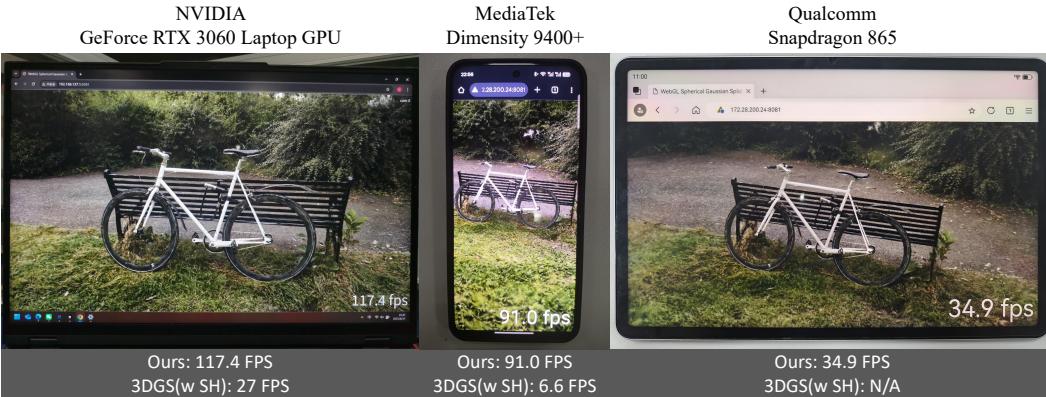


000 001 002 003 004 005 MEGS²: MEMORY-EFFICIENT GAUSSIAN SPLATTING 006 VIA SPHERICAL GAUSSIANS AND UNIFIED PRUNING 007

008 **Anonymous authors**
009
010
011
012
013
014
015
016
017
018
019
020



021 Figure 1: We present **MEGS²**, a memory-efficient framework designed to solve the rendering mem-
022 ory bottleneck of 3D Gaussian Splatting and enable high-quality, real-time rendering on edge de-
023 vices. As demonstrated in our WebGL-based viewer, 3D Gaussian Splatting(3DGS) with Spherical
024 Harmonics (SH) exhibits low frame rates on desktop GPU and fails to run on some mobile plat-
025 forms. In contrast, MEGS² achieves interactive frame rates across all tested devices, significantly
026 expanding the applicability of 3DGS. The detailed result is in Appendix A.3 and A.4.

027 ABSTRACT

030 3D Gaussian Splatting (3DGS) has emerged as a dominant novel-view synthesis
031 technique, but its high memory consumption severely limits its applicability
032 on edge devices. A growing number of 3DGS compression methods have been
033 proposed to make 3DGS more efficient, yet most only focus on storage compres-
034 sion and fail to address the critical bottleneck of rendering memory. To address
035 this problem, we introduce **MEGS²**, a novel memory-efficient framework that
036 tackles this challenge by jointly optimizing two key factors: the total primitive
037 number and the parameters per primitive, achieving unprecedented memory com-
038 pression. Specifically, we **fully** replace the memory-intensive Spherical Har-
039 monics with lightweight, arbitrarily oriented **and prunable** Spherical Gaussian lobes
040 as our color representations. More importantly, we propose a unified soft prun-
041 ing framework that models primitive-number and lobe-number pruning as a sin-
042 gle constrained optimization problem. Experiments show that MEGS² achieves a
043 50% static VRAM reduction and a **40%** rendering VRAM reduction compared to
044 existing methods, while maintaining comparable rendering quality.

045 1 INTRODUCTION

048 Although 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) has been rapidly replacing implicit
049 neural radiance fields (NeRF) (Mildenhall et al., 2020) as the dominant paradigm in neural rendering,
050 thanks to its fast reconstruction, real-time performance, and high-quality outputs, supporting its
051 application across devices with varying constraints has become increasingly important. As a result,
052 compression of 3DGS (Bagdasarian et al., 2025) is gaining significantly more attention, with the
053 goal of enabling real-world use cases on edge devices, such as mobile 3D scanning and previewing,
virtual try-on, and real-time rendering in video games.

054 Nevertheless, existing compression methods only focus on storage compression rather than memory
 055 compression. While the former storage compression speeds up the one-time data transfer of 3DGS
 056 files, the rendering memory dictates whether the 3DGS rendering process can run smoothly on edge
 057 devices. This limits the applicability of 3DGS applications to low-end devices like cell phones.
 058 For example, some methods based on neural compression, vector quantization, and hash grid com-
 059 pression (Chen et al., 2024; Girish et al., 2024; Lee et al., 2024) achieve high storage compression
 060 rates. However, these methods require decoding the full Gaussian parameters from a compressed
 061 state before rendering, even resulting in a larger rendering memory than the 3DGS methods without
 062 compression. Another branch of works, i.e., primitive pruning methods (Zhang et al., 2025b; Fang
 063 & Wang, 2024), is effective at reducing both storage and rendering memory due to the reduced primi-
 064 tives. However, these pruning methods still need to retain a substantial number of 3D Gaussians
 065 for a reasonable rendering quality, which still consumes a large rendering memory. How to further
 066 reduce the rendering memory of 3DGS-based rendering methods remains an open question.
 067

068 In this paper, we observe that the overall memory consumption for 3DGS rendering is intrinsically
 069 tied to two key factors: the total primitive count and the parameters per primitive. Thus, we propose
 070 a novel framework that simultaneously reduces both factors, rather than only optimizing for the
 071 primitive numbers in previous methods Zhang et al. (2025b); Fang & Wang (2024). Specifically,
 072 the rendering VRAM can be divided into a static component and a dynamic component. The static
 073 part relates directly to the total number of Gaussian primitives loaded into the renderer, which is a
 074 product of the primitive count and the parameters per primitive. On the other hand, the dynamic
 075 part, which consists of intermediate data like projected 2D Gaussian parameters and the tile-depth-
 076 gaussian key-value table, is also related to the primitive count in the specific camera viewpoint. This
 077 inherent relationship underscores that to reduce the overall memory footprint, we need to reduce
 078 both the number of primitives and the per-primitive parameter size.
 079

080 **To address the memory bottleneck, we must reduce both primitive counts and per-primitive param-
 081 eters.** We first analyze the limitations of using Spherical Harmonics (SH) for color representation.
 082 While effective for low-frequency lighting, SH functions are inherently global and require many
 083 high-order coefficients to represent localized, high-frequency details like sharp highlights. This re-
 084 sults in low parameter utilization and makes them difficult to compress due to the varying parameter
 085 counts across different degrees. **While SG-Splatting (Wang et al., 2024a) pioneered using Spheri-
 086 cal Gaussians (SG) to mitigate these issues, they relied on a hybrid model combining SG with SH.**
 087 **Advancing in this direction, we introduce fully standalone, arbitrarily-oriented, and prunable SG as**
 088 **a more memory-efficient alternative. Such representation excel at modeling view-dependent signals**
 089 **with very few parameters, and their complexity can be flexibly controlled by the number of lobes.**
 090 **This inherent locality and sparsity make it a convenient and effective choice for compression.**
 091

092 Based on the SG-based representation, we further introduce a novel unified soft pruning frame-
 093 work. We leverage the favorable properties of Spherical Gaussians to dynamically prune redun-
 094 dant lobes for each primitive, thereby compressing the per-primitive parameter count. To achieve
 095 a globally optimal memory footprint, our framework models the two traditionally separate pruning
 096 problems—primitive-count pruning and per-primitive lobe pruning—as a single constrained opti-
 097 mization problem, with the total parameter budget serving as a unified constraint. Extensive ex-
 098 periments demonstrate that the proposed MEGS² achieves an excellent balance between rendering
 099 quality and VRAM efficiency.
 100

101 Overall, our contributions can be summarized as follows:
 102

- 103 **• Advancing beyond the hybrid approach in SG-Splatting (Wang et al., 2024a), we introduce a**
**104 fully standalone color representation by replacing Spherical Harmonics entirely with arbitrarily-
 105 oriented and prunable Spherical Gaussians.** This significantly reduces the per-primitive parameter
 106 count and thus lowers rendering VRAM with minimal impact on quality.
- 107 **• We propose a unified soft pruning framework that models both primitive-count and lobe-count**
108 as a memory-constrained optimization problem, which yields superior performance compared to
109 existing staged or hard-pruning methods.
- 110 **• We achieve unprecedented memory compression for 3DGS, surpassing both vanilla and state-
 111 of-the-art lightweight methods. Our method delivers over an **8×** static VRAM compression and**
112 nearly a **6× rendering VRAM compression compared to vanilla 3DGS, while maintaining or even**
113 improving rendering quality. Furthermore, it still achieves a **2× static VRAM compression and a**
114 40% rendering VRAM reduction over the SOTA method, GaussianSpa, with comparable quality.

108

2 RELATED WORK

109

110 2.1 EVOLUTION OF SPLATTING-BASED SCENE REPRESENTATIONS

111
112 While Neural Radiance Fields (NeRF) (Mildenhall et al., 2020) and their variants (Barron et al.,
113 Müller et al., 2022; Barron et al., 2023) achieved excellent quality in novel view synthesis,
114 their slow rendering speeds precluded real-time applications. 3D Gaussian Splatting (3DGS) (Kerbl
115 et al., 2023) overcame this limitation by introducing a differentiable rasterizer for 3D Gaussians,
116 enabling unprecedented real-time performance with high visual fidelity. The success of 3DGS, how-
117 ever, highlighted its primary challenge: a substantial memory and storage footprint. This has directly
118 motivated the body of work on 3DGS compression and pruning that we discuss subsequently.
119120

2.2 MEMORY ANALYSIS IN 3DGS RENDERING

121
122 In most 3DGS compression studies (Bagdasarian et al., 2025), the rendering memory footprint has
123 not been thoroughly analyzed or compared. Memory usage can be conceptualized into two main
124 components: a static portion, consisting of the total parameters of all loaded Gaussian primitives,
125 and a dynamic portion, representing the runtime overhead. The dynamic overhead, which is highly
126 dependent on a renderer’s implementation, includes the storage of preprocessed 2D Gaussian
127 attributes—an overhead that scales significantly with the number of rendering channels, as seen in
128 multi-channel applications like Feature Splatting (Qiu et al., 2024). Additionally, it includes data
129 structures required by tile-based renderers, an overhead first discussed and optimized by FlashGS
130 (Feng et al., 2024) for large-scale scenes.
131132

2.3 PRUNING TECHNIQUES FOR 3DGS

133
134 Pruning unimportant primitives is a natural and widely explored approach for 3DGS compression.
135 Some work has focused on refining adaptive density control strategies to achieve more efficient
136 representations (Kheradmand et al., 2025; Cheng et al., 2024; Liu et al., 2024a; Pateux et al., 2025;
137 Mallick et al., 2024; Kim et al., 2024). Other research has concentrated on defining better importance
138 metrics to decide which primitives to remove, with notable examples including LP-3DGS (Zhang
139 et al., 2024), mini-splatting (Fang & Wang, 2024), and Reduced3DGS (Papantonakis et al., 2024),
140 the latter of which also involves spherical harmonic pruning. A recent noteworthy development is
141 GaussianSpa (Zhang et al., 2025b), which models pruning as a constrained optimization problem,
142 offering a novel perspective and good compatibility with other pruning techniques. While effective
143 at reducing the number of primitives, these methods generally achieve a limited compression ratio
144 and are therefore often used as an initial step within a larger, integrated compression pipeline.
145146

