GauStudio: A Modular Framework for 3D Gaussian Splatting and Beyond

Chongjie Ye^{1,2*}, Yinyu Nie^{3*}, Jiahao Chang^{1, 2*}, Yuantao Chen², Yihao Zhi^{1,2}, Xiaoguang Han^{2,1†}
^{*}equal contribution [†]corresponding author

²SSE, CUHKSZ

¹FNii, CUHKSZ

github.com/GAP-LAB-CUHK-SZ/gaustudio

³Technical University of Munich



Figure 1. **GauStudio** is a modular framework that unifies various 3D Gaussian Splatting techniques. It decomposes 3D scenes into components like foreground and background models, represented using specialized techniques on 3D Gaussians. These components can be flexibly combined and rendered to synthesize novel views, enabling tailored modeling pipelines for different tasks.

Abstract

We present GauStudio, a novel modular framework for modeling 3D Gaussian Splatting (3DGS) to provide standardized, plug-and-play components for users to easily customize and implement a 3DGS pipeline. Supported by our framework, we propose a hybrid Gaussian representation with foreground and skyball background models. Experiments demonstrate this representation reduces artifacts in unbounded outdoor scenes and improves novel view synthesis. Finally, we propose Gaussian Splatting Surface Reconstruction (GauS), a novel render-then-fuse approach for high-fidelity mesh reconstruction from 3DGS inputs without fine-tuning. Overall, our GauStudio framework, hybrid representation, and GauS approach enhance 3DGS modeling and rendering capabilities, enabling higher-quality novel view synthesis and surface reconstruction.

1. Introduction

Novel view synthesis - the task of rendering photorealistic novel views of a 3D scene from partial image observations - is a longstanding challenge in computer vision and graphics. Recent advances in neural rendering techniques like NeRF [32] have made significant strides by optimizing multi-layer perceptrons (MLPs) using volumetric raymarching to represent and synthesize novel views of captured scenes. Recently, many efficient radiance field followups are built on continuous representations by interpolating values stored in voxel grids [11, 52], hash tables [34], or point-based data structures [37, 55].

However, these implicit neural representations often struggle to model high-frequency details and lack explicit geometric structures, making them difficult to edit and compress for real-time applications. As an alternative, explicit 3D representations such as 3D Gaussians[23] have recently emerged as an efficient and editable scene modeling approach. Unlike NeRFs[3, 33] which model a continuous volumetric field, 3D Gaussians Splatting(3DGS) represents a scene as a set of anisotropic 3D Gaussian kernels which encode local geometry and view-dependent appearance. During testing time, these Gaussians can reproduce detailed geometry and view effects with a rendering and editing friendly format.

In this work, we introduce GauStudio, a modular framework tailored for 3D Gaussian Splatting (3DGS) with a set of customizable modules for diverse scene reconstruction scenarios. The key advantage of GauStudio lies in its highly modularized design. Its architecture allows different components to be flexibly combined and replaced, enabling the convenient construction of different 3D scene modeling methods. Researchers can freely combine different foreground models (e.g., Gaussian models), background models (e.g., environment maps [18] or our Gaussian Sky models), and other components according to their needs, forming hybrid modeling pipelines tailored for specific scenes and tasks. This composability significantly enhances the framework's flexibility and scalability, facilitating accelerated innovation in the field of 3D scene modeling.

GauStudio consists of several key stages: 1) scene initialization from SFM [39] or from splatting priors [4]; 2) Gaussians' optimization with geometric and sparsity regularizers; 3) an enhancement step to boost 3D Gaussians' representation ability; 4) scene compression via learnable or geometric pruning. We also consider different representations for encoding view-dependent appearance within the Gaussian splats, including the traditional spherical harmonics, learned neural feature vectors, and explicit feature caching structures.

Additionally, we present an efficient surface reconstruction module from optimized 3D Gaussians, namely *GauS*. It uses volumetric fusion to convert 3D Gaussians into textured triangle meshes. Combined with our GauStudio which supports various GS-based methods [7, 31, 60, 65] natively, our Gaus is a plug-and-play module that can easily extract meshes for different types of Gaussians with a single command. For background modeling, we build the sky background using spherical environment maps composed of Gaussians. All these modules are fully customizable in GauStudio.

In summary, we introduce GauStudio, a flexible and modular framework assembled with 3D Gaussian Splatting techniques for different tasks (e.g., scene reconstruction, editing, simulation, and path planning). The key contributions are concluded as follows:

1. A modularized and composable Gaussian Splatting framework for practitioners to seamlessly integrate different components, e.g., foreground Gaussian models, background representations, and other modules based on their specific tasks.

- The framework encapsulates many key stages in 3D Gaussians reconstruction, such as initialization, optimization, regularization, representation enhancement, and compression, with each stage fully customizable.
- 3. We provide an efficient mesh extraction pipeline, namely *GauS*, to convert optimized Gaussians into textured meshes, serving as a versatile plug-and-play solution across different GS-based methods.

