

Article

Complex Illumination-Aware 3D Gaussian Reconstruction for Uncooperative Space Objects

Ziang Qu ^{1,2} , Zhang Zhang ^{1,2}, Ruiqi Xun ^{1,2}, Junlan Zhou ^{1,2} and Liang Chang ^{1,2,*}¹ Innovation Academy for Microsatellites of Chinese Academy of Sciences, Shanghai 201304, China; quziang24@mailsucas.ac.cn (Z.Q.)² University of Chinese Academy of Sciences, Beijing 101408, China

* Correspondence: changl@microsat.com

Abstract

High-precision 3D reconstruction of non-cooperative space targets is a critical technology for on-orbit servicing (OOS) and situational awareness, driven by the growing number of OOS missions. However, traditional visual algorithms struggle to acquire accurate geometric information due to the unique high-dynamic-range lighting and strong specular reflections characteristic of the space environment. This paper proposes Space-Gaussian, a compact 3D Gaussian reconstruction method tailored for complex lighting environments. Built upon the 3D Gaussian Splatting (3DGS) framework, the method incorporates a physically based rendering pipeline and a microfacet bidirectional reflectance distribution function model. By decoupling geometric structure from material properties and utilizing deferred rendering, it effectively suppresses geometric artifacts and specular highlights arising from non-Lambertian surface reflections. Comparative experiments on a high-fidelity simulation dataset demonstrate that Space-Gaussian outperforms mainstream methods—including Neural Radiance Fields (NeRF), Instant-NGP, GaussianShader, and 3DGS—in geometric reconstruction accuracy, novel view synthesis quality, and real-time rendering. On our self-created dataset, our approach achieves a significant performance boost over existing 3DGS methods. The results highlight its potential for high-fidelity, real-time 3D perception on resource-constrained spacecraft platforms.

Keywords: 3D Gaussian Splatting; 3D reconstruction; non-cooperative space targets; on-orbit servicing

1. Introduction

With the boom in commercial aerospace and deep-space exploration, the increasing number of space objects has dramatically escalated the potential for spacecraft collisions, presenting significant challenges for space situational awareness systems [1]. Consequently, high-precision 3D perception of non-cooperative space targets has become critical to the success of on-orbit servicing (OOS) and active debris removal (ADR) missions [2]. However, the unique high-dynamic-range (HDR) lighting in outer space, combined with the strong non-Lambertian reflectance characteristics of spacecraft surfaces such as multi-layer insulation (MLI) and solar panels, results in a deep coupling between geometric structure and surface appearance [3]. When facing such extreme lighting conditions, traditional vision algorithms often fail to distinguish between physical structures and specular artifacts, and there is currently a lack of standard datasets and testing benchmarks specifically designed



Academic Editor: Hyun-Ung Oh

Received: 31 January 2026

Revised: 1 March 2026

Accepted: 9 March 2026

Published: 10 March 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

for illumination perception, making it difficult to quantify the robustness of different algorithms under extreme lighting conditions [4]. Consequently, the reconstructed models not only lack photorealism but also suffer from severe geometric distortions, failing to meet the stringent requirements of space operations.

In the current landscape of reconstruction technologies, neural radiance fields (NeRF) [5] have achieved impressive photorealistic rendering, yet their substantial training and inference overheads struggle to meet the strict real-time constraints of spaceborne platforms. As an alternative, 3D Gaussian Splatting (3DGS) [6] has emerged as a new paradigm that facilitates real-time rendering through explicit representations, although it remains dependent on spherical harmonics for color fitting. Because this approach inherently assumes low-frequency lighting, it struggles when encountering the sharp specular highlights typical of space targets. In these instances, 3DGS tends to generate numerous fragmented pseudo-geometries to forcibly fit high-frequency signals, leading to geometric overfitting that destroys surface smoothness and renders reconstruction results misleading for physical interactions. Therefore, integrating physically grounded lighting processing into Gaussian splatting while maintaining its real-time performance represents a critical bottleneck that must be addressed.

In light of this, we propose Space-Gaussian, a physics-aware reconstruction framework tailored for complex lighting environments. Our core motivation is to explicitly decouple geometry, material, and lighting within the lightweight Gaussian representation by introducing physical constraints, thereby eliminating geometric ambiguities caused by specularities. Specifically, we deeply integrate a physically based rendering (PBR) [7] pipeline into the 3DGS framework. We design an anisotropic bidirectional reflectance distribution function (BRDF) model based on the microfacet theory to replace spherical harmonics [8], endowing Gaussian primitives with physical material properties. Additionally, we introduce an efficient deferred rendering strategy [9], achieving precise computation of interactions between environmental lighting and materials while preserving real-time performance.

The main contributions of this paper are summarized as follows:

1. We propose a physically consistent reconstruction approach specifically tailored for non-cooperative space targets. By decoupling geometry from material properties within the 3DGS framework, our method effectively mitigates geometric artifacts caused by intense specular reflections.
2. We establish a robust PBR neural rendering scheme for 3DGS. By incorporating a microfacet BRDF model, we enable explicit Gaussian representations to accurately simulate complex light-material interactions on highly reflective surfaces.
3. We introduce a solar orientation initialization strategy as a directional physical prior. This approach leverages the known sun direction in space to provide a deterministic constraint for light-source modeling, significantly improving optimization stability and convergence.
4. We design an efficient deferred rendering pipeline. This architecture allows for the high-fidelity integration of complex shading and PBR attributes while preserving the real-time rendering advantages of the 3DGS paradigm.

2. Related Work

With the increasingly complex requirements of on-orbit perception tasks, 3D reconstruction technology for non-cooperative space targets has undergone a paradigm shift from geometry-driven to data-driven approaches. This section systematically reviews research progress in related fields from three main dimensions: traditional geometric vision methods, implicit representations based on NeRF, and the emerging 3DGS technology. We focus on

analyzing the applicability and limitations of each method under special constraints such as extreme space lighting, strong specular reflections, and limited computational resources.

2.1. Classical Geometric 3D Reconstruction

The core objective of 3D reconstruction is to recover the 3D geometric structure of a scene from multiview images or sensor data. As early as the 1960s, Roberts first proposed a theoretical framework for inferring the 3D structure of objects using 2D images [10]. Traditional mainstream methods relied primarily on the multiple view geometry theory, specifically including structured light methods, stereo matching, and structure from motion (SfM) techniques [11]. With the development of computer vision, scholars such as Hartley and Zisserman systematically expounded the theory of multiple view geometry [12], laying a solid mathematical foundation for sparse reconstruction based on feature matching. Stanford University's "Digital Michelangelo Project" [13], by utilizing laser scanners of high precision and arrays of multiple cameras to achieve modeling of high fidelity at the level of artifacts, became a landmark achievement of the period.

In the field of space applications, the reconstruction of non-cooperative space targets is mainly applied to tasks such as OOS, ADR, and rendezvous and proximity operations (RPO). Space agencies represented by NASA and ESA initially adopted methods with active sensors or strong prior constraints. For instance, NASA's Orbital Express [14] mission and ESA's RemoveDEBRIS [15] mission employed a combination of binocular vision systems and LiDAR to achieve a 3D shape estimation of targets. Subsequently, in the ClearSpace-1 [16] mission, ESA further introduced model-based image registration techniques, performing pose estimation using pre-stored target CAD models to enhance on-orbit capture precision. Traditional SfM and multiview stereo (MVS) methods encounter several formidable obstacles under passive optical observation conditions. Initially, the lack of distinctive textures poses a significant hurdle. The multi-layer insulation typically found on spacecraft surfaces exhibits a uniform gold or silver finish with a scarcity of extractable corner features. Simultaneously, solar panels often feature periodic, repetitive patterns that frequently trigger feature-matching ambiguities. Furthermore, the prevalence of intense specular reflections in the space environment complicates the process. Because metallic surfaces and solar cells are non-Lambertian, they disrupt the brightness constancy assumption essential for multiview geometry. This violation leads to substantial noise, large-scale holes, or even total failure in the resulting point clouds. High computational complexity remains a critical constraint. The heavy processing load required for dense, pixel-level depth estimation and fusion makes it nearly impossible to achieve real-time performance on resource-constrained onboard processing units.

