

Fourier Splatting: Generalized Fourier encoded primitives for scalable radiance fields

Mihnea-Bogdan Jurca^{1,2*}, Bert Van hauermeiren^{1*}, and
Adrian Munteanu¹

¹ Vrije Universiteit Brussel, Brussels, Belgium

² Technical University of Cluj-Napoca, Cluj-Napoca, Romania

{Mihnea-Bogdan.Jurca, Bert.Karel.Van.hauermeiren,
Adrian.Munteanu}@vub.be

Abstract. Novel view synthesis has recently been revolutionized by 3D Gaussian Splatting (3DGS), which enables real-time rendering through explicit primitive rasterization. However, existing methods tie visual fidelity strictly to the number of primitives: quality downscaling is achieved only through pruning primitives. We propose the first inherently scalable primitive for radiance field rendering. **Fourier Splatting** employs scalable primitives with arbitrary closed shapes obtained by parameterizing planar surfels with Fourier encoded descriptors. This formulation allows a single trained model to be rendered at varying levels of detail simply by truncating Fourier coefficients at runtime. To facilitate stable optimization, we employ a straight-through estimator for gradient extension beyond the primitive boundary, and introduce HYDRA, a densification strategy that decomposes complex primitives into simpler constituents within the MCMC framework. Our method achieves state-of-the-art rendering quality among planar-primitive frameworks and comparable perceptual metrics compared to leading volumetric representations on standard benchmarks, providing a versatile solution for bandwidth-constrained high-fidelity rendering.

Keywords: novel view synthesis · radiance fields · splatting · scalability

1 Introduction

Photorealistic novel view synthesis has emerged as a cornerstone of computer vision, underpinning downstream tasks from content generation [24, 28] to semantic scene understanding [10, 20, 26], etc. Neural Radiance Fields (NeRF) [18] established that continuous volumetric representations can produce high-fidelity views, but their reliance on dense ray marching confines rendering to offline settings. 3D Gaussian Splatting (3DGS) [11] offered a compelling alternative: explicit Gaussian primitives combined with tile-based rasterization enable real-time rendering, making splatting-based methods the preferred framework for interactive radiance fields. Given their growing impact across applications, scalability of these representations has become one of the most pressing open problems.

* Equal contributions.



Fig. 1: Overview of Fourier Splatting using generalized Fourier encoded primitives. Varying the number K of retained frequencies progressively improves the shape of the primitives. This provides an inherently scalable representation of the radiance field with progressively improved rendering quality.

Recent efforts to refine radiance field primitives have moved beyond the standard ellipsoidal 3D Gaussian, seeking representations that better capture complex geometries and sharp boundaries. These include volumetric approaches [5, 7, 16, 30], with higher quality image generation and planar approaches [6, 8, 23] that enforce surface consistency and resolve the depth ambiguities inherent in volumetric formulations. However, while these designs improve per-primitive fidelity, they remain fundamentally static in their capacity; they rely on traditional density-driven optimization, where high-frequency details are resolved through the proliferation of more primitives rather than an expansion of the primitive’s own geometric complexity.

To manage the computational and memory overhead of increasingly dense splat collections, contemporary research has turned toward Level-of-Detail (LoD) hierarchies. State-of-the-art scalable frameworks typically organize Gaussians into hierarchical layers where distant or less significant primitives are pruned or filtered based on importance metrics. While effective for reducing bitrate and improving real-time performance, these methods achieve scalability only through pruning. By focusing exclusively on reducing the primitive count.

To address these limitations, we present **Fourier Splatting**, the first inherently scalable primitive for real-time radiance field rendering. Rather than relying solely on primitive population growth to increase fidelity, we propose a novel primitive that can approximate any arbitrary shape and is inherently scalable. Each surfel’s boundary is parameterized as a Fourier encoded descriptors, where the coefficients control progressively finer shape detail: at the lowest order the

primitive reduces to a standard disc, and each additional frequency component enables it to express increasingly complex closed curves within a single, compact parameterization. The design yields a natural level-of-detail (LoD) hierarchy: the same trained model can be rendered at any target complexity simply by truncating the Fourier coefficients, without retraining or post-hoc compression. We achieve state-of-the-art rendering among planar methods, and achieve at least similar performance to volumetric 3D representations on standard benchmarks. Our primary contributions are summarized as follows:

- **Arbitrary-Shape Planar Primitives:** We introduce a novel planar primitive for splatting that leverages a Fourier-based boundary parameterization. Unlike other splats, our primitives can assume arbitrary shapes, enabling the representation of complex geometry and fine-grained details with significantly higher fidelity compared to 2D and 3D Gaussians.
- **Inherent Scalability:** We present the first radiance field primitive that is inherently scalable. By truncating the set of active Fourier coefficients, the level-of-detail (LoD) can be continuously adjusted at runtime without additional training or post-processing, providing direct control over the rate-distortion trade-off.
- **MCMC-Compatible Densification:** We propose HYDRA, a densification strategy that enables generalizable planar primitives to operate within the MCMC framework, together with a straight-through estimator that extends gradients beyond the primitive boundary to prevent optimization stalling.
- **State-of-the-Art Performance:** We demonstrate that our method surpasses existing planar-primitive approaches across standard benchmarks. Furthermore, our representation achieves competitive results against leading volumetric methods, outperforming in some categories. Finally, we demonstrate that our scalable primitive exhibits more graceful qualitative degradation during downscaling compared to the state-of-the-art

2 Related Work

Neural scene representations. Neural radiance fields [18] established image-based 3D reconstruction through implicit volumetric representations. While subsequent work has significantly reduced training times and improved quality [1, 2, 19], the underlying reliance on dense per-ray volume integration remains a fundamental computational bottleneck. To circumvent this, 3D Gaussian Splatting (3DGS) [11] introduced an explicit representation composed of millions of anisotropic Gaussians. By replacing volumetric sampling with highly optimized tile-based rasterization, 3DGS achieves real-time frame rates with state-of-the-art visual fidelity. Recent advancements have further refined the 3DGS framework by optimizing densification through MCMC sampling [12], absolute gradient accumulation [27], or by introducing structured spatial organization [17, 21]. All these approaches use the same smooth elliptical Gaussian primitive.

