

MagnetGS-Mesh: High-Quality Multi-Object Mesh Reconstruction via Adaptive Surface Optimization

Min-Su Park¹, Yeonho Han¹, Uijoon Jeong¹,
Jun-Hyeong Park¹, and Eun-Seok Ryu^{1*}

Sungkyunkwan University, Republic of Korea
{hskiok, dusgh9259, uiuh4569, joke0702, esryu}@skku.edu

Abstract. We present **MagnetGS-Mesh**, an integrated framework that directly recovers high-quality, object-wise surface meshes from 3D Gaussian Splatting (3DGS) while preserving photorealistic novel-view synthesis. Prior 3DGS-to-mesh methods generate meshes by considering all Gaussians without distinguishing objects from the background, which inevitably yields excessive and unnecessary vertices in irrelevant regions. To address this efficiency bottleneck, we adopt an object-wise reconstruction approach. However, decomposing 3DGS representations into individual objects introduces distinct geometric challenges: numerous outlier Gaussians appear far from true surfaces, and the remaining surface Gaussians often exhibit highly non-uniform density. Since robust mesh extraction fundamentally requires a uniform distribution of Gaussians to prevent surface holes and artifacts, these irregularities severely degrade geometric fidelity. To resolve these issues, we propose a geometry-aligned two-stage training procedure. In Stage 1, we employ Enhanced Occupancy Learning to produce crisp boundaries, ensuring accurate 3D segmentation. In Stage 2, to address the uneven distribution, we introduce the **Adaptive Local Normal Consistency (ALNC) Loss**. This mechanism acts as a “magnet,” strategically relocating outlier Gaussians to sparse surface regions to ensure uniform coverage. By adaptively aligning these relocated Gaussians with local surface normals, **ALNC** effectively fills geometric gaps while maintaining structural coherence. Experimentally, MagnetGS-Mesh achieves competitive or superior mesh quality compared to state-of-the-art methods while simultaneously reducing model size by substantial margins. The code is available at <https://github.com/MinsuPark0752/MagnetGS-Mesh>.

Keywords: 3D Gaussian Splatting · Mesh Reconstruction · Object Segmentation · Surface Optimization · Normal Consistency

1 Introduction

3DGS [20] has revolutionized real-time novel view synthesis by representing scenes as collections of millions of optimized 3D Gaussians. Building upon this

* Corresponding author: Eun-Seok Ryu.



Fig. 1: MagnetGS-Mesh delivers high-quality meshes with low storage footprint through Stage 1 (Enhanced Occupancy Learning) and Stage 2 (ALNC-Guided Optimization). After 3D segmentation, outlier Gaussians appear far from true surfaces. ALNC Loss relocates these outliers to sparse surface regions by computing the average of local surface normals from neighboring surface points (adaptively selected based on local density: small k for dense regions, large k for sparse regions) and aligning the outlier-to-surface movement direction with this averaged normal, ensuring natural surface attachment and balanced coverage for high-quality mesh extraction.

foundation, recent advances have introduced complementary capabilities: MILO [14] incorporates auxiliary occupancy fields to bridge explicit point-based representations and implicit surface geometry for better geometric reconstruction, while 3D segmentation [3, 4, 47] methods enable real-time object-level scene decomposition.

Despite these advances, existing 3DGS-to-mesh methods [14, 15, 18, 50] process all Gaussians without distinguishing objects from background, yielding bloated meshes unsuitable for per-object editing or animation. Our object-wise approach segments scenes before mesh extraction, but this introduces geometric challenges unique to per-object decomposition that existing methods cannot handle: outlier Gaussians scattered far from true surfaces, and severe non-uniformity in the remaining surface Gaussians—irregularities that fundamentally prevent high-quality mesh extraction.

Our approach. We present **MagnetGS-Mesh**, whose name reflects its core mechanism: like a magnet attracting scattered particles into a coherent structure, our method magnetically draws outlier Gaussians toward sparse surface regions. The key technical contribution is the **Adaptive Local Normal Consistency (ALNC) Loss**, which uniquely resolves the non-uniform Gaussian distribution by relocating outlier Gaussians into sparse surface areas while respecting local surface geometry through density-adaptive normal alignment. Unlike prior

normal consistency approaches [28,29] that apply a fixed neighborhood size uniformly, ALNC adapts its neighborhood based on local density: small neighborhoods preserve sharp details in dense regions, while large neighborhoods provide noise-robust guidance in sparse regions. To ensure ALNC operates on a clean, well-separated surface representation, Stage 1 first applies **Enhanced Occupancy Learning**—five synergistic regularization losses that sharpen occupancy boundaries from the ambiguous 0.3–0.7 range into crisp binary values, enabling reliable surface-outlier separation as a prerequisite for ALNC. Figure 2 illustrates the complete pipeline.

Our framework achieves competitive or superior mesh quality compared to state-of-the-art methods while simultaneously reducing model size by up to 27%, with significant improvements in mesh completeness and detail preservation for complex geometries with curved surfaces and sharp features.

The contributions of our paper are as follows:

- **Adaptive Local Normal Consistency (ALNC) Loss:** Our core contribution—a novel density-adaptive surface optimization that relocates outlier Gaussians to sparse surface regions, resolving the non-uniform Gaussian distribution caused by object-wise segmentation and enabling topology-complete, artifact-free mesh extraction.
- **Enhanced Occupancy Learning:** Five synergistic regularization losses that sharpen occupancy boundaries into crisp binary values, providing the clean surface-outlier separation required for effective ALNC optimization.
- **MagnetGS-Mesh Pipeline:** An end-to-end framework for object-centric scene decomposition and high-quality per-object mesh extraction, delivering superior geometric fidelity at substantially lower storage costs.

