SAME BYTES, DIFFERENT SCENE: A SEMANTIC CONTRACT FOR PORTABLE 3D GAUSSIAN SPLATTING

Anonymous authors
Paper under double-blind review

ABSTRACT

Learned three-dimensional representations increasingly leave the system that trained them and are consumed by heterogeneous converters, importers, and renderers. For Gaussian splatting, serialized parameters do not alone determine a scene because consumer code completes their meaning along two coupled dimensions: *bounded admission*, covering schema, extent, and resource commitments, and *semantic portability*, covering rotation, coordinates, and radiance. The same accepted asset can therefore denote different scenes across consumers: a *semantic polyglot*. We formalize this problem with a typed cross-implementation observation model and root-adjudicated differential semantics, which promotes a disagreement only after invariant reduction, source localization, and stable replay. SPLATCONTRACT then enforces a format-explicit boundary: it proves bounded structure before size-proportional allocation, constructs a canonical Bounded Gaussian Intermediate Representation, and removes source-byte authority at handoff. Across five formats and six implementation families, adjudication isolates four causal mechanisms spanning both dimensions. SPLATCONTRACT blocks all four while admitting 99.975% of $8{,}000$ sealed legal assets and it recovers matched scene interpretations on calibrated real scenes.

1 INTRODUCTION

A scene optimized in one framework is rarely consumed there. It is serialized as PLY, packed as SPZ or SPLAT, pushed through an interchange tool, embedded in glTF, and rendered by a browser or a game engine. 3D Gaussian splatting (3DGS) makes portability unusually sharp: every primitive carries position, scale, orientation, opacity, and spherical-harmonic (SH) coefficients whose meaning is fixed only by activation functions, coordinate conventions, quantization schedules, and rasterization code (Kerbl et al., 2023; Yu et al., 2024; Lu et al., 2024; Huang et al., 2024). PLY is a generic property container rather than a Gaussian schema (Turk, 1994); compact splat formats add versioned, lossy compression; and the KHR extension places Gaussian semantics inside glTF's typed-buffer and scene model (The Khronos 3D Formats Working Group, 2021; Khronos Group, 2026).

Consider a captured scene exported as a GLB asset under that extension and opened in a browser engine. The asset declares linear radiance for its SH coefficients; the pinned importer retains those values but evaluates them as sRGB, so an asset that every stage accepts resolves to a visibly different scene (Figure 1). This importer defect exposes a broader portability gap. The baseline extension stores splats as point primitives, defines current kernel, color-space, projection and sorting semantics while remaining extensible, and lets consumers without extension support fall back to a sparse point cloud (Khronos Group, 2026). Legacy paths defer more: over a generic property container (Turk, 1994), deployed converters recover absent semantics from incidental evidence such as which property columns are present, and settle conflicting provenance by local convention. Across these paths, asset bytes encode different subsets of the semantics that consumers need, leaving downstream behavior dependent on local conventions.

A learned representation is therefore *not defined by its serialized parameters alone: the consumer completes its semantics.* Two applicable consumers may accept the same bytes and still disagree about which property is authoritative, how much storage a declaration justifies, whether a rotation is valid, or which color space an SH coefficient inhabits. An asset that multiple applicable consumers admit but resolve to different scenes is a *semantic polyglot*. These disagreements fall into two coupled classes.

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.

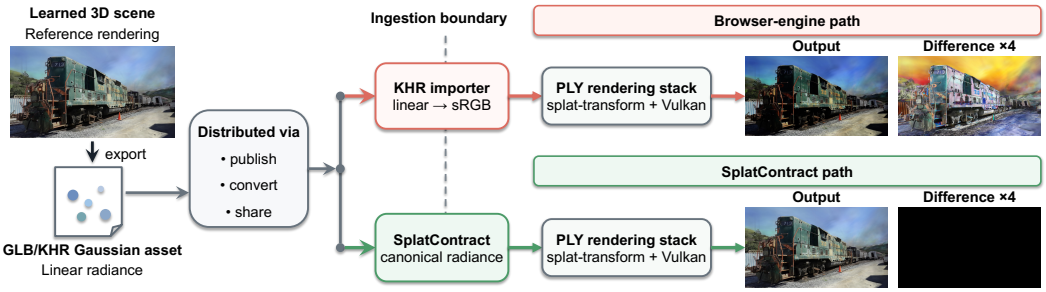


Figure 1: SPLATCONTRACT mediates a serialized Gaussian scene before a new consumer renders it. The panels show controlled renders of the Train scene from a GLB/KHR asset that declares linear radiance. The browser-engine KHR importer applies an sRGB interpretation to those values, while SPLATCONTRACT canonicalizes radiance at ingestion. Both paths are then materialized as PLY and rendered by the same `splat-transform` + Vulkan stack.

Failures of **bounded admission** let caller-controlled structural declarations authorize inconsistent state or unproven resource commitments; failures of **semantic portability** assign different numeric or rendered meanings to a tensor that every party has already accepted. A memory-safe parser addresses part of the first class, and canonicalizing the second is ill-defined until structure and cardinality are proven. A portable asset boundary must therefore mediate both, in that order, before any implementation-specific execution.

Existing boundaries stop short of this guarantee. Language-theoretic security and generated or verified parsers make recognition precise and memory safe (Sassaman et al., 2013; Bangert & Zeldovich, 2014; Ramananandro et al., 2019; Cai et al., 2025); differential studies of ZIP and URL parsers show that a conforming-looking byte string can receive context-dependent meanings (You et al., 2025; Ajmani et al., 2022); and cross-format polyglots exploit acceptance under several format languages at once (Magazinius et al., 2013). We study the complementary case, in which *one* format language is admitted by several consumers that complete it with different semantics. In deployment, a conformance validator establishes that an asset obeys a serialization specification, a strict parser decodes safely, and a converter shows that one source-to-target pair round-trips under its own conventions. Taken separately, these interfaces neither establish one canonical Gaussian tensor across a mixed-format pipeline nor prevent a later stage from re-consulting the original property order, declared count, rotation-repair rule or color convention.

Closing this gap is difficult for three reasons. First, no single normative Gaussian profile spans the ecosystem. Consequences rarely appear in the parse tree, surfacing instead as a size-proportional allocation, a silently repaired covariance, or an image produced only after quantization, activation, coordinate change and rasterization. Disagreement must therefore be stated against consumer-independent invariants observed beyond acceptance codes, with legitimate version and feature gaps separated from semantic failure rather than counted as it. Second, attribution is harder than detection: many candidates share one mechanism, one asset shows many non-causal differences, and deliberate permissiveness or an unsupported feature is indistinguishable from a finding until causality is established. Third, enforcement must be ordered, bounded and non-bypassable, since normalizing a tensor presupposes proven cardinality and exact extents, the structural proof must precede any size-proportional allocation, and canonicalization is useless if a later stage may re-read the source bytes, all while admitting the legal assets the ecosystem ships.

SPLATCONTRACT answers these with three layers. A *typed observation model* makes consumer-completed Gaussian semantics explicit across structure, numbers, resources and rendering, and declares applicability and materiality before outcomes are seen. *Root-adjudicated differential semantics* reduces matched disagreements to independently stated invariants and promotes a mechanism only after source localization and stable minimized replay. A *bounded deployment contract* performs format-explicit preflight, constructs a Bounded Gaussian Intermediate Representation (BGIR) under a declared profile and resource policy, and exposes only its canonical tensors through a trusted handoff. This yields a checkable postcondition: once two admitted encodings produce the same canonical object, a conforming consumer cannot recover a different meaning from the original bytes. The contract fixes asset semantics; rasterization and declared consumer configuration remain explicit downstream choices.

We evaluate five formats, six implementation families, 8,000 sealed legal assets, and fourteen public files spanning six scenes and four captured objects. The study isolates four causal mechanisms across bounded admission and semantic portability. Matched compatibility, visual, resource and component evidence shows that the contract prevents each mechanism while retaining near-complete legal and public compatibility; generator comparisons locate assurance in adjudication and the contract rather than one candidate source. Our contributions are as follows:

- **Consumer-completed Gaussian semantics.** An applicability-aware typed observation model that separates bounded admission from semantic portability, making portability a property of the representation *and* its consumers.
- **Root-adjudicated differential semantics.** A generator-agnostic causal procedure that reduces matched disagreements to invariants, localizes the responsible source decision, and requires stable minimized replay before a mechanism is counted.
- **A bounded canonical asset contract.** Format-explicit preflight proves schema and extent before allocation, BGIR fixes one declared tensor interpretation, and a trusted handoff removes source-byte authority, so both classes hold under one enforceable postcondition with proved noninterference.