2.2. Advances in Neural Implicit Fields and Volume Rendering

Early data-driven 3D reconstruction methods, such as 3D-R2N2 [17] and PointNet [18], primarily used deep neural networks to directly regress 3D geometric structures such as voxels [19] or point clouds, laying the foundation for geometric perception using deep learning.

However, it was not until the introduction of NeRF [5] in 2020 that 3D reconstruction truly marked the transition from discrete geometric representations to a continuous neural rendering paradigm. NeRF abandons explicit structures like meshes or voxels, instead implicitly modeling the volume density and color of any point in the scene space through a multi-layer perceptron (MLP). Combined with differentiable volume-rendering technology, NeRF can self-supervise the learning of scene representations from sparse 2D images via backpropagation, achieving photorealistic novel view synthesis, with the image processing pipeline shown in Figure 1.




Figure 1. Illustration of the image processing pipeline in NeRF. The diagonal line represents a camera ray, with the dark blue circles denoting discrete 3D points sampled along this path. The thin blue arrows indicate that the 5D input vector, comprising 3D spatial coordinates (x, y, z) and 2D viewing direction (θ, ϕ) .

Despite the immense success of NeRF, subsequent research identified limitations in the original formulation, including aliasing effects, slow training speeds, and weak lighting decoupling capabilities. To address these issues, numerous improvements have been proposed. For instance, Mip-NeRF 360 [20] replaced ray sampling with cone casting to effectively handle anti-aliasing across multiple scales, while Ref-NeRF [21] introduced reflection direction parameterization and integrated directional encoding, significantly enhancing the rendering quality of specular surfaces.

Although NeRF and its variants have achieved breakthroughs in visual quality, applying them to on-orbit space perception still faces significant bottlenecks: First is slow inference speed. NeRF imaging relies on ray marching, where rendering a single pixel necessitates hundreds of MLP queries. Realizing real-time rendering at high resolutions typically demands substantial GPU computational power. Second is prolonged training time. Traditional NeRF training often requires hours or even days, failing to meet the strict timeliness requirements of instant perception and decision making for space missions. Although Instant-NGP [22], proposed by Google, significantly enhances training efficiency using multi-resolution hash encoding, it still necessitates a large amount of VRAM for training. A final concern pertains to geometric uncertainty. In the absence of explicit geometric constraints, NeRF is susceptible to producing spurious volumetric artifacts when observations are limited by sparse views or extensive baselines. Such inaccuracies jeopardize the safety of proximity operations for spacecrafts.

In the space domain, the CS Group, in collaboration with CNES, has developed a neural rendering algorithm optimized for space scenarios [23]. Targeting non-cooperative space objects, this study proposes a NeRF method that jointly estimates attitude and performs 3D reconstruction from simulated monochromatic optical image sequences. By incorporating multiscale hash encoding to accelerate training and logarithmic tone mapping to handle high-dynamic-range monochromatic images, as well as employing a coarse-to-fine attitude registration strategy and a unified rotation parameterization, the algorithm successfully achieves robust reconstruction of targets such as the Galileo satellite, even under adverse conditions lacking diffuse lighting and prior attitude information.

2.3. 3D Gaussian Splatting and Its Space Applications

To balance reconstruction quality and rendering speed, 3DGS [6], proposed by the INRIA team in 2023, has garnered widespread attention and rapidly emerged as a mainstream paradigm for next-generation real-time radiance field reconstruction. Unlike NeRF, which relies on MLP to implicitly model volume density, 3DGS employs massive, dis-

crete, anisotropic Gaussian ellipsoids as an explicit scene representation, as illustrated in Figure 2. Each Gaussian primitive carries not only geometric attributes such as position, rotation, and scale but also opacity and spherical harmonics coefficients used to fit view-dependent colors.




Figure 2. Representation of 3D Gaussian.

In terms of rendering mechanism, 3DGS discards the computationally expensive ray marching volume integration, instead designing a fully differentiable point-based splatting pipeline. This pipeline utilizes efficient tile-based rasterization technology to project 3D Gaussians onto the 2D image plane and performs α -blending based on depth sorting, thereby achieving real-time rendering speeds of over 100 FPS at 1080p resolution. In addition, 3DGS introduces an adaptive density control strategy, which, guided by gradients during training, automatically executes split or clone operations in regions rich in geometric details to increase the number of Gaussians and performs pruning in empty or occluded areas. This achieves a high-fidelity reconstruction of complex scenes while maintaining a high efficiency. The emergence of this approach addresses the persistent computational challenges of implicit fields by accelerating training and inference speeds. This breakthrough paves the way for autonomous perception and navigation on spacecrafts where time is a primary constraint.

However, applying standard 3DGS directly to the reconstruction of non-cooperative space targets is challenging. The difficulty is in fitting high-frequency reflections. Original 3DGS employs spherical harmonics to fit view-dependent color variations, while SH excels at simulating low-frequency diffuse reflections, approximating the sharp high-frequency specular highlights generated by space targets like solar panels and metal antennas that require extremely high-order coefficients. This not only leads to explosive parameter growth but also tends to induce ringing artifacts similar to those in signal processing. The coupling of geometry and appearance is another difficulty. To forcibly fit sudden color changes in highlight regions, 3DGS often generates numerous fragmented Gaussian ellipsoids with elongated shapes and chaotic distributions. This geometric overfitting results in the reconstructed satellite surfaces being riddled with artifacts, severely disrupting the object's true geometric structure.

Although recent advancements in 3DGS have integrated PBR, such as GaussianShader [24], most remain restricted to terrestrial indoor scenes and fail to account for the unique illumination characteristics of space, including HDR environment and low visibility. The Space-Gaussian method proposed in this paper aims to fill this gap. By constructing a complete PBR-3DGS pipeline combined with sparsification and compression strategies, it achieves an optimal balance among model lightweighting, real-time rendering, and high-fidelity reconstruction.

3. Method

3.1. System Overview

The Space-Gaussian method proposed in this paper aims to address the geometry–appearance coupling problem of space targets under strong illumination through explicit physical constraints. The overall workflow of the system consists of the following main stages: First, the physically-aware Gaussian representation upgrade, where we extend standard 3D Gaussian primitives into entities with physical properties, endowing them with microfacet material attributes such as albedo, metallicity, roughness, and normals to replace spherical harmonic color coefficients; second, G-Buffer-based deferred rendering, where we design a lightweight deferred rendering pipeline that rasterizes Gaussian attributes into screen-space geometry buffers, followed by pixel-level shading combined with a microfacet BRDF model and environmental lighting; thirdly, initialization of sunlight direction since sunlight is the primary light source in the scene; and finally, a two-stage robust optimization, employing a geometry warm-up followed by joint optimization strategy, utilizing the geometric fitting capability of 3DGS to provide initial values, and then correcting specular artifacts through the physical rendering pipeline. The flow mentioned above is presented in Figure 3.




Figure 3. An overview of the proposed framework.

