# MoE-GS: Shared Appearance and Selection for Quantized Gaussian Images

**Anonymous authors**
Paper under double-blind review

## Abstract

Explicit Gaussian image representations support direct spatial rendering, but their quality–storage trade-off depends on the precision of every stored attribute. We investigate the interaction between shared appearance, quantization, and truncated rendering in MoE-GS. Each Gaussian combines a small shared color table through continuous mixture coordinates, while a learned positive amplitude participates in both top-$K$ selection and normalized blending. We formulate the storage break-even condition, characterize the color and selection constraints, and optimize the representation through its quantized inference path. A self-contained bitstream makes the reported quality reproducible from the stored bytes alone. Our evaluation distinguishes 16-bit uniform quantization from IEEE FP16, combines fixed-budget joint fitting with fixed-geometry conditional controls, and measures component interventions without changing geometry. On 45 matched images at a common 0.30-bpp payload cap, the four-expert model has a +0.55 dB mean PSNR difference relative to uniform-16 RGB, with rates and qualities measured from decoded files. On a predefined five-image diagnostic, direct RGB10 with the same amplitude exceeds the primary shared model by 0.36 dB mean PSNR. The results delimit when appearance sharing is useful and when its quantization error outweighs its storage savings. Region queries retain the global tile partition; their cost includes candidate discovery rather than only the final rasterization.

## 1 Introduction

An explicit image representation is useful when a system needs to reconstruct selected spatial regions without evaluating a coordinate network. Two-dimensional Gaussian primitives provide one such representation: their geometry determines spatial support, their features determine color, and a tile rasterizer reconstructs the requested pixels. Image-GS (Zhang et al., 2025) demonstrates this approach with content-adaptive primitive placement. The question we study is whether shared appearance and learned response amplitudes can improve its quality–storage trade-off *after the parameters are actually quantized and decoded*.

This qualification matters. A parameter-count comparison can assign a low bit width to a tensor while evaluating its full-precision values. Such a comparison measures a hypothetical capacity, not an operational representation. Even a correctly packed file is insufficient evidence if its reconstruction quality is taken from a different evaluation path. Precision also changes the appropriate baseline: replacing 32-bit RGB with a compact mixture may appear to create a large budget advantage, but Image-GS with all eight attributes stored at 16 bits already uses only 128 bits per primitive.

MoE-GS couples two changes to the representation. First, instead of storing an independent RGB vector, a primitive stores mixture coordinates over a shared color table. Second, a positive amplitude multiplies its spatial response *before* top-$K$ truncation and remains in the normalized blend. The amplitude therefore changes both the selected set and the relative contribution within that set. These changes are optimized jointly with Gaussian geometry, but the target primitive count and field precisions remain external operating-point choices.

Our emphasis is on the interaction, rather than on claiming that a soft dictionary or a scalar amplitude is individually new. A palette reduces storage only if its coordinates are cheaper than the RGB features it replaces; a coarsely quantized mixture can erase the benefit of a larger primitive count.



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Similarly, an amplitude can change membership, but this fact alone does not imply improved quality or more stable neighboring selections. We make these restrictions explicit and test the representation under a common decoder and optimization protocol.

The study makes three contributions. (i) A quantization-aware shared-appearance Gaussian representation with a self-contained storage format and a common inference path for training, evaluation, and decoding. (ii) A storage and expressivity analysis that distinguishes color-table capacity, coordinate precision, and pre-truncation amplitudes. (iii) An evaluation protocol with strong low-precision RGB controls, exact byte counts, paired image-level comparisons, and regional-access measurements that include the work needed to find candidates. Positive and negative results are both retained so that the scope of the representation advantage is explicit.

## 2 Related work

**Gaussian image representations.** GaussianImage (Zhang et al., 2024) and Image-GS (Zhang et al., 2025) establish efficient explicit 2D Gaussian representations. Large Images Are Gaussians (Zhu et al., 2025) studies high-resolution fitting. GaussianImage++ (Li et al., 2026) combines content-dependent densification with attribute-specific compression, making it particularly relevant to precision-aware comparisons. SGI (Pan et al., 2026) generates structured Gaussians from spatial seeds and lightweight networks, then compresses the seed representation. SAD (Iinbor et al., 2026) represents images with soft anisotropic site diagrams, top-$K$ propagation, and learned site-specific ownership parameters; its normalized competition is especially close to the selection component studied here. These methods address storage at different levels: individual attributes, spatial structure, and shared generation. MoE-GS shares a small appearance table while retaining explicit per-primitive geometry and a normalized tile renderer. We do not infer superiority over a system from results obtained with different datasets, bit accounting, or decoders.

**Dictionary and mixture representations.** The color table is a soft dictionary: it represents a Gaussian's color as a convex combination of shared vectors. Hard VQ (van den Oord et al., 2017) restricts the combination to a single codeword; residual VQ represents a sum of codewords from successive dictionaries. These are meaningful alternatives, not interchangeable names for the same storage assignment. A post-hoc VQ diagnostic at fixed geometry answers whether a fitted appearance field can be replaced; it does not establish the performance of a jointly trained VQ model. We distinguish both settings. Sparse steered mixtures have long been used for image coding (Verhack et al., 2016); rasterized mixtures (Li et al., 2025) further connect this formulation to efficient normalized local aggregation.

**Learned image and texture compression.** Learned codecs (Ballé et al., 2016; Cheng et al., 2020) optimize image-level compression; implicit representations (Sitzmann et al., 2020; Dupont et al., 2021) store an image in network parameters. Random-access neural texture compression (Vaidyanathan et al., 2023) explicitly targets regional access with a structured latent representation and neural decoding. Our experiments do not equate fast warm rasterization with a competitive complete compression system. Conventional AVIF results are included under actual-file accounting, while likelihood-based codec results are not labeled as file sizes.

## 3 Representation and storage

### 3.1 Baseline renderer and the proposed coupling

A primitive has a center $\mu_i$, spatial scale factors, a rotation, and a color $f_i \in \mathbb{R}^3$. Centers and query positions in the equations are in pixels; storage uses $(\mu_{ix}/W, \mu_{iy}/H)$. The covariance is $\Sigma_i = R_i \operatorname{diag}(s_{i1}^2, s_{i2}^2) R_i^\top$; the inherited unconstrained scale factors may be signed. Its local response is

$$\alpha_i(p) = \exp[-\tfrac{1}{2}(p-\mu_i)^\top \Sigma_i^{-1}(p-\mu_i)]. \tag{1}$$

The implementation first assigns finite footprints to a fixed global grid of $16 \times 16$ tiles. Let $\mathcal{C}(p)$ be the candidates indexed by the tile containing $p$. The inherited renderer selects $\mathcal{S}_0(p) = \mathrm{TopK}_{i \in \mathcal{C}(p)}\, \alpha_i(p)$

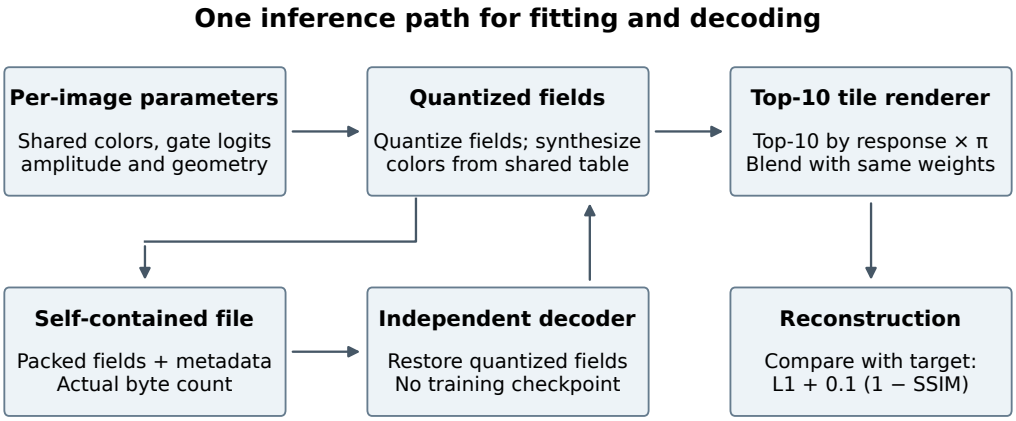



Figure 1: The training and independent decoding paths share the same quantized inference representation. Mixture logits and inverse scales are optimization variables; the file instead stores quantized mixture coordinates and physical scales.

and computes

$$I_0(p) = \frac{\sum_{i \in \mathcal{S}_0(p)} \alpha_i(p) f_i}{\epsilon + \sum_{i \in \mathcal{S}_0(p)} \alpha_i(p)}, \qquad K = 10, \quad \epsilon = 10^{-4}. \tag{2}$$

Footprint truncation and tile membership are part of the operational renderer; we do not describe the infinite-support Gaussian function as an exact spatial index.

**Shared continuous appearance.** Let $M \in \mathbb{R}^{E \times 3}$ contain shared expert colors. During fitting,

$$g_i = \mathrm{softmax}(z_i/T), \qquad f_i = g_i M. \tag{3}$$

The stored representation contains quantized *mixture coordinates*, not the logits $z_i$. Coordinates are quantized in $[0, 1]$ and renormalized before color synthesis. The expert values are learned and are not constrained to $[0, 1]$; clipping is applied to reconstructed images for evaluation. Initial logits favor nearby sampled colors, but finite-temperature softmax does not assign a primitive exclusively to its nearest expert.

**Nonredundant simplex encoding.** For four experts, the sum-to-one constraint permits storing only three coordinates. As an additional diagnostic, let $L = 2^{10} - 1$ and $q_j = \mathrm{round}(L g_j)$ for $j = 1, 2, 3$. If their sum exceeds $L$, we subtract the excess from the largest $q_j$; the excess is at most one. Set $q_4 = L - \sum_{j=1}^{3} q_j$ and reconstruct $f_i = \sum_{j=1}^{4} (q_j/L) M_j$. This uses 30 coordinate bits and exactly preserves simplex membership. It defines a learned affine color lattice inside a tetrahedron, rather than an unconstrained RGB grid. The complete payload is $120N + 192$ bits with 16-bit geometry and a 10-bit amplitude. RGB10 plus the same amplitude uses $120N$ bits, providing a direct control for whether the learned color lattice is useful. The independent-coordinate four-expert configuration remains the primary comparison.

**Pre-truncation amplitude.** For a learned raw variable $u_i$, let

$$\pi_i = \mathrm{softplus}(u_i) + 10^{-6}, \quad s_i(p) = \pi_i \alpha_i(p), \quad \mathcal{S}_\pi(p) = \mathrm{TopK}_{i \in \mathcal{C}(p)}\, s_i(p). \tag{4}$$

MoE-GS reconstructs $I(p)$ by replacing $\alpha_i$ and $\mathcal{S}_0$ in Eq. 2 with $s_i$ and $\mathcal{S}_\pi$. The stored field is the quantized positive amplitude, not $u_i$. This is a learned response amplitude in a normalized RBF-like renderer. It is distinct from depth-ordered opacity and transmittance, but does not constitute a new general weighting mechanism.