2.4 OTHER COMPRESSION SCHEMES FOR 3DGS

147
148 Besides pruning, researchers have applied common compression techniques like vector quantization,
149 scalar quantization, neural, and hash grid compression to 3DGS (Girish et al., 2024; Lee et al., 2024;
150 Fan et al., 2024; Lee et al., 2025; Liu et al., 2024b; Navaneet et al., 2024; Xie et al., 2024; Shin et al.,
151 2025). While methods with entropy encoding (Chen et al., 2024; 2025; Liu et al., 2025a; Wang
152 et al., 2024b) achieve the highest storage compression ratios, they are often limited to structured
153 (anchor-based) Gaussian representations (Lu et al., 2023; Ren et al., 2024; Zhang et al., 2025a).
154 These techniques may drastically reduce file size, but they fail to deliver a comparable reduction in
155 rendering memory. This is because they cannot render directly from a compressed state; instead,
156 they must first decode the full Gaussian parameters. The resulting memory footprint is, therefore,
157 never smaller than that required for uncompressed Gaussian primitives, and the decoding process
158 itself, particularly when involving neural networks, can introduce significant additional memory
159 overhead.
160

3 METHOD

161
162 We present MEGS², Memory-Efficient Gaussian Splatting via Spherical GUnified Pruning,
163 a novel Gaussian compression framework designed from a VRAM-centric perspective. Our

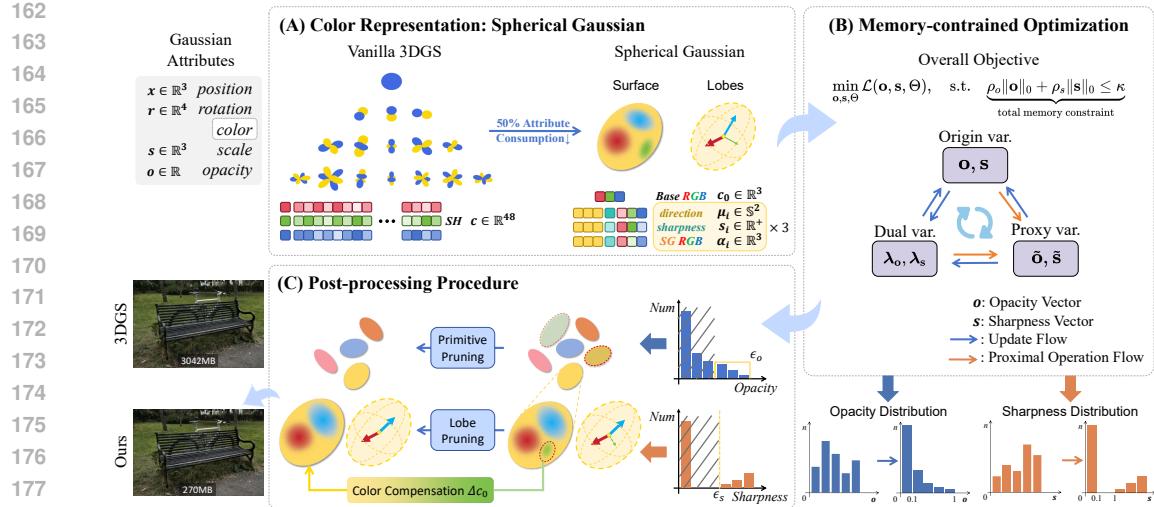


Figure 2: **Overview of our proposed MEGS².** (A) In Section 3.2, we first replace the Spherical Harmonics with [arbitrarily-oriented and prunable](#) Spherical Gaussians. (B) In Section 3.3.1, we formulate the compression as a memory-constrained optimization problem, which is solved using an ADMM-inspired approach (Section 3.3.2). (C) In Section 3.3.3, near-invalid primitives and low-sharpness lobes are removed, and a color compensation term (Eq.12) is introduced to recover the energy of the removed lobes.

method achieves an excellent balance between rendering quality and memory footprint by simultaneously optimizing both the number of Gaussian primitives and the average parameters per primitive.

3.1 PRELIMINARIES FOR 3D GAUSSIAN SPLATTING

3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) represents a scene using a set of 3D Gaussian primitives. Each primitive is defined by a center $\mu \in \mathbb{R}^3$, a covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$, a scalar opacity o , and a view-dependent color model. The color is determined by a function $c(\mathbf{v}) : \mathbb{S}^2 \rightarrow \mathbb{R}^3$, which maps a viewing direction \mathbf{v} to an RGB color. In the original work, this function is modeled using Spherical Harmonics (SH).

To render an image, these 3D primitives are projected onto the 2D image plane. The final pixel color is computed by alpha-blending the primitives that overlap with the pixel in front-to-back order, according to the equation: $C = \sum_i c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)$. Here, the rendering opacity α_i is a product of the primitive's scalar opacity o_i and its 2D Gaussian value at the pixel location.

3.2 SPHERICAL GAUSSIANS

Although the original 3D Gaussian Splatting (3DGS) utilized Spherical Harmonics (SH) for color modeling, as we discussed in Section 1, SH functions suffer from limitations in parameter efficiency and the ability to model local and high-frequency signals.

In contrast, Spherical Gaussians (SG) offer a more compact representation, where the parameter count can be controlled by adjusting the number of lobes. As we demonstrated in Figure 4, a three-lobe SG requires only about half the parameters of a 3rd-order SH while achieving comparable expressive power and superior high-frequency detail capture because, in most cases, we only need to model highlights. Furthermore, as shown in Table 2, SG is more amenable to pruning than SH.

SG was first introduced in real-time rendering to support all-frequency shadows from both point lights and environment lights (Wang et al., 2009). Specifically, a lobe of SG has the form

$$G(\mathbf{v}; \mu, s, a) = a e^{s(\mu \cdot \mathbf{v} - 1)}, \quad (1)$$

where $\mu \in \mathbb{S}^2$ is the unit-length lobe axis, $s \in \mathbb{R}^+$ controls the sharpness, and $a \in \mathbb{R}^3$ is the RGB amplitude vector, with $\mathbf{v} \in \mathbb{S}^2$ being the viewing direction.

216 We use the sum of all lobes of SG for color modeling and introduce a diffuse term to model the
 217 direction-independent component. The view-dependent color $c(\mathbf{v})$ is thus computed as:
 218

$$219 \quad 220 \quad 221 \quad c(\mathbf{v}) = c_0 + \sum_{i=1}^n G(\mathbf{v}; \mu_i, s_i, a_i). \quad (2)$$

222 where $c(\mathbf{v}) \in \mathbb{R}^3$ is the final color from the viewing direction $\mathbf{v} \in \mathbb{S}^2$, $c_0 \in \mathbb{R}^3$ is the direction-
 223 independent diffuse color, and the sum is over n all lobes of SG, with $\mu_i \in \mathbb{S}^2$ being the lobe axis,
 224 $s_i \in \mathbb{R}^+$ the sharpness, and $a_i \in \mathbb{R}^3$ the RGB amplitude for the i -th spherical lobe.

225 **Choice of arbitrarily-oriented SG lobes** It is particularly notable that the lobe axes of different
 226 SG are not constrained to be orthogonal, nor is any regularization term introduced to enforce or-
 227 thogonality. Instead, each SG lobe is allowed to have an arbitrary direction. This flexibility grants
 228 the SG model a higher degree of freedom, leading to greater representation capability compared to
 229 models with fixed orthogonal axes. As shown in Figure 4 (in Appendix), SG-Splatting (Wang et al.,
 230 2024a) with fixed orthogonal axes demonstrated significant rendering performance degradation, fur-
 231 ther underscoring the importance of supporting arbitrarily oriented SG lobes.

232 3.3 UNIFIED SOFT-PRUNING FRAMEWORK

233 Benefiting from the property of Spherical Gaussians to control the number of parameters with the
 234 number of lobes, they provide an ideal object for pruning. Given that most Gaussian primitives
 235 in a scene require only a few lobes to effectively model their color, and that the number of Gaus-
 236 sian primitives is itself often redundant, we propose a novel unified soft pruning framework. In
 237 the framework, we first redefine the pruning of both Gaussian primitives and spherical lobes as a
 238 unified optimization problem with a total memory-overhead constraint (Section 3.3.1). To handle
 239 the non-differentiable nature of this constraint, we introduce an ADMM-inspired (Boyd et al., 2011)
 240 algorithm to efficiently solve the problem (Section 3.3.2). After optimization, the model enters a
 241 post-processing procedure where the final primitives and spherical lobes are removed, with render-
 242 ing quality subsequently recovered through a color compensation strategy and minor fine-tuning
 243 (Section 3.3.3).

244 3.3.1 PROBLEM FORMULATION

245 We unify primitive count pruning and spherical lobe pruning into a memory-constrained optimiza-
 246 tion problem. The L_0 norm of the opacity vector represents the number of active primitives, as
 247 primitives with zero opacity do not contribute to rendering. Similarly, the L_0 norm of the sharp-
 248 ness vector denotes the number of active spherical lobes, since lobes with zero sharpness exhibit
 249 no view-dependent effect and their color can be directly added to the diffuse term c_0 . Building
 250 upon GaussianSpa (Zhang et al., 2025b)’s sparsification framework and our analysis, we extend the
 251 pruning objective from primitive count optimization to total memory budget control, formalized as:
 252

$$253 \quad \min_{\mathbf{o}, \mathbf{s}, \Theta} \mathcal{L}(\mathbf{o}, \mathbf{s}, \Theta), \quad \text{s.t.} \quad \underbrace{\rho_o \|\mathbf{o}\|_0 + \rho_s \|\mathbf{s}\|_0}_{\text{total memory constraint}} \leq \kappa \quad (3)$$

254 where $\mathbf{o} \in \mathbb{R}^{N \times 1}$ means opacity vector for N Gaussian primitives, $\mathbf{s} \in \mathbb{R}^{N \times 3}$ means flattened
 255 sharpness vector for N Gaussians, $\Theta \in \mathbb{R}^{N \times 13}$ (other Gaussian variables), $\rho_o = 11$ and $\rho_s = 7$
 256 count base parameters for a single Gaussian primitive and a single SG lobe respectively, and κ is the
 257 total parameter budget. $\mathcal{L}(\mathbf{o}, \mathbf{s}, \Theta)$ is the reconstruction loss function.

258 3.3.2 MEMORY-CONSTRAINED OPTIMIZATION

259 In this subsection, we aim to solve the memory-constrained optimization defined in Section 3.3.1.
 260 Due to the introduction of a non-differentiable component (the L_0 norm) into the constraint, we can-
 261 not directly optimize it with regularized stochastic gradient descent. Instead, we adopt an ADMM-
 262 inspired approach by introducing proxy variables $\tilde{\mathbf{o}}$ and $\tilde{\mathbf{s}}$ for \mathbf{o} and \mathbf{s} , respectively. This decomposi-
 263 tion strategy enables us to split the original optimization problem into two tractable subproblems:
 264

$$265 \quad \min_{\mathbf{o}, \mathbf{s}, \Theta} \mathcal{L}(\mathbf{o}, \mathbf{s}, \Theta) + \frac{\delta}{2} (\rho_o \|\mathbf{o} - \tilde{\mathbf{o}} + \lambda_o\|^2 + \rho_s \|\mathbf{s} - \tilde{\mathbf{s}} + \lambda_s\|^2) \quad (4)$$

