The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XLIX-B2-2026
XXV ISPRS Congress 2026 “From Imagery to Understanding”, Commission II, 4–11 July 2026, Toronto, Canada

Bundle-Adjusted Initialization for Efficient Earth Observation Gaussian Splatting

Jiyong Kim[1], Shuang Song[1,3], Rongjun Qin[1,2]*

[1] Department of Civil, Environmental and Geodetic Engineering, The Ohio State University, USA
[2] Department of Electrical and Computer Engineering, The Ohio State University, USA
[3] Translational Data Analytics Institute, The Ohio State University, Columbus, USA
{kim.9854, song.1634, qin.324}@osu.edu

Keywords: 3DGS, Initialization, Bundle-Adjustment, 3D Reconstruction, Digital Surface Model

Abstract

Satellite imagery offers a distinct advantage in Earth observation by providing expansive coverage and enabling the monitoring of inaccessible regions without physical on-site intervention, serving as a significantly more cost-effective and scalable alternative to traditional aerial or ground-based surveys. The task of 3D reconstruction from multi-view satellite images has therefore been a pivotal point of research at the intersection of photogrammetry and remote sensing. Recently, novel-view synthesis techniques such as Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS) have accelerated the accuracy and speed of topographic modeling. Among these, Earth Observation Gaussian Splatting (EOGS) has emerged as a state-of-the-art approach by adapting 3DGS to handle the unique geometric and radiometric characteristics of satellite data, including Rational Polynomial Coefficients (RPCs) and varying solar conditions. However, the standard EOGS pipeline relies on stochastic initialization, where Gaussians are distributed uniformly within a volumetric bounding box, leading to high computational overhead and dependency on aggressive pruning that can inadvertently remove critical geometric features, particularly in areas with complex urban structures. To address these limitations, we propose Bundle-Adjusted Initialization for Earth Observation Gaussian Splatting, which leverages sparse point clouds from bundle adjustment as geometric priors for Gaussian initialization. Combined with an adaptive densification strategy, our method achieves faster convergence and improved DSM accuracy on the DFC2019 dataset compared to the EOGS baseline.

1. Introduction

Over the past few decades, satellite imagery has become increasingly abundant, and its use in 3D reconstruction has attracted significant attention (Schuegraf et al., 2025). Traditional photogrammetric methods, such as stereo matching and triangulation, have proven effective for generating 3D terrain models from satellite images (Qin et al., 2022, Sadeq, 2025). In these pipelines, RPCs serve as a key camera model, and various approaches have been developed to enhance both the efficiency and accuracy of RPC-based reconstruction (Qin, 2016).

Meanwhile, Earth Observation Gaussian Splatting (EOGS) (Aira et al., 2025) was proposed to adapt Gaussian Splatting techniques specifically for satellite imagery. This method utilizes the bundle-adjusted RPCs to generate affine camera models, and simulates building shadows based on the projection directions of both the sun and the camera, providing guidance for height estimation and improving rendering quality. Taking less than 5 minutes to process $256 \times 256m^2$ region, EOGS demonstrated the capability to generate DSMs from multiple satellite images. To foster the learning process and dynamically control the number of Gaussians, EOGS initializes Gaussians randomly in the volumetric space and prunes Gaussians when their opacity values fall below a threshold.

However, this initialization strategy introduces several issues during optimization. First, it generates a large number of Gaussians at the initial stage and gradually reduces them over time, leading to computational inefficiency as the system must handle excessive memory overhead rather than maintaining a stable footprint throughout training. Second, since the initialization does not incorporate a splitting mechanism for large-scale primitives, the optimization tends to generate oversized Gaussians across the scene. These large primitives manifest as blurs in the final rendered output, degrading reconstruction quality. Finally, Gaussians tend to spread into irrelevant regions such as the atmosphere or deep water bodies, where geometric representation is not required. This unintended expansion wastes memory and introduces artifacts that degrade the accuracy of the reconstructed DSMs.