Overall, GauStudio presents a comprehensive yet customizable platform to drive innovation in 3D scene modeling tasks like reconstruction, editing, and simulation.

2. Preliminaries

Neural volume rendering (NeRF) NeRF [33] represents a scene through continuous volumetric density and color fields. To render a view, NeRF first casts the rays from a posed camera and samples 3D points along the rays. Then, it predicts the volume density σ and RGB color *c* at 3D positions *x* by querying the multi-layer perceptrons (MLPs), and the color is integrated along each ray. Formally, The rendered color $\hat{c}(o, d)$ for a pixel approximates a Riemann sum over N samples:

$$\hat{\mathbf{c}}(\mathbf{o}, \mathbf{d}) = \sum_{i=1}^{N} w_i \mathbf{c}_i,\tag{1}$$

where the weight w_i is determined by the accumulated opacity along each ray. The color loss \mathcal{L}_{RGB} which measures the difference between renderings and input views provides supervision for the MLPs predicting σ and c:

$$\mathcal{L}_{\text{RGB}} = |\hat{\mathbf{c}} - \mathbf{c}|_1. \tag{2}$$

3D Gaussian Splatting (**3DGS**) Recent works have developed Gaussians as an explicit 3D scene representation for novel view synthesis [7, 23, 31, 60]. In contrast to implicit NeRFs [33], 3DGS represents a scene using a set of 3D anisotropic Gaussian kernels. Each kernel is defined by: Mean position μ , Opacity σ , Covariance matrix \sum , and Spherical harmonic (SH) coefficients *C* for representing angular color variations. To render a view, the Gaussian splats are first projected onto 2D screen space based on sorted depth. Then, differentiable volumetric rendering combines the kernels' color and density properties, which can be formulated as:

$$C = \sum_{k \in \mathcal{P}} \alpha_k \operatorname{SH}(d_k; \mathbf{C}k) \prod_{j=1}^{k-1} (1 - \alpha_j)$$
(3)

3. GauStudio

In this section, we describe the components of GauStudio and how they are integrated into a uniform pipeline. An overview of our framework is exhibited in Fig.1. In particular, our framework enables a hybrid scene representation comprising Gaussian foreground models and background models. As for the foreground Gaussians, GauStudio can be readily converted them to textured meshes for explicit surface modeling, obviating the need for further fine-tuning or NeRF-based [33] re-training. Meanwhile, the background models support diverse backends like environment maps, textured meshes, and other neural representations [34].

3.1. Problem setting

Given a set of calibrated RGB images $I = \{I_1, \ldots, I_N\}$ and the corresponding camera parameters $\Theta = \{\theta_1, \ldots, \theta_N\}$, where $\theta_i = \{R_i, T_i, K_i\}$ represents the extrinsic rotation $R_i \in \mathbb{R}^{3\times 3}$ (an element of the Special Orthogonal group SO(3)), extrinsic translation $T_i \in \mathbb{R}^3$, and intrinsic camera parameters $K_i = [c_x, c_y, f_x, f_y]^\top$ with principal point (c_x, c_y) and focal lengths (f_x, f_y) , GauStudio aims to learn an optimized radiance field F^* that models the scene's appearance and geometry.

We define the 3D Gaussian Splatting (3DGS) function Fas a mapping from a 3D spatial position $\mathbf{x} \in \mathbb{R}^3$, view direction θ and ϕ , and camera parameters $\Theta = \{\mathbf{R}, \mathbf{T}, \mathbf{K}\} \in \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^4$ to a color $\mathbf{c} \in \mathbb{R}^3$ and density $\sigma \in \mathbb{R}$:

$$\begin{split} F: \mathbb{R}^3 \times \mathbb{R}^2 \times (\mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^4) \to \mathbb{R}^3 \times \mathbb{R} \\ F(\mathbf{x}, \mathbf{d}, \Theta) = (\mathbf{c}, \sigma) \end{split}$$

The goal is to find the optimal radiance field F^* that minimizes the image reconstruction loss and regularization terms:

$$F^* = \operatorname{argmin} F\mathcal{L}_{rgb}(F; I, \Theta) + \lambda \mathcal{L}_{reg}(F)$$

Where \mathcal{L}_{rgb} is the image reconstruction loss, and \mathcal{L}_{reg} represents additional regularization losses. We postpone the details of $\mathcal{L}_{rgb}(F; I, \Theta)$ and $\mathcal{L}_{reg}(F)$ to Section 3.5.

3.2. Framework overview

GauStudio implements 3D Gaussian Splatting (3DGS) through a unified reconstruction framework φ that takes input images I and camera projection matrices P calculated from camera parameters $\Theta = \theta_1, \ldots, \theta_N$ as input and outputs optimized 3D Gaussian Kernels X and their parameters γ :

$$\varphi(I,P) = X,\theta \tag{4}$$

$$F(\mathbf{x}, \mathbf{d}, \Theta; X, \gamma) = (\mathbf{c}, \sigma) \tag{5}$$

Our framework φ comprises four components: Initialization, Optimization, Enhancement and Compression.

