

NeuralGS: Bridging Neural Fields and 3D Gaussian Splatting for Compact 3D Representations

Zhenyu Tang¹, Chaoran Feng¹, Xinhua Cheng¹, Wangbo Yu¹, Junwu Zhang¹,
Yuan Liu^{2*}, Xiao-Xiao Long², Wenping Wang³, Li Yuan^{1*},

¹School of Electronic and Computer Engineering, Peking University, China

²Hong Kong University of Science and Technology

³Texas A&M University

{zhenyutang, chaoran.feng}@stu.pku.edu.cn, yuanly@ust.hk, yuanli-ece@pku.edu.cn

Abstract

3D Gaussian Splatting (3DGS) achieves impressive quality and rendering speed, but with millions of 3D Gaussians and significant storage and transmission costs. In this paper, we aim to develop a simple yet effective method called **NeuralGS** that compresses the original 3DGS into a compact representation. Our observation is that neural fields like NeRF can represent complex 3D scenes with Multi-Layer Perceptron (MLP) neural networks using only a few megabytes. Thus, NeuralGS effectively adopts the neural field representation to encode the attributes of 3D Gaussians with MLPs, only requiring a small storage size even for a large-scale scene. To achieve this, we adopt a clustering strategy and fit the Gaussians within each cluster using different tiny MLPs, based on importance scores of Gaussians as fitting weights. We experiment on multiple datasets, achieving a $91\times$ average model size reduction without harming the visual quality.

Code — <https://github.com/PKU-YuanGroup/NeuralGS>

1 Introduction

Novel view synthesis (NVS) is a fundamental task in 3D vision, with substantial applications across fields such as virtual reality (Dai et al. 2019), augmented reality (Zhou et al. 2018; Ma et al. 2025b), and media generation (Poole et al. 2022; Ma et al. 2025a). This task aims to render photo-realistic novel-view images, given limited multi-view input images. Neural radiance field (NeRF) (Mildenhall et al. 2020) has gained significant attention as a 3D scene representation for its compact structure and exceptional capability to reconstruct large-scale scenes (Barron et al. 2021, 2022; Niemeyer et al. 2022; Tancik et al. 2023). However, a persistent challenge hindering the widespread usage of NeRF lies in the computational bottlenecks imposed by volumetric rendering (Drebin, Carpenter, and Hanrahan 1988), which limit the utilization in real scenes that require fast renderings.

3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) has emerged as an alternative representation, utilizing a efficient point-based representation associated with several

explicit attributes. Unlike the slow volume rendering of NeRFs, 3DGS utilizes a fast differentiable splatting technique, achieving exceptionally fast rendering speeds and promising image quality. However, employing point-based representations inherently leads to substantial storage demands, as millions of points and their attributes are stored independently, which significantly hinders the compactness of 3DGS as a practical 3D representation.

To address the above size issue, some 3DGS compression methods (Fan et al. 2023; Xie et al. 2024; Niedermayr et al. 2024) mainly adopt pruning and quantization on Gaussian. Another noticeable direction of works (Liu et al. 2024; Chen et al. 2024), achieves impressive compression ratio based on Scaffold-GS (Lu et al. 2024) which adopts anchors to predict local Gaussians by view-dependent neural networks. These methods typically require per-view attribute prediction by neural networks for rendering and longer rendering time compared to original 3DGS, which makes them less suitable for applications demanding high-speed rendering.

In this paper, we explore how to achieve high compression ratio and maintain rendering efficiency of original 3DGS. Our method is based on the observation that neural fields like NeRF are able to represent complex scenes with small sizes. Thus, rather than proposing complex quantization like previous works, our target is to combine the neural fields and point-based representation for original 3DGS compression.

Adopting neural fields in compression is not trivial. A straightforward solution is to directly employ a multi-layer perceptron (MLP) to map the positions of Gaussians to their attributes, which could represent all attributes with a compact neural field. However, only fitting a single MLP to represent all Gaussian attributes leads to large fitting errors, severely degenerating the rendering quality, because the Gaussians show strong spatial variations. Even nearby 3D Gaussians have totally different attributes, resulting in a significant difficulty in fitting with a single MLP.

To address the aforementioned issues, we propose **NeuralGS**, a novel framework designed for the post-training compression of original 3DGS. We adopt three strategies to facilitate the effective encoding of 3D Gaussian attributes with neural fields as follows:

First, instead of fitting all Gaussians equally, we compute

*Corresponding author.

the importance of each Gaussian according to their contributions to the renderings. Gaussians with low importance are first pruned to reduce the Gaussian numbers. More importantly, we introduce a novel use of these importance scores as weighting factors in the fitting process, which ensures that important Gaussians are fitted with high accuracy.

Second, an important observation from us is that the attributes of Gaussians do not change smoothly with their positions. For example, a Gaussian with a small scale factor could have a neighboring Gaussian with an extremely large scale factor, which prevents the neural fields from accurately fitting them due to the smoothness nature of neural fields. To reduce attribute variability among Gaussians, we cluster 3D Gaussians based on their attributes to preserve similarity among Gaussians within the same cluster. For different clusters, we use different tiny neural fields (MLPs) to map the positions of Gaussians to the remaining attributes, which significantly reduces the fitting errors.

Third, we further fine-tune the learned NeuralGS representation with training images and propose a frequency loss to improve the reconstruction quality. We find that the MLPs often have difficulty in learning the high-frequency signals of Gaussian attributes during fitting. Thus, we incorporate a frequency loss, that puts emphasis on the high-frequency details of renderings, along with the original rendering loss in the fine-tuning process to recover fine details.

In the end, our NeuralGS only needs to store the positions of important Gaussians and the weights of the corresponding tiny MLPs for all clusters, substantially reducing the storage compared to the original 3DGS. NeuralGS achieves about **87×** and **117×** model size reduction compared to 3DGS on Mip-NeRF360 dataset (Barron et al. 2022) and Deep Blending dataset (Hedman et al. 2018), while delivering superior rendering quality than existing compression works. Meanwhile, NeuralGS achieves an average 1.9× rendering speed than the state-of-the-art compression work HAC (Chen et al. 2024), while maintaining comparable performance.