2 Related Work

Gaussian representations. 3DGS introduced explicit anisotropic Gaussians for real-time radiance-field rendering (Kerbl et al., 2023). Mip-Splatting and Multi-Scale GS address sampling and scale (Yu et al., 2024; Yan et al., 2024), Scaffold-GS structures view-adaptive primitives (Lu et al., 2024), and 2DGS changes the primitive geometry to improve surface recovery (Huang et al., 2024). Such diversity is why interchange needs an explicit profile: a serialized splat cannot be assumed to share one activation, coordinate or covariance convention with every consumer. A parallel line studies the integrity of Gaussian scenes through training-cost poisoning, watermarking, steganography, view-dependent manipulation and downstream perception (Lu et al., 2025; Li et al., 2026; Chen et al., 2025; Jang et al., 2025; Guo et al., 2026; Hull et al., 2025a;b), addressing training inputs, model ownership or rendered outputs. Neither line specifies the asset boundary itself: we study what a serialized Gaussian asset *means* when importer, converter and renderer disagree.

Parser and representation consistency. Language-theoretic security treats overly permissive recognizers and parser compositions as an input-language problem (Sassaman et al., 2013), and generated or verified parsers make precise boundaries practical (Bangert & Zeldovich, 2014; Ramananandro et al., 2019; Cai et al., 2025). ZIP and URL parser studies show context-dependent meanings for conforming byte strings (You et al., 2025; Ajmani et al., 2022), and cross-format polyglots exploit acceptance under several format languages at once (Magazinius et al., 2013). We extend this line from syntax to typed Gaussian tensors, resource commitments and rendered semantics, and, rather than stopping at detection, pair observation with an enforceable canonical downstream contract.

Structured input generation. Grammar- and structure-aware systems generate or mutate inputs from grammars, fragments, inferred structure or semantic validity (Holler et al., 2012; Aschermann et al., 2019a; Wang et al., 2019; Padhye et al., 2019; Dutra et al., 2024; Godefroid et al., 2017; Blazytko et al., 2019). Other systems exploit input-to-state or comparison dependencies, mutate data-generation programs, or use domain-specific feedback (Aschermann et al., 2019b; Fioraldi et al., 2020; Bars et al., 2023; Peng et al., 2023). Multi-run benchmarking shows that target choice, censoring, variance and metric choice can reverse fuzzer comparisons (Klees et al., 2018; Metzman et al., 2021; Manès et al., 2021; Schlögel et al., 2024). We adopt that discipline and treat candidate generation as a modular, replaceable component.

3 Gaussian Assets as Semantic Polyglots

3.1 A semantic, not merely numeric, object

A Gaussian primitive is usually summarized as

$$g_i = (\boldsymbol{\mu}_i, \boldsymbol{s}_i, \boldsymbol{q}_i, \alpha_i, \boldsymbol{h}_{i,0}, \boldsymbol{h}_{i,>0}), \tag{1}$$

with center $\boldsymbol{\mu}_i \in \mathbb{R}^3$, scale $\boldsymbol{s}_i \in \mathbb{R}^3_{>0}$, rotation $\boldsymbol{q}_i$, opacity α_i, and DC/residual SH coefficients $\boldsymbol{h}_{i,0}, \boldsymbol{h}_{i,>0}$. Equation 1 hides every serialization choice that actually decides the scene: property identity and order, pre- or post-activation values, quaternion layout, coordinate handedness, SH basis and coefficient ordering, and color space.

Let b be an asset byte string and j an applicable implementation path. We record the typed observation

$$\Omega_j(b) = (a_j, \mathcal{S}_j, \mathcal{N}_j, \mathcal{R}_j, \mathcal{V}_j), \tag{2}$$

with terminal acceptance a_j and structural, numeric, resource and (when applicable) rendered components $\mathcal{S}_j, \mathcal{N}_j, \mathcal{R}_j, \mathcal{V}_j$, detailed next. A *semantic differential* exists when applicable paths accept b (or at least one executes far enough to cause an effect) and then disagree in a consequential component of Ω, after documented format and version support has been accounted for.

3.2 Four observation dimensions serving two objectives

We keep four components separate because a return code collapses semantically distinct behavior: 1) **Structure** records counts, property identity and order, chunk coverage, accessors, archive members and byte extents; 2) **Numbers** records activation domains, finiteness, quaternion convention, coordinates, units, SH basis and color transformation; 3) **Resources** records justified bytes, expansion, point count, peak memory and elapsed time; and 4) **Rendering** records matched-camera images, invalid values, foreground occupancy, nondeterminism and perceptual materiality. Visual and resource materiality criteria are frozen before aggregation; Appendix C.3 records their thresholds and measurement boundaries.

These four dimensions serve two objectives. **Bounded admission** covers structure and resources: caller-controlled declarations must not authorize inconsistent state or size-proportional work before proof. **Semantic portability** covers numbers and rendering: every admitted tensor must carry exactly one declared interpretation. They meet at the same boundary because canonicalization needs proven cardinality and extents, while structural proof alone cannot decide rotation, coordinate or radiance semantics. SplatContract therefore orders bounded admission before canonical interpretation and exposes both through one typed handoff.

3.3 Setting and objective

The study covers local asset ingestion, conversion and rendering on pinned, applicable implementation paths. Asset bytes and declared metadata are caller controlled; implementation binaries, policy and evaluation cameras are fixed. Format applicability follows each implementation's documented support.

The objective is a partial function C_π parameterized by a declared profile and resource policy π:

$$C_\pi(b) = \begin{cases} G, & \text{if structure and semantics admit one bounded Gaussian object,} \\ \bot, & \text{otherwise,} \end{cases} \tag{3}$$

where the downstream path may consume G but may not re-read b as an independent authority. This converts parser acceptance into a typed semantic postcondition. Each admitted encoding must name exactly one typed Gaussian state under π, and equivalent encodings must map to the same state; rasterization and explicitly declared consumer configuration stay outside the contract.

4 SplatContract

Figure 2 summarizes the two paths. Analysis draws candidates from structured mutation, byte mutation, corpus replay or targeted probes and turns matched observations into causal mechanisms. Deployment is independent of the generator and admits an asset only after structural proof and semantic canonicalization.

Validation, memory-safe parsing and pairwise round-tripping each leave the downstream interpretation open. The unit of assurance in SplatContract is therefore the *composition* of (i) a bounded structural proof, (ii) a canonical typed object with explicit semantics, and (iii) an authority-revoking handoff. Preflight alone is retained as a simpler boundary baseline, so that the value of (ii) and (iii) is measured rather than assumed.

4.1 Root-adjudicated differential semantics

The analysis path uses four gates. *(1) Observation*: applicable paths execute matched assets and emit the typed records of Equation 2. *(2) Invariant reduction*: provisional signals are grouped by an

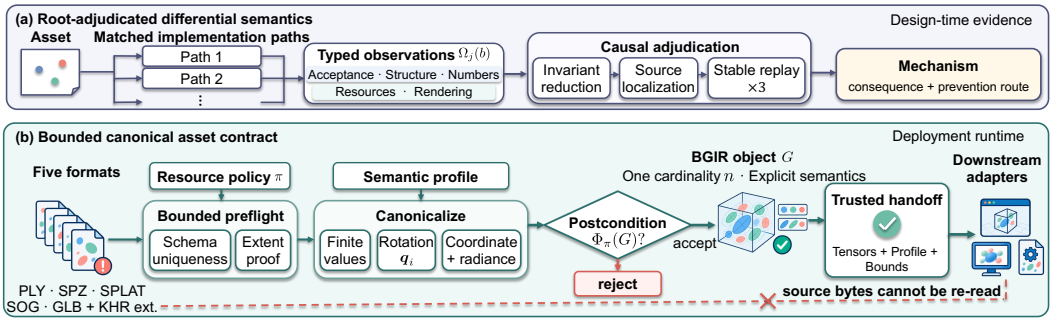


Figure 2: Method overview. **(a)** The same candidate asset traverses matched implementation paths, producing typed observations of terminal acceptance, structure, numbers, resources and rendering. A mechanism is retained only after invariant reduction, source localization and three stable replays, which also establish its prevention route. **(b)** Resource policy π controls the two-module bounded preflight, while the declared semantic profile controls the three-module canonicalization. The separate postcondition $\Phi_\pi(G)$ yields atomic rejection or an admitted BGIR object G; only its trusted handoff reaches downstream adapters, and original source bytes have no bypass.

independently stated invariant rather than by stack trace, return code or mutation operator. *(3) Source localization*: pinned source, or a source-faithful harness, must isolate the decision that causally links the invariant to the observation. *(4) Stable replay*: a minimized private case must replay identically three times, and its impact is summarized without releasing reconstructive content.

