

Topology-Aware 3D Gaussian Splatting: Leveraging Persistent Homology for Optimized Structural Integrity

Tianqi Shen^{1,4}, Shaohua Liu^{2,3}, Jiaqi Feng², Ziyi Ma¹, Ning An^{4,5*}

¹Department of Computer Science, City University of Hong Kong

²Image Processing Center, Beihang University

³Shen Yuan Honors College, Beihang University

⁴Research Institute of Mine Artificial Intelligence, China Coal Research Institute

⁵State Key Laboratory of Intelligent Coal Mining and Strata Control

tianqshen5-c@my.cityu.edu.hk, liushaohua@buaa.edu.cn, fengjiaqi@buaa.edu.cn,
ziyema@cityu.edu.hk, ning.an@ccteg-bigdata.com

Abstract

Gaussian Splatting (GS) has emerged as a crucial technique for representing discrete volumetric radiance fields. It leverages unique parametrization to mitigate computational demands in scene optimization. This work introduces Topology-Aware 3D Gaussian Splatting (Topology-GS), which addresses two key limitations in current approaches: compromised pixel-level structural integrity due to incomplete initial geometric coverage, and inadequate feature-level integrity from insufficient topological constraints during optimization. To overcome these limitations, Topology-GS incorporates a novel interpolation strategy, Local Persistent Voronoi Interpolation (LPVI), and a topology-focused regularization term based on persistent barcodes, named PersLoss. LPVI utilizes persistent homology to guide adaptive interpolation, enhancing point coverage in low-curvature areas while preserving topological structure. PersLoss aligns the visual perceptual similarity of rendered images with ground truth by constraining distances between their topological features. Comprehensive experiments on three novel-view synthesis benchmarks demonstrate that Topology-GS outperforms existing methods in terms of PSNR, SSIM, and LPIPS metrics, while maintaining efficient memory usage. This study pioneers the integration of topology with 3D-GS, laying the groundwork for future research in this area.

Code — <https://github.com/AmadeusSTQ/Topology-GS>

Introduction

Novel-view synthesis (NVS) is a critical area in computer vision and graphics, aimed at generating new views of a scene from a limited set of input images (Yan et al. 2021). Significant advancements in this area have been driven by techniques such as Neural Radiance Fields (NeRF) (Mildenhall et al. 2021; Zhang et al. 2020; Barron et al. 2021; Xu et al. 2023) and 3D Gaussian Splatting (3D-GS) (Kerbl et al. 2023; Yu et al. 2024; Li et al. 2024; Guédon and Lepetit 2024). NeRF employs an implicit representation that models a scene as a continuous volumetric field, optimized through

neural networks. Despite its impressive results, NeRF is computationally intensive due to volumetric ray-marching (Fridovich-Keil et al. 2022; Müller et al. 2022). In contrast, 3D-GS offers an efficient point-based representation by utilizing Gaussian primitives (hereafter referred to as Gaussians), parametrized by their positions μ and covariance matrices Σ , along with learnable attributes for color c and opacity o . Combined with a tile-based rasterization pipeline, 3D-GS achieves high-fidelity rendering results and real-time rendering speeds (Wu et al. 2024).

The density function of a 3D Gaussian in 3D-GS is defined as:

$$f(\mathbf{x} \mid \mu, \Sigma) = \exp\left\{-\frac{1}{2}(\mathbf{x} - \mu)^T \Sigma^{-1}(\mathbf{x} - \mu)\right\}, \quad (1)$$

where \mathbf{x} is an arbitrary 3D point. To ensure the covariance matrix is positive semi-definite, (Kerbl et al. 2023) proposed decomposing it using rotation and scaling: $\Sigma = \mathbf{R}\mathbf{S}\mathbf{S}^T\mathbf{R}^T$, with the rotation represented using quaternions.

Finally, the 3D Gaussians are projected onto a 2D image, and their corresponding 2D pixel colors are determined by blending the N ordered Gaussians at the queried pixel as:

$$\mathbf{C} = \sum_{i=1}^N o_i \prod_{j=1}^{i-1} (1 - o_j) \mathbf{c}_i, \quad (2)$$

where \mathbf{c}_i and o_i represent the color and opacity of Gaussian i , respectively.

Typically, 3D-GS initializes *plain* Gaussian structures with a *sparse* point cloud generated by Structure from Motion (SfM) (Schönberger and Frahm 2016), a photogrammetric technique that reconstructs 3D structures from sequences of 2D images by matching features across multiple images. During optimization, an Adaptive Density Control (ADC) mechanism (Kerbl et al. 2023) manages the growth or pruning of each Gaussian *independently* to model *complex* structures. As illustrated in Figure 1, these design choices present two primary limitations that hinder the rendered images from accurately reflecting the structural integrity of the scenes. Here, structural integrity refers to *the consistency of geometric shapes and topological features throughout the 3D reconstruction or image rendering process*, ensuring the preservation of local texture patterns and

*Corresponding author: Ning An (ning.an@ccteg-bigdata.com)
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

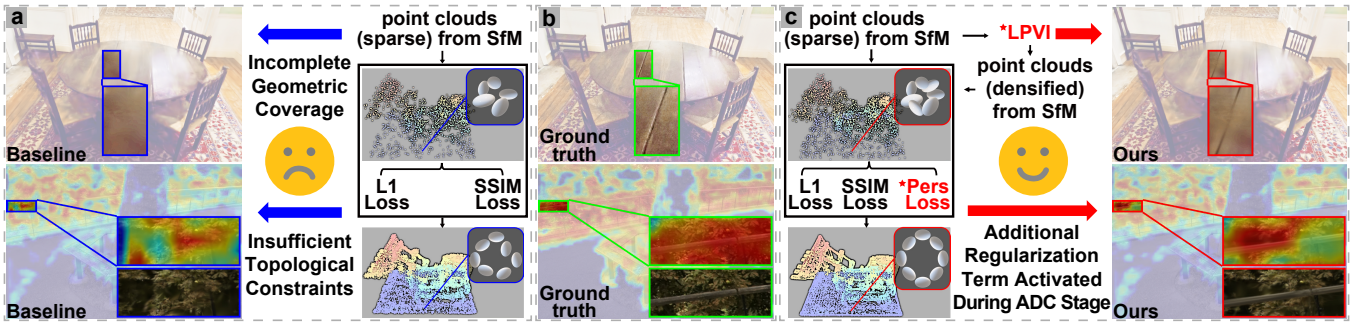


Figure 1: Illustration of main motivations and our contributions. **a** Typical pipeline of state-of-the-art (SOTA) methods and visualization of their two main limitations. **b** Visualization of the ground truth. **c** Pipeline of our method (i.e., Topology-GS) and visualization showing improvements over existing methods. Feature maps, visualized using EigenCAM (Muhammad and Yeasin 2020), are from the last layer of VGG16 (Simonyan 2014), one of the layers used to compute the LPIPS (Zhang et al. 2018) metric. The differences in point clouds & Gaussians before and after training for both the baseline and our method are illustrated using Lego from (Mildenhall et al. 2021) as a conceptual example.

global semantic relationships. Below, we provide a detailed analysis of these limitations and propose solutions using Persistent Homology (PH) (Huber 2021) from topological data analysis (TDA) (Chazal and Michel 2021).