270 **Algorithm 1** Memory-constrained optimizing271 **Require:**272 Initial parameters: $\Theta^0, \mathbf{o}^0, \mathbf{s}^0$
273 Proxy variables: $\tilde{\mathbf{o}}^0 = \mathbf{o}^0, \tilde{\mathbf{s}}^0 = \mathbf{s}^0$
274 Dual variables: $\lambda_{\mathbf{o}}^0 = \mathbf{0}, \lambda_{\mathbf{s}}^0 = \mathbf{0}$
275 Learning rate η , penalty δ_o, δ_s , budget κ 276 **Ensure:**277 Optimized parameters: $\Theta^K, \mathbf{o}^K, \mathbf{s}^K$ 278 1: **for** $k = 0$ **to** $K - 1$ **do**279 2: **Gradient Step:**280 3: $\Theta^{k+1} \leftarrow \Theta^k - \eta \nabla_{\Theta} \mathcal{L}(\Theta^k, \mathbf{o}^k, \mathbf{s}^k)$
281 4: $\mathbf{o}^{k+1} \leftarrow \mathbf{o}^k - \eta [\nabla_{\mathbf{o}} \mathcal{L}(\Theta^k, \mathbf{o}^k, \mathbf{s}^k) + \delta_o (\mathbf{o}^k - \tilde{\mathbf{o}}^k + \lambda_{\mathbf{o}}^k)]$
282 5: $\mathbf{s}^{k+1} \leftarrow \mathbf{s}^k - \eta [\nabla_{\mathbf{s}} \mathcal{L}(\Theta^k, \mathbf{o}^{k+1}, \mathbf{s}^k) + \delta_s (\mathbf{s}^k - \tilde{\mathbf{s}}^k + \lambda_{\mathbf{s}}^k)]$ 283 6: **Proximal Step:**284 7: $(\tilde{\mathbf{o}}^{k+1}, \tilde{\mathbf{s}}^{k+1}) \leftarrow \text{prox}_h(\mathbf{o}^{k+1} + \lambda_{\mathbf{o}}, \mathbf{s}^{k+1} + \lambda_{\mathbf{s}})$ 285 8: **Dual Update:**286 9: $\lambda_{\mathbf{o}}^{k+1} \leftarrow \lambda_{\mathbf{o}}^k + (\mathbf{o}^{k+1} - \tilde{\mathbf{o}}^{k+1})$ 287 10: $\lambda_{\mathbf{s}}^{k+1} \leftarrow \lambda_{\mathbf{s}}^k + (\mathbf{s}^{k+1} - \tilde{\mathbf{s}}^{k+1})$ 288 11: **end for**289 12: **Return** $\Theta^K, \mathbf{o}^K, \mathbf{s}^K$

290

291

$$\min_{\tilde{\mathbf{o}}, \tilde{\mathbf{s}}} h(\tilde{\mathbf{o}}, \tilde{\mathbf{s}}) + \frac{\delta}{2} (\rho_o \|\mathbf{o} - \tilde{\mathbf{o}} + \lambda_{\mathbf{o}}\|^2 + \rho_s \|\mathbf{s} - \tilde{\mathbf{s}} + \lambda_{\mathbf{s}}\|^2) \quad (5)$$

292 where $h(\cdot, \cdot)$ is an indicator function enforcing the sparsity constraint:

293

$$h(\mathbf{o}, \mathbf{s}) = \begin{cases} 0 & \text{if } \rho_o \|\mathbf{o}\|_0 + \rho_s \|\mathbf{s}\|_0 \leq \kappa, \\ +\infty & \text{otherwise} \end{cases} \quad (6)$$

294

295 and the proxy variables are updated via the proximal operator:

296

$$(\tilde{\mathbf{o}}, \tilde{\mathbf{s}}) = \text{prox}_h(\mathbf{o} + \lambda_{\mathbf{o}}, \mathbf{s} + \lambda_{\mathbf{s}}). \quad (7)$$

297

298 Here, $\lambda_{\mathbf{o}}$ and $\lambda_{\mathbf{s}}$ denote the corresponding Lagrange multipliers for the constraints. The specific
299 algorithm flow can be found in Algorithm 1, and the detailed derivation is in Appendix A.1.1.

300

301 **Proximal Operator Implementation** The proximal operator projects \mathbf{o} and \mathbf{s} onto the constraint
302 $\rho_o \|\mathbf{o}\|_0 + \rho_s \|\mathbf{s}\|_0 \leq \kappa$, compatible with any importance metric. To simplify the design and leverage
303 importance criteria from prior primitive-pruning research, we choose to factorize the projection:

304

$$\tilde{\mathbf{o}} = \text{prox}_h(\mathbf{o} + \lambda_{\mathbf{o}}), \|\tilde{\mathbf{o}}\|_0 < \kappa_{\mathbf{o}} \quad (8)$$

305

$$\tilde{\mathbf{s}} = \text{prox}_h(\mathbf{s} + \lambda_{\mathbf{s}}), \|\tilde{\mathbf{s}}\|_0 < \kappa_{\mathbf{s}} \quad (9)$$

306

307 For opacity projection, Zhang et al. (2025b) has established multiple proximal operators that are
308 selected based on scene characteristics. Our approach reuses these operators (See Appendix A.2).

309

310 For sharpness projection, we project the attribute onto the constraint space by retaining only the κ_s
311 most important spherical lobes. This retention is based on a dynamic range metric, which quantifies
312 the view-dependent color change contributed by the i -th lobe:

313

$$D_i = |\max_{\mathbf{v}}(c_i) - \min_{\mathbf{v}}(c_i)| = |a_i|(1 - e^{-2s_i}) \quad (10)$$

314

315 where a_i is the RGB amplitude for the i -th spherical lobe, and κ_s controls the total number of active
316 lobes. The κ_s elements corresponding to the highest D_i values in $(\mathbf{s} + \lambda_{\mathbf{s}})$ are preserved.

317

318 **Discussion: Sequential vs. Unified Pruning** While sequential pruning tackles memory reduction
319 in a two-stage process—first reducing primitive count, then per-primitive parameters—our unified
320 framework jointly optimizes both factors as a single problem. This approach finds a better trade-off
321 between primitive count and per-primitive complexity, thus avoiding the sub-optimal solutions often
322 found by sequential methods. For a detailed experimental validation, please refer to Section 4.4.2.

324 3.3.3 POST-PROCESSING PROCEDURE
325

326 After completing the constrained optimization process, we obtain a large number of primitives and
327 spherical lobes with near-zero opacity and sharpness values, respectively. As the hard constraints
328 are enforced on the proxy variables rather than the original ones, these near-zero values must still
329 be pruned to achieve an actual memory and storage benefit. We address this with a three-step
330 post-processing strategy: first, we remove Gaussian primitives whose opacity falls below a certain
331 threshold. Then, for spherical lobes, we remove those with negligible sharpness and introduce a
332 simple color compensation method to minimize the average per-view color variation and mitigate
333 performance degradation caused by their removal.

334 This compensation is achieved by finding an optimal term, Δc_0 , that minimizes the integral of the
335 squared color difference over all view directions on the unit sphere:

$$336 \min_{\Delta c_0} \int_{\mathbb{S}^2} ((c_0 + \Delta c_0) - (c_0 + G(\mathbf{v}; \mu_i, s_i, a_i)))^2 d\mathbf{v} \quad (11)$$

337 Solving equation 11 yields an exact compensation term that preserves the color of the removed lobe:
338

$$339 \Delta c_0 = a_i \cdot \frac{1 - e^{-2s_i}}{2s_i} \quad (12)$$

340 The diffuse term of the parent primitive is then updated with this value:
341

$$342 c'_0 = c_0 + \Delta c_0 \quad (13)$$

343 Finally, after removing some primitives and spherical lobes, we continue to fine-tune for a small
344 number of steps to recover the rendering quality. The detailed derivation is in Appendix A.1.2. [Figure 5](#) presents the distribution of the spherical lobes after post-processing procedure across scenes.
345

346 4 EXPERIMENTS
347348 4.1 EXPERIMENTAL SETTINGS
349

350 **Datasets and Metrics** Following 3DGS (Kerbl et al., 2023), we evaluate the rendering performance
351 on real-world datasets, including all indoor scenes and outdoor scenes in Mip-NeRF360
352 (Barron et al., 2022), two scenes from Tanks & Temples (Knapitsch et al., 2017), and two scenes
353 from Deep Blending (Hedman et al., 2018). Additionally, we measure the VRAM overhead during
354 the Gaussian loading (static) and rendering stages on these scenes.

355 **Baselines** On all datasets, we benchmark our method against four categories of baselines: (1)
356 vanilla 3DGS (Kerbl et al., 2023) and SG-Splatting (Wang et al., 2024a); (2) lightweight 3DGS
357 methods based on primitive pruning: LP-3DGS (Zhang et al., 2024), Mini-Splatting (Fang & Wang,
358 2024), [MaskGaussian](#) (Liu et al., 2025b) and GaussianSpa (Zhang et al., 2025b); (3) lightweight
359 3DGS methods based on pruning spherical harmonic coefficients: Reduced3DGS (Papantonakis
360 et al., 2024); (4) 3DGS compression methods utilizing orthogonal techniques such as neural
361 compression and vector quantization: CompactGaussian (Lee et al., 2024), EAGLES (Girish et al.,
362 2024), [LightGaussian](#) (Fan et al., 2024), [LocoGS](#) (Shin et al., 2025) and MesonGS (Xie et al., 2024).
363

364 **Experimental Setup** To ensure a fair comparison, we follow the core training procedures estab-
365 lished by 3DGS (Kerbl et al., 2023) and GaussianSpa (Zhang et al., 2025b). We evaluate rendering
366 quality using PSNR, SSIM, and LPIPS, and measure efficiency via static and rendering VRAM
367 consumption. For a detailed description of our specific implementation, as well as the definitions and
368 measurement procedures for both static and rendering VRAM, please refer to Appendix A.2.
369

370 4.2 MAIN RESULTS AND ANALYSIS
371

372 The quantitative results are summarized in Table 1, comparing our approach with several existing
373 lightweight 3DGS methods. Overall, our method achieves comparable or slightly superior rendering
374 quality to the current state-of-the-art lightweight 3DGS methods, while demonstrating signifi-
375 cantly lower VRAM overhead across all existing methods. Compared to vanilla 3DGS, our method
376 achieves a 0.4 dB PSNR improvement on the DeepBlending dataset and a 0.012 LPIPS reduction on
377

378
 379
 380
 381
Table 1: Quantitative comparison across three datasets. Best results are in red region, second
 382 best are in orange region. Memory (VRAM) values are in MB. HQ denotes the version prioritizing
 383 high rendering quality, and LM denotes the version prioritizing lower VRAM consumption.

Method	Mip-NeRF 360						Tanks&Temples						DeepBlending					
	PSNR↑	SSIM↑	LPIPS↓	VRAM↓	Stat.	Rend.	PSNR↑	SSIM↑	LPIPS↓	VRAM↓	Stat.	Rend.	PSNR↑	SSIM↑	LPIPS↓	VRAM↓	Stat.	Rend.
3DGS	27.48	0.813	0.217	648	1717		23.68	0.849	0.171	370	1021		29.71	0.902	0.242	582	1569	
SG-Splatting	27.27	0.813	0.218	416	1327		23.47	0.840	0.177	229	822		29.57	0.901	0.247	357	983	
Reduced3DGS	27.21	0.810	0.225	191	493		23.51	0.839	0.187	90	252		29.60	0.902	0.248	121	340	
CompactGaussian	27.08	0.798	0.247	267	838		23.32	0.831	0.201	177	469		29.79	0.901	0.258	187	532	
EAGLES	27.23	0.810	0.240	—	—		23.37	0.840	0.200	333	965		29.86	0.910	0.250	510	1550	
LightGaussian	27.13	0.806	0.237	290	640		23.44	0.832	0.202	168	437		—	—	—	—	—	
LocOGS	27.28	0.809	0.231	609	994		23.62	0.846	0.182	458	732		30.05	0.905	0.247	643	1040	
MesonGS-FT	26.98	0.801	0.233	709	1336		23.32	0.837	0.193	424	783		29.51	0.901	0.251	656	1202	
LP-3DGS	27.12	0.805	0.239	420	1239		23.41	0.834	0.198	251	769		—	—	—	—	—	
Mini-Splatting	27.40	0.821	0.219	125	477		23.45	0.841	0.186	72	253		30.05	0.909	0.254	89	324	
MaskGaussian	27.43	0.811	0.227	271	799		23.72	0.847	0.181	132	517		29.69	0.907	0.244	156	501	
GaussianSpa	27.56	0.824	0.215	115	448		23.73	0.857	0.162	106	336		30.00	0.912	0.239	104	372	
Ours(HQ)	27.54	0.824	0.209	55	265		23.45	0.853	0.159	51	211		30.17	0.912	0.233	54	243	
Ours(LM)	27.21	0.814	0.227	40	224		23.27	0.851	0.167	37	163		30.01	0.908	0.246	33	193	

398
 399 the Tanks & Temples dataset. Furthermore, it achieves more than an $8\times$ compression rate for static
 400 VRAM and nearly a $6\times$ compression rate for rendering VRAM across all datasets. Even when com-
 401 pared to the current state-of-the-art lightweight 3DGS method, GaussianSpa (Zhang et al., 2025b),
 402 we still achieve a nearly $2\times$ compression rate for static VRAM on the Mip-NeRF360 dataset and
 403 reduce rendering VRAM by approximately 40% with comparable rendering quality.