Alternative primitive representations. An emerging branch of research proposes to move beyond the standard Gaussian kernel as the building block

for splatting. GES [5] generalizes the radial falloff to an exponential family with a learnable shape parameter, enabling sharper profiles. Deformable Beta Splatting [16] replaces Gaussians with compact-support Beta kernels whose shape parameters control the falloff from uniform to peaked. 3D-HGS [15] splits each Gaussian along a plane to create asymmetric half-kernels for discontinuities. DRK [9] parameterizes the radial profile with learnable basis functions whose scale and sharpness vary per angular direction, producing star-shaped footprints. Beyond kernel modifications, 3D Convex Splatting [7] replaces Gaussians with smooth convex bodies that can represent hard edges and flat surfaces, Quadratic Gaussian Splatting [30] deforms a surfel from a flat disc to a paraboloid. While these methods significantly expand the design space of the splatting unit, they share a common constraint: the expressiveness of each primitive is determined by its architectural design and remains static once trained.

Surface primitives. Another line of research is surface-based splatting methods, which better aligns primitives with underlying scene geometry to facilitate high-fidelity reconstruction. 2D Gaussian Splatting (2DGS) [8] collapses 3D Gaussians to oriented planar surfels. Gaussian Surfels [4] similarly flatten ellipsoids to discs and enforce normal-depth consistency through monocular priors. Triangle Splatting [6] departs from smooth kernels entirely, rendering optimizable triangles as differentiable splats with hard boundaries. BBSplat [23] equips each billboard plane with a learned texture and alpha mask, encoding spatially-varying appearance within a single flat primitive.

Scalable representations. Orthogonal to primitive design, a separate line of work focuses on managing the complexity of Gaussian-based scenes. Octreegs [21] and LODGE [14] maintain different sets of LoD Gaussians, pruning and blending Gaussians based on rendering distance. LOD-GS [22] learns triangle soup representation with 3D gaussians embedded in them, scalability is achieved through downsampling the gaussians in each triangle. While adaptive densification strategies [11, 12, 27] control how many primitives are spawned during optimization. In all cases, the axis of scalability is the number of primitives: representations are made smaller by removing Gaussians and larger by adding them, while each individual primitive remains unchanged. Our Fourier Splatting approach differs fundamentally by embedding scalability within the primitive itself; rather than discarding entire spatial regions or points, we reduce the precision of the primitives themselves.

3 Preliminaries

3.1 3D Gaussian Splatting

3DGS [11] represents a scene as a collection of anisotropic Gaussian primitives. Each primitive is parameterized by a mean $\boldsymbol{\mu} \in \mathbb{R}^3$, a covariance $\boldsymbol{\Sigma} = \mathbf{R}\mathbf{S}\mathbf{S}^\top\mathbf{R}^\top$ factored into rotation $\mathbf{R}$ and diagonal scale $\mathbf{S}$, an opacity $o \in (0, 1)$, and spherical harmonic coefficients for view-dependent color $\mathbf{c}$. The 3D covariance is projected to 2D via $\boldsymbol{\Sigma}' = \mathbf{J}\mathbf{V}\boldsymbol{\Sigma}\mathbf{V}^\top\mathbf{J}^\top$, where $\mathbf{V}$ is the viewing transformation (world-to-

camera) and $\mathbf{J}$ the Jacobian of the local affine approximation of the projective transformation. Primitives are sorted by depth and accumulated front-to-back:

$$\mathbf{C}(\mathbf{x}) = \sum_{i=1}^N \mathbf{c}_i \alpha_i T_i = \sum_{i=1}^N \mathbf{c}_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j). \quad (1)$$

where α_i is the product of opacity and the evaluated 2D Gaussian $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}'_i, \boldsymbol{\Sigma}'_i)$. This enables real-time tile-based rasterization, but each primitive is rendered as a smooth ellipsoidal kernel with infinite support and fixed shape.

3.2 2D Gaussian Splatting

2DGS [8] replaces volumetric Gaussians with oriented planar surfels. Each surfel is defined by a center $\mathbf{p}_k \in \mathbb{R}^3$, tangent vectors $\mathbf{t}_u, \mathbf{t}_v$ with scales $s_u, s_v > 0$, and normal $\mathbf{t}_w = \mathbf{t}_u \times \mathbf{t}_v$. A point on the surfel is:

$$\mathbf{P}(u, v) = \mathbf{p}_k + s_u \mathbf{t}_u u + s_v \mathbf{t}_v v = \mathbf{H}(u, v, 1, 1)^\top, \quad (2)$$

where $\mathbf{H} \in \mathbb{R}^{4 \times 4}$ is the homogeneous surfel-to-world matrix.

Ray-surfel intersection. Instead of projecting a covariance, 2DGS intersects each pixel ray with the surfel plane. Let $\mathbf{W} \in \mathbb{R}^{4 \times 4}$ be the world-to-screen transformation. A pixel at (x_p, y_p) defines two screen-space planes $\mathbf{h}_x = (-1, 0, 0, x_p)^\top$ and $\mathbf{h}_y = (0, -1, 0, y_p)^\top$. Transforming these into the surfel-local frame gives $\mathbf{h}_u = (\mathbf{W}\mathbf{H})^\top \mathbf{h}_x$ and $\mathbf{h}_v = (\mathbf{W}\mathbf{H})^\top \mathbf{h}_y$. The intersection coordinates follow from solving the 2×2 system:

$$u_j = \frac{h_u^2 h_v^4 - h_u^4 h_v^2}{h_u^1 h_v^2 - h_u^2 h_v^1}, \quad v_j = \frac{h_u^4 h_v^1 - h_u^1 h_v^4}{h_u^1 h_v^2 - h_u^2 h_v^1}. \quad (3)$$

The per-pixel opacity is then $\alpha_i = o_i \exp(-(u_j^2 + v_j^2)/2)$, a 2D Gaussian evaluated in the surfel plane. Our method inherits the surfel parameterization, ray-plane intersection, and tile-based compositing from 2DGS, but replaces the fixed Gaussian falloff with the Fourier boundary and power window described in Secs. 4.1 and 4.2.