Firstly, since SfM relies heavily on the density and quality of feature matches, which can fluctuate substantially across different image regions, the resulting point clouds often exhibit notably sparse *geometric coverage* in areas with low curvature (Ververas et al. 2025). As shown in Figure 1 (upper left, blue arrow), this sparsity compromises 3D-GS rendering quality with the loss of planar texture details, undermining **pixel-level structural integrity**. Furthermore, the ADC mechanism in 3D-GS can cause significant Gaussian displacement in highly sparse areas (Ververas et al. 2025). Similarly, other point cloud enhancement methods (Niemeyer et al. 2024; Foroutan et al. 2024; Lee et al. 2025; Ververas et al. 2025) also carry a common risk of disrupting the topological structure.

To tackle the specific challenge in 3D-GS, we propose a novel approach called Local Persistent Voronoi Interpolation (LPVI). LPVI introduces a unique TopoDiff check, which leverages differences in persistent homology as a guiding criterion. This mechanism enables adaptive switching between preserving topological structures in 3D space and performing interpolation in 2D space, effectively mitigating the sparsity issue in low-curvature regions. By adaptively increasing point coverage in these areas, LPVI enhances surface detail without compromising overall structural integrity. As shown in Figure 1 (upper right, red arrow), the interpolation process significantly improves the rendering of low-curvature surfaces, recovering previously missing textures and ultimately refining pixel-level structural consistency.

Secondly, the lack of *topological constraints* during training results in poor **feature-level structural integrity**. As shown in Figure 1 (lower left, indicated by the blue arrow), existing methods often use pixel-level loss functions (L1 loss and SSIM loss) to guide the independent optimization of Gaussians. This results in high PSNR and SSIM (Wang et al. 2004) metrics but also high LPIPS (Zhang et al.

2018) metrics, which measure visual perceptual similarity by calculating the distances between the feature maps of the rendered and ground truth images. Some alternative approaches (Lu et al. 2024; Yang et al. 2024a; Turkulainen et al. 2024; Guédon and Lepetit 2024) impose structural constraints on Gaussians but do not ensure the reduction of distances between feature maps of rendered images, while relying on deeper networks to reduce LPIPS is not practical for lightweight applications.

To address the specific challenge in 3D-GS, we introduce a topology-based regularization term inspired by persistent homology, termed PersLoss. This term assesses the persistent homology of both rendered and ground truth images by analyzing their truncated persistence barcodes, which represent the birth and death times of topological features with top-k lifespans, with its differentiability enabled by techniques in (Gabrielsson et al. 2020; Zhang et al. 2022). Activated during the ADC phase, it ensures the alignment of 2D rendered images with the abstract features of the ground truth, directly guiding the visual perceptual similarity between rendered images and the ground truth. As shown in Figure 1 (lower right, red arrow), post-training, the rendered images exhibit improved structural & semantic information and reduced LPIPS, enhancing overall visual quality without incurring extra computational overhead during rendering.

As illustrated in Figure 1&5, extensive experiments across three NVS benchmarks demonstrate that our method (i.e. Topology-GS), which incorporates the proposed LPVI and PersLoss, significantly enhances PSNR and SSIM metrics, while reducing LPIPS. This study represents the first successful integration of topology with 3D-GS, highlighting a substantial synergy between the two.

Notations

The rendered image obtained from 3D-GS is denoted by $y \in \mathbb{R}^{3 \times H \times W}$, where H represents the height and W represents the width of the image. Correspondingly, the ground truth image is represented by $\hat{y} \in \mathbb{R}^{3 \times H \times W}$. Let $Vor^{nD}(x)$ denote the vertices of the n -dimensional Voronoi region of

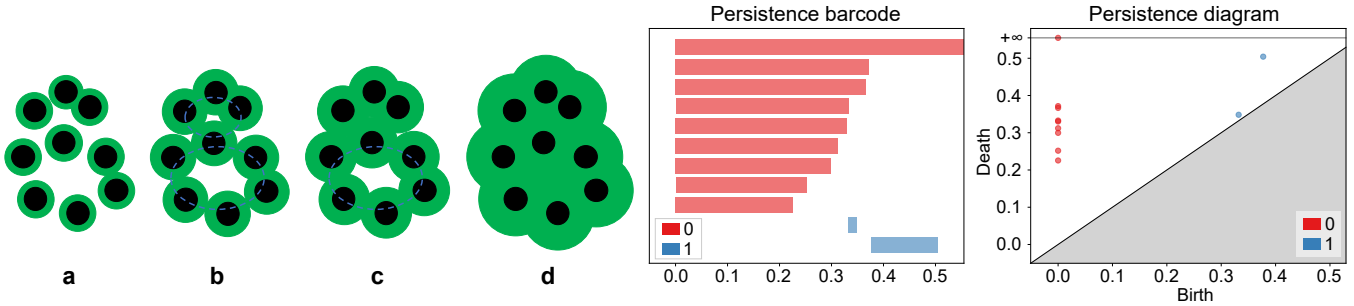


Figure 2: An illustration of persistent homology. Panels a to d depict the evolution of the topological structure when a filtration based on the Euclidean distance (indicated in green) is applied to a set of black points. The two images on the right display the corresponding persistence barcode and persistence diagram. Blue bars or points represent 0-dimensional homology, i.e., connected components, while red bars or points represent 1-dimensional homology, i.e., loops.

a point x . The symbol H^k represents the k -dimensional homology group, which quantifies various topological features such as connected components (H^0), loops (H^1), and voids (H^2). The Betti numbers, denoted by β^i , indicate the rank of the i -th homology group, reflecting the count of k -dimensional features. The operation $PD(\cdot)$ denotes the process of obtaining the persistence diagram from a three-dimensional point set. The scalar-valued function f_W is a filter function with model parameters W .

Related Works

NeRF and 3D-GS

NeRF (Mildenhall et al. 2021) uses a Multi-Layer Perceptron (MLP) to represent the radiance field and volume rendering for realistic images. However, it faces high computational costs, long training times, and slow rendering speeds due to volumetric ray-marching. Methods like Plenoxels (Fridovich-Keil et al. 2022) and InstantNGP (Müller et al. 2022) aim to address these issues by using sparse voxel grids and hash-grid encodings, respectively, but the need for sample queries still hampers rendering speed.

3D-GS extends NeRF by incorporating Gaussians with learnable attributes and a tile-based rasterization pipeline, enabling faster and higher-fidelity rendering. While advancements have been achieved in areas such as anti-aliasing (Yu et al. 2024; Song et al. 2024), high-frequency signal modeling (Yang et al. 2024b), few-shot view synthesis (Zhu et al. 2025; Li et al. 2024), and 3D mesh reconstruction (Guédon and Lepetit 2024), two main challenges persist.