Therefore, we present a faster and practical approach to initialize and optimize Gaussians from sparse points. Our method can be summarized into two major steps:

1. **Initialization guided by sparse points:** We propose method to leverage sparse points generated during the bundle-adjustment process on initialization. These points are geometrically well aligned with bundle-adjusted RPCs, allowing the Gaussians to expand from precise points.
2. **Optimization incorporating Gaussian densification:** To further improve the quality of reconstructed 3D models, we selectively densify Gaussian, ensuring detailed reconstruction while maintaining overall computational efficiency.

2. Related Works

2.1 Novel-View Synthesis-Based 3D Reconstruction

Recently, novel-view synthesis-based 3D reconstruction has emerged and greatly evolved the 3D reconstructoin pipeline.

* Corresponding author

Starting from NeRF-based methods, S-NeRF (Derksen and Izzo, 2021) extended the NeRF approach in two ways: direct illumination from the Sun was modeled via a local light source visibility field, and indirect illumination from a diffuse light source is learned as a non-local color field as a function of the position of the sun. Similarly, Sat-NeRF (Marí et al., 2022) applied RPC-based warping to satellite imagery and introduced shadow-aware irradiance model and uncertainty weighting to deal with transient phenomena that cannot be explained by the position of the sun. EO-NeRF (Marí et al., 2023) expanded this pipeline into exploit varying appearances of raw satellite images. Sat-DN (Liu et al., 2025) induced depth and normal supervision to improve surface reconstruction from multi-view satellite images. Sat-Mesh (Qu and Deng, 2023) represented the scene as continuous signed distance function (SDF) and leveraged a volume rendering framework to learn the SDF value. Since NeRF-based methods required several hours or even a day to be trained, some works concentrated on reducing the time consumption. SatelliteRF (Zhou et al., 2024) introduced multi-resolution hash coding, enabling the model to greatly increase the training speed while maintaining high reconstruction quality. Similarly, Sat-NGP (Billouard et al., 2024) introduced multi-resolution hash encoding for faster NeRF-based 3D reconstruction.

Furthermore, 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) has emerged as a powerful tool for real-time view synthesis, with subsequent improvements like Mip-Splatting (Yu et al., 2024) addressing aliasing issues. EOGS (Aira et al., 2025) introduced the 3DGS-version of 3D reconstruction with satellite images. Incorporating the shadow modeling pipeline into the process, EOGS achieved state-of-the-art 3D reconstruction performance when compared with ground truth LiDAR DSM. Subsequently, EOGS++ (Bournez et al., 2025) extended EOGS framework by leveraging optical flow techniques and embedding bundle adjustment directly within the training process, avoiding reliance on external optimization tools while improving camera pose estimation.

2.2 Gaussian Optimization

The original 3DGS (Kerbl et al., 2023) proposed adaptive density control (ADC) for optimization. ADC starts with the SfM point cloud and creates a set of 3D Gaussians. After initialization, 3DGS densify and remove Gaussians that are transparent. The criteria for removing Gaussians is the opacity value, in which if the opacity value $\alpha < 0.005$. Furthermore, since 3DGS is initialized from SfM points, the locations where the Gaussians need densification need to be optimized. ADC introduced scale-based densification strategy, in which the Gaussians with high view-space positional gradients are either cloned or split depending on their scale. Small Gaussians in under-reconstructed areas are cloned and shifted along the gradient, while large Gaussians in over-reconstructed regions are split into smaller primitives with reduced scales to better capture geometric details. Specifically, Gaussians exceeding a threshold $\tau_{pos} = 0.0002$ are identified as either under-reconstructed or over-reconstructed based on their spatial extent. For small Gaussians in under-reconstructed regions, the method clones the primitive and moves it along the positional gradient to cover new geometry. Conversely, large Gaussians in over-reconstructed areas are split into two smaller Gaussians, with their scales divided by a factor of $\phi = 1.6$ and positions re-initialized through PDF sampling. To further regulate the population and eliminate floaters, the opacity α of all Gaussians is reset to near zero every 3,000 iterations, forcing the optimization to re-estimate the contribution of each primitive.