4 Method

Fourier Splatting builds on the surfel formalism [8, 32] but departs from the fixed disc and ellipsoid assumptions. Given a scene represented by a set of N primitives $\mathcal{S} = \{s_i\}_{i=0}^{N-1}$, we equip each s_i with a Fourier encoded descriptors that defines an arbitrary closed boundary within its tangent plane, making primitive expressiveness a continuous, controllable quantity rather than a fixed design choice. We first describe the primitive parameterization (Sec. 4.1), then its differentiable rendering (Sec. 4.2), and the densification, pruning and training strategies (Sec. 4.3).




Fig. 2: Overview of the Fourier splatting primitive.

4.1 Fourier Splatting Primitive

Each primitive s_i is parameterized by a center $\mathbf{p}_i \in \mathbb{R}^3$, a rotation quaternion $\mathbf{q}_i \in \mathbb{R}^4$ defining the local tangent frame $(\mathbf{t}_u, \mathbf{t}_v, \mathbf{t}_w)$, an opacity $o_i \in (0, 1)$, spherical harmonic coefficients $\mathbf{c}_i \in \mathbb{R}^{48}$ for view-dependent appearance, a circumradius $R_i \in \mathbb{R}_{>0}$ bounding the primitive extent, Fourier amplitudes $\{r_{i,k}\}_{k=0}^{K-1} \in \mathbb{R}^K$ and phases $\{\phi_{i,k}\}_{k=0}^{K-1} \in \mathbb{R}^K$ defining the boundary shape, and a sharpness control parameter $\sigma_i \in \mathbb{R}_{>0}$.

Fourier encoded boundary. We define the boundary of each primitive s_i as a closed curve $r_i(\theta)$ in its local tangent plane (Fig. 2a). Given a point (u_j, v_j) in the tangent plane obtained via ray-plane intersection (Eq. (3)), we compute its polar coordinates relative to the primitive center μ_i : the radial distance $\rho_j = \sqrt{u_j^2 + v_j^2}$ and the angle $\theta_j = \arctan\left(\frac{v_j}{u_j}\right)$. The boundary radius at angle θ_j is given by the modulus of a Fourier polynomial:

$$r_i(\theta_j) = R_i \cdot \left| \sum_{k=0}^{K-1} c_{i,k} e^{ik\theta_j} \right|, \quad (4)$$

where $c_{i,k} = \bar{r}_{i,k} e^{i\phi_{i,k}}$ are complex coefficients with normalized amplitudes $\bar{r}_{i,k}$ and phases $\phi_{i,k}$ (note: i in the exponent denotes the imaginary unit, while the subscript i indexes the primitive). The $k=0$ term acts as the DC component, setting a uniform base radius, while each higher frequency k introduces k -fold angular variation. At $K=1$ only the DC term remains and the boundary reduces to a circle of radius $R_i \bar{r}_{i,0}$, recovering the standard disc primitive.

Amplitude normalization. Since each frequency k contributes a radius $\bar{r}_{i,k}$ to the boundary, the total reach of the primitive is bounded by the sum of

all amplitudes. Without normalization, these amplitudes can sum to exceed the circumradius R_i , causing the boundary to overshoot the circumscribed circle. We enforce this constraint via squared- ℓ_1 normalization:

$$\bar{r}_{i,k} = \frac{r_{i,k}^2}{\sum_{k'=0}^{K-1} r_{i,k'}^2} \quad \text{such that} \quad \sum_{k=0}^{K-1} \bar{r}_{i,k} = 1. \quad (5)$$

Since each $\bar{r}_{i,k} \geq 0$ and their sum equals one, the maximum constructive interference across all frequencies cannot exceed R_i , guaranteeing $r_i(\theta) \leq R_i$ for all θ .

4.2 Differentiable Rendering

Forward pass. We implement a tile-based rasterizer for the Fourier primitive, adapting the power window function of Triangle Splatting [6] to our polar coordinate formulation. For each pixel j with polar coordinates (ρ_j, θ_j) in the tangent plane of primitive s_i , the signed distance to the Fourier boundary is:

$$\varphi_j = \rho_j - r_i(\theta_j), \quad (6)$$

which is negative inside and positive outside the primitive; points with $\varphi_j \geq 0$ are discarded, providing compact support. The per-pixel opacity is:

$$\alpha_{i,j} = o_i \cdot \max\left(0, \frac{r_i(\theta_j) - \rho_j}{r_i(\theta_j)}\right)^{\sigma_i}. \quad (7)$$

The window peaks at the primitive center and decays to zero at the boundary, with σ_i controlling the decay profile. Colors are accumulated via front-to-back compositing (Eq. (1)).

Backward pass. The hard clamp in Eq. (7) creates a zero-gradient region outside the primitive boundary. For any pixel with $\varphi_j \geq 0$, the rendered contribution is exactly zero and the gradient with respect to the Fourier coefficients $c_{i,k}$ vanishes. The boundary therefore contracts, as most gradient signal comes from interior pixels. To compensate, the optimization inflates the circumradius R_i , scaling the entire shape uniformly rather than letting the Fourier coefficients learn more complex geometries.