404 **Analysis on Existing Compression Methods** Many recent methods achieve high storage com-
 405 pression but fail to reduce VRAM. Techniques based on hash grid compression, vector quantization
 406 or neural networks (e.g., CompactGaussian (Lee et al., 2024), EAGLES (Girish et al., 2024)) must
 407 first decompress their parameters into a renderable state. This requirement, coupled with the decod-
 408 ing process itself, results in a VRAM footprint that can exceed even vanilla 3DGS (Appendix, Tab.
 409 3). Similarly, HAC++ (Chen et al., 2025), a SOTA method designed specifically for anchor-based
 410 3DGS, is not a general solution and offers minimal VRAM reduction for the same reason. In con-
 411 trast, our method achieves superior perceptual quality (SSIM/LPIPS) with 50-60% less VRAM and
 412 a 1.5-1.7x rendering speedup.

413 **Analysis on Pruning-based Methods** Pruning-based methods (Zhang et al., 2025b; Fang &
 414 Wang, 2024; Zhang et al., 2024) effectively reduce VRAM but face a bottleneck, as aggressively
 415 reducing primitive count degrades quality. Our method overcomes this by not only pruning primitives
 416 but also reducing per-primitive costs. When compared to methods that also compress color,
 417 such as Reduced3DGS (Papantonakis et al., 2024) which prunes SH coefficients, our advantage is
 418 twofold: the superior amenability of our Spherical Gaussians to pruning, and our more effective
 419 unified soft pruning framework, as validated in Section 4.4.

420 421 4.3 QUALITATIVE RESULTS

423 Figure 3 presents a qualitative comparison of MEGS² against baselines, including GaussianSpa and
 424 3DGS on various scenes. In Bicycle and Truck, our method accurately recovers specular reflections
 425 on smooth surfaces and mirrors, details that other methods fail to fully capture. In Bonsai, MEGS²
 426 faithfully captures high-contrast lighting in both contours and brightness. Furthermore, MEGS²
 427 delivers a cleaner and more complete reconstruction in Playroom. Notably, MEGS² achieves these
 428 high-quality results with a VRAM footprint of only 50-60% compared to GaussianSpa, demon-
 429 strating a substantial enhancement in memory efficiency. These visual results indicate that our Spherical
 430 Gaussian based representation has a stronger capability for fitting local view-dependent signals than
 431 Spherical Harmonics, leading to sharper and more photorealistic images. [More qualitative results](#)
 432 [are in Appendix, Figure 7](#).

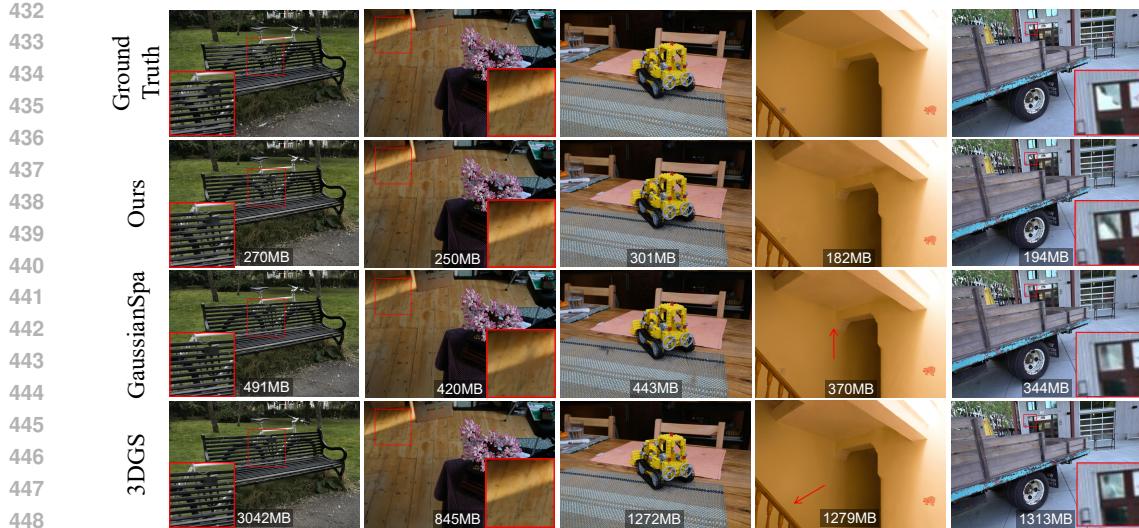


Figure 3: Qualitative results on the Bicycle, Bonsai, Kitchen, Playroom and Truck scenes comparing to previous methods and the corresponding ground truth images from test views. The rendering VRAM consumption for the corresponding method is annotated at the bottom of each image.

4.4 ABLATION STUDIES

In this section, we conduct a series of ablation studies to validate the effectiveness of our proposed designs and answer two key questions: 1. Is the replacement of spherical harmonics with spherical Gaussians necessary? And why not use the orthogonal-axis spherical variants ([fixed-axis SG](#))? (Section 4.4.1) 2. Is our proposed unified soft pruning framework effective? And does it offer a significant advantage over simply combining existing pruning methods? (Section 4.4.2)

4.4.1 ABLATION ON COLOR MODELING

To validate our color model, we replace the spherical harmonics (SH) in 3DGS with our arbitrarily-oriented spherical Gaussians (SG) while keeping training settings identical. Our SG model achieves a superior quality-VRAM trade-off (see Figure 4). Furthermore, we show that the fixed-axis SG from SG-Splatting (Wang et al., 2024a) struggles to capture complex view-dependent effects, leading to a sharp quality drop (about 0.6 dB PSNR). This rigidity also makes it incompatible with our lobe-pruning strategy.

4.4.2 ABLATION ON PRUNING STRATEGIES

We evaluate our unified soft pruning framework against several ablations and baselines, with results in Table 2. First, replacing our soft pruning with a hard-pruning variant (soft \rightarrow hard) that only optimizes opacity results in a significant performance drop. Second, performing primitive and lobe pruning sequentially (unified \rightarrow sequential) is inferior to our joint optimization, confirming the advantage of a unified approach. We also found that removing our color compensation harms performance (w/o color comp.). Finally, we test two naive baselines. Combining GaussianSpa with Reduced3DGS leads to severe quality degradation, showing that existing SH pruning metrics fail on fewer primitives. Simply replacing SH with SG in GaussianSpa is also outperformed by our method. This demonstrates that our pruning strategy not only reduces memory but also acts as a regularizer, improving rendering quality. [Visual comparison are in Appendix, Figure 9.](#)

Table 2: **Ablation studies for different pruning strategies** on Mip-NeRF360 dataset.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	VRAM \downarrow
Spa+Redu.	26.05	0.776	0.280	402
Spa(SH \rightarrow SG)	27.01	0.808	0.230	339
soft \rightarrow hard	27.23	0.814	0.228	288
unified \rightarrow seq.	27.33	0.818	0.222	328
w/o color comp.	27.46	0.822	0.213	265
full model	27.54	0.824	0.209	265

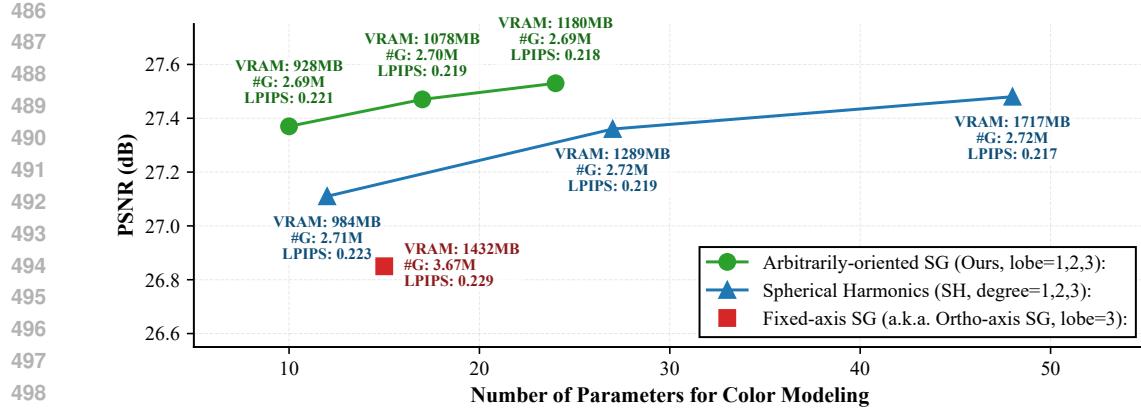
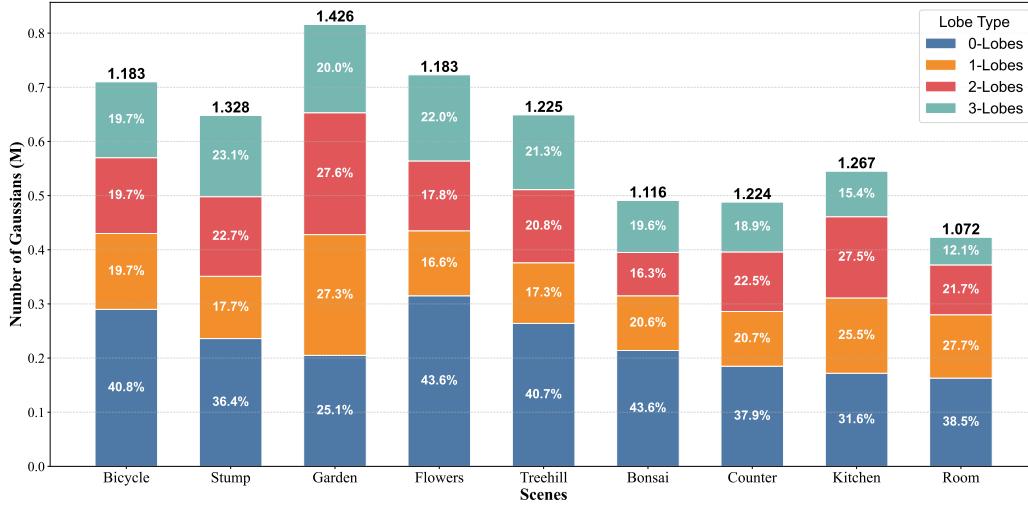


Figure 4: Comparison of different 3DGS color representations on Mip-NeRF360 dataset.

Figure 5: **Distribution of Gaussian lobes across scenes.** We report the distribution of Gaussians using 0 to 3 lobes for each scene in the Mip-NeRF 360 dataset. The value above each bar represents the average lobe count per Gaussian.

