

# SizeGS: Size-aware Compression of 3D Gaussian Splatting via Mixed Integer Programming

Shuzhao Xie\*

SIGS, Tsinghua University  
Shenzhen, China  
xsx24@mails.tsinghua.edu.cn

Shijia Ge

SIGS, Tsinghua University  
Shenzhen, China  
gsj23@mails.tsinghua.edu.cn

Yunpeng Bai

The University of Texas at Austin  
Texas, USA  
byp215@gmail.com

Jiahang Liu\*

Harbin Institute of Technology  
Shenzhen, China  
liu030526@gmail.com

Sicheng Pan

SIGS, Tsinghua University  
Shenzhen, China  
sccpan0306@gmail.com

Cong Zhang

Jiangxing Intelligence Inc.  
Shenzhen, China  
congzh@ieee.org

Zhi Wang<sup>†</sup>

SIGS, Tsinghua University  
Shenzhen, China  
wangzhi@sz.tsinghua.edu.cn

Weixiang Zhang

SIGS, Tsinghua University  
Shenzhen, China  
zwx.healthy@gmail.com

Chen Tang

MMLab, The Chinese University of  
Hong Kong  
Hong Kong, China  
chentang@link.cuhk.edu.hk

Xiaoyi Fan

Jiangxing Intelligence Inc.  
Shenzhen, China  
xiaoyi.fan@ieee.org

## Abstract

Recent advances in 3D Gaussian Splatting (3DGS) have greatly improved 3D reconstruction. However, its substantial data size poses a significant challenge for transmission and storage. While many compression techniques have been proposed, they fail to efficiently adapt to fluctuating network bandwidth, leading to resource wastage. We address this issue from the perspective of size-aware compression, where we aim to compress 3DGS to a desired size by quickly searching for suitable hyperparameters. Through a measurement study, we identify key hyperparameters that affect size—namely, the reserve ratio and bit-width settings. Then, we formulate this hyperparameter optimization problem as a mixed-integer nonlinear programming (MINLP) problem, with the goal of maximizing visual quality while respecting the size budget constraint. To solve the MINLP, we decouple this problem into two parts: discretely sampling the reserve ratio and determining the bit-width settings using integer linear programming (ILP). To solve the ILP more quickly and accurately, we design a quality loss estimator and a calibrated size estimator, as well as implement a CUDA kernel. Extensive experiments on multiple 3DGS variants demonstrate that our method achieves state-of-the-art performance in

post-training compression. Furthermore, our method can achieve comparable quality to leading training-required methods after fine-tuning. Project page & Code: [shuzhaoxie.github.io/sizogs](https://shuzhaoxie.github.io/sizogs).

## CCS Concepts

• **Computing methodologies** → *Image compression; Point-based models; Discrete space search; 3D imaging.*

## Keywords

3D Gaussian Splatting, Compression, Integer Linear Programming

## ACM Reference Format:

Shuzhao Xie\*, Jiahang Liu\*, Weixiang Zhang, Shijia Ge, Sicheng Pan, Chen Tang, Yunpeng Bai, Cong Zhang, Xiaoyi Fan, and Zhi Wang<sup>†</sup>. 2025. SizeGS: Size-aware Compression of 3D Gaussian Splatting via Mixed Integer Programming. In *Proceedings of the 33rd ACM International Conference on Multimedia (MM '25)*, October 27–31, 2025, Dublin, Ireland. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3746027.3755370>

## 1 Introduction

In recent years, 3D Gaussian Splatting (3DGS) [24] has revolutionized 3D scene representation and has been widely adopted in a variety of applications [7, 20, 28, 30, 32, 50, 56, 63, 70]. Despite its success, 3DGS still faces limitations in transmission efficiency due to its giant number of points and complex attributes. Though various 3DGS compression methods have been proposed [5], they overlook requirements arising from on-demand applications such as volumetric video streaming [56, 66], live streaming [21], and remote teleoperation [28]. These applications frequently encounter fluctuating network bandwidth, leading to jitter and blurriness that

\*Shuzhao Xie and Jiahang Liu contributed equally to this research.

<sup>†</sup>Corresponding author.



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

MM '25, Dublin, Ireland

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2035-2/2025/10

<https://doi.org/10.1145/3746027.3755370>

significantly impact user experience. These inefficiencies are further exacerbated in large-scale 3DGS scenes, which often exceed 20 megabytes (MB) even after compression. To overcome this challenge, size-aware 3DGS compression presents a promising direction. The key idea is to *compress the 3DGS to desired size by automatically searching the suitable hyperparameter set*.

Based on the way of hyperparameter configuration, current 3DGS compression methods fall into two categories. One is the *offline* method [16, 39, 41, 43, 58–60], which requires manually setting hyperparameters to adjust the compressed file size. Moreover, there are a variety of hyperparameters to consider, which demands a significant amount of human effort for adjustment. The other is the *online* method [10, 27, 35, 54, 57], which selects hyperparameters based on a context model and imports an extra hyperparameter  $\lambda$  to balance rate and bit consumption. While  $\lambda$  can adjust the final size, each adjustment requires retraining the context model from scratch, which is time-consuming. Additionally, it is difficult to predict the final size given a certain  $\lambda$ . Therefore, existing methods are unable to search for feasible hyperparameters under the constraint of the size budget in a short amount of time. Given the fast speed of recent integer programming solvers [64], we propose to achieve fast size-aware compression via mixed integer programming.

However, searching for appropriate 3DGS compression hyperparameters from an optimization perspective is non-trivial. First, 3DGS compression pipeline has lots hyperparameters. Some hyperparameters have a significant impact on size, while others are not. It is necessary to filter out a set of hyperparameters that significantly influence size. Second, to formulate an integer programming model, an explicit analytical relationship needs to be established between size and hyperparameters, as well as between quality and hyperparameters. This model can then use existing solvers to determine the appropriate hyperparameter settings. Finally, to fastly achieve the accurate results, an accurate approximation function for the actual size and acceleration techniques should be propose, which will allow the solver determine whether a candidate hyperparameter solution satisfies the constraints faster.

To tackle these challenges, we propose *SizeGS*, a codec that can compression 3DGS to a desired file size while maximizing visual quality. First, we summarize and analyze the pipeline of 3DGS compression and identify the two most influential hyperparameters on size: the **reserve ratio**  $\tau$  and **bit-width setting**  $Q$ . Then, we formulate the problem as a Mixed Integer Nonlinear Programming (MINLP) model, where the objective is visual quality and the constraint is size. To solve it, we decompose the MINLP into two subproblems: one is discrete sampling for  $\tau$ , and the other is Integer Linear Programming (ILP) for  $Q$ , which can be quickly solved using existing solvers [46]. Specifically, we perform multiple rounds of search to find suitable hyperparameters. In each round, we sample a value for  $\tau$ , solve for  $Q$  using ILP, compute the quality. We select the hyperparameters with the highest quality as the search result. To ensure that the bit-width settings search aligns with the ILP definition (i.e., both objective and constraint are linear functions of  $Q$ ), we propose using quantization loss to replace visual quality and construct a linear analytical relationship between size and  $Q$  for accurate size estimation. To accelerate the solution process, we implement a CUDA kernel for parallel quantization and reuse the search results from the previous round as the solver's search

starting point. Finally, we carefully design a three-step fine-tuning method to improve accuracy.

In summary, our contributions are as follows: **1)** We formulate the hyperparameter search problem in 3DGS compression as a Mixed Integer Nonlinear Programming (MINLP) model. To make this optimization tractable, we decouple this problem into two sub-problems: discrete sampling for the reserve ratio and Integer Linear Programming (ILP) for bit-width setting. **2)** To enable precise size estimation, we establish an analytical linear relationship between the bit-width setting and the resulting file size. To accelerate the ILP solver, we utilize the attribute loss to estimate the quality loss and implement a CUDA kernel to parallelize the quantization process. **3)** Extensive experiments on multiple 3DGS variants reveal that our method can search a set of hyperparameters to compress 3D Gaussians to desired size in a minute and achieves SOTA performance on offline compression. With finetuning, our method can achieve the comparable performance as the SOTA online compression methods.