To address this, we draw on the straight-through estimator (STE) [3] and decouple the forward and backward passes. The forward rendering is unchanged: the hard window of Eq. (7) determines the rendered image. In the backward pass, we replace the hard window with a smooth surrogate $\tilde{w}$ whose gradient is non-zero on both sides of the boundary (Fig. 3):

$$\tilde{w}(x) = \text{sigmoid}(\beta x) \cdot \left(\frac{\text{softplus}(\beta x)}{\beta} \right)^{\sigma_i} + \gamma \cdot \text{sigmoid}(\beta x), \quad (8)$$

where $x = (r_i(\theta_j) - \rho_j)/r_i(\theta_j)$ is the unclamped normalized distance ($x > 0$ inside, $x < 0$ outside), and β, γ are fixed hyperparameters. The first term is




Fig. 3: Straight-Through Estimator (STE) for gradient extension. (1) Forward rendering with the hard cutoff window (Eq. (7)). (2) Gradient without STE: dead zone outside the boundary. (3) Smooth STE surrogate (Eq. (8)). (4) Gradient with STE: signal flows across the full primitive extent. Solid line: $r_i(\theta)$; dashed: R_i . $\beta=3$, $\gamma=0.5$, $\sigma=1$.

a smooth approximation of $\max(0, x)^{\sigma_i}$, matching the forward window for interior points. The additive sigmoid term ensures $\partial\tilde{w}/\partial x \neq 0$ even when $x < 0$, providing gradient signal to the Fourier coefficients from pixels just outside the boundary. This allows the boundary to both expand and contract during optimization, letting the Fourier coefficients learn complex shapes instead of relying solely on the circumradius R_i for spatial coverage. Since exterior pixels have zero contribution in the forward pass, we restrict their STE gradient to the Fourier coefficients $c_{i,k}$. Position and orientation gradients are computed from interior pixels where the photometric loss is well-defined. The effect of β and γ on the gradient coverage area is illustrated in Fig. 3.

4.3 Optimization

We adopt the MCMC framework of 3DGS-MCMC [12], which interprets SGD updates as Stochastic Gradient Langevin Dynamics (SGLD) and treats primitives as samples from a target distribution $P(\mathbf{g}) \propto \exp(-\mathcal{L})$, where $\mathbf{g}$ denotes the full primitive state. Densification is a deterministic state transition where dead primitives are relocated to live ones with parameters adjusted to preserve $P(\mathbf{g}^{\text{new}}) = P(\mathbf{g}^{\text{old}})$. A key ingredient of MCMC [12] is the SGLD noise added to primitive positions, scaled proportionally to the covariance Σ and inversely to the opacity, encouraging spatial exploration. For planar primitives this is ill-suited: positional noise scaled by a 3D covariance would push surfels out of their tangent plane when they have not yet converged or lie in high volumetric-frequency regions, and since children inherit the parent’s orientation, such perturbations violate the planar constraint and degrade $P(\mathbf{g}^{\text{new}})$.

Relocation. When a primitive with opacity o , circumradius R , and sharpness σ is replaced by N co-located copies, we require the rendered image to remain unchanged. The new opacity follows the kernel-independent formula $o_{\text{new}} = 1 - (1 - o)^{1/N}$ [12]. For the scale, we equate the 2D polar integral of the parent’s power window $w(\rho) = \max(0, (R - \rho)/R)^\sigma$ with the composited integral of N children as in [12]. Expanding the per-child transmittance $(1 - o_{\text{new}} w)^{i-1}$

via the binomial theorem and integrating each term $w^{k+1} = ((R - \rho)/R)^{\sigma(k+1)}$ in polar coordinates gives:

$$D = \sum_{i=1}^N \sum_{k=0}^{i-1} \binom{i-1}{k} (-1)^k \frac{o_{\text{new}}^{k+1}}{(\sigma(k+1) + 1)(\sigma(k+1) + 2)}, \quad (9)$$

where the denominator factors replace the Gaussian integral of [12]. The corrected circumradius is:

$$R_{\text{new}} = R \sqrt{\frac{o}{(\sigma + 1)(\sigma + 2) \cdot D}}. \quad (10)$$

Death criterion. Primitives are removed if: (i) opacity falls below τ_o , (ii) the maximum blending weight $\max_j (\alpha_{i,j} T_{i,j})$ over all pixels j across recent views falls below τ_I (Eq. (1)), or (iii) primitive s_i is observed in fewer than V_{min} frames [6]. Newly created primitives are protected by a one-epoch grace period.

Primitive relocation and addition. Dead primitives, identified by the death criterion, are relocated to live ones by sampling primitives weighted by opacity or inverse sharpness $1/\sigma_i$ in alternation due to different local optimization convergence [6]. Each sampled primitive produces two children via the relocation formulas of Eqs. (9) and (10) with $N=2$: one at the sampled position, the other offset in the primitive’s tangent plane by $\eta\epsilon$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, with noise scale η . To reach the budget N_{max} , primitives are added either indirectly through HYDRA decomposition, or by sampling additional primitives for relocation. In all cases $P(\mathbf{g}^{\text{new}}) \approx P(\mathbf{g}^{\text{old}})$ is maintained.

HYDRA: Hybrid Decomposition for Rendering-Aware Preservation. Under-reconstructed regions accumulate high positional gradients, as observed in [11, 27]. Since we do not apply positional SGLD noise, these regions risk being undersampled across densification steps. The most effective solution is to subdivide the primitives covering such regions, increasing the number of independent samples that can explore the local geometry. HYDRA is our mechanism for performing this subdivision while preserving $P(\mathbf{g}^{\text{new}}) \approx P(\mathbf{g}^{\text{old}})$. We identify high-gradient primitives via absolute gradient accumulation [27] and distinguish two cases:

(i) *Scale-preserving split.* For smaller primitives, we reduce the circumradius and place two children symmetrically offset in the parent’s tangent plane, applying the mass-preserving relocation formulas of Eqs. (9) and (10) with $N=2$. (ii) *Learned lobe decomposition.* For large primitives that have developed complex, multi-lobed boundaries, a geometric split is insufficient. We identify the S deepest valleys along the boundary $r_i(\theta)$ sampled at M angles, partitioning it into S angular segments. A pre-trained MLP f_ψ generates a child primitive for each segment:

$$f_\psi: (\sigma, \{\phi_k\}, \{r_k\}, r_M, \theta_M, \theta_{\min}, \theta_{\max}) \mapsto (\sigma', \{\phi'_k\}, \{r'_k\}, \Delta r, \Delta\theta, o', R'), \quad (11)$$

where (r_M, θ_M) and $(\theta_{\min}, \theta_{\max})$ denote the center of mass and angular bounds of the segment respectively and $(\Delta r, \Delta\theta)$ is the polar offset used to compute the