This changes the counting unit. Many mutated files can exercise one mechanism, and one asset can expose several non-causal differences. A mechanism is retained only when the source decision, the violated invariant, the stable observation and the prevention route all agree. AddressSanitizer supplies diagnostic evidence for memory violations (Serebryany et al., 2012), not a semantic oracle.

4.2 Format-explicit bounded preflight

Preflight recognizes the declared format and performs only bounded work, without invoking a full external decoder. For a fixed-record format it requires $|b| - h = nr$ with header size h, declared record count n and record width r, all parsed with checked arithmetic. For typed buffers, every accessor interval must lie inside its buffer view, each view inside its buffer and chunk, and required semantic cardinalities must agree. For compressed or member-based formats, preflight checks member uniqueness, declared expanded size and policy limits before decompression.

Two components establish the structural postcondition. **Schema Uniqueness** rejects duplicate semantic names, missing fields, partial SH bands, conflicting accessors and unsupported required modes. **Extent Proof** rejects truncated or trailing records, inconsistent counts, uncovered chunks and policy overflow. No caller-controlled count reaches a size-proportional allocation before both exact extent and policy proof.

4.3 Bounded Gaussian intermediate representation

After structural proof, a format-specific decoder constructs BGIR with one shared cardinality n:

$$M \in \mathbb{R}^{n\times 3}, \quad S \in \mathbb{R}^{n\times 3}_{>0}, \quad Q \in \mathbb{R}^{n\times 4}, \quad A \in [0,1]^n, \quad H_0 \in \mathbb{R}^{n\times 3}, \quad H_r \in \mathbb{R}^{n\times k}. \tag{4}$$

All arrays are finite and shape exact. **Finite Numeric Domain** enforces finiteness and post-activation ranges. **Rotation Validity** requires $\|q_i\|_2^2 \geq 10^{-12}$, rejects non-finite or degenerate inputs, and normalizes accepted quaternions to one declared $xyzw$ convention; comparisons are sign invariant because q and $-q$ denote the same rotation. The **Coordinate and Radiance Canonicalizer** converts declared axes, scale and opacity activation, SH basis and ordering, and color-space semantics into the explicit profile. Lossy-format tolerances $\tau_f = \max(\tau_{\text{mechanism}},\ 1.25\, e_f^{\text{dev}} + 10^{-12})$ are set from a separate development split and frozen before the sealed legal test is inspected. Scale is required to be finite while memory and point count remain policy bounded, so the contract never depends on a unit ceiling that the supported legacy formats do not define.

4.4 SEMANTIC POSTCONDITION AND TRUSTED HANDOFF

Let $\Phi_\pi(G)$ be the conjunction of the contract invariants: one cardinality and exact tensor shapes; finite post-activation values; policy-bounded extents; nondegenerate unit rotations in one declared layout; and explicit coordinate, SH and color-space semantics. Equation 3 returns G only when $\Phi_\pi(G)$ holds, and successful admission carries the profile, the tensors and their proven bounds.

A downstream adapter A is *contract conforming* when, for fixed explicit configuration x, its asset-dependent observation factors through the handoff:

$$\Omega_A(b; x) = F_A(C_\pi(b); x) \qquad \text{for every admitted } b. \tag{5}$$

Proposition 1 (Encoding noninterference)**.** *Let A be deterministic and contract conforming with fixed configuration x. If $C_\pi(b) = C_\pi(b') = G \neq \bot$, then $\Omega_A(b; x) = \Omega_A(b'; x)$, and every asset-dependent allocation that A performs through the handoff is bounded by the certified cardinality and tensor extents in π.*

Proof. By Equation 5, $\Omega_A(b; x) = F_A(G; x) = \Omega_A(b'; x)$. The bound follows from $\Phi_\pi(G)$, which certifies a single cardinality and exact tensor shapes, and from the fact that the handoff exposes no other asset-derived size. □

Proposition 1 is the downstream invariant we want: after canonicalization, encoding-specific names, ordering, counts, repairs and color conventions cannot regain authority, while renderer choice stays explicit in x. Invalid rotations and unproven extents fail closed and only declared convertible conventions are normalized, so canonicalization never silently repairs an input with no unique bounded meaning.

4.5 IMPLEMENTATION

The reference implementation separates preflight, canonical decoders and downstream adapters. The PLY path provides materialized and field-wise disk-backed stores with the same typed schema and content fingerprint; exercised external paths use pinned public implementations or a source-faithful game-engine importer harness. Appendix C.3 records the runtime and measurement configuration.

5 EXPERIMENTS

We test whether typed observations expose failures beyond return codes, root adjudication isolates causal mechanisms across candidate sources, and the unified contract prevents those mechanisms while retaining compatibility at measured cost. Matched visual and resource controls establish causal consequence.

5.1 SETUP

Data. A *generated legal corpus* covers the registered semantic profile across all five formats, with a 2,000-asset development split for policy and tolerance selection and a sealed 8,000-asset test split. A *public real-asset corpus* contains fourteen source-native files from six scenes and four independently captured objects in PLY, SPZ and SPLAT, including a CC BY 4.0 collection by Marcel Padilla. Same-content SOG and GLB/KHR conversions exercise route applicability and remain correlated derivatives. Generated adversarial cases target structure, numbers, resources and rendering. Calibrated causal rendering uses Train and Truck from Tanks and Temples (Knapitsch et al., 2017) and Dr Johnson and Playroom from Deep Blending (Hedman et al., 2018) under fixed views and repeated execution. Appendix C specifies file units, splits, correlations, cameras and repeats.

Evaluated systems. The surface contains six families: a native compressed-splat library, a multi-format converter and renderer, a standalone WebGL viewer, a game-engine importer harness, a browser rendering engine, and a glTF conformance validator. Preflight and BGIR are proposed components. Applicability is declared before outcomes are observed, and an unsupported feature is never reclassified as an unsafe rejection.

Table 1: Causal mechanisms across the contract's two objectives.

Mechanism	Objective	Isolated consequence	Boundary response
Duplicate semantic property	Bounded admission	Unique-name storage and declaration-order offsets disagree; release builds accept while diagnostic builds expose an out-of-bounds access.	Schema Uniqueness
Unproven record cardinality	Bounded admission	A declared count reaches allocation before the payload proves that many records, exhausting a bounded process.	Extent Proof
Degenerate rotation repair	Semantic portability	One path repairs a zero rotation to identity; another normalizes it to a different finite orientation.	Rotation Validity
Radiance-space reinterpretation	Semantic portability	Retained SH coefficients traverse a color interpretation inconsistent with their declared linear semantics.	Coordinate and Radiance Canonicalizer

Baselines. The comparison separates candidate discovery from boundary enforcement. (1) **Matched byte mutation** uses the same seeds and active budget as structured mutation. (2) **Parser-only differential replay** applies acceptance and deep-decode observations to every structured candidate without a semantic materiality oracle. (3) **Unmediated implementation paths** consume the original asset directly. (4) **Preflight-only admission** applies the registered structural and numeric checks without canonical tensors or the trusted handoff. (5) **Ecosystem validation and conversion** exercise a conformance validator and ordinary converters or importers on applicable legal formats. The last three distinguish format acceptance, bounded admission and the full representation postcondition and are not pooled into one ranking.