Initialization The Initialization module ψ takes the input image *I*, the projection function *P*, and an optional set of

parameters Ω_{ψ} (e.g., initialization strategy, desired number of Gaussians) to generate the initial Gaussians X and their parameters θ :

$$X^{(0)}, \gamma^{(0)} = \psi(I, P; \Omega_{\psi})$$

Optimization The Optimization stage ω takes the initial Gaussians X, parameters θ , input image I, projection function P, and an optional set of parameters Ω_{ω} (e.g., regularization terms, depth priors) to produce the optimized Gaussians: \hat{X} and parameters $\hat{\theta}$.

$$X^{(1)}, \gamma^{(1)} = \omega(X^{(0)}, \theta^{(0)}, I, P; \Omega_{\omega})$$

Enhancement Unlike neural radiance fields, which implicitly capture the entire scene, the Gaussian Kernels representation requires the existence of corresponding Gaussian Kernels to accurately model specific parts of the scene. Additionally, gaussian's representation ability is bounded by the amounts of kernels and the method to model view-dependent appearance. By densifying and completing the optimized Gaussians \hat{X} and parameters $\hat{\theta}$, the Enhancement stage augments the representation with additional Gaussian Kernels, enabling it to model high-frequency details and provide a more complete representation of the scene:

$$X^{(2)}, \gamma^{(2)} = \delta(X^{(1)}, \theta^{(1)}, P; \Omega_{\delta})$$

Compression However, introducing a large number of Gaussian Kernels can lead to computational inefficiency and increased memory requirements, which can be detrimental to the overall performance of the system. To address this concern, the Compression stage becomes essential. The Compression stage aims to strike a balance between the representation's ability to capture scene details and its efficiency in terms of computational resources and memory usage. It achieves this by identifying and removing insignificant Gaussian Kernels that contribute minimally to the overall representation, or by compacting the parameters of the Gaussian Kernels through techniques such as quantization or pruning: $X^{(3)}, \gamma^{(3)} = \pi(X^{(2)}, \theta^{(2)}, P; \Omega_{\pi})$

3.3. Representation

To modeling the view-dependent appearance of each 3D Gaussian, we consider various representations in our GauStudio framework as illustrated in Fig.1. In particular, we consider Fixed Basis Function, Neural Feature Vectors, Triplane, Hash Grid, Codebook.

Fixed Basis Functions: Traditional Gaussian splattingç utilizes spherical harmonics (SH) to model the viewdependent color. Additionally, we consider anisotropic spherical Gaussians (ASG) [58] for modeling specular components within scenes and Spherical Gaussians (SG)[59] for balancing image quality and rendering speed. To achieve material and lighting decomposition, we further consider associating each point with additional PBR parameters[13, 18, 28, 36, 41] and indirect illumination[18, 41] within our framework.

Neural Feature Vectors: To model the view-dependent color in real-time view synthesis systems, each point is associated with a compact Neural Feature Vector, which is decoded using a small MLP[14, 25, 31, 40].

Triplane, Hash Grid: Instead of indexing neural feature vectors with point cloud IDs, alternative approaches cache the neural features in explicit data structures, such as triplanes[54, 65] and hash grids[26]. For any given position x, the corresponding feature vector can be queried from these explicit structures. In the triplane representation, the position is projected onto axis-aligned feature planes, and the final feature f_t is obtained by concatenating three trilinear interpolated features. For hash grids, the feature is interpolated from a multi-resolution grid stored using instant-ngp[34]. To support unbounded scenes, we consider employing the contraction technique[59] to compactly represent background features.

Codebook Quantization Motivated by the balance between compact storage and rendering efficiency, we also consider learning a discrete codebook of vector codes during training[9, 35]. This is crucial as only a simple lookup into the codebook is required using the stored indices at rendering time. We seek to exploit lower bandwidth and computational costs compared to conventional feature fetching and processing.

3.4. Initialization

ciTraditional 3D Gaussian splatting pipelines [23] often rely on a sparse point cloud initialized from correspondence-rich image pairs. However, this approach can be sensitive to the quality of feature matches and may fail to reconstruct texture-less surfaces accurately. Our GauStudio framework, besides traditional initializations from SfM point clouds [39], MVS point clouds [2, 10, 13], depth maps [12], and LiDAR[56, 63], we consider alternative strategies that better leverage off-the-shelf prior network.

Semi-Dense Point Cloud Initialization Inspired by recent advances in semi-dense feature matching [8, 42] and point tracking [19, 50], we propose initializing the point cloud using dense feature matches instead of traditional detector-based matches. Specifically, we extract and aggregate dense feature correspondences using hierarchical localization from[38], then form the point cloud via multi-view 3D triangulation using the Direct Linear Transform (DLT) [16]. Semi-dense initialization strikes a balance between training time (which can be prohibitively high for dense MVS point initialization) and rendering quality on texture-less surfaces and unbouneded background.