## 2 Preliminary and Motivation

In this section, we first introduce the the data composition of 3DGS. Then, we summarize the general techniques of online and offline 3DGS compression methods. Finally, we identify the hyperparameters that greatly influence the compressed file size.

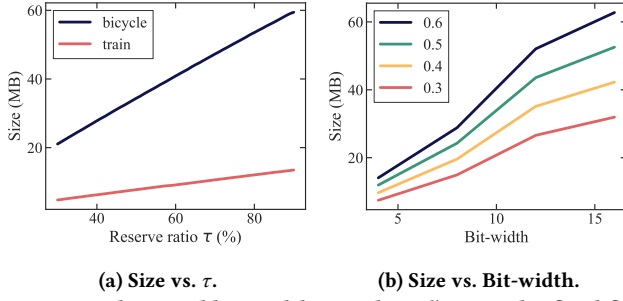
### 2.1 Base Models for 3DGS Compression

**3DGS** [24] consists of mulitple 3D Gaussian functions. Each Gaussian is characterized by a covariance matrix  $\Sigma$  and a center point  $\mu$ , which is referred to as the mean value of the Gaussian:  $G(x) = e^{-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)}$ . To maintain the positive definiteness of the covariance matrix  $\Sigma$ , 3DGS decomposes  $\Sigma$  into a scaling matrix  $S = \text{diag}(s)$ ,  $s \in \mathbb{R}^3$  and a rotation matrix  $R$ :  $\Sigma = RSS^T R^T$ . The rotation matrix  $R$  is parameterized by a rotation quaternion  $q \in \mathbb{R}^4$ . In summary, each element of 3D Gaussians contains: (1) a 3D center  $\mu \in \mathbb{R}^3$ ; (2) a rotation quaternion  $q \in \mathbb{R}^4$ ; (3) a scale vector  $s \in \mathbb{R}^3$ ; (4) a color feature defined by spherical harmonics (SH) coefficients  $SH \in \mathbb{R}^h$ , with  $h = 3(d+1)^2$ , where  $d$  is the harmonics degree; and (5) an opacity logit  $o \in \mathbb{R}$ .

**Scaffold-GS** [36] is a variant of 3DGS, widely adopted in 3DGS compression [10, 57]. It introduces *anchor points* to capture common attributes of local 3D Gaussians. Specifically, the *anchor points* are initialized from neural Gaussians by voxelizing the 3D scenes. Each anchor point has a context feature  $f \in \mathbb{R}^{32}$ , a location  $x \in \mathbb{R}^3$ , a scaling factor  $l \in \mathbb{R}^6$  and  $k$  learnable offset  $O \in \mathbb{R}^{k \times 3}$ .

**4DGS** [61] models dynamic scenes as spatio-temporal 4D volumes composed of 4D Gaussian primitives. Each 4D Gaussian consists of (1) a 3D center  $\mu \in \mathbb{R}^3$ ; (2) a rotation quaternion  $q \in \mathbb{R}^4$ ; (3) a scale vector  $s \in \mathbb{R}^3$ ; (4) a color feature defined by SH coefficients  $SH \in \mathbb{R}^h$ , with  $h = 3(d+1)^2$ , where  $d$  is the harmonics degree; (5) an opacity logit  $o \in \mathbb{R}$ ; (6) a time coordinate  $t \in \mathbb{R}$ ; (7) a time scale component  $s_t \in \mathbb{R}$ ; (8) a time-level rotation vector  $r_t \in \mathbb{R}^4$ .

**Summary.** Elements of base 3D Gaussian models can be divided into two components: *geometry* and *attributes*. The *geometry* consists of the coordinates of 3DGS. The *attributes* consists of the other components. For example, for ScaffoldGS, the *geometry* refers to



**Figure 1: The  $\tau$  and bit-width greatly influence the final file size and establish a linear relationship with the size.**

the coordinates of the anchor points, and the *attributes* refers to the context feature, scaling factors and learnable offsets.

## 2.2 Universal Pipeline of 3DGS Compression

3DGS compression pipeline consists of four steps: 1) pruning the unimportant Gaussian points, 2) Transforming the attributes (i.e., quaternions) to reduce data entropy; 3) Quantizing the attributes in a group-wise way; 4) Entropy coding the coordinates and attributes.

**Pruning.** Offline methods first prune the unimportant Gaussians based on an importance score [16, 41, 59], derived from the volume rendering function:  $C = \sum_{i \in N} c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)$ . Here,  $c_i$  and  $\alpha_i$  are the density and color of this pixel, computed by a Gaussian with covariance  $\Sigma$  multiplied by a per-point opacity and SH coefficients. Based on this function, the importance score is given by:

$$I_d = \sum_{p \in \mathcal{P}} \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (1)$$

$\mathcal{P}$  is the pixel set that is overlapped by the Gaussian  $g$ , and  $i$  is the rank of Gaussian  $g$  in a set of Gaussians that overlap with the pixel  $p$ . It uses a threshold  $\tau$  to prune Gaussians, which means that it retains the percent of  $\tau$  of the sorted Gaussians. Some works [16, 59] use volume to weight  $I_d$  to obtain a more precise importance estimation, but this improvement is negligible. Here the volume means the product of the scale vector.

**Transformation.** This technique is usually employed by offline compression method. A typical method is region-adaptive hierarchical transform (RAHT) [13], which is utilized by MesonGS [59] to decompose *attributes* into low-frequency and high-frequency coefficients. This step can reduce the entropy of the attributes, resulting in less information loss during subsequent quantization and leading to better entropy encoding results.

**Quantization.** Due to the extreme sensitivity of geometry, previous works typically quantize it to 16-bit. For attributes, previous methods partition each channel of attributes into multiple groups. For each group of attributes, there are currently two types of quantization approaches. The first is based on traditional compression models, where the quantization step  $Q_s \in [0, 1]$  are used to convert the float attributes into integers:  $\hat{x} = \lfloor x/Q_s \rfloor$ . The second approach adopts model-based quantization, which quantize a group of attributes  $x$  with the quantization bit-width  $Q_b \in [1, 32] \cap \mathbb{Z}$ :

$$\hat{x} = \lfloor \text{clamp}(x/s_x + z_x, 0, 2^{Q_b} - 1) \rfloor, \quad (2)$$

where  $s_x = [\max(x) - \min(x)]/2^{Q_b}$  and  $z_x = \lfloor 2^{Q_b} - \max(x)/s_x \rfloor$ . Here,  $\lfloor \cdot \rfloor$  represents the rounding-to-nearest function, and  $\hat{x}$  refers

to the quantized attributes. Besides, function  $\text{clamp}(\cdot)$  specifies a range of values. Values below the minimum are set to the minimum. Values above the maximum are set to the maximum. Without specification, we use the notation  $Q$  to denote the bit-width settings and  $Q$  to denote the number of possible bit-widths.

**Entropy Coding.** For quantized geometry, most existing methods utilize G-PCC [42] for entropy coding. For quantized attributes, some approaches [16, 41, 59] directly compress them using LZ77 codecs. Other methods [11, 35] model the probability distributions of quantized attributes as Gaussian Mixture Model (GMM) to provide additional information for entropy coding and thereby achieve higher compression rates. Specifically, they estimate the mean and variance for each quantization group to construct the GMM.

## 2.3 Motivation

**Size-sensitive Hyperparameters.** Among the above four steps, three of them require configuring the hyperparameters. These hyperparameters includes the reserve ratio  $\tau$ , bit-width setting  $Q$ , the number of quantization groups  $B$ , and the mean and variance used for constructing the probability distribution of entropy coding. Among these hyperparameters, as depicted in Fig. 1, the reserve ratio  $\tau$  and the bit-width setting  $Q$  show great impact on the final file size. The number of quantization groups  $B$  and the bit-width setting  $Q$  are actually inversely related, so only one needs to be retained for the search. Entropy encoding is used as the final step of the entire pipeline, and its corresponding hyperparameters have minimal impact on the final size [11, 65]. Therefore, we do not consider the hyperparameters of entropy encoding or the number of blocks during the hyperparameters optimization.

