

KISS-GS: 3D Gaussian Splatting Compression Kept Simple

Wieland Morgenstern^{1,2}, Friedrich Elias Branschke^{1,3},
Florian Fleischmann^{1,3}, Adrian Szatmari^{1,3}, Paul Schlack¹,
Florian Barthel^{1,2}, Peter Eisert^{1,2}, and Anna Hilsmann^{1,2}

¹ Fraunhofer HHI, Germany

`{first}.{last}@hhi.fraunhofer.de`

² Humboldt-Universität zu Berlin, Germany

³ Technische Universität Berlin, Germany

Abstract. Scene reconstruction with 3D Gaussian Splatting (3DGS) has become common, however deployment remains painful as the uncompressed file sizes can be massive. Current 3DGS compression systems combine multiple strategies for file size reduction, which can obscure where gains come from and limit component reuse across training pipelines. To make the gains more transparent, we propose KISS-GS, a modular compression pipeline named after the principle of keeping things simple, designed to decouple compression entirely from training. Given a 3DGS scene reconstructed with vanilla 3DGS, we are able to reduce it through compaction by 15.7x using a combination of state-of-the-art pruning schemes. Then we encode it into an image-based format designed for simple, ubiquitous decoding. With the SOG-XT format, we propose a novel extension to Self-Organizing Gaussians with two main contributions: (i) Self-organizing 2D Codebooks and (ii) Parallel Representative Assignment Smoothing (PRAS), which leverages the symmetry of quaternion and scale parameterizations to produce 2D attribute grids more amenable to encoding. This encoding reduces scene size by 6.6x. We show that optional encoding-aware fine-tuning yields a further 2.2x. Across standard 3DGS benchmarks, our simple and modular approach thus achieves a total of 85x to 319x reductions in the size of the scene over uncompressed vanilla 3DGS, setting new benchmarks for real-world scenes and surpassing tightly integrated methods in rate-distortion. Decoding relies solely on web-native image formats, and the modular design makes each stage easy to combine with future advances in reconstruction and compaction.

Code and project page: <https://fraunhoferhhi.github.io/KISS-GS/>

1 Introduction

“Simplicity is a great virtue but it requires hard work to achieve it and education to appreciate it. And to make matters worse: complexity sells better.”

— Edsger W. Dijkstra (1984), On the nature of Computing Science [6]




Fig. 1: Rate distortion on *Mip-NeRF 360* for KISS-GS, with quality measured by peak signal-to-noise ratio (PSNR). At the INRIA 3DGS quality reference (INRIA-Q), KISS-GS reduces file size by 228 $\times$ from 738 MB to 3.2 MB. From the uncompressed INRIA .ply baseline, POPSpa matches INRIA-Q with a 15.7 $\times$ reduction, our SOG-XT encoding adds 6.6 $\times$, and encoding-aware fine-tuning adds a further 2.2 $\times$, outperforming HAC++ [5].

In practice, 3D Gaussian Splatting (3DGS) scenes are still commonly exchanged as raw .ply exports. These files are easy to produce, but large and inefficient for storage, streaming, and deployment. As 3DGS moves beyond research prototypes into hardware-constrained web viewers, mobile renderers, and AR/VR applications, file size, decoding cost, and memory footprint become obstacles.

Storage efficiency was not a focus of the original 3DGS formulation [16], leading to a rapidly growing body of work on compression. The recent 3DGS.zip survey [1] distinguishes two storage-reduction strategies, *compaction*, which reduces the number of primitives, and *compression*, which reduces the number of bits per primitive. To avoid overloading the term *compression*, we refer to this bits-per-primitive component as *encoding*. We add *adaptation* as a third strategy, which uses encoding-aware training or fine-tuning to steer parameters toward configurations that are friendly to quantization and coding. KISS-GS turns these three strategies into explicit pipeline stages, whose separate rate-distortion gains are previewed in Fig. 1. Many methods combine multiple strategies to achieve their final rate-distortion scores. While effective, this obscures where gains actually originate, creating an *attribution gap*. Without clear attribution, it is difficult to compare methods fairly and reuse components across pipelines. Compaction alone can already yield large gains. Densification and pruning improvements can reduce primitive counts by an order of magnitude in some settings without large drops in reconstruction quality, as demonstrated by Mini-Splatting [7, 8]. Some approaches, such as anchor-based families of methods, restructure the representation so fundamentally that compression becomes inseparable from the

reconstruction regime and data format. This *training-format coupling* effectively locks the codec to a particular 3DGS training pipeline: when new reconstruction methods or parameterizations emerge, their improvements cannot be adopted without reworking the compression scheme itself.

We call our approach *KISS-GS*, inspired by the KISS principle (“keep it simple, stupid” [37]) and restrict ourselves to a deliberately simple, browser-compatible decoding pipeline backed by widely supported image codecs.

Guided by this decoder-first stance, we ask three research questions (RQs):

1. **(RQ1)** Can we close the **attribution gap** by modularizing the compression pipeline and measuring the contribution of each stage in turn?
2. **(RQ2)** Can a **simple image-based encoding format** match more complex, integrated methods like HAC++ [5] across real-world 3DGS compression benchmarks?
3. **(RQ3)** Does **encoding-aware fine-tuning** still deliver gains when constrained to a simple, web-friendly codec?

To close the **attribution gap** (RQ1), we propose KISS-GS, a modular compression framework that separates compaction, encoding, and adaptation into distinct stages, enabling each stage to be replaced, ablated, or analyzed independently. We instantiate the compaction stage with POPSpa, a reconstruction-independent module that combines state-of-the-art pruning schemes after the initial reconstruction, ensuring that gains from fewer primitives are never conflated with gains from better encoding. To test whether a **simple image-based encoding format** can match more integrated methods (RQ2), we introduce SOG-XT, an image-based codec and storage format designed for web-friendly decoding. We extend the Self-Organizing Gaussians (SOG) [27] approach with (i) Self-organizing 2D codebooks for spherical harmonics and (ii) Parallel Representative Assignment Smoothing (PRAS) to improve compression rates while keeping the decoder simple. To isolate **encoding-aware fine-tuning** under the constrained decoder model (RQ3), we systematically evaluate it with encoding-in-the-loop as a separate stage.