Meanwhile Markov Chain Monte Carlo (MCMC) (Kheradmand et al., 2024) was proposed to adaptively densify Gaussians. MCMC majorly differs with ADC on densification strategy. ADC prune Gaussians with the opacity value of $\alpha < 0.005$. However, MCMC just moves these Gaussians to the location of live Gaussians. This relocatoin strategy ensures that the number of Gaussians remains efficient while effectively reusing dead primitives to better represent complex geometry. This stochastic perturbation allows the Gaussians to explore the parameter space more thoroughly, effectively escaping local minima that often trap gradient-based methods.

EOGS, on the other hand, simplified these pipelines. For initialization, EOGS first utilized the Gaussians through spreading the points evenly on the volumetric space. Furthermore, instead of adding points on the volumetric space, EOGS used fixed number of points and pruned the points when the opacity value α goes under the value 0.0025. Since the pipeline doesn't split the Gaussians that goes too big, this pipeline often cause large Gaussians to dominate the entire scene, leading to visual artifacts. Due to the absence of an adaptive densification mechanism.

3. Proposed Methods

3.1 Sparse Point Initialization

Our initialization method could be described as Figure 3. To generate reliable correspondences for initialization, SIFT-sparse points were extracted from each image. Using these sparse points, we performed bundle-adjustment with the sat-bundleadjust library (Marí et al., 2021), refining the RPCs for EOGS input. The adjusted RPCs were then converted into affine camera models and provided as input to the EOGS framework. The sparse points associated with the bundle-adjustment process were directly used as initialization seeds for Gaussian Splatting.

The number of generated sparse points typically ranges from 4,000 to 8,000, which is insufficient to fully cover the scene at the initial stage. To address this, we generated four additional points around each point at a small distance of 0.02m, effectively densifying the initial set. Additionally, we added a "carpet" of Gaussians to define the ground-level scene, enabling accurate casting of shadow masks. These carpet points were randomly distributed with a density of $1\ \mathrm{m}^{-2}$ on the altitude plane where lowest 98% of sparse points are placed. For the sparse points, which are geometrically precise, we set the color value into white, and black for the carpet points to correctly overlay the shadow masks. The opacity value was set to 1.0 for both points since both should be considered as solid starting points.

3.2 Shadow Mapping

To account for solar effects, we adopt the shadow mapping formulation from the EOGS (Aira et al., 2025) baseline. We first render the elevation map $E(\mathbf{u})$ at pixel coordinates $\mathbf{u}$ by aggregating the height contributions of K Gaussians:

$$E(\mathbf{u}) = \sum_{k=1}^{K} h_k \alpha_k(\mathbf{u}), \tag{1}$$

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-445-2026

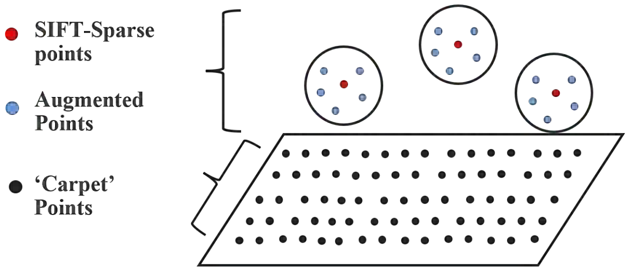

Figure 1. Gaussian initialization.

where h_k denotes the elevation (center height) of the k-th Gaussian, and $\alpha_k(\mathbf{u})$ represents its contributed opacity at pixel $\mathbf{u}$.