2 Related Works

2.1 Novel View Synthesis

Neural radiance field (NeRF) (Mildenhall et al. 2020) proposes to use MLPs to represent a scene, and this compact representation has brought view synthesis quality to a new stage. However, NeRF-based methods (Müller et al. 2022; Barron et al. 2022; Niemeyer et al. 2022; Reiser et al. 2021; Govindarajan et al. 2024; Lee et al. 2024b; Hu et al. 2023) struggle to achieve real-time rendering speed in large-scale scenes, limiting their practical use. The idea of utilizing multiple MLPs is also explored by KiloNeRF (Reiser et al. 2021) for efficient rendering. Recently, 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) and its variants (Yu et al. 2024; Lu et al. 2024; Liang et al. 2024; Liu and Banerjee 2024; Sun et al. 2024; Zhang et al. 2024; Cao et al. 2024; Zhan et al. 2025; Salman Ali et al. 2025), offer state-of-the-art scene reconstruction by utilizing a set of optimized 3D Gaussians that can be rendered efficiently.

2.2 Compression of 3D Gaussian Splatting

Although 3DGS achieves superior performance and higher rendering speed compared to NeRF-based methods, it typically requires hundreds of megabytes to store 3D Gaussian attributes, posing challenges for its practical application in large-scale scenes. Several existing works (Fan et al. 2023; Xie et al. 2024; Lee et al. 2024a; Niedermayr et al. 2024; Fan et al. 2024; Ali, Bae, and Tartaglione 2024) have made initial attempts to compress 3DGS models, primarily using pruning to reduce the number of 3D Gaussians, vector quantization to discretize Gaussian attributes into shared codebooks, and context-aware entropy encoding. Specifically, CompressGS (Niedermayr et al. 2024) utilizes vector quantization to discretize Gaussian parameters into the codebooks through clustering and entropy coding to minimize statistical redundancies in the codebooks. LightGS (Fan et al. 2023) reduces the number of Gaussians through pruning and lower Spherical Harmonics degree through distillation. Compact3DGS (Lee et al. 2024a) employs a hash grid to encode view-adaptive colors and vector quantization for geometric attributes, achieving a remarkable compression ratio for color but only 6× for geometric attributes. Moreover, SOG (Morgenstern et al. 2023) explore a image-codec-based compression to reduce the model size. Another line of work is anchor-based compression built upon ScaffoldGS (Lu et al. 2024) which adopts anchor to predict local Gaussians with attributes changing based on view directions. CompGS (Liu et al. 2024) shares a similar anchor-based structure with entropy optimization and HAC (Chen et al. 2024) uses a hash grid to further compress anchors. These works achieve high rendering quality and compression ratio but require slow per-view processing for rendering. In contrast, our method employ compact neural fields to encode all Gaussian attributes within each cluster with tiny MLPs and integrates seamlessly the original 3DGS rendering pipeline with high rendering speed compared to anchor-based works.

3 Method

3.1 Overview

General idea. The general idea of **NeuralGS** is to adopt compact neural fields to compress original 3DGS. Specifically, given 3D Gaussians reconstructed from multi-view images, we learn MLP networks to map the 3D positions of Gaussians to their attributes including opacities, spherical harmonic coefficients, scales, and rotations. These MLP networks can be regarded as a set of neural attribute fields. In this case, we only need to store the 3D positions of all Gaussians and the MLP parameters, which are highly compact thanks to the compactness of neural field representations. When we need to render from NeuralGS, the attributes only need to be decoded once from the corresponding MLPs.

Challenges and solutions. However, fitting the neural field is not a trivial task because naively fitting a compact MLP network on all Gaussian attributes leads to severe fitting errors and inferior rendering quality. To improve the fitting process, we propose three essential strategies, as shown in Fig. 1. First, not all Gaussians contribute equally to the final renderings and some of them are entirely redundant without