2 Related work

The original 3DGS paper [16] focused on reconstruction quality and real-time rendering performance, without explicitly addressing storage or memory constraints. Due to its strong reconstruction quality with explicit primitives (in contrast to the implicit nature of NeRF-based approaches [26, 28]), the community quickly adopted 3DGS. Many follow-up works have since investigated storage reduction, memory efficiency, and rendering performance, introducing modifications at different pipeline stages which can broadly be categorized into *compaction*, *encoding*, and *adaptation* [1].

Compaction reduces the number of active Gaussians while preserving or improving reconstruction quality, directly impacting storage footprint and rendering cost. Many methods perform sensitivity-based pruning after or during training by estimating the importance of individual Gaussians. Various criteria have

been proposed, including strict primitive budgets via score-based sampling [25], reconstruction error [11, 12, 20, 33], as well as primitive transmittance, opacity, or rendering contribution [9, 19, 22]. Another class of methods controls the Gaussian population through sampling-based strategies. For example, MCMC-GS [17] removes low-opacity primitives and relocates them to high-opacity regions to improve reconstruction under a fixed primitive budget. Finally, some approaches formulate compaction as an explicit sparsity-constrained optimization problem: RDO-GS [41] learns binary pruning masks jointly with the reconstruction objective using sparsity-inducing regularization, whereas GaussianSpa [47] introduces an auxiliary sparsity variable and uses an augmented-Lagrangian formulation that alternates between standard reconstruction optimization and a projection step that enforces a target sparsity level.

Encoding reduces the number of bits required to represent each primitive without necessarily changing the primitive count. A common strategy is numerical compression of Gaussian attributes, often using vector quantization (VQ), particularly for view-dependent spherical harmonics [29, 30, 32]. These approaches reduce bitwidth without altering primitive layout through quantization and entropy coding.

Another class of methods replaces explicit attribute storage with structured representations that exploit spatial correlations between primitives. These include anchors, codebooks, multi-layer perceptrons (MLPs), octrees, triplanes, or hash grids to encode Gaussian attributes [3, 21, 24, 36, 44]. Scaffold-GS [24], for example, predicts attributes from neighboring anchors via a lightweight MLP. Extensions include learning sparse primary and dense secondary anchors [23], hierarchical pyramids where coarser anchors predict finer ones [42], and querying hash-grid features at anchor locations [3, 5].

Another direction organizes Gaussian attributes into **structured 2D layouts** to leverage standard image codecs. Mapping 3D primitives to 2D grids via Morton curves or Self-Organizing Gaussians (SOG) [27] produces locally smooth attribute maps optimized for conventional encoders. Although encoding requires an initial sorting step, decoding reduces to standard image decompression, which is widely hardware-accelerated and readily available in browser and mobile environments. This decoding simplicity has enabled **practical deployment** on mobile devices, AR/VR platforms, and browser-based viewers, leading to formalization of the `.sog` format [35], which combines Self-Organizing Gaussians with VQ strategies for spherical harmonics. The format has been adopted in research and practical web pipelines [15, 34, 39]. Alternative formats like `.spz` [31] use structured quantization with a custom binary format.

Adaptation modifies scene parameters to make them more compatible with the target encoding. This can be achieved either during training or through post-training fine-tuning. RDO-GS [41] adds VQ-loss and rate-loss to the objective function, FCGS [4] incorporates entropy loss accounting for bit count, while HAC [3] and HAC++ [5] combine rendering fidelity, entropy, masking, and hash-grid size losses.




Fig. 2: The KISS-GS pipeline with four distinct stages. (1) A standard 3DGS reconstruction yields a .ply, with primitive count setting the rate-distortion point. (2) POPSpa compacts the scene. (3) SOG-XT encodes the compacted primitives into attribute images. (4) Optional codec-aware fine-tuning improves rate-distortion without changing the output format. Below the dashed separator, we show per-stage alternatives to the KISS-GS stages and contrast them with integrated methods like HAC++ spanning all stages as a single non-modular unit.

3 Modules of the KISS-GS pipeline

Figure 2 shows how the modular KISS-GS pipeline takes any off-the-shelf 3DGS scene as input, applies compaction, encodes the result with SOG-XT, and optionally fine-tunes it with the codec in the loop. When competing methods differ in both file size and reconstruction quality, individual operating points listed in tables are difficult to compare. We therefore use rate-distortion curves to compare this tradeoff across operating points. When one curve dominates another in the Pareto sense (higher quality at the same file size, or smaller file size at the same quality) the dominant method is strictly better at every operating point. We control the rate-distortion tradeoff primarily through primitive count, which has the additional benefit of reducing rendering cost and the work required for encoding and decoding. Using more primitives with heavier compression to reach the same operating point would instead increase all three, with no gain in rate or quality. We describe each stage in turn.

3.1 3D Reconstruction

KISS-GS accepts any standard 3DGS .ply file as input, regardless of which implementation produced it. This is a deliberate design choice: by decoupling compression from reconstruction, we avoid training-format coupling and ensure that improvements in reconstruction quality can be adopted without changes

to the compression pipeline. The standard 3DGS parameterization (means, covariances represented as quaternions and scales, spherical harmonic colors up to degree 3, and scalar opacity) is shared across all major implementations, and KISS-GS preserves this representation throughout. Importantly, this parameterization relies on no MLPs, hash grids, or anchor structures, keeping the decoder simple. The original training pipelines apply non-linear mappings to scales and opacities to constrain parameter ranges. We use the same mappings because the resulting values are well suited to uniform quantization.

The 3DGS reconstruction ecosystem is actively evolving, with implementations varying in speed, quality, and license, including INRIA [16], gsplat [45], Lichtfeld Studio [40], and Hahlbohm et al. [10]. Being decoupled from any specific implementation means KISS-GS can benefit from future reconstruction advances without modification. Our main experiments use the gsplat MCMC reconstruction [17] with default parameters for 30k iterations; compatibility with alternative implementations is demonstrated in Sec. 4.