Firstly, the quality of the initial point cloud coverage significantly impacts the reconstruction quality. However, the ADC mechanism, which aims to mitigate this issue, often introduces significant Gaussian displacement in sparse areas (Ververas et al. 2025). Existing solutions, such as optimizing NeRF as a prior (Niemeyer et al. 2024; Foroutan et al. 2024), expanding point clouds (Lee et al. 2025), and curvature-based enhancements (Ververas et al. 2025), risk disrupting the topology of the scene. To address this, we introduce LPVI, a novel approach that refines sparse coverage while preserving the topology of the scene. *Secondly*, the heuristic optimization of 3D-GS often neglects scene structural con-

straints, leading to poor feature-level integrity. While integrating with deep neural networks can improve ground truth similarity, it comes at the cost of computational efficiency (Zou et al. 2024). Recent approaches, such as hierarchical scene representation (Lu et al. 2024), graph-based structural modeling (Yang et al. 2024a), and geometric priors (Turkulainen et al. 2024), have improved the structural quality of the reconstructed scenes. However, these methods do not directly guide semantic feature similarity. To bridge this gap, we develop PersLoss, a novel topological constraint that optimizes feature map distances, enhancing visual perceptual similarity, which represents a semantic-level similarity.

Topological Data Analysis

Recent advances in topological data analysis (TDA) have enhanced topology-based machine learning techniques (Chazal and Michel 2021). Persistent homology (PH) (Huber 2021), a key method in TDA, captures topological features like "holes" in data, which are invariant to noise and scale variations, and analogous to high-dimensional abstract features in topological neural networks (Reinauer, Caorsi, and Berkouk 2021; de Surrel et al. 2022; Nishikawa, Ike, and Yamanishi 2024). In our study, using alpha complexes is suitable for point cloud representation (DiFrancesco, Bonneau, and Hutchinson 2020).

An alpha complex is derived from the Delaunay triangulation (Preparata and Shamos 2012) of a point set, filtered by a parameter α . For points $P \subset \mathbb{R}^d$, the Delaunay triangulation yields simplices included in the alpha complex if circumscribed by a sphere of radius $\sqrt{\alpha}$. The nested sequence of simplicial complexes forms a filtration:

$$\emptyset = K_0 \subseteq K_1 \subseteq \dots \subseteq K_n = K, \quad (3)$$

where each K_i relates to an alpha complex for some α_i .

As illustrated in Figure 2, PH tracks the birth and death of topological features as α changes, visualized through persistence diagrams (PD) and persistence barcodes (PB). The PD, a scatter plot, maps each point $(\alpha_{birth}, \alpha_{death})$ to the birth and death of a topological feature. Points further from the diagonal $\alpha_{birth} = \alpha_{death}$ indicate more significant features. The PB consists of a series of intervals, each representing the lifespan of a topological feature. Longer bars denote features




Figure 3: An illustration of LPVI. Figure b shows the principle of LPVI, which involves switching between 3D and 2D local Voronoi interpolation based on topological differences. Figures a and c show point clouds of the table, with a displaying the initial sparse point cloud and c showing the point cloud enhanced by LPVI. Figure d presents the final rendered result, where the green shading visualizes the effect of the interpolated point cloud rendered onto the planar image.

that endure across a broad range of α values, denoting their significance, while shorter bars typically suggest noise.

PH has facilitated various tasks (Pun, Lee, and Xia 2022; Zia et al. 2024), integrating topological features into machine learning models to capture data structures (Wang et al. 2020; Hu et al. 2024), and employing PH-based loss functions to train neural networks preserving topological characteristics (Hu et al. 2019; Clough et al. 2020; Qi et al. 2023). Incorporating TDA into 3D-GS, as shown in this work, leverages PH to preserve the structural integrity of point clouds during the initialization and optimization of Gaussians. Recent studies (Jignasu et al. 2024; Nishikawa, Ike, and Yamanishi 2024) highlight the potential of PH in this context, demonstrating its effectiveness in capturing the structural properties of point clouds. Building on these advancements, our approach further enhances the rendering quality of 3D scene representations.

Methodology

Local Persistent Voronoi Interpolation

In this study, we propose the Local Persistent Voronoi Interpolation (LPVI), as illustrated in Figure 3.b, which visually demonstrates the mechanism of LPVI. This algorithm interpolates sparse 3D point clouds while preserving the local topological structure by adaptively switching between 3D and 2D Voronoi Interpolation based on changes in topological features before and after interpolation. The pseudocode is presented in Algorithm 1. To further illustrate the ideas behind LPVI, we break down the procedures in Algorithm 1 and explain them step by step.

3D Voronoi Interpolation The execution of LPVI begins by defining an empty set for processed point indices and another for interpolated points. The main loop iterates over each point x_l in the initial sparse point cloud, processing each point only once. For each unprocessed point, the algorithm identifies its K -nearest neighbors and performs a 3D Voronoi tessellation (Tanemura, Ogawa, and Ogita 1983). The vertices of this tessellation, shown as the orange points in Figure 3.b, are potential candidates for interpolation.

Topological Assessment A critical step in the algorithm is the evaluation of topological changes, facilitated by the *TopoDiff* function, which computes the Wasserstein distance

W_Dist between PDs of two point sets \mathcal{X}_l and $\hat{\mathcal{X}}_l$:

$$W_Dist(PD(\mathcal{X}_l), PD(\hat{\mathcal{X}}_l)), \quad (4)$$

where \mathcal{X}_l is the K -nearest neighbor set centered at x_l , and $\hat{\mathcal{X}}_l$ is this set plus the interpolation points from the 3D Voronoi tessellation.

Algorithm 1: Local Persistent Voronoi Interpolation (LPVI)

Require: $\mathcal{X} = \{x_1, \dots, x_m\} \subset \mathbb{R}^3$: a sparse point set
Require: K : the number of neighbors for 3D Voronoi
Require: K' : the number of neighbors for 2D Voronoi
Require: τ : the threshold of topological differences
Ensure: $\hat{\mathcal{X}} \subset \mathbb{R}^3$: an interpolated point set

- 1: $I \leftarrow \emptyset, \hat{\mathcal{X}} \leftarrow \emptyset, \mathcal{M} = \{1, \dots, m\} \leftarrow$ indices of \mathcal{X}
- 2: **for** $l \leftarrow 1$ **to** m **do**
- 3: **if** $l \notin I$ **then**
- 4: $I \leftarrow I \cup \{l\}$
- 5: $\{\mu_1, \dots, \mu_K\}_{\mu_i \in \mathcal{M}} \leftarrow$ the indices of the K nearest neighbors of $x_l \in \mathcal{X}$
- 6: $\{v_1, \dots, v_p\} \leftarrow Vor^{3D}(x_l)$
- 7: $\mathcal{X}_l \leftarrow \{x_l, x_{\mu_1}, \dots, x_{\mu_K}\}$
- 8: $\hat{\mathcal{X}}_l \leftarrow \{x_l, x_{\mu_1}, \dots, x_{\mu_K}, v_1, \dots, v_p\}$
- 9: **if** $TopoDiff(\mathcal{X}_l, \hat{\mathcal{X}}_l) < \tau$ **then**
- 10: $I \leftarrow I \cup \{\mu_1, \dots, \mu_K\}$
- 11: $\hat{\mathcal{X}} \leftarrow \hat{\mathcal{X}} \cup \{v_1, \dots, v_p\}$
- 12: **continue**
- 13: **end if**
- 14: $\{\lambda_1, \dots, \lambda_{K'}\}_{\lambda_i \in \mathcal{M}} \leftarrow$ the indices of the K' nearest neighbors of $x_l \in \mathcal{X}$
- 15: $\{x'_0, x'_1, \dots, x'_{K'}\}, U \leftarrow$ Dimensionality Reduction Using $PCA_{3 \rightarrow 2}(\{x_l, x_{\lambda_1}, \dots, x_{\lambda_{K'}}\})$
- 16: $\{v'_1, \dots, v'_{p'}\} \leftarrow Vor^{2D}(x'_0)$
- 17: $\{\hat{x}_1, \dots, \hat{x}_{p'}\} \leftarrow \{\hat{x}_i = Uv'_i + x_l\}_{i=1, \dots, p'}$
- 18: $I \leftarrow I \cup \{\lambda_1, \dots, \lambda_{K'}\}$
- 19: $\hat{\mathcal{X}} \leftarrow \hat{\mathcal{X}} \cup \{\hat{x}_1, \dots, \hat{x}_{p'}\}$
- 20: **end if**
- 21: **end for**