This method is built upon the 3DGS [6] framework. For clarity, we first briefly review the mathematical representation of standard 3DGS. In standard 3DGS, each Gaussian is modeled as an ellipsoid characterized by a 3D position p_k and a covariance matrix Σ , defined as:

$$\mathcal{G}(p) = \exp\left(-\frac{1}{2}(p - p_k)^T \Sigma^{-1}(p - p_k)\right). \quad (1)$$

During the rendering process, 3D Gaussians are projected into 2D Gaussians in the camera space. The transformed 2D covariance matrix is approximated as: $\Sigma' = JW\Sigma W^T J^T$, where W is the view transformation matrix and J is the Jacobian of the perspective projection. Each Gaussian is also assigned an opacity o and a view-dependent color c_i represented by Spherical Harmonics (SH). The final pixel color C is computed by blending the Gaussians in front-to-back depth order:

$$C = \sum_{i=1}^N c_i T_i \alpha_i, \quad \text{where} \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (2)$$

where α_i is computed by multiplying the opacity o_i with the Gaussian weight, which is determined by the 2D covariance matrix Σ'_i .

3.2. Physically Based Deferred Rendering

In space environments, lighting conditions are complex and variable, and target surfaces often exhibit high-frequency reflection characteristics. To more accurately simulate such complex lighting phenomena, we introduce physically based deferred rendering technology. Adopting the BRDF model proposed by Disney in 2012 [8], each Gaussian is associated with a set of material-related properties, including albedo $\lambda \in [0, 1]^3$, metallicity $m \in [0, 1]$, and roughness $r \in [0, 1]$. The normal vector $\mathbf{n} \in [0, 1]^3$ can be computed from the tangent vectors of each 2D Gaussian via $\mathbf{n} = \mathbf{t}_u \times \mathbf{t}_v$.

3.2.1. Microfacet BRDF Modeling

We adopt the common microfacet form of the Disney principled BRDF, decomposing reflection into diffuse and specular components. For any pixel with an outgoing direction $\mathbf{v}$ that is pointing to the camera and incident direction $\mathbf{l}$ that is pointing to the light source, let the half-vector be $\mathbf{h} = \frac{\mathbf{v} + \mathbf{l}}{\|\mathbf{v} + \mathbf{l}\|}$. Then, the BRDF is written as:

$$f_r(\mathbf{l}, \mathbf{v}) = k_d f_d(\mathbf{l}, \mathbf{v}) + k_s f_s(\mathbf{l}, \mathbf{v}), \quad (3)$$

where k_d, k_s are energy distribution coefficients. To represent the transition between dielectric and metallic states, we employ the metallicity parameter m to interpolate the material properties. Consequently, the specular base reflectivity is defined as follows:

$$F_0 = (1 - m) \cdot 0.04 \cdot \mathbf{1} + m \cdot \lambda, \quad (4)$$

where λ is the albedo. The specular term employs the GGX normal distribution and Smith geometric shadowing term, and uses the Schlick approximation to calculate the Fresnel term:

$$f_s(\mathbf{l}, \mathbf{v}) = \frac{D(\mathbf{n}, \mathbf{h}) G(\mathbf{n}, \mathbf{l}, \mathbf{v}) F(\mathbf{v}, \mathbf{h})}{4(\mathbf{n} \cdot \mathbf{l})(\mathbf{n} \cdot \mathbf{v}) + \epsilon}, \quad (5)$$

$$F(\mathbf{v}, \mathbf{h}) = F_0 + (1 - F_0)(1 - \mathbf{v} \cdot \mathbf{h})^5, \quad (6)$$

where ϵ is a small constant to prevent numerical instability. Roughness r controls the profile of the microfacet distribution; we set $\alpha = r^2$ and calculate $D(\cdot)$ and $G(\cdot)$ accordingly.

The diffuse term uses the energy-conserving Lambertian form:

$$f_d(\mathbf{l}, \mathbf{v}) = \frac{\lambda}{\pi}. \quad (7)$$

Simultaneously, to ensure energy conservation, we set $k_s = F(\mathbf{v}, \mathbf{h})$ and $k_d = (1 - k_s)(1 - m)$.

3.2.2. Lighting Discretization and Outgoing Radiance

The rendering equation is given by:

$$L_o(\mathbf{v}) = \int_{\Omega} f_r(\mathbf{l}, \mathbf{v}) L_i(\mathbf{l}) (\mathbf{n} \cdot \mathbf{l}) d\omega. \quad (8)$$

Considering that lighting in space scenes is typically dominated by directional sunlight, we discretize the integral into a summation over a finite number of light sources:

$$L_o(\mathbf{v}) \approx \sum_{q=1}^Q f_r(\mathbf{l}_q, \mathbf{v}) \mathbf{I}_q \max(0, \mathbf{n} \cdot \mathbf{l}_q) + L_{\text{amb}}, \quad (9)$$

where $\mathbf{I}_q$ represents the RGB intensity of the q -th light source, and L_{amb} represents the ambient and indirect light term.

3.2.3. Deferred Rendering and Gaussian Blending

To decouple physically based shading from the forward splatting process of 3DGS, we construct a material G-buffer in the screen space. Specifically, for each Gaussian i , in addition to position and covariance, we maintain material and geometric attributes:

$$\mathcal{A}_i = \{\lambda_i, m_i, r_i, \mathbf{n}_i, o_i\}. \quad (10)$$

During splatting, we follow the forward alpha blending of Equation (2), accumulating the contributions of different Gaussians to the same pixel in front-to-back depth order. Unlike direct color blending, the deferred rendering stage first blends to obtain pixel-level material/normal attributes (denoted as $\bar{\lambda}, \bar{m}, \bar{r}, \bar{\mathbf{n}}$), and then performs unified BRDF shading calculations at the pixel level:

$$\bar{x} = \frac{\sum_{i=1}^N x_i w_i}{\sum_{i=1}^N w_i + \epsilon}, \quad w_i = T_i \alpha_i, \quad (11)$$

where x represents any attribute requiring blending (e.g., λ, m, r), and w_i is the effective weight of that Gaussian for the pixel. For normals, we normalize them after blending:

$$\bar{\mathbf{n}} \leftarrow \frac{\bar{\mathbf{n}}}{\|\bar{\mathbf{n}}\| + \epsilon}. \quad (12)$$

Finally, the blended material attributes are substituted back into the discretized rendering equation to compute the final pixel color $\mathbf{C}_{\text{final}}$:

$$\mathbf{C}_{\text{final}} = \sum_{q=1}^Q f_r(l_q, \mathbf{v}; \bar{\lambda}, \bar{m}, \bar{r}, \bar{\mathbf{n}}) I_q \max(0, \bar{\mathbf{n}} \cdot \mathbf{l}_q) + L_{\text{amb}}. \quad (13)$$

3.2.4. Parameter Constraints and Differentiable Optimization

Since material parameters must satisfy value range constraints ($\lambda \in [0, 1]^3, m \in [0, 1], r \in [0, 1]$), we represent them as Sigmoid mappings of unconstrained variables:

$$\lambda = \sigma(\hat{\lambda}), \quad m = \sigma(\hat{m}), \quad r = \sigma(\hat{r}), \quad (14)$$

This mapping mechanism effectively guarantees numerical stability during the optimization process. For the normal $\mathbf{n}$, we ensure it strictly satisfies the unit geometric constraint through cross-product of tangent vectors followed by normalization.