Metrics and statistics. We evaluate discovery, compatibility, visual consequence and deployment cost. *Discovery* uses deep-decode fraction, source-adjudicated mechanism yield and censored time to first mechanism, summarized with the product-limit estimator (Kaplan & Meier, 1958). *Compatibility* reports accepted over applicable assets with Wilson intervals (Wilson, 1927). *Visual consequence* uses SSIM, LPIPS, foreground-fraction change and non-finite values under matched asset, camera and repeat tuples. *Deployment* uses wall time and process peak RSS under matched execution. Appendix C specifies pairing, aggregation, clustering, repeat and execution-order protocols.

5.2 Adjudicated cross-implementation mechanisms

Table 1 lists the four source-adjudicated mechanisms and the objective each one tests. Every row names one invariant, one isolated interpretation decision and one exact boundary response.

The two bounded-admission mechanisms show two forms of caller-controlled authority: a duplicate property name is read once as a unique field identity and again as an offset schedule, and a declared record count is read as an allocation size before exact payload proof. The two semantic-portability mechanisms show why syntactic validity is insufficient: both rotation paths produce finite images yet select different orientations, and typed KHR buffers can validate and render while the radiance convention changes downstream.

Adjudication is what makes these numbers interpretable. The analysis starts from 582 provisional signatures, and exactly four distinct mechanisms survive invariant reduction, source localization, stable replay and causal impact. Counting crashes, rejection strings or mutation clusters would overstate the result by two orders of magnitude.

5.3 Prevention and compatibility

With the complete boundary enabled, all four mechanisms are rejected or normalized before the affected implementation-specific interpretation. Structural preflight accepts $8{,}000/8{,}000$ sealed legal assets, and full BGIR accepts $7{,}998/8{,}000$ (99.975%, Wilson 95% interval $[99.909\%, 99.993\%]$), with both misses being SH-residual tolerance cases in SOG.

Preflight-only admission establishes bounded decoding but constructs neither the canonical tensor object nor the trusted handoff. Both the ecosystem conformance validator and the browser engine accept the applicable generated GLB/KHR assets, while the controlled study below shows that

Table 2: Public-asset compatibility by applicable format.

Asset stratum	Units	Preflight	BGIR	External routes
Native PLY	4 scenes / 8 files	8/8	8/8	32/32
Native SPZ	2 scenes / 2 files	2/2	n/a	4/6
Native SPLAT	4 objects / 4 files	4/4	4/4	8/8
Derived SOG†	correlated / 4 files	4/4	n/a	4/4
Derived GLB†	correlated / 4 files	4/4	4/4	8/8

†Converted from the four SPLAT objects; correlated route-applicability probes rather than independent native samples. External counts pool applicable file-by-route cells.

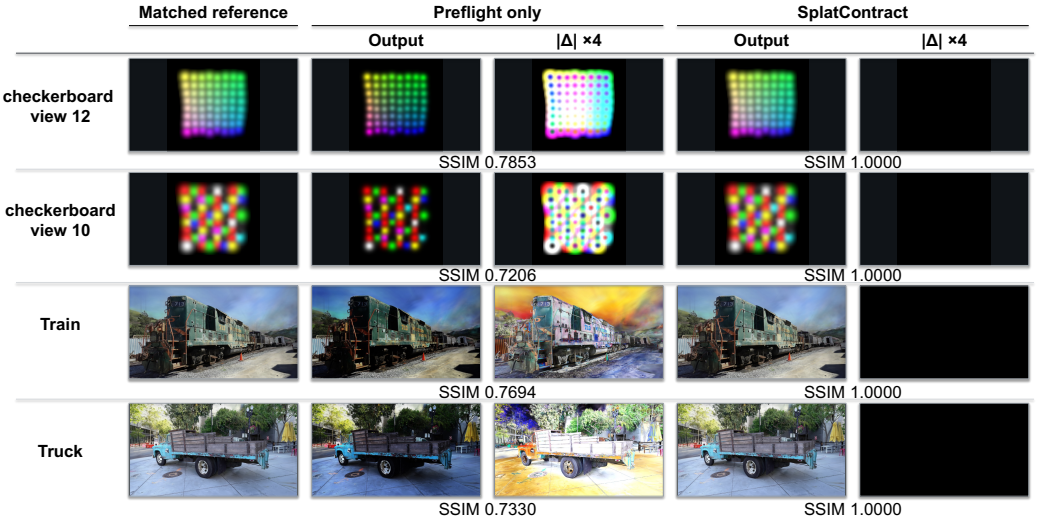


Figure 3: Matched radiance comparison across generated checkerboards and public scenes, in dataset order. Each row fixes the asset, view, renderer and post-processing. The preflight-only baseline and SPLATCONTRACT each pair their output with $|\Delta| = \mathrm{clip}(4|I - I_{\mathrm{ref}}|)$ under the same channel range; darker residuals are lower, while higher SSIM is better.

declared-linear radiance can still be reinterpreted after acceptance. SPLATCONTRACT therefore supplies a downstream postcondition beyond format validation.

Table 2 separates source-native evidence, correlated conversions, applicability and acceptance. Preflight accepts every native file, and full BGIR accepts every applicable PLY and SPLAT file; the source-native SPLAT objects and their correlated derivatives complete all applicable registered routes. The standalone viewer's SPZ result records its pinned support boundary.

5.4 CAUSAL VISUAL EFFECTS

Radiance (KHR). All renders in the frozen KHR matrix complete repeat-stably, and $120/240$ matched pairs are material. Scene-level SSIM and LPIPS estimates support the expected view-dependent SH-radiance effect; Appendix G reports the full distributions and intervals.

Rotation. All 12 matched pairs are material and deterministic under a control that changes only the interpretation of the invalid encoded rotation. Non-finite and negative-scale controls remain non-material, separating this mechanism from generic malformed numeric input; Appendix G gives the complete perceptual and foreground-shift measurements.

Calibrated real scenes. All calibrated renders complete deterministically. Mediation brings every main and sensitivity pair within the frozen raster-recovery bound, including 18 pixel-exact pairs. Train and Truck supply all 18 material public-scene pairs and the examples in Figure 3; the complete matrix, including the indoor scenes, appears in Table 11.

Figure 3 combines two frozen generated checkerboards with the two eligible public scenes. The preflight-only baseline has passed the same format-explicit structural checks but lacks radiance canonicalization and the trusted handoff. Matched outputs and residuals show the resulting appearance change and its recovery under identical downstream rendering.

Table 4: Component removal and claim-relevant effect.

Removed component	Newly forwarded	Downstream Accepts	Lost prevention or evidence
Extent Proof	56,541	22,914	cardinality mechanism
Schema Uniqueness	40,434	32,641	duplicate-property mechanism
Finite Numeric Domain	34,161	32,161	finite-domain postcondition
Rotation Validity	5,739	5,739	degenerate-rotation mechanism
Coordinate and Radiance Canonicalizer	n/a	n/a	radiance mechanism
Causal Materiality Oracle	240 pairs	n/a	120 material labels

5.5 Resource behavior and admission cost

Preflight rejects all 22 resource-policy cases before an external decoder runs, with stable process RSS under large declared expansions. A separate cardinality case produces a tens-of-GiB commitment under both instrumented-process and whole-step estimators before rejection. Appendix G reports the input, expansion and estimator strata separately.

Table 3 shows that full materialization dominates the offline cost, reaching 9.24 s and 5.43 GiB at p95. The full guard adds a median $0.206\times$ over existing canonicalization, while sequential peak RSS remains canonicalization-dominated at ratio 1.0. A field-wise disk-backed construction reproduces the materialized fingerprint in $40/40$ blocks with p95 peak RSS 76.2% lower (Appendix I, Figure 4).

Table 3: Registered paired efficiency on the public corpus.