5 CONCLUSION

We introduced MEGS², a framework that addresses the 3D Gaussian Splatting memory bottleneck by jointly optimizing primitive count and per-primitive parameters. Our approach combines a lightweight spherical Gaussian color representation with a unified soft pruning method to achieve state-of-the-art memory compression. MEGS² shifts the focus of 3DGS compression towards rendering VRAM efficiency, paving the way for high-quality rendering on edge devices.

Limitation Our method focuses on compressing static memory for broad, renderer-agnostic applicability. We leave the optimization of implementation-specific dynamic VRAM for future work, and note that the model’s performance on highly complex highlights warrants further investigation.

ETHICS STATEMENT

Our work presents no direct ethical concerns. The primary application of our method is for novel view synthesis from captured data.

540

REPRODUCIBILITY STATEMENT

541

542

To ensure reproducibility: (1) While not included with this submission, our full project (including the WebGL renderer, training, and evaluation scripts) will be released on GitHub upon publication. (2) All experimental details, including our VRAM measurement protocol, are provided in Appendix A.2. (3) Our data and preprocessing follow the official 3DGS implementation. (4) Derivations and proofs for our core algorithms are in Appendix A.1.1 and A.1.2.

547

548

REFERENCES

549

550

Milena T. Bagdasarian, Paul Knoll, Yi-Hsin Li, Florian Barthel, Anna Hilsmann, Peter Eisert, and Wieland Morgenstern. 3dgs.zip: A survey on 3d gaussian splatting compression methods, 2025. URL <https://arxiv.org/abs/2407.09510>.

551

552

Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields, 2021. URL <https://arxiv.org/abs/2103.13415>.

553

554

Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022.

555

556

Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields, 2023. URL <https://arxiv.org/abs/2304.06706>.

557

558

Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, January 2011. ISSN 1935-8237. doi: 10.1561/2200000016. URL <https://doi.org/10.1561/2200000016>.

559

560

Yihang Chen, Qianyi Wu, Weiyao Lin, Mehrtash Harandi, and Jianfei Cai. Hac: Hash-grid assisted context for 3d gaussian splatting compression, 2024. URL <https://arxiv.org/abs/2403.14530>.

561

562

Yihang Chen, Qianyi Wu, Weiyao Lin, Mehrtash Harandi, and Jianfei Cai. Hac++: Towards 100x compression of 3d gaussian splatting, 2025. URL <https://arxiv.org/abs/2501.12255>.

563

564

Kai Cheng, Xiaoxiao Long, Kaizhi Yang, Yao Yao, Wei Yin, Yuexin Ma, Wenping Wang, and Xuejin Chen. Gaussianpro: 3d gaussian splatting with progressive propagation, 2024. URL <https://arxiv.org/abs/2402.14650>.

565

566

Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, Dejia Xu, and Zhangyang Wang. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps, 2024. URL <https://arxiv.org/abs/2311.17245>.

567

568

Guangchi Fang and Bing Wang. Mini-splatting: Representing scenes with a constrained number of gaussians, 2024. URL <https://arxiv.org/abs/2403.14166>.

569

570

Guofeng Feng, Siyan Chen, Rong Fu, Zimu Liao, Yi Wang, Tao Liu, Zhilin Pei, Hengjie Li, Xingcheng Zhang, and Bo Dai. Flashgs: Efficient 3d gaussian splatting for large-scale and high-resolution rendering, 2024. URL <https://arxiv.org/abs/2408.07967>.

571

572

Sharath Girish, Kamal Gupta, and Abhinav Shrivastava. Eagles: Efficient accelerated 3d gaussians with lightweight encodings, 2024. URL <https://arxiv.org/abs/2312.04564>.

573

574

Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brossow. Deep blending for free-viewpoint image-based rendering. 37(6):257:1–257:15, 2018.

575

576

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023. URL <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>.

594 Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Jeff Tseng, Hossam Isack, Ab-
 595 hishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3d gaussian splatting as markov chain
 596 monte carlo, 2025. URL <https://arxiv.org/abs/2404.09591>.

597

598 Sieun Kim, Kyungjin Lee, and Youngki Lee. Color-cued efficient densification method for 3d gaus-
 599 sian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
 600 *Recognition (CVPR) Workshops*, pp. 775–783, June 2024.

601 Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking
 602 large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4), 2017.

603

604 Kevin Kwok and Haoyang Ye. splat, 2023. URL <https://github.com/antimatter15/splat>.

605

606 Joo Chan Lee, Daniel Rho, Xiangyu Sun, Jong Hwan Ko, and Eunbyung Park. Compact 3d gaussian
 607 representation for radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision*
 608 *and Pattern Recognition (CVPR)*, pp. 21719–21728, 2024.

609

610 Soonbin Lee, Fangwen Shu, Yago Sanchez, Thomas Schierl, and Cornelius Hellge. Compression
 611 of 3d gaussian splatting with optimized feature planes and standard video codecs, 2025. URL
 612 <https://arxiv.org/abs/2501.03399>.

613

614 Lei Liu, Zhenghao Chen, Wei Jiang, Wei Wang, and Dong Xu. Hemgs: A hybrid entropy model
 615 for 3d gaussian splatting data compression, 2025a. URL <https://arxiv.org/abs/2411.18473>.

616

617 Rong Liu, Rui Xu, Yue Hu, Meida Chen, and Andrew Feng. Atomgs: Atomizing gaussian splatting
 618 for high-fidelity radiance field, 2024a. URL <https://arxiv.org/abs/2405.12369>.

619

620 Xiangrui Liu, Xinju Wu, Pingping Zhang, Shiqi Wang, Zhu Li, and Sam Kwong. Compgs: Efficient
 621 3d scene representation via compressed gaussian splatting, 2024b. URL <https://arxiv.org/abs/2404.09458>.

622

623 Yifei Liu, Zhihang Zhong, Yifan Zhan, Sheng Xu, and Xiao Sun. Maskgaussian: Adaptive 3d
 624 gaussian representation from probabilistic masks, 2025b. URL <https://arxiv.org/abs/2412.20522>.

625

626 Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-
 627 gs: Structured 3d gaussians for view-adaptive rendering, 2023. URL <https://arxiv.org/abs/2312.00109>.

628

629 Saswat Subhajyoti Mallick, Rahul Goel, Bernhard Kerbl, Francisco Vicente Carrasco, Markus Stein-
 630 berger, and Fernando De La Torre. Taming 3dgs: High-quality radiance fields with limited re-
 631 sources, 2024. URL <https://arxiv.org/abs/2406.15643>.

632

633 Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and
 634 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis, 2020. URL
 635 <https://arxiv.org/abs/2003.08934>.

636

637 Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics prim-
 638 itives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):1–15, July
 639 2022. ISSN 1557-7368. doi: 10.1145/3528223.3530127. URL <http://dx.doi.org/10.1145/3528223.3530127>.

638

639

640 KL Navaneet, Kossar Pourahmadi Meibodi, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash.
 641 Compgs: Smaller and faster gaussian splatting with vector quantization, 2024. URL <https://arxiv.org/abs/2311.18159>.

642

643

644 Panagiotis Papantonakis, Georgios Kopanas, Bernhard Kerbl, Alexandre Lanvin, and George Dret-
 645 takis. Reducing the memory footprint of 3d gaussian splatting. *Proceedings of the ACM on*
 646 *Computer Graphics and Interactive Techniques*, 7(1), May 2024. URL https://repo-sam.inria.fr/fungraph/reduced_3dgs/.

647

648 Stéphane Pateux, Matthieu Gendrin, Luce Morin, Théo Ladune, and Xiaoran Jiang. Bogauss: Better
 649 optimized gaussian splatting, 2025. URL <https://arxiv.org/abs/2504.01844>.
 650

651 Ri-Zhao Qiu, Ge Yang, Weijia Zeng, and Xiaolong Wang. Feature splatting: Language-driven
 652 physics-based scene synthesis and editing, 2024. URL <https://arxiv.org/abs/2404.01223>.
 653

654 Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Lining Xu, Zhangkai Ni, and Bo Dai. Octree-gs:
 655 Towards consistent real-time rendering with lod-structured 3d gaussians, 2024. URL <https://arxiv.org/abs/2403.17898>.
 656

657 Seungjoo Shin, Jaesik Park, and Sunghyun Cho. Locality-aware gaussian compression for fast and
 658 high-quality rendering, 2025. URL <https://arxiv.org/abs/2501.05757>.
 659

660 Jiaping Wang, Peiran Ren, Minmin Gong, John Snyder, and Baining Guo. All-frequency rendering
 661 of dynamic, spatially-varying reflectance. *ACM Trans. Graph.*, 28(5):1–10, December 2009.
 662 ISSN 0730-0301. doi: 10.1145/1618452.1618479. URL <https://doi.org/10.1145/1618452.1618479>.
 663

664 Yiwen Wang, Siyuan Chen, and Ran Yi. Sg-splatting: Accelerating 3d gaussian splatting with
 665 spherical gaussians, 2024a. URL <https://arxiv.org/abs/2501.00342>.
 666

667 Yufei Wang, Zhihao Li, Lanqing Guo, Wenhan Yang, Alex C. Kot, and Bihan Wen. Contextgs:
 668 Compact 3d gaussian splatting with anchor level context model, 2024b. URL <https://arxiv.org/abs/2405.20721>.
 669

670 Shuzhao Xie, Weixiang Zhang, Chen Tang, Yunpeng Bai, Rongwei Lu, Shijia Ge, and Zhi Wang.
 671 Mesongs: Post-training compression of 3d gaussians via efficient attribute transformation, 2024.
 672 URL <https://arxiv.org/abs/2409.09756>.
 673

674 Jiahui Zhang, Fangneng Zhan, Ling Shao, and Shijian Lu. Sogs: Second-order anchor for advanced
 675 3d gaussian splatting, 2025a. URL <https://arxiv.org/abs/2503.07476>.
 676

677 Yangming Zhang, Wenqi Jia, Wei Niu, and Miao Yin. Gaussianspa: An "optimizing-sparsifying"
 678 simplification framework for compact and high-quality 3d gaussian splatting. In *Proceedings of
 679 the Computer Vision and Pattern Recognition Conference (CVPR)*, pp. 26673–26682, June 2025b.
 680

681 Zhaoliang Zhang, Tianchen Song, Yongjae Lee, Li Yang, Cheng Peng, Rama Chellappa, and Deliang
 682 Fan. Lp-3dgs: Learning to prune 3d gaussian splatting, 2024. URL <https://arxiv.org/abs/2405.18784>.
 683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A APPENDIX**
703704 **A.1 ALGORITHMIC DERIVATION**
705706 **A.1.1 DERIVATION FOR MEMORY-CONSTRAINED OPTIMIZING**
707

708 This appendix provides the detailed derivation of Algorithm 1 presented in the main text. The
709 algorithm solves the memory-constrained optimization problem with L_0 norm constraints using an
710 ADMM-inspired approach.