Generalized Gaussian Splatting Initialization In scenarios with limited scene observations, the optimized gaussian shapes may deviate significantly from the actual surface geometry. This deviation leads to a decline in the rendering quality when viewed from a new viewpoint. To alleviate this issue, we consider initializing the Gaussians using pre-computed properties θ from a generalized Gaussian splatting model[4, 43, 44, 54, 65] for further optimization.

3.5. Optimization and Regularization

In the simplest case, the Gaussian kernels are optimized to minimize the image reconstruction loss (Eq.2). Additionally, auxiliary regularization can guide the optimization towards desired properties. We classify the regularization based on the optimization target: Geometry Modeling, Floater Cleaning, and Efficiency Improvement.

Geometry Modeling Extracting a mesh from millions of tiny 3D Gaussians is challenging due to their arbitrary spatial distribution and anisotropic kernels. One approach is to represent surfaces with Gaussians is to flatten them into 3D "surfels" aligned with the surface. For instance, adopting scaling loss [5] could minimize the smallest component of the scaling factor $s \in \mathbb{R}^3$ for each Gaussian towards zero. This regularization also aids in determining Gaussian normals by selecting the shortest scaled edge. To align Gaussians with the surface, we consider regularizing the rendered normals against planar constraints [2, 7], rendered depth buffers [13, 18, 28], or normal priors [27, 61].

Floaters Cleaning We employ the entropy loss[15] to encourage the Gaussians' opacities to be either 0 or 1. Then, with the pruning mechanism pipeline, floaters will be automatically removed during training. However, we found that such regularization could lead to a trade-off situation, as it may reduce the number of remaining points, potentially compromising the rendering quality and negatively impacting transparent components.

Efficiency Improvement Since small Gaussians have a negligible contribution to the overall rendering quality, it is natural to consider removing less-contributing Gaussians for efficiency with little rendering quality loss. We consider the masking loss[26] to encourage Gaussian sparsity. Learnable masking is utilized to eliminate redundant Gaussians based on their volume (scales s) and transparency (opacities o). Binary masks M generated from a mask parameter m are applied to s and o, effectively removing Gaussians with small volumes or low opacities.

3.6. Enhancement

Gaussians Enhancement is critical, especially given sparse point cloud initialization. We follow [23] and [64] and apply rule-based adaptive control during optimization. To better model textureless surfaces and improve convergence speed, monocular depth from an RGB-D sensor or a prior network[57] could be taken into account for complete the gaussians during training[12, 22]. Additionally, a novel completing method has been proposed to utilize propagated depth, which serves the same effort as depth-guided densification while in a self-supervised way. For scaffold-based representations[31, 58], a different densification method is adopted to control the number of anchor points according to the accumulated gradients within the voxel instead of screen spaces.

Besides traditional splitting, copying, and inserting methods, we tend to explore point cloud completion and upsampling methods to empower Gaussian densification in the future [65].

3.7. Compression

We consider a simple opacity-based pruning strategy for general learning. Additionally, if a mask is provided, a visible mask will be calculated to prune the Gaussians projected onto the masked pixels. We could also use depth for prune the 3D gaussians in the free space or control the density of gaussians according to the distance[25].



SuGaR (30min)

Ours (1.5min)

Figure 2. Comparsion with other mesh extraction strategies. Utilizing the Poisson reconstruction directly on 3DGS leads to a bad performance. Due to the special design of GaussianPro on normal and depth, the performance of Poisson reconstruction on GaussianPro is slightly better. SuGaR always tends to produce larger ellipses. Compared with the above methods, our method can generate high-quality mesh efficiently.

4. Gaussian Surface Reconstrution(GauS)

Converting 3D Gaussians into triangle meshes presents a non-trivial task, and several strategies have been proposed to tackle this challenge. Among these, the most basic approach involves applying Poisson surface reconstruction [20, 21] to the means of the 3D Gaussians, with point normals computed using the shortest axis of the Gaussians [15]. Nonetheless, this method often yields poor results due to the noisy Gaussians around the surface and leads to floater-like Gaussians in the free space.

Recently, Sugar [15] proposed several geometryoriented regularizations to align the Gaussians with the surface of the scene. Despite yielding promising results, this approach is plagued by long refinement periods because it relies on a computationally expensive sorting algorithm to determine the connectivity of neighboring Gaussians.

In our experiments with the Sugar [15], we noticed that even without fine-tuning and regularization, most points were well-aligned with the surface when dense training views are provided. However, due to noisy estimated normals, the resulting mesh quality suffered, displaying numerous holes and artifacts. While some of the regularization methods discussed in Sec.3.5 could potentially enhance normal quality, they incur a 1-2 dB lower PSNR in object-centric and indoor scenes, representing an undesirable trade-off.