### 3.2 WHAT THE PARAMETERIZATION CAN AND CANNOT PROVIDE

**Proposition 1** (Color restriction)**.** *For nonnegative coordinates normalized to one, every synthesized color lies in* $\mathrm{conv}\{M_1, \ldots, M_E\}$*. With at most three experts, this set has affine dimension at most two. Four affinely independent experts can span a three-dimensional color volume.*

This observation motivates examining four higher-precision coordinates alongside eight lower-precision coordinates under the same total coordinate budget. It does not guarantee that a particular learned tetrahedron covers the target colors or that four experts are empirically optimal. The restriction is conditional on the table: with unrestricted learned expert values, four experts can represent any finite collection of RGB feature vectors before quantization (Appendix proof). Thus a continuous capacity argument alone cannot justify a sharing advantage; coordinate precision and optimization matter.

**Proposition 2** (Storage break-even)**.** *Let geometry have the same precision in both methods, let independent RGB require* $3b_f$ *bits per primitive, and let a mixture use* $Eb_g$ *coordinate bits and* $b_\pi$ *amplitude bits. Including an* $E \times 3$ *expert table with precision* $b_e$*, the complete field payload is smaller exactly when*

$$N(3b_f - Eb_g - b_\pi) > 3Eb_e. \tag{5}$$

For $b_f = 16$, $E = 8$, $b_g = 4$, $b_\pi = 10$, and $b_e = 16$, the saving is $6N - 384$ bits. Thus the asymptotic saving relative to an all-16-bit RGB Gaussian is $6/128 = 4.69\%$ of the *total primitive payload*. A $66.7\%$ figure applies only to replacing a 96-bit RGB field with 32 coordinate bits, before table and amplitude costs. Four experts with eight-bit coordinates retain the same per-primitive assignment and halve the shared-table overhead. Headers are additional and are included in the measured file rate.

**Proposition 3** (Pairwise selection condition)**.** *For two candidates with positive responses,* $s_i(p) > s_j(p)$ *exactly when*

$$\log \pi_i - \log \pi_j > \log \alpha_j(p) - \log \alpha_i(p). \tag{6}$$

*A pairwise reversal changes top-*$K$ *membership only if it crosses the selected-set boundary.*

The amplitude is an additive bias in log-response space. This characterization does not imply that learned margins grow, or that neighboring pixels have more stable selected sets. We measure membership changes directly. Multiplying all amplitudes by a common positive constant leaves the ranking unchanged; the blend is exactly invariant only when $\epsilon = 0$, and approximately invariant when its denominator dominates $\epsilon$.

### 3.3 QUANTIZATION-AWARE FITTING AND INDEPENDENT DECODING

For an assigned range $[a, b]$, uniform quantization maps $x$ to

$$Q_q(x) = a + \frac{b-a}{2^q - 1} \,\mathrm{round}\left(\frac{(2^q - 1)(x - a)}{b - a}\right). \tag{7}$$

Constant fields are stored without division by a zero range. Thirty-two-bit fields are raw IEEE float32; FP16 controls use an actual float16 round trip. A straight-through estimator (Bengio et al., 2013) differentiates through the corresponding decoded value. Mixture coordinates use the fixed range $[0, 1]$; other uniform fields use their tensor range. The min/max values needed by the decoder are serialized.

We optimize $\mathcal{L} = \|I - \mathrm{GT}\|_1 + 0.1(1 - \mathrm{SSIM}(I, \mathrm{GT}))$. All new joint-fitting runs use 10,000 steps; continuous fields have a 5,000-step unquantized warm-up, rather than a stopping criterion that can terminate before quantization. Progressive placement begins with half the final count and adds primitives at steps 500, 1000, 1500, and 2000. Positions are sampled from an error map following the Image-GS procedure. Geometry, table values, coordinates, and amplitudes are fitted jointly. Precision, expert count, and final primitive count are not learned discrete decisions.

The file stores physical pixel scales, normalized centers, rotations, optional amplitudes, and either RGB values or a color table and coordinates. It also stores image dimensions, gamma, shapes, precision tags, field offsets, quantization ranges, and renderer version. Decoding never reads a checkpoint or source image. The file rate is $8 \times$ file bytes$/(HW)$; fixed-width payload rate is separately reported. A generic lossless backend is an additional storage result, not a learned entropy model. Every accepted experimental row must pass an encode–decode rendering comparison against the quantized inference path.

| Method | P-bpp | File bpp | PSNR ↑ | MS-SSIM ↑ | Edge ↓ |
|---|---|---|---|---|---|
| Image-GS (uniform 16) | 0.29999 | 0.30132 | 33.23 | 0.9657 | 0.0729 |
| Image-GS (IEEE FP16) | 0.29998 | 0.30092 | 31.75 | 0.9599 | 0.0847 |
| MoE-GS (4 experts, 8-bit) | 0.29999 | 0.30181 | 33.78 | 0.9689 | 0.0681 |

Table 1: Decoded reconstruction on the same 45 images. P-bpp is fixed-width payload; file bpp includes the complete self-contained header. Common target payload: 0.30 bpp.

# 4 EXPERIMENTS

## 4.1 PROTOCOL AND EVALUATION SCOPE

The source benchmark contains 45 RGB images at $2048^2$ resolution across anime, art, painting, photograph, and vector categories, following Image-GS. We predefine a diagnostic set by selecting the first filename in each category: anime-1, art-1, painting-1, photo-1, and vector-1. Configuration diagnostics and dataset evaluation are labeled separately; incomplete dataset runs are never presented as a 45-image average. The central comparison contains 45 completed matched images. Averages use the intersection shared by every listed method.

The primary dataset comparison uses seed 123. The component study has complete three-seed joint fits on 5 predefined images (seeds 123–125); seed means are formed within image before bootstrapping. The primary image-bootstrap interval does not quantify fitting-seed uncertainty. Conditional fixed-geometry controls also use seeds 123–125.

We report PSNR, MS-SSIM (Wang et al., 2003), and edge-band error. The latter sums channel-mean absolute error inside a $7 \times 7$ dilation of strong Sobel-luma edges and divides by the gradient magnitude summed over the same band. It is not a metric that uniquely distinguishes ringing from blur or displacement. Training uses single-scale SSIM with a valid $11 \times 11$ Gaussian window; evaluation uses MS-SSIM. We evaluate decoded floating-point reconstructions, rather than JPEG previews. All controlled methods use the same renderer, input preprocessing, optimizer family, and scheduled learning-rate decay. The Image-GS rows are controlled reimplementations: they retain its RGB representation and renderer but use the common QAT schedule and canonical physical-scale storage. Its official source implements uniform quantization; IEEE FP16 is evaluated as a separate storage choice. Appendix A provides exact settings and limitations.

**Comparison axes.** The primary RGB controls use uniform 16-bit attributes and IEEE FP16 attributes. Reducing RGB precision while leaving geometry at 32 bits is a separate diagnostic, not a replacement for these baselines. At a common target payload $B$, the integer count is computed from the complete field assignment, including the shared table. Exact payload and file rates remain visible; nearly equal targets are not described as identical byte counts. Fixed-count comparisons hold $N$ constant and explicitly expose the resulting rate differences.

## 4.2 DECODED RATE AND RECONSTRUCTION QUALITY

Relative to uniform-16 RGB, the paired mean PSNR difference of the four-expert representation is +0.55 dB (95% image-bootstrap interval [+0.42, +0.67]); it wins on 41/45 images. Excluding the five configuration-diagnostic images leaves 40 images: +0.58 dB (95% image-bootstrap interval [+0.46, +0.71]); it wins on 37/40 images. On the same decoded images, mean LPIPS changes from 0.1344 for uniform-16 RGB to 0.1189 for MoE-GS, while mean FLIP changes from 0.07382 to 0.07435. Both metrics are lower-is-better; the perceptual measures therefore do not show a uniform advantage (Appendix Tables 10–11). The comparison must be read jointly with actual file rate: headers produce small method-dependent offsets even when the payload cap is common. Zlib rates are provided separately in the appendix. The FP16 and uniform-16 controls are both retained because equal nominal precision does not imply equal quantization error. In particular, normalized positions near one have coarser absolute spacing in FP16 than in a 16-bit uniform grid.

On five fixed diagnostic checkpoints, replacing only center storage with FP16 changes mean PSNR by -1.14 dB, versus -1.14 dB when all fields are replaced (Appendix Table 20). These are inference-only substitutions, separate from the independently trained controls. At 2048 pixels, FP16 spacing

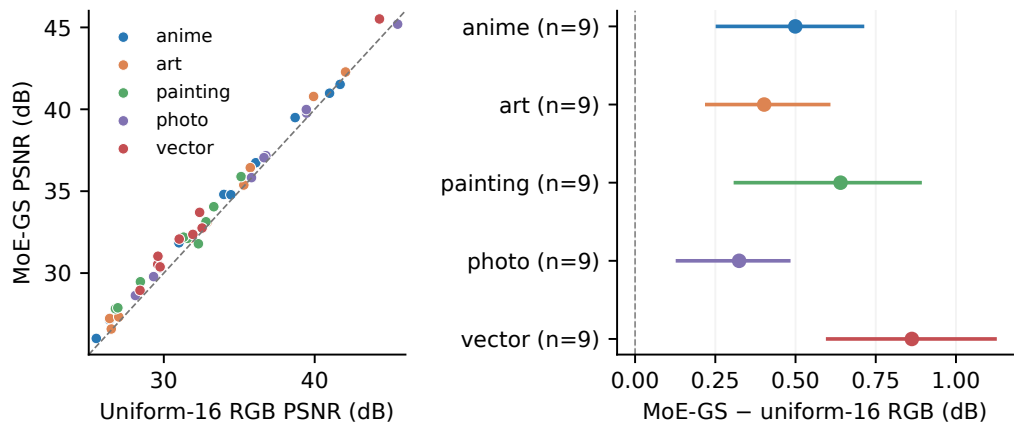



Figure 2: Paired decoded quality on all 45 images. Left: one point per image, with the diagonal indicating equal PSNR. Right: category means and paired image-bootstrap 95% intervals, with nine images per category. Detailed per-image rates and the lowest differences are retained in the appendix.

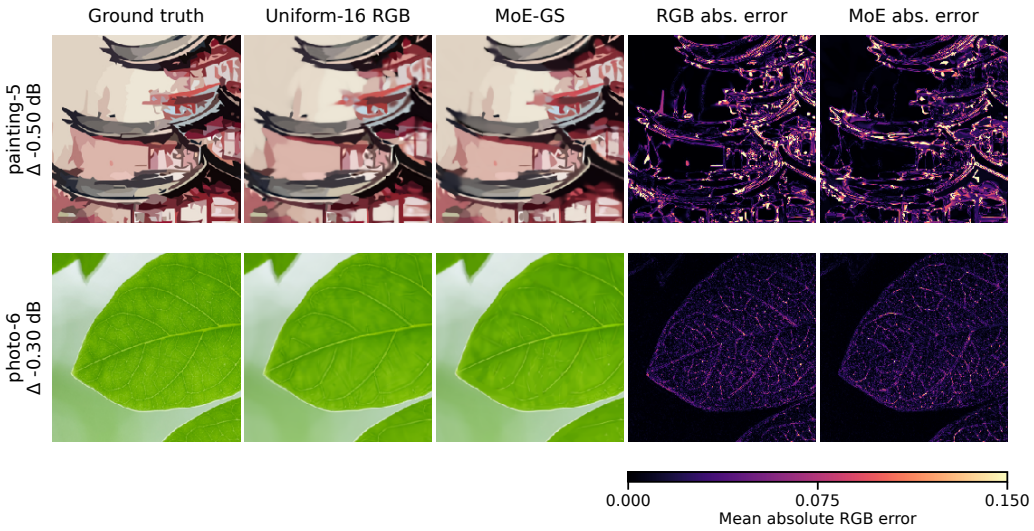



Figure 3: The two lowest full-image MoE-GS minus uniform-16 PSNR differences. For each image we show the nonoverlapping $256 \times 256$ cell with the largest MoE-GS absolute RGB error. Both error maps use the same scale, from black (zero) to bright (0.15). Selection rules and crop coordinates are included with the results.

for normalized centers in $[1/2, 1)$ is one pixel, whereas uniform-16 spacing over a unit interval is approximately $0.031$ pixels.

