. This confirms that 3DGS successfully bridges the gap between computationally heavy parametric modelling and the necessity for interoperable, real-time visualisation. Finally, this methodology yields a highly optimised digital model, ensuring fluid navigation and rendering it perfectly suited for interactive public dissemination and immersive virtual reality applications.

Acknowledgement

This work is supported by the ”Châteaux Rhénans: Burgen am Oberrhein” Interreg VI Project (2023-2025), coordinated by the European Collectivity of Alsace, and the Digital Twins for Cultural Heritage (Twin4CH) Chair, funded by the French National Agency for Research (ANR).

References

Apopei, A. I., 2025. 3D Gaussian Splatting in Geosciences: A Novel High-Fidelity Approach for Digitizing Geoheritage from Minerals to Immersive Virtual Tours. *Geosciences*, 15(10).

Bolognesi, C. M., Manfredi, V., 2024. Optimization of VR Application in Texturing Cultural Heritage. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-2/W4-2024, 73–78.

Bonopera, S., Hedman, P., Esnault, J., Prakash, S., Rodriguez, S., Thonat, T., Benadel, M., Chaurasia, G., Philip, J., Drettakis, G., 2020. SIBR: A system for image based rendering.

Clini, P., Angeloni, R., D’Alessio, M., Coppetta, L., Galli, I., 2025. Digital Representation and AI-driven Virtual Experience for Historic Houses. The Case Study of Borgo Storico Seghetti Panichi. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-M-9-2025, 337–344.

Croce, V., Forleo, G., Billi, D., Bevilacqua, M. G., Piemonte, A., Caroti, G., 2024. Neural Radiance Fields (NeRF) for Multi-Scale 3D Modeling of Cultural Heritage Artifacts. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-2/W4-2024, 165–171.

Fadilah, W. A., Murtiyoso, A., Landes, T., Grussenmeyer, P., 2026. Metric Assessment of 3D Gaussian Splatting for UAV-based Urban Heritage Reconstruction. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-2/W12-2026, 143–150.

Franczuk, J., 2025. Information Processes in Virtual 3D Reconstruction of Roman Three-Bay Double Arch of Musti (Tunisia). *Virtual Archaeology Review*, 16(32), 1–16.

Frommholz, D., 2019. A Synthetic 3d Scene for the Validation of Photogrammetric Algorithms. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W13, 1221–1228.

Gonizzi Barsanti, S., Guagliano, M., Rossi, A., 2022. 3D Reality-Based Survey and Retopology for Structural Analysis of Cultural Heritage. *Sensors*, 22(9593).

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. *ACM Transactions on Graphics*.

Guo, Z., Su, J., Qian, C., Wang, S., Fan, J., Zhang, J., Zhou, W., Amirpour, H., Zhao, Y., Han, L., Wang, P., 2026. GP-GS: Gaussian Processes Densification for 3D Gaussian Splatting. _eprint: 2502.02283. [Online] Available at https://arxiv.org/abs/2502.02283.

Innocente, C., Nonis, F., Lo Faro, A., Ruggieri, R., Ulrich, L., Vezzetti, E., 2024. A Metaverse Platform for Preserving and Promoting Intangible Cultural Heritage. *Applied Sciences*, 14(8).

Kerbl, B., Kopanas, G., Leimkuehler, T., Drettakis, G., 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.*, 42(4).

Kheradmand, S., Rebain, D., Sharma, G., Sun, W., Tseng, Y.-C., Isack, H., Kar, A., Tagliasacchi, A., Yi, K. M., 2024. 3D Gaussian Splatting as Markov Chain Monte Carlo. *Advances in Neural Information Processing Systems (NeurIPS)*.

Koehl, M., Heitz, J.-E., Koch, J., Sommer, E., 2025. Managing 4D Modelling Supported by Modern Archaeology for the In-Depth Study of a Ruined Castle – Oedenbourg Castle in Alsace, France. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-M-9-2025, 749–756.

Kraus, K., 2007. *Photogrammetry: Geometry from Images and Laser Scans*. De Gruyter Textbook, 2, Walter De Gruyter.

LichtFeld Studio, 2025. A high-performance C++ and CUDA implementation of 3D gaussian splatting.

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., Ng, R., 2021. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *Communications of The Acm*, 65(1), 99–106.

Mouaddib, E. M., Pamart, A., Pierrot-Deseilligny, M., Girardeau-Montaut, D., 2024. 2D/3D Data Fusion for the Comparative Analysis of the Vaults of Notre-Dame de Paris before and after the Fire. *Journal of Cultural Heritage*, 65, 221–231.

Murtiyoso, A., Grussenmeyer, P., 2023. Initial Assessment on the Use of State-of-the-Art NeRF Neural Network 3D Reconstruction for Heritage Documentation. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-M-2-2023, 1113–1118.

Rocha, G., Mateus, L., Fernández, J., Ferreira, V., 2020. A Scan-to-BIM Methodology Applied to Heritage Buildings. *Heritage*, 3(1), 47–67.

Sommer, E., Koehl, M., Grussenmeyer, P., 2025a. Comparative Study of Path Tracing Rendering Parameters and Visual Quality in Blender. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-M-9-2025, 1387–1394.

Sommer, E., Koehl, M., Grussenmeyer, P., 2025b. Parametric Modelling Techniques for Rhine Castle Models in Blender. *Heritage*, 8(1).

Sommer, E., Murtiyoso, A., Koehl, M., Grussenmeyer, P., 2025c. Radiance fields for archaeological documentation of medieval rhine valley castle ruins. *IEEE International Conference on Imaging Systems and Techniques (IST)*, Strasbourg, France, 1–6.

Suwardhi, D., Agustina, L. K., Andika, W., Widyastuti, R., Fadilah, W. A., Murtiyoso, A., Remondino, F., 2026. Integrating Terrestrial Laser Scanning and 3D Gaussian Splatting for Heritage Building Documentation and Monitoring. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-2/W12-2026, 463–470.