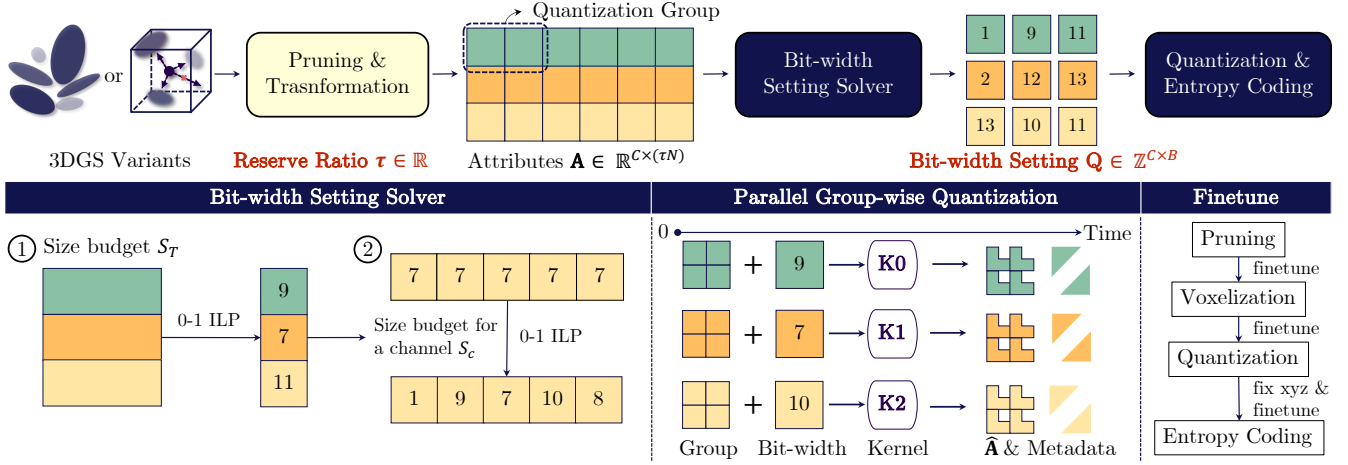
For *online compression*, hyperparameters are configured via a learned mask and context model:  $\tau$  from a learned mask,  $Q$  and distribution parameters from a context model. The online method balances rate and quality via  $\lambda$ , but adjusting hyperparameters requires retraining the context model. Hence, searching by size is time-consuming. For *offline compression*, all these hyperparameters need to be manually configured. Inspired by linear programming solvers [3, 46], we frame hyperparameter search as a planning problem and introduce techniques to speed up verification, enabling rapid iterative search.

## 3 Methodology

**Overview.** We present the compression pipeline and formulate size-aware 3DGS compression as a MINLP to find optimal reserve ratio  $\tau$  and bit-width  $Q$ . We split it into two parts: (1) discrete  $\tau$  sampling, and (2) solving  $Q$  via ILP with fixed  $\tau$ . We derive an analytical  $Q$ -to-size mapping, model it as a 0-1 ILP for accuracy, and refine size estimates iteratively. To speed up ILP, we use quantization loss to estimate quality drop and accelerate group-wise quantization via CUDA. Finally, we describe the finetuning process.

### 3.1 Problem Formulation and Decoupling

**Compression pipeline.** As shown in Fig. 2, given a pre-trained 3DGS, we prune the Gaussian points based on the importance score and percentage  $\tau$ , which means that  $\tau N$  of the most important points are reserved. To calculate the importance of anchor



**Figure 2: Compression pipeline and acceleration techniques.** We achieve size-aware 3DGS compression via hyperparameter search. As illustrated above, the compression process can be controlled by adjusting the **reserve ratio**  $\tau$  and **bit-width setting**  $\mathbf{Q}$ . Algo. 1 presents our hyperparameter optimization algorithm, which aims to meet the size target while maximizing quality. To accelerate the bit-width setting solver, we use 0-1 ILP to compute the bit-width for each channel, and then construct a 0-1 ILP per channel to solve for the bit-width of each group. Additionally, we implement parallel group-wise quantization to further speed up the process. Finally, we restore quality via piecewise fine-tuning after pruning, voxelization, and quantization.

points in ScaffoldGS, we average the importance of the corresponding generated Gaussian splats. Given the anchor point  $a$ , the set of training viewpoints  $\mathcal{V}_a$  from which  $a$  is visible, and the corresponding generated Gaussian points  $\mathcal{G}_a$ , the importance score  $I_a$  is:  $I_a = \sum_{v \in \mathcal{V}_a} \sum_{g \in \mathcal{G}_a} m_g I_d$ , where  $m_g \in \{1, 0\}$  reflects whether Gaussian  $g$  is retained or discarded under viewpoint  $v$ .  $I_d$  aggregates contributions across pixels overlapped by  $g$ , as shown in Eq. 1. For 3DGS and 4DGS, we use the Eq. 1.

We then quantize geometry and attributes. Following previous works [10, 11, 35], coordinates are set to 16-bit and compressed using G-PCC [42]. Attribute tensor  $\mathbf{A} \in \mathbb{R}^{C \times (\tau N)}$  is quantized group-wise by splitting the first dimension into  $C$  parts and the second into  $B$  nearly equal parts, forming  $C \times B$  **quantization groups**. We adopt the bit-width to quantize the attribute. Compared to the quantization step, a continuous variable, the bit-width is discrete, resulting in a smaller search space that is easier for optimization. Quantized attributes are then compressed with LZ77 or torchac [37]. The final file includes geometry, attributes, and metadata (voxel size, block count, and per-group scale/zero points).

**Problem Formulation.** Given the above pipeline, our goal is to search for suitable values of the  $\tau$  and  $\mathbf{Q}$  under a given size constraint, while maximizing visual quality. Hence, we formulate this problem as a mixed integer nonlinear programming (MINLP) model:

$$\begin{aligned}
 & \underset{\tau, \mathbf{Q}}{\text{minimize}} && \mathcal{M}(\tau, \mathbf{Q}), \\
 & \text{subject to} && \mathcal{S}(\tau, \mathbf{Q}) \leq \text{Size Budget}, \\
 & && \tau \in [0, 1], \\
 & && \mathbf{Q} \in [1, 32]^{C \times B} \cap \mathbb{Z}^{C \times B}.
 \end{aligned} \tag{3}$$

Here,  $\mathcal{M}(\cdot)$  and  $\mathcal{S}(\cdot)$  are the quality loss and the estimated size of compressed model under the configuration of  $\{\tau, \mathbf{Q}\}$ , respectively.

“Mixed” in MINLP means that the variable types include both discrete ( $\mathbf{Q}$ ) and continuous ( $\tau$ ) variables. A detailed introduction of MINLP are proposed in the supplementary material.

**Structure Decoupling.** Since both  $\mathcal{M}(\cdot)$  and  $\mathcal{S}(\cdot)$  are nonlinear functions of  $\{\tau, \mathbf{Q}\}$ , and the search space is extremely large, solving this problem directly is highly challenging. For example, even when  $C = 8$  and  $B = 8$ , the discrete combinatorial space already contains  $32^{64}$  possibilities—far beyond astronomical scale! Moreover, optimizing the continuous variable  $\tau$  within this vast space further increases the difficulty. Besides, existing MINLP solvers [1, 2] often require hours to find a reasonable solution.

However, we observe only  $\tau$  is continuous among these variables, while others are integers. Besides, by fixing  $\tau$ , we find a linear relationship between  $\mathcal{S}(\cdot)$  and  $\mathbf{Q}$ . E.g., as shown in Fig. 1b, using 8-bit across all groups yields proportional size, and 16-bit files are roughly twice as large as 8-bit ones—allowing size estimation from lower-bit results. Hence, if we fix the value of  $\tau$  and approximate  $\mathcal{M}$  and  $\mathcal{S}$  as functions that are linearizable with respect to  $\mathbf{Q}$ , then this problem can be decoupled into two problems: 1) using discrete sampling to search and fix the value of  $\tau$ , and then 2) solve the  $\mathbf{Q}$  via integer linear programming (ILP). The advantage of such a decoupling lies in the efficiency of recent ILP solvers [3, 46], which requires a few seconds to obtain a solution.