3.3. Sunlight Direction Initialization

In the space environment, the scene is illuminated by multiple sources, including the sun, Earth-reflected light, and starlight. Among these, the sun acts as the dominant directional light source with extremely high intensity. Therefore, in our discretized lighting model, we designate the primary light-source component ($q = 1$) to explicitly represent the sunlight direction $\mathbf{l}_{\text{sun}}$ in the world coordinate system. To maintain the unit vector constraint during optimization, $\mathbf{l}_{\text{sun}}$ is parameterized using spherical coordinates (θ, ϕ) :

$$\mathbf{l}_{\text{sun}} = [\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta]^T. \quad (15)$$

Before the joint optimization, we perform an initialization of the sunlight direction to provide a reliable starting point, which is crucial for decoupling geometry and material properties. The initialization leverages the intense specular highlights observed in the captured images and the known camera poses.

First, we identify the pixel region Ω_{spec} exhibiting high intensities across the multiview image set $\mathcal{I}$. For an image $I_k \in \mathcal{I}$ captured at camera pose k , we extract the 2D pixel coordinates $\mathbf{u}_p = [u, v]^T$ of the specular highlights by applying an intensity threshold.

Given the camera intrinsic matrix $\mathbf{K}$ and the extrinsic matrix $\mathbf{T}_k = [\mathbf{R}_k | \mathbf{t}_k]$, where $\mathbf{R}_k \in SO(3)$ is the rotation matrix and $\mathbf{t}_k \in \mathbb{R}^3$ is the translation vector from the world coordinate system to the camera coordinate system, we can determine the view direction $\mathbf{v}_p$ in the world space. The 3D ray direction $\mathbf{d}_c$ in the camera coordinate system corresponding to pixel $\mathbf{u}_p$ is computed as:

$$\mathbf{d}_c = \mathbf{K}^{-1}[u, v, 1]^T. \quad (16)$$

The view direction $\mathbf{v}_p$ pointing from the target surface to the camera optical center in the world coordinate system is then given by:

$$\mathbf{v}_p = -\frac{\mathbf{R}_k^{-1}\mathbf{d}_c}{\|\mathbf{R}_k^{-1}\mathbf{d}_c\|}. \quad (17)$$

Simultaneously, we approximate the surface normal $\mathbf{n}_p$ at the corresponding 3D point from the coarse geometry obtained during the warm-up stage. Assuming the specular highlight is generated by an ideal mirror reflection on the microfacet, the incident sunlight direction $\mathbf{l}_p$ for a specific pixel p can be derived using the reflection law:

$$\mathbf{l}_p = 2(\mathbf{n}_p \cdot \mathbf{v}_p)\mathbf{n}_p - \mathbf{v}_p. \quad (18)$$

To obtain a robust global estimation of the sunlight direction, we aggregate the calculated incident directions from all valid specular pixels across multiple views. The initial sunlight direction $\mathbf{l}_{\text{init}}$ is computed as the normalized average of these vectors:

$$\mathbf{l}_{\text{init}} = \frac{\sum_k \sum_{p \in \Omega_{\text{spec}}^k} \mathbf{l}_p}{\left\| \sum_k \sum_{p \in \Omega_{\text{spec}}^k} \mathbf{l}_p \right\|}, \quad (19)$$

where Ω_{spec}^k denotes the set of specular pixels in the k -th image. This physics-based geometric initialization effectively mitigates the ambiguity between geometry and lighting during the early stages of training, ensuring a physically plausible convergence.

3.4. Loss Function

To learn the geometric structure and surface material properties of the scene, we design a total loss function incorporating geometric constraints:

$$\mathcal{L} = \mathcal{L}_c + \lambda_n \mathcal{L}_n + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}} + \lambda_{\text{tone}} \mathcal{L}_{\text{tone}}. \quad (20)$$

$\mathcal{L}_c$ is the RGB reconstruction loss, defined as a weighted combination of the $\mathcal{L}_1$ photometric loss and the D-SSIM structural similarity loss:

$$\mathcal{L}_c = (1 - \lambda) \mathcal{L}_1 + \lambda \mathcal{L}_{\text{D-SSIM}}, \quad (21)$$

where λ is a balancing coefficient. Following the original 3DGS implementation [6], we set $\lambda = 0.2$.

$\mathcal{L}_n = 1 - \tilde{\mathbf{n}}^T \mathbf{n}$ is the normal consistency loss. This term aims to promote alignment between Gaussian distributions and the geometric surface by minimizing the cosine difference between the rendered normal $\mathbf{n}$ obtained via alpha blending and the pseudo-ground

truth surface normal $\tilde{n}$. To obtain $\tilde{n}$, we first render the depth map D using the same alpha-blending scheme as color rendering:

$$D = \sum_{i=1}^N d_i T_i \alpha_i, \quad (22)$$

where d_i is the depth of the i -th Gaussian center relative to the camera. The pseudo-ground truth normal $\tilde{n}$ is then derived from the spatial gradients of the depth map in screen space (u, v) :

$$\tilde{n} = \frac{\left[-\frac{\partial D}{\partial u}, -\frac{\partial D}{\partial v}, 1 \right]^T}{\left\| \left[-\frac{\partial D}{\partial u}, -\frac{\partial D}{\partial v}, 1 \right]^T \right\|}. \quad (23)$$

$\mathcal{L}_{\text{smooth}} = \|\nabla n\| \exp(-\|\nabla C_{\text{gt}}\|)$ is the edge-aware normal smoothness loss, intended to regularize normal variations in weakly textured regions.

Additionally, to address HDR issues in strong space lighting environments, we introduce a tone-mapping loss $\mathcal{L}_{\text{tone}}$. The standard $\mathcal{L}_1$ loss is overly sensitive to numerical errors in highlight regions, which can lead to the optimization being dominated by specular highlights and overexposure artifacts. We apply a tone-mapping operator before calculation to balance gradients between highlight and shadow regions:

$$\mathcal{L}_{\text{tone}} = \|\mathcal{T}(C) - \mathcal{T}(C_{\text{gt}})\|_1, \quad \text{where} \quad \mathcal{T}(x) = \frac{x}{1+x}. \quad (24)$$

The optimization process utilizes the Adam optimizer to jointly refine the complete set of Gaussian parameters, encompassing geometric attributes—position, orientation, scale, and opacity—as well as physical material properties including diffuse albedo, roughness, metallicity, and surface normals. This holistic approach ensures that all spatial and material components converge within a unified framework.

4. Experiments

4.1. Dataset and Experimental Setup

4.1.1. Dataset

To validate the proposed method, we constructed a high-fidelity, physics-based simulation dataset as shown in Figure 4. We selected high-precision 3D satellite models—specifically CubeSat, Sentinel-6, and Starlink—publicly released by agencies such as the ESA and the NASA and utilized the Blender rendering engine to generate synthetic image data with physical lighting attributes. To ensure the physical accuracy of the rendering parameters, the material properties used in our Blender simulation are directly sourced from the official 3D spacecraft models provided by the NASA and ESA. These official models are meticulously configured to reflect the authentic optical characteristics and physical materials of real-world satellites, thereby providing a reliable physical foundation for our synthetic dataset and minimizing the domain gap with actual space observations. In terms of scene configuration, we simulated realistic on-orbit observation tasks, constructing a LEO environment containing an Earth background and setting the observer satellite to execute fly-around observation tasks of the target satellite along specific orbits. During data acquisition, the system precisely recorded the absolute trajectory of the observer in the world coordinate system and the camera intrinsic matrix. However, given that this study aims to solve the 3D reconstruction problem of non-cooperative space targets, the relative pose between the two satellites is difficult to obtain in space, so the relative pose is not treated as known data.




Figure 4. Some image collages from our dataset.