Using the rendered elevation map, the sun visibility mask $V(\mathbf{u})$ is defined to represent the shadow intensity at each coordinate:

$$V(\mathbf{u}) = \exp\left(-\alpha \cdot \max(0, E(\mathbf{u}) - \hat{E}_{sun}(\mathbf{u}))\right), \quad (2)$$

where α is a shadow density parameter scaled by the site's spatial resolution, and $\hat{E}_{sun}(\mathbf{u})$ denotes the reference elevation value sampled from the shadow map along the solar ray direction.

Finally, the shaded appearance $I_{shade}(\mathbf{u})$ is computed by modulating the original color $I(\mathbf{u})$ with the sun visibility mask and adding an ambient light component ψ:

$$I_{shade}(\mathbf{u}) = I(\mathbf{u}) \cdot V(\mathbf{u}) + (1 - V(\mathbf{u}))\psi, \quad (3)$$

where ψ represents the ambient term accounting for indirect illumination in shadowed regions.

3.3 Gaussian Optimization

To optimize the scene representation, we adopted a densification strategy rather than pruning (Figure 2). Following the Adaptive Density Control (ADC) pipeline of 3DGS (Kerbl et al., 2023), we densified the initial sparse point cloud based on view-space positional gradients and scale metrics.

Specifically, we integrated the gsplat (Ye et al., 2025) library to guide the duplication and splitting of Gaussians. Since the scene coordinates were normalized to the range $[-1, 1]$, all scale-related thresholds are defined relative to this unit cube. Gaussians are selected for densification if they meet any of the following criteria: a spatial scale exceeding 0.03, a projected pixel scale greater than 0.15, or a pixel-scale gradient surpassing 0.0006. This multi-criteria approach ensures that complex structures are sufficiently represented. Conversely, we performed pruning based on an opacity threshold of 0.0025 to eliminate redundant or negligible Gaussians.

Furthermore, consistent with ADC practices, we implemented an opacity reset every 3,000 iterations. During this process, opacities are clamped to twice the pruning threshold to facilitate further optimization. This periodic reset effectively reduces floaters, which are semi-transparent Gaussians that typically accumulate around the reconstructed surfaces.

4. Experiments

4.1 Experimental Settings

For comparison, we utilized the DFC2019 dataset (Bosch et al., 2019, Le Saux et al., 2019), a benchmark widely recognized for urban semantic labeling and 3D reconstruction tasks. This dataset provides LiDAR ground-truth DSM data with a 0.5m resolution. For generating consistent inputs, we processed the multispectral (MSI) data to generate representative RGB imager. Specifically, we extracted and combined the 5 (red), 3 (green) 2 (blue) bands from pan-sharpene 8-band MSI products, and applied a linear contrast stretching from 0 to 255 after excluding the top and bottom 2% of pixel values.

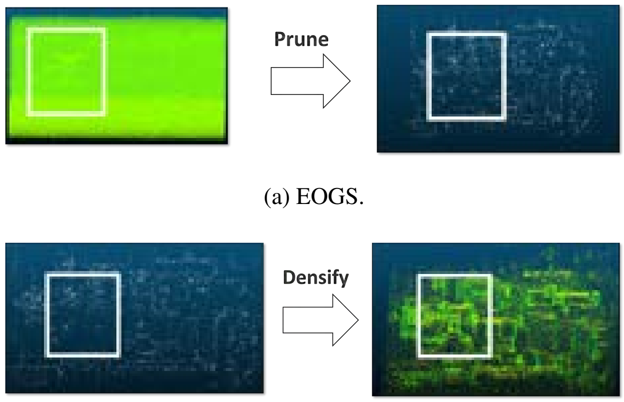

(a) EOGS.

(b) BAEOGS.

Figure 2. Gaussian optimization and densification strategy. (a) EOGS, (b) proposed BAEOGS.