## 3.2 Solve the MINLP

**Discrete Sampling of  $\tau$ .** Based on the above observations, we propose to traverse the value of reserve ratio  $\tau$ . Then, we can search the bit-width settings  $\mathbf{Q}$  by solving the ILP. Hence, the hyperparameter searching algorithm for size-aware 3D Gaussian compression are proposed in Algo. 1. The input is a target size and a pretrained 3DGS, and the output is the optimal hyperparameters set  $\Phi^* = \{\tau^*, \mathbf{Q}^*\}$  that can satisfy the size constraint while maximizing visual quality. We iterate over all possible values of  $\tau$ . Here,  $2 \times S_a < S_T$  refers

**Algorithm 1: Hyperparameter Optimization**


---

**Input:** Size budget  $S_T$  and a pre-trained 3DGS model  
**Output:** Hyperparameter set  $\Phi^* = \{\tau^*, Q^*\}$

```

1 Initialize bit-width setting  $Q \leftarrow \{8\}^{C \times B}$ ;
2 Initialize the best quality  $M^* \leftarrow 0$ ;
3 for  $\tau \in \{\tau_1, \tau_2, \tau_3, \dots\}$  do
4   Compress model with current hyperparameters, obtain size  $S_a$ ;
5   if  $2 \times S_a < S_T$  then
6     continue;
7   while true do
8      $S_\Delta \leftarrow S_a - S_T$ ;
9     /* 01-ILP ( $\cdot$ ) solver takes  $Q$  and  $S_\Delta$  as input. */
10    Search bit-widths via 01-ILP:  $Q \leftarrow \text{01-ILP}(Q, S_\Delta)$ ;
11    Compress model with  $\tau$  and updated  $Q$ , obtain new size  $S_a$ ;
12    if  $\frac{|S_a - S_T|}{|S_T|} < 0.05$  then
13      break;
14    /* Retain the one with the best quality. */
15    if  $\Omega(Q) > M^*$  then
16       $M^* \leftarrow \Omega(Q)$ ;
17       $\Phi^* \leftarrow \{\tau, Q\}$ ;
18 return  $\Phi^*$ ;
```

---

to the situation where, if doubling the bit-width setting under the current configuration (up to a maximum value of 16) cannot achieve the target size, it indicates that  $\tau$  is too small, and thus, more Gaussian points need to be retained. For the bit-width setting search, we formulate the problem as a 0-1 Integer Linear Programming (ILP) task. The ILP solver takes  $Q$  and  $S_\Delta$  as inputs, where  $Q$  serves as the initialization for each search iteration, enabling faster convergence, and  $S_\Delta$  is used to calibrate the estimated size for more accurate results. After each iteration, the newly determined bit-width setting is used to store the results. If the stored result size  $S_a$  is sufficiently close to the target size  $S_T$ , the search is terminated. For each pair of  $\{\tau, Q\}$ , we compute the estimated quality  $\Omega$  and retain the pair that yields the highest quality. We will introduce  $\Omega$  below.

**Binary ILP for  $Q$ .** The input of mixed precision quantization is the important attributes  $\mathcal{A}$ , which can be seen as a 2D-matrix. We divide this 2D-matrix into  $C \times B$  blocks and there are  $Q$  quantization options for each attribute block (e.g., 16 options for 1–16 bits). The search space of the ILP problem is  $(C \times B)^Q$ . The objective of solving the ILP is to find the best bit configuration in this search space that optimally balances quality loss  $\Omega$  and the size limit  $S$ . Besides, we define the bit-width variables as  $Q \in \{0, 1\}^{C \times B \times Q}$ , which means that we use a one-hot vector  $v \in \{0, 1\}^Q, |v| = 1$  to represent the bit-width setting for an attribute block. In all, the 0-1 ILP model tries to find the right bit-width setting  $Q$  can be formulated as

$$\begin{aligned}
& \underset{Q}{\text{minimize}} && \Omega(Q), \\
& \text{subject to} && S(Q) \leq \text{Size Budget}, \\
& && \forall (i, j) \in [0, C) \times [0, B), \sum_{q=1}^Q Q_{i,j,q} = 1.
\end{aligned} \tag{4}$$

Here,  $S(Q)$  and  $\Omega(Q)$  denote the estimated file size and estimated

quality loss under a bit-width setting  $Q$ , respectively. Next, we introduce the details of the  $\Omega$  and  $S$ .

**Acceleration.** We propose the following techniques to accelerate the problem solving. Since computing the quality metric like PSNR requires traversing the entire training set and evaluating the metric, it takes at least 10 seconds, which significantly slows down the search process. To solve this, we use quantization loss  $\Omega(\cdot)$  to replace the metric function  $\mathcal{M}(\cdot)$ . Here, we assume that the bit-width of each group are independent of one another. This allows us to precompute the quantization loss of each group of attributes separately, and it only requires  $Q$  times quantizations. As for the metric of quality loss, we use the distance between the original attributes and the restored attributes<sup>1</sup>. Formally, we can precompute the estimated quality loss matrix  $\Omega \in \mathbb{R}^{C \times B \times Q}$  with:

$$\Omega(i, j, b) = |\hat{\mathcal{A}}_{i,j}^b - \mathcal{A}_{i,j}|. \tag{5}$$

The  $|\cdot|$  can be 1-norm, 2-norm, or  $\infty$ -norm. We plot the relationship between PSNR and  $\Omega$  in the supplementary material, which reveals that minimizing  $\Omega$  is equal to maximizing PSNR. Besides, we set  $Q$  as 16 to prune the search space. Finally, we implement a CUDA kernel to accelerate the quantization process, in which each quantization group is quantized in parallel, as shown in Fig. 2.

**Size Estimator.** A compressed 3DGS file contains the following components: 1) voxelized coordinates, 2) quantized attributes, and 3) metadata, which is used to restore the coordinates and attributes. Fortunately, we can store the voxelized coordinates and metadata to obtain the accurate compressed size in a few seconds. Then the challenge of size estimator lies in estimating the size of quantized attributes. This is to say, we only have to establish a analytical relationship between the compressed file size and the bit-width settings, the actual size of other components can be obtained by saving them to storage. Besides, the size estimator must be a linear function of the bit-width variables.

According to information theory [12], the lower bound of bit consumption can be calculated by  $\tau N \times (-\sum_i p_i \log_2 p_i)$ . However, such a size estimator is not suitable for the formulation of ILP. The reason is that, for each searching iteration, we have to quantize the attributes into integers with the candidate bit-width setting and calculate the probability of each values to derive the bit consumption, which costs a lot of time. Moreover, the relationship between the bit-width settings and the estimated size is non-linear, which cannot satisfy the linear requirement of the ILP. Hence, an explicit and linear relationship between the size  $S$  and the bit-width setting  $Q$  must be established. Thus, we estimate the size by

$$S(Q) = \sum_{i,j} P_{i,j} Q_{i,j} + C + S_\Delta. \tag{6}$$

Here,  $P \in \mathbb{R}^{C \times B}$  refers to size of quantization groups.  $C$  refers to the accurate storage consumption of the metadata and the coordinates, which can be obtained by storing them to the disk directly. This process is very fast. Of course, such a estimation for the compressed file size is not accurate. To calibrate it, we update the  $S_\Delta$  multiple times, as shown in Algo. 1.


**Hierarchical Solver.** As there are  $(C \times B)^Q$  options for  $Q$ , if we set  $(C, B, Q)$  to  $(73, 60, 16)$ , a typically setting for compressing ScaffoldGS, then there are  $1.8 \times 10^{58}$  options. In such a giant search

<sup>1</sup>Similar assumption can be found in [14, 15].