3.2 Compaction with POPSpa

Densely trained 3DGS scenes typically contain substantial redundancy. Since storage scales approximately linearly with the primitive count, compaction is essential to achieve high compression ratios. We instantiate this stage with *POPSpa*, a reconstruction-independent compaction recipe that starts from a trained 3DGS scene and combines score-based pruning from GaussianPOP [20] with the alternating optimize-sparsify scheme of GaussianSpa [47], augmented by effective-rank regularization. Its role is not to introduce a new pruning objective, but to provide a strong, replaceable module whose primitive-count gains can be measured separately from encoding gains.

Effective Rank Regularization. Although all covariance matrices maintain a full rank of 3, they exhibit widely different eigenvalue spectra, dictating the shape of the Gaussians. To quantify these shapes during training, we utilize the effective rank (erank) as a differentiable relaxation that measures the dimensions effectively occupied by the distribution. Formally, for a 3D Gaussian with ordered scale attributes $s_1 \geq s_2 \geq s_3$, the normalized eigenvalues of its covariance matrix Σ are defined as $p_i = \frac{s_i^2}{\sum_{j=1}^3 s_j^2}$. The effective rank is formulated as the exponential of the Shannon entropy of the probability distribution induced by the eigenvalues of the covariance matrix: $\text{erank}(\Sigma) = \exp\left(-\sum_{i=1}^3 p_i \cdot \log p_i\right)$ [38]. We employ the effective rank regularization (Eq. (1)) proposed by Hyung *et al.* [14] to discourage the formation of extremely needle-shaped Gaussians, characterized by $\text{erank}(\Sigma_k) \approx 1$, which can overfit training views and produce artifacts from novel viewpoints. At the same time, this regularization favors disk-shaped Gaussians exhibiting an $\text{erank}(\Sigma_k) \approx 2$, which are better suited for surface reconstruction.

$$\mathcal{L}_{\text{erank}} = \sum_k \lambda_{\text{erank}} \max(-\log(\text{erank}(\Sigma_k) - 1 + \epsilon), 0) + s_3 \quad (1)$$

First Pruning Stage. For each Gaussian primitive k , we use the error criterion of GaussianPOP [20] to estimate how much the rendered image changes when k is removed from the scene representation. Let C_{render} be the original rendered color and C'_{render} the color obtained after omitting Gaussian k ; the per-primitive error $\Delta SE_k = \|C_{\text{render}} - C'_{\text{render}}\|^2$ then quantifies the change in pixel colors, accumulated over all training views. We compute ΔSE_k for all primitives using the efficient single-pass algorithm of [20] and prune a fixed fraction with the lowest scores.

Optimize-Sparsify Refinement. To counteract potential degradation caused by large pruning ratios, we refine the remaining primitives using the alternating optimize-sparsify scheme of GaussianSpa [47]. GaussianSpa formulates primitive selection as a constrained optimization problem with an ℓ_0 sparsity objective on primitive opacities:

$$\min_{\mathbf{a}, \Theta} \mathcal{L}(\mathbf{a}, \Theta), \quad \text{s.t. } \|\mathbf{a}\|_0 \leq \kappa \quad (2)$$

where Θ denotes all primitive parameters except opacity $\mathbf{a}$, and κ is the primitive budget. The optimization alternates between a standard reconstruction step and a sparsification step that gradually reduces the opacity of expendable primitives. GaussianSpa applies this scheme during training; we repurpose it as a post-training refinement stage of 5k iterations. This allows the remaining primitives to compensate for the reduced primitive set by redistributing opacity and geometry among nearby Gaussians. In this stage, we additionally apply the effective rank regularization term $\mathcal{L}_{\text{erank}}$ defined in Eq. (1).

Second Pruning and Fine-Tuning. For a second pruning pass, we recompute ΔSE_k to remove primitives that have become redundant. Finally, another 5k iterations of fine-tuning recover quality lost in the previous steps, for a total of 10k post-processing iterations.

3.3 Encoding with SOG-XT

Following compaction with POPSpa, the remaining primitives are encoded by SOG-XT, an image-based codec extending the Self-Organizing Gaussians (SOG) approach [27] with two novel contributions described below.

The core idea of SOG [27] is to arrange the N primitives into a $\lceil \sqrt{N} \rceil \times \lceil \sqrt{N} \rceil$ 2D grid using PLAS, a parallel assignment algorithm that preserves spatial locality in the grid layout. The resulting attribute maps are locally smooth, making them well-suited for standard image codecs. SOG stores the individual per-attribute slices as lossless JPEG-XL images, with quantization applied post-training. In the original SOG method, densification parameters, training-time smoothness regularization, and quantization are coupled, which makes its encoding stage difficult to apply in isolation. The `.sog` format keeps the image-based storage idea of SOG [27], but makes it usable as a post-training format by using

Morton-order curves to arrange primitives and k-means clustering for spherical harmonics dimensionality reduction, with centroids sorted lexicographically.

Like `.sog`, SOG-XT decouples encoding from reconstruction and compaction entirely, operating as a pure post-training step on any compacted scene. SOG-XT retains the PLAS-based grid layout for all attributes but replaces the 48 raw SH slices of SOG [27] with vector quantization: k-means clustering reduces the high-dimensional SH representation to a compact codebook, with only per-primitive cluster indices stored. Rather than sorting these centroids lexicographically as in `.sog`, SOG-XT sorts the 2D cluster centroids with PLAS, yielding a codebook that is itself locally smooth in 2D. Cluster assignments are stored as two 8-bit image channels (u, v) indexing the PLAS-sorted centroid grid. This deliberately caps the SH codebook at $256 \times 256 = 65,536$ entries, which was the point where increasing the codebook size no longer improved PSNR in our experiments. This 2D structure additionally enables fine-tuning to refine the codebook and provides a natural foundation for future vector quantization methods that exploit spatial coherence more aggressively.

3.4 Covariance Symmetry and PRAS

The mapping from a 3D covariance matrix to the parameterization of tuples of quaternions and scales (q, s) is one-to-many. A unique covariance matrix maps to an entire equivalence class of discrete parameter choices that all yield the identical spatial distribution. This symmetry stems from three independent sources of redundancy in the decomposition of Σ : **Scale Permutations**: The three principal radii can be assigned to the coordinate axes in $3! = 6$ ways, provided they are compensated by 90° rotations around principal axes to arrive at the identical covariance. **Eigenvector Orientations**: Flipping eigenvector directions yields 2^3 combinations, of which 4 represent valid rotations satisfying $\det(R) = 1$. **Quaternion Double-Cover**: For every valid rotation matrix R , there exist two quaternions, q and $-q$, which produce that exact rotation.

