OBSERVATION-AWARE QIM FOR GAUSSIAN SPLATS

Anonymous authors
Paper under double-blind review

ABSTRACT

When one 3D Gaussian-splatting asset is released to several recipients, a leaked copy should reveal which release it came from. Releases share geometry exactly, so the identifier can live only in the spherical-harmonic (SH) colour coefficients, where it must stay invisible yet survive re-export. Existing marks need trained decoders, and a direct write budgets its change in coefficient distance although camera sensitivity spans a median 7.2 orders of magnitude, leaving margins too thin for 8-bit export. We instead spend the budget in the renderer's own error metric, placing the mark where cameras barely look. Observation-aware quantization index modulation (QIM) solves each keyed quantization write in closed form in a metric built from compositing weights, and a compositing inequality bounds the image error of the complete write on the declared cameras, certifying each write before release. On fourteen scenes at matched training error, exact recovery after 8-bit export rises from 10/70 to 70/70 over Euclidean and diagonal controls, at higher held-out error; at a matched 48-bit payload, every identifier survives half-pruning at a marking cost of .0005 dB, against 3.99 dB for GaussianMarker. Float16 coordinates, SH truncation, and a keyless recipient holding the declared cameras remain failure modes. These results show that solving the write in the renderer's metric lets a training-free mark survive export, and they give native 3D assets a certified release identifier with measured limits.

1 INTRODUCTION

When one 3D Gaussian-splatting (3DGS) asset is released to several recipients, a leaked copy should reveal which release it came from. 3DGS stores translucent ellipsoids with view-dependent spherical-harmonic (SH) colours, rendered by projection and compositing (Kerbl et al., 2023). Releases of one asset share identical geometry, so a keyed release identifier can live only in the SH coefficients, where it must stay invisible and survive re-export, cropping, and an unknown similarity transform. Recovering it links a copy to a committed release; it proves neither ownership nor priority nor collusion resistance.

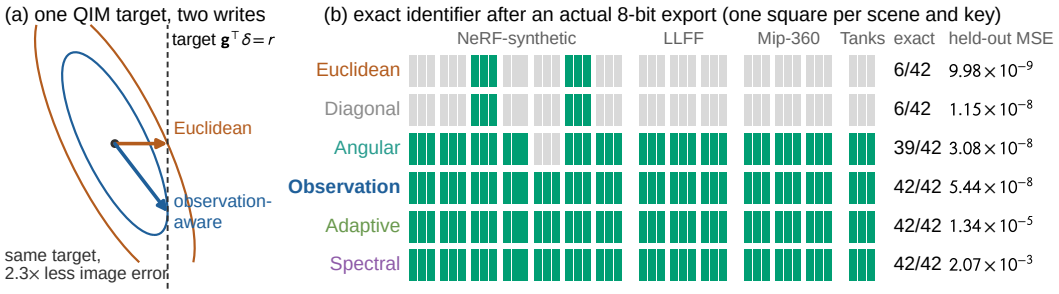

Figure 1: (a) Equal-image-error ellipses under an illustrative metric $\mathbf{H}$, elongated where the cameras barely observe. Both writes reach the same QIM target; the observation-aware write moves along $\mathbf{H}^{-1}\mathbf{g}$ at 2.3 times less image error than the Euclidean write along $\mathbf{g}$. (b) Exact identifier after an actual 8-bit export at matched training-camera mean squared error (MSE), one square per scene and key, for the six writes of Section 3.4; dark squares mark recovery, and rows are ordered by clean held-out (test-view) MSE (Table 2). Only the four camera-aware writes survive the export.

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.

The obstacle is the budget. A direct write limits a coefficient change $\boldsymbol{\delta}$ by $\|\boldsymbol{\delta}\|_2$, as our Euclidean control does, whereas the image error a viewer sees is at most $\boldsymbol{\delta}^\top \mathbf{M}_i \boldsymbol{\delta}$, with $\mathbf{M}_i$ accumulating contribution-weighted SH outer products over the cameras the owner declares (the training cameras here; Eq. 2). The ratio $\boldsymbol{\delta}^\top \mathbf{M}_i \boldsymbol{\delta} / \boldsymbol{\delta}^\top \boldsymbol{\delta}$ is a Rayleigh quotient over the spectrum of $\mathbf{M}_i$, which spans a median 7.2 orders of magnitude across our scenes: full-sphere rigs resolve all sixteen SH directions and spread five orders, while forward-facing and unbounded captures resolve a median of six and spread nearly sixteen, the machine-precision floor (Table 7). A Euclidean budget therefore overspends where cameras resolve and underspends elsewhere, and its quantization margins do not survive an 8-bit export, which stores each SH coefficient at one of 256 levels (Fig. 1).

Quantization index modulation (QIM) writes a bit by moving a scalar projection of the coefficients to the nearest point of one of two interleaved grids (Chen and Wornell, 2001). We solve that move in the renderer's metric: for a displacement r along a unit measurement $\mathbf{g}$, minimize $\boldsymbol{\delta}^\top \mathbf{H} \boldsymbol{\delta}$ subject to $\mathbf{g}^\top \boldsymbol{\delta} = r$ (Eq. 4), with $\mathbf{H}$ the ridge-regularized metric. Theorem 1 gives the minimizer in closed form at cost r^2/κ with $\kappa = \mathbf{g}^\top \mathbf{H}^{-1} \mathbf{g}$, so the grid spacing a per-slot budget D buys grows as $\sqrt{D\kappa}$ (Eq. 7), and the spectral form of κ (Section 3.4) places the mark on weakly observed eigendirections of $\mathbf{M}_i$. Our contributions are:

- **A complete-write image bound** (Theorem 2): nonnegative compositing masses and nonexpansive clamps turn a coefficient budget into a bound on the image error of the simultaneous write on the declared cameras; an audit at tolerance 10^{-10} finds 0 excesses (Fig. 4).
- **The metric, not the pipeline, buys recovery:** Euclidean and diagonal controls that share our pipeline recover 6/42 identifiers after an 8-bit export, against 42/42 for the metric (Table 2).
- **A reference-assisted protocol for native assets:** alignment frozen before any keyed read, committed strengths, separate consistency and identifier decisions, and an erasure-aware null at $\alpha = 10^{-6}$ keep every decision from being tuned to the key (Section 4).
- **Measured boundaries:** GuardSplat and GaussianMarker fail on edits complementary to ours (Table 3); the certificate covers only the declared cameras, and a keyless recipient holding them can strip the mark (Table 1).

2 Related work

Quantization watermarking. Dithered spread-transform QIM removes host interference (Chen and Wornell, 2001; 2000), and writing under a linear constraint is established (Altun et al., 2006); our projection inherits that theory. New is the metric the write is solved against and the bound it yields: Eq. 6 converts a coefficient-space cost into an image-space bound on a declared camera set. Angular concentration (Simons et al., 2006) explains why SH sensitivity is anisotropic under finite coverage.

Watermarking Gaussian splats. GuardSplat and NGS-Marker (Chen et al., 2025; Qin et al., 2026) optimize $O(Nd)$ scene parameters against fixed-size networks; 3D-GSW (Jang et al., 2025) and GaussianMarker (Huang et al., 2024) also change the primitive set; GaussianMarker and Robust3DGSW (Wang et al., 2026) support parameter and image decoders; and CompMarkGS and W2M (In et al., 2026; Li et al., 2026) read rendered images. An image reader survives any appearance-preserving edit and fails when appearance changes; a parameter reader has the complementary failure set and needs alignment. GuardSplat is the nearest neighbour in carrier: it also writes only SH coefficients, but its learned decoder reads renders. NGS-Marker, which also reads native fragments, is nearest in reader contract but could not be run (Appendix G). GaussianMarker pays for its mark by densifying roughly $1.8\times$. Splats-in-splats (Guo et al., 2026) hides content rather than a label, and GSPure (Huang et al., 2026) removes watermarks. None of these schemes certifies its distortion, although Theorem 2 would bound any that changes only SH coefficients. Table 14 compares the contracts and Table 18 the asymptotic costs; QIM trains no network and stores only $48N$ reference bytes, plus $128N$ for the spectral carrier.

Importance-weighted placement. W2M (Li et al., 2026) guides 3DGS placement by learned importance: a policy selects *which primitives* to write and is validated by the recovery it produces.

Our metric selects *which direction* at each primitive and is validated by a bound computed before any recovery is measured, so the two compose. TEASER++ (Yang et al., 2020) contextualizes our random sample consensus (RANSAC) alignment; passport commitments (Fan et al., 2019) address ambiguity, not authorship.

3 OBSERVATION-AWARE QUANTIZATION

3.1 RENDERER AND OBSERVATION DOMAIN

Bold denotes vectors and matrices, and Table 5 collects all symbols. Primitive i has a fixed centre $\mathbf{x}_i \in \mathbb{R}^3$, covariance, opacity, and SH coefficients $\mathbf{s}_{ic} \in \mathbb{R}^{16}$ per channel c; camera v has centre $\mathbf{o}_v$, and $\mathbf{d}_{iv} = (\mathbf{x}_i - \mathbf{o}_v)/\|\mathbf{x}_i - \mathbf{o}_v\|$ follows the audited 3DGS convention (Kerbl et al., 2023). $\mathbf{y}_{iv} \in \mathbb{R}^{16}$ evaluates real SH through degree three, and $f_{icv} = \max(0, \mathbf{y}_{iv}^\top \mathbf{s}_{ic} + 1/2)$. At pixel p the compositing weight $a_{ivp} = \alpha_{ivp} T_{ivp} \geq 0$ satisfies $\sum_i a_{ivp} \leq 1$ and absorbs culling, alpha thresholds, and early termination. With fixed background β_c and residual transmission T_{vp}^{end},

$$I_{vpc} = \mathrm{clip}_{[0,1]} \left(\sum_i a_{ivp} f_{icv} + T_{vp}^{\mathrm{end}} \beta_c \right). \tag{1}$$

The certificate holds geometry, opacity, background, camera parameters, and compositing order fixed, excludes nontrivial exposure transforms, and covers only the declared cameras $\mathcal{V}$ on which the owner builds the metric. With contribution mass $w_{iv} = \sum_p a_{ivp}$,

$$\mathbf{M}_i = \sum_{v \in \mathcal{V}} w_{iv} \mathbf{y}_{iv} \mathbf{y}_{iv}^\top, \qquad \mathbf{H}_i = \mathbf{M}_i + \rho_i \mathbf{I}_{16}, \quad \rho_i > 0. \tag{2}$$

Equation 2 is the only place the renderer enters the write, and a per-primitive sensitivity cannot supply it. Primitives overlap, so writing all of them at once produces cross terms of no fixed sign, and the renderer clamps twice, so a host derivative underestimates a finite write that crosses a clamp; a block Hessian (second-derivative matrix) therefore certifies nothing. Two facts restore a separable bound: masses are nonnegative and sum to at most one along a ray, which turns the cross terms into a weighted Cauchy–Schwarz bound, and clamps are nonexpansive, so replacing them by identity only increases the error. The result is Eq. 6 (Theorem 2), whose per-primitive factor is exactly $\mathbf{M}_i$. The ridge, $\rho_i = 10^{-3} \operatorname{tr}(\mathbf{M}_i)/16 + 10^{-6}$ in our runs (Table 6), penalizes unobserved changes and makes the solve unique; Appendix A describes how masses are measured.