**Coordinate precision and expert count.** A larger color table does not by itself provide a better low-bit representation. At a fixed 32-bit coordinate budget, eight four-bit coordinates trade precision for more vertices, while four eight-bit coordinates trade vertices for resolution. On 5 common diagnostic images, changing eight 4-bit coordinates to four 8-bit coordinates changes mean PSNR by +2.06 dB, at the same 32-bit coordinate budget. The three-coordinate simplex diagnostic further removes the redundant coordinate: Against the RGB10-plus-amplitude control on 5 diagnostic images, the simplex encoding changes mean PSNR by -0.06 dB. Its direct RGB control is essential because compact storage alone does not establish a benefit from sharing.

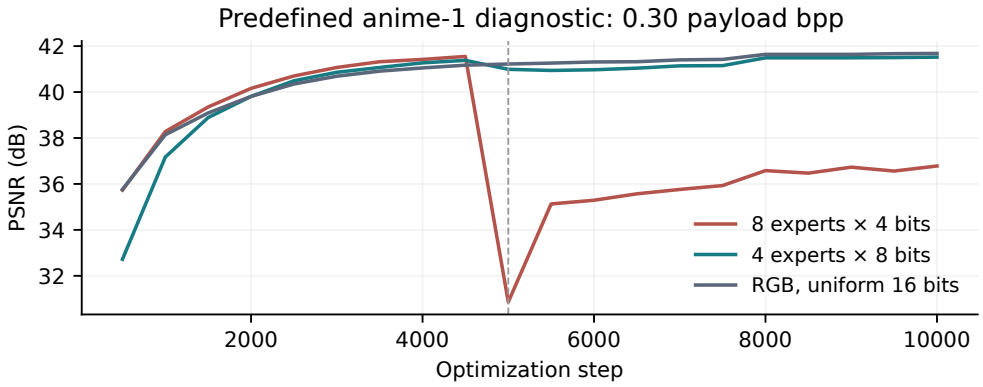



Figure 4: Measured optimization traces on the predefined anime-1 diagnostic image. The dashed line enables quantization at step 5000. Both factored variants assign 32 bits per primitive to mixture coordinates. These single-image traces are diagnostic, not dataset averages.

| Method | P-bpp | File bpp | PSNR ↑ | MS-SSIM ↑ | Edge ↓ |
|---|---|---|---|---|---|
| Image-GS (uniform 16) | 0.29999 | 0.30132 | 33.87 | 0.9643 | 0.0729 |
| Shared appearance only | 0.30000 | 0.30155 | 34.04 | 0.9667 | 0.0707 |
| RGB16 + amplitude | 0.30000 | 0.30160 | 34.08 | 0.9651 | 0.0711 |
| MoE-GS (4 experts, 8-bit) | 0.29999 | 0.30181 | 34.16 | 0.9674 | 0.0697 |

Table 2: Complete joint-fitting component controls: 5 predefined images, three seeds (123–125), 10K steps per run, and a common 0.30-bpp payload cap. Primitive counts differ by field assignment. Means first average seeds within image.

### 4.3 Attribution and Regional Access

With full joint retraining across three seeds, the complete model relative to RGB16-plus-amplitude has a paired PSNR difference of +0.08 dB (95% image-bootstrap interval [-0.20, +0.33]); it wins on 3/5 images. This comparison holds the payload cap fixed; it does not hold primitive count or fitted geometry fixed. Separately, in the seed-123 precision diagnostic, the primary shared model differs from RGB10-plus-amplitude by -0.36 dB on 5 images. Thus the current evidence does not establish shared appearance as necessary for the observed fixed-budget gain. At a fixed count of 9830 primitives with full 10K-step joint retraining on 4 completed images from the five-image diagnostic set and three seeds, the full model relative to RGB16-plus-amplitude has a paired difference of -0.34 dB (95% image-bootstrap interval [-0.58, -0.16]); it wins on 0/4 images (Appendix Table 15). Geometry and growth locations are independently optimized; conditional controls separately hold geometry fixed. The descriptive PSNR interaction $P_{\text{full}} - P_{\pi} - P_{\text{shared}} + P_{\text{base}}$ is -0.09 dB (95% interval [-0.15, -0.03]); the interval is negative under this protocol (Appendix Table 17). When geometry is fixed, switching the amplitude off or restricting it to selection or blending tests a property of a fitted representation. These interventions are not substitutes for independently trained ablations. Likewise, post-hoc VQ/RVQ diagnostics are reported as replacements of appearance with unchanged geometry, rather than as end-to-end trained competitors.

Across 45 fitted images, amplitude changes the nonzero top-10 set at 18.8% of sampled pixels. Applying amplitude in blending alone gives 33.91 dB sampled PSNR, versus 33.93 dB when it also changes selection. Allowing selection to change as well alters sampled PSNR by +0.02 dB. This measures a fixed-checkpoint effect; independently retrained controls are needed to assess necessity. The mean fraction of original selected members replaced is 2.1%.

The decoder scans and projects all primitives for each query; no sublinear candidate-discovery claim is made. Resident unaligned crops agree with full-frame crops within 3.6e-07 maximum absolute error. Separately, CPU dequantization and transfer were checked against CUDA decoding on 135 full frames and 675 queries: maximum pixel difference 4.8e-07. Small field-level floating-point rounding

| Method | P-bpp | File bpp | PSNR ↑ | MS-SSIM ↑ | Edge ↓ |
|---|---|---|---|---|---|
| MoE-GS (4 experts, 8-bit) | 0.29999 | 0.30181 | 34.11 | 0.9670 | 0.0701 |
| MoE-GS (simplex) | 0.30000 | 0.30182 | 34.42 | 0.9686 | 0.0681 |
| RGB10 + amplitude | 0.29998 | 0.30159 | 34.47 | 0.9675 | 0.0680 |

Table 3: Representation diagnostic on 5 predefined images under the same payload cap. The primary model assigns 122 bits per primitive; the last two assign 120. Both shared models additionally store a 192-bit color table.

| Method | Side (px) | Resident GPU (ms) | Cold+readback (ms) |
|---|---|---|---|
| Image-GS (uniform 16) | 16 | 0.81 | 1.90 |
| Image-GS (uniform 16) | 256 | 0.94 | 6.37 |
| Image-GS (uniform 16) | 2048 | 3.76 | 56.43 |
| Image-GS (IEEE FP16) | 16 | 0.83 | 2.09 |
| Image-GS (IEEE FP16) | 256 | 0.92 | 6.47 |
| Image-GS (IEEE FP16) | 2048 | 3.98 | 57.76 |
| MoE-GS (4 experts, 8-bit) | 16 | 0.89 | 7.89 |
| MoE-GS (4 experts, 8-bit) | 256 | 0.97 | 6.14 |
| MoE-GS (4 experts, 8-bit) | 2048 | 3.89 | 70.66 |

Table 4: Median per-image query times on 45 matched images. Resident queries include appearance synthesis, all-primitive projection, candidate construction and rasterization, ending on GPU. Cold+readback also reads the complete file, dequantizes on CPU, transfers all fields to GPU and returns the float32 RGB query to CPU. OS cache is not evicted; external workloads are preserved. Separate timings without readback and per-query p90 values are supplied.

differences are retained in this check. For $16 \times 16$ queries, the measured resident latency is 0.89 ms for MoE-GS and 0.81 ms for uniform-16 RGB; the corresponding complete cold paths with readback take 7.89 and 1.90 ms. The current implementation therefore has an additional cold-access cost for the shared representation. Regional outputs are compared with crops from the *same decoded full image*. For a non-aligned query, we expand to the surrounding global tiles and crop the rendered result. Our current implementation projects all $N$ primitives to discover candidates; its cost therefore contains an $O(N)$ term. Cold measurements mean attributes are not resident on the GPU: they include complete file reading, CPU dequantization, device transfer, and rendering. The operating-system file cache is not evicted. Warm measurements use resident decoded attributes. Neither includes network download.

**Conventional compression reference.** At fixed quality 50 on 45 images, AVIF gives 37.32 dB at 0.29692 actual-file bpp (Appendix Table 5). AVIF is a useful actual-file reference, but it addresses a different complete decoder. We retain its reconstruction results and do not claim that Gaussian representations outperform conventional codecs in compression efficiency. End-to-end random-access comparisons against neural texture codecs would require their complete data layout, cold decode, and hardware-matched measurements; these are outside the validated experiments here.

## 5 Limitations and conclusion

Shared appearance creates a tension between parameter reuse and quantization error. The storage analysis and actual decoder narrow the conditions under which a palette helps: the baseline precision, coordinate resolution, expert overhead, and amplitude cost all matter. Our results are specific to per-image fitting and the tested operating points; they are not evidence of learned generalization to unseen images. A small color table can restrict color variation, and the normalized renderer can lose accuracy when its coordinates are too coarsely quantized. The method incurs per-image optimization and does not yet include a learned entropy model or a persistent sublinear spatial index. The central study covers one payload operating point; rate-dependent crossovers and learned attribute quantizers remain to be evaluated.

At the tested payload cap, the complete representation improves on uniform-16 RGB across most images, while the stronger five-image RGB10-plus-amplitude diagnostic shows that sharing is not necessary to obtain a fixed-budget gain on that set. These findings separate an overall baseline improvement from evidence for a particular component. Operational byte accounting and independent decoding make that distinction testable; fixed-budget, fixed-count, and conditional controls expose the remaining attribution limits.

### AI USE STATEMENT

Generative AI tools were used only to edit the manuscript for readability and language. We have not used generative AI tools to generate synthetic data, develop theoretical models, formulate mathematical claims, assist with proofs, propose or refine hypotheses, design experiments, implement methods, translate text, clean or reformat data, analyze results, interpret findings, draft the paper, create figures, or write code. The authors reviewed the language edits and take responsibility for the final content of this work.

### REPRODUCIBILITY STATEMENT

The supplementary code includes the quantized training path, a versioned decoder, field-level and rendering round-trip checks, experiment configurations, aggregation scripts, and per-image results. Appendix A specifies the datasets, optimization schedule, precision assignments, and measurement boundaries. Source images are not redistributed.

## REFERENCES

Pontus Andersson, Jim Nilsson, Tomas Akenine-Möller, Magnus Oskarsson, Kalle Åström, and Mark D Fairchild. FLIP: A difference evaluator for alternating images. *Proc. ACM Comput. Graph. Interact. Tech.*, 3(2):1–23, 2020. doi: 10.1145/3406183.

Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704*, 2016.

Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 7939–7948, 2020.

Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. COIN: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021.

Laki Iinbor, Zhiyang Dou, and Wojciech Matusik. Soft anisotropic diagrams for differentiable image representation. *ACM Transactions on Graphics*, 2026. URL https://arxiv.org/abs/2604.21984. SIGGRAPH 2026.

Tiantian Li, Xinjie Zhang, Xingtong Ge, Tongda Xu, Dailan He, Jun Zhang, and Yan Wang. GaussianImage++: Boosted image representation and compression with 2D gaussian splatting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 6442–6449, 2026. doi: 10.1609/aaai.v40i8.37572.

Yi-Hsin Li, Mårten Sjöström, Sebastian Knorr, and Thomas Sikora. Rasterized steered mixture of experts for efficient 2D image regression. *arXiv preprint arXiv:2510.05814*, 2025.

Zixuan Pan, Kaiyuan Tang, Jun Xia, Yifan Qin, Lin Gu, Chaoli Wang, Jianxu Chen, and Yiyu Shi. SGI: Structured 2D gaussians for efficient and compact large image representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12162–12172, 2026. URL https://arxiv.org/abs/2603.07789.

Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.

Karthik Vaidyanathan, Marco Salvi, Bartlomiej Wronski, Tomas Akenine-Möller, Pontus Ebelin, and Aaron Lefohn. Random-access neural compression of material textures. *arXiv preprint arXiv:2305.17105*, 2023.

Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.

Ruben Verhack, Thomas Sikora, Lieven Lange, Glenn Van Wallendael, and Peter Lambert. A universal image coding approach using sparse steered mixture-of-experts regression. In *2016 IEEE International Conference on Image Processing (ICIP)*, pp. 2142–2146. IEEE, 2016. doi: 10.1109/ICIP.2016.7532737.

Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers*, volume 2, pp. 1398–1402. IEEE, 2003.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.

Xinjie Zhang, Xingtong Ge, Tongda Xu, Dailan He, Yan Wang, Hongwei Qin, Guo Lu, Jing Geng, and Jun Zhang. GaussianImage: 1000 FPS image representation and compression by 2D gaussian splatting. In *European Conference on Computer Vision*, pp. 327–345. Springer, 2024.

Yunxiang Zhang, Bingxuan Li, Alexandr Kuznetsov, Akshay Jindal, Stavros Diolatzis, Kenneth Chen, Anton Sochenov, Anton Kaplanyan, and Qi Sun. Image-GS: Content-adaptive image representation via 2D gaussians. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pp. 1–11, 2025.

Lingting Zhu, Guying Lin, Jinnan Chen, Xinjie Zhang, Zhenchao Jin, Zhao Wang, and Lequan Yu. Large images are gaussians: High-quality large image representation with levels of 2D gaussian splatting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 10977–10985, 2025.

# A Reproducibility details

**Inputs.** We use the 45 supplied $2048 \times 2048$ PNG images without further resizing, cropping, or augmentation. RGB samples are divided by 255; gamma is one, so no additional linear-light conversion is applied. Each image is fitted independently. Input-file hashes, category names, and run configurations accompany the results.

**Optimization.** The common loss is mean absolute RGB error plus $0.1$ times one minus single-scale SSIM. Adam learning rates are $5 \times 10^{-4}$ for centers, $2 \times 10^{-3}$ for scales and rotation, and $5 \times 10^{-3}$ for RGB, table values, coordinate logits, and amplitude logits. At step 7500 all learning rates are multiplied by $0.1$. Training always reaches step 10,000 unless a run fails; failed runs are excluded and recorded. The forward pass enables quantization at step 5000. The default seed is 123; any repeated-seed studies list their seeds explicitly. Gamma is one. Geometry is optimized through inverse scales but serialized as physical scales. Stored parameters are never reconstructed from optimizer state.

**SSIM implementation.** Early diagnostic fits use the `pytorch_msssim` reference implementation; subsequent fits use a pinned CUDA implementation of the same valid-window SSIM. Each saved configuration identifies the backend. In the primary comparison, all three anime-1 fits use the reference backend and the remaining 44 images use the CUDA backend for all three methods. On the supplied numerical checks at sizes 64, 256 and 2048, the maximum value discrepancy is $5.96 \times 10^{-8}$ and the maximum relative gradient $\ell_2$ discrepancy is $2.08 \times 10^{-6}$. These checks validate the loss implementation, not bitwise-identical optimization trajectories. The source, checks and measurements accompany the code.

**Initialization and growth.** The initial positions follow the same Image-GS sampling rule and NumPy seed: 30% uniform sampling and the remainder sampled using the image-gradient map. At a fixed primitive count this gives the same initial positions; different budget-derived counts do not give identical arrays. Initial rotations are zero and scale factors are five pixels. Independent RGB starts from the target colors at these positions. Shared tables sample these colors and initialize mixture logits to negative squared color distances, at temperature $T = 0.5$. During growth, independent RGB uses the sampled residual color, whereas soft-mixture additions start from zero logits (uniform coordinates); the hard-VQ control initializes new assignment logits from target-to-table distances. Amplitude logits start at zero. Thus the native appearance initialization is parameterization-dependent. The inherited implementation recreates Adam after each growth event for every method. Growth times are shared, but error-driven locations can diverge. The fixed-geometry conditional study additionally holds the decoded geometry and pre-quantization initial color field equal; it does not replay identical growth trajectories from scratch.

**Hardware and timing.** Experiments use NVIDIA RTX 4090 D GPUs with PyTorch 2.4.1 and CUDA runtime 12.4, on an AMD EPYC 7K62 host (48 physical cores, 96 logical CPUs). Evaluation scripts use four PyTorch CPU threads. Training runs may share a host with unrelated jobs, so training wall time is a resource observation, not an isolated throughput benchmark. The fitting timer covers the optimization loop and final coded evaluation; data loading, initialization, model import, and CUDA extension compilation are excluded. Peak memory is the maximum PyTorch allocation in the training process, including initialization but excluding driver overhead and other processes. Query times use wall-clock timers synchronized with CUDA before and after each query. Timing tables state their resident-data and synchronization assumptions. The decoder is versioned with the tile size, top-$K$, normalizer epsilon, and field semantics.

**Scale convention.** The stored scale field is the reciprocal of the optimization variable, in pixel units, and may have either sign in the inherited unconstrained optimizer. The renderer constructs the covariance from $RS(RS)^\top$, so its effective axis lengths are the magnitudes of these factors. We retain their signed values for both training and serialization; taking absolute values only at export would change the uniform quantization grid. This convention is common to every controlled method. Large scale outliers can widen the field's quantization range, which remains a limitation of tensor-range uniform quantization.

**Numerical failures and repeatability.** 3 initial IEEE-FP16 fits (art-2, vector-9, painting-5) reached 10,000 steps but failed export because geometry fields were nonfinite. The recovery rule permits at most one same-seed repeat with unchanged fitting and no quality-based selection. 2 repeats were attempted; 1 completed successfully. The repeat also failed for vector-9. For recoverable original first-attempt states, nonfinite primitives are omitted only if their original tile-hit counts are zero and independent encoding and decoding change no full-frame pixel (maximum absolute difference zero). Allocated/stored primitive counts are vector-9 9830/9829, painting-5 9830/9826. Actual stored counts determine rate; these results do not select retried weights. All failed attempts remain recorded. Each repeat adds 10,000 steps, for 20,000 extra steps excluded from successful-attempt resource summaries. Original wall-time and peak-memory measurements are unavailable for vector-9. Their quality remains included and resource coverage is explicit. For recovered states with resource measurements, the timer ends at the original failed export; the separate recovery time is recorded in metadata. The failed state was not retained for art-2. The numerical root cause remains unresolved. CUDA atomic accumulation is not bitwise deterministic; seed-based retraining can vary. Encoders reject nonfinite values and FP16 overflow, and the independent decoder accepts only finite stored fields.

**Fixed payload targets.** For RGB with 16-bit geometry and features, the payload is $128N$ bits. The original eight-expert complete representation with 16-bit geometry assigns $122N + 384$ bits; its 32-bit geometry counterpart assigns $202N + 384$ bits. The four-expert/eight-bit-coordinate counterpart assigns $122N + 192$ bits. At a target payload $0.3 \times 2048^2$ bits the corresponding primitive counts are 9830, 10310, 6227, and 10312, respectively. The serializer additionally stores its header and per-field byte alignment. These are included in file bpp, not silently absorbed into rounded parameter bpp.

The three-coordinate simplex diagnostic uses $120N + 192$ bits and 10,484 primitives. Its RGB10-plus-amplitude control uses $120N$ bits and 10,485 primitives at the same cap. The one-primitive difference

follows integer budgeting; neither table overhead nor unused budget is counted as effective appearance capacity. The hard-VQ control uses hard assignments throughout fitting, with straight-through gradients through its assignment logits; the 5000-step warm-up applies to its continuous geometry, table and amplitude fields. This is a basic end-to-end assignment control, not an optimized literature VQ codec or a jointly trained RVQ baseline. The end-to-end VQ256 control uses $98N + 12{,}288$ bits: 80 geometry bits, an 8-bit hard index, a 10-bit amplitude, and a 256-entry RGB table with 16-bit values.

**Supplementary completion boundary.** All five predefined images have complete fixed-budget component fits at seeds 123–125. The fixed-count queue stopped starting new fits at the scheduling cutoff; 37 of 45 additional fits completed. Four complete image blocks (anime-1, art-1, painting-1 and photo-1) enter its aggregate. The one completed vector-1 fit remains in the per-image records, but its incomplete block is excluded from that comparison. This boundary follows scheduling order rather than quality. The fixed-geometry conditional study completed all 60 fits.

**Image-level uncertainty.** Paired comparisons use identical images. Bootstrap confidence intervals resample images with replacement, with 10,000 draws and a fixed analysis seed. Repeated training seeds, when available, are summarized within image before resampling; individual images and seeds are not treated as independent samples from the same population. For the five-image diagnostic set the confidence interval is descriptive of that predefined set and is not a population claim. Category differences and worst cases are retained.

**Perceptual evaluation.** Decoded floating-point RGB images are clipped to $[0, 1]$. LPIPS (Zhang et al., 2018) uses the AlexNet backbone and the library's `normalize=True` option, which maps this input to its required range. LDR-FLIP (Andersson et al., 2020) receives $[0, 1]$ RGB with the implementation default of approximately 67.0 pixels per degree. Before perceptual evaluation, the independent decoder must reproduce PSNR within 0.001 dB, MS-SSIM within $2 \times 10^{-5}$, and edge-band error within $10^{-6}$ of the saved metrics. This tolerance concerns metric replay; the encode–decode rendering tolerance remains $2 \times 10^{-6}$ in maximum absolute pixel error.

**Edge-band error.** Let $Y = 0.299R + 0.587G + 0.114B$ be target luma and let $G(p) = \sqrt{(K_x * Y)(p)^2 + (K_y * Y)(p)^2 + 10^{-12}}$, using the standard $3 \times 3$ Sobel kernels with zero padding. Threshold $G$ at its image-wide 90th percentile (including ties), then apply a $7 \times 7$ max-pool dilation to obtain the mask $B(p)$. The reported score is

$$\mathrm{Edge} = \frac{\sum_p B(p)\, \frac{1}{3} \sum_{c=1}^{3} |I_c(p) - \mathrm{GT}_c(p)|}{\sum_p B(p) G(p) + 10^{-12}}.$$

The channel aggregation is a mean, while the spatial aggregation is a sum. This proxy measures edge-band error, including blur and displacement, rather than isolating oscillatory ringing.