Combined, this results in $6 \cdot 4 \cdot 2 = 48$ distinct parameter combinations within an equivalence class that describe a single Gaussian. To exploit this symmetry, we propose **Parallel Representative Assignment Smoothing (PRAS)**. Similarly to *PLAS*, *PRAS* employs coarse to fine optimization over 2D parameter grids. First, to ensure comparable magnitudes during optimization, the floating-point quaternions and log-scales are mapped to 8-bit quantized proxy grids. The algorithm then iterates over a sequence of exponentially decreasing blur radii. At each scale, a low-pass filter is applied to the current proxies to generate spatially smoothed target grids. For every primitive, *PRAS* evaluates all candidates within its equivalence class and assigns the representative that minimizes the sum of absolute channel differences, i.e. the ℓ_1 distance, to the smoothed targets. Because the search space is strictly restricted to valid equivalence classes, *PRAS* effectively removes artificial variance to improve the compressibility of scale and quaternion grids while leaving the resulting 3D covariances and grid coordinates unchanged.




Fig. 3: Encoding overview for SOG-XT. PLAS sorts primitives into 2D attribute grids using both mean byte planes. View dependent color is reduced with a k-means codebook whose centroids are PLAS sorted again, and assignments are stored as UV indices. Bars show the relative on-disk size contribution of each component.

3.5 SOG-XT format

Figure 3 shows an example of an encoded SOG-XT container with the file size and share of the total for each stored component.

Following SOG [27], we arrange primitives into a 2D grid with PLAS and store the attribute grids as images. SOG uses JPEG-XL with higher bit depth support, while we restrict storage to browser-decodable 8-bit images. We avoid truncation by padding primitives to a square grid with side length aligned to a multiple of 8 and storing an *active mask*. Padded primitives use fixed defaults and are dropped at decoding time, which adds little metadata overhead.

We apply per-channel min-max quantization with uniform step sizes and store each channel’s minimum and maximum as float32 values in a YAML metadata file. In the example container shown in Fig. 3, all metadata including these ranges accounts for only 0.2% of the total size on disk. These values are therefore negligible for compression, but necessary for exact dequantization.

For the Gaussian means, which define primitive positions, we follow SOG and apply a signed log remapping for space contraction before 16-bit quantization. As in browser-oriented `.sog`, we store the quantized positions as two 8-bit byte planes. The upper byte captures coarse spatial structure, while the lower byte contains finer residual variation and is less smooth under image coding. Our contribution is to expose this split to the layout optimization: we run PLAS on both byte planes, with a stronger weight on the upper byte and a weaker

weight on the lower byte, and apply the resulting permutation to all remaining attributes.

For view-dependent color, we split spherical harmonics into the DC term f_{dc} and the remaining coefficients f_{rest} . We store f_{dc} directly as an 8-bit image. We compress f_{rest} with a k-means codebook of size K that scales with the number of active primitives. In our experiments, we choose the SH codebook size as $K = S^2$. We set the codebook side length S from $\sqrt{N/8}$, clamp it to $[8, 256]$, and round it down to a multiple of 8. As a further enhancement, we PLAS-sort the SH codebook centroids to obtain a 2D centroid grid and store assignments as two 8-bit UV channels that index this grid. We quantize the SH codebook centroids with 100 levels using per-channel min-max ranges. Each centroid contains 15 non-DC SH coefficients with RGB values. We store the resulting centroid grid as a single RGB image by arranging the 15 coefficient images in a 3×5 grid of tiles.

Scales use a log mapping and 8-bit quantization, and opacities use an inverse sigmoid mapping and 8-bit quantization. We optimize quaternions jointly with scales using PRAS, which enables quaternion quantization with $q = 99$. Without PRAS, we find that $q = 255$ is needed for comparable quality. The encoder does not require camera parameters or a renderer.

3.6 Adapting the SOG-XT Format in Encoding-aware Fine-tuning

Optionally, we fine-tune with SOG-XT in the loop. At each training step, we map the current Gaussian attributes to the quantized image tensors used by SOG-XT, reconstruct a Gaussian set from these tensors, render it, and backpropagate the rendering loss through this round trip. The discrete codec structure is fixed at the start, including the active mask, PLAS order, SH assignments, and codebook layout. During training, SH assignments stay fixed while centroid values are recomputed from the current f_{rest} coefficients. Rounding is handled with a straight-through estimator [2].

We optimize the standard rendering objective and add lightweight smoothness losses on the quantized images. We also re-run PRAS during fine-tuning to preserve the compressibility benefits of covariance symmetry resolution as the parameters change. The output format is identical to SOG-XT without fine-tuning, so the decoder remains unchanged. Implementation details are provided in the appendix.

4 Evaluation

Let us demonstrate the effectiveness of keeping things simple: we will show how POPSpa improves reconstruction quality with a reduced number of primitives compared to INRIA 3DGS and competing compaction methods. Then we'll look into our individual contributions that make up SOG-XT: 2D sorted codebooks and our PRAS covariance smoothing. Finally, we show how our results compare to the previous best-in-class, HAC++ [5], over the standard datasets for benchmarking 3DGS compression.

4.1 Evaluation protocol for 3DGS compression

The original 3DGS paper [16] established a de-facto evaluation protocol through its public codebase, which is the baseline for many follow-up works. As explicitly stated in the 3DGS.zip survey [1], this protocol specifies the datasets and splits (21 scenes across *Tanks and Temples*, *Mip-NeRF 360*, *Deep Blending*, and *Synthetic NeRF*), scene-dependent training resolutions ($4\times$ downsampling for *Mip-NeRF 360* outdoor scenes, $2\times$ for indoor), a fixed evaluation resolution (images downsampled to 1600 px on the longest side), and evaluation with peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM) [43], and learned perceptual image patch similarity (LPIPS) [46], using the VGG backbone for LPIPS and images in their $[0, 1]$ range. Full details are provided in the appendix.

