# UNIVERSAL BETA SPLATTING

## **Anonymous authors**

Paper under double-blind review

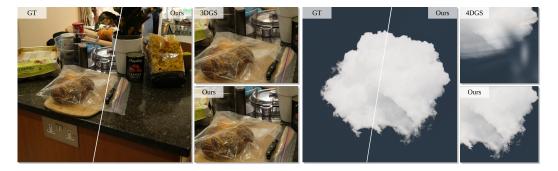


Figure 1: Visualization of UBS rendering quality. For static real-world scenes (left), UBS achieves superior rendering of reflective and specular materials compared to 3DGS (Kerbl et al., 2023). For dynamic volumetric scenes (right), UBS maintains high visual fidelity in complex spatio-temporal scenarios where 4DGS (Yang et al., 2024a) produces blurring artifacts.

#### ABSTRACT

We introduce Universal Beta Splatting (UBS), a unified framework that generalizes 3D Gaussian Splatting to N-dimensional anisotropic Beta kernels for explicit radiance field rendering. Unlike fixed Gaussian primitives, Beta kernels enable controllable dependency modeling across spatial, angular, and temporal dimensions within a single representation. Our unified approach captures complex light transport effects, handles anisotropic view-dependent appearance, and models scene dynamics without requiring auxiliary networks or specific color encodings. UBS maintains backward compatibility by approximating to Gaussian Splatting as a special case, guaranteeing plug-in usability and lower performance bounds. The learned Beta parameters naturally decompose scene properties into interpretable without explicit supervision: spatial (surface vs. texture), angular (diffuse vs. specular), and temporal (static vs. dynamic). Our CUDA-accelerated implementation achieves real-time rendering while consistently outperforming existing methods across static, view-dependent, and dynamic benchmarks, establishing Beta kernels as a scalable universal primitive for radiance field rendering.

## 1 Introduction

Real-time photorealistic rendering of complex scenes remains a fundamental challenge in computer vision and graphics. While Neural Radiance Fields (NeRF) (Mildenhall et al., 2020) achieve exceptional visual quality through continuous volumetric representations, their reliance on dense ray sampling imposes prohibitive computational costs. 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) addresses efficiency through explicit primitive-based rendering, achieving real-time performance with competitive fidelity. However, Gaussian kernels limit representational capacity: their fixed bell-shaped profiles struggle with sharp boundaries, while view-dependent effects require auxiliary spherical harmonic encodings that fragment the representation. Dynamic scene extensions (Yang et al., 2024b; Luiten et al., 2024) further increase complexity through additional deformation networks.

Recent advances partially address these limitations through specialized designs. Deformable Beta Splatting (DBS) (Liu et al., 2025) improves geometric fidelity using spatially adaptive Beta kernels,

but remains limited to 3D spatial dimensions and still requires separate Spherical Beta functions for view-dependent color. N-dimensional Gaussian extensions like 6DGS (Gao et al., 2025a) and 7DGS (Gao et al., 2025b) incorporate view and temporal dimensions through conditional distributions, enabling them to capture scattering effects and dynamic changes. However, these approaches remain constrained by the Gaussian kernel's symmetric profile across dimensions—preventing independent control of spatial sharpness, angular specularity, and temporal dynamics.

The core insight is that different scene properties require different kernel behaviors. Spatial geometry benefits from adaptive sharpness—flat for surfaces, peaked for textures. Angular appearance spans from diffuse to specular responses. Temporal dynamics range from static (constant support) to rapid motion (localized activation). While Gaussian kernels couple all dimensions with identical bell-shaped profiles, real scenes exhibit independent variations across space, angle, and time.

We introduce Universal Beta Splatting (UBS), a unified framework that generalizes radiance field rendering to N-dimensional anisotropic Beta kernels. Unlike fixed Gaussian profiles, Beta kernels provide per-dimension shape control through learnable parameters, enabling each dimension to adopt its optimal form. This allows a single primitive type to simultaneously model spatial geometry, view-dependent appearance, and temporal dynamics while maintaining computational efficiency.

Our approach addresses three technical challenges: (1) spatial-orthogonal Cholesky parameterization preserving orthonormal spatial structure while enabling flexible cross-dimensional correlations, (2) Beta-modulated conditional slicing transforming high-dimensional primitives into renderable 3D representations with dimension-specific anisotropy, and (3) fully accelerated CUDA implementation ensuring real-time performance.

Crucially, UBS maintains backward compatibility: when Beta parameters equal zero, the kernel approximates to Gaussian, recovering 3DGS, 6DGS, or 7DGS depending on dimensionality. This ensures seamless drop-in replacement with guaranteed performance lower bounds. Beyond rendering improvements, learned Beta parameters naturally decompose scenes into interpretable components without explicit supervision: spatial parameters separate geometry from textures, angular parameters distinguish diffuse from specular, and temporal parameters isolate static from dynamic elements.

We validate UBS across static, view-dependent, and dynamic benchmarks, achieving state-of-the-art visual quality—improving PSNR by up to +8.27 dB on static scenes and +2.78 dB on dynamic scenes compared to corresponding Gaussian baselines. Our contributions are:

- Universal Beta Splatting: A unified N-dimensional representation with per-dimension shape control, enabling simultaneous modeling of spatial, angular, and temporal properties through anisotropic Beta kernels with spatial-orthogonal Cholesky parameterization.
- **Efficient CUDA implementation:** Fully accelerated differentiable rendering with custom kernels achieving real-time performance.
- Interpretable scene decomposition: Emergent separation of geometric, appearance, and motion components through learned Beta parameters without supervision.
- Backward compatibility: Reduction to approximate existing methods as special cases, ensuring
  plug-in usability with performance lower bounds while enabling substantial improvements.