The accuracy of the reconstructed geometry is primarily assesed using the registrated mean absolute Error (MAE_{reg}):

$$\mathrm{MAE}_{reg} = \min_{\mathbf{o}} \sum_{i=1}^{N} \frac{\left|\hat{M}(\hat{\mathbf{x}}_i) - o_z - M(\mathbf{x}_i)\right|}{N} \quad (4)$$

where N denotes the total number of valid pixels, $\hat{M}(\hat{\mathbf{x}}_i)$ represents the predicted elevation at coordinate i, and $M(\mathbf{x}_i)$ is the corresponding LiDAR ground truth. The term o_z is a scalar offset optimized to minimize the absolute difference, ensuring that the evaluation focuses on the structural fidelity and relative height accuracy of the reconstruction rather than absolute geo-rectification shifts.

4.2 Results

To assess training efficiency, we analyzed the densification process throughout the optimization (Figure 3). The baseline EOGS initializes with 261k Gaussians, which steadily decreases to 72k by the end of training, which is only 38% of the initial count. This reduction indicates that a large portion of the initially placed Gaussians are redundant or incorrectly positioned, eventually being pruned during optimization.

In contrast, our method initializes with a smaller set of 93k Gaussians and gradually densifies to 166k by iteration 3,000. Following the opacity reset at iteration 3,000, the count stabilizes at approximately 119k. Unlike the baseline, our approach maintains controlled densification throughout the optimization process, indicating that bundle-adjusted initialization provides a more effective geometric prior that reduces reliance on pruning and leads to a more stable training trajectory.

The quantitative advantages of our method are summarized in Table 1. Across most DFC2019 test sites, our method outper-

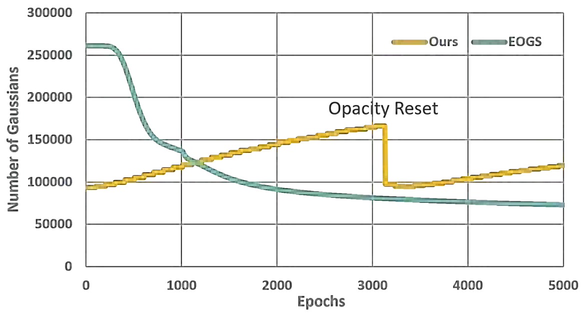

Figure 3. Number of Gaussians on training epochs.

formed EOGS in terms of MAE_{reg}. For instance, in the JAX-004 and JAX-214 tiles, our method achieved lower error rates (1.58m and 1.61m, respectively), demonstrating better geometric accuracy.

Method	JAX-004	JAX-068	JAX-214	JAX-260
EOGS	1.65	**0.99**	1.65	1.44
Ours	**1.58**	1.09	**1.61**	**1.41**

Table 1. MAE_{reg} comparison across different sites of DFC2019.

The visual comparison in Figure 4 further supports these findings. The DSMs generated by our method exhibit finer details compared to the baseline. We attribute this to the sparse points obtained from bundle adjustment serving as geometric anchors, enabling Gaussian densification to prioritize complex structural regions while suppressing redundant points in flat or featureless areas.

Beyond accuracy, our method reduces training latency. As shown in Table 2, our approach consistently achieves faster convergence across all sites, with training time reduced by approximately 20–30% compared to EOGS. This efficiency gain is a direct result of the optimized initial Gaussian distribution, which reduces the computational overhead associated with processing and pruning redundant Gaussians.

Method	JAX-004	JAX-068	JAX-214	JAX-260
EOGS	120s	133s	163s	127s
BAEOGS	**94s**	**117s**	**103s**	**95s**

Table 2. Training time comparison across different sites of DFC2019.

Figure 5 provides empirical evidence supporting our hypothesis regarding Gaussian distribution. In the baseline EOGS, optimized Gaussians tend to be distributed across background areas or outside actual building structures. This occurs because the random initialization fails to provide adequate geometric constraints, leading to over-dispersion of Gaussians during optimization.