Fig. 4: Learned lobe decomposition overview

child center in the parent’s tangent plane. Children inherit color and orientation from the parent, the parent is removed, yielding a net increase of $S - 1$ primitives. Fig. 4

Loss. We optimize a photometric objective combining ℓ_1 and D-SSIM:

$$\mathcal{L}_{\text{photo}} = (1 - \lambda) \mathcal{L}_1 + \lambda \mathcal{L}_{\text{D-SSIM}}. \quad (12)$$

Following 2DGS [8], we add a depth distortion loss $\mathcal{L}_{\text{dist}}$ penalizing spread in rendered depth distributions, and a normal consistency loss $\mathcal{L}_{\text{normal}}$ aligning each primitive’s normal $\mathbf{t}_w$ with the depth-derived surface normal:

$$\mathcal{L} = \mathcal{L}_{\text{photo}} + \lambda_{\text{dist}} \mathcal{L}_{\text{dist}} + \lambda_{\text{normal}} \mathcal{L}_{\text{normal}}. \quad (13)$$

5 Experiments

We evaluate Fourier Splatting against state-of-the-art methods in novel view synthesis across standard benchmarks, and then analyze the scalability properties that distinguish our approach.

5.1 Experimental Setup

Datasets, metrics, and baselines. We evaluate on Mip-NeRF 360 [1] (5 outdoor, 4 indoor scenes) and Tanks and Temples [13] (2 large-scale scenes), measuring PSNR, SSIM [25], and LPIPS [29]. We compare against implicit methods (Mip-NeRF 360 [1], Zip-NeRF [2]), volumetric primitives (3DGS [11], 3DGS-MCMC [12], DBS [16], 3DCS [7], Gabor Splatting [31], GES [5]), and planar primitives (2DGS [8], Triangle Splatting [6], BBSplat [23]). We follow the evaluation protocol of Triangle Splatting [6] and report their published baseline numbers, additional results are taken from DBS [16].

Implementation details. We use $K = 6$ Fourier frequencies and spherical harmonics of degree three, yielding 70 parameters per primitive. The Fourier boundary (Eq. (4)) is evaluated via Horner’s method. We compute $w = e^{i\theta} = (\cos \theta, \sin \theta)$ directly from the tangent-plane coordinates as $\cos \theta = u_j / \rho_j$, $\sin \theta =$




Fig. 5: Qualitative Scalability Analysis. Comparison of the different LoDs between the proposed method and Octree-GS [21] on the Truck from Tanks and Temples.

v_j/ρ_j , avoiding all transcendental calls. The polynomial is then accumulated as a recurrence from $k=K-1$ down to 0:

$$z_{K-1} = c_{i,K-1}, \quad z_k = z_{k+1} \cdot w + c_{i,k}, \quad (14)$$

yielding $r_i(\theta) = R_i|z_0|$ in $K-1$ complex multiply-adds, without evaluating $e^{ik\theta}$ for any $k > 1$. The optimization process begins by activating only the lowest frequency components $K_{\text{active}} = 1$ to establish a stable global structure. At iteration 600, the remaining frequency coefficients are unfrozen, preventing low-frequency dominance from obstructing the acquisition of finer textural details. Due to the orthonormality of the Fourier basis, the optimal L^2 approximation of an K -term polynomial using $K-1$ coefficients is obtained by simple truncation. Because the expansion coefficients are independent, the remaining terms require no adjustment to minimize the residual energy.

5.2 Novel View Synthesis

Quantitative results. Tab. 1 presents a comprehensive evaluation on the Mip-NeRF 360 and Tanks and Temples datasets. Among methods using non-volumetric

Table 1: Quantitative comparison on Mip-NeRF 360 [1] and Tanks & Temples [13]. Among planar primitives: **best**, **second best**. We follow the evaluation protocol of Triangle Splatting [6]. Our method ranks first among planar methods on Mip-NeRF 360 across all metrics and achieves the best SSIM and LPIPS on Tanks & Temples.

Method	Outdoor Mip-NeRF 360			Indoor Mip-NeRF 360			Avg. Mip-NeRF 360			Tanks & Temples		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
<i>Implicit Methods</i>												
Mip-NeRF 360 [1]	24.47	0.691	0.283	31.72	0.917	0.179	27.35	0.792	0.237	22.22	0.759	0.257
Zip-NeRF [2]	25.58	0.750	0.207	32.25	0.926	0.167	28.42	0.826	0.189	—	—	—
<i>Volumetric Primitives</i>												
3DGS [11]	24.64	0.731	0.234	30.41	0.920	0.189	26.98	0.813	0.214	23.14	0.841	0.183
3DGS-MCMC [†] [12]	25.51	0.760	0.210	31.08	0.917	0.208	27.84	0.850	0.210	24.29	0.860	0.190
DBS [†] [16]	25.10	0.751	0.246	32.29	0.936	0.220	28.13	0.827	0.234	24.85	0.870	0.140
GES [5]	24.89	0.717	0.288	27.60	0.856	0.233	26.91	0.794	0.250	23.35	0.836	0.198
3DCS [7]	24.07	0.700	0.238	31.33	0.927	0.166	26.99	0.800	0.207	23.94	0.851	0.156
Octree-GS-Scaffold [21]	24.83	0.724	0.262	32.08	0.937	0.155	28.05	0.819	0.214	24.68	0.866	0.153
<i>Planar Primitives</i>												
2DGS [8]	24.34	0.717	0.246	30.40	0.916	0.195	26.84	0.804	0.252	23.13	0.831	0.212
BBSplat [†] [23]	23.55	0.669	0.281	30.62	0.921	0.178	26.49	0.778	0.236	25.12	0.868	0.172
Tri. Splatting [6]	24.27	0.722	0.217	30.80	0.928	0.160	26.98	0.812	0.191	23.14	0.857	0.143
Ours	24.62	0.739	0.217	31.44	0.930	0.162	27.65	0.824	0.193	24.15	0.868	0.137