2 Related Work

Novel View Synthesis Novel view synthesis aims to generate photorealistic images from arbitrary viewpoints given a set of input views. Neural Radiance Fields (NeRF [31]) pioneered this domain by representing scenes as continuous 5D functions optimized via volumetric rendering, achieving impressive photorealistic quality. However, NeRF [31] suffers from slow rendering speeds due to dense ray marching through implicit neural representations. Subsequent works have improved efficiency through spatial data structures [32,49], factorized representations [5,9], and explicit voxel grids [10,41], yet these implicit methods remain computationally expensive for real-time applications. To address this limitation, 3D Gaussian Splatting [20] (3DGS [20]) introduced explicit point-based representations, modeling scenes as collections of anisotropic 3D Gaussians rendered via differentiable splatting, achieving real-time rendering at over 100 FPS while maintaining photorealistic quality.

Mesh Extraction from 3DGS Converting 3DGS [20] representations to explicit surface meshes remains challenging due to their point-based nature. SuGaR [15] binds Gaussians to mesh surfaces via Poisson reconstruction [19] but requires

careful regularization to prevent degeneracies. 2DGS [18] constrains Gaussians to 2D oriented disks for direct mesh extraction but sacrifices representational flexibility. GOF [50] leverages Gaussian opacity fields to identify surfaces yet produces noisy meshes with inconsistent vertex distributions. MLO [14] introduces auxiliary occupancy fields trained jointly with Gaussians in a mesh-in-the-loop framework, achieving improved geometric fidelity. However, all these methods train entire scenes holistically, computing meshes from all Gaussians without distinguishing objects from background. This scene-level approach yields excessive vertices in irrelevant regions and produces outlier Gaussians far from true surfaces, degrading per-object mesh quality. Our approach builds upon MLO [14]’s occupancy learning framework and introduces the **Adaptive Local Normal Consistency (ALNC) Loss** to resolve the non-uniform Gaussian distribution unique to object-wise segmentation, enabling topology-complete mesh extraction with reduced storage requirements.

3D Segmentation in Radiance Fields Decomposing 3D scenes into meaningful objects has been explored extensively in both NeRF [31] and 3DGS [20] contexts. Early NeRF [31]-based methods [11, 24, 53] require per-view semantic labels or multi-view masks, limiting scalability. Recent works like SA3D [4] and SAM3D [47] leverage foundation models (SAM [22]) for interactive 3D segmentation [17, 43] but operate on implicit representations. For 3DGS [20], LangSplat [36] and LEGaussians [39] embed language features into Gaussians for open-vocabulary queries, while Gaussian Grouping [48] clusters Gaussians via identity encoding. SAGA [3] attaches scale-gated affinity features to each Gaussian, enabling efficient multi-granularity segmentation with minimal overhead. However, these segmentation methods focus on accurately partitioning scenes rather than addressing the downstream geometric artifacts (outliers, uneven density) that emerge when extracting per-object meshes.

Surface Regularization and Normal Consistency Enforcing geometric priors during optimization has proven effective for surface reconstruction. Normal consistency losses have been widely used in multi-view stereo [25, 27] and neural implicit surfaces [8, 34, 45] to encourage smoothness and coherence. In 3DGS [20], GeoGaussian [28] incorporates depth and normal cues from monocular estimators, while GOF [50] introduces opacity-based regularization for better surface alignment [7]. Our ALNC Loss differs fundamentally by adapting to local surface density: rather than applying uniform constraints, we dynamically adjust neighborhood sizes to handle both dense and sparse regions, guiding outlier Gaussians to naturally consolidate onto surfaces while preserving geometric coherence. This density-adaptive mechanism is crucial for object-wise scenarios where segmentation [3] introduces highly variable surface coverage.

3 Preliminary

3D Gaussian Splatting A 3DGS [20] scene is represented as a collection of N Gaussians $\{\mathcal{G}_i\}_{i=1}^N$, where each Gaussian is parameterized by position $\mu_i \in$




Fig. 2: Overview of the MagnetGS-Mesh pipeline. **Stage 1 (Enhanced Occupancy Learning):** We extend MILO [14] with five regularization losses to produce crisp occupancy boundaries suitable for segmentation. **Stage 2 (ALNC-Guided Optimization):** After segmenting objects via SAGA [3], we detect surface and outlier Gaussians, then optimize outliers toward sparse surface regions using Adaptive Local Normal Consistency Loss, which adaptively adjusts neighborhood sizes based on local density for natural surface attachment. The total loss combines photometric loss (rendering quality), proximity loss (surface attachment), and ALNC loss (normal consistency), backpropagates gradients, and updates Gaussian positions via optimizer. After optimization, per-object meshes are extracted from the occupancy field, and all object and background meshes are combined to produce the final complete scene mesh.

$\mathbb{R}^3$, covariance Σ_i , opacity $\alpha_i \in [0, 1]$, and spherical harmonic coefficients for view-dependent color. Rendering is performed via differentiable splatting [20], producing photorealistic novel views at real-time rates.

Mesh-In-the-Loop (MILO) MILO [14] introduces an auxiliary occupancy field $o : \mathbb{R}^3 \rightarrow [0, 1]$ trained jointly with Gaussians. Each Gaussian $\mathcal{G}_i$ is assigned an occupancy value $o_i = o(\mu_i)$, and regularization losses encourage consistency between rendering (opacity α_i) and geometry (occupancy o_i). This hybrid representation enables mesh extraction via marching tetrahedra [30] while preserving rendering quality. However, vanilla MILO [14] produces soft, ambiguous occupancy boundaries insufficient for clean object segmentation.