## 2 RELATED WORK

**Neural Radiance Fields and Gaussian Splatting.** Neural Radiance Fields (NeRF) (Mildenhall et al., 2020) revolutionized novel view synthesis through continuous volumetric representations. While extensions like Instant-NGP (Müller et al., 2022), Mip-NeRF (Barron et al., 2022), and Zip-NeRF (Barron et al., 2023) improved training speed and quality, NeRF methods remain limited by slow inference and inability to model off-ray transport phenomena.

3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) achieved real-time rendering through explicit anisotropic Gaussian primitives. Extensions address anti-aliasing (Yu et al., 2024; Yan et al., 2024), compression (Lee et al., 2024; Morgenstern et al., 2024), and semantic applications (Shi et al., 2024; Zhou et al., 2024). However, 3DGS's spherical harmonic encoding struggles with complex view-dependent effects. To better capture view dependence, *N*-dimensional Gaussians (N-DG) (Diolatzis et al., 2024) extended to higher dimensions, refined by 6D Gaussian Splatting (6DGS) (Gao et al.,

2025a) using conditional slicing. While elegant, Gaussian distributions cannot produce sharp angular cutoffs without many small primitives, limiting efficiency for specular materials.

Alternative Kernel Designs. Recent work explores kernels beyond Gaussians: GES (Hamdi et al., 2024) uses generalized exponentials for edge representation; Deformable Beta Splatting (DBS) (Liu et al., 2025) introduces adaptive Beta kernels with learnable shape parameters; 3D Convex Splatting (Held et al., 2025) handles flat surfaces; 3D Student Splatting (Zhu et al., 2025) leverages heavier-tailed distributions; and Gabor Splats (Zhou et al., 2025) captures high-frequency details. While these improve specific aspects, they remain confined to 3D spatial modeling without addressing view-dependent or temporal dimensions.

**Dynamic Neural Radiance Fields and Gaussian Splatting.** Early NeRF extensions modeled dynamics through deformation fields: D-NeRF (Pumarola et al., 2021) learns canonical-to-deformed mappings, Nerfies (Park et al., 2021a) handles non-rigid deformations, HyperNeRF (Park et al., 2021b) models topological changes, and HexPlane (Cao & Johnson, 2023) uses factorized representations.

For Gaussian Splatting, 4D Gaussian Splatting (4DGS) (Yang et al., 2024a) extends primitives temporally for real-time dynamic rendering. Dynamic 3D Gaussians (Luiten et al., 2024) and related works (Wu et al., 2024; Yang et al., 2024b) combine canonical representations with learned deformations. 7D Gaussian Splatting (7DGS) (Gao et al., 2025b) unifies spatial, temporal, and angular dimensions without separate deformation networks. Despite this unification, Gaussian's symmetric profile enforces smooth transitions—a single primitive cannot simultaneously represent sharp spatial edges, abrupt temporal motion, and narrow specular highlights.

Our Universal Beta Splatting addresses these limitations through controllable Beta kernels with per-dimension parameters, enabling dimension-specific adaptation and computational efficiency.

#### 3 Preliminary

**3D Gaussian Splatting (3DGS).** 3DGS (Kerbl et al., 2023) represents scenes as collections of anisotropic 3D Gaussians parameterized by position  $\mu \in \mathbb{R}^3$ , opacity  $\alpha$ , rotation quaternion q, scale s, and color features f. The covariance matrix  $\Sigma = RSS^{\top}R^{\top}$  is constructed from rotation and scale matrices. Rendering uses alpha-composited 2D projections:

$$C(\mathbf{x}) = \sum_{i=1}^{N} c_i \alpha_i \mathcal{G}(r_i(\mathbf{x})) \prod_{j=1}^{i-1} (1 - \alpha_j \mathcal{G}(r_j(\mathbf{x}))),$$
(1)

where  $\mathcal{G}(r)$  is the 2D Gaussian and  $r_i(x)$  is the Mahalanobis distance.

**Deformable Beta Splatting (DBS).** DBS (Liu et al., 2025) replaces Gaussians with adaptive Beta kernels:

$$\mathcal{B}(x;b) = (1-x)^{\beta(b)}, \quad \beta(b) = 4\exp(b), \quad x \in [0,1], \tag{2}$$

where b controls kernel shape (negative: flat surfaces, positive: peaked details, zero: Gaussian-like).

N-Dimensional Gaussian. 6DGS (Gao et al., 2025a) and 7DGS (Gao et al., 2025b) extend 3DGS through conditional distributions. 6DGS incorporates viewing direction into a 6D joint distribution:

$$X = \begin{pmatrix} X_p \\ X_d \end{pmatrix} \sim \mathcal{N} \left( \begin{pmatrix} \boldsymbol{\mu}_p \\ \boldsymbol{\mu}_d \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Sigma}_p & \boldsymbol{\Sigma}_{pd} \\ \boldsymbol{\Sigma}_{pd}^{\dagger} & \boldsymbol{\Sigma}_d \end{pmatrix} \right). \tag{3}$$

Conditioning on view d yields conditional covariance  $\Sigma_{\text{cond}} = \Sigma_p - \Sigma_{pd} \Sigma_d^{-1} \Sigma_{pd}^{\top}$ , conditional mean  $\mu_{\text{cond}} = \mu_p + \Sigma_{pd} \Sigma_d^{-1} (d - \mu_d)$ , and modulated opacity  $\alpha_{\text{cond}} = \alpha \cdot \exp(-\lambda (d - \mu_d)^{\top} \Sigma_d^{-1} (d - \mu_d))$ .