primitives, our approach achieves state-of-the-art performance, delivering the best PSNR and SSIM on both indoor and outdoor Mip-NeRF scenes. Specifically, outperforming Triangle Splatting [6] by 0.67 dB on the Mip-NeRF 360 average. Compared to 2DGS [8], which shares our planar surfel formulation, the gains are more pronounced (+0.81 dB). These significant gains validate that the generalised boundary is able to achieve higher reconstruction quality. Furthermore, Fourier Splatting competes favorably against volumetric methods, we match or exceed DBS [16] on average LPIPS despite our purely planar primitive, and surpass 3DGS [11] across all metrics. On Tanks and Temples, our method achieves the top LPIPS score among all tested baselines while remaining highly competitive in PSNR and SSIM.

Qualitative results. Fig. 6 shows visual comparisons on representative scenes from Mip-NeRF 360 [1] and Tanks and Temples [13]. Fourier Splatting recovers fine details—thin structures, high-frequency textures, and sharp silhouettes—more faithfully than 3DGS, which tends to blur these regions due to the limited expressiveness of its geometry. Compared to Triangle Splatting, our method is less prone to artifacts where the underlying primitive shape is clearly visible. We refer to the supplementary material for per-scene breakdowns and additional visual comparisons.

5.3 Scalability Analysis

A key contribution of our Fourier Splatting approach is the ability to scale primitive expressiveness independently of primitive count. This scalability requires no additional training or computation; image quality is adjusted simply by modulating the number of coefficients transmitted, up to the maximum K used during training. In Fig. 7, we demonstrate the monotonic decrease in quality as coeffi-




Fig. 6: Qualitative Analysis. Comparison between the ground truth, state-of-the-art methods and our proposed approach on 2 scenes from Mip-NeRF 360 and Tanks and Temples.

cients are truncated. Fig. 8 provides a quantitative rate-distortion comparison against surface based primitive level-of-detail representation [21], with Fig. 5 showing the corresponding qualitative degradation. While current scalable splatting methods like Octree-GS degrade quality by reducing the primitive count, which can cause focal objects to vanish, our method scales the expressiveness of the primitives themselves. This results in a more graceful and uniform degradation of visual quality.


Compression. Another benefit of the per-primitive scalability is the compatibility with compression codecs, which often code attributes independently. In contrast, scaling by primitive count complicates efficient geometry compression. We demonstrate this benefit by compressing a scene using the ISO V3C codec, resulting in independent bitstreams per primitive, facilitating efficient scalable streaming (Tab. 2).

Table 2: Quantitative results for compressed scalability with V3C.


K	1	2	3	4	5	6
Bytes/primitive	61.92	64.52	67.13	69.67	72.21	74.89
PSNR	12.66	16.74	19.18	20.77	21.81	22.40
SSIM	0.53	0.66	0.74	0.79	0.82	0.84
LPIPS	0.45	0.35	0.27	0.22	0.19	0.17

Table 3: Comparison at matched 960K primitive budget (T&T).

Method	N_{prim}	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FPS	VRAM	Train
2DGS	960K	22.94	0.827	0.160	182	1.4 GB	12 min
Ours ($K=6$)	960K	23.79	0.853	0.156	135	2.08 GB	41 min

**Fig. 7:** Ablation on trained K .**Table 4:** Ablation on Tanks and Temples.

Configuration	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Full model	24.15	0.868	0.137
w/o lobe decomposition	23.93	0.847	0.141
w/o STE	23.81	0.842	0.147
w/o cloning noise	24.10	0.845	0.138

**Fig. 8:** Rate-distortion comparison.

5.4 Ablation Studies

We ablate the key components of Fourier Splatting on Tanks and Temples (Tab. 4). Removing **automatic lobe decomposition** degrades PSNR by 0.22 dB, as naive splitting cannot decompose multi-lobed primitives into simpler constituents that better fit the local geometry. Disabling the **straight-through estimator** causes the largest drop of 0.34 dB: without synthetic gradients outside the boundary, primitives cannot grow to cover under-reconstructed regions and the optimization stalls near circular shapes. By removing the **randomized cloning offset** the representation loses 0.05 dB in PSNR, as the stochastic angular placement of added primitives in the surfel plane encourages spatial diversity.

6 Limitations and Future Work

Bitrate. The proposed radiance field representation prioritizes achieving state-of-the-art fidelity for a given primitive budget rather than minimizing bitrate required to compress and transmit such representations. A promising avenue for future research involves the integration of scale regularization to incentivize the formation of fewer, more spatially expansive, and complex primitives. By leveraging the high-frequency expressiveness of high- K boundaries to represent regions typically requiring dense distributions of simple surfels, such an approach could yield significant improvements in compression efficiency and overall bitrate reduction.

Computational Complexity. Despite their richer geometry, our primitives remain real-time: Tab. 3 reports 135 FPS at a matched 960K budget. The Horner recurrence keeps the overhead small, adding only 8 operations per frequency.

7 Conclusion

Fourier Splatting introduces a novel geometric primitive capable of approximating arbitrary planar shapes, marking the first representation for real-time radiance field rendering with inherently scalable expressiveness. By parameterizing surfel boundaries via a truncatable Fourier encoded descriptors, our method effectively decouples reconstruction fidelity from primitive density. This allows dynamic downscaling of visual quality and bitrate via coefficient truncation rather than primitive removal. To enable stable optimization of these expressive primitives within a fixed budget, we employ a straight-through estimator that extends gradients beyond the primitive boundary and introduce HYDRA, a learned decomposition strategy that splits complex primitives into simpler constituents within the MCMC densification framework. Extensive evaluations on Mip-NeRF 360 and Tanks and Temples demonstrate that Fourier Splatting achieves state-of-the-art rendering quality among surface-based methods and remains competitive with volumetric approaches. Furthermore, our scalability analysis underscores the potential of per-primitive refinement as a robust alternative to traditional pruning. When integrated with classic pruning techniques, our approach provides a versatile framework for high-fidelity rendering in bandwidth-constrained environments.