4 Proposed Method

Our method addresses two fundamental problems introduced by object-wise 3DGS segmentation: outlier Gaussians far from true surfaces, and non-uniform Gaussian density causing surface holes. The core of our solution is the **Adaptive Local Normal Consistency (ALNC) Loss** (Section 4.3), which relocates outliers to sparse surface regions via density-adaptive normal alignment. For ALNC to operate effectively, it requires a clean separation between surface

and outlier Gaussians. This prerequisite is fulfilled by **Enhanced Occupancy Learning** (Section 4.1), which sharpens ambiguous occupancy boundaries into reliable binary values. Figure 2 illustrates the complete two-stage pipeline.

4.1 Enhanced Occupancy Learning

Vanilla MILO [14] produces soft, ambiguous occupancy boundaries (values around 0.3–0.7) that hinder clean object segmentation. For per-object mesh extraction, we require *crisp* binary boundaries transitioning sharply from 0 to 1. We extend MILO [14]’s occupancy regularization with five geometric priors tailored for segmentation-friendly representations. The total Stage 1 loss combines MILO [14]’s original objectives with our five complementary terms:

$$\begin{aligned}\mathcal{L}_{\text{Stage1}} = & \mathcal{L}_{\text{MILO}} + \lambda_1 \mathcal{L}_{\text{entropy}} \\ & + \lambda_2 \mathcal{L}_{\text{smooth}} + \lambda_3 \mathcal{L}_{\text{consist}} \\ & + \lambda_4 \mathcal{L}_{\text{boundary}} + \lambda_5 \mathcal{L}_{\text{temporal}}.\end{aligned}\tag{1}$$

where $\mathcal{L}_{\text{entropy}}$ minimizes binary entropy [12] to push values toward 0 or 1 for clean segmentation. $\mathcal{L}_{\text{smooth}}$ penalizes occupancy [33] differences between k -nearest neighbors ($k=3$) for spatial coherence. $\mathcal{L}_{\text{consist}}$ aligns opacity and occupancy [45] to maintain appearance-geometry coherence. $\mathcal{L}_{\text{boundary}}$ maximizes occupancy gradients near boundaries [13] ($o_i \approx 0.5$) while minimizing them elsewhere, producing sharp transitions for accurate surface detection. $\mathcal{L}_{\text{temporal}}$ stabilizes optimization via exponential moving average (EMA [42]). After Stage 1 training, approximately 80–90% of Gaussians (varies by scene complexity) have occupancy values in definite ranges ($o < 0.1$ or $o > 0.9$), while the remaining 10–20% lie in the boundary region. Detailed formulations are in the supplementary material.

4.2 Object Segmentation and Surface Detection

We perform multi-object segmentation using SAGA [3]. However, segmentation [3] artifacts manifest as scattered outlier Gaussians far from true surfaces. We therefore perform explicit surface detection [26, 44] as a prerequisite for subsequent outlier relocation (Section 4.3), enabling targeted optimization that produces a uniform Gaussian distribution essential for high-quality mesh extraction.

For each segmented object [3], we detect surface Gaussians from the boundary region (Gaussians with occupancy values not pushed to 0 or 1 by Stage 1) via two complementary criteria. **Occupancy-based scoring** assigns higher scores to Gaussians near the isosurface ($o_i \approx 0.5$):

$$s_i^{\text{occ}} = \exp\left(-\frac{(o_i - 0.5)^2}{2\sigma_{\text{occ}}^2}\right).\tag{2}$$




Fig. 3: ALNC Loss Computation Pipeline. (a) Detect nearest surface point (yellow) for each outlier (red). (b) Estimate local density (purple=dense, beige=sparse). (c) Select adaptive k_i neighbors (small k for dense, large k for sparse). (d) Average k_i surface normals (black arrows). (e) Measure alignment between movement direction (blue) and averaged normal (black). (f–g) Optimize with combined losses, guiding outliers along coherent directions. (h) Final surface-attached Gaussians with balanced coverage.

Photometric consistency compares rendered masks with ground-truth masks (from SAGA [3] segments or SAM2 [37]):

$$s_i^{\text{photo}} = \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \mathbb{I}[\text{IoU}(R_v, M_v) > \tau], \quad (3)$$

where $\mathcal{V}$ is the set of views, R_v is the rendered mask, M_v is the ground-truth mask, and τ is a threshold. The final surface score combines both:

$$s_i = w_{\text{occ}} \cdot s_i^{\text{occ}} + w_{\text{photo}} \cdot s_i^{\text{photo}}. \quad (4)$$

We classify Gaussians as *surface* if $s_i > \theta$ and *outlier* otherwise.

4.3 Adaptive Local Normal Consistency Loss

After surface detection, outlier Gaussians often lie far from true surfaces or cluster in already-dense regions, degrading mesh quality. We introduce Adaptive Local Normal Consistency (ALNC) Loss, our core technical contribution, to relocate outliers toward sparse surface regions while maintaining geometric coherence. As illustrated in Figure 3, our ALNC computation begins by identifying the nearest surface point for each outlier Gaussian, then estimates local surface density to distinguish dense and sparse regions. Based on this density, we adaptively select the neighborhood size k_i —using smaller k for dense areas to preserve high-frequency details and larger k for sparse areas to ensure robust geometric guidance. We compute the average surface normal from these k_i neighbors and measure the alignment between the outlier’s relocation direction and this averaged normal. The combined optimization objectives utilize these cues to guide outliers along geometrically coherent paths, resulting in a surface-attached Gaussian representation with balanced and uniform coverage. We optimize outliers in two phases:

Phase 1: Proximity-Driven Attachment. In the first 30% of iterations, we aggressively pull outliers toward surfaces using proximity loss:

$$\mathcal{L}_{\text{prox}} = \frac{1}{N_{\text{out}}} \sum_{i \in \mathcal{O}} \min_{j \in \mathcal{S}} \|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j\|^2, \quad (5)$$

where $\mathcal{O}$ and $\mathcal{S}$ denote outlier and surface Gaussians, respectively. This phase forcefully moves outliers closer to surfaces.