7DGS adds temporal dimension  $X_t$  for dynamic scenes, with conditioning on both time and view producing spatio-temporal-angular modulation. These methods establish high-dimensional representations but remain limited by fixed Gaussian profiles that cannot independently control cross-dimensional dependencies—a limitation UBS addresses through adaptive Beta kernels.

## 4 METHOD

In this section, we introduce Universal Beta Splatting (UBS), a unified framework that extends explicit radiance field rendering to N-dimensional anisotropic Beta kernels.

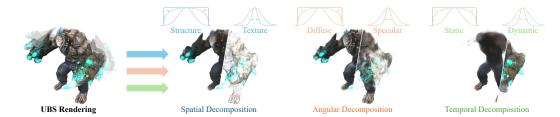


Figure 2: Visualization of decomposition. Our learned Beta parameters provide interpretable scene decomposition across both spatial, angular, and temporal dimension without explicit supervision.

#### 4.1 N-DIMENSIONAL BETA KERNEL

The core innovation of UBS is the N-dimensional anisotropic Beta kernel that models the universal radiance function:

$$\sigma(\boldsymbol{x}, \boldsymbol{q}) = B(\boldsymbol{x}, \boldsymbol{q}; \, \boldsymbol{\mu}, \, \boldsymbol{\Sigma}, \, \boldsymbol{b}), \tag{4}$$

where  $\boldsymbol{x} \in \mathbb{R}^3$  represents spatial coordinates, the query  $\boldsymbol{q} \in \mathbb{R}^{N-3}$  encodes additional dimensions (view direction, time, etc.),  $\boldsymbol{\mu} \in \mathbb{R}^N$  is the mean vector,  $\boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}$  is the covariance matrix, and  $\boldsymbol{b} \in \mathbb{R}^{N-2}$  controls the Beta shape parameters across dimensions (Betas are shared in spatial domain). The shape parameters  $\boldsymbol{b}$  are transformed to ensure positive Beta exponents:  $\beta_i = 4\exp(b_i)$ , enabling each dimension to adopt optimal kernel forms—from flat (negative  $b_i$ ) to peaked (positive  $b_i$ ).

In addition to the primitive shape, each kernel primitive also carries an opacity  $o \in [0,1]$  and an RGB color  $c \in [0,1]^3$ . Crucially, because the Beta kernel can natively incorporate view and temporal dimensions through q, no auxiliary color encoding is required to model view-dependent appearance or dynamics. This unified representation significantly reduces the per-primitive parameter count: for static scenes, UBS achieves a 41% parameter reduction compared to 3DGS, and for dynamic scenes, UBS uses 73% fewer parameters than 4DGS, as detailed in Table 4 in the Appendix. In contrast, existing methods rely on auxiliary encodings: 3DGS, 6DGS, and 7DGS employ 48-parameter spherical harmonics, 4DGS requires 144-parameter 4D spherical harmonics, and DBS needs additional spherical Beta functions for color modeling.

To maintain valid spatial geometry and guarantee backward compatibility while preserving high-dimensional expressiveness, we partition the N-dimensional space as:

$$\mu = \begin{pmatrix} \mu_x \\ \mu_q \end{pmatrix}, \quad \Sigma = \begin{pmatrix} \Sigma_x & \Sigma_{xq} \\ \Sigma_{qx} & \Sigma_q \end{pmatrix}, \tag{5}$$

where subscript x denotes spatial components and q denotes the additional dimensions. This partitioning allows conditioning on non-spatial dimensions to obtain renderable 3D representations.

#### 4.2 Spatial-Orthogonal Cholesky Parameterization

While mean vectors, beta shapes, and colors are straightforward to parameterize, the covariance matrix requires additional care as it must be symmetric and positive semi-definite. In 3DGS, this is achieved by factorizing the covariance as  $\Sigma = R \operatorname{diag}(s)^2 R^{\top}$ , where R is a rotation matrix parameterized by quaternions and s are per-axis scales. However, quaternions do not generalize to higher dimensions (N > 3). N-dimensional Gaussian methods therefore resort to Cholesky factorization  $\Sigma = LL^{\top}$ , which loses explicit orthogonality and often leads to suboptimal solutions in the spatial subspace.

To address this limitation, we introduce the *Spatial-Orthogonal Cholesky* parameterization that preserves rotation-scale structure in the spatial subspace while enabling flexible cross-dimensional correlations. For N=3+C dimensions, we construct:

$$L = \begin{pmatrix} R_x \operatorname{diag}(s_x) & \mathbf{0} \\ L_{qx} & \operatorname{diag}(s_q) \end{pmatrix}, \tag{6}$$

where  $R_x \in SO(3)$  is the spatial rotation computed via first-order Taylor approximation  $R_x \approx I + A_x$  of the matrix exponential (with  $A_x$  skew-symmetric),  $s_x$  are spatial scales,  $s_q$  are scales for other dimensions, and  $L_{qx}$  encodes cross-correlations. As shown in Table 5 (Appendix), this first-order

approximation provides nearly identical fidelity while greatly reducing computation. The covariance  $\Sigma = LL^{\top}$  maintains explicit 3D geometric structure while allowing expressive correlations across view and time.