3.2 THE HOST-PRESERVING WRITE

A 256-bit key K and a persistent reference identity generate, for each slot (i, c), a unit measurement $\mathbf{g}_{ic} \in \mathbb{R}^{16}$ (sixteen signs of magnitude $1/4$), a dither ξ_{ic}, and a committed strength $h_{ic} > 0$. For a padded bit b_{ic}, an identifier bit masked by a keyed pad (Section 4), let q_{ic} be the point of

$$\mathcal{Q}_b = \{2h_{ic} z + b h_{ic} + \xi_{ic} : z \in \mathbb{Z}\} \tag{3}$$

closest to $t_{ic} = \mathbf{g}_{ic}^\top \mathbf{s}_{ic}$. With $r_{ic} = q_{ic} - t_{ic}$, the update solves

$$\min_{\boldsymbol{\delta} \in \mathbb{R}^{16}} \boldsymbol{\delta}^\top \mathbf{H}_i \boldsymbol{\delta} \quad \text{subject to} \quad \mathbf{g}_{ic}^\top \boldsymbol{\delta} = r_{ic}. \tag{4}$$

Theorem 1 (Standard constrained projection)**.** *For $\kappa_{ic} = \mathbf{g}_{ic}^\top \mathbf{H}_i^{-1} \mathbf{g}_{ic} > 0$, the unique solution and its cost are*

$$\boldsymbol{\delta}_{ic}^* = \frac{r_{ic}}{\kappa_{ic}} \mathbf{H}_i^{-1} \mathbf{g}_{ic}, \qquad D_{ic}^* = \frac{r_{ic}^2}{\kappa_{ic}} \leq \frac{h_{ic}^2}{\kappa_{ic}}. \tag{5}$$

Stationarity gives $\mathbf{H}_i \boldsymbol{\delta} = \mu \mathbf{g}_{ic}$ for a multiplier μ, the constraint fixes $\mu = r_{ic}/\kappa_{ic}$, positive definiteness gives uniqueness, and $|r_{ic}| \leq h_{ic}$ the bound.

The naive combination fails, and we build it to show that. Keyed code, Hamming redundancy, pads, primitive cap, and geometry-first registration are standard and shared by all our writes, and none of them controls the decoding margin: the spacing a per-slot budget D buys scales as $\sqrt{D\kappa}$, and for the shared measurement κ is set by the metric. Section 5.3 builds that combination and measures it; Algorithm 1 (Section 4) states the complete procedure as implemented.

3.3 WHAT IS ACTUALLY CERTIFIED

Theorem 2 (Finite-camera image distortion)**.** *Under the renderer assumptions above, for any complete coefficient update,*

$$\sum_{v\in\mathcal{V},p,c} |\Delta I_{vpc}|^2 \le \sum_{i,c} \boldsymbol{\delta}_{ic}^\top \mathbf{M}_i \boldsymbol{\delta}_{ic} \le \sum_{i,c} \boldsymbol{\delta}_{ic}^\top \mathbf{H}_i \boldsymbol{\delta}_{ic}. \tag{6}$$

For exact updates from Eq. 5, the last quantity is $\sum_{i,c} D_{ic}^*$.

Appendix A gives the proof. For RGB mean squared error (MSE), divide each term by $P = 3\sum_{v\in\mathcal{V}} Q_v$, with Q_v the pixel count of camera v; peak signal-to-noise ratio (PSNR) is $-10\log_{10}(\mathrm{MSE})$ on unit range. Ground-truth (GT) PSNR and the error on held-out cameras (test views) and off-path (rotated-camera) views, which have no ground truth, are measured, not bounded.

3.4 OBSERVABILITY AND RECOVERABILITY

For inspected coefficients $\mathbf{s}_{ic} + \boldsymbol{\delta}_{ic}^* + \mathbf{e}_{ic}$ with $\|\mathbf{e}_{ic}\|_2 \le \epsilon_{ic}$, the scalar error is at most ϵ_{ic} because $\|\mathbf{g}_{ic}\|_2 = 1$, and opposite-bit points are h_{ic} apart; for a worst-case per-slot budget D_{ic}, sufficient simultaneous conditions are therefore

$$2\epsilon_{ic} < h_{ic} \le \sqrt{D_{ic}\kappa_{ic}}. \tag{7}$$

The interval is sufficient, not necessary, and assumes correct registration (Appendix A). With eigenpairs $\mathbf{M}_i\mathbf{u}_{ij} = \lambda_{ij}\mathbf{u}_{ij}$, $\kappa_{ic} = \sum_j |\mathbf{u}_{ij}^\top \mathbf{g}_{ic}|^2/(\lambda_{ij} + \rho_i)$, so weakly observed directions help in proportion to their overlap with the measurement.

Six writes. All six writes share the keyed code, pads, primitive cap, and registration and differ in the direction and strength of each update. *Observation* QIM, our primary write, solves Eq. 4 at the uniform strength $h = \sqrt{B/\sum_{ic} 1/\kappa_{ic}}$ for a total budget B, quoted as training-camera MSE; *adaptive* QIM sets $h_{ic} = \sqrt{B\kappa_{ic}/(3N)}$ over N written primitives; *angular* QIM sets every $w_{iv} = 1$ in Eq. 2, rescaled to the trace of the weighted $\mathbf{M}_i$, which keeps camera geometry but not contribution weights; and *spectral* QIM measures along the weakest eigenvector of each $\mathbf{M}_i$ and stores that carrier, $128N$ bytes. The controls are *Euclidean* QIM, which writes along $\mathbf{g}_{ic}$ with sufficient common strength $\sqrt{B/\sum_{ic} \mathbf{g}_{ic}^\top \mathbf{H}_i \mathbf{g}_{ic}}$, and *diagonal* QIM, which writes along $\mathrm{diag}(\mathbf{H}_i)^{-1}\mathbf{g}_{ic}$. Equal budgets need not give equal image MSE, and a small ρ_i can permit changes that harm extrapolated views.

4 A REFERENCE-ASSISTED FRAGMENT PROTOCOL

Registration. Figure 2 shows the order of operations. The verifier registers the inspected copy to the committed reference centres by RANSAC (Fischler and Bolles, 1981) on eight-neighbour distances; repeated structure and reparameterization defeat this prototype. For $\mathbf{x}_{\mathrm{ref}} = \gamma\mathbf{R}^\top\mathbf{x}_{\mathrm{inspected}} + \mathbf{t}$ with $\gamma > 0$ and rotation $\mathbf{R}$, SH transport satisfies $f_{\mathrm{ref}}(\mathbf{d}) = f_{\mathrm{inspected}}(\mathbf{R}\mathbf{d})$, and translation and positive scale preserve directional colour. Once alignment freezes, an independent keyed domain ranks matches and keeps at most 8192 primitives (details in Appendix B).

Keyed coding and contract. HMAC-SHA256 is the keyed pseudorandom function (PRF) (Krawczyk et al., 1997), with a separate domain for each keyed quantity. Hamming $(7, 4)$ coding (Hamming, 1950) expands 256 identifier bits to 448; keyed placement assigns them to slots, and pads mask them by exclusive-or. A codeword decodes only when twice its disagreements plus its erasures stay below three, so each block corrects one error or two erasures. Spatial redundancy is probabilistic, not a per-crop guarantee. Only owner and verifier hold the key, fresh per claim, and claim, reference, strengths, registration, cap, threshold, family size, asset, and implementation are committed before inspection (Appendix A). Recipients get neither key nor query oracle; priority needs external timestamping.

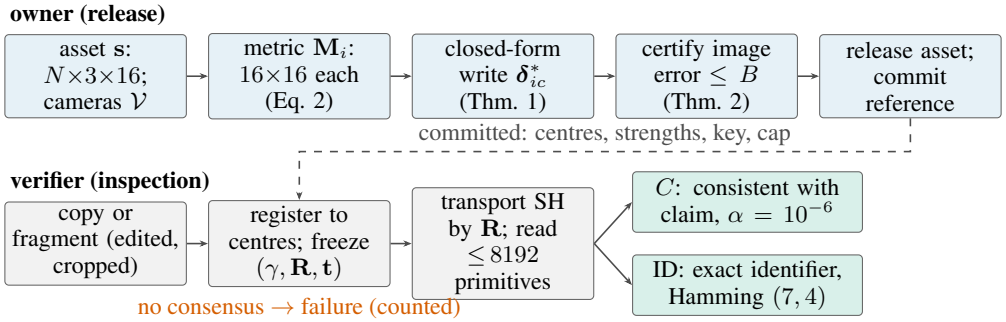

Figure 2: Release and inspection. The owner writes and certifies the identifier and commits a reference (dashed); the verifier registers a copy to that reference, freezes the alignment, and only then decides consistency C and exact recovery ID, attributing only if both hold, while a failed registration counts as a failure. Because alignment, strengths, and threshold are committed before any keyed read, no decision can be tuned to the key.

Algorithm 1 Observation-aware QIM as implemented. $\mathcal{V}$ declared cameras, K key, B certified budget, π_{ic} keyed placement of code bits; C and ID are the decisions below.

procedure EMBED($\{\mathbf{x}_i, \mathbf{s}_{ic}\}, \mathcal{V}, K, B,$ id)
  $\mathbf{M}_i \leftarrow \sum_{v\in\mathcal{V}} w_{iv}\mathbf{y}_{iv}\mathbf{y}_{iv}^\top$; $\mathbf{H}_i \leftarrow \mathbf{M}_i + \rho_i \mathbf{I}_{16}$ ▷ Eq. 2
  $(\mathbf{g}_{ic}, \xi_{ic}, \pi_{ic}, \mathrm{pad}_{ic}) \leftarrow \mathrm{HMAC}_K$(reference identity, channel)
  $\kappa_{ic} \leftarrow \mathbf{g}_{ic}^\top \mathbf{H}_i^{-1} \mathbf{g}_{ic}$ by Cholesky solve; $h \leftarrow \sqrt{B/\sum_{ic} 1/\kappa_{ic}}$
  $b_{ic} \leftarrow c_{\pi_{ic}} \oplus \mathrm{pad}_{ic}$ for the codeword $\mathbf{c} = \mathrm{HAMMING}_{7,4}$(id)
  for all written slots (i, c) **do** ▷ Thm. 1
    $r_{ic} \leftarrow q_{ic} - \mathbf{g}_{ic}^\top \mathbf{s}_{ic}$ for q_{ic} the nearest point of $\mathcal{Q}_{b_{ic}}$
    $\mathbf{s}_{ic} \leftarrow \mathbf{s}_{ic} + (r_{ic}/\kappa_{ic})\, \mathbf{H}_i^{-1} \mathbf{g}_{ic}$
  certify $\sum \boldsymbol{\delta}^\top \mathbf{H} \boldsymbol{\delta} \le B$ on the serialized update ▷ Eq. 6
  commit hashes of centres, identities, strengths, key, policy, and implementation
procedure VERIFY(inspected asset, commitment, K)
  check key, reference, strengths, and implementation against the commitment
  $(\gamma, \mathbf{R}, \mathbf{t}) \leftarrow$ RANSAC on eight-neighbour distances ▷ geometry only
  if no consensus **then return** FAIL ▷ counted, not skipped
  keep ≤ 8192 matches by keyed rank; transport SH by $\mathbf{R}$
  read the nearer grid per slot; ties erase
  return C (binomial tail at α/J) and ID (bounded-distance Hamming decode)

Two decisions and their null. Consistency C rejects the unrelated-asset null; ID is exact agreement with the committed identifier. Attribution requires both, because a nearby wrong identifier under a reused key passes C and fails ID; two independently keyed identifiers sharing geometry each attribute 4/4, so geometry alone cannot distinguish labels. Under the null, asset and geometry selection are independent of the hidden pads, so non-erased raw agreements are fair coins and erasures count as failures: the agreement count is dominated by $\mathrm{Binomial}(n, 1/2)$, tested at $\alpha = 10^{-6}$ with tail $2^{-n}\sum_{j=k}^{n}\binom{n}{j}$. Unmatched primitives supply no evidence, and a family of J tests uses α/J. HMAC's PRF assumption, adaptive pad queries, and recipient unforgeability lie outside this null.

