



HAL
open science

ROI-GS: Interest-based Local Quality 3D Gaussian Splatting

Quoc Anh Bui, Gilles Rougeron, Géraldine Morin, Simone Gasparini

► **To cite this version:**

Quoc Anh Bui, Gilles Rougeron, Géraldine Morin, Simone Gasparini. ROI-GS: Interest-based Local Quality 3D Gaussian Splatting. Proceedings of the International Conference on Visual Communications and Image Processing (VCIP), Dec 2025, Klagenfurt, Austria. <10.1109/VCIP67698.2025.11396827>. <hal-05489655>

HAL Id: hal-05489655

<https://hal.science/hal-05489655v1>

Submitted on 2 Feb 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Copyright - All rights reserved

ROI-GS: Interest-based Local Quality 3D Gaussian Splatting

Quoc-Anh Bui^{1,2}, Gilles Rougeron¹, Géraldine Morin², and Simone Gasparini²

¹Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

²Université de Toulouse, Toulouse INP – IRT, France

Abstract—We tackle the challenge of efficiently reconstructing 3D scenes with high detail on objects of interest. Existing 3D Gaussian Splatting (3DGS) methods allocate resources uniformly across the scene, limiting fine detail to Regions Of Interest (ROIs) and leading to inflated model size. We propose ROI-GS, an object-aware framework that enhances local details through object-guided camera selection, targeted Object training, and seamless integration of high-fidelity object of interest reconstructions into the global scene. Our method prioritizes higher resolution details on chosen objects while maintaining real-time performance. Experiments show that ROI-GS significantly improves local quality (up to 2.96 dB PSNR), while reducing overall model size by $\approx 17\%$ of baseline and achieving faster training for a scene with a single object of interest, outperforming existing methods.

Index Terms—Cultural heritage, 3D Gaussian Splatting, level of detail, scene composition

I. INTRODUCTION

Reconstructing complex 3D scenes with high-quality details for objects of interest remains a long-standing goal in computer vision and graphics. Like human vision, which naturally focuses on salient objects or depth-varying areas while leaving peripheral regions less resolved, this selective fidelity enables efficient perception. Similarly, computational methods benefit from concentrating resources on semantically or geometrically important regions. This capability is particularly vital in industrial applications and cultural heritage contexts. Here, our work aims to create 3D digital twins of cultural heritage sites, facilitating preservation [1], virtual exploration [2], and interactive museum experiences [3], thus supporting archiving, education, and public participation while minimizing physical handling and degradation risks.

Recent advances in 3D Gaussian Splatting (3DGS) [4] have rapidly gained traction as a powerful alternative for neural scene representation, offering explicit scene models, real-time rendering, and state-of-the-art visual quality. However, its explicit nature comes at the cost of a large memory footprint that can become prohibitive in large-scale scenes that require high detail in the internal contents of interest, as reported in [4–6]. This challenge is exacerbated by uniform scene reconstruction in most existing 3DGS methods, which allocates equal resources to all regions regardless of their importance. Prior work on level of detail (LOD) Gaussian Splatting [6–14] or scene chunking [15, 16] improves overall quality and scalability, but does not account for the presence




Fig. 1. A comparison of our proposed method against the classic 3D Gaussian Splatting baseline shows that our method improved object detail with minimal increase in overall scene memory.

of content of interest within the scene. Efforts in object-centric [17] and segmentation-aware [18–20] modeling focus capacity on individual scene elements. While these methods enable precise object handling, they require coherent object masks across images and do not adaptively select or schedule views based on objects of interest.

We address these gaps with **ROI-GS**, an object-aware 3DGS framework that reconstructs user-specified objects of interest using optimized image subsets and reintegrates them into the overall scene model. Our pipeline includes a scene decomposition module based on object-specific automatic camera selection, enabling independent training of scene components. By focusing computational resources and details on user-selected ROIs, the approach achieves fine-grained ROI optimization without significantly increasing the overall memory footprint. Leveraging the explicit nature of 3DGS representation, we seamlessly replace lower-LOD Gaussians in the base scene with refined ROI Gaussians, producing a unified model that combines global context at moderate quality with high-quality local regions, all rendered in real-time.