The Wasserstein distance is a measure of the difference between two probability distributions, capturing both the magnitude and the geometry of the changes in topological features. In this paper, we adopt the Wasserstein distance as

defined in (Maria et al. 2014; Kerber, Morozov, and Nigmatov 2017). If the topological difference remains below a predefined threshold τ , the vertices are added to the interpolated set. Conversely, a significant topological change triggers a switch to 2D interpolation.

2D Voronoi Interpolation In cases requiring 2D interpolation, the algorithm selects a smaller set of neighbors, applies Principal Component Analysis (PCA) to project these points onto a plane, and performs a 2D Voronoi tessellation. The vertices from this tessellation are then mapped back to the 3D space using the principal components (step 17 in Algorithm 1), ensuring the preservation of local geometric structures. Here, the notation $PCA_{3 \rightarrow 2}$ refers to the process of reducing dimensionality from 3D to 2D using PCA (Yamada and Shibuya 2020), yielding the reduced point set and the first two principal vectors, denoted by $U = [u_1, u_2]$.

We utilize 3D-to-2D Voronoi Interpolation for several reasons. First, manifold learning fundamentally assumes high-dimensional data lie on a low-dimensional manifold (Fefferman, Mitter, and Narayanan 2016), with the minimal neighborhood of any non-isolated point approximated by a tangent hyperplane (or tangent plane for 3D point clouds). Second, point clouds from SfM are often sparse on planar surfaces (Ververas et al. 2025), making 2D embedding better suited for interpolating low-curvature regions. Third, embedding in low-dimensional manifolds is computationally more efficient for interpolation than in high dimensions (Melodia and Lenz 2020).

Persistent Homology Loss Function

As outlined in the introduction, NVS evaluation metrics include PSNR, SSIM, and LPIPS. Unlike PSNR and SSIM, LPIPS evaluates perceptual similarity by comparing semantic feature maps extracted from a pretrained vision network. It measures the feature-level structural integrity of the scene, where smaller distances indicate higher rendering quality.

To maintain efficiency, we avoid using deeper networks in 3D-GS, as they would increase memory usage and reduce rendering speed. Instead, we propose a persistent homology-based loss function, PersLoss, as an additional term in the total loss to guide the Gaussians in learning abstract structural and semantic features. The computation of our loss function involves the following steps:

2D Image and 3D RGB Space After the rasterization step of 3D-GS, we obtain the rendered image \mathbf{y} , along with the ground truth image $\hat{\mathbf{y}}$. Given that the true ground truth for 3D point clouds (or 3D Gaussians) is unknown, we do not utilize them for subsequent persistent homology calculation. Instead, we rely on the rendered image obtained from differentiable rasterization, which has a known ground truth, for loss calculation. The gradient information from the 2D image propagates back through the differentiable rasterization of 3D-GS, optimizing the training parameters.

To facilitate persistent homology calculations, we reshape both images:

$$\mathbf{y}' = S(\mathbf{y}), \quad \hat{\mathbf{y}}' = S(\hat{\mathbf{y}}), \quad (5)$$

where S reshapes both \mathbf{y} and $\hat{\mathbf{y}}$ to the size of $HW \times 3$.

This transformation treats the 2D image data as data in a 3D RGB space, allowing us to perform persistent homology calculations in this space.

Filtering and Truncated Persistent Barcode We utilize the alpha complex to filter \mathbf{y}' and $\hat{\mathbf{y}}'$ for PH computation, obtaining topological features expressed in the form of PD and PB. For each feature, its birth and death values are denoted as b^k and d^k , respectively. The lifespan of a feature in the rendered image is represented as (b^k, d^k) , while the lifespan of a feature in the ground truth image is represented as (\hat{b}^k, \hat{d}^k) . In the context of PB, longer bars indicate features that persist over a wide range of alpha values, signifying their importance, while shorter bars often correspond to noise. Consequently, we set three parameters, k_0 , k_1 , and k_2 , to select the top- k longest lifespans of topological features for 0-dimensional, 1-dimensional, and 2-dimensional homology, respectively. We then obtain the truncated PBs for both images, denoted as pb and \hat{pb} :

$$pb = \left\{ \begin{array}{l} (b_1^0, d_1^0), \dots, (b_{k_0}^0, d_{k_0}^0) \\ (b_1^1, d_1^1), \dots, (b_{k_1}^1, d_{k_1}^1) \\ (b_1^2, d_1^2), \dots, (b_{k_2}^2, d_{k_2}^2) \end{array} \right\}, \quad (6)$$

$$\hat{pb} = \left\{ \begin{array}{l} (\hat{b}_1^0, \hat{d}_1^0), \dots, (\hat{b}_{k_0}^0, \hat{d}_{k_0}^0) \\ (\hat{b}_1^1, \hat{d}_1^1), \dots, (\hat{b}_{k_1}^1, \hat{d}_{k_1}^1) \\ (\hat{b}_1^2, \hat{d}_1^2), \dots, (\hat{b}_{k_2}^2, \hat{d}_{k_2}^2) \end{array} \right\}. \quad (7)$$

Since we are only interested in high-dimensional features, we discard the topological features beyond the top- k .

PersLoss We then calculate the total difference between the birth and death values of the features in each topological dimension to serve as our regularization term:

$$\text{PersLoss} = \sum_{i=0}^2 \frac{\beta^i}{\sum_{k=0}^2 \beta^k} \sum_{j=1}^{k_i} \left(\left| b_j^i - \hat{b}_j^i \right|^2 + \left| d_j^i - \hat{d}_j^i \right|^2 \right), \quad (8)$$

where β^i represents the Betti numbers, indicating the number of i -dimensional features.

Finally, our topology-aware total loss function comprises a standard supervision loss term, L_{supv} , which combines L1 loss and SSIM loss, and a topology-focused regularization term, PersLoss. The regularization term quantifies the similarity between the predicted topology and the desired topology. As discussed in the introduction, both terms are differentiable, allowing gradients from the topology-aware total loss to propagate through the entire pipeline. This effectively constrains and guides the optimization of Gaussian parameters, thereby improving the feature-level structural integrity of the rendered output.