5 EXPERIMENTS

RQ1 Does the metric buy recovery against controls at matched training MSE and against published baselines at a matched payload (Sections 5.3 and 5.5)? **RQ2** How does it change what is written, and is that detectable (Section 5.2)? **RQ3** Is each part necessary (Section 5.3)? **RQ4** Where does the channel break, and does it transfer (Section 5.4)?

Table 1: Cost and survival at the common budget $B = 10^{-6}$ (fourteen scenes, three keys). Left: mean image error by camera set and its bound on training cameras (Eq. 6; per scene in Table 9). Right: exact identifiers of 42 after SH noise σ, rounding step .1, truncation to degree 2, and the keyless strip of the four least-observed directions and ridge refit (Section 5.4); all 42 recover unattacked, and Table 12 adds step .01. Bold: best; underline: second best, where it differs from the worst. Observation QIM keeps held-out error near the Euclidean control's and pays for its margin off-path.

	Image error (MSE)				Exact identifiers of 42 after					
Write	Training	Bound	Held-out	Off-path	noise .01	noise .05	step .1	SH deg. 2	strip 4 dirs	ridge refit
Euclidean	2.99e-8	3.34e-7	**2.99e-8**	**3.25e-8**	0	0	0	1	0	3
Observation	7.17e-9	1.32e-7	3.09e-8	2.61e-5	41	9	24	15	0	3
Adaptive	**2.50e-9**	**9.45e-8**	5.00e-6	1.15e-3	**42**	**37**	**39**	**37**	**17**	**6**

5.1 Setup

The native renderer evaluates eight NeRF synthetic objects (Mildenhall et al., 2020), three Local Light Field Fusion (LLFF) scenes (Mildenhall et al., 2019), three Mip-NeRF 360 captures (Barron et al., 2022), and Tanks and Temples truck (Knapitsch et al., 2017); Lego is development, and 14 scenes enter evaluation under the protocol of Table 6. Each write declares the training cameras, and trials use keys 0 to 4. Per scene, the mean write takes 0.03 to 0.94 s after 1.2 to 414.2 s of preprocessing and 10.5 to 243.2 s of key setup, and uncached verification takes 2.9 to 73.5 s (Table 17).

5.2 Certificates, extrapolation, and recovery

This section asks what the certificate covers and what the write costs outside it (RQ2). Observation QIM cuts mean training MSE by 4.18 times (95% interval 2.98–6.38) at 802 times greater off-path MSE, with held-out MSE nearly equal (Table 1). The certificate audit of Fig. 4 finds 0 excesses at tolerance 10^{-10}; the bound is loose by 11.1, 18.5, and 37.9 times for the Euclidean, observation, and adaptive writes (ratios of unrounded means), plausibly because discarding the clamps assumes nothing saturates and bounding $\sum_i a_{ivp}$ by one assumes a fully occluded ray, and the gap widens as a metric concentrates writes in weakly observed directions. In structural similarity (SSIM) and learned perceptual image patch similarity (LPIPS) against the unmarked render (Table 8), off-path error stays small: mean LPIPS .0085 at SSIM .9983, worst view .0406 at SSIM .9816, while GT-PSNR moves at most .0014 dB (Table 9). The restriction to declared cameras remains a real gap in what is *proved*, and we claim no perceptual threshold for what is *rendered*. A coefficient-statistics detector is at chance without a clean copy of the asset (AUC .490) and reaches .929 with one (Appendix H). Adaptive QIM matches or beats observation QIM on every attacked row of Table 12 but pays far more held-out and off-path error, so observation QIM is our primary write (Appendix D). **Key Takeaway: the certificate binds where it is declared, and the off-path cost it leaves uncovered is measurable and, for observation QIM, small at this budget.**

5.3 Competition at matched training MSE

This section asks whether the metric buys recovery once training error is matched, and which of its parts matter (RQ1, RQ3). Table 2 targets training MSE 10^{-8} on fourteen scenes; calibration key 99033 freezes an amplitude factor from training renders, so matching is approximate (Appendix E). Spectral QIM's recovery comes at greater clean held-out MSE than observation QIM's on every scene (median ratio 2.05; fern, trex, and room dominate the mean). Observation QIM's post-export held-out MSE averages 3.24×10^{-6}, never exceeds 1.80×10^{-5}, and costs at most .030 dB of GT-PSNR; spectral's mean is 6.97×10^{-2}.

Over five keys, paired per scene against the Euclidean control (Wilcoxon signed-rank primary; the paired t-test agrees except where noted), training MSE is equivalent within the 3% calibration tolerance (paired difference in $[-1.1, 1.2]\%$ of the target at 95%, $p = 3.5\times10^{-5}$, two one-sided tests), exact recovery after 8-bit export rises from 10/70 to 70/70 (paired gain .857 per scene, $p = 5.3\times10^{-4}$),

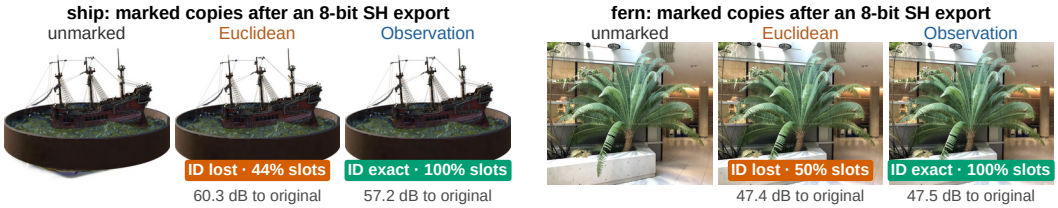

Figure 3: Marked copies of ship (a full-sphere scene) and fern (a forward-facing scene) after an actual 8-bit export at budget 10^{-6}, key 0, middle held-out camera. Badges give exact recovery and the fraction of raw slots read; grey text gives PSNR to the unmarked, unexported render. Over all held-out cameras both writes keep the same PSNR to ground truth (31.01 dB on ship, 24.28 dB on fern), yet only the observation-aware write reads its identifier.

Table 2: All six writes at matched training MSE (10^{-8}; fourteen scenes, three keys, all test cameras) under noise $\sigma = .01$, SH step .1, and an actual 8-bit export; the right block repeats the export over five keys, with the Wilcoxon p against the Euclidean control paired over scenes (fidelity in Table 11). Bold: best; underline: second best, where it differs from the worst. Training MSE is matched and not ranked. The four camera-aware writes keep through export the identifiers the controls lose and pay in held-out error.

	Image error (MSE)			Exact identifiers of 42 after			five keys	
Write	Training	Held-out	Off-path	noise	step	8-bit	exact of 70	Wilcoxon p
Euclidean	1.00e-8	**9.98e-9**	**1.10e-8**	0	0	6	10	—
Diagonal	9.97e-9	1.15e-8	1.78e-7	0	0	6	10	n/a
Angular	9.87e-9	3.08e-8	6.14e-5	36	9	39	65	9.1×10^{-4}
Observation	1.00e-8	5.44e-8	8.94e-5	39	36	**42**	**70**	5.3×10^{-4}
Adaptive	1.00e-8	1.34e-5	2.80e-3	**42**	**39**	**42**	**70**	5.3×10^{-4}
Spectral	9.89e-9	2.07e-3	1.06e-1	**42**	**39**	**42**	**70**	5.3×10^{-4}

and clean held-out MSE is significantly worse ($p = 1.2 \times 10^{-4}$). Diagonal and Euclidean recover identically, so that pair has no statistic, and spectral's held-out penalty fails the t-test.

Which part of the metric matters. Camera geometry carries most of the gain, and contribution weights add the rounding margin: angular QIM, which drops the weights, nearly matches observation QIM through export (39/42 against 42/42) but trails it under step rounding (9/42 against 36/42; Wilcoxon $p = 2.7 \times 10^{-3}$; Table 2) at a median 1.50 times lower held-out MSE. A four-scene sweep of five methods at four budgets agrees (Table 10, Fig. 5; renders in Fig. 3), and Appendix E adds the angular arms, a ridge sweep, and a frozen Lego-only budget rule that fails and is reported without retuning. One certified budget, $B = 10^{-6}$ (a training-camera PSNR of at least 60 dB), serves every scene without per-scene tuning and transfers to three unseen captures (Section 5.4). **Key Takeaway: the metric converts equal training MSE into a margin that survives 8-bit export, 42/42 against 6/42 for coefficient-space budgeting, paid for in held-out error.**

5.4 Fragments and attacks

This section asks where the channel breaks and whether it transfers (RQ4). Under unknown rotation, translation, and log-uniform scale in $[.25, 4]$, sphere, box, and halfspace fragments recover 108/108 identifiers at 1% and 10% retention and 85/108 at 0.1% (interval 36.1–100%), and centre jitter of 10^{-4} of scene extent cuts the 1% result to 3/108. Loss tracks coordinate error against primitive spacing: in a separate study of sixteen cases (twelve 1% crops and four full scenes; Appendix F), jitter of .009, .09, and .90 nearest-neighbour spacings leaves 16, 13, and 0 identifiers, each loss a failed or wrong registration that C rejects, and float16 coordinates keep six of the twelve crops. With the declared cameras, keyless edits remove the mark cheaply: a ridge refit of the coefficients in the cameras' metric under a norm penalty changes held-out GT-PSNR by $+.014$ dB and leaves 3/42 identifiers, and stripping four low-cost directions costs .19 dB and leaves 0/42 (Table 12; Table 6

Table 3: Both baselines against observation QIM on four shared checkpoints and one attack battery at matched 48-bit payloads. Bit accuracy and PSNR are four-scene means; for ours, C and ID count consistency and bit-exact recoveries over four scenes and three keys, and raw is the slot accuracy. The floor row reads each decoder without a watermark; †GuardSplat's .500 is chance, not measured. GaussianMarker's PSNR is on its own densified asset. The readers fail on different edits.