711 **Problem Formulation** Consider the constrained optimization problem:
712

$$\begin{aligned} 713 \min_{\mathbf{o}, \mathbf{s}, \Theta} \quad & \mathcal{L}(\mathbf{o}, \mathbf{s}, \Theta) \\ 714 \text{s.t.} \quad & \rho_o \|\mathbf{o}\|_0 + \rho_s \|\mathbf{s}\|_0 \leq \kappa \end{aligned} \quad (\text{A.1})$$

715 where \mathcal{L} is the loss function, $\|\cdot\|_0$ denotes the L_0 norm (count of non-zero elements), and κ represents the memory budget.
716

717 **Augmented Lagrangian Formulation** Introduce proxy variables $\tilde{\mathbf{o}}, \tilde{\mathbf{s}}$ to reformulate the problem:
718

$$\begin{aligned} 719 \min_{\mathbf{o}, \mathbf{s}, \Theta, \tilde{\mathbf{o}}, \tilde{\mathbf{s}}} \quad & \mathcal{L}(\mathbf{o}, \mathbf{s}, \Theta) \\ 720 \text{s.t.} \quad & \mathbf{o} = \tilde{\mathbf{o}}, \quad \mathbf{s} = \tilde{\mathbf{s}} \\ 721 \quad & \rho_o \|\tilde{\mathbf{o}}\|_0 + \rho_s \|\tilde{\mathbf{s}}\|_0 \leq \kappa \end{aligned} \quad (\text{A.2})$$

722 The augmented Lagrangian function is constructed as:
723

$$\begin{aligned} 724 \mathcal{L}_\delta(\mathbf{o}, \mathbf{s}, \Theta, \tilde{\mathbf{o}}, \tilde{\mathbf{s}}, \boldsymbol{\lambda}_o, \boldsymbol{\lambda}_s) = & \mathcal{L}(\mathbf{o}, \mathbf{s}, \Theta) \\ 725 \quad & + \boldsymbol{\lambda}_o^\top (\mathbf{o} - \tilde{\mathbf{o}}) + \boldsymbol{\lambda}_s^\top (\mathbf{s} - \tilde{\mathbf{s}}) \\ 726 \quad & + \frac{\delta}{2} (\rho_o \|\mathbf{o} - \tilde{\mathbf{o}}\|^2 + \rho_s \|\mathbf{s} - \tilde{\mathbf{s}}\|^2) \end{aligned} \quad (\text{A.3})$$

727 where $\boldsymbol{\lambda}_o$ and $\boldsymbol{\lambda}_s$ are Lagrange multipliers, and $\delta > 0$ is the penalty parameter.
728

729 **ADMM Alternating Minimization** The ADMM framework decomposes the optimization into
730 three subproblems:
731

732 **Primal Variables Update** With proxy variables and multipliers fixed, update the primal variables:
733

$$\begin{aligned} 734 (\Theta^{k+1}, \mathbf{o}^{k+1}, \mathbf{s}^{k+1}) = & \arg \min_{\Theta, \mathbf{o}, \mathbf{s}} \mathcal{L}(\mathbf{o}, \mathbf{s}, \Theta) \\ 735 \quad & + \frac{\delta}{2} [\rho_o \|\mathbf{o} - \tilde{\mathbf{o}}^k + \boldsymbol{\lambda}_o^k\|^2 + \rho_s \|\mathbf{s} - \tilde{\mathbf{s}}^k + \boldsymbol{\lambda}_s^k\|^2] \end{aligned} \quad (\text{A.4})$$

736 Using gradient descent approximation:
737

$$\begin{aligned} 738 \Theta^{k+1} = & \Theta^k - \eta \nabla_\Theta \mathcal{L}(\Theta^k, \mathbf{o}^k, \mathbf{s}^k) \\ 739 \mathbf{o}^{k+1} = & \mathbf{o}^k - \eta [\nabla_\mathbf{o} \mathcal{L}(\Theta^k, \mathbf{o}^k, \mathbf{s}^k) + \delta_o (\mathbf{o}^k - \tilde{\mathbf{o}}^k + \boldsymbol{\lambda}_o^k)] \\ 740 \mathbf{s}^{k+1} = & \mathbf{s}^k - \eta [\nabla_\mathbf{s} \mathcal{L}(\Theta^k, \mathbf{o}^{k+1}, \mathbf{s}^k) + \delta_s (\mathbf{s}^k - \tilde{\mathbf{s}}^k + \boldsymbol{\lambda}_s^k)] \end{aligned} \quad (\text{A.5})$$

741 where $\delta_o = \delta \rho_o, \delta_s = \delta \rho_s$.
742

743 **Proxy Variables Update** With primal variables and multipliers fixed, update the proxy variables:
744

$$745 (\tilde{\mathbf{o}}^{k+1}, \tilde{\mathbf{s}}^{k+1}) = \arg \min_{\tilde{\mathbf{o}}, \tilde{\mathbf{s}}} h(\tilde{\mathbf{o}}, \tilde{\mathbf{s}}) + \frac{\delta}{2} [\rho_o \|\mathbf{o}^{k+1} - \tilde{\mathbf{o}} + \boldsymbol{\lambda}_o^k\|^2 + \rho_s \|\mathbf{s}^{k+1} - \tilde{\mathbf{s}} + \boldsymbol{\lambda}_s^k\|^2] \quad (\text{A.6})$$

756 where $h(\tilde{\mathbf{o}}, \tilde{\mathbf{s}})$ is the indicator function enforcing the sparsity constraint:
 757

$$758 \quad 759 \quad h(\tilde{\mathbf{o}}, \tilde{\mathbf{s}}) = \begin{cases} 0 & \text{if } \rho_o \|\tilde{\mathbf{o}}\|_0 + \rho_s \|\tilde{\mathbf{s}}\|_0 \leq \kappa \\ +\infty & \text{otherwise} \end{cases} \quad 760 \quad (\text{A.7})$$

761 The solution is given by the proximal operator:
 762

$$763 \quad 764 \quad (\tilde{\mathbf{o}}^{k+1}, \tilde{\mathbf{s}}^{k+1}) = \mathbf{prox}_h(\mathbf{o}^{k+1} + \boldsymbol{\lambda}_{\mathbf{o}}^k, \mathbf{s}^{k+1} + \boldsymbol{\lambda}_{\mathbf{s}}^k) \quad (\text{A.8})$$

765 **Dual Variables Update** Update the Lagrange multipliers:
 766

$$767 \quad 768 \quad \boldsymbol{\lambda}_{\mathbf{o}}^{k+1} = \boldsymbol{\lambda}_{\mathbf{o}}^k + (\mathbf{o}^{k+1} - \tilde{\mathbf{o}}^{k+1}) \quad (\text{A.9})$$

$$769 \quad \boldsymbol{\lambda}_{\mathbf{s}}^{k+1} = \boldsymbol{\lambda}_{\mathbf{s}}^k + (\mathbf{s}^{k+1} - \tilde{\mathbf{s}}^{k+1})$$

771 A.1.2 DERIVATION FOR COLOR COMPENSATION

$$772 \quad 773 \quad \min_{\Delta c_0} \int_{\mathbb{S}^2} ((c_0 + \Delta c_0) - (c_0 + G(\mathbf{v}; \boldsymbol{\mu}_i, s_i, a_i)))^2 d\mathbf{v} \quad (\text{A.10})$$

774 where $G(\mathbf{v}; \boldsymbol{\mu}, s, a) = ae^{s(\boldsymbol{\mu} \cdot \mathbf{v} - 1)}$, $\boldsymbol{\mu} \in \mathbb{S}^2$ is the unit-length lobe axis, $s \in \mathbb{R}^+$ controls the sharpness, and $a \in \mathbb{R}^3$ is the RGB amplitude vector, with $\mathbf{v} \in \mathbb{S}^2$ being the viewing direction.
 775

776 Simplifying the objective function:
 777

$$778 \quad 779 \quad \min_{\Delta c_0} \int_{\mathbb{S}^2} (\Delta c_0 - G(\mathbf{v}; \boldsymbol{\mu}_i, s_i, a_i))^2 d\mathbf{v} \quad (\text{A.11})$$

780 Expanding the squared term:
 781

$$782 \quad 783 \quad \min_{\Delta c_0} \int_{\mathbb{S}^2} (\Delta c_0^2 - 2\Delta c_0 G(\mathbf{v}; \boldsymbol{\mu}_i, s_i, a_i) + G(\mathbf{v}; \boldsymbol{\mu}_i, s_i, a_i)^2) d\mathbf{v} \quad (\text{A.12})$$

784 Taking derivative with respect to Δc_0 and setting to zero:
 785

$$786 \quad 787 \quad \frac{\partial}{\partial \Delta c_0} = 2 \int_{\mathbb{S}^2} (\Delta c_0 - G(\mathbf{v}; \boldsymbol{\mu}_i, s_i, a_i)) d\mathbf{v} = 0 \quad (\text{A.13})$$

788 Solving for the optimal Δc_0 :
 789

$$790 \quad 791 \quad \Delta c_0 = \frac{\int_{\mathbb{S}^2} G(\mathbf{v}; \boldsymbol{\mu}_i, s_i, a_i) d\mathbf{v}}{\int_{\mathbb{S}^2} d\mathbf{v}} = \frac{1}{4\pi} \int_{\mathbb{S}^2} G(\mathbf{v}; \boldsymbol{\mu}_i, s_i, a_i) d\mathbf{v} \quad (\text{A.14})$$

792 Evaluating the spherical Gaussian integral. Without loss of generality, align $\boldsymbol{\mu}_i$ with the z-axis:
 793

$$794 \quad 795 \quad \int_{\mathbb{S}^2} G(\mathbf{v}; \boldsymbol{\mu}_i, s_i, a_i) d\mathbf{v} = a_i \int_0^{2\pi} \int_0^\pi e^{s_i(\cos \theta - 1)} \sin \theta d\theta d\phi$$

$$796 \quad 797 \quad = 2\pi a_i \int_0^\pi e^{s_i(\cos \theta - 1)} \sin \theta d\theta$$

$$798 \quad \text{(substitute } u = \cos \theta, du = -\sin \theta d\theta\text{)}$$

$$799 \quad 800 \quad = 2\pi a_i \int_{-1}^1 e^{s_i(u-1)} du$$

$$801 \quad 802 \quad = 2\pi a_i \cdot \frac{1 - e^{-2s_i}}{s_i}.$$

803 Substituting back:
 804

$$805 \quad 806 \quad \Delta c_0 = \frac{1}{4\pi} \cdot 2\pi a_i \cdot \frac{1 - e^{-2s_i}}{s_i} = a_i \cdot \frac{1 - e^{-2s_i}}{2s_i} \quad (\text{A.15})$$

807 The diffuse term is updated as:
 808

$$809 \quad c'_0 = c_0 + \Delta c_0 \quad (\text{A.16})$$

810 A.2 EXPERIMENTAL DETAILS
811

812 **Implementation Details** Our experiments are conducted on a single NVIDIA RTX 3090 GPU
813 (24GB). For a fair comparison, we adopt the same hyperparameters and optimizer used by vanilla
814 3DGS (Kerbl et al., 2023) on all datasets. In terms of the hyperparameters for our unified soft
815 pruning framework, we conduct the memory-constrained optimization for a total of 10,000 iterations, ap-
816 plied at an interval of 50 iterations. We set the penalty parameter δ to 0.0005. For the post-processing
817 procedure defined in Section 3.3.3, we set the sharpness threshold to 1 to prune spherical lobes with
818 negligible contribution. Regarding the training strategy, we employ the same densification, pruning,
819 and fine-tuning procedures as GaussianSpa, conducting densification training, sparse training, and
820 fine-tuning for the same number of steps. To ensure the accuracy of our VRAM measurements, we
821 reproduce 3DGS, GaussianSpa (Zhang et al., 2025b), EAGLES (Girish et al., 2024), CompactGaus-
822 sian (Lee et al., 2024), LightGaussian (Fan et al., 2024), LocoGS (Shin et al., 2025), MesonGS (Xie
823 et al., 2024), and Reduced3DGS (Papantonakis et al., 2024) for direct VRAM measurement. For
824 other pruning-based methods, we utilize both the rendering quality metrics and primitive counts re-
825 ported in the GaussianSpa paper. We then estimate their corresponding VRAM consumption based
826 on our own VRAM measurements for 3DGS and GaussianSpa. EAGLES cannot be trained on a
827 single NVIDIA RTX 3090 GPU with 24GB on some scenes in Mip-NeRF360, so we do not report
828 VRAM consumption. For fairness, MesonGS uses the version that continues fine-tuning after com-
829 pression, and LocoGS utilizes point clouds from COLMAP for initialization to maintain consistency
830 with other methods.