Deviating from this evaluation protocol can lead to considerable differences in reported metrics. The effects of training and test resolutions are discussed in 3DGS.zip survey [1], while a discussion of the incorrect LPIPS range being used in 3DGS literature can be found in the NeRF baselines benchmark [18].

In all tables that have color-highlighted values, the best results in each category are highlighted **first**, **second**, and **third**.

4.2 Baselines: INRIA for Quality and HAC++ for Compression

For an uncompressed quality reference, we use INRIA 3DGS [16], which we denote INRIA-Q. The original 3DGS protocol trains for 30k iterations, which matches the reconstruction stage of our pipeline. Our compaction stage then adds 10k refinement iterations before compression. We therefore train the INRIA quality reference for 40k iterations as well, matching the total pre-compression optimization budget of our method.

As the compression baseline, we use HAC++ [5], the highest ranked method in the 3DGS.zip survey. Recomputing all published 3DGS compression methods under a single protocol is infeasible, but HAC++ is the strongest available reference point in the survey. Only the preprint HEMGS [22] achieves higher PSNR on several datasets, while HAC++ is more balanced across metrics and has public code available. We therefore recompute HAC++ under our evaluation protocol and compare to the remaining methods using their self-reported values from the survey. We run HAC++ for 44k iterations as a compute-matched baseline. This matches the total budget of the full KISS-GS pipeline, whose stages use 30k iterations for reconstruction, 10k iterations for POPSpa compaction, and 4k iterations for optional encoding-aware fine-tuning. We use the parameters provided by the authors for their high- and low-rate versions, but follow our evaluation protocol throughout, which leads to slight differences from the numbers reported in their paper. The appendix ablates HAC++ across the reported 30k setting and our 30k, 40k, and 44k recomputations under both evaluation protocols.

4.3 Evaluating Compaction with POPSpa

Figure 1 shows results for POPSpa using seven Gaussian primitive budgets. For the smallest model, we set the target to 64 k Gaussians for all datasets, except for *Synthetic NeRF*, where the lower scene complexity motivates a target of 16 k. To obtain an intended pruning ratio of approximately 80%, we first train an over-parameterized 3DGS model with *gsplat-MCMC* [45] using 320 k primitives (80 k for *Synthetic NeRF*).

For each subsequent point in Fig. 1, we double both the compacted and the corresponding initial budget. The initial (uncompacted) models are capped by a dataset-specific maximum of 3000 k Gaussians for *Tanks and Temples* and *Mip-NeRF 360*, 2048 k for *Deep Blending*, and 640 k for *Synthetic NeRF*. The right-most point corresponds to the maximal model size without pruning and only applies 10 k fine-tuning iterations.

Table 1: Comparison of 3DGS compaction methods across four standard benchmark datasets. Ranks represent average rankings across all available datasets, with model size and overall quality weighted equally: size contributes half the weight, and PSNR, SSIM, and LPIPS each contribute one-sixth. In the top half, post-training pruning is applied to a 30k-step vanilla 3DGS checkpoint, removing 90% of Gaussians. The bottom half compares against compaction-aware training methods, where each method is trained for 10k steps beyond its original implementation, resulting in 40k steps for all methods except Mini-Splatting 2 which was originally trained for only 18k steps. Our method is post-training only, applied to a checkpoint trained with *gsplat*’s MCMC implementation.

	Method	Rank	Tanks and Temples				Mip-NeRF 360				Deep Blending				Synthetic NeRF			
			PSNR↑	SSIM↑	LPIPS↓	k Gauss	PSNR↑	SSIM↑	LPIPS↓	k Gauss	PSNR↑	SSIM↑	LPIPS↓	k Gauss	PSNR↑	SSIM↑	LPIPS↓	k Gauss
post training	Inria [16] SIGGRAPH ’23 (30k)		23.77	.847	.170	1,563	27.28	.807	.222	2,723	29.69	.905	.242	2,463	33.46	.970	.030	259
	+ Ours (+10k)	1.1	23.37	.819	.231	156	26.78	.781	.289	272	29.51	.903	.265	246	32.46	.964	.041	25
	+ Speedy-Splat [11] CVPR ’25 (+10k)	1.5	23.52	.813	.245	156	26.87	.772	.305	272	29.45	.902	.267	246	32.40	.962	.045	25
	+ PUP 3D-GS [12] CVPR ’25 (+10k)	1.8	22.62	.798	.251	156	25.58	.767	.300	272	29.26	.903	.264	246	31.49	.958	.049	25
	gsplat-MCMC [45] (30k)		24.54	.869	.150	1,280	27.95	.829	.196	2,560	30.08	.907	.244	1,280	33.82	.972	.028	320
compaction aware	+ Ours (+10k)	1.7	24.47	.856	.179	256	27.79	.811	.239	512	29.94	.904	.262	256	33.41	.969	.032	64
	gsplat-MCMC [45] (40k)	2.5	23.90	.841	.205	256	27.52	.809	.243	512	29.43	.899	.270	256	32.73	.967	.036	64
	GaussianSpa [47] CVPR ’25 (40k)	2.8	23.63	.849	.171	349	27.40	.817	.227	561	30.09	.913	.245	457	33.34	.968	.029	212
	Mini-Splatting2 [8] (28k)	3.3	23.57	.843	.178	357	27.33	.815	.223	619	30.05	.911	.241	658	32.69	.964	.037	54
	Taming 3DGS [25] SIGGRAPH Asia ’24 (40k) 3.9		24.10	.836	.203	318	27.33	.795	.258	684	29.77	.903	.273	294				

Table 1 presents results for both post-training pruning and compaction-aware training. In the post-training setting, our method consistently outperforms all competing approaches across all benchmarks.

The compaction-aware comparison is more nuanced: since most methods cannot easily predetermine the final scene size or quality, the compared samples sit at slightly different positions on each method’s quality-size trade-off curve. Despite this, our method achieves the best average rank of 1.7 while operating at significantly fewer Gaussians than competing methods across most datasets. The exception is *Synthetic NeRF*, where model sizes are more comparable. In *Deep Blending*, we achieve competitive quality using only 56% of the primitives of the best performing method. Full quality-size trade-off curves for all datasets are

provided in the appendix. Across these curves, our method consistently achieves higher quality than competing methods at a given size.