Figure 3. Our Gaussian Surface Reconstruction can be integrated into existing 3DGS pipelines. Before GS rasterization, 3DGSbased frameworks will convert Gaussian Splatting into a unified representation that can be used for our Gaussian Surface Reconstruction.

To circumvent the dependency on Gaussian normals and their associated challenges, we employed a volumetric fusion approach [47]. Specifically, we rendered the median depth as proposed in [22] and fused it into a mesh using the VDBFusion [47]. Fig.2 demonstrates the effectiveness of volumetric fusion compared to naive Poisson reconstruction and the Sugar method [15], showcasing significantly improved quality and efficiency.

Notably, our surface reconstruction method is compatible with nearly all 3DGS-based methods, as illustrated in Fig.3 by utilizing a customized GS rasterizer. This versatility ensures seamless integration with existing 3DGS pipelines and frameworks, facilitating widespread adoption and application. Moreover, a compelling motivation for employing volumetric fusion lies in its scalability, memory efficiency, and time efficiency, particularly for super large-scale scenes like city districts and buildings. This attribute makes our approach ideally suited for handling intricate, large-scale environments, which is essential for numerous real-world applications across domains such as urban modeling, virtual reality, and digital twins.

5. Gaussian Sky Modeling

Outdoor scenes often contain sky regions located at an infinite distance from the camera. The 3DGS model tends to generate cloud-like artifacts in the foreground region when attempting to model the sky. These artifacts significantly influence the quality of novel view synthesis. To address this issue, our rendering model incorporates a spherical environment map composed of gaussians to model the sky separately from the foreground. We render the color for each pixel using compositional rendering similar to [53]:

$$C(p) = \sum_{i} T_i \alpha_i c_i + (1 - O(p)) \cdot C_{sky}(p)$$
(6)

$$O(p) = \sum_{i} T_{i} \alpha_{i} \qquad (7)$$

$$T_i = \prod_{k=1}^{n-1} (1 - \alpha_i) \qquad (8)$$

Here, $C_{sky}(p)$ is the rendered pixel color from the sky model, O(p) is the foreground opacity value obtained by summing the sample weights, and T_i represents the accumulated transmittance.

However, we found that simply blending the sky color and foreground color could not guarantee the clearance of artifacts when sparse points existed in the free space initially. To address this, we incorporate semantic priors from a semantic segmentation network [6] to generate sky masks. We then penalize the opacity values O(p) for the sky region to be close to zero, enforcing a clear separation between the sky and foreground regions:

$$L_{sky} = \lambda_{sky} \sum_{p \in \mathcal{R}_{sky}} O(p)$$

where λ_{sky} is the weighting factor.

Experiments demonstrate that penalizing the mask loss in the sky region helps remove cloud artifacts, even with an imperfect sky mask. Although alternative approaches, such as a sky depth loss to penalize the sky region depth to be as far as possible or a binary cross-entropy (BCE) semantic regularization, were explored, we found that they tended to require longer convergence times and could adversely impact the foreground modeling quality. Nonetheless, the sky depth loss and BCE semantic regularization could potentially be beneficial in specific scenarios or as complementary losses. The sky depth loss may aid in enhancing the perceived depth separation between the sky and foreground elements, while the BCE semantic regularization could further refine the sky-foreground boundary based on semantic cues.

6. Experiments

6.1. Experimental Setup

Dataset and Metric In order to assess the effectiveness of our gaussian-based environment map, we conducted an evaluation on commonly used Tanks and Temples [24] dataset. We did not perform experiments with Mip-NeRF360 [3] and DeepBlending [17] since our method focuses on scenes with unbounded skies. We quantified the image synthesis quality with standard metrics such as Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) [51], and Learned Perceptual Image Patch Similarity (LPIPS) [62] to measure the effectiveness of the method. Additionally, we also reported the memory footprint and runtime for each method to assess their computational efficiency.

In order to evaluate our gaussian surface reconstruction, we conducted expriments on the DTU [1] and Blender [32] datasets. For the assessment of surface reconstruction, we present lots of qualitative comparisons.

Baselines and Implementation For our gaussian-based environment map, we selected 3D-GS [23] as our main baseline and also included some latest 3DGS-based works, including Scaffold-GS [31], GaussianPro [7], and Mip-Splatting [60]. We initialized the gaussian environment map by sampling 100,000 points from a Fibonacci sphere. We set λ_{sky} as 10 to remove cloud artifacts in the beginning 7,000 steps.

6.2. Effectiveness of Gaussian Surface Reconstruction(GauS)

For our gaussian surface reconstruction, we selected the SOTA method NeuS [49], SuGaR [15], GaMeS [48], and mesh extraction used in LGM [45] for comparison. We utilized the camera poses of training images to render depth maps from GS and removed the pixels whose depth values are larger than 10. To adjust the parameters of VDBFusion [47] according to the size of input GS, we first calculate a bounding sphere and denote its radius as r. We set voxel size as r/256, SDF truncation value as r/64, and space carving as False in VDBFusion [47]. It is noted that if the camera poses of training images are unavailable, we render depth maps corresponding to the camera trajectory which circles the center of the scene.