In contrast, bundle-adjusted initialization ensures that Gaussians are placed within geometrically relevant regions. The sparse points from bundle adjustment act as geometric priors, guiding Gaussian densification toward areas with structural complexity. As a result, our method suppresses redundant Gaussians in flat or featureless areas, which is reflected in the performance metrics.

5. Conclusion

In this study, we presented a novel approach to satellite-based 3D reconstruction by integrating sparse point-guided initialization with an optimized Gaussian densification process within the EOGS framework. Our research was motivated by the inherent limitations of random or poorly structured initializations in Gaussian Splatting, which often lead to redundant computational overhead and suboptimal geometric fidelity. By leveraging sparse points derived from bundle adjustment as high-quality geometric anchors, our method establishes a more robust starting point for the optimization process. Experimental evaluations on the DFC2019 dataset demonstrate that this strategy not only accelerates convergence but also improves the efficiency of memory usage by maintaining a controlled and purposeful distribution of Gaussians.

The quantitative and qualitative results further demonstrates the advantages of our proposed method. Compared to the baseline EOGS, our approach consistently produces DSMs with finer structural details that are often smoothed out by EOGS. The stability in Gaussian count throughout training indicates that the model prioritizes densification in regions of structural complexity while suppressing redundant points in flat or featureless areas, leading to a more compact scene representation.

However, our method exhibits dependencies on the density and quality of the initial bundle-adjusted points. In scenarios where satellite imagery lacks sufficient texture or overlap, the optimization can become unstable or fail to capture the full geometric complexity of the scene. This suggests that bundle-adjusted initialization may require further augmentation in challenging environments. Future work will focus on developing hybrid initialization strategies that adaptively supplement sparse points with synthetic or learned priors, ensuring robust reconstruction across diverse geographic regions.

Acknowledgements

This work is supported by the Office of Naval Research (Award No. N000142312670) and Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DOI/IBC) contract number 140D0423C0075. This work also used High Performance Computing resources provided by the Ohio Supercomputer Center (Ohio Supercomputer Center, 1987).

References

Aira, L. S., Facciolo, G., Ehret, T., 2025. Gaussian splatting for efficient satellite image photogrammetry. *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5959–5969.

Billouard, C., Derksen, D., Sarrazin, E., Vallet, B., 2024. Satngp: Unleashing neural graphics primitives for fast relightable transient-free 3d reconstruction from satellite imagery. *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, 8749–8753.

Bosch, M., Foster, K., Christie, G., Wang, S., Hager, G. D., Brown, M., 2019. Semantic stereo for incidental satellite images. *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, IEEE, 1524–1532.