#### 4.3 Beta-Modulated Conditional Slicing

In N-dimensional Gaussians, all dimensions are coupled: when the input view or time changes, the primitive is inevitably dragged along—it shifts in space, stretches in shape, and fades in opacity. To break this constraint, we introduce Beta-Modulated Conditional Slicing.

To render at query condition q, we condition the N-dimensional Beta kernel to obtain a 3D spatial representation. We introduce Beta modulation to the standard conditional Gaussian formulation:

Beta Parameters. We define per-dimension Beta exponents with spatial parameter sharing:  $\beta = [\beta_x, \beta_x, \beta_x, \beta_{q_1}, \dots, \beta_{q_C}]$  where  $\beta_x = 4 \exp(b_x)$  is shared across the first three spatial dimensions and  $\beta_{q_i} = \exp(b_{q_i})$  are individual parameters for the C non-spatial dimensions. We share the spatial Beta parameter across all three spatial dimensions to ensure that splatted kernels maintain consistent spatial geometry. This design choice prevents inconsistent spatial artifacts while still allowing per-dimension control in angular and temporal domains.

**Beta-Modulated Conditioning.** The conditional mean and covariance incorporate dimension-specific Beta modulation:

$$\boldsymbol{\mu}_{x|q} = \boldsymbol{\mu}_x + \boldsymbol{\Sigma}_{xq} \, \boldsymbol{\Sigma}_q^{-1} \operatorname{Diag}(\boldsymbol{\beta}_q) \, (\boldsymbol{q} - \boldsymbol{\mu}_q), \tag{7}$$

$$\Sigma_{x|q} = \Sigma_x - \Sigma_{xq} \, \Sigma_q^{-1} \operatorname{Diag}(\beta_q) \, \Sigma_{qx}, \tag{8}$$

where  $\mathrm{Diag}(\beta_q)=\mathrm{Diag}([\beta_{q_1},\ldots,\beta_{q_C}])$  applies dimension-specific Beta modulation to the non-spatial conditioning terms while preserving isotropic spatial structure.

**Product-Form Opacity Gate.** We adopt a factorized opacity with inherent multidimensional support:

$$d^{\text{raw}} = \Sigma_q^{-1}(q - \mu_q), \quad d = \max(\tanh(\frac{d^{\text{raw}}}{2}), \ 0) \in [0, 1)^C, \quad o(q) = o \prod_{i=1}^C (1 - d_i)^{4\beta_{q_i}}, \quad (9)$$

where  $\tanh$  and  $\max$  transformations ensure bounded distance measures  $d \in [0,1)^C$ , satisfying Beta kernel input definition. The product form allows independent control over each non-spatial dimension's contribution.

Final Beta-Splat Density. The conditioned 3D Beta kernel becomes:

$$\sigma(\boldsymbol{x}, \boldsymbol{q}) = B(\boldsymbol{x}; \, \boldsymbol{\mu}_{x|q}, \, \boldsymbol{\Sigma}_{x|q}, \, b_x) \cdot o(\boldsymbol{q}), \tag{10}$$

where the spatial Beta  $b_x$  controls the 3D kernel shape uniformly across all spatial dimensions.

### 4.4 Universal Compatibility and Interpretability

**Adaptive Specialization.** UBS naturally specializes for different applications through its query dimensions: (1) q = d (view direction, N = 6) for view-dependent rendering with anisotropic BRDF modeling, denoted as UBS-6D; (2) q = [t, d] (time + view, N = 7) for dynamic scenes capturing spatio-temporal-angular correlations, denoted as UBS-7D; and (3) seamless extension to higher dimensions for additional modalities like lighting or material properties.

Universal Compatibility. UBS provides backward compatibility by design, reducing to existing methods under specific parameter settings. For N=3 (pure spatial), UBS reduces to Deformable Beta Splatting with  $\sigma(\boldsymbol{x}) = B(\boldsymbol{x}; \boldsymbol{\mu}_x, \boldsymbol{\Sigma}_x, b_x) \cdot o$ , and further approximates to standard 3DGS when  $b_x=0$  (Gaussian limit). For higher dimensions with all Beta parameters zero ( $\boldsymbol{b}=0$ ), UBS approximates corresponding Gaussian methods: 6DGS for view-dependent rendering (N=6) and 7DGS for dynamic scenes (N=7). This guaranteed compatibility ensures UBS serves as a seamless drop-in replacement while providing performance lower bounds.

Interpretable Decomposition. Beyond compatibility, learned Beta parameters provide interpretable scene decomposition without explicit supervision. Spatial parameters distinguish geometry types: negative  $b_x$  produces flat kernels for smooth surfaces, while positive  $b_x$  creates peaked kernels for

fine textures. Angular parameters separate appearance models: negative  $b_d$  yields broad responses for diffuse surfaces, while positive  $b_d$  generates sharp peaks for specular reflections. Temporal parameters isolate motion patterns: negative  $b_t$  maintains broad support for static elements, while positive  $b_t$  provides localized activation for dynamic components, as illustrated in Figure 2. This decomposition enables applications like relighting and motion analysis without requiring separate encodings or auxiliary networks.