	GuardSplat		GaussianMarker		Ours			
Attack	bit acc.	PSNR	bit acc.	PSNR	C of 12	ID of 12	raw	PSNR
Unattacked	0.935	31.34	0.894	28.22	12	12	1.000	32.21
Keep 50%	0.620	27.02	0.721	27.71	12	12	1.000	27.21
Quantise 8 bit	0.935	31.34	0.894	28.22	12	12	1.000	32.21
Quantise 4 bit	0.868	31.11	0.886	28.13	12	12	0.718	31.94
SH degree ≤ 1	0.523	29.46	0.827	27.21	9	1	0.524	29.51
SH degree 0	0.510	28.11	0.795	26.42	0	0	0.504	28.11
Noise $\sigma = .05$	0.778	30.75	0.872	28.00	6	3	0.522	31.42
Strip 4 low-cost directions	0.923	31.16	0.890	28.16	0	0	0.505	31.99
Ridge coefficient fit	0.936	31.38	0.893	28.25	9	0	0.520	32.24
Informed strip	0.935	31.34	0.894	28.22	0	0	0.000	32.13
Unwatermarked floor	.500†	—	0.600	32.21	—	—	—	32.21

ablates the cap). Estimated cameras also suffice on fern, and a disclosed key defeats confirmation while blind targeting does not (Appendix F). On three unseen captures (kitchen, bicycle, stump) at the fixed budget, composing 8-bit export, cropping, unknown similarity, and float32 serialization over three keys and three 1% shapes recovers all 27 identifiers, and 18/27 under float16 coordinates (Table 13); this is transfer within Mip-NeRF 360, not to a new population. Null controls yield 0 consistencies in 300 trials, too few to validate a 10^{-6} tail. **Key Takeaway: the channel survives 1% crops and unknown similarity but fails as coordinate error grows from a tenth of the primitive spacing toward a full spacing, which confines it to pipelines that keep float32 centres.**

5.5 Learned baselines and reconstruction

This section compares against published learned watermarks (RQ1). Table 3 runs GuardSplat (Chen et al., 2025) and GaussianMarker (Huang et al., 2024) against observation QIM at matched 48-bit payloads, GaussianMarker's own and within GuardSplat's 72-bit tokenizer ceiling; our column uses three keys per scene through the verifier of Section 4, and its counts cluster by scene. The other five schemes of Table 14 are compared by contract only (Appendix G). Every mark survives quantization, and our exact identifiers at 4 bits rest on coding and redundancy, since our raw accuracy there is below both baselines'. Half-pruning costs both baselines bits while we keep every identifier; GuardSplat's drums checkpoint converged to 78.3% against 97.1–98.4% elsewhere, lowering its unattacked mean. SH truncation, strong noise, and the two keyless edits go against us. GaussianMarker stays above its measured .600 floor on those rows because its mark rides geometry; against that floor, its .619 at keep-10% means the message is gone (row omitted; Appendix G); GuardSplat keeps its bits under the strip and ridge fit but not under truncation, and the informed strip, which uses our key, leaves both baselines unmodified. Every exact identifier here also passed C. GaussianMarker densifies by $1.8\times$ at 3.99 dB against our .0005 dB, so the quality columns compare unequal assets; we omit its parameter-space decoder, which returns 1.000 on marked and unmarked Lego alike (Appendix G). Image-only reconstruction registers but loses confirmation, which alone does not show useful removal (Appendix F). **Key Takeaway: our parameter reader and the baselines' image readers fail in complementary places, so the choice is one of carrier, not quality.**

6 Limitations

The identifier rides the SH coefficients but is located by matching primitive centres to a float64 reference, so it breaks under lossy coordinates, not 8-bit colour. FP16 positions, octree quantization, and

Table 4: Where the channel loses the identifier (upper rows) and what every result assumes (lower rows). The operational boundary is coordinate precision, not colour precision.

Condition	Measured
Coordinate requantization	float16: 9/27 lost; jitter 10^{-4}: 108/108 $\rightarrow$ 3/108
Rig-informed removal	ridge fit: ID 3/42 at $+.014$ dB; strip: 0/42 at $-.19$ dB
Retention below 1%	108/108 at 1–10%, 85/108 at 0.1% (interval 36.1–100%)
Near-isotropic rigs	full-sphere scenes spread < 6 orders; little to exploit
Deployment needs	key, strengths, float64 reference, float32 centres; C forgeable on key reuse
Scale reached	≤ 3.64M primitives; 14 scenes, 4 datasets; fern has 3 test views
Assumed	audited renderer, fixed geometry and cameras; nothing outside $\mathcal{V}$

Draco-style compression are routine in production and can each defeat registration; the supported deployment therefore keeps float32 centres, under which all 27 composed recoveries hold. Keyless removal is the second boundary: the declared cameras suffice to strip the write from the directions they barely observe, which is where the metric places it (Table 1). Table 4 measures both and lists the remaining requirements.

7 Conclusion

Observation-aware QIM gives a compositing-derived finite-camera bound and a keyed, training-free write that beats the Euclidean and diagonal controls at matched training MSE and, under coarse rounding, the angular metric. Stored spectral QIM leads observation QIM only at a common certified budget and only on arms without keyless edits (Table 10), so avoiding its $128N$-byte carrier is the narrower advantage, and neither write establishes ownership.

AI use statement

The authors used LLMs solely for proofreading, improving language and clarity, and organizing the LaTeX source. The research methodology, experiments, results, figures, and scientific content are the authors' own contributions. The authors remain fully responsible for the accuracy, originality, and integrity of the manuscript.

Ethics statement

Dual use. Tracing a leak to its recipient equally traces lawful redistribution, a use we do not endorse; recovery shows only that a copy descends from a labelled release, not ownership, infringement, or wrongdoing. *Data and privacy.* All datasets are cited, obtained from their authors, and used under their non-commercial research terms, with a further non-commercial condition for Tanks and Temples; we redistribute no scene data. The captures involve no human subjects and show no people, faces, licence plates, or readable documents, so no personal data is processed; user-created assets would differ. A fresh key per asset–identifier claim limits cross-asset linkage without proving unlinkability, and a holder of the unmarked original detects the mark at AUC .929 (Appendix H), so undetectability holds only against the public. *Removal research.* We publish attacks that destroy the mark; they help an adversary, but a robustness claim no one can attack is not a claim. *Environment.* The reported per-scene runs take about 1090 GPU and CPU seconds and .04 kg CO_2e, a floor on the project total (Appendix I); the method trains nothing.

Reproducibility statement

The anonymized code in the supplementary material contains the implementation, tests, exact commands, source and checkpoint hashes, recorded failures, reference packages, and the generators of every data-derived table and figure. Theorems 1 and 2 are proved in Section 3 and Appendix A, which lists every assumption; Section 5 and Table 6 give scene sources, splits, keys, budgets, and

the frozen selection rule, and Appendix I gives hardware, environment, and timings. The audit of Fig. 4 and the spectrum of Table 7 regenerate from per-scene probes in the supplementary material.

REFERENCES

Bernhard Kerbl et al. 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023.

Brian Chen and Gregory W. Wornell. Quantization index modulation: A class of provably good methods for digital watermarking and information embedding. *IEEE Transactions on Information Theory*, 47(4):1423–1443, 2001.

Brian Chen and Gregory W. Wornell. Preprocessed and postprocessed quantization index modulation methods for digital watermarking. *SPIE Security and Watermarking of Multimedia Contents II*, 3971, 2000.

Oktay Altun, Gaurav Sharma, and Mark Bocko. Set theoretic quantization index modulation watermarking. *ICASSP*, 2006.

Zixuan Chen et al. GuardSplat: Efficient and robust watermarking for 3D Gaussian splatting. *CVPR*, pp. 16325–16335, 2025.

Yijia Guo et al. Splats in Splats: Robust and effective 3D steganography towards Gaussian splatting. *AAAI*, 40(6):4485–4493, 2026.

Frederik J. Simons, F. A. Dahlen, and Mark A. Wieczorek. Spatiospectral concentration on a sphere. *SIAM Review*, 48(3):504–536, 2006.

Hao Qin et al. NGS-Marker: Robust native watermarking for 3D Gaussian splatting. *ICLR*, 2026.

Xiufeng Huang et al. GaussianMarker: Uncertainty-aware copyright protection of 3D Gaussian splatting. *NeurIPS*, 2024.

Sumin In et al. CompMarkGS: Robust watermarking for compressed 3D Gaussian splatting. *ICLR*, 2026.

Nan Li et al. Write where it matters: Policy-guided watermarks for 3D Gaussian splatting. *CVPR*, pp. 6580–6590, 2026.

Wenkai Huang et al. Can protective watermarking safeguard the copyright of 3D Gaussian splatting? *AAAI*, 2026.

Youngdong Jang, Hyunje Park, Feng Yang, Heeju Ko, Euijin Choo, and Sangpil Kim. 3D-GSW: 3D Gaussian splatting for robust watermarking. *CVPR*, 2025.

Boyu Wang, Jun Xia, and Mingsong Chen. Robust3DGSW: Toward robust watermarking for quantization-aware 3D Gaussian splatting. *CVPR*, pp. 19076–19084, 2026.

Heng Yang, Jingnan Shi, and Luca Carlone. TEASER: Fast and certifiable point cloud registration. *IEEE Transactions on Robotics*, online publication 2020.

Lixin Fan, Kam Woh Ng, and Chee Seng Chan. Rethinking deep neural network ownership verification: Embedding passports to defeat ambiguity attacks. *NeurIPS*, 2019.

Richard W. Hamming. Error detecting and error correcting codes. *Bell System Technical Journal*, 29(2):147–160, 1950.

Hugo Krawczyk, Mihir Bellare, and Ran Canetti. HMAC: Keyed-hashing for message authentication. RFC 2104, 1997.

Ben Mildenhall et al. NeRF: Representing scenes as neural radiance fields for view synthesis. *ECCV*, 2020.

Ben Mildenhall et al. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics*, 2019.

Jonathan T. Barron et al. Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022.

Arno Knapitsch et al. Tanks and Temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4), 2017.

Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.

Ethan Rublee et al. ORB: An efficient alternative to SIFT or SURF. *ICCV*, pp. 2564–2571, 2011.

Charles R. Qi et al. PointNet: Deep learning on point sets for 3D classification and segmentation. *CVPR*, pp. 652–660, 2017.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *ICML*, pp. 448–456, 2015.

A NOTATION, PROOFS, AND ANALYSIS

Table 5 collects the symbols.

Implementation. Eq. 5 runs inside the pipeline of Section 4, and each write is reread and certified after serialization. Float64 Cholesky factors solve without inversion, and certificates use the float32-rounded update. Masses are measured by differentiating the raw rasterizer sum against a positive override colour, excluding output clamps and saturation masks. The 8-bit (uint8) export stores each SH column at 256 levels with float32 minimum and step, $48N + 384$ bytes against $192N$. Ties, nonfinite values, and uncovered bins erase, and failed registration fails. The portable degree-three reference package binds archive and cap and needs a validated $\mathbf{M}_i$. Commitments cover identity, unattested time, and source, renderer, camera, and checkpoint hashes and are renewed on any source change; production keys must be restricted, and the released environment pins NumPy 1.26.4.

Assumptions. Inputs are finite and $\mathbf{x}_i \neq \mathbf{o}_v$, so every $\mathbf{d}_{iv}$ and $\mathbf{y}_{iv}$ is defined. Masses satisfy $w_{iv} \geq 0$ and $\sum_i a_{ivp} \leq 1$, which makes the weighted Cauchy–Schwarz step of Theorem 2 an upper bound. Fixed geometry, opacity, background, camera parameters, and compositing order, with trivial exposure, keep the compositing weights identical across the write. The ridge $\rho_i > 0$ gives $\mathbf{H}_i \succeq \rho_i \mathbf{I}_{16} \succ 0$, so Eq. 4 minimizes a positive definite quadratic over a nonempty affine set, which gives existence and uniqueness in Theorem 1; any feasible perturbation $\boldsymbol{\psi}$ of the optimum adds exactly $\boldsymbol{\psi}^\top \mathbf{H}_i \boldsymbol{\psi} \geq 0$, so the optimum is global. Both clamps are 1-Lipschitz, so the bound holds for finite writes without corner derivatives. Nothing is assumed about the renderer's convexity, the certificate is not claimed outside $\mathcal{V}$, and serialized costs are recomputed.