**Recent-method protocols.** GaussianImage++ (Li et al., 2026) reports experiments on 24 Kodak images and 100 DIV2K validation images, with 50,000-step Adam fitting and an $L_2$ reconstruction loss. Its compression protocol uses a 6,000-step warm-up and learned LSQ+ quantizers with 12-bit positions, 10-bit covariance attributes and 6-bit colors. These learned quantizer offsets and scales differ from our tensor-range uniform controls. Its published results therefore cannot be inserted into the present 45-image, 10,000-step table as matched measurements.

SGI (Pan et al., 2026) primarily evaluates four FGF2 satellite images (approximately 51 megapixels each), two ICB natural images (27.7 and 39.1 megapixels), and three STimage histopathology images (approximately 76 megapixels), with additional DIV2K experiments. It uses multiscale seed fitting, quantization noise during training, GPCC for seed positions, and arithmetic coding for seed attributes and a binary hash grid. The stored decoder also includes appearance/covariance MLPs and a probability-context MLP. Its complete encoding/decoding pipeline and large-image protocol differ from independently stored primitives; we do not claim a hardware-matched comparison to SGI.

The SAD reference uses the official native implementation and its 4,000-iteration schedule. Its inspected code assigns a nominal 128 bits per site while exporting ten floating-point columns as text. We recompute quality from the saved PNG, retain the original protocol, and label this rate as nominal

rather than a validated low-bit bitstream. This reference is separate from the controlled quantized comparison. Source URLs, code revision and protocol provenance accompany the supplementary results.

The AVIF reference uses Pillow 12.1.1 with pillow-avif-plugin 1.6.0, speed 6, default 4:2:0 subsampling, and four encoding threads. Decoding uses the plugin's automatic codec and thread settings. Its reported CPU time covers complete file reading, decoding and RGB conversion, with five repetitions and no OS-cache eviction. It is a separate decoder measurement, not an isolated four-thread or hardware-matched ROI comparison. Actual AVIF files and codec-version metadata are supplied.

## B PROOFS AND INTERPRETATION

**Color restriction.** Softmax coordinates are nonnegative and sum to one. Quantized coordinates remain nonnegative; after a positive-sum renormalization they again sum to one. Their product with $M$ is therefore a convex combination of table rows. The affine hull of $E$ points has dimension at most $E-1$. Four affinely independent colors form a tetrahedron with nonzero three-dimensional volume. This permits but does not ensure adequate coverage of any target color distribution. No claim requires the table entries themselves to be displayable RGB values.

The proposition concerns primitive colors. Before display clipping, the normalized pixel output with $\epsilon > 0$ belongs to $\mathrm{conv}(\{0\} \cup \{M_e\})$: its expert coefficients sum to the response mass divided by that mass plus $\epsilon$, and the remaining weight multiplies black. Thus four experts are sufficient for a full-dimensional primitive-color hull, not a necessary condition for a full-dimensional set of pixel outputs.

**Unquantized four-expert equivalence.** For any finite collection $\{f_i\} \subset \mathbb{R}^3$, choose $a > \max_i \|f_i\|_\infty$ with $a > 0$. Take the four table rows

$$M_0 = (-3a, -3a, -3a), \quad M_1 = (9a, -3a, -3a),$$
$$M_2 = (-3a, 9a, -3a), \quad M_3 = (-3a, -3a, 9a).$$

For $j = 1, 2, 3$ set $g_{ij} = (f_{ij} + 3a)/(12a)$ and $g_{i0} = 1 - \sum_{j=1}^{3} g_{ij}$. Every coordinate is strictly positive: $f_{ij} > -a$ and $\sum_j f_{ij} < 3a$. They sum to one and direct substitution gives $g_i M = f_i$. For any finite $T > 0$, the finite logits $z_{ij} = T \log g_{ij}$ realize these coordinates through softmax. Therefore, at fixed geometry and without quantization, four unrestricted shared experts can express every finite independent-RGB feature assignment. This is an existence statement, not a guarantee that fitting finds the enclosing table. It also does not guarantee a good quantized representation: larger table extents amplify coordinate-rounding error, and stored expert values have finite precision.

**Storage condition.** Write the shared geometry payload as $NB_g$. Independent RGB consumes $N(B_g + 3b_f)$ bits; the complete factored representation consumes $N(B_g + Eb_g + b_\pi) + 3Eb_e$. Subtracting and requiring a positive difference yields Eq. 5. If $3b_f \le Eb_g + b_\pi$, increasing the primitive count can never amortize the overhead into a payload saving. The proof concerns fixed-width fields; file headers, entropy distributions, spatial indexing, and decoder memory require separate measurements.

**Selection condition.** Both $\alpha$ and $\pi$ are positive for the candidates under consideration. Taking logarithms of $\pi_i \alpha_i > \pi_j \alpha_j$ and rearranging gives Eq. 6. A changed relative order between two already selected candidates, or two excluded candidates, need not change membership. This is why both membership and reconstruction must be measured. If $K$ includes all candidates, the selection effect disappears, although the blending effect remains.

**Coordinate quantization error.** Let $g$ lie on the simplex and let $\widehat{g} = g + e$ be independently rounded coordinates. For quantization step $\delta$, $\|e\|_1 \le E\delta/2$. Put $t = \sum_e e_e$. If $|t| < 1$, renormalization gives $\widetilde{g} - g = (e - tg)/(1 + t)$ and hence

$$\|\widetilde{g} - g\|_1 \le \frac{2\|e\|_1}{1 - |t|} \le \frac{E\delta}{1 - E\delta/2}, \qquad E\delta/2 < 1.$$

| Quality | File bpp | PSNR | CPU decode (ms) |
|---|---|---|---|
| 20 | 0.11062 | 32.48 | 72.0 |
| 30 | 0.14821 | 33.99 | 68.3 |
| 40 | 0.20689 | 35.63 | 74.2 |
| 50 | 0.29692 | 37.32 | 73.8 |
| 60 | 0.41074 | 38.68 | 72.7 |

Table 5: AVIF actual-file reference on 45 images, fixed quality settings and speed 6. Rates vary by image; no per-image quality selection or interpolation is used. CPU time includes complete file reading and RGB conversion, and is not a hardware-matched Gaussian ROI comparison.

| Comparison | $n$ | $\Delta$ dB | 95% interval | Wins |
|---|---|---|---|---|
| Full4 versus RGB-U16 | 45 | +0.55 | [+0.42, +0.67] | 41/45 |
| Full4 versus RGB-FP16 | 45 | +2.02 | [+1.75, +2.35] | 45/45 |
| RGB-U16 versus RGB-FP16 | 45 | +1.48 | [+1.24, +1.77] | 45/45 |

Table 6: Paired image-bootstrap PSNR comparisons, seed123. The payload cap is shared; measured file rates differ slightly and are reported in the main table.

For any reference color $m_0$, the color error is bounded by this quantity times $\max_e \|M_e - m_0\|$, using the zero sum of the coordinate error. This bound is conservative, but it makes the precision–table-size trade-off explicit. It is not a rate–distortion guarantee.

**Simplex remainder error.** Let $\delta = 1/L$. Without overflow correction the first three coordinate errors each have magnitude at most $\delta/2$; the fourth is their negative sum. Thus the total $\ell_1$ error is at most $3\delta$. If overflow occurs, the rounded sum is exactly $L + 1$, because the unrounded sum is at most $L$ and three rounding errors sum to at most $3/2$. Subtracting one from a largest entry preserves nonnegativity. The corrected first-three error has $\ell_1$ norm at most $5\delta/2$, and overflow requires $g_4 \leq \delta/2$. The fourth coordinate is then zero, so again the total error is at most $3\delta$. The color-error bound follows by multiplying by $\max_e \|M_e - m_0\|$. This bound isolates coordinate quantization; it does not cover quantization of the expert table or changes in fitted geometry.

## C Additional numerical results

| Category | $n$ | RGB16 | MoE-GS | Δ dB | Wins |
|---|---|---|---|---|---|
| anime | 9 | 34.31 | 34.81 | +0.50 | 7/9 |
| art | 9 | 33.07 | 33.47 | +0.40 | 9/9 |
| painting | 9 | 30.95 | 31.59 | +0.64 | 8/9 |
| photo | 9 | 35.65 | 35.97 | +0.32 | 8/9 |
| vector | 9 | 32.17 | 33.03 | +0.86 | 9/9 |

Table 7: Category results on the central matched set. Uniform 16-bit RGB is the baseline.

| Image | RGB16 | MoE-GS | Δ dB |
|---|---|---|---|
| painting-5_2k | 32.29 | 31.79 | -0.50 |
| photo-6_2k | 45.50 | 45.20 | -0.30 |
| anime-1_2k | 41.68 | 41.52 | -0.16 |
| anime-4_2k | 40.99 | 40.98 | -0.00 |
| photo-1_2k | 35.81 | 35.83 | +0.02 |

Table 8: Five lowest MoE-GS minus uniform-16 PSNR differences, selected mechanically from all matched images.

| Method | Zlib bpp | Fit (s) | Peak MiB | Coverage |
|---|---|---|---|---|
| Image-GS (uniform 16) | 0.28321 | 362.5 | 1432 | 45/45 |
| Image-GS (IEEE FP16) | 0.26590 | 351.5 | 1432 | 44/45 |
| MoE-GS (4 experts, 8-bit) | 0.27938 | 391.8 | 1441 | 45/45 |

Table 9: Zlib compresses the complete stream. Fit time is the median; peak memory is the maximum process allocation over measured images. Coverage gives images with both resource measurements. Shared-GPU times are observations, not isolated speed comparisons. Failed-attempt costs are disclosed separately; the recovered vector-9 export lacks original wall-time/peak-memory measurements and is excluded only from these resource summaries.

| Method | LPIPS ↓ | FLIP ↓ |
|---|---|---|
| Image-GS (uniform 16) | 0.1344 | 0.0738 |
| Image-GS (IEEE FP16) | 0.1415 | 0.0828 |
| MoE-GS (4 experts, 8-bit) | 0.1189 | 0.0743 |

Table 10: Decoded perceptual metrics on the same 45 central images. LPIPS uses AlexNet and its documented input normalization.

| Metric | Mean difference | 95% interval | Wins |
|---|---|---|---|
| LPIPS | -0.01555 | [-0.01906, -0.01211] | 42/45 |
| FLIP | +0.00053 | [-0.00055, +0.00170] | 18/45 |

Table 11: Paired perceptual differences: complete model minus uniform-16 RGB. Negative differences are better for both metrics; intervals resample images at fixed training seed123.

| Method | P-bpp | File bpp | PSNR ↑ | MS-SSIM ↑ | Edge ↓ |
|---|---|---|---|---|---|
| MoE-GS (32-bit geometry) | 0.29999 | 0.30152 | 30.39 | 0.9332 | 0.1039 |
| MoE-GS (8 experts, 4-bit) | 0.29998 | 0.30180 | 32.05 | 0.9513 | 0.0865 |
| MoE-GS (4 experts, 8-bit) | 0.29999 | 0.30181 | 34.11 | 0.9670 | 0.0701 |