Addressing the background interference in on-orbit images of non-cooperative targets, this study implemented a background suppression and target extraction strategy. Considering the diversity and complexity of space backgrounds like Earth, deep space and stars, we employed the classic GrabCut image segmentation algorithm [25] to separate the foreground as shown in Figure 5. GrabCut is an iterative foreground extraction technique based on Graph Cut. Its core idea is to model the color distributions of the foreground and background separately by constructing Gaussian mixture models. The algorithm first uses the initial bounding box—which, in our implementation, is automatically generated using a threshold-based saliency detection method to ensure reproducibility—as a weak supervisory signal to partition the image region into definite background, probable foreground, and undefined regions. Subsequently, by constructing a Markov random field energy function, the pixel region assignment problem is transformed into a minimum cut optimization problem in the graph theory. During the iteration process, the algorithm utilizes alternating updates of GMM parameters and pixel labels to progressively optimize the segmentation boundary until the energy function converges. Compared with dee-learning methods, GrabCut achieves precise fitting of the satellite body edges using statistical features of color and texture without requiring large-scale annotated data training. Through this preprocessing step, background-removed satellite images were generated, retaining only the texture and geometric information of the target body, which significantly improved the reconstruction effects of subsequent algorithms.

4.1.2. Baselines

To comprehensively evaluate the performance of the proposed algorithm, we selected the following four representative 3D reconstruction methods as comparison baselines:

- NeRF [5]: As a pioneering work based on implicit neural representation, NeRF utilizes MLP to map spatial coordinates and viewing directions to volume density and color, synthesizing novel view images through differentiable volume rendering techniques. Although it established high-quality standards for continuous scene representation, its computational cost is extremely high due to expensive viewing ray sampling during training and inference, resulting in slow speeds.
- Instant-NGP [22]: This method introduces a hybrid architecture combining multi-resolution hash encoding with tiny MLPs. By leveraging the advantages of GPU parallel computing, Instant-NGP greatly alleviates the memory access bottleneck of neural networks, increasing training and inference speeds by several orders of magnitude while maintaining high reconstruction quality. It is a representative of fast neural reconstruction frameworks.

- 3DGS [6]: This is a real-time rendering method based on explicit representation. Unlike implicit fields, 3DGS uses a set of anisotropic 3D Gaussian ellipsoids to represent the scene, combined with an efficient differentiable rasterization pipeline, achieving real-time rendering frame rates and photorealistic synthesis quality. It is currently the mainstream solution balancing reconstruction quality and rendering speed and also the primary baseline for improvement in this paper.
- GaussianShader [24]: This method integrates a simplified physically based rendering pipeline into the 3DGS framework to handle reflective surfaces. It serves as a representative baseline for PBR-enhanced 3DGS methods, allowing us to evaluate the effectiveness of our space-specific material and lighting modeling compared with general-purpose PBR extensions.




Figure 5. Data preprocessing pipeline.

All baseline methods are implemented based on official open-source code. To ensure a fair and objective comparison, all baseline methods were configured using the optimal hyperparameter settings recommended by their respective authors for novel view synthesis tasks. This approach aims to evaluate each method at its intended peak performance, thereby providing a reliable baseline for assessing the improvements introduced by our proposed framework.

4.1.3. Evaluation Metrics

To comprehensively measure the performance of 3D reconstruction algorithms in non-cooperative target scenarios, this paper employs metrics from two dimensions: image synthesis quality and resource computational overhead.

In terms of image synthesis quality, we selected three mainstream metrics. First is the peak signal-to-noise ratio (PSNR), which measures the mean squared error between the reconstructed image and the ground truth reference image at the pixel level; a higher value (unit: dB) indicates less distortion in the reconstruction and higher pixel-level restoration accuracy. Second is the structural similarity index measure (SSIM), which evaluates image quality comprehensively from three dimensions: luminance, contrast, and structure. Its value ranges from $[0, 1]$, with values closer to 1 indicating that the reconstructed image is more consistent with the ground truth in terms of structural texture, aligning better with the human perception of structural information. Finally, learned perceptual image patch similarity (LPIPS) utilizes a pre-trained VGG deep neural network to extract features and calculate the distance between the reconstructed image and the ground truth in the

feature space. A lower LPIPS value indicates that the two are closer in visual perception; compared with traditional metrics, it better reflects the recovery of high-frequency textures and details.

Targeting the resource-constrained nature of satellite computing platforms, we also focused on resource overhead and efficiency. This primarily involves two key metrics relevant to deployment feasibility: First is peak VRAM, which tracks the maximum video memory usage during both training and inference, directly reflecting the algorithm's hardware memory requirements. The other one is execution time; since missions involving non-cooperative targets are typically time-critical, the algorithm's computational latency is of vital importance.

4.1.4. Implementation Details

The proposed method is implemented based on the PyTorch2.6 framework, with the core rendering pipeline accelerated using CUDA. The experimental platform is configured with $12 \times$ Intel(R) Xeon(R) Platinum 8260C CPU processors and an NVIDIA GeForce RTX 3090 GPU. Specific training parameters are set as follows:

Optimization Strategy: The total number of iterations is set to 20,000, using the Adam optimizer throughout. The initial learning rates for each learnable parameter are set as follows: position parameters at 1.6×10^{-4} ; scaling component of 5×10^{-3} and rotation component of 1×10^{-3} for covariance parameters; and an opacity parameter of 0.05. For physical material attributes, we set the learning rates for roughness, metallic, reflectance, and normal to 5×10^{-3} , 0.01, 5×10^{-3} , and 6×10^{-3} respectively.

Initialization and SH Settings: To fully fit the space lighting environment, the order of Spherical Harmonics is set to 3. The Gaussian point cloud initialization utilizes the sparse point cloud generated by the COLMAP SfM algorithm for a warm start, providing a good geometric prior for the optimization process.

Sparsification and Density Control: To ensure the compactness of the scene representation, we set the opacity threshold to 0.05 to prune redundant Gaussians with low contributions and execute gradient-based density control every 100 iterations. No additional model quantization compression strategy is applied in this experimental stage.

4.2. Results and Analysis

4.2.1. Quantitative Evaluation

As illustrated in Figure 6, our method achieves an optimal balance between rendering quality and memory consumption, significantly outperforming state-of-the-art NeRF and 3DGS techniques in both geometric accuracy and view synthesis quality under complex lighting conditions.

Table 1 presents the quantitative evaluation results of each method on the test set. The data indicates that in test scenes containing strong specular reflections and HDR lighting, traditional implicit methods including NeRF and Instant-NGP often struggle to accurately model drastic view-dependent lighting changes, leading to geometric divergence or foggy artifacts. Although standard 3DGS has an advantage in inference speed, it still suffers from specular artifacts and detail loss when dealing with non-Lambertian surfaces. While GaussianShader improves upon standard 3DGS by incorporating PBR, it still falls short in extreme space lighting conditions due to its reliance on general environment maps rather than explicit directional sunlight modeling. In contrast, the proposed method, benefiting from explicit material parameter modeling and physical rendering constraints, achieved significant improvements in both PSNR and SSIM metrics, proving its superior image synthesis capability under complex lighting. Meanwhile, the advantage in geometric accuracy metrics also verifies the positive role of physical constraints in geometric structure recovery.




Figure 6. Performance comparison of different methods in terms of rendering quality (PSNR) versus memory consumption.

Table 1. Quantitative results and mean values for various datasets. The intensity of the red color corresponds to performance, with deeper shades indicating superior results. An upward arrow indicates that a higher value is better, while a downward arrow indicates that a lower value is better.