Proof of Theorem 2. The write changes SH coefficients only. Every factor of $a_{ivp} = \alpha_{ivp} T_{ivp}$ depends on geometry and opacity, which the certificate holds fixed, so the compositing weights, the culling and alpha thresholds, and the early-termination point are identical before and after the write. We may therefore compare the two renders weight by weight. Both clamps are nonexpansive: restricting a value to its allowed interval cannot increase the distance between two values. Let υ_i be the clamped primitive-colour change with $a_i = a_{ivp}$ at fixed v, p, c. Then $|\upsilon_i| \leq |\mathbf{y}_{iv}^\top \boldsymbol{\delta}_{ic}|$, and the fixed background term $T_{vp}^{\mathrm{end}} \beta_c$ cancels in the difference. Weighted Cauchy–Schwarz followed by $\sum_i a_i \leq 1$ from Section 3.1 gives $|\sum_i a_i \upsilon_i|^2 \leq (\sum_i a_i) \sum_i a_i |\upsilon_i|^2 \leq \sum_i a_i |\mathbf{y}_{iv}^\top \boldsymbol{\delta}_{ic}|^2$, where the nonnegativity of the masses is what makes the middle step an inequality rather than an estimate. Summing over pixels collapses $\sum_p a_{ivp}$ to w_{iv}, so summing over pixels, cameras, and channels yields $\sum_{i,c} \boldsymbol{\delta}_{ic}^\top \mathbf{M}_i \boldsymbol{\delta}_{ic}$, the first bound. Finally $\mathbf{H}_i - \mathbf{M}_i = \rho_i \mathbf{I}_{16} \succ 0$ for $\rho_i > 0$ gives the second. □

Euclidean comparison. Cauchy–Schwarz on $\mathbf{H}^{\pm 1/2}\mathbf{g}$ gives $(\mathbf{g}^\top \mathbf{H}^{-1} \mathbf{g})(\mathbf{g}^\top \mathbf{H} \mathbf{g}) \geq 1$ for unit $\mathbf{g}$, hence $D^* \leq D_{\mathrm{Euc}} = r^2 \mathbf{g}^\top \mathbf{H} \mathbf{g}$. Equality holds for an eigenvector; at $r = 0$ both vanish. This compares identical targets and metrics, not off-domain error or robustness.

Gaussian prediction. For $h, \sigma > 0$, independent isotropic noise yields $\mathsf{E} \sim \mathcal{N}(0, \sigma^2)$ under unit projection. Correct cells $[(2z - 1/2)h, (2z + 1/2)h]$ give

$$p_{\text{correct}}(h, \sigma) = \sum_{z \in \mathbb{Z}} \left[\Phi\left(\frac{(2z + 1/2)h}{\sigma} \right) - \Phi\left(\frac{(2z - 1/2)h}{\sigma} \right) \right]. \tag{8}$$

Equation 8 is what Fig. 4 audits; it predicts raw accuracy under independent coefficient noise with mean absolute error 0.21 points. At $\sigma = 0$, ideal raw accuracy is one. Tie tolerance $0 \le \tau < 1$ shrinks the sufficient radius to $h(1-\tau)/2$; serialization and coded recovery are measured separately.

Table 5: Notation. Bold lower case marks vectors, bold upper case matrices, calligraphic letters sets and costs, sans serif random variables, and upright letters architecture sizes; $*$ marks an optimum, a prime an edited value, and a hat an estimate.

Indices and sets			
$i,\ N$	primitive; number written	c	RGB channel
$v,\ \mathcal{V}$	camera; declared set, $V = \lvert\mathcal{V}\rvert$	$p,\ Q_v$	pixel; pixels of camera v
j	eigen index	z	grid integer
Scene and renderer (Section 3.1)			
$\mathbf{x}_i,\ \mathbf{o}_v$	primitive and camera centres	$\mathbf{d}_{iv}$	unit viewing direction
$\mathbf{s}_{ic},\ d$	SH coefficients in $\mathbb{R}^d$, $d = 16$	$\mathbf{y}_{iv}$	SH basis values, degree ≤ 3
$f_{icv},\ I_{vpc}$	primitive colour; rendered value	$\alpha_{ivp},\ T_{ivp}$	opacity; transmittance
a_{ivp}	compositing weight $\alpha_{ivp} T_{ivp}$	$T_{vp}^{\text{end}},\ \beta_c$	residual transmission; background
w_{iv}	contribution mass $\sum_p a_{ivp}$	P	pixel-channel count of the MSE
Metric and write (Section 3)			
$\mathbf{M}_i,\ \mathbf{H}_i$	contribution metric; $\mathbf{M}_i + \rho_i \mathbf{I}_{16}$	ρ_i	ridge
$\lambda_{ij},\ \mathbf{u}_{ij}$	eigenpairs of $\mathbf{M}_i$	κ_{ic}	inverse sensitivity $\mathbf{g}^\top \mathbf{H}^{-1} \mathbf{g}$
$\boldsymbol{\delta}_{ic},\ \boldsymbol{\delta}_{ic}^*$	write; optimal write	μ	Lagrange multiplier
$D_{ic},\ D^*$	slot cost; optimal cost	B	total certified budget
Quantizer and code (Sections 3 and 4)			
K	256-bit key	$\mathbf{g}_{ic},\ \xi_{ic}$	unit keyed measurement; dither
$h_{ic},\ \mathcal{Q}_b$	grid spacing (strength); grid of bit b	b_{ic}	padded bit
$t_{ic},\ q_{ic},\ r_{ic}$	host projection $\mathbf{g}^\top \mathbf{s}$; target; residual	$\mathbf{e}_{ic},\ \epsilon_{ic}$	inspection error; its bound
η	measurement-frame error	τ	tie tolerance
Protocol and statistics (Section 4, Appendix A)			
$\gamma,\ \mathbf{R},\ \mathbf{t}$	similarity scale, rotation, translation	$C,\ \mathrm{ID}$	consistency; exact identifier
$n,\ k$	non-erased slots; agreements	$\alpha,\ J$	test level; family size
ζ	ratio of successive tail masses	$\sigma,\ \Phi$	noise level; normal CDF
$\mathsf{B}, \mathsf{X}, \mathsf{Z}$	payload, images, reconstruction	$L,\ \mathsf{E}$	payload bits; projected noise
Controls and costs (Appendices G and I)			
$\mathbf{U}_i,\ k_{\mathrm{c}}$	carrier basis; carrier count	$\mathbf{z}_{ic},\ \nu_i$	carrier signs; magnitude
$\omega,\ \mathbf{S}_i$	carrier strength; host coefficients	$A_0,\ A_1$	carrier write and clearing costs
ϕ	smallest recovering retained fraction	$\mathcal{R}_f, \mathcal{R}_b, \mathcal{R}_a$	renderer passes
$\mathcal{T}_j,\ \mathcal{A}_j$	operator work; activations	$\mathcal{W}_j,\ \mathcal{S}_j$	weights; scene storage
$L_{\mathrm{t}},\ E_{\mathrm{t}}$	tile entries; pixel-primitive visits	$N_{\mathrm{a}},\ f_{\mathrm{a}}$	anchors; anchor width
$\mathrm{b}, \mathrm{m}, \mathrm{k}, \mathrm{t}$	batch, sites, kernel, tokens	u, v	layer widths

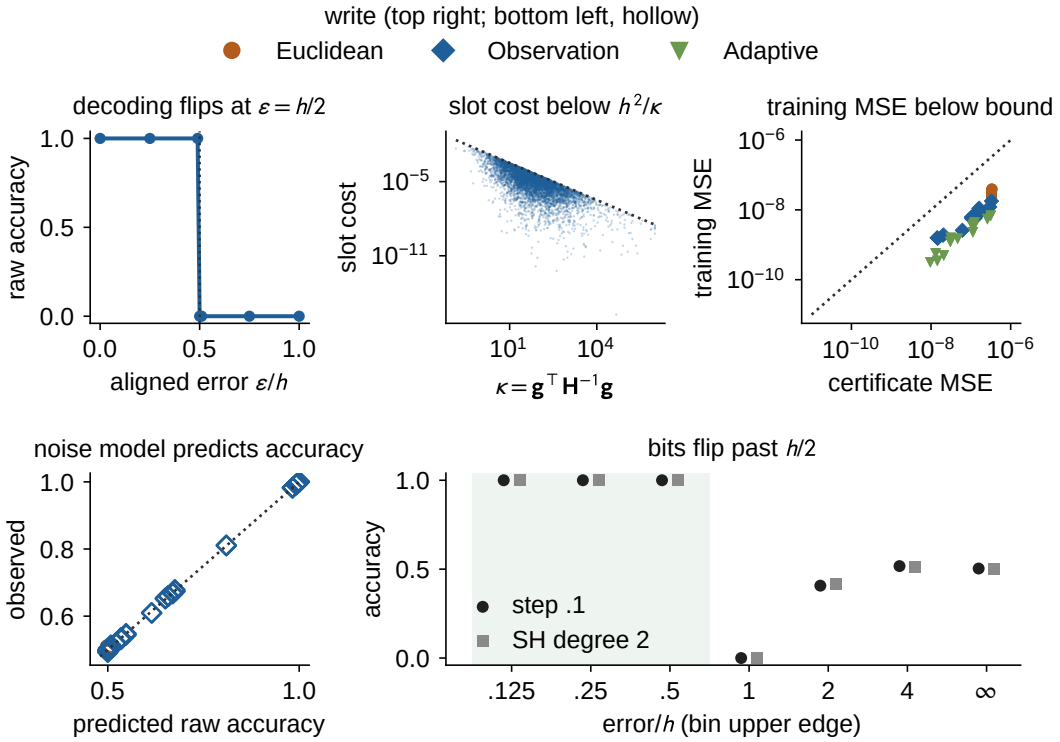

Figure 4: Audits of the analysis. Top: raw accuracy is 1 up to aligned error $\epsilon/h = .49$ and 0 from .50 under all four relative ridge weights (10^{-4} to 10^{-1}), ties at $\epsilon/h = 1/2$ being erased; on Lego every slot's cost stays below the h^2/κ envelope; and every primary write's training MSE lies below its certificate. Bottom: Eq. 8 predicts raw accuracy for the Euclidean and observation writes, and slots pooled by projected-error bin flip once the error passes $h/2$ (shaded). Dotted lines mark $\epsilon = h/2$, the h^2/κ envelope, and equality; marker shapes identify the write. Each audit agrees with its analytic prediction.

Registration errors. The interval of Eq. 7 excludes missing primitives and incorrect registration: a frame change $\|\widehat{\mathbf{g}} - \mathbf{g}\| \le \eta$ grows the scalar error to $\epsilon + \eta\|\mathbf{s} + \boldsymbol{\delta} + \mathbf{e}\|$, so a small frame error matters for a large host norm. Committed strengths stop a verifier silently recomputing a different quantizer.

Null with erasures and search. Condition on the pad-independent choices of Section 4. Distinct primitive/channel ideal pads give fair raw agreements even with repeated code indices; erasures only lower counts. Union-bound J binomial tests; HMAC adds PRF advantage, excluding pad-dependent selection. For $k > n/2$, successive masses have ratio $\le \zeta = (n-k)/(k+1) < 1$, hence $\Pr(X \ge k) \le \Pr(X = k)/(1-\zeta)$ by geometric summation.

B Registration and Protocol Details

Section 4 summarizes these details. SHA256 hashes of float64 reference centres give persistent identities independent of row order. Eight-neighbour normalized distances initialize RANSAC, fitting similarities from three pairs by geometric support with identity as a candidate; refinement, mutual matching, and best/second-best ratios reject ambiguity. SH transport is realized by degree-wise collocation at fixed Fibonacci directions and checked numerically. The keyed PRF uses separate domains for measurements, dithers, placement, sampling, and pads. Table 6 lists the experimental protocol.