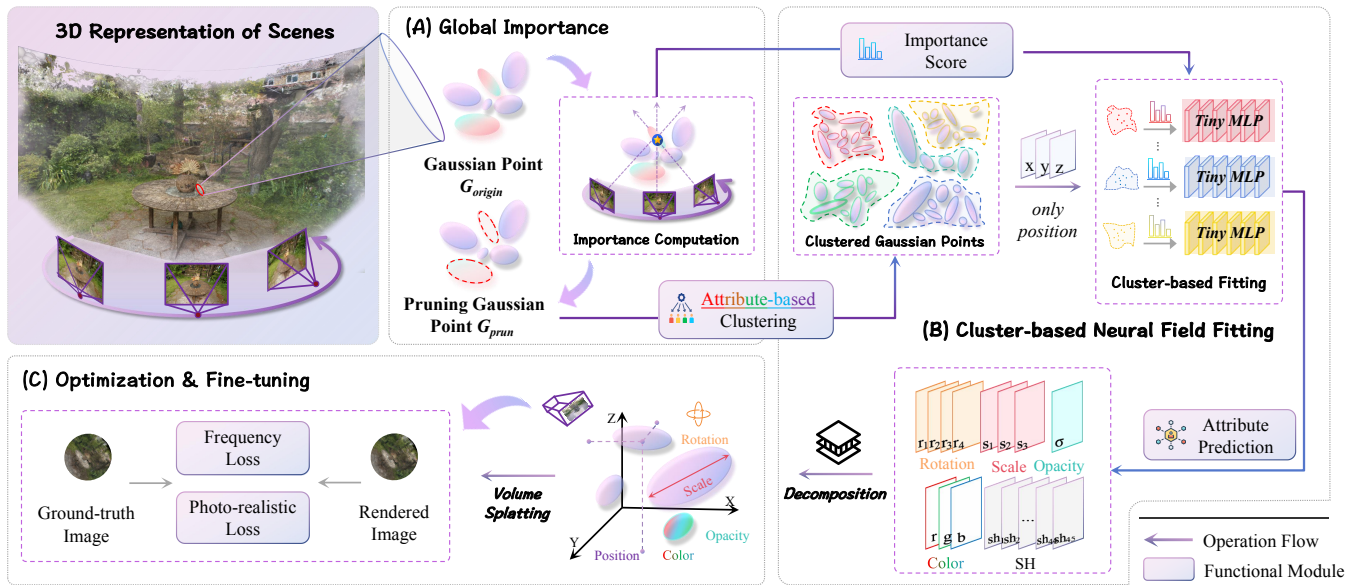


Figure 1: Overview of NeuralGS. (A) In Sec. 3.2, for each Gaussian GS_j in the scene, we first calculate its global importance score S_j (Eq.1) and prune unimportant Gaussians. (B) In Sec. 3.3, we cluster the retained Gaussians and use different tiny MLPs to map the positions to Gaussian attributes of different clusters with the loss (Eq.3) using the importance score as weights. (C) In Sec. 3.4, we fine-tune the tiny MLPs of all clusters with photorealistic loss (Eq.4) and frequency loss (Eq.5) to restore quality.

any effects on the rendering quality. Thus, this motivates us to compute the importance of all 3D Gaussians in Sec. 3.2, which is used in pruning and in a novel importance-aware fitting process. Second, we find that the Gaussian attributes do not distribute evenly or smoothly in the 3D space, where a small-scale Gaussian could have a large-scale neighbor. This uneven distribution severely hinders the fitting process because the MLPs naturally fit into smooth fields but have difficulty handling abrupt changes. Thus, in Sec. 3.3, we propose to first cluster Gaussians based on attributes and then fit a neural field for each cluster instead of solely using a single neural field for each Gaussian. Third, in Sec. 3.4, we further improve the rendering quality of NeuralGS by fine-tuning on input images with photorealistic loss and frequency loss.

3.2 Global Importance

Importance Computation. Each Gaussian contributes differently to the final renderings in 3DGS (Kerbl et al. 2023). To quantify this, we define a global importance score for each Gaussian, representing its contribution to the rendering result. Inspired by (Fan et al. 2023; Xie et al. 2024), we can calculate the importance based on each Gaussian’s contribution to every pixel p_i across all training views. We use the criterion $1(GS_j, p_i)$ to determine whether a Gaussian GS_j overlaps with pixel p_i after projection onto the 2D plane. At last, we can iterate over all training pixels and sum up the accumulated opacity of GS_j , denoted as $\alpha_k \prod_{l=1}^{k-1} (1 - \alpha_l)$, to compute each Gaussian’s contribution to the rendering result. Here, k is the index of the Gaussian GS_j in the depth ordering for pixel p_i and α is the opacity. This importance score can be further refined by incorporating the 3D Gaus-

sian’s normalized volume V_{norm} . Finally, the global importance score can be expressed as:

$$S_j = \sum_{i=1}^{MHW} 1(GS_j, p_i) \cdot (V_{\text{norm}})^\beta \cdot \alpha_k \prod_{l=1}^{k-1} (1 - \alpha_l), \quad (1)$$

$$V_{\text{norm}} = \min \left(\max \left(\frac{V}{V_{\text{max90}}}, 0 \right), 1 \right). \quad (2)$$

Here, S , M , H , and W represent the importance score, the number of training views, the image height, and the image width, respectively. V_{max90} denotes the 90% largest volume of all sorted Gaussians, and β is the hyperparameter to enhance the score’s flexibility.

Importance-based Pruning and Weighting. Thus, we rank each Gaussian based on its importance score, allowing us to prune out redundant Gaussians, thereby reducing the total number of Gaussians. Beyond pruning, we propose a novel use of the importance scores as weighting factors in the subsequent fitting process, which ensure that important Gaussians are fitted with higher accuracy.

3.3 Cluster-based Neural Field Fitting

The original 3DGS does not ensure any attribute similarities between neighboring Gaussians. Two neighboring Gaussians could have totally different colors or scales, which poses challenges in the neural field fitting. To address this, we propose an attribute-based clustering strategy to ensure attribute similarity within the same cluster and fit separate neural fields for different clusters as shown in Fig. 2.

Attribute-based Clustering. Specifically, we employ K-means (Likas et al. 2003) to cluster 3D Gaussians into K

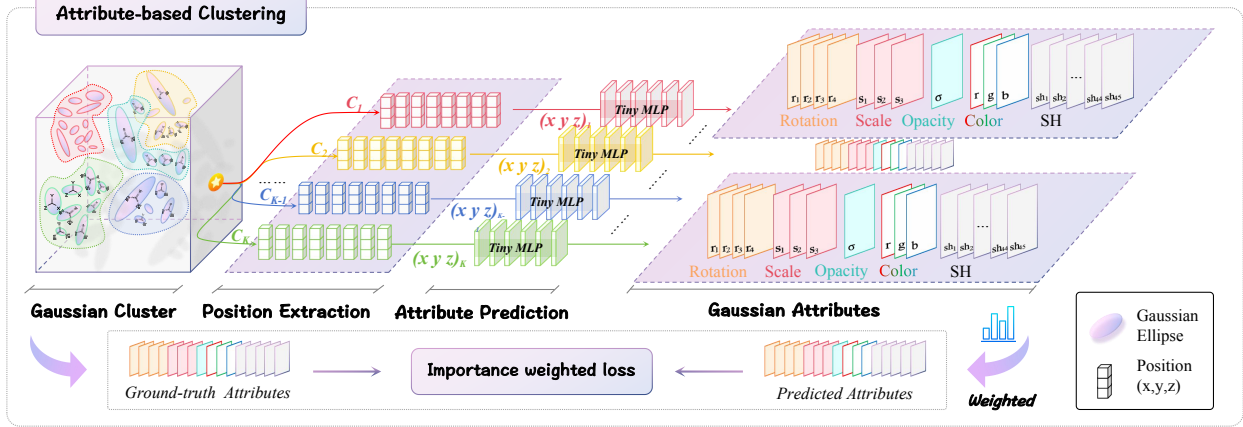


Figure 2: Details of Cluster-based Neural Field Fitting. The positions of the 3D Gaussians within each cluster are fed into the corresponding tiny MLP to fit the attributes with the importance-weighted loss. During rendering, the predicted outputs are then split into the respective attributes of the Gaussians, i.e., rotation, scale, opacity, color, and SH coefficients.