Acknowledgments

The work is supported by FWO, project Panoptes (G0ACE26N) and BELSPO, project ATTENTION.

References

1. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. pp. 5470–5479 (2022), https://openaccess.thecvf.com/content/CVPR2022/html/Barron_Mip-NeRF_360_Unbounded_Anti-Aliased_Neural_Radiance_Fields_CVPR_2022_paper.html
2. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. pp. 19697–19705 (2023), https://openaccess.thecvf.com/content/ICCV2023/html/Barron_Zip-NeRF_Anti-Aliased_Grid-Based_Neural_Radiance_Fields_ICCV_2023_paper.html
3. Bengio, Y., Léonard, N., Courville, A.: Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation (Aug 2013). <https://doi.org/10.48550/arXiv.1308.3432>, <http://arxiv.org/abs/1308.3432>, arXiv:1308.3432 [cs]

4. Dai, P., Xu, J., Xie, W., Liu, X., Wang, H., Xu, W.: High-quality Surface Reconstruction using Gaussian Surfels. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–11. SIGGRAPH '24, Association for Computing Machinery, New York, NY, USA (Jul 2024). <https://doi.org/10.1145/3641519.3657441>, <https://dl.acm.org/doi/10.1145/3641519.3657441>
5. Hamdi, A., Melas-Kyriazi, L., Mai, J., Qian, G., Liu, R., Vondrick, C., Ghanem, B., Vedaldi, A.: GES : Generalized Exponential Splatting for Efficient Radiance Field Rendering. pp. 19812–19822 (2024), https://openaccess.thecvf.com/content/CVPR2024/html/Hamdi_GES__Generalized_Exponential_Splatting_for_Efficient_Radiance_Field_Rendering_CVPR_2024_paper.html
6. Held, J., Vandeghen, R., Deliege, A., Hamdi, A., Giancola, S., Cioppa, A., Vedaldi, A., Ghanem, B., Tagliasacchi, A., Droogenbroeck, M.V.: Triangle Splatting for Real-Time Radiance Field Rendering (May 2025). <https://doi.org/10.48550/arXiv.2505.19175>, <http://arxiv.org/abs/2505.19175>, arXiv:2505.19175 [cs]
7. Held, J., Vandeghen, R., Hamdi, A., Deliege, A., Cioppa, A., Giancola, S., Vedaldi, A., Ghanem, B., Van Droogenbroeck, M.: 3D Convex Splatting: Radiance Field Rendering with 3D Smooth Convexes. pp. 21360–21369 (2025), https://openaccess.thecvf.com/content/CVPR2025/html/Held_3D_Convex_Splatting_Radiance_Field_Rendering_with_3D_Smooth_Convexes_CVPR_2025_paper.html
8. Huang, B., Yu, Z., Chen, A., Geiger, A., Gao, S.: 2D Gaussian Splatting for Geometrically Accurate Radiance Fields. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–11. SIGGRAPH '24, Association for Computing Machinery, New York, NY, USA (Jul 2024). <https://doi.org/10.1145/3641519.3657428>, <https://dl.acm.org/doi/10.1145/3641519.3657428>
9. Huang, Y.H., Lin, M.X., Sun, Y.T., Yang, Z., Lyu, X., Cao, Y.P., Qi, X.: Deformable Radial Kernel Splatting. pp. 21513–21523 (2025), https://openaccess.thecvf.com/content/CVPR2025/html/Huang_Deformable_Radial_Kernel_Splatting_CVPR_2025_paper.html
10. Jurca, M.B., Royen, R., Giosan, I., Munteanu, A.: RT-GS2: Real-Time Generalizable Semantic Segmentation for 3D Gaussian Representations of Radiance Fields (Aug 2024). <https://doi.org/10.48550/arXiv.2405.18033>, <http://arxiv.org/abs/2405.18033>, arXiv:2405.18033 [cs]
11. Kerbl, B., Kopanas, G., Leimkuehler, T., Drettakis, G.: 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.* **42**(4), 139:1–139:14 (Jul 2023). <https://doi.org/10.1145/3592433>, <https://dl.acm.org/doi/10.1145/3592433>
12. Kheradmand, S., Rebain, D., Sharma, G., Sun, W., Tseng, Y.C., Isack, H., Kar, A., Tagliasacchi, A., Yi, K.M.: 3D Gaussian Splatting as Markov Chain Monte Carlo. *Advances in Neural Information Processing Systems* **37**, 80965–80986 (Dec 2024). <https://doi.org/10.52202/079017-2573>, https://proceedings.neurips.cc/paper_files/paper/2024/hash/93be245fcea0a9bb2333c17ceae4b732-Abstract-Conference.html
13. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: benchmarking large-scale scene reconstruction. *ACM Trans. Graph.* **36**(4), 78:1–78:13 (Jul 2017). <https://doi.org/10.1145/3072959.3073599>, <https://dl.acm.org/doi/10.1145/3072959.3073599>
14. Kulhanek, J., Rakotosaona, M.J., Manhardt, F., Tsalicoglou, C., Niemeyer, M., Sattler, T., Peng, S., Tombari, F.: LODGE: Level-of-Detail Large-Scale Gaussian Splatting with Efficient Rendering (Oct 2025). <https://doi.org/10.48550/arXiv.2505.23158>, <http://arxiv.org/abs/2505.23158>, arXiv:2505.23158 [cs]