However, as we theoretically demonstrate in *Section B of the appendix*, the optimization of the topology-aware total loss is prone to fluctuations. Specifically, minimizing PersLoss may decrease its value, but the optimization of L_{supv} can perturb the model parameters W , causing an increase in the loss value. In our Topology-GS approach, we selectively activate PersLoss during specific optimization

Dataset Metrics Method	Year	Paper	Mip-NeRF360			Tanks & Temples			Deep Blending		
			PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
3D-GS (Kerbl et al. 2023)	2023	TOG	27.21	0.815	0.214	23.14	0.841	0.183	29.41	0.903	0.243
Scaffold-GS (Lu et al. 2024) (baseline)	2024	CVPR	28.84	0.848	0.220	23.96	0.853	0.177	<u>30.21</u>	<u>0.906</u>	0.254
Mip-Splatting (Yu et al. 2024)	2024	CVPR	27.79	0.827	0.203	23.65	0.849	0.211	29.68	0.903	0.309
2D-GS (Huang et al. 2024)	2024	SIGGRAPH	28.75	0.870	0.213	22.96	0.825	0.217	29.49	0.899	0.259
Revised-GS (Rota Bulò, Porzi, and Kotschieder 2025)	2024	ECCV	27.70	0.823	0.223	<u>24.10</u>	<u>0.857</u>	0.183	29.64	0.905	0.303
Pixel-GS (Zhang et al. 2025)	2024	ECCV	<u>29.11</u>	<u>0.872</u>	0.165	23.61	0.851	<u>0.161</u>	28.83	0.892	0.251
Topology-GS (ours)	2025	AAAI	29.50	0.874	<u>0.179</u>	24.26	0.860	0.160	30.45	0.911	<u>0.245</u>

Table 1: Rendering quality on Mip-NeRF360, Tanks & Temples, and Deep Blending.

* **Bold** indicates the best results, while underline represents the second-best results.

rounds in the ADC stage, which reduces the computational demand of PH and directly guides the adjustment of Gaussians. Nevertheless, due to these fluctuations, we cannot guarantee the convergence of PersLoss at the end of the ADC stage, making it challenging to ensure that the topological features of the final rendered image approximate those of the ground truth image.

To address this issue, we introduce the following assumptions (Zhang et al. 2022) and theorem:

- **Assumpt. 0 (A0):** $B = k_0 + k_1 + k_2$ is sufficiently small such that *Assumptions 1 and 2 in the appendix* hold.
- **Assumpt. 1 (A1):** f is 1-bounded, 1-Lipschitz continuous and 1-Lipschitz smooth with respect to W .
- **Assumpt. 2 (A2):** $L_{\text{supv}}(W)$ is ℓ^0 -bounded, ℓ^1 -Lipschitz continuous and ℓ^2 -Lipschitz smooth with respect to W .

Theorem 1. *Under Assumptions A0, A1, and A2, and a given stopping condition ϵ , the optimization algorithm using our topology aware total loss stops in $O(\frac{1}{\epsilon})$ iterations if step-size η is chosen to be:*

$$\eta \leq \min \left\{ \frac{1}{2\ell^2 + 10\lambda_{\text{topo}}B}, \frac{\epsilon}{4096\lambda_{\text{topo}}^2B^2} \right\}. \quad (9)$$

Here, λ_{topo} is the coefficient of PersLoss in the total loss, and B represents the cardinality of the ground truth PD (excluding the diagonal). *For a detailed proof, see Section C of the appendix.*

This theorem establishes that our topology-aware total loss not only converges despite fluctuations but also converges rapidly with a finite number of training iterations by adjusting hyperparameters. This property makes PersLoss suitable for use with L1 loss and SSIM loss, as well as the ADC trick, rendering it suitable for 3D-GS tasks.

Experiments

Datasets. To evaluate the proposed method, consistent with 3D-GS (Kerbl et al. 2023), we utilized 11 scenes: seven from the Mip-NeRF360 dataset (Barron et al. 2022), two from the Tanks & Temples dataset (Knapitsch et al. 2017), and two from the Deep Blending dataset (Hedman et al. 2018).

Comparison Methods. We compared the proposed method with state-of-the-art (SOTA) 3D-GS-based approaches for novel-view synthesis, including 3D-GS, Mip-Splatting (Yu et al. 2024), Scaffold-GS (Lu et al. 2024), and 2D-GS (Huang et al. 2024), as well as the recent Revised-GS

(Rota Bulò, Porzi, and Kotschieder 2025) and Pixel-GS (Zhang et al. 2025).

Metrics. The Topology-GS model was evaluated based on rendering quality, structure preservation (pixels and feature maps), and visual comparisons. Rendering quality was assessed using PSNR, SSIM, and LPIPS metrics.

For additional experimental results, see Section D of the appendix.

Overall Rendering Quality

Table 1 presents the rendering quality of various algorithms across three NVS benchmarks. Scaffold-GS serves as the baseline for our comparisons. Our proposed method, Topology-GS, consistently outperforms Scaffold-GS and nearly all other methods across the three datasets, achieving SOTA results. The introduction of our LPVI significantly enhances PSNR and SSIM metrics by improving initial geometric coverage. Additionally, incorporating our PersLoss as a topological constraint leads to a substantial reduction in LPIPS across all datasets.

Structural Preservation

As illustrated in Figure 4, we rendered depth maps from perspectives in the test set to visually assess the impact of two proposed enhancements: the LPVI interpolation method and the PersLoss regularization term.




Figure 4: Depth map comparison between a SOTA method (i.e. Scaffold-GS (Lu et al. 2024)) and our Topology-GS. Columns 1 and 3 show results from Scaffold-GS, while Columns 2 and 4 show results from Topology-GS. Depth maps are visualized using COLORMAP_JET.

Comparing Columns 1 and 2, we observe that the baseline (Scaffold-GS) exhibits discontinuous color changes in low-curvature areas on the depth map (e.g., flat ground showing noticeable depth variations). This often indicates significant