clusters, denoted as C_1, C_2, \dots, C_K . In this case, the attributes of Gaussians in the same cluster will be similar and easy to fit for neural fields. Given the significant distributional differences across attributes, we first normalize each attribute to the range $[-1, 1]$ by computing its maximum and minimum values to unify the scales of different attributes and avoid over-reliance on certain attributes for clustering.

Neural Fields. After assigning each 3D Gaussian to a cluster, we use different tiny MLPs for different clusters to map Gaussian positions within each cluster to the normalized attributes of these Gaussians. Each tiny MLP consists of five layers with positional encoding, followed by a tanh activation function (Fan 2000). The fitting processes for different clusters are conducted in parallel for efficiency.

Importance-Weighted Fitting Loss. We apply mean squared error (MSE) loss when fitting Gaussian attributes. Considering that each Gaussian contributes differently to the renderings, we further propose a novel use of importance scores as per-Gaussian fitting weight, which ensures that Gaussians with higher importance are fitted more accurately. Our loss function is defined as:

$$Loss = \frac{1}{\sum_{j \in \mathcal{P}} S_j} \sum_{j \in \mathcal{P}} S_j \cdot \|\mathcal{F}(x_j) - \hat{y}_j\|_2. \quad (3)$$

Here, \mathcal{P} represents the Gaussian index set of a cluster, S denotes the importance score, $\mathcal{F}(\cdot)$ is the tiny MLP corresponding to the cluster, x is the spatial position of the Gaussian, and \hat{y} is the normalized Gaussian attributes.

3.4 Fine-tuning

After fitting, there still remain some residuals which degrade rendering quality. To address this, we incorporate a fine-tuning stage to restore the image quality. In this process, we

fix spatial positions of the 3D Gaussians and only fine-tune the tiny MLPs corresponding to each cluster. The photorealistic loss $\mathcal{L}_{\text{render}}$, is then computed by combining the mean absolute error (MAE) loss \mathcal{L}_1 and the SSIM loss $\mathcal{L}_{\text{SSIM}}$ with the weight λ as follows:

$$\mathcal{L}_{\text{render}} = (1 - \lambda)\mathcal{L}_1 + \lambda\mathcal{L}_{\text{SSIM}}. \quad (4)$$

Frequency Loss. We observe that the fitted attributes often lose high-frequency details, such as dense grass. Thus, we introduce a frequency loss to emphasize these high-frequency details for fine-tuning. Specifically, we use a fourier transform to convert the rendered image I and the ground truth I_{gt} into frequency representations F and F_{gt} . $F(u, v)$ consists of amplitude $|F(u, v)|$ and phase $\angle F(u, v)$, where (u, v) denotes the coordinates in the frequency spectrum. We then introduce a high-pass filter with fixed bandwidth to extract high-frequency information, denoted as $\hat{F}(u, v)$ and $\hat{F}_{\text{gt}}(u, v)$. We define $\Delta|\hat{F}(u, v)| = |\hat{F}(u, v)| - |\hat{F}_{\text{gt}}(u, v)|$ and $\Delta\angle\hat{F}(u, v) = \angle\hat{F}(u, v) - \angle\hat{F}_{\text{gt}}(u, v)$. Thus, the frequency loss $\mathcal{L}_{\text{freq}}$ and the total loss $\mathcal{L}_{\text{Total}}$ can be formulated as follows:

$$\mathcal{L}_{\text{freq}} = \sum_{u=0}^{H-1} \sum_{v=0}^{W-1} \left| \Delta|\hat{F}(u, v)| \right| + \left| \Delta\angle\hat{F}(u, v) \right|, \quad (5)$$

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{render}} + \lambda_{\text{freq}}\mathcal{L}_{\text{freq}}. \quad (6)$$

Here, H , W and λ_{freq} denote the image height, width, and the hyperparameter to balance the loss.

Model Parameters. In the end, we only need to store the positions of the retained 3D Gaussians and the fine-tuned MLP weights for each cluster, significantly reducing the model size. Min-max values for normalization are shared across all clusters with only negligible 116 floating numbers needed.