Phase 2: ALNC-Guided Fine-Tuning. In the remaining 70% of iterations, we combine three losses for natural, geometry-aware attachment:

Photometric Loss To preserve rendering quality, we minimize pixel-wise difference between rendered images and ground-truth segmentation masks:

$$\mathcal{L}_{\text{photo}} = \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \|R_v(\{\boldsymbol{\mu}_i\}) - M_v\|_1, \quad (6)$$

where R_v is the rendering function.

Proximity Loss We continue applying $\mathcal{L}_{\text{prox}}$ but with reduced weight to maintain surface proximity.

Adaptive Local Normal Consistency (ALNC) Loss This is our key contribution. Unlike uniform normal consistency constraints used in prior work [28, 29], ALNC dynamically adapts neighborhood sizes based on local surface density, enabling robust guidance in both dense and sparse regions. For each outlier i , we compute:

Step 1: Local Density Estimation Let $d_i = \frac{1}{k} \sum_{j=1}^k \|\boldsymbol{\mu}_i - \boldsymbol{\mu}_j^{\text{surf}}\|$ be the mean distance to k -nearest surface points. Local density is:

$$\rho_i = \frac{1}{d_i + \epsilon}. \quad (7)$$

ρ_i is used *solely* to determine the adaptive neighborhood size k_i —not to guide the relocation direction itself, which is governed by surface normals and proximity loss. Consequently, the method remains robust even when ρ_i is imprecisely estimated, since any miscalibration only affects *how many* neighbors are used for normal averaging, not *where* the outlier is moved.

Step 2: Adaptive Neighborhood Size We adjust the neighborhood size k_i based on density:

$$k_i = k_{\text{max}} - (k_{\text{max}} - k_{\text{min}}) \cdot \frac{\rho_i - \rho_{\text{min}}}{\rho_{\text{max}} - \rho_{\text{min}}}, \quad (8)$$

where $k_{\text{min}} = 3$ is the theoretical minimum for stable plane fitting [16], and $k_{\text{max}} = 15$ follows the empirically validated upper bound from PCPNet [16], which demonstrated that $k \in [10, 20]$ yields the optimal trade-off between noise robustness and geometric sensitivity for PCA-based local geometry estimation.

Dense regions use small k_i for precision; sparse regions use large k_i for robustness. This adaptive mechanism is crucial for handling uneven density distributions in segmented objects.

Step 3: Surface Normal Estimation For all surface Gaussians, we pre-compute normals via PCA on local neighborhoods (performed once before optimization). For each surface point j , we find its 10 nearest neighbors and compute the covariance matrix $\mathbf{C}_j$. The normal $\mathbf{n}_j$ is the eigenvector corresponding to the smallest eigenvalue of $\mathbf{C}_j$. Following PCPNet [16], $k = 10$ represents the sweet spot between noise robustness and geometric sensitivity for PCA-based normal estimation. Although each Gaussian’s covariance matrix encodes a local orientation via its shortest axis [15, 18], we use PCA-based normals computed from neighboring *surface* Gaussians $\mathcal{S}$. This provides stable estimates regardless of outlier position or shape: even with disk enforcement, Gaussians floating outside the surface still produce incorrect normals, whereas PCA over surface neighbors remains accurate.

Step 4: Averaged K-Nearest Normals For each outlier i , we select the k_i -nearest surface points and average their normals. Before averaging, each neighbor normal $\mathbf{n}_j^{\text{surf}}$ is oriented consistently by flipping its sign if $\mathbf{n}_j^{\text{surf}} \cdot \mathbf{n}_1^{\text{surf}} < 0$, ensuring all normals lie in the same hemisphere and preventing destructive cancellation:

$$\bar{\mathbf{n}}_i = \frac{1}{k_i} \sum_{j=1}^{k_i} s_j \mathbf{n}_j^{\text{surf}}, \quad s_j = \text{sgn}(\mathbf{n}_j^{\text{surf}} \cdot \mathbf{n}_1^{\text{surf}}). \quad (9)$$

By averaging over adaptive neighborhoods, we obtain robust normal estimates even in sparse or noisy regions, unlike fixed- k approaches [16]. Note that the absolute value in $a_i = |\mathbf{d}_i \cdot \bar{\mathbf{n}}_i|$ (Step 5) makes the ALNC loss invariant to any remaining global sign ambiguity.

Step 5: Direction-Normal Consistency Let $\mathbf{d}_i = \frac{\boldsymbol{\mu}_i^{\text{nearest}} - \boldsymbol{\mu}_i}{\|\boldsymbol{\mu}_i^{\text{nearest}} - \boldsymbol{\mu}_i\|}$ be the normalized direction from outlier i to its nearest surface point. We measure alignment with the averaged normal:

$$a_i = |\mathbf{d}_i \cdot \bar{\mathbf{n}}_i|. \quad (10)$$

The ALNC loss encourages alignment:

$$\mathcal{L}_{\text{ALNC}} = \frac{1}{N_{\text{out}}} \sum_{i \in \mathcal{O}} (1 - a_i). \quad (11)$$

The total Phase 2 loss is:

$$\mathcal{L}_{\text{Phase2}} = \mathcal{L}_{\text{photo}} + \lambda_{\text{prox}} \mathcal{L}_{\text{prox}} + \lambda_{\text{ALNC}} \mathcal{L}_{\text{ALNC}}. \quad (12)$$

Through gradient descent, outliers migrate toward sparse surface regions along geometrically coherent directions, consolidating into complete, artifact-free surfaces. The magnetic consolidation effect—where scattered Gaussians are drawn to underrepresented areas while respecting local geometry—naturally emerges from the interplay of proximity, photometric, and ALNC losses.