Method	Mip-NeRF 360 - 18.33 MB					Tank&Temples - 11 MB					Deep Blending - 8 MB				
	PSNR (dB)↑	SSIM ↑	LPIPS ↓	Size (MB) ↓	Time (s) ↓	PSNR (dB)↑	SSIM ↑	LPIPS ↓	Size (MB) ↓	Time (s) ↓	PSNR (dB)↑	SSIM ↑	LPIPS ↓	Size (MB) ↓	Time (s) ↓
HAC	27.17	0.789	0.261	18.29	16627	24.45	0.854	0.179	10.92	10978	30.27	0.910	0.254	8.29	9993
Our	27.48	0.806	0.240	18.17	1328	24.04	0.840	0.200	10.93	1381	30.24	0.903	0.271	7.92	1263

**Table 1: Performance on size-aware compression. “18.33 MB” refers to the average size budget of the scenes in Mip-NeRF 360. The detailed size target of scenes are listed in supplementary material. The best results are highlighted in red cells.**



**Figure 3: Qualitative results. We present the rendering results (rows 1) along with the corresponding error maps (rows 2) from a randomly selected viewpoint of the bicycle scene.**

space, figuring out a suitable bit-width setting is time-consuming. Neither of GPU-based [3] or CPU-based [46] solver can solve this in minutes. Hence, to accelerate the solving process, we solve this 0-1 ILP problem with two steps. As shown in the left bottom of Fig. 2, at the first step, we search the channel-level bit-width setting  $Q_c \in [1, 16]^C$ . Then, we calculate the size budget for each channel of attributes based on the  $Q_c$ . For example, the size budget of channel  $i$  can be computed by:  $S_c = S_T \frac{Q_{c,i}}{\sum Q_{c,i}}$ . At the second step, we solve the group-level bit-width setting  $Q_g \in [0, 16]^B$  for each channel based on the size limit  $S_c$ .

### 3.3 Piecewise Finetuning

For compression with finetuning, we directly quantize the coordinates and attributes without applying any transformations to them. To obtain a better compression quality, after the point pruning and coordinate quantization steps, we finetune for multiple epochs respectively to restore the reconstruction quality, as depicted in the right bottom of Fig. 2.

During fine-tuning, we fix coordinates because G-PCC decomposition yields unordered points, misaligned with attributes. To align them, we sort both by Morton order after second-stage finetuning and build quantization groups on attributes. Changing the coordinates afterward may change the Morton order, which reorders both coordinates and attributes. This reordering breaks the original grouping of attributes used for quantization, making the previously searched bit-widths no longer valid.

## 4 Experiments

**Datasets.** We conduct experiments on four datasets: 1) Mip-NeRF 360 [6]. This dataset contains five outdoor and four indoor scenes. Each scene contains 100 to 300 images. We use the images at  $1600 \times 1063$ . 2) Tank & Temples [25]. This dataset contains the *train* and *truck* scenes. 3) Deep Blending [23]. This dataset contains the *drjohnson* and *playroom* scenes. 4) Synthetic-NeRF [38]. This is a view synthesis dataset consisting of 8 synthetic scans, with 100 views used for training and 200 views for testing.

Method	PSNR	SSIM	LPIPS	Size	Finetune Time (s)
MesonGS	25.95	0.7706	0.2679	35.41 MB	0
Our+MesonGS	26.06	0.7743	0.2642	35.36 MB	0
4DGS	32.06	-	-	5.10 GB	0
Our+4DGS	32.07	-	-	198.64 MB	28

**Table 2: Quantitative results on 3DGS variants.**

**Baselines.** We compare our method with the following baselines: 3DGS [24], ScaffoldGS [36], C3DGS [41], Lee *et al.* [27], LightGaussian [16], EAGLES [19], SOGS [39], Compact3D [40], ReduGS [43], MesonGS [59], HAC [10], DVGO [49], VQRF [29], and ACRF [17]. Some of results are derived from HAC [10], 3DGS.zip [5], and MesonGS [59], while the visual results are produced by ourselves.

### 4.1 Experimental Results

**Size-aware Compression on Static Gaussians.** We evaluate end-to-end performance of our method via latency and quality metrics. In Tab. 1, our method is at most  $8\times$  faster than the baseline and achieves better or comparable quality across three datasets. For comparison with the baseline method, we set the number of finetune iterations as 4000. In practical applications, however, fine-tuning for only 500 steps is expected to achieve satisfactory visual quality. We perform a binary search on the hyperparameter  $\lambda$  of HAC to find a configuration that meets the size budget. The search stops when the difference between the obtained size and the size budget is within 5%. The reason why HAC takes so much longer than ours is that, after each hyperparameter adjustment, it requires 10-20 minutes to retrain the mask and context models in order to converge to a specific size. In contrast, our method first searches for the appropriate hyperparameters based on the size, and then only requires a single retraining. Since the hyperparameters are fixed, our method has almost no impact on the file size during retraining.

**Size-aware Compression on Dynamic Gaussians.** In the bottom part Tab. 2, we applied the SizeGS framework on 4DGS [62]. SizeGS can achieve the target size in 51s and improved visual quality after 28s of finetuning. Experiments are conducted on the *flame\_steak* scene of N3DV dataset [31].




Figure 4: Rate-distortion curves for quantitative comparison. Note that our goal is not to improve the marginal performance and defeat the existing compression works. Instead, we aim to design a hyperparameter parameter searching algorithm to compress the 3DGS model into the desired size while maximizing visual quality.

Base model	$\Omega$	Save	01-ILP	Finetune	End2end
ScaffoldGS	0.11	32.78	145.22	857.71	1328.82
3DGS	0.49	15.10	58.52	-	104.51

Table 3: Decomposition of time. The unit of time is second. From the left to right: time to calculate the estimated quality loss  $\Omega$ , time used to save to storage, time used by ILP solver, time used to finetune the model, and the end-to-end time.

**Performance on Offline Compression.** We implemented our search algorithm on top of the SOTA post-training compression – MesonGS [59]. The results on Mip-NeRF 360 are shown in Tab. 2. We can observe that our method enhance the compression quality of MesonGS. For each scene, the size target is set to the compressed file size generated by the official MesonGS configuration.

**Time Breakdown.** Tab. 3 shows the time breakdown of a round of search. First, the time spent on quality estimation is very short due to parallel quantization. Next, we search  $Q$  based on the estimated quality, where size calibration and 0-1 ILP take similar time. After searching for the hyperparameters, if we fine-tune for 6000 iterations, the average time required is approximately 1000s. We set the time limitation for 0-1 ILP to 50s. For ScaffoldGS, we fix the value of  $\tau$  to 0.6. For 3DGS, we use the configuration from MesonGS. We can see that it takes 1 minute for ILP to solve the bit-width searching. This is because 3DGS only has 10 channels that are involved in the search, while ScaffoldGS has 73.

**Qualitative Evaluation.** In Fig. 3, we present the rendering results and the corresponding error maps. From the error maps, it is evident that our method handles chair reflections better than other methods while achieving rendering results that are comparable to ScaffoldGS.

**Compared to Online Methods.** The quantitative compression results of different methods are presented in Fig. 4. With enough finetune iterations, our method outperforms most others across all three datasets and achieves performance comparable to leading methods. As implicit neural representations [67–69] serve as an important base model for 3D compression, we also compare with NeRF compression in supplementary material.

## 4.2 Ablation Study

Unless stated otherwise, all experiments use the *bicycle* scene from Mip-NeRF 360 and ScaffoldGS as the base model.

**0-1 ILP Superiority in Searching Bit-widths.** In solving the optimal bit-width setting for different attribute channels, we also

Method	Budget (B)	Searched (B)	$\Delta$ size (B)	Information loss
GA	$3 \times 10^7$	21,833,128	8,166,872	42,821,038
Vanilla ILP		28,934,805	1,065,195	1,258,394
0-1 ILP (Our)		29,831,203	168,797	11,826

Table 4: Superiority of 0-1 ILP. “GA”: Genetic Algorithm. With up to 16 bit-width choices, the Vanilla ILP and GA that are widely adopted in model quantization methods are unable to quickly search for suitable mixed-precision settings.