4.4 Evaluating Compression with SOG-XT

As shown in Fig. 1 for *Mip-NeRF 360* and in the appendix for other datasets, our encoding method SOG-XT is able to compress the test scenes to single-digit Megabyte sizes. With a single parameter choice, as presented in Sec. 3.3, our method works well across a large range of primitive counts, from several thousand to millions. The same encoding method works on the INRIA 40k baseline, highlighting the modularity of the approach. Two popular post-training encoding methods in `.spz` and `.sog` are benchmarked against our method using their official implementations (`splat-transform` v1.6.0 [35] for `.sog`, the official `spz` PyPI package v0.0.6 [31] for `.spz`). With `.sog`, the files are even larger, and this format appears sensitive to the number of primitives, producing a jagged rate-distortion curve.

Table 2: Ablation of contributions to the SOG-XT format on the Bicycle scene with 256 k primitives. The first row contains the absolute values after encoding. The following rows disable individual features of SOG-XT, ordered by size penalty. Most removals increase the final size considerably at small quality gains. Using image encoding, our novel codebook sorting and its UV label packing bring clean size wins without incurring any quality loss. Find additional rows in the appendix.

	Size [MB]	Effect of Leaving Feature Out			
		Δ Size [Bytes] ↓	Δ PSNR ↑	Δ SSIM ↑	Δ LPIPS ↓
SOG-XT (Full Pipeline)	3.878	(3,877,944)	(24.2281)	(0.68176)	(0.35404)
w/o SH AC codebooks	7.313	+3,434,816	+0.2320	+0.00790	-0.00579
w/o WebP image compression	6.908	+3,030,148	+0.0000	+0.00000	+0.00000
w/o quaternion u8 quantization	6.427	+2,548,718	+0.0317	+0.00147	-0.00067
w/o primitive PLAS sorting	4.493	+614,851	+0.0059	+0.00022	-0.00036
w/o centroid PLAS sorting	3.920	+42,416	+0.0000	+0.00000	+0.00000
w/o label UV packing	3.890	+12,142	+0.0000	+0.00000	+0.00000
w/o PRAS	3.866	-11,810	-0.1104	-0.00538	+0.00273
w/o means signed-log remapping	3.359	-519,272	-7.5508	-0.37864	+0.22151

In Table 2, we show how the individual components of SOG-XT contribute to its high compression ratio. In these leave-one-out experiments, we can see that most components provide size wins against small quality losses, like the quantization operations. Our 2D codebooks (centroids PLAS sorting and label UV packing) bring free size wins without any quality loss. Our PRAS method incurs a small size hit, but provides large quality gains. The highest quality loss occurs from building codebooks for SH AC, but otherwise, the files would nearly triple in size. Without the log-space contraction of means, file sizes would shrink. However, with 16-bit quantization for the means, reconstruction quality would degrade severely.

4.5 Adaptation and Comparison with State of the Art

In Figure 1, we show that encoding-aware fine-tuning clearly improves the rate-distortion curve over post-training compression, achieving another $2.2\times$ reduction at INRIA quality for the *Mip-NeRF 360* dataset. While SOG-XT by itself can already achieve great results without feedback from the renderer, running some adaptation can be worth it if a GPU and training views are available at encoding time. A sensitivity analysis in the appendix shows that the default fine-tuning setting is not a fragile optimum; the SH codebook size is the clearest rate-control knob, while additional fine-tuning steps and stronger TV regularization do not improve the trade-off.

Table 3: File-size reduction factors over INRIA 3DGS after 40k iterations. Each entry is the interpolated point where a method matches INRIA-Q for the corresponding metric. KISS-GS achieves the highest reductions on *Tanks and Temples* and *Mip-NeRF 360*, and on two of the three *Synthetic NeRF* metrics.

Method	Tanks and Temples			Mip-NeRF 360			Deep Blending			Synthetic NeRF		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
CodecGS [21] ICCV '25 [†]	53*	-	-	72*	72*	-	76*	76*	-	n/a	n/a	n/a
ContextGS [42] NeurIPS '24 [†]	42*	42*	-	55*	55*	-	187*	187*	-	n/a	n/a	n/a
gsplat-1M [13] arXiv '24 [†]	60*	44	26	52	57	28	n/a	n/a	n/a	n/a	n/a	n/a
HAC [3] ECCV '24 [†]	49*	48	-	46*	46*	-	150*	129	-	47	-	-
HEMGS [22] arXiv '24 [†]	69*	69*	-	59*	59*	-	229*	229*	-	57*	-	-
HAC++ [5] TPAMI '25	77*	68	-	70	36	-	156	156	-	-	-	-
Ours (SOG-XT)	227	104	60	104	73	40	-	-	-	10	-	-
Ours+FT	319	131	72	228	109	64	85	98	-	21	14	14

* smallest available point already reaches INRIA-Q, so the value is a lower bound.

- INRIA-Q not reached within available points.

[†] self-reported result from 3DGS.zip, not recomputed under our protocol.

In Table 3, each entry is the interpolated file-size reduction at which a method matches the INRIA 3DGS 40k quality reference. We use recomputed HAC++ as the main baseline, since it is the strongest reproducible published method in the 3DGS.zip survey. HEMGS reports higher reductions in some cases, but it is a preprint without public code, so we include it only as self-reported context. On *Tanks and Temples* and *Mip-NeRF 360*, KISS-GS reaches INRIA-Q at substantially higher reduction factors than HAC++ across all reported metrics. On *Deep Blending*, the gap is already visible before SOG-XT is applied: compressed HAC++ exceeds uncompressed gsplat-MCMC and INRIA 3DGS on *Playroom* (30.63 vs. 30.42 and 30.09) and *Dr. Johnson* (29.60 vs. 29.48 and 29.43). This points to HAC++’s anchor-based learned representation and regularization rather than to SOG-XT quantization or POPSpa pruning. To keep the comparison compact, we include literature methods only if they provide at least four finite entries in this table.

Notably, many competing methods from the literature do not reach INRIA quality in the perceptual LPIPS metric at any of their available operating points. This may point to overfitting the PSNR metric for the training views and losing generalization capabilities in rendering novel views or novel resolutions (for the *Mip-NeRF 360* dataset).