Table 6: Experimental protocol. Every quantity influencing an embedding is fitted on training data alone, the three variance sources are never pooled, and the 8192-primitive cap and Hamming coding are ablated.

Aspect	Setting
Split discipline	$\mathbf{M}_i$, h_{ic}, amplitude fitted on training cameras only
	calibration key 99033 disjoint from evaluation keys 0–4
	8-bit codec applied after embedding, so it crosses no split
Scene variance	95% paired bootstrap, 10^4 resamples, scenes drawn, keys held
Ablated	metric (Euclidean, diagonal, angular); measurement (spectral)
	ridge (swept); carrier (additive, clearing)
Negative controls	unmarked returns MSE and LPIPS exactly 0
	unmarked, wrong-key, cross-copy claims 0/4 each; null 0/300
	reused-key wrong identifier: C 4/4, ID 0/4
Scale	all unique-centre primitives written, 158,483–3,640,214 per scene
	$\rho_i = 10^{-3}\,\mathrm{tr}(\mathbf{M}_i)/16 + 10^{-6}$, budget 10^{-6}
	float64 solve, float32 output
Key-draw variance	5 keys $\times$ 14 scenes; spread sits inside a count out of 70
Checkpoint variance	isolated on ship and fern, two 15000-step optimizer seeds
Baseline parity	each baseline keeps its published schedule
	GuardSplat 100 epochs, GaussianMarker 5000 steps
	final checkpoint, never best-on-held-out; no per-scene tuning
Compute	not equalized: baselines train minutes per scene
	and store 3.9–27.9 MB of weights; our write is a solve
Cap and coding	every matched primitive read, not 8192: observation IDs
	after SH2, ridge fit, strip 38, 15, 21 of 42 (capped 15, 3, 0)
	Hamming correction adds 23 exact observation IDs over Table 12's capped reads
Arms at ceiling	unattacked 42/42 for every method
	every observation, angular, and spectral arm at budget 10^{-5}
Separating arms	noise, step rounding, SH2, and actual 8-bit export
	counts there spread 0/42 to 42/42

C SPECTRUM OF THE CONTRIBUTION METRIC

Table 7 measures the conditioning the introduction asserts. The seven full-sphere NeRF synthetic objects resolve all sixteen SH directions and spread 1.7–5.7 orders; the forward-facing LLFF scenes and the unbounded captures resolve a median of six and spread 8.7–15.7 orders, so over half their SH directions carry no observed evidence. One Euclidean budget cannot serve both regimes, and the renderer metric gains most on the second.

D SCENE-LEVEL MEASUREMENTS

Table 8 adds perceptual metrics to the common-budget comparison of Section 5.2, and Table 9 breaks Table 1 down by scene. At key 0 alone the off-path inflation reproduces at 764 times in MSE and 541 times in LPIPS, so LPIPS sees it; held-out cameras separate the methods by LPIPS 2.0×10^{-5} against 1.8×10^{-5}.

Where the adaptive allocation sits. Adaptive QIM matches or beats observation QIM on every attacked row of Table 12, bought not free: a common budget fixes the certified quantity, not rendered quality, and adaptive writes where it is cheap, which is where the cameras do not look, so its held-out and off-path MSE far exceed observation QIM's (Table 1). Adaptive off-path LPIPS is .0207 against .0085 for observation QIM (Table 8), with worst views .290 and .0406. We therefore report observation QIM as the variant that buys recovery while holding held-out quality near the Euclidean control, not as the most recoverable one we measured.

Table 7: Spectrum of $\mathbf{M}_i$ per scene: medians over a fixed sample of up to 4000 primitives; effective rank counts eigenvalues above 10^{-6} of the leading one, and spread is $\log_{10}(\lambda_1/\lambda_{16})$, floored at machine precision. Full-sphere rigs resolve all sixteen directions; the other captures resolve a median of six.

Scene	Primitives	Median effective rank / 16	Spectrum spread (orders)
bonsai	3898	6.0	15.7
chair	3979	16.0	5.3
counter	3903	13.0	8.7
drums	3920	16.0	5.4
fern	3975	6.0	15.7
ficus	3977	16.0	1.7
garden	3986	9.0	13.9
hotdog	3901	16.0	5.3
materials	3968	16.0	3.3
mic	3875	16.0	5.7
room	3691	6.0	15.7
ship	3935	16.0	5.5
trex	3709	6.0	15.7
truck	3913	5.0	15.7
Median over 14 scenes	—	14.5	7.2

Table 8: SSIM and LPIPS against the unmarked render at the common budget, all fourteen scenes, key 0; off-path cameras have no ground truth. Observation QIM's held-out cost matches the Euclidean control's, and its off-path cost stays small in SSIM.

	Held-out cameras			Off-path cameras		
Method	MSE	SSIM	LPIPS	MSE	SSIM	LPIPS
Unmarked control	0.00e+00	1.000000	0.00e+00	0.00e+00	1.000000	0.00e+00
Euclidean QIM	3.01e-08	0.999996	1.76e-05	3.23e-08	0.999997	1.58e-05
Observation QIM	3.07e-08	0.999996	2.00e-05	2.47e-05	0.998272	8.53e-03
Adaptive QIM	4.20e-06	0.999953	2.03e-04	1.25e-03	0.992853	2.07e-02

E MATCHED-MSE CONTROLS AND DISTORTION–RECOVERY FRONTS

Three-key sweeps use budgets 10^{-8}–10^{-5} on Lego and four evaluation scenes, and Tables 2 and 11 give the controls at matched training MSE and their spread. Calibration key 99033 uses up to eight square-root amplitude corrections and a 3% tolerance, and the frozen factor multiplies each evaluation base strength. Additive controls share capped sampling and the raw test; two-carrier votes differ.

Matching training MSE. Calibration key 99033 freezes an amplitude factor from training renders; the angular and adaptive writes land in $[0.906, 1.065] \times 10^{-8}$, 75 of 84 within 5%, and the other 168 writes in $[0.736, 1.120] \times 10^{-8}$, 157 within 5%, all mismatches included.

Angular arms, ridge weight, and a frozen budget rule. A certified budget bounds the spectral write through the ridge, which matching training MSE does not (Table 2). At 10^{-6}, angular QIM recovers 12, 3, 1, and 12 in the noise, step, SH2, and 8-bit arms of Table 10 and, like observation QIM, every arm at 10^{-5}. The four-scene observation-minus-angular gaps have scene-bootstrap intervals including zero; the Euclidean and diagonal contrasts under noise and 8-bit export do not (completion controls below). A larger ridge trades attacked recovery for off-path fidelity: on the sweep at 10^{-6} with key 0, relative weight .01 keeps noise 4/4 but recovers none under step or SH2, at 10 times lower off-path MSE than the default 10^{-3}, and weights .1 and 1 keep only clean recovery. A frozen Lego-only rule picks 10^{-8}, 3/3 on Lego but 0/12 on the four sweep scenes, a failure we report without retuning.

Table 9: Per scene at the common budget of Table 1. Left: training MSE in units of 10^{-9}, with the pooled 14-scene mean and its 95% scene-bootstrap interval, which exclude the development scene lego. Right: observation QIM's GT-PSNR change (ΔPSNR, dB) on the four-view diagnostic, and the smallest tested retained fraction ϕ that still recovers the identifier (retained primitives in brackets) under cropping and under an unknown similarity (Sim(3)); ϕ is a minimum over tested fractions, not a threshold. Observation QIM's training error is below the Euclidean control's on every scene, its GT-PSNR moves by at most .0014 dB, and at most 1% of the primitives suffices on every scene.

	Training MSE ($\times 10^{-9}$)			Observation QIM		
Scene	Euclidean	Observation	Adaptive	ΔPSNR	Crop ϕ (count)	Sim(3) ϕ (count)
bonsai	26.2	5.72	1.20	−0.0006	0.001 (1031)	0.001 (1031)
chair	37.9	11.1	4.32	−0.0014	0.003 (898)	0.003 (898)
counter	27.8	5.69	1.48	−0.0000	0.001 (919)	0.001 (919)
drums	31.4	9.62	3.21	−0.0003	0.003 (959)	0.003 (959)
fern	39.9	1.58	0.563	−0.0000	0.001 (753)	0.001 (753)
ficus	22.8	17.8	6.76	−0.0005	0.003 (532)	0.003 (532)
garden	29.8	6.22	1.57	−0.0001	0.0003 (1092)	0.0003 (1092)
hotdog	22.2	7.59	3.72	−0.0001	0.003 (475)	0.003 (475)
lego (dev.)	43.6	13.5	4.59	−0.0003	0.003 (899)	0.01 (2998)
materials	32.1	12.1	5.40	−0.0002	0.003 (710)	0.003 (710)
mic	25.3	8.85	2.27	−0.0006	0.003 (828)	0.003 (828)
room	37.4	1.89	0.343	−0.0004	0.003 (693)	0.003 (693)
ship	27.7	8.09	3.32	−0.0003	0.003 (928)	0.003 (928)
trex	36.9	1.54	0.312	+0.0001	0.001 (551)	0.001 (551)
truck	21.8	2.60	0.489	+0.0005	0.001 (1139)	0.001 (1139)
14-scene mean	29.9	7.17	2.50			
95% CI	[27, 33]	[4.9, 9.6]	[1.5, 3.6]			

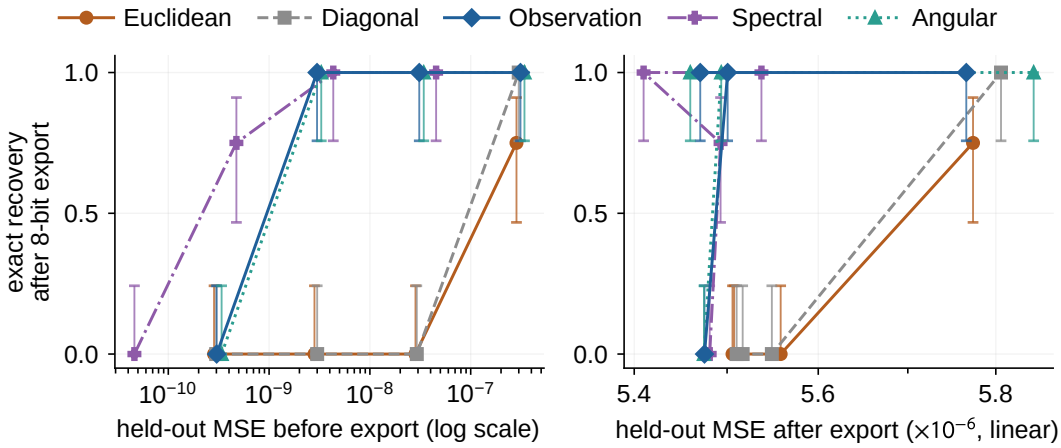

Figure 5: Exact recovery after 8-bit export against held-out MSE (all held-out cameras) before export (left, log axis) and after it (right, a linear axis over a narrow range, so small MSE differences look large), for five writes at four budgets on four evaluation scenes and three keys. Each point is a rate over twelve trials, connected points share a method, colour, marker, and line style identify it, and error bars are 95% Wilson score intervals. Similar post-export MSE can coexist with very different recovery.