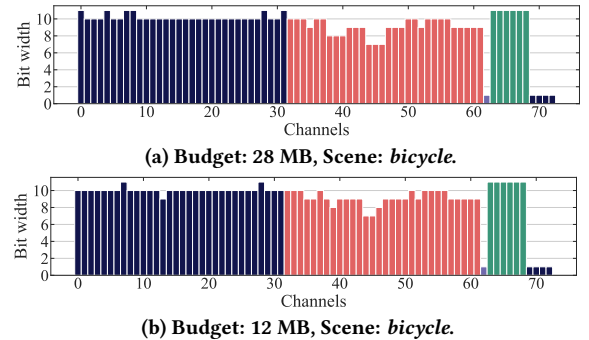


Figure 5: Bit widths of each channel. We use different colors to represent different kinds of attributes. From left to right are: features  $f$ , offsets  $O$ , opacity  $o$ , scaling  $l$ , and rotation  $r$ .

demonstrate the superiority of the 0-1 ILP. As shown in Tab. 4, we experiment with widely-used General ILP [64] and genetic algorithms [22, 51], both of which proved inferior. The 0-1 ILP fully utilizes the size budget while minimizing information loss. General ILP involves variables ranging from 1 to 16. In contrast, 0-1 ILP’s binary values offer finer control, easier integration of constraints, and more efficient solution techniques. Genetic algorithms, though suited for black-box problems, handle constraints less efficiently, making them unsuitable for our problem.

**Bit-widths.** In Fig. 5, we show the bit-widths for different attribute channels, where different colors represent different attributes. Here, we assume that the quantization groups that belong to the same channel share the same bit-width setting. From left to right in order: features  $f$ , offsets  $O$ , opacity  $o$ , scaling  $l$ , and rotation  $r$ . We can see that under different size budgets, the choices of bit-widths are generally the same.

**Effectiveness of Mixed Bit-width Setting.** An alternative is letting the groups that belong to the same channel share the same

Prune	Voxel	Quant	PNSR (dB) $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Size (MB) $\downarrow$
0	0	6000	24.95	0.7331	0.2709	24.20
1000	0	5000	24.98	0.7339	0.2706	24.19
0	1000	5000	24.99	0.734	0.2707	24.12
1000	1000	4000	25.19	0.7451	0.2619	24.25

**Table 5: Efficiency of piecewise finetuning. “Prune”, “Voxel”, and “Quant” refer to the finetune iterations after the three part. We fix the total iterations as 6000.**

$K$	Budget (MB)	PNSR (dB) $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Searched Size (MB) $\downarrow$
40	30	25.13	0.7410	0.2684	29.85
30		25.15	0.7411	0.2685	29.91
50		25.14	0.7413	0.2686	29.86
40	20	25.07	0.7353	0.2752	19.83
30		25.07	0.7357	0.2757	19.85
50		25.12	0.7368	0.2743	19.92

**Table 6: Robustness Evaluation. The bit-width setting can adapt to different values of the number of blocks  $K$ , ensuring the visual quality within a given size is not affected by  $K$ .**

bit-width, like the first step of hierarchical solver in Sec. 3.2. To verify the necessity of group-wise bit-width setting, we compare the effectiveness of channel-wise and group-wise mixed bit-widths setting in Fig. 6a, both with sufficient finetuning. Results show that finer granularity yields better performance.

**Effectiveness of Piecewise Finetune.** In Tab. 5, we evaluate performance of finetuned 3DGS under four iteration settings. We set the target size as 24 MB. Piecewise fine-tuning helps improve the upper bound of compression quality.


**Estimated Quality Loss  $\Omega$ .** There are many ways to calculate the distance between the original attributes and the restored attributes. To investigate which metric is better, in Fig. 6b, we show the impact of different metrics on the information loss of the final rendering results. We present the PNSR-Size curves under three metrics, including 1-norm, 2-norm, and  $\infty$ -norm. It can be seen that the performance of 2-norm and  $\infty$ -norm are nearly the same, both of which outperform 1-norm by a significant margin.

**Robustness Evaluation.** We evaluated the size and corresponding performance of the searching algorithm under different numbers of blocks. As shown in Tab. 6, for varying numbers of blocks and different target sizes, our method consistently finds appropriate bit-width settings, ensuring that the final file size is close to the target while maintaining optimal visual quality. Regardless of the block number setting, the final file size and performance are similar, indicating that our method is robust to the number of blocks.

## 5 Related Work

### 5.1 3D Gaussian Splatting and Its Compression

3D Gaussian Splatting (3DGS) [24] achieves excellent 3D reconstruction but suffers from large model size. Early efforts [16, 19, 27, 35, 39–41, 43, 53, 54, 58, 59] focused on compressing the original 3DGS. Later, works [10, 26, 33] targeted more efficient 3DGS variants [4, 18, 36, 45, 65], especially ScaffoldGS [36], which groups anchors into voxels and uses per-voxel features to predict Gaussians. HAC [10, 11] leverages 3D coordinates to guide quantization and entropy coding. ContextGS [57] encodes anchors hierarchically. HEMGS [33] uses a pretrained PointNet++[44] for better context modeling. FCGS[9] compresses 3DGS via one-shot inference but



**Figure 6: Ablation studies. (a) Results were conducted with sufficient finetuning, confirming that mixed bit-width settings enhances the upper bound of compression quality. (b) 2- and  $\infty$ -norms significantly outperform 1-norm.**

supports only fixed-size output and lacks adaptability to other 3DGS forms. Some works [8, 47, 48] address bitrate fluctuation via layered coding, which is orthogonal to our approach.

Unlike the previous works, our work propose to configure the hyperparameters of 3DGS from the perspective of combinatorial optimization and design many techniques to accelerate the process.

### 5.2 Mixed Precision Quantization

Our work compress the 3DGS to desired size by searching for two hyperparameters: one is  $\tau$ , and the other is the bit-width setting  $Q$ . The optimization for the bit-width setting is closely related to mixed precision quantization for deep learning models.

Mixed Precision quantization (MPQ) is a widely-used technique to improve the trade-off between the accuracy and efficiency of neural networks [14, 15, 52, 55, 64]. The challenge with this approach is to find the right mixed-precision setting for the different layers of neural networks. A brute force approach is not feasible since the search space is exponentially large in the number of layers. HAQ [55] employed reinforcement learning to search this space. However, this RL-based solution requires tremendous computational resources. HAWQ [14, 15, 64] proposed to assign each layer a sensitivity score with the Hessian spectrum and then use an ILP solver to generate mixed-precision settings with various constraints (such as model size and latency). Though CA-NeRF [34] uses the MPQ scheme, their method cannot be applied to 3DGS.

Unlike the HAWQ3 [64] that uses the ILP, we use binary ILP to search a better results for the bit-width settings of 3DGS compression. Though some works [10, 35] employed MPQ, their required retraining for configuring the quantization settings.

## 6 Conclusion

In this paper, we present *SizeGS*, a method for automatically selecting hyperparameters to compress 3D Gaussians to a target file size while maximizing visual quality. We formulate this problem as a mixed integer non-linear programming model and decouples this problem into two steps: discrete sampling of reserve ratio  $\tau$  and ILP for bit-width settings  $Q$ . To accelerate the ILP process, we use quantization loss to replace the quality metric and write a CUDA kernel to parallelize the quantization. We also design a size estimator to help search out a more accurate hyperparameter set. Experiments show *SizeGS* effectively controls file size while preserving quality.



## Acknowledgments

We sincerely thank the anonymous reviewers from ICLR, CVPR, and ACM MM for their valuable feedback and suggestions. We also thank our lab mates for their help in improving the manuscript. This work is supported in part by National Key Research and Development Project of China (Grant No. 2023YFF0905502), National Natural Science Foundation of China (Grant No. 92467204 and 62472249), and Shenzhen Science and Technology Program (Grant No. JCYJ20220818101014030 and KJZD20240903102300001).