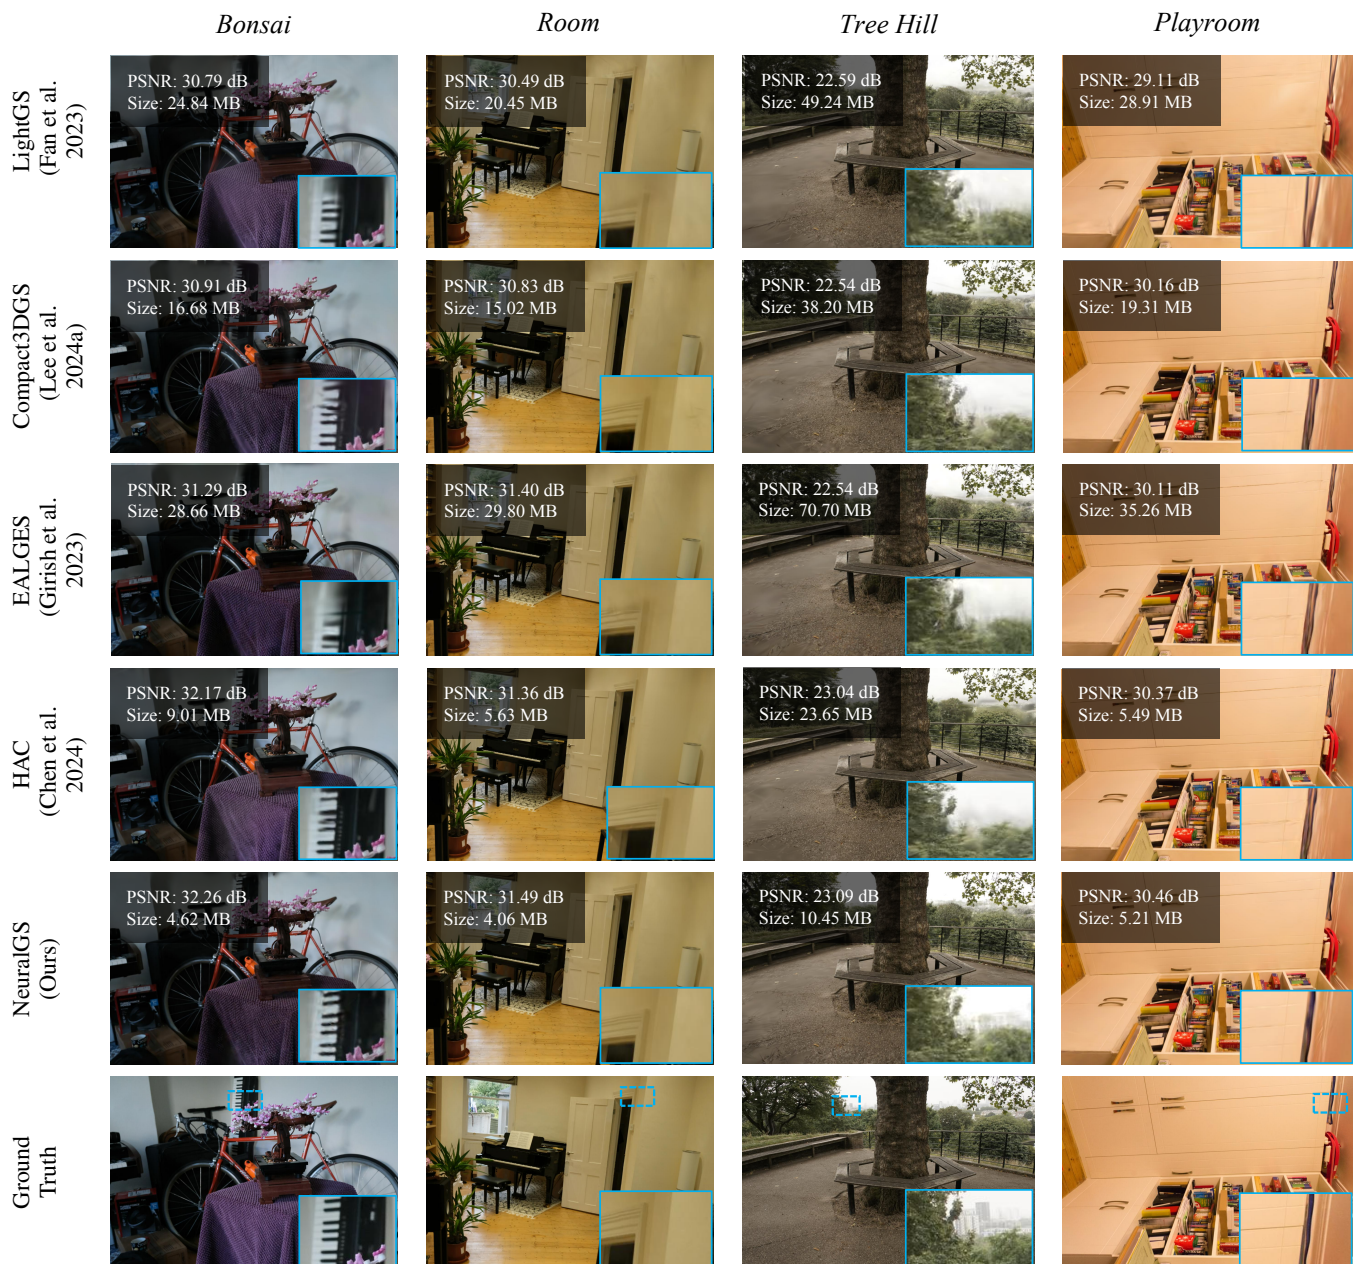


Figure 3: Qualitative results of the proposed method compared to existing compression methods.

4 Experiments

4.1 Experimental Settings

Evaluation Datasets and Metrics. We adopt three datasets for comparison. (1) *Mip-NeRF360* (Barron et al. 2022) offers scene-scale data for view synthesis, containing nine real-world large-scale scenes: five unbounded outdoor scenes and four indoor scenes with complex backgrounds. (2) *Tank and Temple* (Knapitsch et al. 2017) is a unbounded dataset that includes two scenes: *train* and *truck*. (3) *Deep Blending* (Hedman et al. 2018) contains two indoor scenes: *drjohnson* and *playroom*. For all datasets, we maintain the

same train-test splits as the official setting of 3DGS (Kerbl et al. 2023) and utilize PSNR, SSIM, LPIPS, and model size to evaluate image quality and compression ratio.