### 4.5 OPTIMIZATION AND IMPLEMENTATION

For primitive management, we adopt the kernel-agnostic MCMC optimization from DBS (Liu et al., 2025), which provides stable, distribution-preserving updates independent of kernel form.

**Loss Function.** We optimize with reconstruction losses and regularization:

$$\mathcal{L} = (1 - \lambda_{\text{SSIM}})\mathcal{L}_1 + \lambda_{\text{SSIM}}\mathcal{L}_{\text{SSIM}} + \lambda_o \sum_i |o_i| + \lambda_{\Sigma} \sum_i ||s_i||_1,$$
(11)

where opacity regularization  $\lambda_o$  ensures valid MCMC densification and scale penalty  $\lambda_{\Sigma}$  encourages primitive relocation. Spatial noise injection  $\mu_x \leftarrow \mu_x - \lambda_{lr} \nabla_{\mu} \mathcal{L} + \lambda_{\epsilon} \epsilon$  promotes exploration.

**Densification.** Cloned primitive opacity follows  $o' = 1 - \sqrt[N]{1-o}$  to preserve distribution. Under opacity regularization, adjusting opacity maintains density distribution regardless of kernel type.

**CUDA Implementation.** We provide fused CUDA kernels for Beta evaluation, conditional slicing, and spatial-orthogonal operations. The pipeline maintains compatibility with existing frameworks: *N*-dimensional kernels are conditioned to 3D then rasterized using standard splatting. This modular design enables real-time performance while allowing future integration with acceleration techniques like FlashGS (Feng et al., 2025) and extended camera models from 3DGUT (Wu et al., 2025).

## 5 EXPERIMENTS

### 5.1 EXPERIMENTAL SETUP

**Datasets.** We evaluate UBS across five diverse datasets covering both static and dynamic scenes:

- Static: we use (1) NeRF Synthetic (Mildenhall et al., 2020) with 8 synthetic scenes under controlled lighting for geometric reconstruction assessment; (2) Mip-NeRF 360 (Barron et al., 2022) sharing 9 unbounded real-world scenes challenging anti-aliasing and view synthesis; and (3) 6DGS-PBR (Gao et al., 2025a) featuring 7 physically-based scenes with complex view-dependent effects including volumetric scattering, translucent materials, subsurface scattering, and medical volumetrics.
- *Dynamic:* we evaluate on: (4) *D-NeRF* (Pumarola et al., 2021) with 8 synthetic sequences featuring various deformation types; and (5) 7DGS-PBR (Gao et al., 2025b) containing 6 physically-based dynamic scenes with complex spatio-temporal-angular correlations including cardiac motion, daylight transitions, animated volumetric effects, and translucent deformations.

**Evaluation Metrics.** We assess reconstruction quality using standard metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) (Wang et al., 2004), and Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al., 2018). Additionally, we report rendering speed (FPS) and training time to evaluate computational efficiency. All metrics are computed on held-out test views following standard evaluation protocols.

Implementation Details. UBS is implemented in PyTorch with custom CUDA kernels for Beta kernel evaluation and conditional slicing operations. We conduct on a single NVIDIA RTX 4090 24GB GPU for static and ablation experiments and a single NVIDIA Tesla V100 16GB GPU for dynamic experiments to be consistent 7DGS. Training employs the Adam optimizer for 30K iterations with learning rates of  $1.6 \times 10^{-4}$  for positions,  $5 \times 10^{-2}$  for opacity,  $5 \times 10^{-3}$  for scales, and  $1 \times 10^{-3}$  for other parameters. The regularization weights are set to  $\lambda_o = 0.01$ ,  $\lambda_\Sigma = 0.01$ , and  $\lambda_\epsilon = 1$ . For N-dimensional cases, we randomly initialize temporal and directional means ( $\mu_t$ ,  $\mu_d$ ) within [0, 1]. Beta parameters are initialized to zero, corresponding to the Gaussian limit for guaranteed convergence. For dynamic scenes, we use a batch size of 4 to maintain consistency with 4DGS and 7DGS baselines. Unlike these methods which evaluate on test sets every 500 iterations and report best results in their papers, we report at the final iteration (30K) to demonstrate convergence stability.

Table 1: Comparison with 3DGS and 6DGS on static benchmarks of NeRF Synthetic (Mildenhall et al., 2020), Mip-NeRF (Barron et al., 2022), 6DGS-PBR (Gao et al., 2025a), and with 4DGS and 7DGS on dynamic benchmarks of 7DGS-PBR (Gao et al., 2025b) and D-NeRF (Pumarola et al., 2021). 'Train' means time in mins.