	Wall time (s)		Peak RSS (KiB)	
Operation	Median	p95	Median	p95
Structural preflight	0.033	0.066	15,304	15,312
Full semantic guard	0.817	1.996	526,984	1,028,172
External decode	2.595	3.899	623,930	670,149
BGIR construction	5.049	9.241	3,095,434	5,430,744

5.6 Component ablations

Table 4 removes each component under fixed routing. Four boundary components each carry one observed prevention route, Finite Numeric Domain maintains the representation postcondition, and the materiality oracle carries causal evidence. No removal adds a false rejection on the development routing set. The removals forward many distinct candidates, showing profile-wide effects beyond the four minimized witnesses. Policy and lossy tolerances were frozen before the sealed test, and a later source-native SPLAT extension completed every applicable terminal under the same profile. Together, these observations support invariant-level enforcement after the design was frozen.

5.7 Candidate-generation comparison

Matched structured and byte mutation have comparable deep-decode reach: the paired interval spans zero and the effect size is near zero, so the registered broad structured-reach advantage is not supported. Structured mutation isolates two adjudicated mechanisms and byte mutation one, while parser-only replay detects both structured mechanisms without a semantic oracle. Assurance therefore comes from generator-agnostic adjudication and the contract; Appendix F reports the full reach, root-yield and time-to-mechanism results.

6 Conclusion

A portable learned representation is defined jointly by its serialized parameters and consumer semantics; SplatContract makes that boundary explicit for Gaussian assets, adjudicating cross-implementation disagreements at the mechanism level and coupling bounded admission to one canonical tensor interpretation behind a provably noninterfering handoff. It prevents all four source-adjudicated mechanisms while admitting $7{,}998/8{,}000$ sealed legal assets canonically and recovering the calibrated semantic reference to within $1/255$.

Limitations. The empirical results cover the pinned formats, implementation paths, revisions and assets studied here. Future work can extend the semantic profile to new learned-scene encodings and examine canonical handoffs across larger creation-to-rendering pipelines.

AI USE STATEMENT

In this work, we have not used generative AI tools for tasks with required disclosure: they were not used to propose or refine hypotheses, develop theoretical models or conceptual frameworks, formulate mathematical claims or assist in writing proofs, design or provide feedback on methodology or experiments, implement methods, generate or reformat datasets, or interpret results. Assistance with translation and with qualitative and thematic data analysis is not applicable to this work. Additionally, we used generative AI tools for tasks with recommended disclosure: searching for and identifying relevant literature, and editing the manuscript to improve readability. We have reviewed all AI-assisted work: every AI-suggested reference was checked against primary publication records before citation, and all AI-edited text was verified by the authors against the underlying artifacts so that no quantitative value, experimental observation or rendered output depends on generative AI. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work studies untrusted local asset processing and can create disclosure risk if minimized inputs or reconstructive instructions are released prematurely. The paper therefore reports mechanism taxonomies, aggregate resource and render evidence, prevention routes, and non-weaponized executable mechanism laws, while excluding disclosure-sensitive bytes, hashes, offsets and operational reconstruction steps. No remote service or third-party system was tested. Experiments used local builds of public source, generated assets and license-audited public scenes under bounded processes, and involve no human subjects or personal data.

REPRODUCIBILITY STATEMENT

Appendix C documents formats, applicability, data units, implementation modes, metrics and statistical procedures, and Appendices D through I provide the complete paper-facing mechanism routes, conformance, discovery, rendering, resource, ablation, scaling and efficiency evidence. The Supplementary Materials include our source code together with a detailed README covering environment setup, configuration files and the commands needed to reproduce the reported results.

REFERENCES

Dashmeet Kaur Ajmani, Igibek Koishybayev, and Alexandros Kapravelos. yoU aRe a Liar://A Unified Framework for Cross-Testing URL Parsers. In *IEEE Security and Privacy Workshops*, 2022.

Cornelius Aschermann, Tommaso Frassetto, Thorsten Holz, Patrick Jauernig, Ahmad-Reza Sadeghi, and Daniel Teuchert. NAUTILUS: Fishing for deep bugs with grammars. In *Network and Distributed System Security Symposium*, 2019a.

Cornelius Aschermann, Sergej Schumilo, Tim Blazytko, Robert Gawlik, and Thorsten Holz. REDQUEEN: Fuzzing with input-to-state correspondence. In *Network and Distributed System Security Symposium*, 2019b.

Julian Bangert and Nickolai Zeldovich. Nail: A practical tool for parsing and generating data formats. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2014.

Nils Bars, Moritz Schlögel, Tobias Scharnowski, Nico Schiller, and Thorsten Holz. Fuzztruction: Using fault injection-based fuzzing to leverage implicit domain knowledge. In *32nd USENIX Security Symposium*, 2023.

Tim Blazytko, Cornelius Aschermann, Moritz Schlögel, Ali Abbasi, Sergej Schumilo, Simon Wörner, and Thorsten Holz. GRIMOIRE: Synthesizing structure while fuzzing. In *28th USENIX Security Symposium*, 2019.

Yi Cai, Pratap Singh, Zhengyao Lin, Jay Bosamiya, Joshua Gancher, Milijana Surbatovich, and Bryan Parno. Vest: Verified, secure, High-Performance parsing and serialization for Rust. In *34th USENIX Security Symposium*, 2025.

Zixuan Chen, Guangcong Wang, Jiahao Zhu, Jianhuang Lai, and Xiaohua Xie. GuardSplat: Efficient and robust watermarking for 3D gaussian splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.

Rafael Dutra, Rahul Gopinath, and Andreas Zeller. FormatFuzzer: Effective fuzzing of binary file formats. *ACM Transactions on Software Engineering and Methodology*, 2024.

Andrea Fioraldi, Daniele Cono D'Elia, and Emilio Coppa. WEIZZ: Automatic grey-box fuzzing for structured binary formats. In *ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2020.

Patrice Godefroid, Hila Peleg, and Rishabh Singh. Learn&Fuzz: Machine learning for input fuzzing. In *IEEE/ACM International Conference on Automated Software Engineering*, 2017.

Yijia Guo, Wenkai Huang, Yang Li, Gaolei Li, Hang Zhang, Liwen Hu, Jianhua Li, Tiejun Huang, and Lei Ma. Splats in splats: Robust and effective 3D steganography towards gaussian splatting. In *AAAI Conference on Artificial Intelligence*, 2026.

Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics*, 2018.

Christian Holler, Kim Herzig, and Andreas Zeller. Fuzzing with code fragments. In *21st USENIX Security Symposium*, 2012.

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D gaussian splatting for geometrically accurate radiance fields. In *SIGGRAPH 2024 Conference Papers*, 2024.

Matthew Hull, Haoyang Yang, Pratham Mehta, Mansi Phute, Aeree Cho, Haoran Wang, Matthew Lau, Wenke Lee, Willian T. Lunardi, Martin Andreoni, and Polo Chau. 3D gaussian splat vulnerabilities. *arXiv preprint arXiv:2506.00280*, 2025a.

Matthew Hull, Haoyang Yang, Pratham Mehta, Mansi Phute, Aeree Cho, Haorang Wang, Matthew Lau, Wenke Lee, Wilian Lunardi, Martin Andreoni, and Duen Horng Chau. ComplicitSplat: Downstream models are vulnerable to blackbox attacks by 3D gaussian splat camouflages. *arXiv preprint arXiv:2508.11854*, 2025b.

Youngdong Jang, Hyunje Park, Feng Yang, Heeju Ko, Euijin Choo, and Sangpil Kim. 3D-GSW: 3D gaussian splatting for robust watermarking. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.

Edward L. Kaplan and Paul Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 1958.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 2023.

Khronos Group. KHR_gaussian_splatting extension specification, release candidate. *Online specification:* `https://github.com/KhronosGroup/glTF/tree/2b29723d025a995971726f2989697cdc49b1222a/extensions/2.0/Khronos/KHR_gaussian_splatting`, 2026.

George Klees, Andrew Ruef, Benji Cooper, Shiyi Wei, and Michael Hicks. Evaluating fuzz testing. In *ACM SIGSAC Conference on Computer and Communications Security*, 2018.

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

Yanping Li, Zhening Liu, Zijian Li, Zehong Lin, and Jun Zhang. RemedyGS: Defend 3D gaussian splatting against computation cost attacks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026.