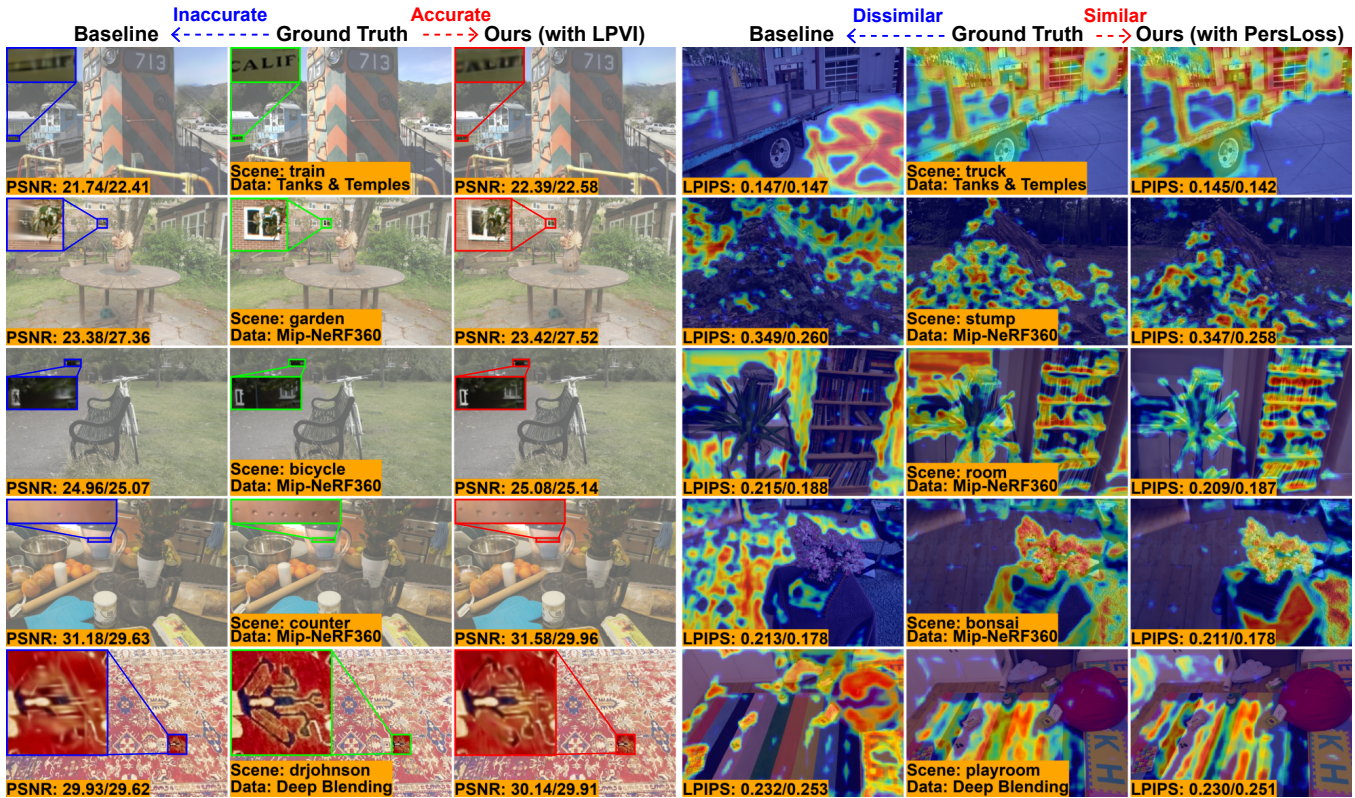


Figure 5: Comparison of Ground Truth (GT) images with those rendered by a SOTA method (i.e. Scaffold-GS (Lu et al. 2024)) and our Topology-GS. Columns 2 and 5 show GT images and feature maps; Columns 1 and 4 show Scaffold-GS results; Columns 3 and 6 show Topology-GS results. PSNR and LPIPS scores are annotated, with “/” separating image and scene scores. The rendered images in Column 3 capture low-curvature regions more accurately compared to the Scaffold-GS in Column 1, and the feature maps in Column 6 closely match the GT feature maps in Column 5.

fluctuations in depth within those regions. In contrast, our method (Topology-GS), enhanced with LPVI, reflects more consistent depth changes in these areas due to the improved pixel-level structural integrity of the initial point cloud in low-curvature regions. Consequently, the depth map’s color variations are smoother. In Columns 3 and 4, the comparison reveals that the baseline model struggles with understanding feature-level structural integrity, particularly in areas with significant structural changes. The baseline’s depth map displays sluggish color transitions, failing to adapt to sharp depth variations caused by structural elements (e.g., vague depth regions between railings and gaps). This results in some structures not being properly rendered. However, guided by PersLoss, our method accurately renders these structural changes, ensuring rapid visual transitions in the depth map.

Visual Comparison

Figure 5 presents a visual comparison of Ground Truth (GT) images with rendered images from Scaffold-GS and our proposed method, Topology-GS. In the first three columns, GT images are compared with those rendered by Scaffold-GS and Topology-GS. The interpolation introduced in Topology-GS significantly enhances pixel-level de-

tails, especially in low-curvature regions, as evidenced by improved PSNR values. The last three columns show feature map comparisons. The topology-focused regularization term in Topology-GS directs the feature maps to better align with the GT feature maps, resulting in lower LPIPS values and greater perceptual similarity to the GT.

Ablation Study

Table 2 summarizes the results of the ablation study on Topology-GS across three NVS benchmarks. To evaluate the contributions of individual components to the overall performance, we independently removed the LPVI interpolation method and the PersLoss regularization term. Furthermore, for comparison, we replaced LPVI with alternative point cloud densification strategies, including random interpolation (denoted as Random), DEBLUR (Lee et al. 2025), which densifies the point cloud by uniformly sampling within its bounding box, and SAGS (Ververas et al. 2025), which constructs a graph over low-curvature points and augments the data using their midpoints.

The results indicate that removing the LPVI interpolation method results in a decline in PSNR and SSIM, along with a slight increase in LPIPS, highlighting the importance of LPVI in enhancing pixel-level rendering quality. Simi-

Dataset Metrics Ablation	Mip-NeRF360				Tanks & Temples				Deep Blending			
	PSNR	SSIM	LPIPS	Avg Memory	PSNR	SSIM	LPIPS	Avg Memory	PSNR	SSIM	LPIPS	Avg Memory
Baseline (Lu et al. 2024)	28.84	0.848	0.220	159.98MB	23.96	0.853	0.177	76.48MB	30.21	0.906	0.254	54.29MB
Baseline+Random	29.15	0.859	0.206	168.35MB	24.13	0.857	0.176	109.71MB	30.10	0.907	0.258	60.43MB
Baseline+DEBLUR (Lee et al. 2025)	29.16	0.861	0.203	150.74MB	24.23	0.861	0.174	100.06MB	30.29	0.909	0.254	55.60MB
Baseline+SAGS (Ververas et al. 2025)	29.16	0.863	0.198	147.17MB	24.39	0.860	0.170	106.51MB	30.38	0.910	0.251	59.01MB
Baseline+LPVI (ours)	<u>29.47</u>	<u>0.873</u>	0.191	173.97MB	<u>24.22</u>	<u>0.858</u>	0.175	89.99MB	<u>30.42</u>	<u>0.911</u>	0.252	64.71MB
Baseline+PersLoss (ours)	29.36	0.870	<u>0.185</u>	160.87MB	24.08	0.853	<u>0.168</u>	76.70MB	30.30	0.909	<u>0.249</u>	54.62MB
Topology-GS (ours)	29.50	0.874	0.179	<u>175.36MB</u>	24.26	0.860	0.160	<u>90.66MB</u>	30.45	0.911	0.245	<u>65.22MB</u>

Table 2: Ablation study on Mip-NeRF360, Tanks & Temples, and Deep Blending.