The main contributions of this work are: (i) we introduce region-aware reconstruction into the 3DGS paradigm through object-guided view selection and targeted Object training; (ii) we improve reconstruction quality using an advanced image selection strategy based on model optimization that maximizes ROI coverage and occupancy; and (iii) we propose a simple yet effective composition strategy that integrates high-fidelity Object Gaussians into the overall scene model. Together, these contributions enable ROI-GS to enhance object-level detail

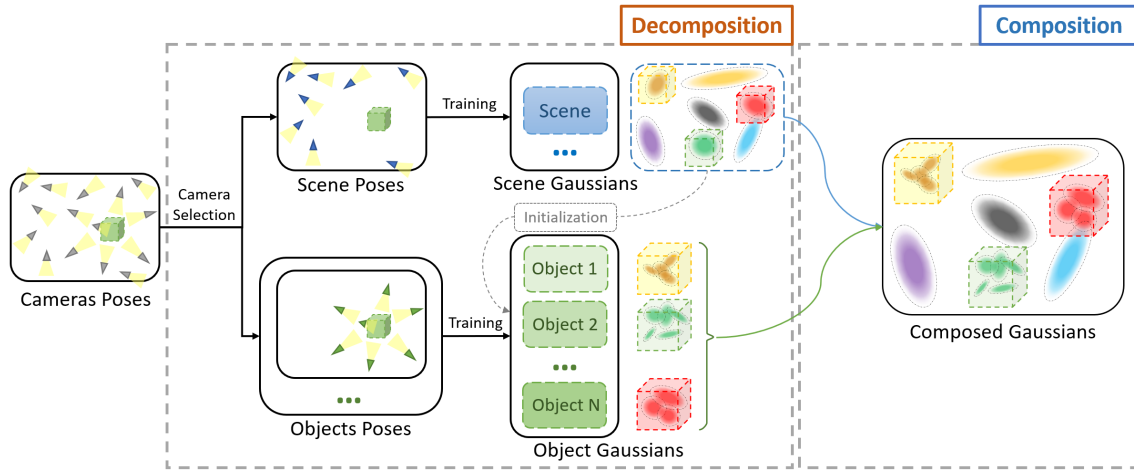


Fig. 2. An overview of the proposed Region of Interest-Focused Gaussian Splatting (ROI-GS) framework, consisting of two stages: scene decomposition and composition. In decomposition, the scene is divided into Scene and Objects groups, with camera sets automatically selected for each GS training. The trained Scene-GS model is used to initialize the Object-GS models. The composition stage integrates high-detail Object-GSs and the global Scene-GS to produce high-quality real-time renderings with enhanced detail for objects in the ROIs.

in complex scenes without sacrificing real-time rendering performance under a practical number of ROIs.

II. METHOD

Our ROI-GS pipeline builds on the standard 3DGS framework, which models the scene as a set of anisotropic 3D Gaussians. It takes as input a set of 2D color images, along with camera poses, intrinsic parameters, and a sparse point cloud, producing highly realistic real-time renderings with enhanced LOD in regions of interest (*cf.* Figure 2).

A. Object-Focused Camera Selection.

We propose a two-step ROI-focused camera selection strategy. First, each user-specified ROI is defined by an axis-aligned bounding box (AABB) enclosing the object of interest. In the initial filtering step, we select all views that observe the object by identifying cameras that capture at least one 3D keypoint within the bounding box, based on point-to-image visibility from SfM data. This first filtering significantly narrows down the set of candidate images. However, using all these views to train an Object-GS may degrade object detail due to low resolution from distant cameras and also reduce computational efficiency. To address this, we introduce a refinement step that assesses and selects the most informative and optimal views for high-quality object reconstruction. We initially used a simple approach based on static criteria, such as the distance between the ROI and viewpoints, the projected area of the ROI’s bounding box on 2D images, or the number of reconstructed keypoints within the box and visible to each camera. While effective in identifying views rich in ROI information, these criteria could introduce spatial bias, favoring regions with dense camera clusters and resulting in incomplete or unbalanced coverage from diverse perspectives. To mitigate this, we adopt an **advanced selection** strategy based on model-driven optimization. Prior works on View selection in GS [21–24] reduce image redundancy and maximize