Baselines. We use 3DGS (Kerbl et al. 2023) as our compression baseline and compare with other original 3DGS-based compression techniques (Girish et al. 2023; Morgenstern et al. 2023; Lee et al. 2024a; Fan et al. 2023; Niedermayr et al. 2024; Xie et al. 2024) along with anchor-based compression works (Liu et al. 2024; Chen et al. 2024). For a fair comparison of rendering quality and model size, we use the official code of each method with the default configura-

Dataset Method	Mip-NeRF 360 (Barron et al. 2022)				Tanks&Temples (Knapitsch et al. 2017)				Deep Blending (Hedman et al. 2018)			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Size \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Size \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Size \downarrow
Mip-NeRF 360 (Barron et al. 2022)	27.69	0.795	0.238	8.5 MB	22.16	0.757	0.261	9.0 MB	29.01	0.895	0.255	8.6 MB
3DGS (Kerbl et al. 2023)	27.51	0.813	0.222	754.6 MB	23.75	0.844	0.178	438.9 MB	29.42	0.900	0.247	672.8 MB
CompressGS (Niedermayr et al. 2024)	26.98	0.801	0.242	28.72 MB	23.32	0.835	0.194	17.73 MB	29.40	0.899	<u>0.252</u>	25.96 MB
Compact3DGS (Lee et al. 2024a)	27.01	0.797	0.248	48.80 MB	23.29	0.829	0.202	39.43 MB	29.71	0.900	0.257	43.21 MB
SOG (Morgenstern et al. 2023)	27.02	0.799	0.232	42.33 MB	23.54	0.833	0.188	19.70 MB	29.21	0.891	0.271	19.28 MB
MesonGS (Xie et al. 2024)	27.08	0.800	0.245	27.51 MB	23.31	0.836	0.195	17.47 MB	29.40	0.903	0.254	25.64 MB
EAGLES (Girish et al. 2023)	27.13	0.809	0.241	60.86 MB	23.27	0.839	0.201	31.05 MB	29.72	0.907	0.250	54.65 MB
LightGS (Fan et al. 2023)	26.95	0.800	0.243	48.71 MB	23.11	0.817	0.231	24.74 MB	29.12	0.892	0.264	45.45 MB
CompGS (Liu et al. 2024)	27.26	0.803	0.240	17.31 MB	23.70	0.837	0.208	10.10 MB	29.69	0.901	0.279	9.20 MB
HAC (Chen et al. 2024)	<u>27.50</u>	0.806	<u>0.238</u>	<u>16.0 MB</u>	24.04	<u>0.846</u>	<u>0.187</u>	<u>8.51 MB</u>	<u>29.98</u>	0.902	0.269	4.62 MB
NeuralGS (Ours)	27.53	<u>0.807</u>	<u>0.238</u>	8.69 MB	<u>23.95</u>	0.847	0.186	6.33 MB	30.09	<u>0.906</u>	<u>0.252</u>	<u>5.76 MB</u>

Table 1: Quantitative results evaluated on Mip-NeRF 360, Tanks&Temples, and Deep Blending datasets. We highlight the best-performing results in **bold** and the second-best results in underlined for all compression methods.

tions for training and rendering.

Implementation Details. We implement our NeuralGS based on the official codes of 3DGS (Kerbl et al. 2023) and conduct training on various scenes using NVIDIA RTX 6000 Ada GPUs. During pruning, we remove 40% of the redundant 3D Gaussians. The number of clusters is determined adaptively, with each cluster containing 20k Gaussians on average. Notably, the total number of clusters does not need to be predefined. Each cluster is assigned a lightweight MLP to fit the corresponding Gaussian attributes for 60k iterations. All MLPs used in our method are 5-layer tiny MLPs with Tanh activation function and positional encoding. To restore rendering quality, we fine-tuned the fitted MLPs for 25k iterations, with λ and λ_{freq} set to 0.2 and 0.01.

4.2 Experimental Results

Quantitative Results. The comparison results for the rendering quality and model size across different datasets are presented in Tables 1. Specifically, compared to other compression works, NeuralGS achieves significant compression ratios while preserving rendering quality. Our method reduces the model size of original 3DGS (Kerbl et al. 2023) by approximately $87\times$, $69\times$ and $117\times$ on the *Mip-NeRF 360* dataset, *Deep Blending* dataset and *Tanks&Temples* dataset, respectively. The compression methods (Niedermayr et al. 2024; Girish et al. 2023; Xie et al. 2024; Fan et al. 2023) employ the quantization techniques and suffer from relatively large storage cost as excessive quantization beyond a certain threshold significantly degrades rendering quality. Compact3DGS (Lee et al. 2024a) uses a hash grid to only encode the color based on 3D positions and view directions, while opting to quantize the geometric attributes, leading to the limited compression ratio. Anchor-based works (Liu et al. 2024; Chen et al. 2024) achieves the impressive rendering quality and compression ratios, but exhibit slower rendering speeds compared to our approach, as shown in the following analysis of rendering time.

Qualitative Results. Figure 3 presents a qualitative comparison between our proposed NeuralGS and other compression methods (Fan et al. 2023; Girish et al. 2023; Lee et al. 2024a; Chen et al. 2024), providing the specific details with zoomed-in views. By leveraging compact cluster-based

FPS\Dataset	Mip-NeRF 360	Tanks&Temples	Deep Blending
3DGS (Kerbl et al. 2023)	112	162	118
CompGS (Liu et al. 2024)	94	105	125
HAC (Chen et al. 2024)	102	117	159
NeuralGS(Ours)	205	279	217

Table 2: Comparison of the rendering speed(FPS \uparrow). The rendering speed of all methods is measured on our machine.

neural fields to encode the Gaussian attributes, our method shows superior rendering quality with clearer textures and sharper edges even using a significantly smaller model size.

Rendering Time. In Table 2, we compare the rendering speed with the original 3DGS (Kerbl et al. 2023) and anchor-based works (Liu et al. 2024; Chen et al. 2024) which achieve SoTA compression performance. Rendering speed is measured in frames per second (FPS), computed based on the total time taken to render all camera views in the dataset. In our approach, multiple neural fields are used to encode Gaussian attributes of different clusters. During rendering, MLPs are used to decode the attributes of all 3D Gaussians in a single forward pass before testing FPS, which constitutes a one-time amortized cost for attribute loading. From Table 2, it is observed that our method achieves an averaged 1.8x rendering speed compared to 3DGS and significantly outperforms anchor-based works which need extra per-view attribute prediction based on view directions for rendering.