Jiahao Lu, Yifan Zhang, Qiuhong Shen, Xinchao Wang, and Shuicheng Yan. Poison-splat: Computation cost attack on 3D gaussian splatting. In *International Conference on Learning Representations*, 2025.

Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-GS: Structured 3D gaussians for view-adaptive rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Jonas Magazinius, Billy K. Rios, and Andrei Sabelfeld. Polyglots: Crossing origins by crossing formats. In *ACM SIGSAC Conference on Computer and Communications Security*, 2013.

Valentin J. M. Manès, HyungSeok Han, Choongwoo Han, Sang Kil Cha, Manuel Egele, Edward J. Schwartz, and Maverick Woo. The art, science, and engineering of fuzzing: A survey. *IEEE Transactions on Software Engineering*, 2021.

Jonathan Metzman, László Szekeres, Laurent Simon, Read Sprabery, and Abhishek Arya. FuzzBench: An open fuzzer benchmarking platform and service. In *ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2021.

Rohan Padhye, Caroline Lemieux, Koushik Sen, Mike Papadakis, and Yves Le Traon. Semantic fuzzing with Zest. In *ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2019.

Hui Peng, Zhihao Yao, Ardalan Amiri Sani, Dave (Jing) Tian, and Mathias Payer. GLeeFuzz: Fuzzing WebGL through error message guided mutation. In *32nd USENIX Security Symposium*, 2023.

Tahina Ramananandro, Antoine Delignat-Lavaud, Cédric Fournet, Nikhil Swamy, Tej Chajed, Nadim Kobeissi, and Jonathan Protzenko. EverParse: Verified secure zero-copy parsers for authenticated message formats. In *28th USENIX Security Symposium*, 2019.

Len Sassaman, Meredith L. Patterson, Sergey Bratus, and Michael E. Locasto. Security applications of formal language theory. *IEEE Systems Journal*, 2013.

Moritz Schlögel, Nils Bars, Nico Schiller, Lukas Bernhard, Tobias Scharnowski, Addison Crump, Arash Ale Ebrahim, Nicolai Bissantz, Marius Muench, and Thorsten Holz. SoK: Prudent evaluation practices for fuzzing. In *IEEE Symposium on Security and Privacy*, 2024.

Konstantin Serebryany, Derek Bruening, Alexander Potapenko, and Dmitriy Vyukov. AddressSanitizer: A fast address sanity checker. In *USENIX Annual Technical Conference*, 2012.

The Khronos 3D Formats Working Group. glTF 2.0 specification, version 2.0.1. *Online specification:* `https://registry.khronos.org/glTF/specs/2.0/glTF-2.0.html`, 2021.

Greg Turk. The PLY polygon file format. *Stanford Graphics Laboratory specification*, 1994.

Junjie Wang, Bihuan Chen, Lei Wei, and Yang Liu. Superion: Grammar-aware greybox fuzzing. In *IEEE/ACM International Conference on Software Engineering*, 2019.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004.

Edwin B. Wilson. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 1927.

Zhiwen Yan, Weng Fei Low, Yu Chen, and Gim Hee Lee. Multi-Scale 3D gaussian splatting for anti-aliased rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Yufan You, Jianjun Chen, Qi Wang, and Haixin Duan. My ZIP isn't your ZIP: Identifying and exploiting semantic gaps between ZIP parsers. In *34th USENIX Security Symposium*, 2025.

Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-Splatting: Alias-free 3D gaussian splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.

A Appendix overview

This appendix is organized as follows.

- Section B: extended discussion of the contract's design rationale and transfer.
- Section C: evaluation surface, applicability, datasets, runtime modes and measurement boundaries.
- Section D: causal adjudication criteria and prevention routes.
- Section E: canonical tolerances and complete legal and public compatibility.
- Section F: full candidate-generation comparison and censored discovery results.
- Section G: rendering matrices, controls, resource strata and measurement scopes.
- Section H: anonymous artifact and executable non-weaponized mechanism laws.
- Section I: component, sensitivity and paired efficiency evidence.

B Extended discussion

Serialization is an authority boundary. The four mechanisms share one pattern even though their outcomes split into two classes. A caller-controlled name or count becomes authority over structure and resource commitment; a tuple or a convention becomes authority over numeric and rendered interpretation. In both cases one layer supplies meaning that the serializer left implicit and another layer consumes it under a different assumption. The proposed boundary removes this ambient authority: downstream code receives arrays with one proven cardinality, one rotation convention and one coordinate and radiance profile.

A representation-level contract. The technical contribution composes schema proof, checked extents, typed tensors, semantic canonicalization and an authority-revoking handoff into one executable consumer interface. Validation, parsing and conversion terminate at narrower boundaries, establishing format conformance, memory-safe native state, or one implementation's chosen interpretation. SplatContract instead exposes Φ_π across five formats and enforces bounded admission together with semantic portability before implementation-specific execution, so that equivalent admitted encodings are noninterfering for a conforming consumer (Proposition 1), while root-adjudicated experiments validate the structural, resource, numeric and rendered consequences that motivate the contract.

Canonicalization is intentionally narrower than the source formats. BGIR retains declared, convertible semantics and rejects cases with no unique bounded meaning. This is why exact extents and zero rotations fail closed while coordinate and radiance conventions are normalized: a more permissive representation would reintroduce the same disagreement one stage downstream. For the same reason, the profile leaves arbitrary glTF node-affine flattening to an explicit future extension rather than assigning it an implicit meaning.

Discovery and prevention should stay separate. Structured operators expose two useful mechanisms, the broad deep-reach hypothesis is not supported, and parser-only replay reaches both formal roots. Candidate generators are therefore interchangeable sources of observations, and the transferable method is causal adjudication plus the canonical asset contract.

Invariant-level rather than witness-level enforcement. The boundary components were motivated by the observed failure classes but check profile predicates rather than retained byte patterns. The large forwarding counts under component removal, the untouched sealed legal split, and the later source-native SPLAT extension test this distinction from complementary directions, supporting enforcement of the declared postcondition across the evaluated profile.

Transfer beyond Gaussian assets. A learned-format family is a candidate for the same treatment when it admits (i) a finite semantic profile for its model state, (ii) bounded structural proof before materialization, (iii) a canonical typed object, and (iv) adapters that consume only that object. Neural fields, learned meshes and compressed feature assets plausibly satisfy these conditions, each with its own profile and conformance study.

C Evaluation surface and protocol

C.1 Formats and applicability

Table 5 lists only the exercised paths. "Production" means unmodified public source at a pinned revision, "harness" a source-faithful extracted importer path, "validator" a conformance checker, and "defense" a proposed boundary component. Additional formats that a production repository may support outside the pinned adapter are not inferred.

Table 5: Format applicability in the evaluated surface.

Path	PLY	SPZ	SPLAT	SOG	GLB/KHR
Native compressed library	production	production	n/a	n/a	n/a
Multi-format converter	production	production	production	production	production
Standalone WebGL viewer	production	production	production	n/a	n/a
Game-engine importer	harness	n/a	n/a	n/a	n/a
Browser rendering engine	n/a	n/a	n/a	n/a	production
glTF conformance validator	n/a	n/a	n/a	n/a	validator
Structural preflight	defense	defense	defense	defense	defense
BGIR	defense	defense	defense	defense	defense

C.2 Data units and splits

Development assets use $2{,}000$ independently seeded scene specifications and select tolerance and policy parameters. The sealed test uses $8{,}000$ distinct seeds contributing $1{,}600$ assets per format. Generated candidates use five seeds per format and method. The source-native public corpus contains two SPZ scenes, four PLY scenes at two optimization checkpoints each, and four independently captured degree-0 SPLAT objects from the attributed CC BY 4.0 collection; file-level acceptance is reported, while checkpoints keep scene-level independence at six and the SPLAT set adds four object-level units. Four same-content conversions per target format exercise the SOG and GLB/KHR routes and remain correlated with the SPLAT objects. The sealed corpus provides controlled profile coverage across five formats, and the public corpus provides source-native production-asset evidence for PLY, SPZ and SPLAT, with BGIR evaluated on its eight PLY and four SPLAT files; these evidentiary roles are kept separate throughout.