scene coverage by selecting globally informative views, but they do not account for object-specific camera relevance. Our advanced selection approach relies on ActiveInitSplat [22], with adaptations for object-focused optimization. The views are ordered based on the ROI’s 3D point quality score: the next view is selected as the one that maximizes the improvement of this score. Specifically, this refinement selection module considers two terms of the sparse point distribution: the point density and voxel occupancy within the ROI bounding box to identify any sparsely covered areas or angles. A Gaussian Process-based (GP) model is used to guide this selection. Furthermore, by incorporating the static criteria mentioned above as parameters, we achieve more accurate and consistent camera selection for the target object, leading to better-aligned Gaussian optimization and ultimately higher-fidelity ROI reconstructions. The original model uses 6 input viewpoint parameters (3 for position and 3 for orientation). Our extended version adds 3 static parameters (distance, projected AABB area, and number of keypoints), totaling 9 inputs. We compare the 6-params and 9-params versions in Section III-C.

B. GS Training

Once the relevant images are selected and grouped, we train separate high-quality Gaussian Splatting models for each object (Object-GS) and one for the global scene (Scene-GS). The Scene-GS model is first trained for 20K iterations using all scene images, excluding the subset of close-up views reserved for ROI training. To maintain consistency, a portion of the ROI images can optionally be retained in the Scene-GS training to give the global model a coarse understanding of the object’s geometry. In our experiments, 50% of the ROI selected images are included in Scene-GS training.

In early experiments, we attempted to initialize Object-GS from the sparse SfM point cloud. However, neglecting the global representation of the trained Scene-GS model proved

inefficient. The next implementation idea was to initialize only Scene-GS Gaussians restricted to the ROI’s volume, and continue optimizing only within this region. However, the lack of out-of-the-box context led to floating artifacts, as the model hallucinated missing geometry beyond the object of interest, significantly degrading reconstruction quality. To overcome this, we initialize each Object-GS directly from the full Scene-GS model. This approach provides both contextual grounding and consistently yields higher-quality object reconstructions, making it a more effective and efficient starting point for ROI refinement. Each Object-GS is then trained independently using only the subset of selected images focused on the target object. The object training follows standard 3DGS optimization for 30K iterations to ensure sufficient reconstruction quality. Optimization is applied to the entire image, updating Gaussians even in the vicinity of the object box and the background, to prevent floating artifacts within the ROI. However, the Gaussian densification process is confined to the ROI volume for 15K iterations, keeping detail enhancement limited within the target region. Additionally, a pruning strategy is applied during training to remove non-contributing Gaussians, many of which were inherited from the full Scene-GS initialisation, thereby reducing the memory footprint of the Object-GS model. Focusing on a bounded region with high-resolution inputs allows for significantly finer LOD in this region than the global model. Object-GS can thus capture high-fidelity representations, including geometric details and texture appearance that could otherwise be blurred or undersampled in a uniformly trained entire-scene model.

C. Scene–Objects Composition with Gaussians.

For the inference step, we integrate the Object-GS models back into the global Scene-GS. This composition is simple and efficient thanks to the explicit nature of 3DGS scene representation. For each object, we replace the original Scene-GS Gaussians that lie inside the object’s bounding box with the corresponding fine details Object-GS ones. Since all Gaussians share a common coordinate system from SfM, the replacement is seamless and co-registered. The result is a combined Gaussian set representing the entire scene, where each ROI is now filled with its high-quality Gaussians.

III. EXPERIMENTS

In this Section, we discuss the implementation details of our ROI-GS method and present its performance in enhancing object-level reconstruction in a complex real-world scene.

A. Experimental setup.

We conduct experiments on two scenes, a *Bureau* and a *Dining room*, captured at the 18th-century cultural heritage site, the *Hôtel de la Marine* in Paris [25]. The *Bureau* capture consists of 761 images at 8192×5464 pixels, covering the working room, including objects on a desk, such as a coffret and statuettes. The *Dining Room* capture includes 916 images at the same resolution, plus 222 close-up 3024×4032 images, featuring interesting objects like ornate bowls, vases, and

TABLE I
EVALUATION OF OUR METHOD COMPARED TO BASELINES ON A ROI.