We evaluates the visual quality of meshes generated by

Gaussian Surface Reconstruction (GauS) compared to three other popular methods: SuGaR, LGM, and GaMes. Here, we assess how well each approach captures the intricate details and overall shape of the target surface.

Our observations suggest that our GauS achieves a good balance between quality and efficiency. While SuGaR could produce detialed surfaces, it suffered from large episollod, GauS maintains a sharp representation without a significant increase in processing time. Conversely, LGM may prioritize coarse shape with limited training time. GaMes often excels at capturing the local shape, but coarse structures might be underrepresented. GauS strikes a middle ground, offering visually compelling meshes that are faithful to the original surface while remaining free from excessive noise or over-smoothing. This balance between quality and efficiency makes GauS a compelling choice for gaussian surface reconstruction tasks.



Figure 4. Qualitative comparisons of Surface Reconstruction on Blender Dataset

Furthermore, we compared GauS with Sugar specifically on phone-captured data from the BlendedMVS dataset. Here, GauS demonstrated a superior ability to preserve both the fine details and the coarse structure of the object. While Sugar produced very fine-grained surfaces, it often lost some of the larger features and introduced additional artifacts. This highlights GauS's advantage in maintaining a balanced representation, particularly for data with varying levels of detail.

Finally, we extended GauS for application to large scenes, a challenging task for traditional NeRF and NeuS methods. NeRFs typically struggle with reconstructing large scenes due to their inherent limitations in handling global coherence. NeuSs, on the other hand, can handle large scenes but often require significant computational resources and may not always achieve the desired level of de-



Figure 5. Qualitative comparisons between SuGaR and GauS on BlendedMVS Dataset

tail. By design, GauS offers a balance between these two approaches, making it suitable for tackling complex reconstruction problems beyond smaller objects. This opens up new possibilities for applications in areas like architectural modeling, autonomous vehicle navigation, and virtual reality scene creation, where reconstructing large and intricate environments is crucial. In Fig 7, we show the visualization of models extracted by GauS from the 3DGS model trained with appearance embedding[29] and our sky modeling on Mill-19 and UrbanScene3D Datasets.

6.3. Effectiveness of Gaussian Sky Modeling

We evaluate the effectiveness of our Gaussians Sky modeling approach by comparing it with two state-of-the-art methods, 3DGS [23] and GaussianPro [7], on three typical outdoor scenes. Fig. 6 shows representative examples from these scenes. Our method achieves a noticeable improvement in reducing noise and artifacts in the sky regions compared to the baseline methods. We further show quantitative comparisons on the Tanks and Temples dataset[] in Table 1.

In the unbounded scenes (left in Fig. 6), 3DGS [23] suffers from prominent cloud-like artifacts in the sky, which can be distracting and detract from the overall realism of the synthesized view. In contrast, our approach effectively separates the sky from the foreground elements, producing a clean and artifact-free gaussian point cloud.

In the aerial scene (top right in Fig. 6), although GaussianPro [7] also integrates sky masks to reduce floater accumulation in the sky region, it did not explicitly regularize the sky region and propose a proper model for the sky. This results in a partially cloudy reconstruction.

In the street scene (bottom right in Fig. 6), our method generates a coherent street reconstruction, while the base-



Figure 6. Qualitative comparisons of Gaussians Sky modeling on three types of outdoor scenes(unbounded, aerial, street). Ours refer to Gaussian Sky Modelling. Our method effectively separates the sky from the foreground, reducing cloud artifacts and producing clean and realistic foreground representations.



Figure 7. Qualitative results on the mesh extraction of large-scale scenes on Mill-19[46] and UrbanScene3D[30] dataset.

line methods exhibit visible artifacts and inconsistencies in the sky region.

	Barn	Caterpillar	Ignatius	Train	Truck
3DGS	28.10	23.56	22.05	21.10	25.19
GaussianPro	27.66	22.25	20.79	20.72	23.58
Scaffold-GS	28.33	24.00	22.78	22.20	25.57
Mip-splatting	28.30	23.64	22.11	22.15	25.25
3DGS+Ours	28.12	23.37	21.93	22.04	24.85

Table 1. Quantitative assessment on the Tanks and Temples dataset (PSNR).

These examples highlight the effectiveness of our Gaussians Sky modeling technique in improving the overall quality and realism of novel view synthesis for outdoor scenes. By explicitly modeling the sky as a separate component and incorporating semantic priors, our approach successfully mitigates common issues such as cloud artifacts, noise, and color inaccuracies, resulting in more natural and visually compelling renderings.