The KHR visual matrix contains five generated scenes, 16 cameras, two semantic roles and three repeats. The rotation matrix contains one asset family, four cameras, two semantic roles and three repeats. Repeats diagnose nondeterminism. The resource matrix covers exact-count, compressed-expansion, accessor and policy strata. The registered efficiency experiment uses ten public files, five repeats and cyclic method order, with BGIR applying to the eight PLY files; the later scaling study keeps those eight checkpoints and repeats four paired implementations five times each.

C.3 Runtime and measurement boundaries

The reference implementation uses Python 3.10+ with NumPy. JavaScript adapters use Node.js 22.22.3, exact rendering uses a pinned software Vulkan backend, and LPIPS uses the official AlexNet v0.1 path on CPU. Native code uses release, assertion-enabled and memory-instrumented builds as required by each path. Every unit records input identity, adapter mode, process outcome, resource measurements and output identity; a resumable driver continues from atomic completion records. Resource numbers retain their estimator: preflight process peak, external-decoder process peak, instrumented-process peak and whole-step accounting remain separate. Timing measures completed invocations on one host from execution start to terminal outcome; setup and pre-execution diagnostics have separate records. A visual pair is material when SSIM is below 0.95, either output contains non-finite values, or foreground fraction changes by at least 0.10; SSIM and LPIPS are complementary perceptual summaries (Wang et al., 2004; Zhang et al., 2018). Resource evidence is material when a bounded process is exhausted or a matched resource observation at least doubles. These criteria are frozen before aggregation.

D CAUSAL MECHANISMS AND PREVENTION ROUTES

D.1 ADJUDICATION CRITERIA

A retained mechanism satisfies seven conditions: a stable differential or impact; an invariant stated independently of the generator; a causal source decision at a pinned revision; three identical private replays; a minimized private input; one final taxonomy entry; and an exact prevention route. All provisional signatures receive one terminal disposition, and mutation operator, stack trace and return-code pattern are not accepted as counting units.

Table 6: Disclosure-safe mechanism-to-boundary audit.

Mechanism	Violated invariant	Source decision isolated	Exact route
Duplicate semantic property	Each required semantic property has one identity, and allocation layout equals record-offset layout.	Unique-name storage disagrees with declaration-order traversal.	Schema Uniqueness
Unproven record cardinality	Exact payload extent is established before size-proportional allocation.	Declared count reaches allocation first.	Extent Proof
Degenerate rotation repair	Rotation is finite and nonzero under one declared convention.	Identity repair and byte normalization select different orientations.	Rotation Validity
Radiance-space reinterpretation	Declared SH and color semantics persist to radiance evaluation.	Linear coefficients traverse a different color interpretation.	Coordinate and Radiance Canonicalizer

The private minimized sizes are 434, 363, 32 and $1{,}260$ bytes in Table 6 order; these characterize the retained cases without providing reconstructive content. The two PLY mechanisms affect one native importer family but violate different invariants with different consequences. The count mechanism persists across an earlier and a current source line, the rotation mechanism is accepted by two viewer families, and the radiance mechanism is specific to the pinned extension and engine path.

E CONFORMANCE AND COMPATIBILITY

E.1 CANONICAL TOLERANCES

Per-field tolerances are selected from development observations and frozen before the sealed test split is evaluated. Values in Table 7 are absolute errors, except the sign-invariant quaternion chordal distance.

Table 7: Fixed BGIR round-trip tolerances.

Format	Position	Scale	Rotation	Opacity	SH0	SH residual
PLY	2×10^{-6}	2×10^{-6}	2×10^{-6}	2×10^{-6}	2×10^{-6}	2×10^{-6}
GLB	2×10^{-6}	2×10^{-6}	2×10^{-6}	2×10^{-6}	2×10^{-6}	2×10^{-6}
SPLAT	2×10^{-6}	2×10^{-6}	0.025	0.0039236	0.0039236	10^{-12}
SPZ	0.001	0.1	0.025	0.0039236	0.01	0.0830
SOG	0.002	0.2158	0.03	0.01	0.1720	0.3514

Structural preflight accepts $8{,}000/8{,}000$ final legal assets with Wilson 95% lower bound 0.9995. BGIR accepts $7{,}998/8{,}000$ with Wilson interval $[0.9991, 0.9999]$. Both misses are SOG SH-residual errors above 0.3514, with test maximum 0.3833; every other format-field maximum stays within its fixed tolerance.

E.2 COMPLETE PUBLIC-FILE MATRIX

Beyond Table 8, sealed generated conformance shows the multi-format converter accepting $5{,}600$ of $6{,}400$ applicable cases and the native compressed path accepting $2{,}400$ of $3{,}200$; the lower totals

Table 8: Complete public-file compatibility by file class.

Path	PLY	SPZ	SPLAT	SOG†	GLB†	Boundary note
Structural preflight	8/8	2/2	4/4	4/4	4/4	bounded format-explicit admission
BGIR	8/8	n/a	4/4	n/a	4/4	registered canonical routes
Multi-format converter	8/8	2/2	4/4	4/4	n/a	format support boundary
Native compressed path	8/8	2/2	n/a	n/a	n/a	PLY conversion and SPZ decode
Game-engine importer	8/8	n/a	n/a	n/a	4/4	PLY harness and browser GLB path
Standalone WebGL viewer	8/8	0/2	4/4	n/a	n/a	pinned path rejects current SPZ

†Each target contains four conversions of the four native SPLAT objects, used as route-applicability probes rather than independent native samples.

Table 9: Complete format and seed candidate-generation matrix. Entries are deep-decode fractions. Mechanism counts are source-adjudicated unique mechanisms, not provisional signatures.

Method	Format	Seed 1	Seed 2	Seed 3	Seed 4	Seed 5	Mean	Mechanisms
Structured	PLY	0.390	0.404	0.437	0.389	0.407	0.405	1
Structured	SPZ	0.143	0.140	0.139	0.141	0.146	0.142	0
Structured	SPLAT	0.693	0.684	0.666	0.706	0.674	0.685	1
Structured	SOG	0.508	0.482	0.507	0.503	0.469	0.494	0
Structured	GLB	1.000	1.000	1.000	1.000	1.000	1.000	0
Byte	PLY	0.889	0.840	0.868	0.861	0.852	0.862	1
Byte	SPZ	0.324	0.271	0.289	0.350	0.262	0.299	0
Byte	SPLAT	0.573	0.591	0.566	0.596	0.589	0.583	0
Byte	SOG	0.156	0.125	0.109	0.150	0.210	0.150	0
Byte	GLB	1.000	1.000	1.000	1.000	1.000	1.000	0

reflect narrower implementation format or version domains rather than SPLATCONTRACT rejections. Other external paths accept all cases in their applicable final strata, and defense policy selection does not use external permissiveness as a target.

F CANDIDATE-GENERATION COMPARISON

In Table 9, structured active time is $90{,}007$ s and matched byte active time is $90{,}009$ s. Averaging the 25 cell fractions gives 0.5451 for structured mutation and 0.5788 for byte mutation, a paired structured-minus-byte difference of -0.0337 with bootstrap 95% interval $[-0.1393, 0.0710]$ and Cliff's $\delta = -0.048$; the ratio of method means is 0.942.

Table 10: Adjudicated mechanism yield and active time to first mechanism.

Format	Method	Positive seeds	Mechanisms	KM median
PLY	Structured	5/5	1	4.06 s
PLY	Byte	1/5	1	not reached; one event at 1,244.0 s
SPLAT	Structured	5/5	1	2,326.3 s
SPLAT	Byte	0/5	0	not reached
SPZ, SOG, GLB	Both	0/15 per method	0	not reached

Table 10 shows that structured mutation reaches the duplicate-property and degenerate-rotation mechanisms while byte mutation reaches only the former. Parser-only replay processes all $150{,}475$ structured candidates, flags $100{,}055$, detects both formal mechanisms, and has zero mechanism-level misses; replay is not a matched discovery budget and therefore does not enter the paired reach comparison, but it shows that these two formal mechanisms do not require a semantic-only discovery oracle.