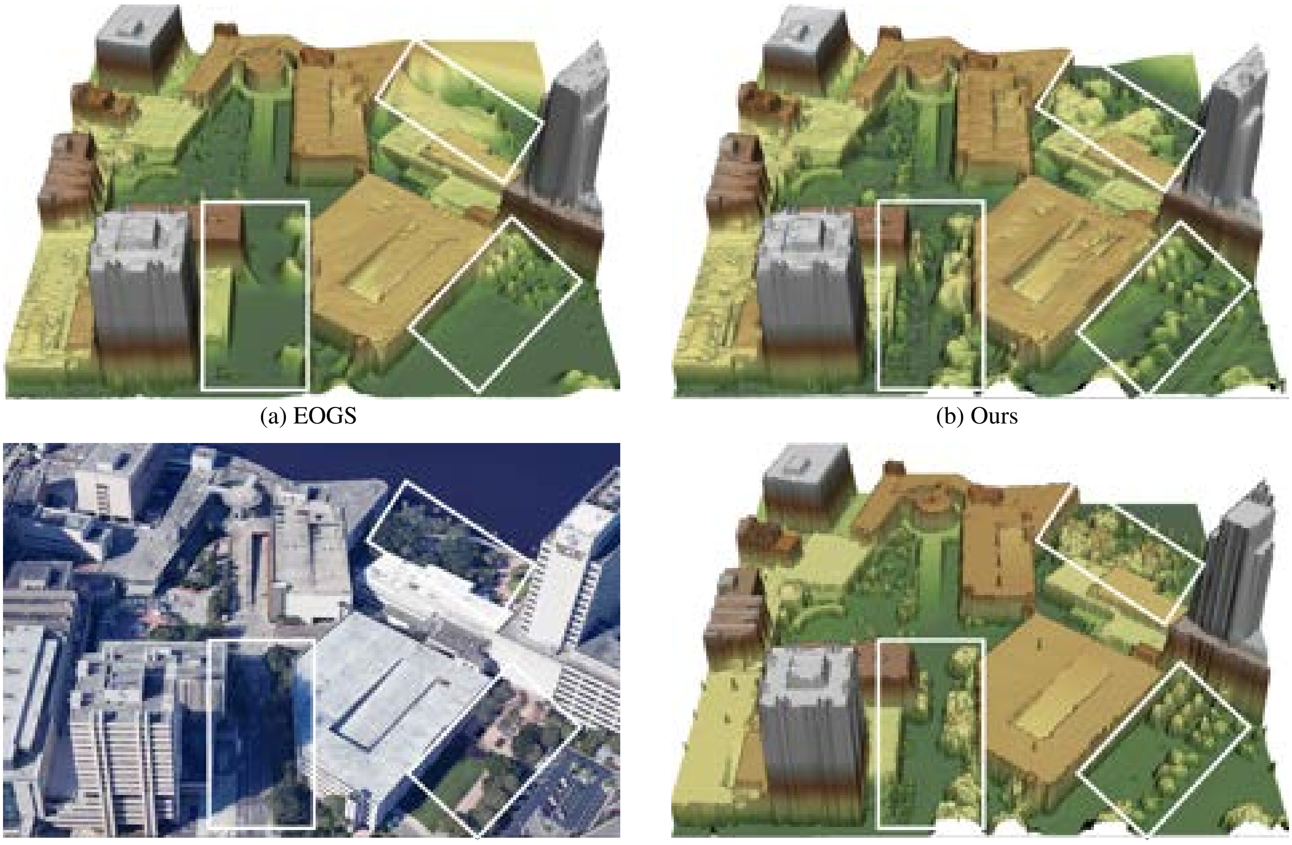

(a) EOGS (b) Ours

(c) Google Earth (d) Ground Truth (LiDAR)

Figure 4. Comparison on generated DSMs and ground truth. JAX-214 tile was selected for comparison

Bournez, P., Aira, L. S., Ehret, T., Facciolo, G., 2025. EOGS++: Earth Observation Gaussian Splatting with Internal Camera Refinement and Direct Panchromatic Rendering. *arXiv preprint arXiv:2511.16542*.

Derksen, D., Izzo, D., 2021. Shadow neural radiance fields for multi-view satellite photogrammetry. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1152–1161.

Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G. et al., 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4), 139–1.

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*, 37, 80965–80986.

Le Saux, B., Yokoya, N., Hansch, R., Brown, M., Hager, G., 2019. 2019 data fusion contest [technical committees]. *IEEE Geoscience and Remote Sensing Magazine*, 7(1), 103–105.

Liu, T., Zhao, S., Jiang, W., Guo, B., 2025. Sat-DN: Implicit Surface Reconstruction from Multi-View Satellite Images with Depth and Normal Supervision. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*.

Marí, R., Facciolo, G., Ehret, T., 2022. Sat-nerf: Learning multi-view satellite photogrammetry with transient objects and shadow modeling using rpc cameras. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1311–1321.

Marí, R., Facciolo, G., Ehret, T., 2023. Multi-date earth observation nerf: The detail is in the shadows. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2035–2045.