Table 12: Geometry and coordinate precision on 5 common predefined images, with the same payload cap and full 10K-step fitting. The 32-bit geometry row uses eight four-bit coordinates and 6227 primitives; the eight-expert/16-bit row uses 10310 primitives; the four-expert row uses 10312. Changing geometry precision therefore also changes count; this is not a fixed-geometry precision substitution.

| Method | P-bpp | File bpp | PSNR ↑ | MS-SSIM ↑ | Edge ↓ |
|---|---|---|---|---|---|
| MoE-GS (8 experts, 4-bit) | 0.29998 | 0.30180 | 32.05 | 0.9513 | 0.0865 |
| MoE-GS (4 experts, 8-bit) | 0.29999 | 0.30181 | 34.11 | 0.9670 | 0.0701 |

Table 13: Coordinate-precision diagnostic. This subset is explicitly restricted to images with both completed configurations.

| Method | Mean seed SD | Max seed SD |
|---|---|---|
| Image-GS (uniform 16) | 0.023 | 0.050 |
| Shared appearance only | 0.091 | 0.124 |
| RGB16 + amplitude | 0.042 | 0.073 |
| MoE-GS (4 experts, 8-bit) | 0.100 | 0.230 |

Table 14: Within-image PSNR standard deviation across the three complete joint-training seeds, in dB. The image set matches the component table.

| Method | P-bpp | File bpp | PSNR ↑ | MS-SSIM ↑ | Edge ↓ |
|---|---|---|---|---|---|
| Image-GS (uniform 16) | 0.29999 | 0.30132 | 35.22 | 0.9676 | 0.0734 |
| Shared appearance only | 0.26254 | 0.26407 | 34.89 | 0.9664 | 0.0753 |
| RGB16 + amplitude | 0.32342 | 0.32502 | 35.64 | 0.9697 | 0.0702 |
| MoE-GS (4 experts, 8-bit) | 0.28597 | 0.28777 | 35.29 | 0.9686 | 0.0721 |

Table 15: Full joint retraining at a fixed count of 9830 primitives on 4 complete images out of the five predefined images, three seeds and 10K steps. RGB uniform16 reuses the identical-count budget runs. Geometry is independently fitted and growth locations can diverge. Actual payloads and file sizes differ by representation; no padding is counted as capacity.

| Conditional variant | P-bpp | File bpp | PSNR | Seed SD |
|---|---|---|---|---|
| rgb | 0.31470 | 0.31604 | 27.33 | 0.000 |
| rgb + pi | 0.33928 | 0.34089 | 33.41 | 0.001 |
| shared | 0.27541 | 0.27695 | 26.71 | 0.008 |
| shared + pi | 0.29999 | 0.30180 | 33.19 | 0.003 |

Table 16: Fixed decoded geometry and primitive count from the complete model, 5 predefined images and three seeds per image, 2000 appearance-only QAT steps. Seed SD is the mean within-image standard deviation. This is conditional retraining initialized from a common fitted color field, not independent 10K-step fitting. Source geometry was learned with sharing and amplitudes, which can favor retaining these components. Actual capacities differ; no padding is counted as capacity.

| Control | Images | Interaction (dB) | 95% interval |
|---|---|---|---|
| Payload cap, seed123 | 5 | -0.14 | [-0.30, +0.00] |
| Payload cap, three seeds | 5 | -0.09 | [-0.15, -0.03] |
| Fixed count, three seeds | 4 | -0.01 | [-0.04, +0.02] |
| Conditional geometry, three seeds | 5 | +0.40 | [+0.19, +0.63] |

Table 17: Descriptive component interaction: $P_{\text{full}} - P_{\pi} - P_{\text{shared}} + P_{\text{base}}$, where $P$ is PSNR. Positive values indicate a larger incremental gain from sharing when amplitude is present under that protocol. Seeds are averaged within image before bootstrap. PSNR is logarithmic; budget controls can change count and geometry, while conditional controls inherit full-model geometry. This contrast does not establish causal synergy or necessity.

| Model | $n$ | Pixel % | Member % | None | Blend | Select | Both |
|---|---|---|---|---|---|---|---|
| Four experts | 45 | 18.8 | 2.1 | 20.99 | 33.91 | 21.03 | 33.93 |

Table 18: Fixed-checkpoint interventions, using 1024 fixed-seed pixels per image. PSNR is on sampled pixels. Pixel is the fraction whose nonzero top-10 set changes; Member is the mean fraction of original nonzero members replaced, read from actual CUDA membership. Recomputed endpoint quality must agree with CUDA within 0.001 dB.

| Appearance replacement | File bpp | PSNR |
|---|---|---|
| RGB-10 | 0.29664 | 34.10 |
| RGB-12 | 0.31139 | 34.11 |
| RGB-16 | 0.34089 | 34.11 |
| RGB-8 | 0.28189 | 34.04 |
| posthoc-RVQ2x256 | 0.26855 | 33.89 |
| posthoc-RVQ4x256 | 0.31431 | 34.11 |
| posthoc-VQ256 | 0.24565 | 31.47 |

Table 19: Post-hoc replacements at identical decoded geometry and amplitudes on 5 matched source images. Dictionaries fit the source colors, not target pixels. These are not end-to-end VQ/RVQ baselines.

| FP16 replacement | Images | File bpp | PSNR | $\Delta$ dB |
|---|---|---|---|---|
| none | 5 | 0.30132 | 33.86 | +0.000 |
| xy | 5 | 0.30122 | 32.72 | -1.137 |
| scale | 5 | 0.30122 | 33.86 | +0.001 |
| rot | 5 | 0.30122 | 33.86 | -0.000 |
| rgb | 5 | 0.30122 | 33.86 | -0.000 |
| all | 5 | 0.30093 | 32.72 | -1.137 |

Table 20: Inference-only precision substitutions in fixed uniform-16 RGB checkpoints, with no retraining or count change. Each variant is independently encoded and decoded. All fields retain a 16-bit payload; header sizes differ by quantizer type.

| Image | Sites | Nominal bpp | PNG PSNR |
|---|---|---|---|
| anime-1_2k | 9830 | 0.29999 | 42.02 |
| art-1_2k | 9830 | 0.29999 | 33.94 |
| painting-1_2k | 9830 | 0.29999 | 33.28 |
| photo-1_2k | 9830 | 0.29999 | 35.97 |
| vector-1_2k | 9830 | 0.29999 | 32.84 |

Table 21: SAD official native implementation under its original optimization protocol. Its nominal 128-bit/site budgeting is not a measured quantized codec rate. Quality is recomputed from its saved 8-bit PNG. This reference is separate from the controlled decoded-file comparison.

| Method | P-bpp | File bpp | PSNR ↑ | MS-SSIM ↑ | Edge ↓ |
|---|---|---|---|---|---|
| Image-GS (uniform 16) | 0.29999 | 0.30132 | 33.86 | 0.9643 | 0.0730 |
| MoE-GS (4 experts, 8-bit) | 0.29999 | 0.30181 | 34.11 | 0.9670 | 0.0701 |
| End-to-end VQ256 | 0.29999 | 0.30178 | 28.69 | 0.9248 | 0.1203 |

Table 22: End-to-end hard VQ256 control on 5 predefined diagnostic images, jointly fitting its geometry, assignments, table and amplitudes for 10K steps. All methods use the same payload cap.

### C.1 COMPLETE PER-IMAGE RECORDS

Table 23: All independently fitted, verified runs. Full4 and Shared4 use four experts with eight-bit coordinates; Full8 uses eight experts with four-bit coordinates. G16/G32 specifies geometry precision. Repeated rows differ in seed or primitive count. Exact run paths and hashes are supplied in the provenance ledger.

| Image | Variant | Seed | $N$ | File bpp | PSNR | Fit (s) |
|---|---|---|---|---|---|---|
| anime-1_2k | Full8-G16 | 123 | 10310 | 0.30179 | 36.78 | 744.0 |
| anime-1_2k | Full8-G32 | 123 | 6227 | 0.30152 | 34.97 | 697.7 |
| anime-1_2k | Full4 | 123 | 10312 | 0.30181 | 41.52 | 692.5 |
| anime-1_2k | Full4 | 123 | 9830 | 0.28778 | 41.47 | 388.1 |
| anime-1_2k | Full4 | 124 | 9830 | 0.28778 | 41.43 | 453.0 |
| anime-1_2k | Full4 | 125 | 9830 | 0.28778 | 41.59 | 388.7 |
| anime-1_2k | Full4 | 124 | 10312 | 0.30181 | 41.40 | 418.0 |
| anime-1_2k | Full4 | 125 | 10312 | 0.30181 | 41.69 | 385.7 |
| anime-1_2k | Shared4 | 123 | 11233 | 0.30155 | 41.32 | 1146.5 |
| anime-1_2k | Shared4 | 123 | 9830 | 0.26407 | 41.02 | 396.1 |
| anime-1_2k | Shared4 | 124 | 9830 | 0.26408 | 40.89 | 400.3 |
| anime-1_2k | Shared4 | 125 | 9830 | 0.26407 | 41.08 | 393.8 |
| anime-1_2k | Shared4 | 124 | 11233 | 0.30155 | 41.55 | 652.5 |
| anime-1_2k | Shared4 | 125 | 11233 | 0.30154 | 41.52 | 377.5 |
| anime-1_2k | RGB16+$\pi$ | 123 | 9118 | 0.30160 | 41.89 | 716.3 |
| anime-1_2k | RGB16+$\pi$ | 123 | 9830 | 0.32503 | 42.23 | 444.3 |
| anime-1_2k | RGB16+$\pi$ | 124 | 9830 | 0.32503 | 42.22 | 351.1 |
| anime-1_2k | RGB16+$\pi$ | 125 | 9830 | 0.32502 | 42.19 | 384.5 |
| anime-1_2k | RGB16+$\pi$ | 124 | 9118 | 0.30161 | 42.02 | 375.3 |
| anime-1_2k | RGB16+$\pi$ | 125 | 9118 | 0.30160 | 41.89 | 405.9 |
| anime-1_2k | RGB10+$\pi$ | 123 | 10485 | 0.30159 | 42.14 | 259.6 |
| anime-1_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 39.65 | 663.0 |
| anime-1_2k | RGB-U16 | 123 | 9830 | 0.30132 | 41.68 | 1063.5 |
| anime-1_2k | RGB-U16 | 124 | 9830 | 0.30132 | 41.70 | 396.9 |
| anime-1_2k | RGB-U16 | 125 | 9830 | 0.30132 | 41.64 | 403.3 |
| anime-1_2k | Simplex4 | 123 | 10484 | 0.30182 | 41.99 | 290.3 |
| anime-1_2k | VQ256 | 123 | 12714 | 0.30177 | 31.49 | 338.0 |
| anime-2_2k | Full4 | 123 | 10312 | 0.30181 | 34.80 | 349.1 |
| anime-2_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 32.38 | 336.4 |
| anime-2_2k | RGB-U16 | 123 | 9830 | 0.30131 | 33.98 | 383.0 |
| anime-3_2k | Full4 | 123 | 10312 | 0.30181 | 34.78 | 278.7 |
| anime-3_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 33.34 | 341.4 |
| anime-3_2k | RGB-U16 | 123 | 9830 | 0.30132 | 34.45 | 241.4 |
| anime-4_2k | Full4 | 123 | 10312 | 0.30181 | 40.98 | 445.0 |
| anime-4_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 36.98 | 395.1 |
| anime-4_2k | RGB-U16 | 123 | 9830 | 0.30132 | 40.99 | 365.5 |
| anime-5_2k | Full4 | 123 | 10312 | 0.30180 | 39.49 | 376.8 |
| anime-5_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 36.50 | 387.1 |
| anime-5_2k | RGB-U16 | 123 | 9830 | 0.30132 | 38.70 | 362.5 |
| anime-6_2k | Full4 | 123 | 10312 | 0.30181 | 31.85 | 399.1 |
| anime-6_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 30.25 | 362.9 |
| anime-6_2k | RGB-U16 | 123 | 9830 | 0.30132 | 31.00 | 360.5 |
| anime-7_2k | Full4 | 123 | 10312 | 0.30181 | 26.01 | 409.6 |
| anime-7_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 25.00 | 370.6 |
| anime-7_2k | RGB-U16 | 123 | 9830 | 0.30132 | 25.52 | 359.9 |
| anime-8_2k | Full4 | 123 | 10312 | 0.30181 | 27.14 | 412.2 |
| anime-8_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 25.57 | 355.3 |
| anime-8_2k | RGB-U16 | 123 | 9830 | 0.30132 | 26.40 | 355.4 |
| anime-9_2k | Full4 | 123 | 10312 | 0.30180 | 36.74 | 380.4 |