Algorithm	PSNR ↑	SSIM ↑	LPIPS ↓	Memory ↓	Time ↓
<i>Average results (avg)</i>					
NeRF	13.92	0.910	0.1900	7051 MiB	25 h
Instant-ngp	14.75	0.870	0.2200	11,301 MiB	19 min
3DGS	25.45	0.947	0.0766	4760 MiB	17 min
GaussianShader	29.67	0.967	0.0457	6541 MiB	33 min
Space-Gaussian	36.10	0.980	0.0170	4989 MiB	22 min
<i>Results for cubesat</i>					
NeRF	11.71	0.875	0.3137	8367 MiB	26 h
Instant-ngp	11.43	0.751	0.4190	11,315 MiB	20 min
3DGS	34.23	0.978	0.0466	2835 MiB	18 min
GaussianShader	33.89	0.981	0.0460	6389 MiB	39 min
Space-Gaussian	36.63	0.985	0.0257	5387 MiB	22.85 min
<i>Results for starlink</i>					
NeRF	15.50	0.916	0.1566	6388 MiB	26 h
Instant-ngp	17.10	0.914	0.1195	11,851 MiB	18 min
3DGS	17.44	0.945	0.1051	6319 MiB	19 min
GaussianShader	26.38	0.967	0.0596	6766 MiB	29 min
Space-Gaussian	37.33	0.987	0.0134	5931 MiB	30.23 min
<i>Results for Sentinel-6</i>					
NeRF	14.56	0.948	0.1005	6400 MiB	24 h
Instant-ngp	15.72	0.937	0.1160	10,739 MiB	18 min
3DGS	24.67	0.918	0.0510	5127 MiB	14 min
GaussianShader	28.75	0.952	0.0315	6467 MiB	31 min
Space-Gaussian	34.36	0.982	0.0129	4329 MiB	21.23 min

4.2.2. Qualitative Comparison

Figure 7 shows the comparison of the qualitative rendering results of different methods under typical test views. We can deeply analyze the pros and cons of each method from two dimensions—geometric structure recovery and appearance detail performance:




Figure 7. Comparison of different algorithms under complex illumination. NeRF-based methods exhibit foggy blur at boundaries, while standard 3DGS produces geometric noise in specular regions. In contrast, our method recovers the cleanest geometry and texture.

Implicit representation methods such as NeRF and Instant-NGP exhibit pronounced volumetric haze when processing the high-frequency reflections typical of space targets. As illustrated in the first and second columns of the figure, this effect is particularly noticeable at the edges of solar panels. This occurs because volume-rendering techniques attempt to reconcile sharp view-dependent color transitions by populating the spatial radiance field with semi transparent density. Consequently, this process blurs the structural boundaries of the object and introduces significant floating noise within the deep space background.

The standard 3DGS method achieves sharp boundaries using explicit Gaussian ellipsoids but introduces significant geometric artifacts in non-Lambertian regions such as golden MLI films. This failure can be mathematically attributed to the frequency mismatch between the SH used for color representation and the high-frequency nature of specular reflections. In standard 3DGS, the view-dependent color $\mathbf{c}(\mathbf{v})$ is parameterized by SH basis functions $Y_l^m(\mathbf{v})$ up to degree l_{max} :

$$\mathbf{c}(\mathbf{v}) = \sum_{l=0}^{l_{max}} \sum_{m=-l}^l c_l^m Y_l^m(\mathbf{v}) \quad (25)$$

where c_l^m are the learnable coefficients. SH functions are inherently low-pass filters. Even with $l_{max} = 3$, they can only capture smooth, low-frequency lighting variations. However, the specular reflection of space materials follows a highly concentrated microfacet distribution, such as the GGX distribution $D(\mathbf{n}, \mathbf{h})$ used in our method, which exhibits a sharp, high-frequency peak when the half-vector $\mathbf{h}$ aligns with the normal $\mathbf{n}$.

Because the low-order SH cannot mathematically fit this high-frequency Dirac-like specular peak, the photometric loss $\mathcal{L}_c$ remains high in highlight regions. To minimize this

loss, the optimization process is forced to compensate for the color error by altering the geometric properties of the Gaussians. Specifically, it distorts the covariance matrix Σ to create numerous elongated, acicular primitives that are only visible from specific narrow viewpoints, thereby artificially mimicking specular highlights. This geometric overfitting manifests as an unnaturally jagged surface topography that lacks physical coherence, causing severe highlight flickering and a loss of visual realism during viewpoint transitions. GaussianShader mitigates some of these artifacts by introducing material properties, but its simplified lighting assumptions still lead to suboptimal highlight shapes and residual noise on highly reflective satellite components. It is worth noting that the geometric artifacts observed in standard 3DGS primarily stem from the inherent mathematical limitations of using low-order spherical harmonics to represent high-frequency signals. Consequently, even if the baseline models were subjected to the same level of optimization scrutiny as our method, the fundamental frequency mismatch would likely persist, making it challenging to completely eliminate erroneous geometries in specular regions.

In contrast, Space-Gaussian proposed in this paper successfully decouples color changes into functions of material attributes and lighting directions, benefiting from the introduction of the BRDF material model. As can be seen from the figure, our method not only reconstructs the smooth and continuous satellite body surface but also effectively suppresses artifacts on solar panels. The results indicate that the reconstruction accuracy of standard 3DGS decreases in direct proportion to the size of a satellite's solar arrays. This pattern is especially visible in configurations like Starlink as shown in Figure 8, where expansive solar panels dominate the structure. When evaluating satellites with smaller panel footprints, such as the Cubesat, the improvement offered by our approach relative to the baseline becomes more modest. Despite these variations in baseline performance, Space-Gaussian exhibits unwavering robustness across all tested satellite configurations.




Figure 8. A Starlink whose structure is mostly composed of solar panels.

4.3. Ablation Study and Analysis

To verify the contribution of each core module proposed in this paper to the final performance, we design a comprehensive set of ablation studies. These studies are both conducted quantitatively, with particular emphasis on analyzing the physical-level effects caused by the removal or alteration of individual components, thereby demonstrating the scientific validity of our architectural choices. In the BRDF setting, the physically based BRDF rendering and material modeling modules are disabled, and the method degenerates into the standard 3DGS framework that directly fits view-dependent colors using SH. This variant is mainly intended to evaluate the importance of PBR for recovering specular highlights and fine-grained material details.

In the material optimization variant, the metallic and roughness parameters of all Gaussian primitives in the scene are fixed to global constants, while only the diffuse albedo is optimized during training. This experiment aims to assess the contribution of spatially

varying material properties to accurately modeling heterogeneous materials commonly observed on complex satellite surfaces.

For the solar orientation initialization configuration, the prior sun direction provided by ephemeris data is not used. Instead, the scene illumination direction is randomly initialized and jointly optimized with other parameters. This ablation is designed to examine the convergence behavior and robustness of the proposed method under unknown or inaccurate lighting conditions.

Finally, in the deferred shading setting, the G-buffer generation and deferred shading pipeline are removed, and color computation is performed directly during the rasterization stage. This comparison validates the advantage of the deferred rendering pipeline in decoupling geometry from illumination and enhancing the reconstruction of high-frequency details.

To quantitatively evaluate these variants, Table 2 reports the results of all ablation configurations on the test set. As can be observed, the complete *Space-Gaussian* model consistently achieves the best performance across all metrics, demonstrating the necessity and effectiveness of each proposed module. Specifically, the removal of the BRDF module leads to a significant drop in PSNR (from 36.10 dB to 32.12 dB), quantitatively confirming its critical role in handling complex specular reflections.

Table 2. Ablation studies on the components of Space-Gaussian. An upward arrow indicates that a higher value is better, while a downward arrow indicates that a lower value is better.