4.4 Mesh Extraction

After Stage 2 optimization, we extract per-object meshes using the occupancy field. Following MILo [14], we construct a Delaunay tetrahedralization of Gaussian positions and apply marching tetrahedra [30] at the $\phi = 0.5$ isosurface. The resulting meshes exhibit high quality with balanced vertex distributions, as ALNC naturally consolidates Gaussians and suppresses outliers, leading to more complete surface coverage and fewer artifacts compared to direct extraction from unoptimized Gaussians.

5 Experiments

We thoroughly evaluate the effectiveness of MagnetGS-Mesh by focusing on three key aspects: (1) surface-preserving mesh extraction quality and surface detection accuracy, (2) storage efficiency through compact mesh representations, and (3) high-quality novel view synthesis preserved after mesh extraction.

5.1 Experimental Setup

Implementation Details We implement our method in PyTorch [35] and PyTorch3D [38], and conduct all experiments on a single NVIDIA RTX 4090 GPU. Our geometry-aligned two-stage pipeline consists of Enhanced Occupancy Learning (Stage 1) followed by ALNC-guided optimization (Stage 2). Stage 1 trains with a learning rate schedule that decays over iterations, incorporating five regularization terms for enhanced occupancy learning. Stage 2 optimization uses cosine annealing for the learning rate and operates in two phases: Phase 1 applies proximity loss, while Phase 2 incorporates both proximity loss and ALNC Loss. We employ the Adam optimizer with gradient clipping for stable training. We use identical parameters across all datasets without per-dataset tuning. For quantitative evaluation, we extract per-object meshes from the occupancy field after Stage 2 optimization and combine them to form complete scene meshes, then render these meshes from all test camera viewpoints using MILo [14]’s ScalableMeshRenderer [14] and compute PSNR [1], SSIM [46], and LPIPS [52] metrics by comparing rendered images with ground-truth images. All metrics are averaged across test views.

Storage Efficiency Table 3 compares vertex counts and mesh file sizes across all methods on the LERF [21] dataset. Our method consistently delivers the most compact representations, significantly reducing both vertex count and storage footprint compared to prior work [14, 18, 50] while maintaining or improving rendering quality (Table 1). This validates that ALNC Loss effectively consolidates scattered Gaussians into dense surface coverage without sacrificing visual fidelity.

Table 1: Ground Truth vs. Mesh-rendered Images. Quantitative comparison on Mip-NeRF 360 [1] and LERF [21] datasets. All metrics are computed by rendering the extracted meshes from test camera viewpoints and comparing against ground-truth images. We report average PSNR (dB), SSIM, and LPIPS across all test views. Results are color-coded: **red** for best, **orange** for 2nd best, and **yellow** for 3rd best.

Method	Mip-NeRF 360			LERF		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
2DGS [18]	15.36	0.499	0.475	14.97	0.4726	0.4457
PGSR [6]	16.43	0.527	0.406	18.55	0.652	0.340
SVRaster [40]	18.72	0.549	0.406	18.71	0.663	0.336
GOF [50]	20.78	0.573	0.465	18.40	0.6425	0.3541
MILo [14]	24.09	0.689	0.324	18.93	0.6820	0.2975
Ours	25.68	0.766	0.312	19.85	0.7516	0.3088

5.2 Quantitative Results

Mip-NeRF 360 Dataset As shown in Table 1, our method achieves state-of-the-art performance on Mip-NeRF 360 [1], outperforming recent methods including PGSR [6] and SVRaster [40]. The storage efficiency gains validate that our magnetic consolidation strategy generalizes to unbounded outdoor scenes with complex geometry.

LERF Dataset Our method achieves the best PSNR [1] and SSIM [46] across all LERF [21] scenes (Table 1), demonstrating superior rendering fidelity compared to baselines [14, 18, 50]. While MILo [14] achieves slightly better LPIPS [52], our method delivers competitive perceptual quality while simultaneously reducing mesh file sizes by 27% on average.

Geometric Evaluation on DTU Table 2 reports Chamfer distance on the DTU benchmark. Although DTU targets object-centric meshes while our pipeline targets large-scale scenes, we achieve competitive Chamfer distance (0.66 without cell culling). Notably, all competing methods rely on cell culling as a post-processing step to remove background artifacts, whereas our method achieves competitive performance without any such post-processing. ALNC applies small k for precision in dense regions and large k for robustness in sparse regions, preserving sharp features better than uniform- k approaches. When cell culling is applied to our method, Chamfer distance slightly worsens (0.66→0.71) because ALNC has already relocated outliers onto surfaces—leaving no floaters to remove—so culling instead removes valid surface Gaussians, including thin structures.

Geometric Evaluation on Tanks & Temples Table 4 reports F1 scores on the Tanks & Temples benchmark, which evaluates large-scale outdoor reconstruction. Our method achieves the highest mean F1 score (0.51) among all compared methods, outperforming MILo [14] (0.49) and GOF [50] (0.46). These results demonstrate that the ALNC Loss generalizes effectively to unbounded

Table 2: Chamfer distance and training time on DTU. Results are color-coded: **red** for best, **orange** for 2nd best, and **yellow** for 3rd best.