831 **Evaluation Details** For rendering performance, we report the peak signal-to-noise ratio (PSNR),
832 structural similarity (SSIM), and learned perceptual image patch similarity (LPIPS). For VRAM
833 consumption, we report both static VRAM overhead and rendering VRAM overhead. The former
834 refers to the VRAM required to load all 3DGS primitives into the renderer and dequantize or decode
835 them into a ready-to-render state. This is not necessarily equal to the file size, as some methods
836 might represent Gaussian attributes in a more compact or lower-precision data format but still require
837 restoring them to full, 32-bit precision Gaussian attributes before rendering. The latter denotes the
838 peak VRAM usage during the rendering process, which is typically larger than the static VRAM
839 due to the introduction of intermediate variables (e.g., projected 2D Gaussian attributes, key-value
840 tables for tile-based rendering) during rendering. We measure this value across all test viewpoints
841 and report the average. To exclude the overhead introduced by the framework itself and the impact of
842 memory fragmentation, we note that the memory allocation and management of the official 3DGS
843 implementation’s renderer are almost entirely handled by PyTorch. Therefore, all VRAM-related
844 metrics are reported using the values provided by the PyTorch framework.

845 **Proximal Operators Implementation** As mentioned in the main text, for opacity projection, we
846 reuse the two proximal operators from GaussianSpa (Zhang et al., 2025b). These operators select
847 κ_o primitives to preserve based on different importance criteria. The first is a **Magnitude-Based**
848 **Selection**, which sorts the input vector $(\mathbf{o} + \lambda_o)$ and preserves the κ_o elements with the highest
849 magnitudes. The second operator employs an **Importance-Based Selection**, adopting the impor-
850 tance score metric from MiniSplatting (Fang & Wang, 2024). The importance score I_i for the i -th
851 Gaussian is the sum of its blending weights across all intersecting camera rays seen during training:
852

$$853 \quad I_i = \sum_{j=1}^K w_{ij} \quad (A.17)$$

854 where the blending weight w_{ij} is the product of the Gaussian’s opacity α_i , its projected 2D value
855 G_i^{2D} , and the accumulated transmittance T_{ij} along the j -th ray. Here, K is the total number of inter-
856 secting rays. The operator then preserves the κ_o primitives with the highest importance scores. While
857 GaussianSpa selects the specific operator based on scene characteristics (e.g., indoor vs. outdoor),
858 we simplify the overall framework by consistently adopting the **Importance-Based Selection** as our
859 unified approach across all scenes.

864 A.3 MORE QUANTITIVE RESULTS
865
866867 Table 3: **Average parameter costs for the color model per Gaussian primitive.** Storage and
868 Rendering costs are measured in the number of float32 parameters. Decode Overhead indicates
869 whether a method introduces significant additional VRAM during calculating color. The values of
870 our method are calculated on the DeepBlending dataset.

Method	Storage	Rendering	Decode Overhead
3DGS	48	48	No
EAGLES	< 1	> 48	Yes
CompactGaussian	< 1	> 32	Yes
ours w/o lobe-pruning	24	24	No
ours	9.7	9.7	No

880 Table 4: **Comparison of anchor-based 3DGS and specialized SOTA compression schemes** on
881 the Mip-NeRF360 dataset. The best result is shown in **bold**.
882

Method	PSNR↑	SSIM↑	LPIPS↓	VRAM↓	FPS↑
Scaffold-GS	27.74	0.811	0.226	612	123
HAC++ (highrate)	27.81	0.811	0.231	637	115
HAC++ (lowrate)	27.60	0.803	0.253	514	132
Ours	27.54	0.824	0.209	265	200

891 Table 5: **Mip-NeRF360 Indoor per scene results.** Best results are in **red region** , second best are
892 in **orange region** . Memory (VRAM) values are in MB.
893

Scene	Method	PSNR↑	SSIM↑	LPIPS↓	Static VRAM↓	Rendering VRAM↓
Bonsai	3DGS	32.23	0.940	0.205	393	798
	Reduced3DGS	31.44	0.933	0.214	66	276
	GaussianSpa	31.76	0.943	0.198	146	420
	Ours(HQ)	31.95	0.943	0.192	44	266
	Ours(LM)	31.51	0.937	0.205	28	227
Counter	3DGS	29.09	0.906	0.201	393	893
	Reduced3DGS	28.57	0.899	0.212	67	331
	GaussianSpa	28.98	0.911	0.191	175	518
	Ours(HQ)	28.78	0.906	0.195	44	326
	Ours(LM)	28.35	0.892	0.220	25	272
Kitchen	3DGS	31.24	0.925	0.126	581	1197
	Reduced3DGS	31.03	0.921	0.133	118	428
	GaussianSpa	31.50	0.929	0.127	163	443
	Ours(HQ)	31.51	0.927	0.127	51	309
	Ours(LM)	31.27	0.923	0.133	36	275
Room	3DGS	31.55	0.918	0.220	474	1007
	Reduced3DGS	31.03	0.914	0.227	68	310
	GaussianSpa	31.61	0.922	0.211	141	401
	Ours(HQ)	31.56	0.923	0.207	36	250
	Ours(LM)	31.16	0.915	0.226	22	214

918
 919 **Table 6: Tanks&Temples per scene results.** Best results are in **red region**, second best are in
 920 **orange region**. Memory (VRAM) values are in MB.
 921

Scene	Method	PSNR↑	SSIM↑	LPIPS↓	Static VRAM↓	Rendering VRAM↓
Train	3DGS	21.97	0.818	0.198	397	773
	Reduced3DGS	21.74	0.804	0.219	74	227
	GaussianSpa	21.81	0.826	0.197	164	339
	Ours(HQ)	21.52	0.820	0.197	49	208
	Ours(LM)	21.32	0.818	0.205	35	160
Truck	3DGS	25.39	0.881	0.143	752	1268
	Reduced3DGS	25.28	0.875	0.154	105	276
	GaussianSpa	25.65	0.887	0.126	115	448
	Ours(HQ)	25.37	0.886	0.121	54	214
	Ours(LM)	25.23	0.883	0.130	39	167

929
 930
 931
 932
 933
 934
 935
 936
 937
 938
 939
 940 **Table 7: Mip-NeRF360 Outdoor per scene results.** Best results are in **red region**, second best are
 941 in **orange region**. Memory (VRAM) values are in MB.
 942

Scene	Method	PSNR↑	SSIM↑	LPIPS↓	Static VRAM↓	Rendering VRAM↓
Bicycle	3DGS	25.20	0.764	0.210	1794	2955
	Reduced3DGS	25.10	0.763	0.219	314	686
	GaussianSpa	25.43	0.779	0.222	240	491
	Ours(HQ)	25.34	0.782	0.206	63	248
	Ours(LM)	25.16	0.771	0.226	48	205
Stump	3DGS	26.61	0.769	0.217	1585	2487
	Reduced3DGS	26.75	0.778	0.218	318	663
	GaussianSpa	27.07	0.802	0.207	222	460
	Ours(HQ)	27.19	0.803	0.198	60	233
	Ours(LM)	27.04	0.797	0.213	47	198
Garden	3DGS	27.36	0.863	0.108	1494	2449
	Reduced3DGS	27.14	0.858	0.116	303	684
	GaussianSpa	27.22	0.854	0.140	227	460
	Ours(HQ)	27.17	0.858	0.127	78	273
	Ours(LM)	26.62	0.846	0.143	63	240
Flowers	3DGS	21.49	0.602	0.338	1066	1716
	Reduced3DGS	21.37	0.598	0.345	221	497
	GaussianSpa	21.60	0.624	0.332	209	451
	Ours(HQ)	21.65	0.625	0.331	66	255
	Ours(LM)	21.37	0.609	0.344	48	203
Treehill	3DGS	22.58	0.633	0.328	1192	1947
	Reduced3DGS	22.41	0.630	0.337	241	563
	GaussianSpa	22.84	0.655	0.312	226	471
	Ours(HQ)	22.74	0.649	0.313	59	233
	Ours(LM)	22.44	0.640	0.335	43	188

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
Table 8: **DeepBlending per scene results.** Best results are in red region, second best are in orange region. Memory (VRAM) values are in MB.

Scene	Method	PSNR↑	SSIM↑	LPIPS↓	Static VRAM↓	Rendering VRAM↓
Playroom	3DGS	30.00	0.903	0.244	678	1197
	Reduced3DGS	29.96	0.904	0.248	90	276
	GaussianSpa	30.48	0.914	0.239	137	320
	Ours(HQ)	30.77	0.914	0.229	52	218
	Ours(LM)	30.58	0.912	0.243	30	167
DrJohnson	3DGS	29.41	0.902	0.239	1129	1941
	Reduced3DGS	29.24	0.901	0.247	152	404
	GaussianSpa	29.50	0.910	0.239	184	424
	Ours(HQ)	29.56	0.909	0.237	56	269
	Ours(LM)	29.45	0.904	0.250	37	219

988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
Table 9: **Real-time rendering performance comparison of different methods across various desktop and mobile devices in Bicycle scene.** All tests were conducted using a WebGL-based viewer, which we modified from the repository (Kwok & Ye, 2023) to support both Spherical Gaussians and 3rd-order Spherical Harmonics. Performance measured in FPS. "cannot render" indicates that the device's browser crashed or showed a black screen, while "render error" signifies incorrect rendering, such as color display issues. The best result is shown in **bold**.

Method	RTX3060	Dimensity 9400+	Snapdragon 8+ Gen 1	Snapdragon 888
3DGS	26.3	6.6	cannot render	cannot render
GaussianSpa	165.0	31.4	render error	render error
Ours	165.0	91.0	120.9	60.1

1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
Table 10: **Storage Size Comparison.** We compare the storage size (MB) of our method against pruning-based methods (MaskGaussian, GaussianSpa) and the compact Gaussian approach (EAGLES). Our method not only surpasses pruning baselines but also achieves superior storage efficiency compared to EAGLES, even without employing explicit compression techniques.

Method	Storage on MipNeRF 360 (MB) ↓		Storage on DeepBlending (MB) ↓	
	MaskGaussian	GaussianSpa	EAGLES	Ours (HQ)
MaskGaussian	271		156	
GaussianSpa	115	31.4	104	
EAGLES	54		52	
Ours (HQ)	55		54	
Ours (LM)	40		33	

1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
Table 11: **Quantitative comparison on the Shiny Blender Real dataset.** **Bold** figures indicate the best results.

Method	PSNR ↑				SSIM ↑				LPIPS ↓			
	Mean	Garden	Sedan	Toycar	Mean	Garden	Sedan	Toycar	Mean	Garden	Sedan	Toycar
Ref-NeRF	23.62	22.01	25.21	23.65	0.646	0.584	0.720	0.633	0.264	0.251	0.234	0.231
ENVIDR	23.00	21.47	24.61	22.92	0.606	0.561	0.707	0.549	0.332	0.263	0.387	0.345
3DGS	23.85	21.75	26.03	23.78	0.662	0.571	0.771	0.637	0.230	0.248	0.206	0.237
2DGS	24.15	22.53	26.23	23.70	0.661	0.609	0.778	0.597	0.292	0.254	0.225	0.396
GShader	23.46	21.74	24.89	23.76	0.647	0.576	0.728	0.637	0.254	0.274	0.259	0.239
R3DG	21.98	21.92	21.18	22.83	0.619	0.556	0.643	0.657	0.349	0.354	0.380	0.312
3DGS-DR	24.00	21.82	26.32	23.83	0.664	0.581	0.773	0.639	0.230	0.247	0.208	0.231
Ref-Gaussian	24.61	22.97	26.60	24.27	0.685	0.617	0.777	0.660	0.246	0.256	0.245	0.256
Ours	24.44	23.01	25.93	24.38	0.695	0.631	0.788	0.666	0.230	0.256	0.175	0.259

Table 12: **Comparison of training time across different datasets and hardware configurations.** We report the training duration (in minutes and seconds) for MEGS² on both RTX 3090 and RTX 4090 GPUs, and compare it with GaussianSpa and 3D Gaussian Splatting (3DGS) on an RTX 4090.