| Image | Variant | Seed | $N$ | File bpp | PSNR | Fit (s) |
|---|---|---|---|---|---|---|
| anime-9_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 34.96 | 226.0 |
| anime-9_2k | RGB-U16 | 123 | 9830 | 0.30132 | 36.09 | 386.3 |
| art-1_2k | Full8-G16 | 123 | 10310 | 0.30180 | 30.50 | 723.6 |
| art-1_2k | Full8-G32 | 123 | 6227 | 0.30152 | 28.75 | 381.3 |
| art-1_2k | Full4 | 123 | 10312 | 0.30181 | 32.13 | 270.3 |
| art-1_2k | Full4 | 123 | 9830 | 0.28777 | 31.92 | 444.4 |
| art-1_2k | Full4 | 124 | 9830 | 0.28777 | 32.08 | 396.7 |
| art-1_2k | Full4 | 125 | 9830 | 0.28777 | 32.08 | 365.3 |
| art-1_2k | Full4 | 124 | 10312 | 0.30180 | 32.10 | 386.7 |
| art-1_2k | Full4 | 125 | 10312 | 0.30180 | 32.20 | 404.4 |
| art-1_2k | Shared4 | 123 | 11233 | 0.30155 | 32.12 | 365.8 |
| art-1_2k | Shared4 | 123 | 9830 | 0.26407 | 31.52 | 324.7 |
| art-1_2k | Shared4 | 124 | 9830 | 0.26407 | 31.58 | 374.2 |
| art-1_2k | Shared4 | 125 | 9830 | 0.26407 | 31.52 | 375.8 |
| art-1_2k | Shared4 | 124 | 11233 | 0.30154 | 32.09 | 391.9 |
| art-1_2k | Shared4 | 125 | 11233 | 0.30155 | 31.94 | 433.0 |
| art-1_2k | RGB16+$\pi$ | 123 | 9118 | 0.30160 | 32.09 | 268.7 |
| art-1_2k | RGB16+$\pi$ | 123 | 9830 | 0.32502 | 32.41 | 360.6 |
| art-1_2k | RGB16+$\pi$ | 124 | 9830 | 0.32502 | 32.33 | 372.1 |
| art-1_2k | RGB16+$\pi$ | 125 | 9830 | 0.32502 | 32.32 | 385.6 |
| art-1_2k | RGB16+$\pi$ | 124 | 9118 | 0.30160 | 32.00 | 404.5 |
| art-1_2k | RGB16+$\pi$ | 125 | 9118 | 0.30159 | 32.05 | 360.3 |
| art-1_2k | RGB10+$\pi$ | 123 | 10485 | 0.30159 | 32.53 | 248.6 |
| art-1_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 30.42 | 341.3 |
| art-1_2k | RGB-U16 | 123 | 9830 | 0.30132 | 31.81 | 395.5 |
| art-1_2k | RGB-U16 | 124 | 9830 | 0.30132 | 31.85 | 332.5 |
| art-1_2k | RGB-U16 | 125 | 9830 | 0.30132 | 31.81 | 421.1 |
| art-1_2k | Simplex4 | 123 | 10484 | 0.30181 | 32.44 | 281.2 |
| art-1_2k | VQ256 | 123 | 12714 | 0.30178 | 27.63 | 443.4 |
| art-2_2k | Full4 | 123 | 10312 | 0.30181 | 42.28 | 390.0 |
| art-2_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 38.84 | 356.3 |
| art-2_2k | RGB-U16 | 123 | 9830 | 0.30132 | 42.04 | 251.9 |
| art-3_2k | Full4 | 123 | 10312 | 0.30181 | 36.44 | 264.4 |
| art-3_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 33.73 | 316.2 |
| art-3_2k | RGB-U16 | 123 | 9830 | 0.30132 | 35.73 | 409.6 |
| art-4_2k | Full4 | 123 | 10312 | 0.30181 | 33.11 | 394.4 |
| art-4_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 31.94 | 347.0 |
| art-4_2k | RGB-U16 | 123 | 9830 | 0.30132 | 32.87 | 391.7 |
| art-5_2k | Full4 | 123 | 10312 | 0.30181 | 26.59 | 399.5 |
| art-5_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 25.77 | 395.0 |
| art-5_2k | RGB-U16 | 123 | 9830 | 0.30132 | 26.52 | 357.3 |
| art-6_2k | Full4 | 123 | 10312 | 0.30181 | 27.23 | 411.0 |
| art-6_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 24.98 | 365.4 |
| art-6_2k | RGB-U16 | 123 | 9830 | 0.30132 | 26.41 | 398.0 |
| art-7_2k | Full4 | 123 | 10312 | 0.30180 | 27.32 | 393.8 |
| art-7_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 26.33 | 361.2 |
| art-7_2k | RGB-U16 | 123 | 9830 | 0.30132 | 27.01 | 409.9 |
| art-8_2k | Full4 | 123 | 10312 | 0.30181 | 35.38 | 391.3 |
| art-8_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 34.49 | 365.3 |
| art-8_2k | RGB-U16 | 123 | 9830 | 0.30132 | 35.30 | 366.3 |
| art-9_2k | Full4 | 123 | 10312 | 0.30181 | 40.78 | 273.6 |
| art-9_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 38.21 | 249.6 |
| art-9_2k | RGB-U16 | 123 | 9830 | 0.30132 | 39.93 | 267.7 |
| painting-1_2k | Full8-G16 | 123 | 10310 | 0.30180 | 30.60 | 377.5 |
| painting-1_2k | Full8-G32 | 123 | 6227 | 0.30152 | 28.86 | 382.7 |
| painting-1_2k | Full4 | 123 | 10312 | 0.30180 | 32.12 | 270.1 |
| painting-1_2k | Full4 | 123 | 9830 | 0.28778 | 31.78 | 366.4 |

| Image | Variant | Seed | $N$ | File bpp | PSNR | Fit (s) |
|---|---|---|---|---|---|---|
| painting-1_2k | Full4 | 124 | 9830 | 0.28778 | 31.82 | 437.2 |
| painting-1_2k | Full4 | 125 | 9830 | 0.28777 | 32.05 | 390.0 |
| painting-1_2k | Full4 | 124 | 10312 | 0.30181 | 32.18 | 388.0 |
| painting-1_2k | Full4 | 125 | 10312 | 0.30180 | 32.21 | 383.8 |
| painting-1_2k | Shared4 | 123 | 11233 | 0.30154 | 31.97 | 423.2 |
| painting-1_2k | Shared4 | 123 | 9830 | 0.26407 | 31.42 | 379.5 |
| painting-1_2k | Shared4 | 124 | 9830 | 0.26408 | 31.35 | 387.0 |
| painting-1_2k | Shared4 | 125 | 9830 | 0.26408 | 31.49 | 352.1 |
| painting-1_2k | Shared4 | 124 | 11233 | 0.30155 | 32.10 | 437.0 |
| painting-1_2k | Shared4 | 125 | 11233 | 0.30155 | 32.10 | 420.1 |
| painting-1_2k | RGB16+$\pi$ | 123 | 9118 | 0.30160 | 31.69 | 276.2 |
| painting-1_2k | RGB16+$\pi$ | 123 | 9830 | 0.32503 | 31.98 | 359.0 |
| painting-1_2k | RGB16+$\pi$ | 124 | 9830 | 0.32502 | 32.13 | 390.8 |
| painting-1_2k | RGB16+$\pi$ | 125 | 9830 | 0.32502 | 32.05 | 341.9 |
| painting-1_2k | RGB16+$\pi$ | 124 | 9118 | 0.30159 | 31.71 | 395.5 |
| painting-1_2k | RGB16+$\pi$ | 125 | 9118 | 0.30160 | 31.66 | 365.5 |
| painting-1_2k | RGB10+$\pi$ | 123 | 10485 | 0.30160 | 32.27 | 236.8 |
| painting-1_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 30.35 | 386.2 |
| painting-1_2k | RGB-U16 | 123 | 9830 | 0.30132 | 31.56 | 362.1 |
| painting-1_2k | RGB-U16 | 124 | 9830 | 0.30132 | 31.58 | 325.6 |
| painting-1_2k | RGB-U16 | 125 | 9830 | 0.30132 | 31.56 | 361.1 |
| painting-1_2k | Simplex4 | 123 | 10484 | 0.30182 | 32.22 | 291.8 |
| painting-1_2k | VQ256 | 123 | 12714 | 0.30178 | 28.75 | 426.4 |
| painting-2_2k | Full4 | 123 | 10312 | 0.30181 | 33.13 | 430.1 |
| painting-2_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 32.23 | 238.5 |
| painting-2_2k | RGB-U16 | 123 | 9830 | 0.30132 | 32.80 | 241.9 |
| painting-3_2k | Full4 | 123 | 10312 | 0.30180 | 35.89 | 444.5 |
| painting-3_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 33.08 | 350.6 |
| painting-3_2k | RGB-U16 | 123 | 9830 | 0.30132 | 35.12 | 357.4 |
| painting-4_2k | Full4 | 123 | 10312 | 0.30181 | 34.05 | 400.7 |
| painting-4_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 31.50 | 439.2 |
| painting-4_2k | RGB-U16 | 123 | 9830 | 0.30131 | 33.31 | 398.7 |
| painting-5_2k | Full4 | 123 | 10312 | 0.30181 | 31.79 | 404.9 |
| painting-5_2k | RGB-FP16 | 123 | 9826 | 0.30081 | 31.22 | 426.5 |
| painting-5_2k | RGB-U16 | 123 | 9830 | 0.30132 | 32.29 | 423.9 |
| painting-6_2k | Full4 | 123 | 10312 | 0.30181 | 32.19 | 419.2 |
| painting-6_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 29.93 | 373.1 |
| painting-6_2k | RGB-U16 | 123 | 9830 | 0.30132 | 31.31 | 362.5 |
| painting-7_2k | Full4 | 123 | 10312 | 0.30180 | 27.82 | 386.3 |
| painting-7_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 25.86 | 345.8 |
| painting-7_2k | RGB-U16 | 123 | 9830 | 0.30132 | 26.79 | 348.5 |
| painting-8_2k | Full4 | 123 | 10312 | 0.30181 | 29.47 | 365.9 |
| painting-8_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 27.63 | 340.9 |
| painting-8_2k | RGB-U16 | 123 | 9830 | 0.30132 | 28.45 | 349.2 |
| painting-9_2k | Full4 | 123 | 10312 | 0.30181 | 27.88 | 237.2 |
| painting-9_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 25.70 | 205.7 |
| painting-9_2k | RGB-U16 | 123 | 9830 | 0.30132 | 26.95 | 214.0 |
| photo-1_2k | Full8-G16 | 123 | 10310 | 0.30179 | 34.55 | 400.8 |
| photo-1_2k | Full8-G32 | 123 | 6227 | 0.30152 | 33.66 | 401.6 |
| photo-1_2k | Full4 | 123 | 10312 | 0.30181 | 35.83 | 290.5 |
| photo-1_2k | Full4 | 123 | 9830 | 0.28778 | 35.76 | 426.6 |
| photo-1_2k | Full4 | 124 | 9830 | 0.28777 | 35.74 | 390.2 |
| photo-1_2k | Full4 | 125 | 9830 | 0.28778 | 35.82 | 371.1 |
| photo-1_2k | Full4 | 124 | 10312 | 0.30181 | 35.80 | 369.0 |
| photo-1_2k | Full4 | 125 | 10312 | 0.30181 | 35.77 | 412.5 |
| photo-1_2k | Shared4 | 123 | 11233 | 0.30155 | 35.84 | 442.8 |
| photo-1_2k | Shared4 | 123 | 9830 | 0.26407 | 35.62 | 391.8 |