Method	Cf.↓	Time↓
2DGS [18]	0.80	8.9m
SVRaster [40]	0.76	6.2m
GOF [50]	0.74	55m
MILo [14]	0.68	25m
PGSR [6]	0.52	0.5h
Ours (culling on)	0.71	28m
Ours (culling off)	0.66	28m

Table 3: Storage efficiency comparison on LERF dataset. Our method achieves the smallest mesh sizes while maintaining high rendering quality. Results are color-coded: **red** for best, **orange** for 2nd best, and **yellow** for 3rd best.

Method	Verts↓	Storage Size↓
2DGS [18]	9.3 M	518.6 MB
GOF [50]	1.8 M	71.8 MB
MILo [14]	0.4 M	17.1 MB
Ours	0.3 M	12.4 MB

Table 4: F1 score on Tanks & Temples [23]. Higher is better. Results are color-coded: **red** for best, **orange** for 2nd best, and **yellow** for 3rd best.

Scene	SuGaR [15]	2DGS [18]	GOF [50]	MILo [14]	Ours
Barn	0.14	0.36	0.51	0.59	0.62
Caterpillar	0.16	0.23	0.41	0.39	0.38
Courthouse	0.08	0.13	0.28	0.29	0.34
Ignatius	0.33	0.44	0.68	0.78	0.76
Meetingroom	0.15	0.16	0.28	0.26	0.30
Truck	0.26	0.26	0.59	0.62	0.66
Mean	0.19	0.30	0.46	0.49	0.51

outdoor scenes. Note that the 2DGS numbers follow the evaluation pipeline of MILo [14].

5.3 Qualitative Results

Visual Quality Figure 4 shows visual comparisons on representative scenes. Our method produces cleaner mesh surfaces with fewer artifacts, particularly in sparse regions (*e.g.*, thin structures, curved surfaces) where baselines [14, 18, 50] exhibit holes or noise. The magnetic consolidation effect is evident in object boundaries, where ALNC-guided optimization naturally aligns Gaussians with local surface normals [16], producing smoother transitions and more faithful geometric reconstruction. Figure 5 further highlights this on the garden sequence (Mip-NeRF 360 [1]): in the flower region, 2DGS [18], GOF [50], and MILo [14] fail to capture fine petal details and leave noticeable surface holes, whereas our method faithfully reconstructs delicate geometric structures.

Surface Detection & Segmentation Quality Fig. 6 presents side-by-side analysis of surface detection and segmentation. In surface-outlier detection (Fig. 6-i), MILo [14]’s approach shows substantial confusion with outlier Gaussians persisting inside object regions. Our method achieves significantly cleaner separation. In segmentation (Fig. 6-ii), our Stage 1 Enhanced Occupancy Learn-




Fig. 4: Qualitative comparison on Mip-NeRF 360 [1] and LERF [21] datasets. Our method produces cleaner mesh surfaces with fewer artifacts compared to baselines while achieving substantial storage reduction. On the MipNeRF 360 bicycle dataset, using GOF (1,457.4 MB) as baseline, 2DGS (670.3 MB) reduces storage by 54.01%, MILO (172.6 MB) by 88.16%, and our method (152.3 MB) achieves the highest reduction of 89.55%.





Fig. 5: Mesh extraction results on the flower region of the garden sequence from Mip-NeRF 360 dataset [1]. Comparison of ground-truth image, 2DGS [18], GOF [50], MILO [14], and Ours methods.

ing alone already produces significantly sharper 3D boundaries compared to SAGA [3] used independently, reducing outlier Gaussians before Stage 2 even begins. MagnetGS-Mesh successfully segments and reconstructs individual objects as independent high-quality meshes, enabling downstream applications such as independent editing and relighting.



(i) **Surface-Outlier detection.** (a) GT segmentation. (b) 3D Gaussians. (c) MILO detection. (d) Ours: cleaner surface-outlier separation.

(ii) **Segmentation quality (Stage 1 only).** SAGA [3] baseline (left) vs. our Stage 1 result (right). Stage 1 Enhanced Occupancy Learning alone produces sharper 3D boundaries and fewer outlier Gaussians than SAGA independently.

Fig. 6: Surface detection and segmentation quality. Our Enhanced Occupancy Learning consistently produces cleaner geometric boundaries and sharper 3D boundaries compared to baseline methods.




Fig. 7: Enhanced Occupancy Learning ablation. (a) Ground truth. (b) MILO baseline. (c) Ours with Stage 1 only (without segmentation or ALNC). Stage 1 alone improves rendering quality over baseline MILO.

5.4 Ablation Studies

We validate each component’s contribution through comprehensive ablation studies, analyzing Stage 1 regularization losses, Stage 2 ALNC Loss variants, hyperparameter sensitivity, and geometric quality.

Stage 1: Enhanced Occupancy Learning As shown in Figure 7 and Tables 5 and 6, each regularization loss progressively improves over baseline MILO [14], validating the design necessity of all five components. Boundary sharpening ($\mathcal{L}_{\text{boundary}}$) contributes most to PSNR improvement, while spatial coherence ($\mathcal{L}_{\text{smooth}}$) most improves SSIM. The full Stage 1 configuration achieves the best overall balance across all metrics.