Dataset	Scene	Ours		GaussianSpa	3DGS
		RTX 3090	RTX 4090	RTX 4090	RTX 4090
Mip-NeRF360	Bicycle	37min51s	23min20s	27min57s	25min28s
	Stump	34min09s	22min34s	26min18s	20min48s
	Garden	38min23s	22min56s	29min21s	24min06s
	Flowers	39min12s	26min57s	30min13s	17min08s
	Treehill	37min34s	24min53s	28min38s	18min43s
	Bonsai	51min49s	28min19s	44min14s	16min36s
	Counter	51min22s	32min06s	51min25s	19min20s
Tanks&Temples	Kitchen	50min33s	33min17s	49min36s	22min59s
	Room	41min05s	28min53s	44min41s	20min23s
DeepBlending	Train	28min39s	19min08s	28min45s	11min06s
	Truck	29min59s	20min22s	27min03s	13min21s
DeepBlending	Playroom	33min01s	22min33s	31min32s	18min21s
	DrJohnson	38min28s	24min57s	36min26s	24min31s

Table 13: **Comparison of the number of Gaussian primitives.** We report the total number of Gaussians (in millions) across three benchmark datasets.

Method	Mip-NeRF360	Tanks & Temples	DeepBlending
3DGS	2.718	1.568	2.461
LP-3DGS	1.866	1.116	-
Mini-splatting	0.559	0.320	0.397
MaskGaussian	1.205	0.590	0.694
GaussianSpa	0.528	0.447	0.409
Ours(HQ)	0.611	0.618	0.598
Ours(LM)	0.462	0.437	0.411

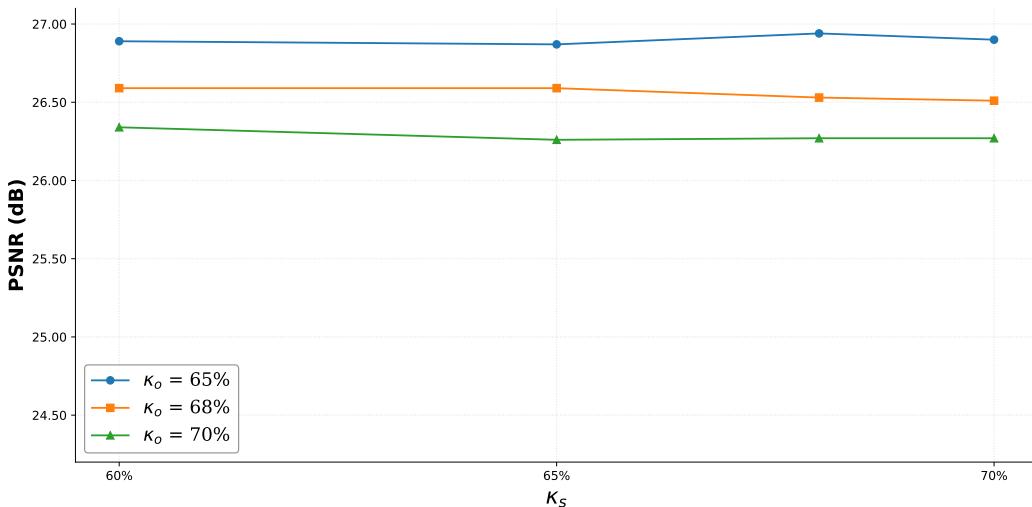
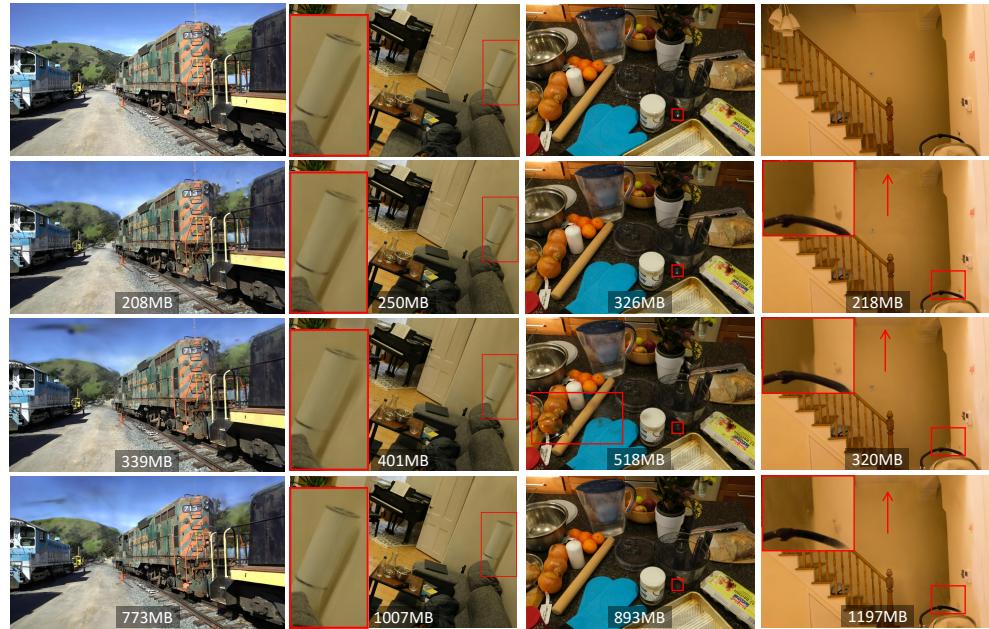


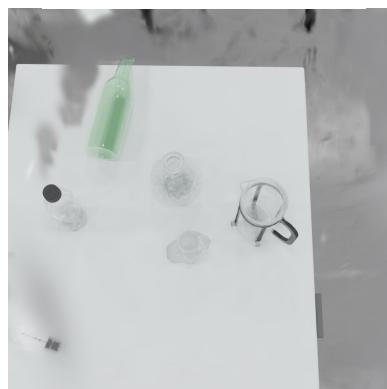
Figure 6: **Sensitivity analysis of pruning hyperparameters.** We investigate the impact of varying the lobe budget κ_s (x-axis) and the primitive budget κ_o (represented by different colored lines) on rendering quality (PSNR) and memory consumption on the garden scene. The values annotated in the boxes indicate the specific rendering VRAM usage for each configuration.

1080
1081 A.4 MORE QUALITATIVE RESULTS
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102

1103 Figure 7: More qualitative results on Train, Room, Kitchen, Counter and Playroom scenes
1104 comparing to previous methods and the corresponding ground truth images. The rendering VRAM
1105 consumption for the corresponding method is annotated at the bottom of each image.

1106
1107 Success Cases

1108 Failure Cases

1109
1110
1111
1112
1113
1114
1115
1116
1117 Specular Highlights
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128 Dense Transparency
1129
1130

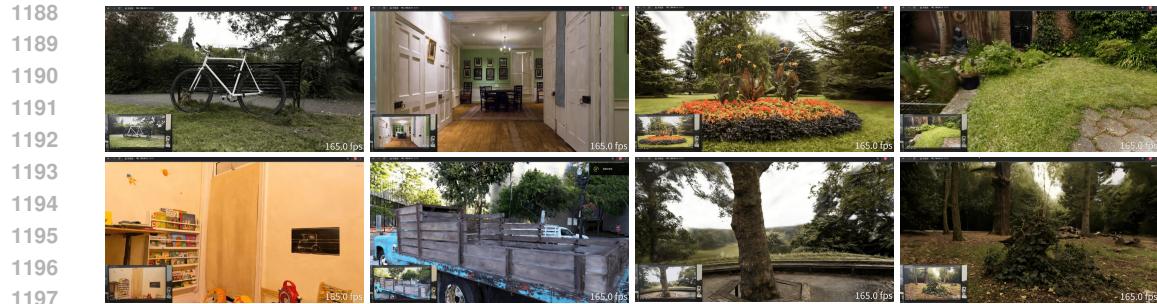
1131 Figure 8: **Visualizations of success and failure cases under challenging material conditions.**
1132 The left column demonstrates successful reconstructions of scenes containing specular highlights
1133 and transparency. Conversely, the right column illustrates specific failure cases where we observe
reconstruction artifacts under conditions of highly specular highlights or dense transparency.



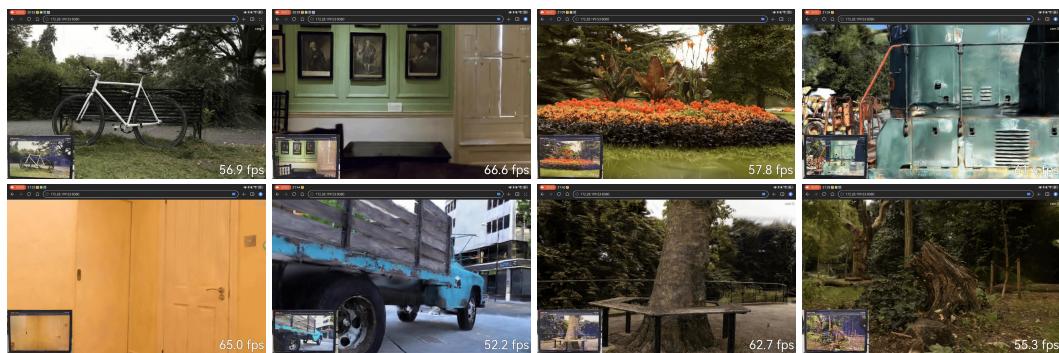
Figure 9: Qualitative ablation results on the Bicycle, Flowers, Counter and Room scenes. The rendering VRAM consumption for the corresponding method is annotated at the bottom of each image.



Figure 10: Real-time rendering performance of MEGS² on various scenes, showcased via a WebGL-based viewer on OnePlus Ace 5 Ultra with **MediaTek Dimensity 9400+**. The live frame rate (FPS) for each view is annotated in the imagery.



1198 Figure 11: Real-time rendering performance of MEGS² on various scenes, showcased via a WebGL-
1199 based viewer on a lenovo laptop with **NVIDIA GeForce RTX3060 Laptop GPU**. The live frame
1200 rate (FPS) for each view is annotated in the imagery.
1201



1217 Figure 12: Real-time rendering performance of MEGS² on various scenes, showcased via a WebGL-
1218 based viewer on Huawei MatePad Air with **Qualcomm Snapdragon 888**. The live frame rate (FPS)
1219 for each view is annotated in the imagery.
1220



1239 Figure 13: Real-time rendering performance of MEGS² on various scenes, showcased via a WebGL-
1240 based viewer on RedMi K60 with **Qualcomm Snapdragon 8+ Gen 1**. The live frame rate (FPS)
1241 for each view is annotated in the imagery.
1242



Figure 14: **Visual comparison in the Playroom scene.** "With SG" indicates the variant where Spherical Harmonics (SH) are replaced by Spherical Gaussians (SG). As observed in the figure, the speckled texture stems from the Ground Truth, which exhibits actual shadow regions alongside inherent image roughness that complicates training. While most methods smooth this detail out entirely, both GaussianSpa and MEGS² (Ours) manage a rough reconstruction.

A.5 THE USE OF LARGE LANGUAGE MODELS (LLMs)

The article was refined using Gemini 2.5 Pro / Gemini 2.5 Flash / DeepSeek R1 to improve phrasing for natural English expression and correct grammatical errors.

1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295