Runtime measurements in the appendix show that post-hoc SOG-XT encoding is inexpensive compared with reconstruction and compaction. Encoding-aware fine-tuning is more expensive, but it can be omitted. A reference Python decoder reconstructs models with up to 1024k Gaussians in well under a second on CPU, supporting our claim that deployment remains simple.

5 Have we kept it simple?

The title of this paper, together with the Dijkstra quote in the introduction, makes a deliberately cheeky promise. We set out to keep 3DGS compression simple. With respect to simplicity, we fell short in one place: the encoder is a sophisticated modular pipeline with modules such as POPSpa and novel components such as PRAS. This deliberate complexity lets KISS-GS separate reconstruction, compaction, encoding, and **encoding-aware fine-tuning**, close the **attribution gap**, and make the source of each gain measurable.

The payoff is substantial: KISS-GS reaches **single-digit megabytes** at vanilla 3DGS reconstruction quality. With $319\times$ size reduction on *Tanks and Temples* in PSNR and $228\times$ on *Mip-NeRF 360*, it sets a new benchmark for 3DGS compression on these two real-world datasets. A pipeline built around a **simple image-based encoding format** can exceed the strongest published baselines on the main real-world benchmarks while keeping deployment simple. SOG-XT reconstructs standard 3D Gaussian Splatting attributes through standard image decoding and small, deterministic rescaling and reshaping operations. This keeps implementation effort low and makes decoders easy to port across languages and platforms. Optional **encoding-aware fine-tuning** improves rate-distortion without changing this decoder. *Deep Blending* remains the main boundary, where anchor-based learned representations better absorb capture inconsistencies than the standard-splat reconstructions used as input to KISS-GS. Its modular design makes future advances in reconstruction, compaction, and vector quantization directly useful for compression. Still, KISS-GS shows that state-of-the-art 3DGS compression does not require a complex deployment stack: the strongest results can come from a sophisticated encoder while keeping the decoder deliberately *simple*.

Acknowledgements

This work has partly been funded by the European Union’s Horizon Europe research and innovation program (Luminous, grant no. 101135724), the German Research Foundation (3DIL, grant no. 502864329), and the Fraunhofer society (Self Structured Gaussians).

References

1. Bagdasarian, M.T., Knoll, P., Li, Y., Barthel, F., Hilsmann, A., Eisert, P., Morgenstern, W.: 3dgs.zip: A survey on 3d gaussian splatting compression methods. *Computer Graphics Forum* **44**(2), e70078 (2025). <https://doi.org/https://doi.org/10.1111/cgf.70078>, <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.70078>, <https://w-m.github.io/3dgs-compression-survey/>
2. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation (2013)
3. Chen, Y., Wu, Q., Cai, J., Harandi, M., Lin, W.: Hac: Hash-grid assisted context for 3d gaussian splatting compression. In: *ECCV*. pp. 422–438 (2024), https://yihangchen-ee.github.io/project_hac/
4. Chen, Y., Wu, Q., Li, M., Lin, W., Harandi, M., Cai, J.: Fast feedforward 3d gaussian splatting compression. In: *ICLR*. pp. 28308–28321 (2025), <https://arxiv.org/pdf/2410.08017>
5. Chen, Y., Wu, Q., Lin, W., Harandi, M., Cai, J.: Hac++: Towards 100x compression of 3d gaussian splatting. *IEEE TPAMI* (2025), <https://arxiv.org/pdf/2501.12255>
6. Dijkstra, E.W.: On the nature of computing science. EWD896 (manuscript) (Jan 1984), <https://www.cs.utexas.edu/users/EWD/ewd08xx/EWD896.PDF>, eWD896
7. Fang, G., Wang, B.: Mini-splatting: Representing scenes with a constrained number of gaussians. In: *ECCV*. pp. 165–181. Springer (2024)
8. Fang, G., Wang, B.: Mini-splatting2: Building 360 scenes within minutes via aggressive gaussian densification (2024), <https://arxiv.org/abs/2411.12788>
9. Girish, S., Gupta, K., Shrivastava, A.: Eagles: Efficient accelerated 3d gaussians with lightweight encodings. In: *ECCV*. pp. 54–71. Springer (2024)
10. Hahlbohm, F., Franke, L., Eisemann, M., Magnor, M.: Faster-gs: Analyzing and improving gaussian splatting optimization (2026), <https://arxiv.org/abs/2602.09999>
11. Hanson, A., Tu, A., Lin, G., Singla, V., Zwicker, M., Goldstein, T.: Speedy-splat: Fast 3d gaussian splatting with sparse pixels and sparse primitives. In: *CVPR*. pp. 21537–21546 (June 2025), <https://speedysplat.github.io/>
12. Hanson, A., Tu, A., Singla, V., Jayawardhana, M., Zwicker, M., Goldstein, T.: Pup 3d-gs: Principled uncertainty pruning for 3d gaussian splatting. In: *CVPR*. pp. 5949–5958 (June 2025), <https://pup3dgs.github.io/>
13. Hu, J., Li, R., Ye, V., Kanazawa, A.: gsplat compression (2024), <https://github.com/w-m/3dgs-compression-survey/pull/7>, accessed: March 5, 2026
14. Hyung, J., Hong, S., Hwang, S., Lee, J., Choo, J., Kim, J.H.: Effective rank analysis and regularization for enhanced 3d gaussian splatting. In: Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., Zhang, C. (eds.) *NeurIPS*. pp. 110412–110435. Curran Associates, Inc. (2024). <https://doi.org/10.52202/079017-3505>, https://proceedings.neurips.cc/paper_files/paper/2024/file/c73052b864497e078db340c9f6fa4ab5-Paper-Conference.pdf
15. Kellogg, M.: 3d Gaussian Splatting with Three.js (2023), <https://github.com/mkkellogg/gaussian-splats-3d>, self-contained renderer for 3D Gaussian Splatting in web environments. Accessed: March 5, 2026.
16. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM TOG* **42**(4) (July 2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>