Methods Metrics	PSNR \uparrow	SSIM \uparrow	#G \downarrow	#G in box	Train
<i>Baselines: 50K iters</i>					
Full Scene images - 3DGS (1)	21.23	0.917	3.72M	8.73K	98m
Full Object images (2)	21.57	0.922	3.28M	16.12K	94m
<i>Ours</i>					
Scene-GS: 20K iters (3)	20.78	0.911	3.02M	6.99K	35m
Object-GS: 30K iters (4)	21.99	0.924	1.04M	74.74K	35m
Composition (3)+(4)	21.76	0.924	3.09M	74.74K	70m

chairs. All images were downsampled by a factor of 4 to reduce storage and computation costs. Our implementation is based on the 3DGS framework [4]. All trainings are performed using a single 95 GB H100 GPU on a DGX workstation.

For fair comparison and objective performance evaluation, we consider a single object of interest in Table I and Table III. Nevertheless, the method is inherently scalable: Table II experiments on multiple objects, and different objects are handled independently. Out of 335 images with at least one visible keypoint in the coffret AABB, 21 are used as a test set to evaluate the model. The baseline models used for comparison are: (1) Full Scene (3DGS), which is trained on all images in the scene, and (2) Full Object, which is trained on all images where the object of interest is visible. Both models are optimized for 50K iterations to ensure maximizing the scene reconstruction quality.

We evaluated the Object-Focused Camera Selection technique on two versions of the Advanced GP-based model. In our experiments, for ROI training we select 150 from 314 non-test ROI-visible images. Selecting more images slightly degraded reconstruction quality, as distant or less informative views reduced effective resolution. Both the original GP-based model (Advanced 6-params GP) and its extended version (Advanced 9-params GP), described in Section II-A, select the first 150 cameras while preserving their order. Unlike the default 3DGS pipeline, where camera order is irrelevant, the advanced selection strategy defines an order for selected views. During training, we retain this order instead of applying random shuffling, which helps the GS model converge faster and achieve higher reconstruction quality as shown in Table III.

B. Quantitative and qualitative evaluation.

We use standard metrics such as PSNR and SSIM for consistent and meaningful comparisons. Since our concern is the quality improvement of the object representation, evaluating the entire image would be less relevant. Therefore, scores are computed only on the pixels within the projected AABB.

TABLE II
EVALUATION OF OUR METHOD ON DIFFERENT OBJECTS OF 2 SCENES.

Dataset Scene	Bureau				Dining room		
	Coffret	Statue	Chair1	Sofa	Vase	Bowl	Chair2
3DGS baseline (1)	21.23	21.88	23.55	23.94	23.65	25.50	26.90
Our composition	21.76	22.31	24.01	24.19	24.81	26.30	27.50
Peak Improvement	0.93	0.93	1.10	0.80	2.96	2.13	2.63

Table I presents quantitative comparisons between our method and baselines on a single ROI. The proposed compo-




Fig. 3. A visual comparison of our method and the baselines. Blurred area are indicated by red arrow.

sition approach, which combines Scene-GS (3) and Object-GS (4), achieves a higher ROI score than all baselines. The table also lists the total number of Gaussians (#G) and those within the ROI box. Table II reports PSNR scores for different objects across the two scenes, highlighting the superiority, flexibility, and scalability of our method. Quality clearly improves across all objects, with a peak improvement per test image up to 2.96 dB for the ROI Vase in the *Dining Room*.

Figure 3 presents qualitative evaluations on three different ROIs. The proposed method yields noticeably sharper object details than both the baseline and Scene-GS. Rows (a) and (b) illustrate the precise reconstruction of complex geometry, while row (c) captures the vase’s intricate surface texture. In Object-GS (4), optimization is applied not only to Gaussians inside the box but also to outer ones, achieving the highest PSNR within the AABB (red box in Table I). This highlights the strength of the ROI reconstruction in capturing objects fine details but limiting scene coverage, leading to reduced overall scene quality (red arrow in Figure 3b).

C. Ablations.

TABLE III
PSNR AND SSIM SCORES FOR ABLATION STUDIES ON A SINGLE ROI.