* Underline highlights the numbers included in the comparison.

larly, replacing LPVI with alternative interpolation densification methods produces inferior outcomes. Specifically, Random and DEBLUR disrupt the topological structure, whereas SAGS proves effective only for interpolating within low-curvature regions, failing to account for region boundaries. Furthermore, excluding PersLoss leads to a significant increase in LPIPS, reflecting a diminished visual perceptual similarity between the rendered images and the ground truth. These findings underscore the critical role of PersLoss in reducing LPIPS and enhancing the perceptual quality of the

1

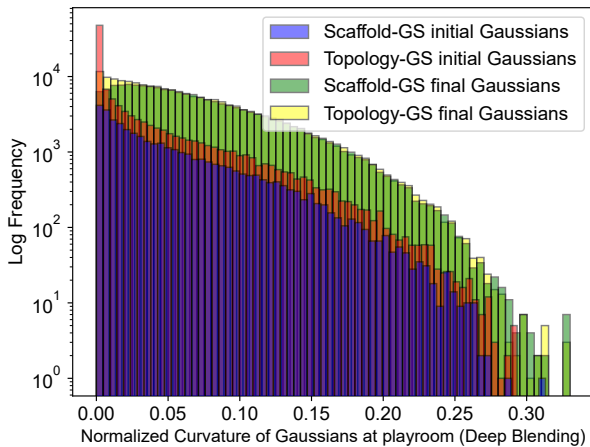


Figure 6: Histogram of Gaussians. Blue and red: the initial point counts for Scaffold-GS (baseline) and Topology-GS, respectively. Green and yellow: the corresponding point counts after training.

Computational Efficiency

In this section, we discuss the computational overhead of our Topology-GS. As mentioned in the introduction, Topology-GS enhances initial sparse point clouds using our LPVI method and incorporates PersLoss as a structural constraint during training. As shown in Table 2, PersLoss acts only during training and does not significantly affect final average memory usage. Therefore, we will primarily discuss the impact of LPVI on final memory usage.

As shown in Table 2, while LPVI slightly increases average memory usage, the impact is minimal and remains within an acceptable range for practical applications, with

increases of less than 20MB across three benchmarks. Figure 6 illustrates point counts for Topology-GS (red and yellow) compared to the baseline (blue and green). The red histogram shows increased points in low-curvature areas due to LPVI interpolation. Post-training, the yellow histogram aligns with the baseline in non-low-curvature areas, indicating no disruption elsewhere. Overall, the total points in the yellow histogram remain similar to the green, avoiding additional computational overhead.

Conclusion

This paper introduced Topology-Aware 3D Gaussian Splatting (Topology-GS), leveraging persistent homology to enhance structural integrity in novel-view synthesis (NVS). By addressing incomplete initial geometric coverage and insufficient topological constraints during training, our approach improves both pixel-level and feature-level integrity. Extensive experiments demonstrate the effectiveness of Topology-GS. The proposed Topology-GS has the potential to significantly impact the field of NVS by providing a more accurate and efficient method for rendering complex scenes. Its integration of topological methods with 3D Gaussian Splatting could lead to more realistic and detailed renderings, and inspire new approaches to other computer vision tasks.

Ethics Statement

This research did not involve human participants, animals, or the use of sensitive data; therefore, ethical approval was not required.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (No. 2023YFC2907600), the Key Science and Technology Innovation Project of CCTEG (No. 2024-TD-ZD016-01), and the City University of Hong Kong (# 9382001). We extend our gratitude to the participants for their time and interest, which has greatly contributed to the research. We are also thankful to Dr. Ali Zia, Postdoctoral Research Fellow at the Australian National University, for his pivotal introduction to topological machine learning, shaping the early stages of our project. Special thanks to Lu Huang, a PhD candidate at Peking University, whose expertise in creating the visual content for this paper was invaluable.

References

- Barron, J. T.; Mildenhall, B.; Tancik, M.; Hedman, P.; Martin-Brualla, R.; and Srinivasan, P. P. 2021. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5855–5864.
- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2022. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5470–5479.
- Chazal, F.; and Michel, B. 2021. An introduction to topological data analysis: fundamental and practical aspects for data scientists. *Frontiers in artificial intelligence*, 4: 108.
- Clough, J. R.; Byrne, N.; Oksuz, I.; Zimmer, V. A.; Schnabel, J. A.; and King, A. P. 2020. A topological loss function for deep-learning based image segmentation using persistent homology. *IEEE transactions on pattern analysis and machine intelligence*, 44(12): 8766–8778.
- de Surrel, T.; Hensel, F.; Carrière, M.; Lacombe, T.; Ike, Y.; Kurihara, H.; Glisse, M.; and Chazal, F. 2022. RipsNet: a general architecture for fast and robust estimation of the persistent homology of point clouds. In *Topological, Algebraic and Geometric Learning Workshops 2022*, 96–106. PMLR.
- DiFrancesco, P.-M.; Bonneau, D.; and Hutchinson, D. J. 2020. The implications of M3C2 projection diameter on 3D semi-automated rockfall extraction from sequential terrestrial laser scanning point clouds. *Remote Sensing*, 12(11): 1885.
- Fefferman, C.; Mitter, S.; and Narayanan, H. 2016. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4): 983–1049.
- Foroutan, Y.; Rebain, D.; Yi, K. M.; and Tagliasacchi, A. 2024. Does Gaussian Splatting need SFM Initialization? *arXiv preprint arXiv:2404.12547*.
- Fridovich-Keil, S.; Yu, A.; Tancik, M.; Chen, Q.; Recht, B.; and Kanazawa, A. 2022. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5501–5510.
- Gabrielsson, R. B.; Nelson, B. J.; Dwarknath, A.; and Skraba, P. 2020. A topology layer for machine learning. In *International Conference on Artificial Intelligence and Statistics*, 1553–1563. PMLR.
- Guédon, A.; and Lepetit, V. 2024. 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*, 5354–5363.
- Hedman, P.; Philip, J.; Price, T.; Frahm, J.-M.; Drettakis, G.; and Brostow, G. 2018. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (ToG)*, 37(6): 1–15.
- Hu, J.; Fei, B.; Xu, B.; Hou, F.; Yang, W.; Wang, S.; Lei, N.; Qian, C.; and He, Y. 2024. Topology-Aware Latent Diffusion for 3D Shape Generation. *arXiv preprint arXiv:2401.17603*.
- Hu, X.; Li, F.; Samaras, D.; and Chen, C. 2019. Topology-preserving deep image segmentation. *Advances in neural information processing systems*, 32.
- Huang, B.; Yu, Z.; Chen, A.; Geiger, A.; and Gao, S. 2024. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, 1–11.
- Huber, S. 2021. Persistent homology in data science. In *Data Science–Analytics and Applications: Proceedings of the 3rd International Data Science Conference–iDSC2020*, 81–88. Springer.
- Jignasu, A.; Balu, A.; Sarkar, S.; Hegde, C.; Ganapathysubramanian, B.; and Krishnamurthy, A. 2024. SDFConnect: Neural Implicit Surface Reconstruction of a Sparse Point Cloud with Topological Constraints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5271–5279.
- Kerber, M.; Morozov, D.; and Nigmatov, A. 2017. Geometry helps to compare persistence diagrams.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4).
- Knapitsch, A.; Park, J.; Zhou, Q.-Y.; and Koltun, V. 2017. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4): 1–13.
- Lee, B.; Lee, H.; Sun, X.; Ali, U.; and Park, E. 2025. Deblurring 3d gaussian splatting. In *European Conference on Computer Vision*, 127–143. Springer.
- Li, J.; Zhang, J.; Bai, X.; Zheng, J.; Ning, X.; Zhou, J.; and Gu, L. 2024. Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20775–20785.
- Lu, T.; Yu, M.; Xu, L.; Xiangli, Y.; Wang, L.; Lin, D.; and Dai, B. 2024. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20654–20664.
- Maria, C.; Boissonnat, J.-D.; Glisse, M.; and Yvinec, M. 2014. The gudhi library: Simplicial complexes and persistent homology. In *Mathematical Software–ICMS 2014: 4th International Congress, Seoul, South Korea, August 5-9, 2014. Proceedings 4*, 167–174. Springer.
- Melodia, L.; and Lenz, R. 2020. Persistent homology as stopping-criterion for voronoi interpolation. In *International Workshop on Combinatorial Image Analysis*, 29–44. Springer.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Muhammad, M. B.; and Yeasin, M. 2020. Eigen-cam: Class activation map using principal components. In *2020 International Joint Conference on Neural Networks (IJCNN)*, 1–7. IEEE.