<u>Schemma</u>			3DGS				6DGS				UBS-6D (Ours)			
Dataset		Scene	   DSNID↑		LPIPS↓	Train	   DSNID↑		LPIPS↓	Train	<u> </u>		LPIPS ↓	Train
			•				•				•		•	
	NeRF Synthetic	chair	35.60		0.010	4.8	35.55	0.986	0.011	19.6	36.72	0.990	0.009	8.7
		drums	26.28 35.49		0.037	4.1 3.0	26.63	0.951 0.984	0.038 0.015	15.4	27.19 36.90	0.960 0.990	0.031 0.009	8.1 8.2
		ficus hotdog	38.07		0.012 0.020	3.5	34.62 37.96	0.984	0.013	11.3 10.0	38.67	0.988	0.009	8.2 9.1
		lego	36.06		0.020	4.0	35.22	0.983		15.3	36.95	0.985	0.010	8.4
		materials	30.50		0.010	2.7	30.63	0.960	0.020	7.8	32.90	0.963	0.014	8.9
		mic	36.67		0.006	3.0	37.10	0.992	0.007	9.4	37.86	0.994	0.005	9.5
		ship	31.68		0.106	4.1	31.09	0.899	0.118	12.7	32.13	0.914	0.100	8.8
Si		avg	33.79	0.970	0.030	3.6	33.60	0.967	0.034	12.7	34.92	0.975	0.026	8.7
STATIC SCENES		bicycle	25.19	0.764	0.212	29.2	22.17	0.535	0.425	114.9	25.64	0.793	0.186	30.2
Ö		flowers	21.39		0.338	19.6	18.99	0.417	0.461	88.8	21.96	0.635	0.313	19.6
()	00	garden	27.27	0.862	0.109	30.4	26.71	0.843	0.142	164.9	28.10	0.883	0.094	36.8
Ĕ	36	stump	26.61	0.772	0.215	22.7	23.83	0.665	0.338	60.2	26.75	0.788	0.199	18.9
Τ	Mip-NeRF 360	treehill	22.47		0.329	20.1	21.51	0.556		94.0	23.57	0.684	0.280	21.9
S		room	31.49	0.917	0.221	18.2	31.34	0.904	0.256	51.4	32.81	0.940	0.173	21.9
		counter	28.96	0.906	0.202	17.6	29.17	0.880	0.253	43.5	31.39	0.937	0.150	27.8
		kitchen	31.40	0.926	0.127	21.5	30.48	0.907	0.158	71.3	32.64	0.940	0.107	26.6
		bonsai	32.05	0.940	0.206	15.7	32.91	0.935	0.228	48.4	35.05	0.962	0.154	22.7
		avg	27.43	0.814	0.218	21.7	26.35	0.738	0.298	81.93	28.66	0.840	0.184	25.2
	6DGS-PBR	bunny	29.19	0.985	0.074	22.8	37.41	0.991	0.052	24.1	45.68	0.996	0.023	32.1
		cloud	34.03	0.985	0.058	23.1	41.00	0.991	0.050	21.6	46.81	0.995	0.027	30.6
		explosion	27.06	0.953	0.097	11.3	41.61	0.990	0.031	18.3	44.54	0.994	0.018	32.4
		smoke	26.82	0.964	0.088	24.2	41.41	0.993	0.041	31.1	44.02	0.995	0.028	33.5
		suzanne	22.45	0.885	0.159	19.5	26.96	0.928	0.106	32.9	27.35	0.931	0.098	35.1
		dragon	26.53	0.812	0.123	12.0	34.95	0.933	0.044	16.3	38.06	0.974	0.034	23.7
		ct-scan	25.69	0.917	0.099	4.5	33.46	0.964	0.058	11.0	34.24	0.969	0.046	6.5
		avg	27.40	0.929	0.100	16.8	36.68	0.970	0.055	22.2	40.10	0.979	0.039	27.70
		_	4DGS				7DGS			UBS-7D (Ours)				
Data	aset	Scene	PSNR↑	$SSIM \!\!\uparrow$	$LPIPS\!\!\downarrow$	Train↓	PSNR↑	SSIM↑	$LPIPS\!\downarrow$	Train $\downarrow$	PSNR↑	$SSIM \!\!\uparrow$	LPIPS↓	Train↓
		heart1	27.23	0.949	0.046	103.0	34.66	0.983	0.023	114.2	37.44	0.990	0.013	38.7
	7DGS-PBR	heart2	25.09	0.919	0.085	103.4	30.99	0.959	0.057	384.6	31.99	0.966	0.043	33.5
	<u>-</u>	cloud	24.63	0.938	0.100	123.7	29.28	0.955	0.075	102.6	30.65	0.955	0.066	50.0
S	Ğ	dust	35.81	0.954	0.037	97.0	36.85	0.955	0.038	69.8	39.17	0.985	0.028	46.5
当	7	flame	29.25		0.068	113.7	31.64	0.937	0.062	74.1	31.70	0.963	0.057	53.7
Ħ		suzanne	24.41	0.912	0.127	222.5	27.09	0.941	0.072	193.9	27.12	0.934	0.069	136.8
DYNAMIC SCENES		avg	27.74	0.933	0.077	127.2	31.75	0.955	0.055	116.7	33.00	0.966	0.046	59.9
M _		b.balls	32.45		0.028	50.7	33.61	0.980		86.6	33.39	0.982	0.022	49.6
Ϋ́		h.warrior	33.69		0.069	35.3	32.30	0.931	0.088	31.3	33.81	0.934	0.086	36.3
×		hook	31.90		0.037	38.0	30.19	0.953	0.047	35.7	30.93	0.956	0.045	34.0
П		j.jacks	28.43	0.962	0.041	66.9	31.37	0.967	0.038	34.1	30.96	0.966	0.040	43.2
	-	lego	24.28		0.098	55.4	27.64	0.935		78.5	27.85	0.925	0.072	31.3
	Ι	mutant	38.64		0.009	39.1	39.21	0.992	0.008	42.5	40.75	0.994	0.006	60.0
		standup	39.00		0.009	34.4	38.42	0.987	0.014	33.5	39.08	0.991	0.009	41.1
		trex	28.35	0.973	0.026	100.7	29.94	0.978		63.4	30.05	0.975	0.023	41.6
		avg	32.09	0.963	0.040	52.6	32.84	0.965	0.039	50.7	33.35	0.965	0.038	42.1

#### 5.2 Comparison with the State of the Art

Table 1 compares our UBS with 3DGS (Kerbl et al., 2023), 4DGS (Yang et al., 2024a), 6DGS (Gao et al., 2025a), and 7DGS (Gao et al., 2025b) on both static and dynamic benchmarks.