7. Conclusion

In summary, GauStudio is a flexible and modular framework tailored for 3D Gaussian Splatting techniques. It allows customizable integration of different components for foreground modeling, background representations, and other modules to construct specialized pipelines for diverse 3D scene modeling tasks like reconstruction, editing, and simulation. The framework's key strength is its composability, enabling rapid innovation. GauStudio is complemented by GauS, an efficient module to extract textured meshes from optimized Gaussians across different methods. Overall, this comprehensive yet customizable platform aims to drive advancements in 3D scene modeling through its modular design and versatility.

References

- Henrik Aanæs, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders Bjorholm Dahl. Large-scale data for multiple-view stereopsis. *International Journal of Computer Vision*, 2016. 6
- [2] Jiayang Bai, Letian Huang, Jie Guo, Wen Gong, Yuanqi Li, and Yanwen Guo. 360-gs: Layout-guided panoramic gaussian splatting for indoor roaming. Feb 2024. 4
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2, 6
- [4] David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. Dec 2023. 2, 4
- [5] Hanlin Chen, Chen Li, and Gim Hee Lee. Neusg: Neural implicit surface reconstruction with 3d gaussian splatting guidance. Dec 2023. 4
- [6] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022. 6
- [7] Kai Cheng, Xiaoxiao Long, Kaizhi Yang, Yao Yao, Wei Yin, Yuexin Ma, Wenping Wang, and Xuejin Chen. Gaussianpro: 3d gaussian splatting with progressive propagation. Feb 2024. 2, 4, 6, 7
- [8] Johan Edstedt, Ioannis Athanasiadis, Mårten Wadenbäck, and Michael Felsberg. Dkm: Dense kernelized feature matching for geometry estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 4
- [9] Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, Dejia Xu, and Zhangyang Wang. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. Nov 2023. 4
- [10] Linus Franke, Darius Rückert, Laura Fink, and Marc Stamminger. Trips: Trilinear point splatting for real-time radiance field rendering. Jan 2024. 4
- [11] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022. 1
- [12] Yang Fu, Sifei Liu, Amey Kulkarni, Jan Kautz, Alexei A Efros, and Xiaolong Wang. Colmap-free 3d gaussian splatting. arXiv preprint arXiv:2312.07504, 2023. 4
- [13] Jian Gao, Chun Gu, Youtian Lin, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. Relightable 3d gaussian: Real-time point cloud relighting with brdf decomposition and ray tracing. Nov 2023. 4
- [14] Sharath Girish, Kamal Gupta, and Abhinav Shrivastava. Eagles: Efficient accelerated 3d gaussians with lightweight encodings. Dec 2023. 4
- [15] Antoine Guédon and Vincent Lepetit. Sugar: Surfacealigned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. Nov 2023. 4, 5, 6
- [16] Richard Hartley and Andrew Zisserman. Multiple view ge-

ometry in computer vision. Cambridge university press, 2003. 4

- [17] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. ACM Transactions on Graphics (ToG), 2018. 6
- [18] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces. Nov 2023. 2, 4
- [19] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Cotracker: It is better to track together. arXiv preprint arXiv:2307.07635, 2023. 4
- [20] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, 2006. 5
- [21] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. ACM Transactions on Graphics (ToG), 2013. 5
- [22] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. *arXiv preprint arXiv:2312.02126*, 2023. 4, 5
- [23] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics, 2023. 1, 2, 4, 6, 7
- [24] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. ACM Transactions on Graphics, 2017. 6
- [25] Byeonghyeon Lee, Howoong Lee, Xiangyu Sun, Usman Ali, and Eunbyung Park. Deblurring 3d gaussian splatting. Jan 2024. 4, 5
- [26] Joo Chan Lee, Daniel Rho, Xiangyu Sun, Jong Hwan Ko, and Eunbyung Park. Compact 3d gaussian representation for radiance field. arXiv preprint arXiv:2311.13681, 2023. 4
- [27] Jiabao Li, Yuqi Li, Ciliang Sun, Chong Wang, and Jinhui Xiang. Spec-nerf: Multi-spectral neural radiance fields. arXiv preprint arXiv:2310.12987, 2023. 4
- [28] Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. Gs-ir: 3d gaussian splatting for inverse rendering. Nov 2023.
- [29] Jiaqi Lin, Zhihao Li, Xiao Tang, Jianzhuang Liu, Shiyong Liu, Jiayue Liu, Yangdi Lu, Xiaofei Wu, Songcen Xu, Youliang Yan, and Wenming Yang. Vastgaussian: Vast 3d gaussians for large scene reconstruction. Feb 2024. 7
- [30] Liqiang Lin, Yilin Liu, Yue Hu, Xingguang Yan, Ke Xie, and Hui Huang. Capturing, reconstructing, and simulating: the urbanscene3d dataset. In *ECCV*, 2022. 8
- [31] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. arXiv preprint arXiv:2312.00109, 2023. 2, 4, 5, 6
- [32] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 2021. 1, 6