- Müller, T.; Evans, A.; Schied, C.; and Keller, A. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4): 1–15.
- Niemeyer, M.; Manhardt, F.; Rakotosaona, M.-J.; Oechsle, M.; Duckworth, D.; Gosula, R.; Tateno, K.; Bates, J.; Kaeser, D.; and Tombari, F. 2024. Radsplat: Radiance field-informed gaussian splatting for robust real-time rendering with 900+ fps. *arXiv preprint arXiv:2403.13806*.
- Nishikawa, N.; Ike, Y.; and Yamanishi, K. 2024. Adaptive Topological Feature via Persistent Homology: Filtration Learning for Point Clouds. *Advances in Neural Information Processing Systems*, 36.
- Preparata, F. P.; and Shamos, M. I. 2012. *Computational geometry: an introduction*. Springer Science & Business Media.
- Pun, C. S.; Lee, S. X.; and Xia, K. 2022. Persistent-homology-based machine learning: a survey and a comparative study. *Artificial Intelligence Review*, 55(7): 5169–5213.
- Qi, Y.; He, Y.; Qi, X.; Zhang, Y.; and Yang, G. 2023. Dynamic snake convolution based on topological geometric constraints for tubular structure segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6070–6079.
- Reinauer, R.; Caorsi, M.; and Berkouk, N. 2021. Persformer: A transformer architecture for topological machine learning. *arXiv preprint arXiv:2112.15210*.
- Rota Bulò, S.; Porzi, L.; and Kotschieder, P. 2025. Revisiting Densification in Gaussian Splatting. In *European Conference on Computer Vision*, 347–362. Springer.
- Schönberger, J. L.; and Frahm, J.-M. 2016. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Simonyan, K. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Song, X.; Zheng, J.; Yuan, S.; Gao, H.-a.; Zhao, J.; He, X.; Gu, W.; and Zhao, H. 2024. SA-GS: Scale-Adaptive Gaussian Splatting for Training-Free Anti-Aliasing. *arXiv preprint arXiv:2403.19615*.
- Tanemura, M.; Ogawa, T.; and Ogita, N. 1983. A new algorithm for three-dimensional Voronoi tessellation. *Journal of Computational Physics*, 51(2): 191–207.
- Turkulainen, M.; Ren, X.; Melekhov, I.; Seiskari, O.; Rahtu, E.; and Kannala, J. 2024. DN-Splatter: Depth and Normal Priors for Gaussian Splatting and Meshing. *arXiv preprint arXiv:2403.17822*.
- Ververas, E.; Potamias, R. A.; Song, J.; Deng, J.; and Zafeiriou, S. 2025. Sags: Structure-aware 3d gaussian splatting. In *European Conference on Computer Vision*, 221–238. Springer.
- Wang, F.; Liu, H.; Samaras, D.; and Chen, C. 2020. Topogan: A topology-aware generative adversarial network. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, 118–136. Springer.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wu, T.; Yuan, Y.-J.; Zhang, L.-X.; Yang, J.; Cao, Y.-P.; Yan, L.-Q.; and Gao, L. 2024. Recent advances in 3d gaussian splatting. *Computational Visual Media*, 1–30.
- Xu, L.; Xiangli, Y.; Peng, S.; Pan, X.; Zhao, N.; Theobalt, C.; Dai, B.; and Lin, D. 2023. Grid-guided neural radiance fields for large urban scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8296–8306.
- Yamada, N.; and Shibuya, T. 2020. Inferring Underlying Manifold of Low Density Data using Adaptive Interpolation. In *ICAART (2)*, 395–402.
- Yan, X.; Hu, S.; Mao, Y.; Ye, Y.; and Yu, H. 2021. Deep multi-view learning methods: A review. *Neurocomputing*, 448: 106–129.
- Yang, R.; Zhu, Z.; Jiang, Z.; Ye, B.; Chen, X.; Zhang, Y.; Chen, Y.; Zhao, J.; and Zhao, H. 2024a. Spectrally Pruned Gaussian Fields with Neural Compensation. *arXiv preprint arXiv:2405.00676*.
- Yang, Z.; Gao, X.; Sun, Y.; Huang, Y.; Lyu, X.; Zhou, W.; Jiao, S.; Qi, X.; and Jin, X. 2024b. Spec-gaussian: Anisotropic view-dependent appearance for 3d gaussian splatting. *arXiv preprint arXiv:2402.15870*.
- Yu, Z.; Chen, A.; Huang, B.; Sattler, T.; and Geiger, A. 2024. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19447–19456.
- Zhang, K.; Riegler, G.; Snavely, N.; and Koltun, V. 2020. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhang, Y.; Yao, J.; Wang, Y.; and Chen, C. 2022. On the convergence of optimizing persistent-homology-based losses. *arXiv preprint arXiv:2206.02946*.
- Zhang, Z.; Hu, W.; Lao, Y.; He, T.; and Zhao, H. 2025. Pixel-GS: Density Control with Pixel-Aware Gradient for 3D Gaussian Splatting. In *European Conference on Computer Vision*, 326–342. Springer.
- Zhu, Z.; Fan, Z.; Jiang, Y.; and Wang, Z. 2025. Fsgs: Real-time few-shot view synthesis using gaussian splatting. In *European Conference on Computer Vision*, 145–163. Springer.
- Zia, A.; Khamis, A.; Nichols, J.; Tayab, U. B.; Hayder, Z.; Rolland, V.; Stone, E.; and Petersson, L. 2024. Topological deep learning: A review of an emerging paradigm. *Artificial Intelligence Review*, 57(4): 77.
- Zou, Z.-X.; Yu, Z.; Guo, Y.-C.; Li, Y.; Liang, D.; Cao, Y.-P.; and Zhang, S.-H. 2024. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10324–10335.