**Static Scenes.** Our method UBS-6D consistently achieves superior reconstruction quality across all datasets while maintaining competitive training efficiency on three static scene benchmarks:

- *NeRF Synthetic:* UBS-6D achieves 34.92 dB PSNR, improving by +1.13 dB over 3DGS and +1.32 dB over 6DGS. Despite lacking strong view-dependent effects, the dataset benefits from Beta kernels' geometric adaptivity, with notable improvements on scenes requiring diverse surface representations (ficus: +1.41 dB over 3DGS, materials: +2.40 dB over 3DGS).
- *Mip-NeRF 360*: UBS-6D reaches 28.66 dB PSNR, surpassing 3DGS by +1.23 dB and 6DGS by +2.31 dB. The real-world unbounded scenes showcase UBS's effectiveness on complex lighting conditions, with substantial gains on indoor environments (counter: +2.43 dB over 3DGS, bonsai: +3.00 dB over 3DGS).

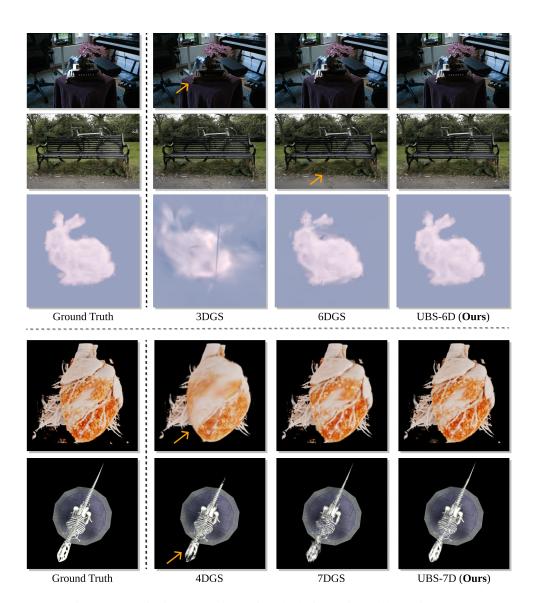


Figure 3: Qualitative comparison of methods for static and dynamic scenes.

• 6DGS-PBR: UBS-6D demonstrates the most significant improvements with 40.10 dB PSNR (+12.70 dB over 3DGS, +3.42 dB over 6DGS average). The physically-based scenes with volumetric scattering and translucency particularly benefit from per-dimension Beta control: volumetric effects like cloud improve by +5.81 dB over 6DGS, while complex materials like bunny achieve a remarkable +8.27 dB improvement over 6DGS.

Overall, UBS-6D demonstrates exceptional performance on challenging view-dependent scenarios in *6DGS-PBR*. In addition, UBS-6D establishes new state-of-the-art results on *NeRF Synthetic* and *Mip-NeRF 360* shown in Table 3 in Appendix.

**Dynamic Scenes.** We compare our method UBS-7D with 4DGS and 7DGS on 2 dynamic benchmarks. Note that we consider the base 7DGS without MLP networks for adaptive Gaussian refinement (AGR) to ensure fair comparison and isolate the benefits of our primitive design from auxiliary networks:

- 7DGS-PBR: UBS-7D achieves 33.00 dB PSNR, outperforming 4DGS by +5.26 dB and 7DGS by +1.25 dB. The physically-based dynamic scenes with spatio-temporal-angular correlations show substantial improvements, particularly on complex volumetric dynamics (dust: +2.32 dB over 7DGS) and cardiac motion sequences (heart1: +2.78 dB over 7DGS).
- *D-NeRF*: UBS-7D reaches 33.35 dB PSNR (+1.26 dB over 4DGS, +0.51 dB over 7DGS). The synthetic monocular sequences benefit from Beta kernels' ability to model diverse temporal

Table 2: Ablation study of UBS across representative static and dynamic scenes.

		Static Scenes						Dynamic Scenes							
	ficus			bonsai		-	∥ clow		oud		lego		Ī	Mean	
	PSNR1	Train↓	FPS↑	PSNR↑	Train↓	FPS↑	PSNR1	`Train↓	$FPS\!\!\uparrow\! $	PSNR↑	Train↓	$\text{FPS}{\uparrow}$	PSNR↑	Train↓	FPS↑
skew-sym. w/o Cholesky	36.34	19.0	207.2	34.24	71.1	50.1	28.49	14.7	142.5	28.42	13.1	241.9	31.87	29.5	160.4
Cholesky w/o spatial-ortho.	35.76	22.8	119.2	34.21	75.0	40.4	28.41	16.3	162.2	28.02	15.0	175.4	31.60	32.3	124.3
w/o Beta condition	36.72	13.7	217.2	34.30	49.3	53.9	28.26	13.7	160.0	27.22	11.4	198.9	31.63	22.0	157.5
w/o Beta opacity	33.76	17.9	213.0	33.17	62.1	57.3	21.93	11.3	260.3	25.68	11.9	209.5	28.64	25.8	185.0
full	36.92	20.7	188.4	34.62	72.9	48.9	28.61	16.1	158.3	28.47	13.1	227.6	32.16	30.7	155.8
full w/ CUDA	36.96	9.3	262.4	34.65	24.1	78.5	28.58	10.2	184.8	28.55	7.8	259.1	32.19	12.9	196.2

behaviors, with notable gains on scenes requiring both spatial detail and temporal precision (mutant: +1.54 dB over 7DGS, h.warrior: +1.51 dB over 7DGS).