Completion controls. Five methods, four budgets, three keys, and five scenes yield 300 writes and 1500 verifications with all-camera quality. The plan follows earlier results but precedes these runs. The rest of this paragraph isolates the conservative budget 10^{-6}; all outcomes remain in Fig. 5. At budget 10^{-6}, median/worst scene-mean clean held-out MSE is Observation: 8.75×10^{-9}, 9.88×10^{-8}; Angular: 1.11×10^{-8}, 1.09×10^{-7}; Spectral: 3.62×10^{-9}, 9.85×10^{-9}. The paired observation-minus-spectral mean MSE difference is 2.66×10^{-8} (95% scene-bootstrap

Table 10: The four-scene sweep at budget 10^{-6} (ship, fern, bonsai, truck; three keys): clean and post-export held-out MSE, exact identifiers of 12 per arm, and stored reference bytes per primitive. Bold: best; underline: second best, where it differs from the worst. Spectral QIM leads or ties every column but storage.

Method	Clean MSE	8-bit MSE	Noise	Step	SH2	8-bit	Ref./N
Euclidean	<u>2.83e-08</u>	5.56e-06	0	0	0	0	**48**
Diagonal	2.92e-08	5.55e-06	0	0	0	0	**48**
Observation	3.09e-08	5.47e-06	**12**	<u>6</u>	<u>6</u>	**12**	**48**
Spectral	**4.33e-09**	**5.41e-06**	**12**	**12**	**12**	**12**	176
Angular	3.42e-08	<u>5.46e-06</u>	**12**	3	1	**12**	**48**

Table 11: Five keys at matched training MSE, 70 trials per write: mean $\pm$ sd across the fourteen scene means and the Wilcoxon p against the Euclidean control. The camera-aware writes keep training MSE matched and recover far more often, at significantly higher held-out MSE.

	Training MSE ($\times 10^{-8}$)		Held-out MSE ($\times 10^{-8}$)		8-bit exact ID rate	
Method	mean $\pm$ sd	p	mean $\pm$ sd	p	mean $\pm$ sd	p
Euclidean	1.00 ± 0.01	—	1.00 ± 0.07	—	0.14 ± 0.36	—
Diagonal	1.00 ± 0.01	0.583	1.16 ± 0.11	<0.001	0.14 ± 0.36	n/a
Angular	0.99 ± 0.03	0.194	3.15 ± 2.92	<0.001	0.93 ± 0.27	<0.001
Observation	1.00 ± 0.02	0.626	5.53 ± 5.03	<0.001	1.00 ± 0.00	<0.001
Adaptive	1.00 ± 0.01	0.952	1.33e3 $\pm$ 3.01e3	<0.001	1.00 ± 0.00	<0.001
Spectral	0.99 ± 0.07	0.903	2.07e5 $\pm$ 5.33e5	<0.001	1.00 ± 0.00	<0.001

interval $[4.29 \times 10^{-9}, 6.81 \times 10^{-8}]$). Observation-minus-angular exact-recovery differences are 25 percentage points for step rounding and 42 for SH2; their 95% intervals are [0.0, 75.0] and [0.0, 83.3]. These 10000 paired resamples preserve scenes, not individual keys, and only four scenes limit precision. Observation-minus-Euclidean differences are 100 points under noise and 8-bit export, with intervals [100.0, 100.0], and 50 under step and SH2, with [0.0, 100.0]; diagonal's counts equal Euclidean's at this budget.

F ATTACKS, FRAGMENTS, AND TRANSFER

Operations and terminology. Noise adds independent Gaussian SH errors; quantization rounds values; codebooks replace them with learned representatives. Truncation zeros degrees; stripping removes low-eigenvalue components; ridge fitting penalizes coefficient norms. These preserve geometry. Pruning deletes primitives; cropping retains regions; insertion mixes scenes; jitter moves centres. Sim(3) combines rotation, translation, and positive scale. Reparameterization replaces primitives; fine-tuning continues optimization; densification creates/splits primitives.

ORB corners and binary descriptors (Rublee et al., 2011) triangulated with granted poses initialize image-only reconstruction: sparse structure from motion (SfM) without pose estimation. PointNet aggregates point features without order dependence (Qi et al., 2017). Batch normalization rescales features using current-batch or stored running statistics (Ioffe and Szegedy, 2015); both modes are tested.

Fragment resolution and correspondence errors. One new key and one 1% sphere, box, or half-space per scene yield twelve crops and four full scenes under a shared unknown similarity. Uint16 uses per-axis scene minimum and step; float16 rounds coordinates. Jitter is relative to global extent. NN scale is the median nearest-other-centre distance from at most 20000 fixed centres. Of the twelve crops, uint16 keeps 10 identifiers and float16 six, the latter at a crop-view MSE of 5.49×10^{-7}; of the four full scenes, uint16 keeps all and float16 three. Jitter of 10^{-6}, 10^{-5}, and 10^{-4} of extent, median RMS .009, .09, and .90 NN spacing, leaves 16, 13, and 0 of 16 identifiers. Across jitter, uint16, and float16, 21 losses fail registration outright and 7 register a wrong frame; C accepts none of them. True correspondence and frame restore perfect raw reading in all 96 diagnostics, never

Table 12: The edits of Table 1 plus SH step .01: consistency (C) and exact identifiers (ID) of 42 scene-key trials. Every attack runs registration, SH transport, and the 8192-primitive cap, and coefficient edits preserve geometry. Cost is the GT-PSNR change (dB) on the four-view held-out diagnostic of the key-0 observation write; the strip and ridge fit use the declared cameras' metric but no key. Strong noise, truncation, and the two keyless edits remove most observation identifiers.

	Euclidean		Observation		Adaptive		
Edit	C	ID	C	ID	C	ID	Cost
Unattacked	42	42	42	42	42	42	—
Noise $\sigma = .01$	0	0	42	41	42	42	-0.033
Noise $\sigma = .05$	0	0	21	9	42	37	-0.738
SH step .01	42	18	42	42	42	42	-0.003
SH step .1	0	0	36	24	42	39	-0.312
Keep SH degree 2	21	1	39	15	42	37	-1.449
Remove 4 low-cost directions	0	0	20	0	36	17	-0.195
Ridge coefficient fit	6	3	18	3	25	6	+0.014

counted as deployable success. Contribution coverage measures rendering sensitivity, not semantic area. Earlier same-frame jitter can use the geometry-only identity candidate. Even estimated cameras suffice on fern: a coarse-rig attacker with 16 perturbed poses, teacher images, and no key removes confirmation below 10^{-6} training MSE among nineteen projection, ridge, and codebook candidates. Disclosed-key flips defeat confirmation; blind targeting does not.

Native training attacks. The seed-zero teacher exports 8-bit training PNGs. Reconstruction submits its complete cloud after 15000 steps. Fine-tuning uses 5000 steps, rate scale .1, all parameters, and no densification. Fern/bonsai random initialization densifies until step 4500, threshold .0004; ship uses native defaults. Teacher ORB initialization densifies until step 10000. Both image-initialized runs register but lose confirmation. Fixed budgets do not imply equal convergence; poor random reconstructions do not establish useful removal. A reconstruction Z computed from rendered images X with no side input on a uniform L-bit payload B forms a Markov chain $\mathsf{B} \to \mathsf{X} \to \mathsf{Z}$, so the data-processing inequality gives $\mathrm{I}(\mathsf{B};\mathsf{Z}) \le \mathrm{I}(\mathsf{B};\mathsf{X})$, and recovery is at most 2^{-L} when the images carry no information about the payload; claims supplied to the verifier are not recovered information, and a nearly invisible mark needs a noise or precision model before this applies.

Composed transfer. Transfer training uses seed 20260913, resolution factor four, and densification until step 4500 (threshold .0004). Training cameras define the metric and all test cameras the quality. Keys 8101–8103 and crop seed 983711 are fixed before evaluation. Sphere, rotated box, and halfspace crops retain 1%. SH, geometry, covariance, and opacity are serialized after export/crop/transform, then verified uncached. True transforms serve paired rendering only, never registration or decoding. The development-selected budget 10^{-8} yields 27/27, 0/27, 0/27 exact IDs for clean, 8-bit, and 8-bit plus float16 workflows. This independently selected setting fails export transfer and is not replaced using these outcomes. At the primary budget, mean 8-bit crop MSE ranges from 4.39×10^{-9} to 1.31×10^{-8}; float16-coordinate arms range from 2.00×10^{-7} to 3.58×10^{-7}. Cropped images can contain substantial background, so records also report error relative to original foreground energy. Three scenes, one checkpoint seed, and one crop per shape do not imply population reliability.

G BASELINES AND CONTROLS

GaussianMarker and Robust3DGSW support parameter and image decoders, NGS-Marker reads native fragments, and 3D-GSW, CompMarkGS, W2M, and GuardSplat read rendered images. We run the two schemes whose code we could obtain and execute, GuardSplat and GaussianMarker; the other five are positioned by contract, because the linked NGS-Marker endpoint could not be retrieved and 3D-GSW, CompMarkGS, W2M, and Robust3DGSW were not integrated, so our comparative claims cover the two measured schemes only. In Table 3 the noise, strip, and ridge rows were rescored where GuardSplat's unattacked read differs from the recorded one by at most 2 of 768 bits,

Table 13: Composed transfer to three unseen captures (three keys $\times$ three 1% crop shapes per scene and setting; every arm includes crop, unknown similarity, and serialization). Counts are exact identifiers; seconds are median uncached verification and its registration stage. At 10^{-6} every 8-bit arm recovers and float16 coordinates cost 9 of 27; the development-selected 10^{-8} loses every 8-bit arm.

Scene	Budget	Clean	8-bit	+float16	8-bit MSE	Verify (s)	Reg. (s)
kitchen	1.00e-06	9/9	9/9	6/9	9.59e-09	4.39	1.63
bicycle	1.00e-06	9/9	9/9	9/9	1.31e-08	6.57	2.70
stump	1.00e-06	9/9	9/9	3/9	4.39e-09	7.00	2.65
kitchen	1.00e-08	9/9	0/9	0/9	7.46e-09	4.31	1.61
bicycle	1.00e-08	9/9	0/9	0/9	1.17e-08	6.14	2.39
stump	1.00e-08	9/9	0/9	0/9	3.73e-09	7.23	2.83

Table 14: Scheme contracts: carrier, reader input, and guarantee ("certified: a proved bound on rendered error over $\mathcal{V}$; "empirical: fidelity measured, not bounded); Bits is the capacity we could verify. Only the two QIM writes bound their rendered error.

Method	Carrier	Reader input	Guarantee	Bits
Ours, observation	SH coeff.	params + key	certified on $\mathcal{V}$	256
Ours, spectral	SH coeff.	params + key	certified on $\mathcal{V}$	256
GuardSplat (Chen et al., 2025)	SH coeff.	renders	empirical	≤ 72
GaussianMarker (Huang et al., 2024)	params + density	params or renders	empirical	48
3D-GSW (Jang et al., 2025)	params + density	renders	empirical	—
NGS-Marker (Qin et al., 2026)	params	native fragments	empirical	—
Robust3DGSW (Wang et al., 2026)	params	params or renders	empirical	—
CompMarkGS (In et al., 2026)	params	renders	empirical	—
W2M (Li et al., 2026)	params	renders	empirical	—

and keep-10% and keep-1% are omitted: at 18.2 and 11.8 dB the asset is destroyed for every scheme. GuardSplat's 72-bit ceiling is set by the tokenizer of its CLIP image-text model.