17. Kheradmand, S., Rebain, D., Sharma, G., Sun, W., Tseng, Y.C., Isack, H., Kar, A., Tagliasacchi, A., Yi, K.M.: 3d gaussian splatting as markov chain monte carlo. *NeurIPS* **37**, 80965–80986 (2024)
18. Kulhanek, J., Sattler, T.: NerfBaselines: Consistent and reproducible evaluation of novel view synthesis methods. In: *NeurIPS* (2025)
19. Lee, J.C., Rho, D., Sun, X., Ko, J.H., Park, E.: Compact 3d gaussian representation for radiance field (2024), <https://maincold2.github.io/c3dgs/>
20. Lee, S., Kim, Y.G., Sasse, S., Borges, T.M., Sanchez, Y., Ryu, E.S., Schierl, T., Hellge, C.: Gaussianpop: Principled simplification framework for compact 3d gaussian splatting via error quantification (2026), <https://arxiv.org/abs/2602.06830>
21. Lee, S., Shu, F., Sanchez, Y., Schierl, T., Hellge, C.: Compression of 3d gaussian splatting with optimized feature planes and standard video codecs. In: *ICCV*. pp. 25496–25505 (2025)
22. Liu, L., Chen, Z., Xu, D.: Hemgs: A hybrid entropy model for 3d gaussian splatting data compression (2024), <https://arxiv.org/pdf/2411.18473>
23. Liu, X., Wu, X., Zhang, P., Wang, S., Li, Z., Kwong, S.: Compgs: Efficient 3d scene representation via compressed gaussian splatting. In: *ACM MM*. pp. 2936–2944 (2024), <https://www.arxiv.org/abs/2404.09458>
24. Lu, T., Yu, M., Xu, L., Xiangli, Y., Wang, L., Lin, D., Dai, B.: Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In: *CVPR*. pp. 20654–20664 (2024), <https://city-super.github.io/scaffold-gs/>
25. Mallick, S.S., Goel, R., Kerbl, B., Carrasco, F.V., Steinberger, M., Torre, F.D.L.: Taming 3dgs: High-quality radiance fields with limited resources. In: *SIGGRAPH Asia Conference Papers* (2024). <https://doi.org/10.1145/3680528.3687694>, <https://humansensinglab.github.io/taming-3dgs/>
26. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: *ECCV* (2020)
27. Morgenstern, W., Barthel, F., Hilsmann, A., Eisert, P.: Compact 3d scene representation via self-organizing gaussian grids (2024), <https://fraunhoferhhi.github.io/Self-Organizing-Gaussians/>
28. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG* **41**(4) (Jul 2022). <https://doi.org/10.1145/3528223.3530127>
29. Navaneet, K., Meibodi, K.P., Koohpayegani, S.A., Pirsiavash, H.: Compact3d: Smaller and faster gaussian splatting with vector quantization (2023), <https://ucdvision.github.io/compact3d/>
30. Navaneet, K., Pourahmadi Meibodi, K., Abbasi Koohpayegani, S., Pirsiavash, H.: Compgs: Smaller and faster gaussian splatting with vector quantization. In: *ECCV*. pp. 330–349. Springer (2024)
31. Niantic Labs: SPZ: Open-sourcing the JPG for 3D Gaussian Splats (2024), <https://github.com/nianticlabs/spz>, open-source file format for compressed 3D Gaussian Splatting. Accessed: March 5, 2026.
32. Niedermayr, S., Stumpfegger, J., Westermann, R.: Compressed 3d gaussian splatting for accelerated novel view synthesis. In: *CVPR*. pp. 10349–10358 (2024), <https://keksboter.github.io/c3dgs/>
33. Papantonakis, P., Kopanas, G., Kerbl, B., Lanvin, A., Drettakis, G.: 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/

34. PlayCanvas Team: PlayCanvas: Web-First 3d Engine (2024), <https://playcanvas.com/>, open-source JavaScript run-time with native 3D Gaussian Splatting support. Accessed: March 5, 2026.
35. PlayCanvas Team: The SOG Format - PlayCanvas Developer Site (2025), <https://developer.playcanvas.com/user-manual/gaussian-splatting/formats/sog/>, technical specification for Spatially Ordered Gaussians (SOG). Accessed: March 5, 2026.
36. Ren, K., Jiang, L., Lu, T., Yu, M., Xu, L., Ni, Z., Dai, B.: Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. PAMI (2025). <https://doi.org/10.1109/TPAMI.2025.3568201>
37. Rich, B.R.: Clarence leonard (kelly) johnson 1910-1990: a biographical memoir. Biographical Memoirs **67**, 221–241 (1995)
38. Roy, O., Vetterli, M.: The effective rank: A measure of effective dimensionality. In: European Signal Processing Conference. pp. 606–610 (2007)
39. SparkJS contributors: Sparkjs. <https://sparkjs.dev/> (2025), <https://sparkjs.dev/>, accessed: March 5, 2026
40. Studio, L.: A high-performance c++ and cuda implementation of 3d gaussian splatting (2025), <https://github.com/MrNeRF/LichtFeld-Studio>, accessed: March 5, 2026
41. Wang, H., Zhu, H., He, T., Feng, R., Deng, J., Bian, J., Chen, Z.: End-to-end rate-distortion optimized 3d gaussian representation. In: ECCV. pp. 76–92 (2024), <https://rdogaussian.github.io/>
42. Wang, Y., Li, Z., Guo, L., Yang, W., Kot, A.C., Wen, B.: Contextgs: Compact 3d gaussian splatting with anchor level context model. In: NeurIPS. pp. 51532–51551 (2024), <https://github.com/wyf0912/ContextGS>
43. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. TIP (2004)
44. Wu, M., Tuytelaars, T.: Implicit gaussian splatting with efficient multi-level tri-plane representation (2024), <https://www.arxiv.org/abs/2408.10041>
45. Ye, V., Li, R., Kerr, J., Turkulainen, M., Yi, B., Pan, Z., Seiskari, O., Ye, J., Hu, J., Tancik, M., Kanazawa, A.: gsplat: An open-source library for gaussian splatting. Journal of Machine Learning Research **26**(34), 1–17 (2025)
46. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)
47. Zhang, Y., Jia, W., Niu, W., Yin, M.: Gaussianspa: An" optimizing-sparsifying" simplification framework for compact and high-quality 3d gaussian splatting. In: CVPR. pp. 26673–26682 (2025)