15. Li, H., Liu, J., Sznaiier, M., Camps, O.: 3D-HGS: 3D Half-Gaussian Splatting. pp. 10996–11005 (2025), https://openaccess.thecvf.com/content/CVPR2025/html/Li_3D-HGS_3D-Half-Gaussian-Splatting_CVPR_2025_paper.html
16. Liu, R., Sun, D., Chen, M., Wang, Y., Feng, A.: Deformable Beta Splatting. In: Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers. pp. 1–11. SIGGRAPH Conference Papers '25, Association for Computing Machinery, New York, NY, USA (Jul 2025). <https://doi.org/10.1145/3721238.3730716>, <https://dl.acm.org/doi/10.1145/3721238.3730716>
17. Lu, T., Yu, M., Xu, L., Xiangli, Y., Wang, L., Lin, D., Dai, B.: Scaffold-GS: Structured 3D Gaussians for View-Adaptive Rendering. pp. 20654–20664 (2024), https://openaccess.thecvf.com/content/CVPR2024/html/Lu_Scaffold-GS_Structured_3D_Gaussians_for_View-Adaptive_Rendering_CVPR_2024_paper.html
18. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: representing scenes as neural radiance fields for view synthesis. *Commun. ACM* **65**(1), 99–106 (Dec 2021). <https://doi.org/10.1145/3503250>, <https://dl.acm.org/doi/10.1145/3503250>
19. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* **41**(4), 102:1–102:15 (Jul 2022). <https://doi.org/10.1145/3528223.3530127>, <https://dl.acm.org/doi/10.1145/3528223.3530127>
20. Qin, M., Li, W., Zhou, J., Wang, H., Pfister, H.: LangSplat: 3D Language Gaussian Splatting. pp. 20051–20060 (2024), https://openaccess.thecvf.com/content/CVPR2024/html/Qin_LangSplat_3D-Language_Gaussian-Splatting_CVPR_2024_paper.html
21. Ren, K., Jiang, L., Lu, T., Yu, M., Xu, L., Ni, Z., Dai, B.: Octree-GS: Towards Consistent Real-time Rendering with LOD-Structured 3D Gaussians. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 1–15 (2025). <https://doi.org/10.1109/TPAMI.2025.3568201>, <https://ieeexplore.ieee.org/abstract/document/10993308>
22. Shen, J., Qian, Y., Zhan, X.: LOD-GS: Achieving Levels of Detail using Scalable Gaussian Soup. pp. 671–680 (2025), https://openaccess.thecvf.com/content/CVPR2025/html/Shen_LOD-GS_Achieving_Levels_of_Detail_using_Scalable_Gaussian_Soup_CVPR_2025_paper.html
23. Svitov, D., Morerio, P., Agapito, L., Del Bue, A.: Billboard Splatting (BB-Splat): Learnable Textured Primitives for Novel View Synthesis. pp. 25029–25039 (2025), https://openaccess.thecvf.com/content/ICCV2025/html/Svitov_BillBoard_Splatting_BB_Splat_Learnable_Textured_Primitives_for_Novel_View_Synthesis_ICCV_2025_paper.html
24. Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: DreamGaussian: Generative Gaussian Splatting for Efficient 3D Content Creation (Mar 2024). <https://doi.org/10.48550/arXiv.2309.16653>, <http://arxiv.org/abs/2309.16653>, arXiv:2309.16653 [cs]
25. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (Apr 2004). <https://doi.org/10.1109/TIP.2003.819861>, <https://ieeexplore.ieee.org/abstract/document/1284395>
26. Ye, M., Danelljan, M., Yu, F., Ke, L.: Gaussian Grouping: Segment and Edit Anything in 3D Scenes. In: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O.,

- Sattler, T., Varol, G. (eds.) *Computer Vision – ECCV 2024*. pp. 162–179. Springer Nature Switzerland, Cham (2025). https://doi.org/10.1007/978-3-031-73397-0_10
27. Ye, Z., Li, W., Liu, S., Qiao, P., Dou, Y.: AbsGS: Recovering Fine Details in 3D Gaussian Splatting. In: *Proceedings of the 32nd ACM International Conference on Multimedia*. pp. 1053–1061. MM ’24, Association for Computing Machinery, New York, NY, USA (Oct 2024). <https://doi.org/10.1145/3664647.3681361>, <https://dl.acm.org/doi/10.1145/3664647.3681361>
 28. Yi, T., Fang, J., Wang, J., Wu, G., Xie, L., Zhang, X., Liu, W., Tian, Q., Wang, X.: GaussianDreamer: Fast Generation from Text to 3D Gaussians by Bridging 2D and 3D Diffusion Models. pp. 6796–6807 (2024), https://openaccess.thecvf.com/content/CVPR2024/html/Yi_GaussianDreamer_Fast_Generation_from_Text_to_3D_Gaussians_by_Bridging_CVPR_2024_paper.html
 29. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. pp. 586–595 (2018), https://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_The_Unreasonable_Effectiveness_CVPR_2018_paper.html
 30. Zhang, Z., Huang, B., Jiang, H., Zhou, L., Xiang, X., Shen, S.: Quadratic Gaussian Splatting: High Quality Surface Reconstruction with Second-order Geometric Primitives. pp. 28260–28270 (2025), https://openaccess.thecvf.com/content/ICCV2025/html/Zhang_Quadratic_Gaussian_Splatting_High_Quality_Surface_Reconstruction_with_Second-order_Geometric_ICCV_2025_paper.html
 31. Zhou, J., Huang, Y., Dai, W., Zou, J., Zheng, Z., Kan, N., Li, C., Xiong, H.: 3DGabSplat: 3D Gabor Splatting for Frequency-adaptive Radiance Field Rendering (Aug 2025). <https://doi.org/10.48550/arXiv.2508.05343>, <http://arxiv.org/abs/2508.05343>, arXiv:2508.05343 [cs]
 32. Zwicker, M., Pfister, H., van Baar, J., Gross, M.: Surface splatting. In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. pp. 371–378. SIGGRAPH ’01, Association for Computing Machinery, New York, NY, USA (Aug 2001). <https://doi.org/10.1145/383259.383300>, <https://dl.acm.org/doi/10.1145/383259.383300>