Method	PSNR ↑	SSIM ↑	LPIPS ↓
<i>w/o</i> BRDF	32.12	0.952	0.645
<i>w/o</i> Material Opt.	34.89	0.976	0.172
<i>w/o</i> Solar Orientation Init.	29.45	0.967	0.298
<i>w/o</i> Deferred Shading	34.10	0.960	0.023
Full Model (Ours)	36.10	0.980	0.017

5. Discussion and Limitations

5.1. The Scarcity of Non-Cooperative Space Target Datasets

While 3D reconstruction datasets for terrestrial scenes such as urban buildings and indoor environments have become increasingly comprehensive through benchmarks like ScanNet and DTU, high-fidelity datasets for non-cooperative space targets remain remarkably scarce. This shortage significantly restricts the advancement of deep-learning algorithms designed for on-orbit perception. The primary obstacle to acquiring real-world data is the high technical and financial threshold associated with space missions.

Beyond the difficulties of data acquisition, existing simulation datasets like SPEED+ [26], SHIRT [27], SPADES [28], SPARK, and URSO [29] often fail to provide sufficient physical realism. Most current synthetic data focus on pose estimation and employ simplified lighting models or low-dynamic-range environment maps that cannot accurately reproduce the extreme contrast of the space environment. More importantly, the surface materials of space targets—including multi-layer insulation, solar arrays, and composite antennas—exhibit complex reflectance properties that are frequently overlooked. Traditional simulations generally ignore strong specular reflections and anisotropic effects, resulting in a substantial domain gap when models trained on synthetic data are applied to real scenarios.

This study addresses these limitations by introducing a simulation dataset specifically engineered to enhance physical accuracy. We have implemented the precise modeling of direct sunlight and Earth albedo while establishing a material library based on actual physical measurements to capture the intricate scattering effects of satellite components. Although this dataset represents a significant step forward in replicating the physical attributes of light and matter, the diversity of target types and the complexity of their

motion states remain areas for future expansion. By refining these simulation techniques, we aim to provide a more robust foundation for the development of autonomous on-orbit perception systems.

5.2. Dilemma of SfM-Based Pose Estimation and Feature Matching

In the reconstruction workflow of implicit neural fields and 3D Gaussian splatting, precise camera pose estimation and sparse point cloud initialization are fundamental to achieving high quality results. While traditional structure from motion algorithms like COLMAP are standard, they encounter significant obstacles when applied to non-cooperative space targets. This failure stems primarily from the extreme optical properties of spacecraft surfaces, particularly the strong specular reflections produced by multi-layer insulation films. These metallic surfaces act as mirrors, causing detected feature points to shift across the object as the viewpoint changes rather than remaining anchored to physical locations. Such violations of multiview consistency prevent bundle adjustment from converging, which introduces substantial pose errors and produces stratified point clouds.

Furthermore, the lack of distinct textures and the prevalence of periodic patterns on the satellite structures exacerbate matching ambiguities. Many spacecraft bodies consist of smooth metal skins that offer minimal corner or edge features, resulting in extremely sparse feature sets. Conversely, the lattice patterns of large solar panels often trigger perceptual aliasing, where the matching algorithm incorrectly identifies correspondences between identical looking patches. In practice, these repetitive structures frequently lead to camera drift or total scene breakage, especially in satellite models characterized by expansive solar arrays. These limitations highlight the inherent fragility of feature-based matching in orbital environments.

Due to the inherent difficulty of matching feature points on certain aircraft surfaces, conventional SfM pipelines often fail [30]. Consequently, exploring COLMAP-free methodologies represents a critical direction for future research.

5.3. Model Limitations and Efficacy Analysis

Space-Gaussian significantly enhances reconstruction quality by effectively decoupling geometry from appearance through a microfacet-based PBR pipeline. Traditional frameworks often produce semi transparent artifacts or spatial noise when attempting to fit specular highlights. In contrast, this approach interprets highlights as a function of material properties and lighting directions, ensuring that the reconstructed surfaces remain dense and geometrically accurate. Consequently, the model successfully captures sharp highlight transitions and realistic shadow dynamics on complex reflective surfaces such as solar panels and thermal insulation.

Despite these advancements, the framework faces limitations regarding lighting complexity and viewpoint generalization. The current reliance on simplified environment maps and spherical harmonics fails to account for intricate global illumination effects, such as inter reflections between satellite components. This deficiency can introduce inaccuracies in albedo estimation and shadow rendering within occluded regions. Moreover, the model remains sensitive to extreme viewpoint deviations. In the absence of comprehensive geometric priors, observing unrecorded surfaces during constrained orbital maneuvers can lead to significant artifacts or structural gaps.

Finally, integrating physically based rendering necessitates a compromise between visual fidelity and computational efficiency. The inclusion of BRDF components and normal estimations increases the processing burden, resulting in a performance reduction of approximately thirty percent compared with standard 3DGS. Such an overhead poses challenges for deployment on hardware with limited onboard processing power. Future development will therefore focus on streamlining the calculation flow to achieve a more

favorable balance between rendering precision and real-time operational requirements in space environments.

6. Conclusion and Future Work

Addressing the challenges of 3D reconstruction for on-orbit non-cooperative targets caused by HDR lighting and strong specular reflections, this paper proposes a physics-aware Gaussian reconstruction framework named Space-Gaussian. This method innovatively integrates a PBR pipeline into the 3DGS architecture. By introducing a microfacet BRDF model and deferred rendering technology, it achieves effective decoupling of geometric structure and surface material attributes. Extensive experiments on a high-fidelity simulation dataset demonstrate that Space-Gaussian significantly outperforms mainstream methods such as NeRF, Instant-NGP, GaussianShader, and standard 3DGS. It not only effectively eliminates geometric artifacts and highlights leakage in strong specular reflection areas but also achieves the best balance among novel view synthesis quality, geometric reconstruction accuracy, and real-time rendering, exhibiting excellent robustness and engineering application potential.

Looking to the future, this study plans to deepen and expand the following dimensions: first, enhancing complex lighting interactions, by introducing ambient occlusion and higher-order lighting estimation models to further improve reconstruction stability under extreme shadow variations and self-occlusion conditions; second, conducting experiments using real-world satellite imagery in the future, although the acquisition of such datasets remains a significant challenge; and third, improving the algorithm to be COLMAP-free, enabling it to adapt to a wider range of scenarios.

Author Contributions: Conceptualization, Z.Q., Z.Z. and L.C.; methodology, Z.Q.; software, Z.Q.; validation, Z.Q., Z.Z., R.X., J.Z. and L.C.; formal analysis, Z.Q.; investigation, R.X.; resources, R.X.; data curation, J.Z.; writing—original draft preparation, Z.Q.; writing—review and editing, Z.Z.; visualization, J.Z.; supervision, L.C.; project administration, L.C.; funding acquisition, L.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding

Data Availability Statement: Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

OOS	On-Orbit Servicing
ADR	Active Debris Removal
HDR	High Dynamic Range
MLI	Multi-layer Insulation
NeRF	Neural Radiance Fields
3DGS	3D Gaussian Splatting
BRDF	Bidirectional Reflectance Distribution Function
PBR	Physically Based Rendering
MVS	Multiview Stereo
LEO	Low Earth Orbit
SfM	Structure from motion
RPO	Rendezvous and Proximity Operations
MLP	Multi-layer Perceptron