Methods Metrics	PSNR \uparrow	SSIM \uparrow
Ours	21.76	0.924
No Scene-GS Init + 9-params GP (1)	21.65	0.919
No 3 static ROI params (6-params GP) (2)	21.20	0.908
No Advanced GP-based model (3)	19.18	0.877

We evaluate the impact of Advanced Camera Selection and Scene-GS initialization for object training in Table III. In ablation setting (1): we use SfM point cloud initialization instead of trained Gaussians from Scene-GS, combined with Advanced 9-parameters GP camera selection. This results in a slight drop in ROI reconstruction quality, compared to the full method. Settings (2) and (3) both use Scene-GS initialization. (2) employs the original GP-based model without the three

static ROI information, while (3) uses the simple selection based on static criteria only instead of the Advanced GP-based selection. The results highlight the effectiveness of the Advanced Selection approaches using GP model-based optimization. Moreover, incorporating additional object-specific information further improves the GP model.

IV. CONCLUSION AND PERSPECTIVES

We presented ROI-GS, a 3DGS method that enhances the level of detail for objects of interest by leveraging camera selection and a simple composition strategy. Our approach highlights its flexibility in enhancing object fidelity when applied to any global optimization model without compromising surrounding scene quality. While ROI-GS enables efficient and high-quality object view synthesis, there are limitations.

First, our method relies on bounding boxes for object delineation, which can yield imprecise boundaries. Incorporating object-level segmentation could improve spatial accuracy and reduce visual artifacts in cluttered or occluded areas.

Second, LOD generation during GS training is currently static at render time. Introducing dynamic LOD adjustment based on the camera viewpoint could enable smarter resource allocation and better performance. Additionally, ROI-GS training time scales linearly with the number of objects, and model size depends on object count and size. An excessive number of Gaussians could negatively impact real-time rendering performance, making dynamic LOD essential. For example, distant objects could simply ignore the Object-GS and fall back to Scene-GS.

Finally, the composition strategy may produce minor inconsistencies near ROI boundaries, as center-inside Gaussians may extend beyond the AABB. Although generally negligible in practice, improved spatial filtering or blending strategies could further enhance compositional accuracy.

ACKNOWLEDGMENTS. We would like to thank the *Centre des Monuments Nationaux* and the *Hôtel de la Marine* team for giving us access to this magnificent place.