Stage 2: ALNC Loss Variants Tables 7 and 8 ablate the key components and hyperparameter sensitivity of our ALNC Loss. Using a single normal direction without averaging degrades performance. Proximity loss alone (without normal consistency) improves surface attachment but lacks geometric coherence. Fixed $k = 10$ without density-adaptive adjustment performs well but falls short of the full adaptive mechanism. The full ALNC with adaptive k achieves the best results across all metrics: by assigning smaller k in dense regions for precise local normal estimation and larger k in sparse regions for noise-robust normal estimation, it

Table 5: Stage 1 ($\mathcal{L}_{\text{enh}}$) ablation on LERF [21]. GT vs. mesh-rendered. Each stage progressively added. Size (MB): mesh storage. **red** best, **orange** 2nd, **yellow** 3rd.

Configuration	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Size $\downarrow$
Base (No Enh.)	18.93	0.6820	0.2975	17.1
+ $\mathcal{L}_{\text{enh}}$	19.24	0.7105	0.3002	15.8
+Seg. [3]	19.61	0.7389	0.3040	14.2
Full (+$\mathcal{L}_{\text{ALNC}}$)	19.85	0.7516	0.3088	12.4

Table 7: Stage 2 ($\mathcal{L}_{\text{ALNC}}$) ablation on LERF [21]. GT vs. mesh-rendered. Each ALNC component ablated. Size (MB): mesh storage. **red** best, **orange** 2nd, **yellow** 3rd.

Configuration	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Size $\downarrow$
Single Normal	19.68	0.743	0.312	14.8
Proximity Only	19.70	0.744	0.310	13.9
Fixed $k = 10$	19.76	0.747	0.310	13.2
Full ALNC	19.85	0.7516	0.3088	12.4

Table 6: Stage 1 detailed loss ablation on LERF [21]. GT vs. mesh-rendered. Each regularization loss progressively added to MILo [14]. **red** best, **orange** 2nd, **yellow** 3rd.

Configuration	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Baseline (MILo) [14]	18.93	0.6820	0.2975
+ $\mathcal{L}_{\text{entropy}}$	19.05	0.6950	0.2988
+ $\mathcal{L}_{\text{smooth}}$	19.18	0.7120	0.3002
+ $\mathcal{L}_{\text{consist}}$	19.32	0.7250	0.3025
+ $\mathcal{L}_{\text{boundary}}$	19.48	0.7310	0.3042
+ $\mathcal{L}_{\text{temporal}}$	19.56	0.7380	0.3058
+Seg. [3]	19.72	0.7450	0.3072
Full (+$\mathcal{L}_{\text{ALNC}}$)	19.85	0.7516	0.3088

Table 8: Stage 2 ALNC adaptive k range sensitivity analysis on LERF [21]. GT vs. mesh-rendered. Default ($k_{\min} = 3$, $k_{\max} = 15$) is best. **red** best, **orange** 2nd, **yellow** 3rd.

$k_{\min}$	$k_{\max}$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
3	20	19.74	0.742	0.316
5	15	19.76	0.744	0.314
3	10	19.78	0.746	0.312
1	15	19.81	0.749	0.315
3	15	19.85	0.7516	0.3088

effectively handles uneven Gaussian distributions that a fixed k cannot address uniformly. Moreover, Table 8 shows that performance remains consistently high across different k range configurations, confirming that adaptivity itself—rather than specific parameter values—is the key driver of performance.

6 Conclusion

We presented MagnetGS-Mesh, a two-stage framework for high-quality object-wise mesh extraction from 3DGS [20]. Our core contribution, the ALNC Loss, consolidates scattered Gaussians onto surface regions via density-adaptive normal alignment, resolving the non-uniform Gaussian distribution introduced by object-wise segmentation. Enhanced Occupancy Learning provides the clean surface-outlier separation required for ALNC to operate effectively. Together, they consistently outperform state-of-the-art methods [14, 18, 50] in rendering quality while achieving up to 27% storage reduction, enabling practical applications in editing, animation, and relighting [2, 51]. While optimized for large-scale multi-object scenes, extending its applicability to object-centric scenarios remains a promising future direction.

Acknowledgements

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No.RS-2026-25517417, Development of Compression, Reconstruction, and Rendering Technologies for Free-viewpoint Media). This work was also supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2024-00457605).

References

1. 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 IEEE/CVF conference on computer vision and pattern recognition. pp. 5470–5479 (2022)
2. Boss, M., Braun, R., Jampani, V., Barron, J.T., Liu, C., Lensch, H.: Nerd: Neural reflectance decomposition from image collections. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12684–12694 (2021)
3. Cen, J., Fang, J., Yang, C., Xie, L., Zhang, X., Shen, W., Tian, Q.: Segment any 3d gaussians. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 1971–1979 (2025)
4. Cen, J., Zhou, Z., Fang, J., Shen, W., Xie, L., Jiang, D., Zhang, X., Tian, Q., et al.: Segment anything in 3d with nerfs. *Advances in Neural Information Processing Systems* **36**, 25971–25990 (2023)
5. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: European conference on computer vision. pp. 333–350. Springer (2022)
6. Chen, D., Li, H., Ye, W., Wang, Y., Xie, W., Zhu, S., Wang, N., Liu, H., Bao, H., Zhang, G.: Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics* (2024)
7. Dai, P., Xu, J., Xie, W., Liu, X., Wang, H., Xu, W.: High-quality surface reconstruction using gaussian surfels. In: ACM SIGGRAPH 2024 conference papers. pp. 1–11 (2024)
8. Darmon, F., Bascle, B., Devaux, J.C., Monasse, P., Aubry, M.: Improving neural implicit surfaces geometry with patch warping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6260–6269 (2022)
9. Fridovich-Keil, S., Meanti, G., Warburg, F.R., Recht, B., Kanazawa, A.: K-planes: Explicit radiance fields in space, time, and appearance. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12479–12488 (2023)
10. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5501–5510 (2022)
11. Fu, X., Zhang, S., Chen, T., Lu, Y., Zhu, L., Zhou, X., Geiger, A., Liao, Y.: Panoptic nerf: 3d-to-2d label transfer for panoptic urban scene segmentation. In: 2022 International Conference on 3D Vision (3DV). pp. 1–11. IEEE (2022)
12. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. *Advances in neural information processing systems* **17** (2004)