References

- Gu, Y.; Li, Z.; Liu, H.; Luo, Q.; Liu, H.; Wu, G. Multi-objective early warning mission planning by multiple satellites using a critical task aggregation-based NSGA-II algorithm. *Adv. Space Res.* 2025, *in press*. [\[CrossRef\]](#)
- Flores-Abad, A.; Ma, O.; Pham, K.; Ulrich, S. A review of space robotics technologies for on-orbit servicing. *Prog. Aerosp. Sci.* **2014**, *68*, 1–26. [\[CrossRef\]](#)
- Park, T.H.; D’Amico, S. Robust multi-task learning and online refinement for spacecraft pose estimation across domain gap. *Adv. Space Res.* **2024**, *73*, 5726–5740. [\[CrossRef\]](#)
- Rangelov, D.; Waanders, S.; Waanders, K.; van Keulen, M.; Miltchev, R. Three-Dimensional Reconstruction Techniques and the Impact of Lighting Conditions on Reconstruction Quality: A Comprehensive Review. *Lights* **2025**, *1*, 1. [\[CrossRef\]](#)
- Mildenhall, B.; Srinivasan, P.P.; Tancik, M.; Barron, J.T.; Ramamoorthi, R.; Ng, R. NeRF: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* **2021**, *65*, 99–106. [\[CrossRef\]](#)
- Kerbl, B.; Kopanas, G.; Leimkuehler, T.; Drettakis, G. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.* **2023**, *42*, 1–14. [\[CrossRef\]](#)
- Cook, R.L.; Torrance, K.E. A reflectance model for computer graphics. *SIGGRAPH Comput. Graph.* **1981**, *15*, 307–316. [\[CrossRef\]](#)
- Burley, B.; Studios, W.D.A. Physically-based shading at disney. In *Proceedings of the ACM Siggraph*; Association for Computing Machinery: New York, NY, USA, 2012; Volume 2012, pp. 1–7.
- Saito, T.; Takahashi, T. Comprehensible rendering of 3-D shapes. In *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques*; SIGGRAPH ’90; Association for Computing Machinery: New York, NY, USA, 1990; pp. 197–206. [\[CrossRef\]](#)
- Roberts, L.G. Machine Perception of Three-Dimensional Solids. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, 1963.
- Longuet-Higgins, H.C. A computer algorithm for reconstructing a scene from two projections. *Nature* **1981**, *293*, 133–135. [\[CrossRef\]](#)
- Hartley, R.; Zisserman, A. *Multiple View Geometry in Computer Vision*; Cambridge University Press: Cambridge, UK, 2003.
- Levoy, M.; Pulli, K.; Curless, B.; Rusinkiewicz, S.; Koller, D.; Pereira, L.; Ginzton, M.; Anderson, S.; Davis, J.; Ginsberg, J.; et al. The digital Michelangelo project: 3D scanning of large statues. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*; SIGGRAPH ’00; Association for Computing Machinery: New York, NY, USA, 2000; pp. 131–144. [\[CrossRef\]](#)
- Friend, R.B. Orbital Express program summary and mission overview. In *Proceedings of the Sensors and Systems for Space Applications II*; SPIE: Bellingham, WA, USA, 2008; Volume 6958, pp. 11–21. [\[CrossRef\]](#)
- Forshaw, J.L.; Aglietti, G.S.; Navarathinam, N.; Kadhemi, H.; Salmon, T.; Pisseloup, A.; Joffe, E.; Chabot, T.; Retat, I.; Axthelm, R.; et al. RemoveDEBRIS: An in-orbit active debris removal demonstration mission. *Acta Astronaut.* **2016**, *127*, 448–463. [\[CrossRef\]](#)
- Biesbroek, R.; Aziz, S.; Wolahan, A.; Cipolla, S.; Richard-noca, M.; Piguat, L. The ClearSpace-1 mission: ESA and ClearSpace team up to remove debris. In *Proceedings of the 8th European Conference on Space Debris*; ESA Space Debris Office: Darmstadt, Germany, 2021.
- Choy, C.B.; Xu, D.; Gwak, J.; Chen, K.; Savarese, S. 3D-R2N2: A Unified Approach for Single and Multi-view 3D Object Reconstruction. In *Proceedings of the Computer Vision—ECCV 2016, Honolulu, HI, USA, 11–14 October*; Leibe, B., Matas, J., Sebe, N., Welling, M., Eds.; Springer: Cham, Switzerland, 2016; pp. 628–644. [\[CrossRef\]](#)
- Charles, R.Q.; Su, H.; Kaichun, M.; Guibas, L.J. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*; IEEE: Piscataway, NJ, USA, 2017; pp. 77–85. [\[CrossRef\]](#)
- Artzy, E.; Frieder, G.; Herman, G.T. The theory, design, implementation and evaluation of a three-dimensional surface detection algorithm. *Comput. Graph. Image Process.* **1981**, *15*, 1–24. [\[CrossRef\]](#)
- Barron, J.T.; Mildenhall, B.; Verbin, D.; Srinivasan, P.P.; Hedman, P. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 2022*; IEEE: Piscataway, NJ, USA, 2022; pp. 5460–5469. [\[CrossRef\]](#)
- Verbin, D.; Hedman, P.; Mildenhall, B.; Zickler, T.; Barron, J.T.; Srinivasan, P.P. Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields. *IEEE Trans. Pattern Anal. Mach. Intell.* **2025**, *47*, 9426–9437. [\[CrossRef\]](#) [\[PubMed\]](#)
- Müller, T.; Evans, A.; Schied, C.; Keller, A. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph. (TOG)* **2022**, *41*, 102. [\[CrossRef\]](#)
- Forray, C.; Delporte, P.; Delaygue, N.; Genin, F.; Derksen, D. Joint attitude estimation and 3D neural reconstruction of non-cooperative space objects. *arXiv* **2025**, arXiv:2506.20638. [\[CrossRef\]](#)
- Jiang, Y.; Tu, J.; Liu, Y.; Gao, X.; Long, X.; Wang, W.; Ma, Y. GaussianShader: 3D Gaussian Splatting with Shading Functions for Reflective Surfaces. In *Proceedings of the 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*; IEEE: Piscataway, NJ, USA, 2024; pp. 5322–5332. [\[CrossRef\]](#)
- Rother, C.; Kolmogorov, V.; Blake, A. “GrabCut”: Interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.* **2004**, *23*, 309–314. [\[CrossRef\]](#)

26. Park, T.H.; Märtens, M.; Lecuyer, G.; Izzo, D.; D’Amico, S. SPEED+: Next-Generation Dataset for Spacecraft Pose Estimation across Domain Gap. In *Proceedings of the 2022 IEEE Aerospace Conference (AERO)*; IEEE: Piscataway, NJ, USA, 2022; pp. 1–15. [CrossRef]
27. Park, T.H.; D’Amico, S. SHIRT: Satellite Hardware-In-the-Loop Rendezvous Trajectories Dataset. Stanford Digital Repository. 2022. Available online: <https://purl.stanford.edu/zq716br5462> (accessed on 31 January 2026).
28. Rathinam, A.; Qadadri, H.; Aouada, D. SPADES: A Realistic Spacecraft Pose Estimation Dataset using Event Sensing. In *Proceedings of the 2024 IEEE International Conference on Robotics and Automation (ICRA)*; IEEE: Piscataway, NJ, USA, 2024; pp. 11760–11766. [CrossRef]
29. Proença, P.F.; Gao, Y. Deep Learning for Spacecraft Pose Estimation from Photorealistic Rendering. In *Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA)*; IEEE: Piscataway, NJ, USA, 2020; pp. 6007–6013. [CrossRef]
30. Issitt, A.; Mahendrakar, T.; Alvarez, A.; White, R.T.; Sizemore, A. On Optimal Observation Orbits for Learning Gaussian Splatting-Based 3D Models of Unknown Resident Space Objects. In *Proceedings of the AIAA SCITECH 2025 Forum*; American Institute for Aeronautics and Astronautics (AIAA): Reston, VA, USA, 2025; p. 1780. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.