Two carrier controls replace the quantizer rather than the metric. With orthonormal $\mathbf{U}_i \in \mathbb{R}^{16\times k_c}$ and $\mathbf{z}_{ic} \in \{\pm\nu_i\}^{k_c}$ for $\nu_i = \|\mathbf{S}_i\|_F/4$, where $\mathbf{S}_i \in \mathbb{R}^{16\times 3}$ holds the host coefficients of primitive i, the additive control writes $\mathbf{s}'_{ic} = \mathbf{s}_{ic} + \omega\mathbf{U}_i\mathbf{z}_{ic}$ and the clearing control writes $\mathbf{s}'_{ic} = (\mathbf{I}_{16} - \mathbf{U}_i\mathbf{U}_i^\top)\mathbf{s}_{ic} + \omega\mathbf{U}_i\mathbf{z}_{ic}$, paying for host deletion; both read signs in stored bases. For a single carrier ($k_c = 1$), $A_0 = 3\sum_i \nu_i^2\,\mathbf{U}_i^\top\mathbf{H}_i\mathbf{U}_i$ and $A_1 = \sum_{ic}(\mathbf{U}_i\mathbf{U}_i^\top\mathbf{s}_{ic})^\top\mathbf{H}_i(\mathbf{U}_i\mathbf{U}_i^\top\mathbf{s}_{ic})$ are positive scalars, and an $\mathbf{H}$-norm triangle inequality bounds the total cost by $(\sqrt{A_1} + \omega\sqrt{A_0})^2$, giving sufficient strengths $\sqrt{B/A_0}$ and $(\sqrt{B} - \sqrt{A_1})/\sqrt{A_0}$; these ignore favourable cross terms and need not be necessary. The anonymized code in the supplementary material computes A_0 and A_1.

Native GaussianMarker reproduction. Its released parameter-space PointNet reads bit accuracy 1.000 on the marked and on the unmarked densified Lego alike, so Table 3 omits it. Removing unused fixed-input gradients preserves decoder outputs and gradients and prevents graph growth; assets, failures, and reloaded checkpoint hashes are retained; and sampling repair keeps full pools and adds identity augmentation. Paper-aligned training uses original negatives, positive-only message loss weighted by two, and Adam at 10^{-4}, whereas the released code uses densified negatives, batch-averaged message loss, and 10^{-5}; both failed adaptations remain reported. Balanced PointNet training draws eight 10000-point subsets (two intact and two augmented positives, two original and two densified negatives); seed 901 fixes 5000 Adam steps at 10^{-4}, halved at 2000 and 4000, and final-weight evaluation freezes running statistics and disables dropout. Training takes 313.5 L40S seconds, and the weights occupy 27.92 MB. Calibration at .00349 accepts 2/60 native and 70/75 foreign negatives, and separation at .5 does not establish general attribution.

H COEFFICIENT-STATISTICS DETECTABILITY

A Kolmogorov-Smirnov (KS) statistic is not a detector, so Table 15 turns the statistics of Table 16 into a decision and scores it against two reference distributions that bracket the adversary. Against a clean corpus pooled leave-one-scene-out from the other thirteen scenes, which is what an adversary holds, the asset-level area under the receiver operating characteristic curve (AUC) is $.490$, chance for this detector, which does not exclude a stronger one. Given a clean sibling of the very asset under test, which the threat model does not grant, AUC rises to .929. Undetectability is therefore claimed only for this detector and only while no clean reference of the same asset exists; a holder of the pre-release original can detect the mark, which concerns the release model, not the public adversary. Every reference primitive is written, so a released asset has no marked/unmarked partition to find; the primitive rows ask the weaker question of whether a primitive came from a marked asset, at AUC $.553$ without an oracle. Both asset ROCs rest on fourteen positives and fourteen negatives, so their resolution is $1/14$ in either rate.

Table 15: Detection decisions. Positives and negatives are odd-indexed halves of the marked and clean arrays, so both classes cover identical primitives and differ only by the mark; TPR is at the stated FPR. Without a clean copy of the same asset the asset-level detector is at chance.

Unit	Score	Reference distribution	$n_+ = n_-$	AUC	FPR	TPR
Asset	Max-band KS	Clean corpus (13 scenes)	14	0.490	0.000	0.071
Asset	Max-band KS	Clean half of same asset	14	0.929	0.000	0.929
Primitive	Mahalanobis	Clean corpus (13 scenes)	5022201	0.553	0.010	0.011
Primitive	Mahalanobis	Clean half of same asset	5022201	0.689	0.010	0.079
Primitive	SH energy	None	5022201	0.512	0.010	0.010

Table 16: Coefficient statistics behind Table 15 over fourteen scenes: KS is the maximum empirical-CDF difference, and the split control uses alternating original coefficients; correlated coefficients are not independent samples. The mark moves SH degrees 1 to 3 far more than the split control.

SH degree	Mean KS	Maximum KS	Split-control KS	Mean energy ratio
0	0.0060	0.0189	0.0012	1.001
1	0.0310	0.0661	0.0007	1.083
2	0.0360	0.1343	0.0005	1.038
3	0.0347	0.1364	0.0005	1.009

I COMPUTATIONAL COST

Timing method. Seconds are wall clock, and GPU blocks synchronize the device before starting and after stopping the clock. Times are cold, one-shot, and end-to-end, because a signing run happens once per asset; the cached and uncached verification columns give warm and cold readings. Verification (registration, transport, decoding) runs on NumPy arrays, so its seconds are CPU time. Metric construction has its own clock, separate from identity selection and clean renders.

Metric storage. The host metric costs $2048N$ bytes, 20.57 GB across the fourteen evaluation scenes and 7.46 GB for garden alone, and need not be stored: rebuilding all fourteen takes 904 s, 97% of it in the contribution-matrix loop (per-scene split in the supplementary material). Because the verifier takes only the reference centres, the committed strengths, and the key and never reads $\mathbf{M}_i$, storage buys faster re-signing, not faster verification.

Throughput and total compute. Across the fourteen evaluation scenes the write sustains a median 5.2×10^6 primitives per second and uncached verification 5.9×10^4. Verification scales linearly: a log-log fit of uncached seconds against primitive count over $158\,483$ to $3\,640\,214$ primitives, a factor of 23, has exponent 1.04. Cost depends on N, not on the budget, which changes the magnitude of a write but not its work, so the only frontier is quality against recovery (Fig. 5). The writes,

Table 17: Seconds on an L40S in the released environment, excluding I/O: preprocessing (Pre.), key setup (Key), three-key mean write, complete uncached (U) and cached (R) verification, recorded GPU peak, and reference size; N in millions. The write takes under a second everywhere; preprocessing and key setup dominate.

Scene	N (M)	Pre.	Key	Write	U	R	GPU (GB)	Ref. (MB)
lego	0.300	1.8	20.2	0.27	5.3	2.7	4.26	14.39
ship	0.310	3.1	20.6	0.12	5.1	2.8	4.47	14.86
chair	0.299	1.9	19.8	0.05	5.0	2.7	4.22	14.38
drums	0.320	1.9	21.1	0.05	5.1	2.8	2.98	15.35
ficus	0.177	1.2	11.7	0.03	3.0	1.8	1.79	8.51
hotdog	0.158	1.4	10.5	0.03	2.9	1.7	1.64	7.61
materials	0.237	1.7	15.6	0.05	4.0	2.3	2.30	11.37
mic	0.276	1.9	18.1	0.05	4.7	2.5	2.63	13.25
fern	0.754	8.6	49.9	0.25	12.5	6.5	1.11	36.18
trex	0.552	18.8	36.4	0.08	8.8	4.6	0.92	26.50
room	0.231	2.3	15.4	0.05	4.0	2.2	2.34	11.09
counter	0.919	101.0	61.6	0.15	16.1	8.2	1.19	44.12
bonsai	1.031	177.9	68.8	0.40	18.1	9.4	1.29	49.49
garden	3.640	414.2	243.2	0.70	73.5	37.6	4.55	174.73
truck	1.139	168.1	76.1	0.94	20.6	10.4	1.38	54.69

Table 18: Asymptotic costs for N primitives, $d = 16$ SH coefficients per channel, and $V = |\mathcal{V}|$ streamed cameras; renderer passes cost $\mathcal{R}_f$, $\mathcal{R}_b$, or $\mathcal{R}_a$ per camera. Learned schemes appear as a class because their operator counts $\mathcal{T}_j$ depend on the architecture (sizes in Table 5). Our write replaces training with a solve and certifies in $O(Nd^2)$ what learned schemes measure.

Cost	Observation-aware QIM	Learned schemes
Write work	$O(VNd^2)$ metric construction, then $O(Nd)$ writes from cached directions; no backward pass	$\mathcal{T}_j$ over all iterations, including sampling, loss and policy work: dense and convolutional $O(\mathrm{b\,m\,u\,v\,k}^2)$, attention $O(\mathrm{b}(\mathrm{tu}^2 + \mathrm{t}^2\mathrm{u}))$, reverse mode at the same order
Certification	$O(Nd^2)$	none; fidelity is measured, not bounded
Working memory	host metric $O(Nd^2)$, scene $O(Nd)$	$\mathcal{A}_j$ outputs and maps, $\mathcal{W}_j$ weights, and gradient and Adam storage a constant multiple of optimized parameters
Held by the verifier	$48N$ reference bytes, $+128N$ for the spectral carrier	$\mathcal{W}_j$ weights, 3.9 to 27.9 MB for the two run baselines
Scene storage $\mathcal{S}_j$	$O(Nd)$	$O(Nd)$ native or $O(N_\mathrm{a} f_\mathrm{a})$ anchor-based, plus decoded buffers

preprocessing, and verification of Table 17 total 1090 s, at most .11 kWh at an L40S board power of 350 W and, at an assumed grid intensity of .369 kg CO_2e per kWh, about .04 kg CO_2e. This is a floor: it excludes budget sweeps, the 300 completion writes and 1500 verifications, baseline training, abandoned decoder adaptations, and exploratory runs, and the energy is inferred from nominal board power, not metered.

Cost model. Costs use renderer work and storage $\mathcal{R}_f, \mathcal{R}_b, \mathcal{R}_a$, non-render work, activations, and weights $\mathcal{T}_j, \mathcal{A}_j, \mathcal{W}_j$, and scene storage $\mathcal{S}_j$ (Table 18). Per camera, Q pixels, L_t tile entries, and E_t pixel–primitive visits give $\mathcal{R}_f = O(Nd + L_\mathrm{t} \log L_\mathrm{t} + E_\mathrm{t} + Q)$, $\mathcal{R}_b = O(Nd + E_\mathrm{t} + Q)$, and $\mathcal{R}_a = O(Nd + L_\mathrm{t} + Q)$; fixed-width radix sorting makes the sort term L_t. Metric and angular construction add $O(VNd^2)$ for V streamed cameras; cached directions give $O(Nd)$ writes plus $O(Nd^2)$ certification; host metrics and scenes cost $O(Nd^2)$ and $O(Nd)$; and QIM writes need no backward pass. Learned $\mathcal{T}_j$ sums executed operators over all iterations, including sampling, loss, and policy work: dense and convolutional layers at $O(\mathrm{b\,m\,u\,v\,k}^2)$ and attention at $O(\mathrm{b}(\mathrm{tu}^2 + \mathrm{t}^2\mathrm{u}))$

for batch b, sites m, widths u, v, kernel k, and tokens t, with reverse mode at the same order; $\mathcal{A}_j$ stores outputs and maps and $\mathcal{W}_j$ weights. Fixed-width PointNet is linear in points. $\mathcal{S}_j = O(Nd)$ for native scenes or $O(N_{\mathrm{a}} f_{\mathrm{a}})$ for N_{a} anchors of width f_{a}, plus decoded buffers, with gradient and Adam storage a constant multiple of optimized parameters. Unexecuted FLOPs are unverified.