REFERENCES

- [1] X. Kong and R. G. Hucks, "Preserving our heritage: A photogrammetry-based digital twin framework for monitoring deteriorations of historic structures," *Automation in Construction*, vol. 152, p. 104928, Aug. 2023, ISSN: 0926-5805. DOI: 10.1016/j.autcon.2023.104928
- [2] M. Haibt, "End-to-end digital twin creation of the archaeological landscape in uruk-warka (iraq)," *International Journal of Digital Earth*, vol. 17, no. 1, Mar. 2024, ISSN: 1753-8955. DOI: 10.1080/17538947.2024.2324964
- [3] Z. Liu and S. Chang, "A study of digital exhibition visual design led by digital twin and vr technology," *Measurement: Sensors*, vol. 31, p. 100970, Feb. 2024, ISSN: 2665-9174. DOI: 10.1016/j.measen.2023.100970
- [4] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3D Gaussian Splatting for Real-Time Radiance Field Rendering," *ACM Transactions on Graphics*, vol. 42, no. 4, Jul. 2023. DOI: 10.1145/3592433
- [5] M. Niemeyer et al., "RadSplat: Radiance Field-Informed Gaussian Splatting for Robust Real-Time Rendering with 900+ FPS," in *International Conference on 3D Vision 2025*, 2025. DOI: 10.1109/3DV66043.2025.00018
- [6] J. Kulhanek et al., "LODGE: Level-of-detail large-scale Gaussian splatting with efficient rendering," in *Proceedings of the 39th International Conference on Neural Information Processing Systems*, 2025. DOI: 10.48550/arXiv.2505.23158
- [7] Z. Yu, A. Chen, B. Huang, T. Sattler, and A. Geiger, "Mip-Splatting: Alias-free 3D Gaussian Splatting," *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. DOI: 10.1109/CVPR52733.2024.01839
- [8] K. Ren et al., "Octree-GS: Towards consistent real-time rendering with lod-structured 3d gaussians," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–15, 2025, ISSN: 1939-3539. DOI: 10.1109/tpami.2025.3568201
- [9] J. Cui et al., "LetsGo: Large-Scale Garage Modeling and Rendering via LiDAR-Assisted Gaussian Primitives," *ACM Trans. Graph.*, vol. 43, no. 6, Nov. 2024, ISSN: 0730-0301. DOI: 10.1145/3687762
- [10] Y. Seo, Y. S. Choi, H. S. Son, and Y. Uh, *Flod: Integrating flexible level of detail into 3d gaussian splatting for customizable rendering*, 2024. arXiv: 2408.12894 [cs.CV].
- [11] N. Milef et al., "Learning fast 3d gaussian splatting rendering using continuous level of detail," *Computer Graphics Forum*, Apr. 2025, ISSN: 1467-8659. DOI: 10.1111/cgf.70069
- [12] Z. Yang, B. Gong, and K. Chen, *LOD-GS: Level-of-detail-sensitive 3d gaussian splatting for detail conserved anti-aliasing*, 2025. arXiv: 2507.00554 [cs.CV].
- [13] F. Windisch, L. Radl, T. Köhler, M. Steiner, D. Schmalstieg, and M. Steinberger, *A lod of gaussians: Unified training and rendering for ultra-large scale reconstruction with external memory*, 2025. arXiv: 2507.01110 [cs.GR].
- [14] J. Shen, Y. Qian, and X. Zhan, "Lod-gs: Achieving levels of detail using scalable gaussian soup," in *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 2025, pp. 671–680. DOI: 10.1109/CVPR52734.2025.00071
- [15] B. Kerbl, A. Meuleman, G. Kopanas, M. Wimmer, A. Lanvin, and G. Drettakis, "A Hierarchical 3D Gaussian Representation for Real-Time Rendering of Very Large Datasets," *ACM Transactions on Graphics*, vol. 43, no. 4, Jul. 2024. DOI: 10.1145/3658160
- [16] J. Cui et al., "LetsGo: Large-Scale Garage Modeling and Rendering via LiDAR-Assisted Gaussian Primitives," *ACM Trans. Graph.*, vol. 43, no. 6, Nov. 2024, ISSN: 0730-0301. DOI: 10.1145/3687762
- [17] M. Rogge and D. Stricker, "Object-Centric 2D Gaussian Splatting: Background Removal and Occlusion-Aware Pruning for Compact Object Models," in *Proceedings of the 14th International Conference on Pattern Recognition Applications and Methods - ICPRAM, INSTICC, SciTePress*, 2025, pp. 519–530, ISBN: 978-989-758-730-6. DOI: 10.5220/0013305500003905
- [18] M. Ye, M. Danelljan, F. Yu, and L. Ke, "Gaussian Grouping: Segment and Edit Anything in 3D Scenes," in *ECCV*, 2024.
- [19] W. Lyu, X. Li, A. Kundu, Y.-H. Tsai, and M.-H. Yang, *Gaga: Group Any Gaussians via 3D-aware Memory Bank*, 2024. arXiv: 2404.07977 [cs.CV].
- [20] X. Zhou, Z. Lin, X. Shan, Y. Wang, D. Sun, and M.-H. Yang, "Drivinggaussian: Composite Gaussian Splatting for surrounding dynamic autonomous driving scenes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 21634–21643. DOI: 10.1109/CVPR52733.2024.02044
- [21] M. M. Q. Li, P.-Y. Lajoie, and G. Beltrame, *Frequency-based View Selection in Gaussian Splatting Reconstruction*, 2024. DOI: 10.48550/arXiv.2409.16470 arXiv: 2409.16470 [cs.CV].
- [22] K. D. Polyzos, A. Bacharis, S. Madhuvarasu, N. Panikolopoulos, and T. Javidi, *ActiveInitSplat: How Active Image Selection Helps Gaussian Splatting*, 2025. arXiv: 2503.06859 [cs.CV].
- [23] M. Strong, B. Lei, A. Swann, W. Jiang, K. Daniilidis, and M. K. III, "Next best sense: Guiding vision and touch with fisherrf for 3d gaussian splatting," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2025, pp. 3204–3210. DOI: 10.1109/icra55743.2025.11127233
- [24] Y. Li et al., "Activesplat: High-fidelity scene reconstruction through active gaussian splatting," *IEEE Robotics and Automation Letters*, vol. 10, no. 8, pp. 8099–8106, 2025. DOI: 10.1109/LRA.2025.3580331
- [25] C. des Monuments Nationaux. "Hôtel de la marine." [Online]. Available: <https://www.hotel-de-la-marine.paris/> (accessed: 01.09.2025).