4.3 Ablation Studies

In this subsection, we conduct ablation studies on the *Deep Blending* dataset to demonstrate the effectiveness of each proposed component. Specifically, our core idea is to use neural fields to encode all Gaussian attributes. Hence, the baseline variant, referred to as “vanilla NeuralGS” in Table 3, employs a single tiny MLP to fit the attributes of all Gaussians in the scene, followed by basic fine-tuning to restore quality. As shown in Table 3, we incrementally incorporate each improvement to validate the effectiveness of our approach. More ablations, including different integration orders, the impact of different cluster numbers and comparisons of clustering algorithms are shown in the *appendix*.

Effectiveness of Cluster-based Fitting. As shown in Table 3, the Vanilla NeuralGS results in significant degradation




Figure 4: NeuralGS allows progressive loading new clusters in the *playroom* scene to obtain more details and sharper texture.

Dataset	Deep Blending Dataset (Hedman et al. 2018)			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Size \downarrow
3DGS (Kerbl et al. 2023)	29.42	0.900	0.247	672.8 MB
Vanilla NeuralGS	23.71	0.798	0.519	2.69 MB
+ Cluster-based fitting	28.98	0.893	0.283	5.79 MB
+ Importance weighting	29.67	0.903	0.264	5.75 MB
+ Frequency loss (Ours)	30.09	0.906	0.252	5.76 MB

Table 3: Quantitative ablation study on the Deep Blending dataset by *progressively* adding our proposed improvement.

of rendering quality compared to the original 3DGS (Kerbl et al. 2023). This is primarily due to the substantial variation among 3D Gaussians, where a single tiny MLP tends to produce substantial fitting errors. To mitigate this issue, we introduce a clustering strategy based on Gaussian attributes to ensure similarity within each cluster and assign a separate tiny MLP to fit the Gaussians of each cluster. As shown in Table 3, utilizing different tiny neural fields for different clusters significantly reduces fitting errors, leading to 5.3 dB improvement in PSNR and 10% increase in SSIM.

Effectiveness of Importance Weighting. Notably, it is unnecessary to equally fit every Gaussian in the scene. Instead, we assign each Gaussian an importance score to represent its contribution to the renderings. These scores are applied as weighting factors for the tiny MLPs during the fitting process, ensuring that important Gaussians are fitted by the neural fields with higher accuracy. As shown in Table 3 and Figure 5, adding importance scores as fitting weights, without introducing additional parameters, can further enhance visual quality and provide detailed textures.

Effectiveness of Frequency Loss. During the fine-tuning stage, we observed that within a limited number of training iterations, MLPs tend to be less sensitive to high-frequency details. As shown in the second row of Figure 5, incorporating the frequency loss makes the blurry edges of leaves sharper. The quantitative results in Table 3 further show the improvement in rendering quality by lastly introducing the frequency loss. We also provide results for only adding the frequency loss in the *appendix*.

4.4 JPEG-like Progressive Loading

Benefiting from our usage of different neural fields to fit the Gaussians within different cluster, we can transmit and decode Gaussian attributes cluster by cluster in a streamable




Figure 5: Ablation study about importance weighting (IW) and frequency loss (FL) in the *bicycle* and *stump* scenes.

manner like JPEG (Skodras et al. 2001). Specifically, we can sort clusters from the largest to the smallest based on the number of Gaussians and progressively transmit the positions along with the corresponding MLP weights. During transmission, Gaussian attributes can be decoded simultaneously, as shown in Figure 4, enabling a progressive loading for the entire scene and making it suitable for streamable applications. From magnified images, it is evident that newly loaded clusters contribute additional details and sharper texture, allowing the scene to gradually become clearer.

5 Conclusion

In this paper, we introduce *NeuralGS*, a novel and effective post-training compression for original 3DGS. To this end, we introduce a clustering strategy and fit all attributes of Gaussians within each cluster using different tiny MLPs, based on importance scores of Gaussians as fitting weights. Additionally, we introduce a frequency loss during the fine-tuning stage to better preserve high-frequency details. Extensive experiments demonstrate that our method achieves superior rendering quality compared to existing compression methods while utilizing significantly less model size.

6 Acknowledgments

This work was supported in part by the Natural Science Foundation of China (No. 62332002, 62425101) and Shenzhen Science and Technology Program (KQTD20240729102051063).