## References

- [1] 2018. *The SCIP Optimization Suite 6.0*. Retrieved April 2, 2025 from <https://www.scipopt.org/>
- [2] 2025. *BARON Solver*. Retrieved April 2, 2025 from <https://minlp.com/baron-solver>
- [3] Ahmed Abbas and Paul Swoboda. 2022. FastDOG: Fast Discrete Optimization on GPU. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*. IEEE, 439–449. doi:10.1109/CVPR52688.2022.00053
- [4] Muhammad Salman Ali, Maryam Qamar, Sung-Ho Bae, and Enzo Tartaglione. 2024. Trimming the Fat: Efficient Compression of 3D Gaussian Splats through Pruning. In *BMVC*.
- [5] Milena T. Bagdasarian, Paul Knoll, Florian Barthel, Anna Hilsman, Peter Eisert, and Wieland Morgenstern. 2024. 3DGS.zip: A survey on 3D Gaussian Splatting Compression Methods. arXiv:2407.09510 [cs.CV] <https://arxiv.org/abs/2407.09510>
- [6] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. 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.
- [7] Anpei Chen, Haoqi Xu, Stefano Esposito, Siyu Tang, and Andreas Geiger. 2025. Lara: Efficient large-baseline radiance fields. In *European Conference on Computer Vision*. Springer, 338–355.
- [8] Yihang Chen, Mengyao Li, Qianyi Wu, Weiyao Lin, Mehrtash Harandi, and Jianfei Cai. 2025. PCGS: Progressive Compression of 3D Gaussian Splatting. arXiv preprint arXiv:2503.08511 (2025).
- [9] Yihang Chen, Qianyi Wu, Mengyao Li, Weiyao Lin, Mehrtash Harandi, and Jianfei Cai. 2024. Fast Feedforward 3D Gaussian Splatting Compression. arXiv preprint arXiv:2410.08017 (2024).
- [10] Yihang Chen, Qianyi Wu, Weiyao Lin, Mehrtash Harandi, and Jianfei Cai. 2024. HAC: Hash-grid Assisted Context for 3D Gaussian Splatting Compression. In *European Conference on Computer Vision*.
- [11] Yihang Chen, Qianyi Wu, Weiyao Lin, Mehrtash Harandi, and Jianfei Cai. 2025. HAC++: Towards 100X Compression of 3D Gaussian Splatting. arXiv preprint arXiv:2501.12255 (2025).
- [12] Thomas M Cover. 1999. *Elements of information theory*. John Wiley & Sons.
- [13] Ricardo L De Queiroz and Philip A Chou. 2016. Compression of 3D point clouds using a region-adaptive hierarchical transform. *IEEE Transactions on Image Processing* 25, 8 (2016), 3947–3956.
- [14] Zhen Dong, Zhewei Yao, Daiyaan Arfeen, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. 2020. HAWQ-V2: Hessian Aware trace-Weighted Quantization of Neural Networks. In *Advances in neural information processing systems (NeurIPS)*.
- [15] Zhen Dong, Zhewei Yao, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. 2019. HAWQ: Hessian AWARE Quantization of Neural Networks With Mixed-Precision. In *The IEEE International Conference on Computer Vision (ICCV)*.
- [16] Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, Dejia Xu, and Zhangyang Wang. 2024. LightGaussian: Unbounded 3D Gaussian Compression with 15x Reduction and 200+ FPS. In *Advances in neural information processing systems (NeurIPS)*.
- [17] Guangchi Fang, Qingyong Hu, Longguang Wang, and Yulan Guo. 2024. ACRF: Compressing Explicit Neural Radiance Fields via Attribute Compression. In *International Conference on Learning Representations (ICLR)*.
- [18] Guangchi Fang and Bing Wang. 2024. Mini-Splatting: Representing Scenes with a Constrained Number of Gaussians. In *European Conference on Computer Vision*.
- [19] Sharath Girish, Kamal Gupta, and Abhinav Shrivastava. 2024. Eagles: Efficient accelerated 3d gaussians with lightweight encodings. In *European Conference on Computer Vision*.
- [20] Chenghao Gu, Zhenzhe Li, Zhengqi Zhang, Yunpeng Bai, Shuzhao Xie, and Zhi Wang. 2024. DragScene: Interactive 3D Scene Editing with Single-view Drag Instructions. arXiv:2412.13552 [cs.CV]
- [21] Yongjie Guan, Xueyu Hou, Nan Wu, Bo Han, and Tao Han. 2023. MetaStream: Live Volumetric Content Capture, Creation, Delivery, and Rendering in Real Time. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking (Madrid, Spain) (ACM MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 29, 15 pages. doi:10.1145/3570361.3592530
- [22] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. 2020. Single path one-shot neural architecture search with uniform sampling. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*. Springer, 544–560.
- [23] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. 2018. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (ToG)* 37, 6 (2018), 1–15.
- [24] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)* 42, 4 (2023), 1–14.
- [25] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. 2017. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)* 36, 4 (2017), 1–13.
- [26] Joo Chan Lee, Jong Hwan Ko, and Eunbyung Park. 2025. Optimized Minimal 3D Gaussian Splatting. arXiv:2503.16924 [cs.CV] <https://arxiv.org/abs/2503.16924>
- [27] Joo Chan Lee, Daniel Rho, Xiangyu Sun, Jong Hwan Ko, and Eunbyung Park. 2024. Compact 3D Gaussian Representation for Radiance Field. *CVPR* (2024).
- [28] Ke Li, Reinhard Bacher, Susanne Schmidt, Wim Leemans, and Frank Steinicke. 2024. Reality Fusion: Robust Real-time Immersive Mobile Robot Teleoperation with Volumetric Visual Data Fusion. arXiv preprint arXiv:2408.01225 (2024).
- [29] Lingzhi Li, Zhen Shen, Zhongshu Wang, Li Shen, and Liefeng Bo. 2023. Compressing volumetric radiance fields to 1 mb. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4222–4231.
- [30] Sicheng Li, Hao Li, Yiyi Liao, and Lu Yu. 2024. NeRFCodec: Neural Feature Compression Meets Neural Radiance Fields for Memory-Efficient Scene Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21274–21283.
- [31] Tianye Li, Mira Slavcheva, Michael Zollhöfer, Simon Green, Christoph Lassner, Changli Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, and Zhao Yang Lv. 2022. Neural 3D Video Synthesis From Multi-View Video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 5521–5531.
- [32] Xiaojie Li, Chenghao Gu, Shuzhao Xie, Yunpeng Bai, Weixiang Zhang, and Zhi Wang. 2024. Tuning-Free Visual Customization via View Iterative Self-Attention Control. arXiv:2406.06258 [cs.CV]
- [33] Lei Liu, Zhenghao Chen, and Dong Xu. 2024. HEMGS: A Hybrid Entropy Model for 3D Gaussian Splatting Data Compression. CoRR abs/2411.18473 (2024). arXiv:2411.18473 doi:10.48550/ARXIV.2411.18473
- [34] Weihang Liu, Xue Xian Zheng, Jingyi Yu, and Xin Lou. 2024. Content-Aware Radiance Fields: Aligning Model Complexity with Scene Intricacy Through Learned Bitwidth Quantization. In *European Conference on Computer Vision*. Springer, 239–256.
- [35] Xiangrui Liu, Xinju Wu, Pingping Zhang, Shiqi Wang, Zhu Li, and Sam Kwong. 2024. CompGS: Efficient 3D Scene Representation via Compressed Gaussian Splatting. In *Proceedings of the 32nd ACM International Conference on Multimedia*.
- [36] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. 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.
- [37] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. 2019. Practical Full Resolution Learned Lossless Image Compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [38] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2021), 99–106.
- [39] Wieland Morgenstern, Florian Barthel, Anna Hilsman, and Peter Eisert. 2024. Compact 3D Scene Representation via Self-Organizing Gaussian Grids. In *European Conference on Computer Vision*. Springer.
- [40] KL Navaneet, Kossar Pourahmadi Meibodi, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. 2024. CompGS: Smaller and Faster Gaussian Splatting with Vector Quantization. *ECCV* (2024).
- [41] Simon Niedermayr, Josef Stumpfegger, and Rüdiger Westermann. 2024. Compressed 3d gaussian splatting for accelerated novel view synthesis.
- [42] R. Pajarola, T. Sattler, M. Behr, L. Xie, and S. Akgul. 2019. An overview of ongoing point cloud compression standardization activities: Video-based (V-PCC) and geometry-based (G-PCC). *IEEE Transactions on Circuits and Systems for Video Technology* 29, 9 (2019), 2672–2687. doi:10.1109/TCSVT.2019.2913190
- [43] Panagiotis Papantonakis, Georgios Kopanas, Bernhard Kerbl, Alexandre Lanvin, and George Drettakis. 2024. Reducing the Memory Footprint of 3D Gaussian Splatting. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 7, 1 (May 2024). <https://repo-sam.inria.fr/fungraph/reduced-3dgs/>
- [44] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. 2017. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus,