| Image | Variant | Seed | $N$ | File bpp | PSNR | Fit (s) |
|---|---|---|---|---|---|---|
| photo-1_2k | Shared4 | 124 | 9830 | 0.26407 | 35.58 | 378.9 |
| photo-1_2k | Shared4 | 125 | 9830 | 0.26407 | 35.64 | 449.2 |
| photo-1_2k | Shared4 | 124 | 11233 | 0.30154 | 35.74 | 429.6 |
| photo-1_2k | Shared4 | 125 | 11233 | 0.30155 | 35.78 | 453.6 |
| photo-1_2k | RGB16+$\pi$ | 123 | 9118 | 0.30160 | 35.84 | 275.2 |
| photo-1_2k | RGB16+$\pi$ | 123 | 9830 | 0.32502 | 35.94 | 427.7 |
| photo-1_2k | RGB16+$\pi$ | 124 | 9830 | 0.32503 | 35.92 | 353.6 |
| photo-1_2k | RGB16+$\pi$ | 125 | 9830 | 0.32503 | 35.93 | 327.3 |
| photo-1_2k | RGB16+$\pi$ | 124 | 9118 | 0.30159 | 35.85 | 367.1 |
| photo-1_2k | RGB16+$\pi$ | 125 | 9118 | 0.30160 | 35.84 | 386.9 |
| photo-1_2k | RGB10+$\pi$ | 123 | 10485 | 0.30160 | 35.97 | 254.6 |
| photo-1_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 35.03 | 363.7 |
| photo-1_2k | RGB-U16 | 123 | 9830 | 0.30132 | 35.81 | 421.6 |
| photo-1_2k | RGB-U16 | 124 | 9830 | 0.30132 | 35.81 | 384.2 |
| photo-1_2k | RGB-U16 | 125 | 9830 | 0.30132 | 35.81 | 410.0 |
| photo-1_2k | Simplex4 | 123 | 10484 | 0.30182 | 35.95 | 297.2 |
| photo-1_2k | VQ256 | 123 | 12714 | 0.30178 | 32.48 | 445.6 |
| photo-2_2k | Full4 | 123 | 10312 | 0.30181 | 37.17 | 459.5 |
| photo-2_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 34.71 | 254.1 |
| photo-2_2k | RGB-U16 | 123 | 9830 | 0.30131 | 36.76 | 256.6 |
| photo-3_2k | Full4 | 123 | 10312 | 0.30181 | 39.77 | 433.4 |
| photo-3_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 37.73 | 404.8 |
| photo-3_2k | RGB-U16 | 123 | 9830 | 0.30132 | 39.46 | 388.0 |
| photo-4_2k | Full4 | 123 | 10312 | 0.30181 | 37.05 | 398.6 |
| photo-4_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 35.32 | 349.4 |
| photo-4_2k | RGB-U16 | 123 | 9830 | 0.30132 | 36.64 | 413.7 |
| photo-5_2k | Full4 | 123 | 10312 | 0.30181 | 29.78 | 453.5 |
| photo-5_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 28.62 | 231.4 |
| photo-5_2k | RGB-U16 | 123 | 9830 | 0.30132 | 29.32 | 351.9 |
| photo-6_2k | Full4 | 123 | 10312 | 0.30181 | 45.20 | 388.2 |
| photo-6_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 44.54 | 370.0 |
| photo-6_2k | RGB-U16 | 123 | 9830 | 0.30132 | 45.50 | 396.4 |
| photo-7_2k | Full4 | 123 | 10312 | 0.30182 | 39.99 | 236.5 |
| photo-7_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 37.19 | 357.6 |
| photo-7_2k | RGB-U16 | 123 | 9830 | 0.30133 | 39.45 | 388.0 |
| photo-8_2k | Full4 | 123 | 10312 | 0.30181 | 28.63 | 388.9 |
| photo-8_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 27.34 | 232.0 |
| photo-8_2k | RGB-U16 | 123 | 9830 | 0.30132 | 28.12 | 365.2 |
| photo-9_2k | Full4 | 123 | 10312 | 0.30181 | 30.33 | 297.2 |
| photo-9_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 28.63 | 354.8 |
| photo-9_2k | RGB-U16 | 123 | 9830 | 0.30131 | 29.76 | 321.3 |
| vector-1_2k | Full8-G16 | 123 | 10310 | 0.30180 | 27.81 | 411.6 |
| vector-1_2k | Full8-G32 | 123 | 6227 | 0.30152 | 25.70 | 426.6 |
| vector-1_2k | Full4 | 123 | 10312 | 0.30181 | 28.94 | 317.1 |
| vector-1_2k | Full4 | 124 | 10312 | 0.30180 | 29.11 | 397.7 |
| vector-1_2k | Full4 | 125 | 10312 | 0.30181 | 29.40 | 376.8 |
| vector-1_2k | Shared4 | 123 | 11233 | 0.30154 | 28.90 | 416.6 |
| vector-1_2k | Shared4 | 123 | 9830 | 0.26407 | 28.01 | 304.5 |
| vector-1_2k | Shared4 | 124 | 11233 | 0.30155 | 28.73 | 381.7 |
| vector-1_2k | Shared4 | 125 | 11233 | 0.30155 | 28.92 | 390.5 |
| vector-1_2k | RGB16+$\pi$ | 123 | 9118 | 0.30160 | 28.87 | 262.3 |
| vector-1_2k | RGB16+$\pi$ | 124 | 9118 | 0.30160 | 28.89 | 420.4 |
| vector-1_2k | RGB16+$\pi$ | 125 | 9118 | 0.30160 | 28.98 | 295.3 |
| vector-1_2k | RGB10+$\pi$ | 123 | 10485 | 0.30158 | 29.44 | 235.3 |
| vector-1_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 27.25 | 346.3 |
| vector-1_2k | RGB-U16 | 123 | 9830 | 0.30132 | 28.42 | 400.2 |
| vector-1_2k | RGB-U16 | 124 | 9830 | 0.30132 | 28.50 | 427.6 |

| Image | Variant | Seed | $N$ | File bpp | PSNR | Fit (s) |
|---|---|---|---|---|---|---|
| vector-1_2k | RGB-U16 | 125 | 9830 | 0.30132 | 28.52 | 344.3 |
| vector-1_2k | Simplex4 | 123 | 10484 | 0.30182 | 29.48 | 285.4 |
| vector-1_2k | VQ256 | 123 | 12714 | 0.30178 | 23.09 | 359.2 |
| vector-2_2k | Full4 | 123 | 10312 | 0.30180 | 32.08 | 449.5 |
| vector-2_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 29.67 | 342.1 |
| vector-2_2k | RGB-U16 | 123 | 9830 | 0.30132 | 31.02 | 250.2 |
| vector-3_2k | Full4 | 123 | 10312 | 0.30181 | 30.53 | 398.5 |
| vector-3_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 28.57 | 348.6 |
| vector-3_2k | RGB-U16 | 123 | 9830 | 0.30132 | 29.60 | 405.6 |
| vector-4_2k | Full4 | 123 | 10312 | 0.30181 | 45.52 | 427.2 |
| vector-4_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 38.88 | 383.6 |
| vector-4_2k | RGB-U16 | 123 | 9830 | 0.30132 | 44.29 | 410.0 |
| vector-5_2k | Full4 | 123 | 10312 | 0.30181 | 32.75 | 367.2 |
| vector-5_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 30.61 | 352.4 |
| vector-5_2k | RGB-U16 | 123 | 9830 | 0.30132 | 32.54 | 235.1 |
| vector-6_2k | Full4 | 123 | 10312 | 0.30181 | 30.38 | 419.6 |
| vector-6_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 28.01 | 342.3 |
| vector-6_2k | RGB-U16 | 123 | 9830 | 0.30132 | 29.75 | 342.3 |
| vector-7_2k | Full4 | 123 | 10312 | 0.30181 | 33.71 | 391.0 |
| vector-7_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 30.90 | 204.0 |
| vector-7_2k | RGB-U16 | 123 | 9830 | 0.30131 | 32.37 | 211.3 |
| vector-8_2k | Full4 | 123 | 10312 | 0.30181 | 31.03 | 391.8 |
| vector-8_2k | RGB-FP16 | 123 | 9830 | 0.30093 | 27.98 | 342.7 |
| vector-8_2k | RGB-U16 | 123 | 9830 | 0.30132 | 29.61 | 348.5 |
| vector-9_2k | Full4 | 123 | 10312 | 0.30181 | 32.36 | 368.1 |
| vector-9_2k | RGB-FP16 | 123 | 9829 | 0.30090 | 31.02 | — |
| vector-9_2k | RGB-U16 | 123 | 9830 | 0.30132 | 31.93 | 354.1 |

## D EVALUATION BOUNDARIES

A Gaussian footprint is used to assign primitives to tiles. Increasing an amplitude does not rebuild that support from infinite Gaussian tails. The same finite candidate convention is used for all methods. Regions expanded to global tile boundaries preserve this convention; resetting tile origins at the requested crop would change the decoder. A future persistent tile index must include its storage, construction cost, and transfer behavior in any random-access claim.

The current byte stream is a reproducibility-oriented fixed-width format. It does not contain a spatial entropy model. Applying zlib to the complete stream preserves reconstructed fields but does not by itself permit direct access to an arbitrary compressed field. Consequently, zlib file size and direct fixed-width regional-access timing are not combined into a single operating point.