13. Gropp, A., Yariv, L., Haim, N., Atzmon, M., Lipman, Y.: Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099* (2020)
14. Guédon, A., Gomez, D., Maruani, N., Gong, B., Drettakis, G., Ovsjanikov, M.: Milo: Mesh-in-the-loop gaussian splatting for detailed and efficient surface reconstruction. *arXiv preprint arXiv:2506.24096* (2025)
15. Guédon, A., Lepetit, V.: Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5354–5363 (2024)
16. Guerrero, P., Kleiman, Y., Ovsjanikov, M., Mitra, N.J.: Pepnet learning local shape properties from raw point clouds. In: *Computer graphics forum*. vol. 37, pp. 75–85. Wiley Online Library (2018)
17. He, Y., Yu, H., Liu, X., Yang, Z., Sun, W., Anwar, S., Mian, A.: Deep learning based 3d segmentation: A survey. *arXiv preprint arXiv:2103.05423* (2021)
18. Huang, B., Yu, Z., Chen, A., Geiger, A., Gao, S.: 2d gaussian splatting for geometrically accurate radiance fields. In: *ACM SIGGRAPH 2024 conference papers*. pp. 1–11 (2024)
19. Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: *Proceedings of the fourth Eurographics symposium on Geometry processing*. vol. 7 (2006)
20. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139–1 (2023)
21. Kerr, J., Kim, C.M., Goldberg, K., Kanazawa, A., Tancik, M.: Lerf: Language embedded radiance fields. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 19729–19739 (2023)
22. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4015–4026 (2023)
23. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)* **36**(4), 1–13 (2017)
24. Kundu, A., Genova, K., Yin, X., Fathi, A., Pantofaru, C., Guibas, L.J., Tagliasacchi, A., Dellaert, F., Funkhouser, T.: Panoptic neural fields: A semantic object-aware neural scene representation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12871–12881 (2022)
25. Kusupati, U., Cheng, S., Chen, R., Su, H.: Normal assisted stereo depth estimation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2189–2199 (2020)
26. Kwak, K., Huber, D.F., Chae, J., Kanade, T.: Boundary detection based on supervised learning. In: *2010 IEEE International Conference on Robotics and Automation*. pp. 3939–3945. IEEE (2010)
27. Li, J., Lu, Z., Wang, Y., Xiao, J., Wang, Y.: Nr-mvsnet: Learning multi-view stereo based on normal consistency and depth refinement. *IEEE Transactions on Image Processing* **32**, 2649–2662 (2023)
28. Li, Y., Lyu, C., Di, Y., Zhai, G., Lee, G.H., Tombari, F.: Geogaussian: Geometry-aware gaussian splatting for scene rendering. In: *European conference on computer vision*. pp. 441–457. Springer (2024)
29. Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8456–8465 (2023)

30. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. *ACM SIGGRAPH Computer Graphics* **21**(4), 163–169 (1987)
31. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
32. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)* **41**(4), 1–15 (2022)
33. Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S., Geiger, A., Radwan, N.: Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5480–5490 (2022)
34. Oechsle, M., Peng, S., Geiger, A.: Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 5589–5599 (2021)
35. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
36. Qin, M., Li, W., Zhou, J., Wang, H., Pfister, H.: Langsplat: 3d language gaussian splatting. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20051–20060 (2024)
37. Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al.: Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714* (2024)
38. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501* (2020)
39. Shi, J.C., Wang, M., Duan, H.B., Guan, S.H.: Language embedded 3d gaussians for open-vocabulary scene understanding. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5333–5343 (2024)
40. Sun, C., Liang, J., Wang, Z., Geiger, A., Tancik, M., Kanazawa, A.: Sparse voxels rasterization: Real-time high-fidelity radiance field rendering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025)
41. Sun, C., Sun, M., Chen, H.T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5459–5469 (2022)
42. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* **30** (2017)
43. Tchapmi, L., Choy, C., Armeni, I., Gwak, J., Savarese, S.: Segcloud: Semantic segmentation of 3d point clouds. In: *2017 international conference on 3D vision (3DV)*. pp. 537–547. *IEEE* (2017)
44. Toshev, A., Taskar, B., Daniilidis, K.: Object detection via boundary structure segmentation. In: *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. pp. 950–957. *IEEE* (2010)
45. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689* (2021)

46. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
47. Yang, Y., Wu, X., He, T., Zhao, H., Liu, X.: Sam3d: Segment anything in 3d scenes. *arXiv preprint arXiv:2306.03908* (2023)
48. Ye, M., Danelljan, M., Yu, F., Ke, L.: Gaussian grouping: Segment and edit anything in 3d scenes. In: *European conference on computer vision*. pp. 162–179. Springer (2024)
49. Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenotrees for real-time rendering of neural radiance fields. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 5752–5761 (2021)
50. Yu, Z., Sattler, T., Geiger, A.: Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Transactions on Graphics (ToG)* **43**(6), 1–13 (2024)
51. Zhang, K., Luan, F., Wang, Q., Bala, K., Snavely, N.: Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5453–5462 (2021)
52. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
53. Zhi, S., Laidlow, T., Leutenegger, S., Davison, A.J.: In-place scene labelling and understanding with implicit scene representation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 15838–15847 (2021)