- S. V. N. Vishwanathan, and Roman Garnett (Eds.). 5099–5108. <https://proceedings.neurips.cc/paper/2017/hash/d8bf84be3800d12f74d8b05e9b89836f-Abstract.html>
- [45] Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Linning Xu, Zhangkai Ni, and Bo Dai. 2024. Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. *arXiv preprint arXiv:2403.17898* (2024).
- [46] J.S. Roy and S.A. Mitchell. 2020. PuLP is an LP modeler written in Python. (2020). <https://github.com/coin-or/pulp>
- [47] Francesco Di Sario, Riccardo Renzulli, Marco Grangetto, Akihiro Sugimoto, and Enzo Tartaglione. 2025. GoDe: Gaussians on Demand for Progressive Level of Detail and Scalable Compression. *arXiv:2501.13558* [cs.CV] <https://arxiv.org/abs/2501.13558>
- [48] Yang Shi, Simone Gasparini, G  r  ldine Morin, and Wei Tsang Ooi. 2025. LapisGS: Layered Progressive 3D Gaussian Splatting for Adaptive Streaming. In *International Conference on 3D Vision, 3DV 2025, Singapore, March 25-28, 2025*. IEEE.
- [49] Cheng Sun, Min Sun, and Hwann-Tzong Chen. 2022. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5459–5469.
- [50] Yuqi Tan, Xiang Liu, Shuzhao Xie, Bin Chen, Shu-Tao Xia, and Zhi Wang. 2024. WATER-GS: Toward Copyright Protection for 3D Gaussian Splatting via Universal Watermarking. *arXiv:2412.05695* [cs.CR]
- [51] Chen Tang, Yuan Meng, Jiacheng Jiang, Shuzhao Xie, Rongwei Lu, Xinzhu Ma, Zhi Wang, and Wenwu Zhu. 2024. Retraining-free model quantization via one-shot weight-coupling learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15855–15865.
- [52] Chen Tang, Kai Ouyang, Zhi Wang, Yifei Zhu, Yaowei Wang, Wen Ji, and Wenwu Zhu. 2022. Mixed-Precision Neural Network Quantization via Learned Layer-wise Importance. In *European Conference on Computer Vision*.
- [53] Zhenyu Tang, Chaoran Feng, Xinhua Cheng, Wangbo Yu, Junwu Zhang, Yuan Liu, Xiaoxiao Long, Wenping Wang, and Li Yuan. 2025. NeuralGS: Bridging Neural Fields and 3D Gaussian Splatting for Compact 3D Representations. *arXiv:2503.23162* [cs.CV] <https://arxiv.org/abs/2503.23162>
- [54] Henan Wang, Hanxin Zhu, Tianyu He, Runsen Feng, Jiajun Deng, Jiang Bian, and Zhibo Chen. 2024. End-to-End Rate-Distortion Optimized 3D Gaussian Representation. In *European Conference on Computer Vision*.
- [55] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. 2019. HAQ: Hardware-Aware Automated Quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [56] Penghao Wang, Zhirui Zhang, Liao Wang, Kaixin Yao, Siyuan Xie, Jingyi Yu, Minye Wu, and Lan Xu. 2024. V<sup>3</sup>: Viewing Volumetric Videos on Mobiles via Streamable 2D Dynamic Gaussians. *CoRR abs/2409.13648* (2024). *arXiv:2409.13648* doi:10.48550/ARXIV.2409.13648
- [57] Yufei Wang, Zhihao Li, Lanqing Guo, Wenhan Yang, Alex C Kot, and Bihan Wen. 2024. ContextGS: Compact 3D Gaussian Splatting with Anchor Level Context Model. In *Advances in neural information processing systems (NeurIPS)*.
- [58] Minye Wu and Tinne Tuytelaars. 2024. Implicit Gaussian Splatting with Efficient Multi-Level Tri-Plane Representation. *arXiv:2408.10041* [cs.CV] <https://arxiv.org/abs/2408.10041>
- [59] Shuzhao Xie, Weixiang Zhang, Chen Tang, Yunpeng Bai, Rongwei Lu, Shijia Ge, and Zhi Wang. 2024. MesonGS: Post-training Compression of 3D Gaussians via Efficient Attribute Transformation. In *European Conference on Computer Vision*. Springer.
- [60] Runyi Yang, Zhenxin Zhu, Zhou Jiang, Baijun Ye, Xiaoxue Chen, Yifei Zhang, Yuntao Chen, Jian Zhao, and Hao Zhao. 2024. Spectrally Pruned Gaussian Fields with Neural Compensation. *arXiv:2405.00676* [cs.CV] <https://arxiv.org/abs/2405.00676>
- [61] Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. 2024. Real-time Photorealistic Dynamic Scene Representation and Rendering with 4D Gaussian Splatting. In *International Conference on Learning Representations (ICLR)*.
- [62] Zeyu Yang, Hongye Yang, Zijie Pan, Xiatian Zhu, and Li Zhang. 2024. Real-time Photorealistic Dynamic Scene Representation and Rendering with 4D Gaussian Splatting. In *ICLR*.
- [63] Wei Yao, Shuzhao Xie, Letian Li, Weixiang Zhang, Zhixin Lai, Shiqi Dai, Ke Zhang, and Zhi Wang. 2025. SD-GS: Structured Deformable 3D Gaussians for Efficient Dynamic Scene Reconstruction. *arXiv:2507.07465* [cs.GR]
- [64] Zhewei Yao, Zhen Dong, Zhangcheng Zheng, Amir Gholami, Jiali Yu, Eric Tan, Leyuan Wang, Qijing Huang, Yida Wang, Michael Mahoney, et al. 2021. Hawq-v3: Dyadic neural network quantization. In *International Conference on Machine Learning*. PMLR, 11875–11886.
- [65] Yu-Ting Zhan, Cheng-Yuan Ho, Hebi Yang, Yi-Hsin Chen, Jui Chiu Chiang, Yu-Lun Liu, and Wen-Hsiao Peng. 2025. CAT-3DGS: A Context-Adaptive Triplane Approach to Rate-Distortion-Optimized 3DGS Compression. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=m3KuuE2ozw>
- [66] Jiakai Zhang, Liao Wang, Xinhao Liu, Fuqiang Zhao, Minzhang Li, Haizhao Dai, Boyuan Zhang, Wei Yang, Lan Xu, and Jingyi Yu. 2022. NeuVV: Neural Volumetric Videos with Immersive Rendering and Editing. *CoRR abs/2202.06088* (2022). *arXiv:2202.06088* <https://arxiv.org/abs/2202.06088>
- [67] Weixiang Zhang, Shuzhao Xie, Chengwei Ren, Shijia Ge, Mingzi Wang, and Zhi Wang. 2025. Enhancing Implicit Neural Representations via Symmetric Power Transformation. In *AAAI*. AAAI Press, 10157–10165.
- [68] Weixiang Zhang, Shuzhao Xie, Chengwei Ren, Siyi Xie, Chen Tang, Shijia Ge, Mingzi Wang, and Zhi Wang. 2025. EVOS: Efficient Implicit Neural Training via EVolutionary Selector. In *CVPR*. Computer Vision Foundation / IEEE, 30472–30482.
- [69] Weixiang Zhang, Wei Yao, Shijia Ge, Shuzhao Xie, Chen Tang, and Zhi Wang. 2025. Expansive Supervision for Neural Radiance Field. *arXiv:2409.08056* [cs.CV] <https://arxiv.org/abs/2409.08056>
- [70] Chen Ziwen, Hao Tan, Kai Zhang, Sai Bi, Fujun Luan, Yicong Hong, Li Fuxin, and Zexiang Xu. 2024. Long-lrm: Long-sequence large reconstruction model for wide-coverage gaussian splats. *arXiv preprint arXiv:2410.12781* (2024).