References

- Ali, M. S.; Bae, S.-H.; and Tartaglione, E. 2024. ELMGS: Enhancing memory and computation scalability through coMpression for 3D Gaussian Splatting. *arXiv preprint arXiv:2410.23213*.
- 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 *International Conference on Computer Vision (ICCV)*.
- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2022. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *Computer Vision and Pattern Recognition (CVPR)*.
- Cao, J.; Goel, V.; Wang, C.; Kag, A.; Hu, J.; Korolev, S.; Jiang, C.; Tulyakov, S.; and Ren, J. 2024. Lightweight Predictive 3D Gaussian Splats. *arXiv preprint arXiv:2406.19434*.
- Chen, Y.; Wu, Q.; Lin, W.; Harandi, M.; and Cai, J. 2024. HAC: Hash-grid Assisted Context for 3D Gaussian Splatting Compression. In *European Conference on Computer Vision*.
- Dai, J.; Zhang, Z.; Mao, S.; and Liu, D. 2019. A view synthesis-based 360° VR caching system over MEC-enabled C-RAN. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(10): 3843–3855.
- Drebin, R. A.; Carpenter, L.; and Hanrahan, P. 1988. Volume rendering. *ACM Siggraph Computer Graphics*, 22(4): 65–74.
- Fan, E. 2000. Extended tanh-function method and its applications to nonlinear equations. *Physics Letters A*, 277(4-5): 212–218.
- Fan, L.; Yang, Y.; Li, M.; Li, H.; and Zhang, Z. 2024. Trim 3D Gaussian Splatting for Accurate Geometry Representation. *arXiv preprint arXiv:2406.07499*.
- Fan, Z.; Wang, K.; Wen, K.; Zhu, Z.; Xu, D.; and Wang, Z. 2023. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. *arXiv preprint arXiv:2311.17245*.
- Girish, S.; Gupta, K.; Gupta, K.; and Shrivastava, A. 2023. Eagles: Efficient accelerated 3d gaussians with lightweight encodings. *arXiv preprint arXiv:2312.04564*.
- Govindarajan, S.; Sambugaro, Z.; Takikawa, T.; Rebain, D.; Sun, W.; Conci, N.; Yi, K. M.; Tagliasacchi, A.; et al. 2024. Lagrangian Hashing for Compressed Neural Field Representations. *arXiv preprint arXiv:2409.05334*.
- 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, W.; Wang, Y.; Ma, L.; Yang, B.; Gao, L.; Liu, X.; and Ma, Y. 2023. Tri-miprf: Tri-mip representation for efficient anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19774–19783.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics (TOG)*.
- 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, J. C.; Rho, D.; Sun, X.; Ko, J. H.; and Park, E. 2024a. Compact 3d gaussian representation for radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21719–21728.
- Lee, J. Y.; Wu, Y.; Zou, C.; Hoiem, D.; and Wang, S. 2024b. Plenoptic PNG: Real-Time Neural Radiance Fields in 150 KB. *arXiv preprint arXiv:2409.15689*.
- Liang, Z.; Zhang, Q.; Hu, W.; Feng, Y.; Zhu, L.; and Jia, K. 2024. Analytic-Splatting: Anti-Aliased 3D Gaussian Splatting via Analytic Integration. *arXiv preprint arXiv:2403.11056*.
- Likas, A.; Vlassis, N.; Verbeek, J. J.; and Verbeek, J. J. 2003. The global k-means clustering algorithm. *Pattern recognition*, 36(2): 451–461.
- Liu, B.; and Banerjee, S. 2024. SwinGS: Sliding Window Gaussian Splatting for Volumetric Video Streaming with Arbitrary Length. *arXiv preprint arXiv:2409.07759*.
- Liu, X.; Wu, X.; Zhang, P.; Wang, S.; Li, Z.; and Kwong, S. 2024. CompGs: Efficient 3d scene representation via compressed gaussian splatting. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 2936–2944.
- 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.
- Ma, Y.; Feng, K.; Zhang, X.; Liu, H.; Zhang, D. J.; Xing, J.; Zhang, Y.; Yang, A.; Wang, Z.; and Chen, Q. 2025a. Follow-Your-Creation: Empowering 4D Creation through Video In-painting. *arXiv preprint arXiv:2506.04590*.
- Ma, Y.; Liu, Y.; Zhu, Q.; Yang, A.; Feng, K.; Zhang, X.; Li, Z.; Han, S.; Qi, C.; and Chen, Q. 2025b. Follow-Your-Motion: Video Motion Transfer via Efficient Spatial-Temporal Decoupled Finetuning. *arXiv preprint arXiv:2506.05207*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *European Conference on Computer Vision (ECCV)*.
- Morgenstern, W.; Barthel, F.; Hilsmann, A.; and Eisert, P. 2023. Compact 3d scene representation via self-organizing gaussian grids. *arXiv preprint arXiv:2312.13299*.
- 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): 102:1–102:15.

Niedermayr, S.; Stumpfegger, J.; Stumpfegger, J.; and Westermann, R. 2024. Compressed 3d gaussian splatting for accelerated novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10349–10358.

Niemeyer, M.; Barron, J. T.; Mildenhall, B.; Sajjadi, M. S.; Geiger, A.; and Radwan, N. 2022. RegNeRF: Regularizing Neural Radiance Fields for View Synthesis from Sparse Inputs. In *Computer Vision and Pattern Recognition (CVPR)*.

Poole, B.; Jain, A.; Barron, J. T.; and Mildenhall, B. 2022. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*.

Reiser, C.; Peng, S.; Liao, Y.; and Geiger, A. 2021. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14335–14345.

Salman Ali, M.; Zhang, C.; Cagnazzo, M.; Valenzise, G.; Tartaglione, E.; and Bae, S.-H. 2025. Compression in 3D Gaussian Splatting: A Survey of Methods, Trends, and Future Directions. *arXiv e-prints*, arXiv–2502.

Skodras, A.; Christopoulos, C.; Ebrahimi, T.; and Ebrahimi, T. 2001. The JPEG 2000 still image compression standard. *IEEE Signal processing magazine*, 18(5): 36–58.

Sun, X.; Lee, J. C.; Rho, D.; Ko, J. H.; Ali, U.; and Park, E. 2024. F-3dgs: Factorized coordinates and representations for 3d gaussian splatting. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 7957–7965.

Tancik, M.; Weber, E.; Ng, E.; Li, R.; Yi, B.; Kerr, J.; Wang, T.; Kristoffersen, A.; Austin, J.; Salahi, K.; Ahuja, A.; McAllister, D.; and Kanazawa, A. 2023. Nerfstudio: A Modular Framework for Neural Radiance Field Development. In *ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH '23*.

Xie, S.; Zhang, W.; Tang, C.; Bai, Y.; Lu, R.; Ge, S.; and Wang, Z. 2024. MesonGS: Post-training Compression of 3D Gaussians via Efficient Attribute Transformation. *arXiv preprint arXiv:2409.09756*.

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.

Zhan, Y.-T.; Ho, C.-Y.; Yang, H.; Chen, Y.-H.; Chiang, J. C.; Liu, Y.-L.; and Peng, W.-H. 2025. CAT-3DGS: A context-adaptive triplane approach to rate-distortion-optimized 3DGS compression. *arXiv preprint arXiv:2503.00357*.

Zhang, J.; Zhan, F.; Xu, M.; Lu, S.; and Xing, E. 2024. Fregs: 3d gaussian splatting with progressive frequency regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21424–21433.

Zhou, T.; Tucker, R.; Flynn, J.; Fyffe, G.; and Snavely, N. 2018. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*.