G RENDERING AND RESOURCE EVIDENCE

G.1 COMPLETE RENDERING MATRICES

The KHR study produces 480 images and 240 matched pairs. SSIM mean, median, minimum and maximum are 0.9460, 0.9572, 0.8568 and 0.9923, with scene-cluster interval $[0.9301, 0.9603]$ for the mean; LPIPS mean is 0.0233 with scene-cluster interval $[0.0162, 0.0282]$. Exactly 120 pairs are material and no repeat group is nondeterministic.

The rotation study produces 24 images and 12 matched pairs. SSIM mean, minimum and maximum are 0.6815, 0.4191 and 0.9052; mean foreground-fraction absolute shift is 0.172 and the maximum shift is 0.297. The interpreted rotations are identity and $(0.5, 0.5, 0.5, 0.5)$, at sign-invariant chordal distance 1.0. Non-finite and negative-scale control families produce $0/48$ material pairs.

The mechanism-level review pool contains 29 repeat-zero cases: five cameras from each of five generated KHR scenes plus four rotation cameras. It supports auditing of the original causal mechanisms and supplies the two exact-reference checkerboards in Figure 3. The calibrated real-scene pool contains all eight 30k scene-and-view candidates; the frozen selection rule requires a material unmediated/reference difference and bounded mediated recovery, prefers exact recovery, and uses capture/source diagnostic SSIM only to choose between eligible views of the same independent scene. It selects one Train and one Truck view, while both indoor scenes fail materiality before layout. The figure follows source order and uses the same fixed-gain error transform for baseline and method; candidate and selected decks preserve image hashes, cameras and metrics, and selection does not alter any aggregate. Table 11 reports the complete matrix rather than only the selected rows.

Table 11: Complete calibrated real-scene matrix. SSIM is rounded to four decimal places; 1.0000 denotes agreement at that reporting precision, while maximum absolute error distinguishes pixel-exact from bounded recovery. "Material" is the frozen causal threshold, not visual preference. The matched reference fixes geometry, camera, renderer, and degree-0 target radiance; public captures provide scene context only.

Split	Scene	Views	Pairs	Material	Unmediated SSIM↑	SplatPolyglot SSIM↑	Max abs↓
30k main	Train	2	6	6/6	0.7639	1.0000	0.003922
30k main	Truck	2	6	6/6	0.7288	1.0000	0.003922
30k main	Dr Johnson	2	6	0/6	0.9875	1.0000	0.000000
30k main	Playroom	2	6	0/6	0.9892	1.0000	0.000000
7k sensitivity	Train	1	3	3/3	0.7694	1.0000	0.003922
7k sensitivity	Truck	1	3	3/3	0.7431	1.0000	0.003922
7k sensitivity	Dr Johnson	1	3	0/3	0.9928	1.0000	0.000000
7k sensitivity	Playroom	1	3	0/3	0.9933	1.0000	0.000000

G.2 RESOURCE STRATA

Preflight rejects every compressed-expansion case in Table 12 and all 22 policy cases before external decode. The complete resource matrix contains 107 valid terminal cases, and program maxima span different inputs and boundaries rather than forming a steady-state ranking. The cardinality case's $61{,}543{,}936$ KiB instrumented-process peak and $49{,}300{,}428$ KiB whole-step peak independently corroborate a tens-of-GiB effect and are reported as separate estimators.

H ARTIFACT SCOPE AND EXECUTABLE MECHANISM LAWS

The anonymous artifact verifies revisions and manifests and runs a resumable driver with append-only cases and atomic completion, failing explicitly on missing builds. Tests exercise canonical invariants, format round trips, fail-closed mutations and row alignment. A safe-reproduction suite executes each abstract mechanism three times and completes $12/12$ deterministic terminals under the disclosure audit; its source-faithful constructions use a toy duplicate-binding schema, a capped logical allocator for count/extent disagreement, a benign zero-rotation pair, and a scalar radiance-transfer pair, and mediation respectively rejects the ambiguous field, the allocation and the rotation, or restores the declared radiance value. Scripts verify aggregate constants before every display, no experimental

Table 12: Compressed expansion cases; expanded size is a declaration, not measured RSS.

Format	Input bytes	Declared expanded bytes	Ratio
SPZ	1,067	1,048,576	982.7×
SOG	12,676	1,048,576	82.7×
SPZ	16,345	16,777,216	1,026.4×
SOG	27,953	16,777,216	600.2×
SPZ	81,573	83,886,080	1,028.4×
SOG	93,182	83,886,080	900.2×

image is generated by a model, and disclosure-sensitive bytes, offsets and operational traces are excluded under the ethics protocol.

I Component, sensitivity and efficiency evidence

I.1 Full component routing

Table 4 reports unique newly forwarded candidates. Downstream acceptance Wilson intervals are $[0.4012, 0.4093]$ for Extent Proof, $[0.8034, 0.8111]$ for Schema Uniqueness, $[0.9389, 0.9439]$ for Finite Numeric Domain and $[0.9993, 1.0000]$ for Rotation Validity; counts and intervals have different denominators because the latter cover downstream adapter classifications.

Two development sensitivities resolve to profile-level choices. The profile bounds scale through finite post-activation values and resource limits; a physical unit ceiling belongs to an explicit unit-bearing extension. Arbitrary node-affine flattening is assigned to an explicit profile extension with its own transform semantics.

I.2 Paired efficiency details

Fifty within-input and repeat blocks run on one host with cyclic method order. Structural preflight, full semantic guard and external decode each have 50 cells; BGIR has 40 because two SPZ files lie outside the public canonicalization cell. Incremental guard wall time over canonicalization has median, p95 and maximum ratios 0.206, 0.215 and 0.219; full-pipeline peak RSS over canonicalization is 1.0 at median, p95 and maximum; and full guard over external-decode wall time has median 0.323, p95 0.515 and maximum 0.515. These are paired observations under one offline-ingestion configuration.

The separate one-host scaling experiment covers eight PLY checkpoints, four methods, five repeats and 0.742 to 3.178 million Gaussians. The field-wise construction accepts $2{,}000/2{,}000$ sealed legal PLYs, rejects both applicable adjudicated mechanisms, and matches the materialized fingerprint in $40/40$ blocks. Its p95 cost is 6.08 s and $1,290,976$ KiB, versus 8.23 s and $5,430,768$ KiB for materialization, giving wall and RSS ratios 0.738 and 0.238. Their per-file median RSS slopes are 0.334 and 1.623 GiB per million Gaussians, respectively; the corresponding fits have $R^2 = 0.9997$ and $R^2 > 0.9999$. Figure 4 plots all four trajectories without extrapolation. Two complete intermediate refinements preserve all 40 fingerprints and semantic-regression outcomes with p95 RSS ratios of 0.437; the final field-wise method reaches 0.238 by mapping one `float64` field at a time and reading at most $65{,}536$ rows per block.

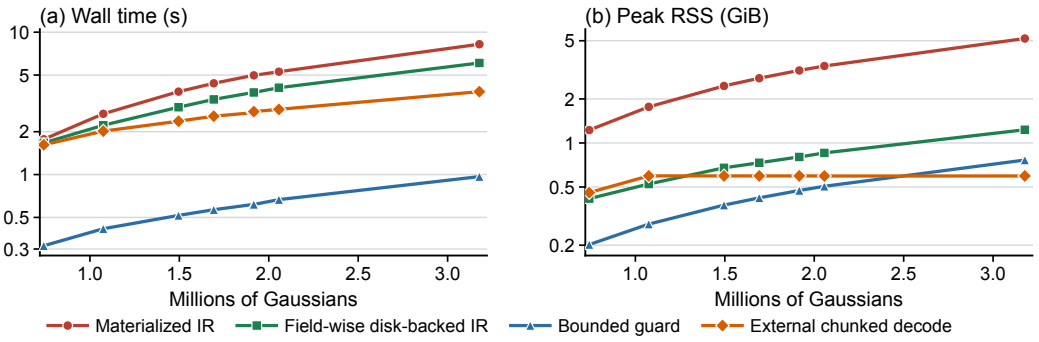


Figure 4: Cardinality scaling on eight PLY checkpoints. Each point is the within-file median of five order-rotated repeats; both vertical axes are logarithmic. Field-wise disk-backed construction preserves the materialized fingerprint while reducing the resident-set slope.