- [33] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 2021. 2, 3
- [34] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. ACM transactions on graphics (TOG), 2022. 1, 3, 4
- [35] KL Navaneet, Kossar Pourahmadi Meibodi, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. Compact3d: Compressing gaussian splat radiance field models with vector quantization. Nov 2023. 4
- [36] Shunsuke Saito, Gabriel Schwartz, Tomas Simon, Junxuan Li, and Giljoo Nam. Relightable gaussian codec avatars. Dec 2023. 4
- [37] Erik Sandström, Yue Li, Luc Van Gool, and Martin R Oswald. Point-slam: Dense neural point cloud-based slam. In Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023. 1
- [38] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition, 2019. 4
- [39] Johannes L Schonberger and Jan-Michael Frahm. Structurefrom-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016. 2, 4
- [40] Jin-Chuan Shi, Miao Wang, Hao-Bin Duan, and Shao-Hua Guan. Language embedded 3d gaussians for openvocabulary scene understanding. Nov 2023. 4
- [41] Yahao Shi, Yanmin Wu, Chenming Wu, Xing Liu, Chen Zhao, Haocheng Feng, Jingtuo Liu, Liangjun Zhang, Jian Zhang, Bin Zhou, Errui Ding, and Jingdong Wang. Gir: 3d gaussian inverse rendering for relightable scene factorization. Dec 2023. 4
- [42] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021. 4
- [43] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Splatter image: Ultra-fast single-view 3d reconstruction. Dec 2023. 4
- [44] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. arXiv preprint arXiv:2402.05054, 2024. 4
- [45] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. arXiv preprint arXiv:2402.05054, 2024. 6
- [46] Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction of largescale nerfs for virtual fly-throughs. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2022. 8
- [47] Ignacio Vizzo, Tiziano Guadagnino, Jens Behley, and Cyrill Stachniss. Vdbfusion: Flexible and efficient tsdf integration of range sensor data. *Sensors*, 2022. 5, 6

- [48] Joanna Waczyńska, Piotr Borycki, Sławomir Tadeja, Jacek Tabor, and Przemysław Spurek. Games: Mesh-based adapting and modification of gaussian splatting. *arXiv preprint arXiv:2402.01459*, 2024. 6
- [49] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689, 2021. 6
- [50] Qianqian Wang, Yen-Yu Chang, Ruojin Cai, Zhengqi Li, Bharath Hariharan, Aleksander Holynski, and Noah Snavely. Tracking everything everywhere all at once. *arXiv preprint arXiv:2306.05422*, 2023. 4
- [51] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 2004. 6
- [52] Tong Wu, Jiaqi Wang, Xingang Pan, Xudong Xu, Christian Theobalt, Ziwei Liu, and Dahua Lin. Voxurf: Voxel-based efficient and accurate neural surface reconstruction. arXiv preprint arXiv:2208.12697, 2022. 1
- [53] Zirui Wu, Tianyu Liu, Liyi Luo, Zhide Zhong, Jianteng Chen, Hongmin Xiao, Chao Hou, Haozhe Lou, Yuantao Chen, Runyi Yang, et al. Mars: An instance-aware, modular and realistic simulator for autonomous driving. In *CAAI International Conference on Artificial Intelligence*. Springer, 2023. 6
- [54] Dejia Xu, Ye Yuan, Morteza Mardani, Sifei Liu, Jiaming Song, Zhangyang Wang, and Arash Vahdat. Agg: Amortized generative 3d gaussians for single image to 3d. Jan 2024. 4
- [55] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Pointnerf: Point-based neural radiance fields. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022. 1
- [56] Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians for modeling dynamic urban scenes. Jan 2024. 4
- [57] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. *arXiv preprint arXiv:2401.10891*, 2024. 4
- [58] Ziyi Yang, Xinyu Gao, Yangtian Sun, Yihua Huang, Xiaoyang Lyu, Wen Zhou, Shaohui Jiao, Xiaojuan Qi, and Xiaogang Jin. Spec-gaussian: Anisotropic view-dependent appearance for 3d gaussian splatting. arXiv preprint arXiv:2402.15870, 2024. 3, 5
- [59] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P Srinivasan, Richard Szeliski, Jonathan T Barron, and Ben Mildenhall. Bakedsdf: Meshing neural sdfs for realtime view synthesis. arXiv preprint arXiv:2302.14859, 2023. 3, 4
- [60] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. Nov 2023. 2, 6
- [61] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. Advances in neural information processing systems, 2022. 4

- [62] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018. 6
- [63] Xiaoyu Zhou, Zhiwei Lin, Xiaojun Shan, Yongtao Wang, Deqing Sun, and Ming-Hsuan Yang. Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. Dec 2023. 4
- [64] Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. Fsgs: Real-time few-shot view synthesis using gaussian splatting. Dec 2023. 4
- [65] Zi-Xin Zou, Zhipeng Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Yan-Pei Cao, and Song-Hai Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. Dec 2023. 2, 4, 5