Marí, R., de Franchis, C., Meinhardt-Llopis, E., Anger, J., Facciolo, G., 2021. A Generic Bundle Adjustment Methodology for Indirect RPC Model Refinement of Satellite Imagery. *Image Processing On Line*, 11, 344–373. `https://doi.org/10.5201/ipol.2021.352`.

Ohio Supercomputer Center, 1987. Ohio supercomputer center.

Qin, R., 2016. Rpc stereo processor (rsp)–a software package for digital surface model and orthophoto generation from satellite stereo imagery. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 3, 77–82.

Qin, R., Ling, X., Farella, E. M., Remondino, F., 2022. Uncertainty-guided depth fusion from multi-view satellite images to improve the accuracy in large-scale DSM generation. *Remote Sensing*, 14(6), 1309.

Qu, Y., Deng, F., 2023. Sat-mesh: Learning neural implicit surfaces for multi-view satellite reconstruction. *Remote Sensing*, 15(17), 4297.

Sadeq, H. A., 2025. Accuracy assessment of dense point cloud generated by deep learning and semiglobal matching. *Photogrammetric Engineering & Remote Sensing*, 91(3), 153–162.

Schuegraf, P., Gui, S., Qin, R., Fraundorfer, F., Bittner, K., 2025. Sat2building: LoD-2 building reconstruction from satellite imagery using spatial embeddings. *Photogrammetric Engineering & Remote Sensing*, 91(4), 203–212.

(a) EOGS (rendered)

(b) BAEOGS (rendered)

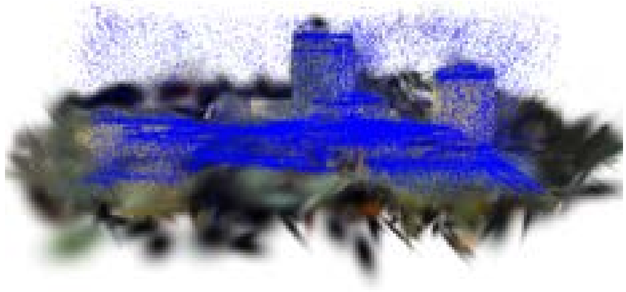

(c) EOGS (result Gaussians)

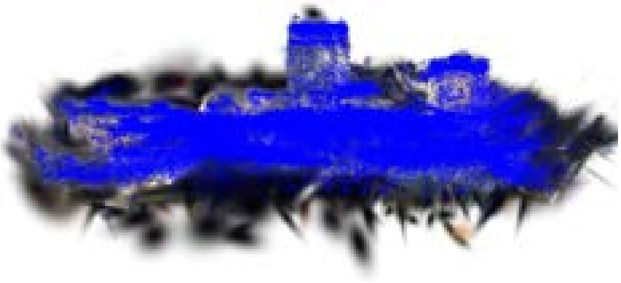

(d) BAEOGS (result Gaussians)

Figure 5. Visual comparison of generated Gaussians and rendered DSMs on the JAX-214 tile. Top row (a, b) illustrates the Gaussian distribution, while the bottom row (c, d) shows the final rendered surface quality. BAEOGS demonstrates significantly reduced ’blur’ artifacts compared to EOGS.

Ye, V., Li, R., Kerr, J., Turkulainen, M., Yi, B., Pan, Z., Seiskari, O., Ye, J., Hu, J., Tancik, M., Kanazawa, A., 2025. gsplat: An open-source library for Gaussian splatting. *Journal of Machine Learning Research*, 26(34), 1–17.

Yu, Z., Chen, A., Huang, B., Sattler, T., Geiger, A., 2024. Mip-splatting: Alias-free 3d gaussian splatting. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19447–19456.

Zhou, X., Wang, Y., Lin, D., Cao, Z., Li, B., Liu, J., 2024. SatelliteRF: Accelerating 3D reconstruction in multi-view satellite images with efficient neural radiance fields. *Applied Sciences*, 14(7), 2729.