The consistent improvements across PSNR, SSIM, and LPIPS metrics confirm that UBS's enhanced representational capacity translates to comprehensive quality gains. Figure 3 demonstrates that UBS achieves better visual quality in both static and dynamic scenes compared with baseline methods.

**Training time.** Beyond reconstruction quality, UBS demonstrates superior computational efficiency through its streamlined primitive design and optimized CUDA implementation. For static scenes, UBS-6D achieves significantly faster training than 6DGS on NeRF Synthetic (8.7 vs 12.7 minutes, 31.5% reduction) and Mip-NeRF 360 (25.2 vs 81.9 minutes, 69.2% reduction). On 6DGS-PBR, UBS-6D requires longer training time (27.7 vs 22.2 minutes) due to our use of higher primitive counts (300K vs 68K) to fully exploit Beta kernels' representational capacity for complex volumetric effects. For dynamic scenes, UBS-7D delivers substantial training time reductions of 48.7% on 7DGS-PBR (59.9 vs 116.7 minutes) and 17.0% on D-NeRF (42.1 vs 50.7 minutes), both compared to 7DGS. These efficiency gains stem from our unified Beta kernel approach, which eliminates auxiliary spherical harmonic encodings while enabling efficient CUDA-accelerated conditional slicing operations.

#### 5.3 ABLATION STUDY

We conduct ablations on representative samples of static scenes in 6DGS-PBR (Gao et al., 2025a) dataset (ficus, bonsai) and of dynamic scenes in 7DGS-PBR (Gao et al., 2025b) dataset (cloud, lego) to validate our design choices and efficient CUDA implementation (Table 2).

**Spatial-Orthogonal Parameterization.** Pure rotation-scale parameterization and full Cholesky without spatial orthogonality reduce quality by 0.29 dB and 0.56 dB, respectively, confirming the importance of preserving the spatial SO(3) structure and maintaining flexibility of other dimensions. **Beta-Modulated Conditioning.** Removing Beta modulation decreases performance by 0.53 dB, with larger drops on dynamic scenes (lego: -1.25 dB), validating dimension-specific shape control for decoupling spatial-angular-temporal responses. **Beta-Modulated Opacity.** Disabling Beta-modulated opacity causes the most significant degradation (-3.52 dB average), particularly on volumetric scenes (cloud: -6.68 dB), proving the product-form opacity crucial for anisotropic falloff. **CUDA Acceleration.** Custom CUDA kernels reduce training time by 58% (12.9 vs 30.7 minutes) and improve rendering speed by 26% (196.2 vs 155.8 FPS), enhancing real-time performance through fused Beta evaluation and conditional slicing operations.

These studies confirm each component's contribution to UBS's performance, with the complete model showing consistent gains across diverse scenarios while maintaining computational efficiency.

#### 6 CONCLUSION

We presented Universal Beta Splatting (UBS), a unified framework extending radiance field rendering to N-dimensional anisotropic Beta kernels. By replacing fixed Gaussian profiles with adaptive per-dimension shape control, UBS simultaneously models spatial geometry, view-dependent appearance, and temporal dynamics within a single primitive. Our spatial-orthogonal parameterization and Beta-modulated conditional slicing enable independent dimensional control while preserving geometric consistency. Experiments demonstrate significant improvements—up to  $+8.27~\mathrm{dB}$  PSNR on volumetric scenes and 48.7% faster training on dynamic sequences—while maintaining backward compatibility by approximating Gaussian methods. The learned Beta parameters provide interpretable scene decomposition without explicit supervision, and our CUDA implementation achieves real-time rendering, establishing Beta kernels as practical universal primitives for radiance fields.

## REFERENCES

- Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *CVPR*, 2022. 2, 6, 7, 12, 13
- Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. In *ICCV*, 2023. 2, 12
- Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In CVPR, 2023. 3
- Stavros Diolatzis, Tobias Zirr, Alexander Kuznetsov, Georgios Kopanas, and Anton Kaplanyan. N-dimensional gaussians for fitting of high dimensional functions. In *SIGGRAPH*, 2024. 2
- Guofeng Feng, Siyan Chen, Rong Fu, Zimu Liao, Yi Wang, Tao Liu, Zhilin Pei, Hengjie Li, Xingcheng Zhang, and Bo Dai. Flashgs: Efficient 3d gaussian splatting for large-scale and high-resolution rendering. In *CVPR*, 2025. 6
- Zhongpai Gao, Benjamin Planche, Meng Zheng, Anwesa Choudhuri, Terrence Chen, and Ziyan Wu. 6dgs: Enhanced direction-aware gaussian splatting for volumetric rendering. In *ICLR*, 2025a. 2, 3, 6, 7, 9
- Zhongpai Gao, Benjamin Planche, Meng Zheng, Anwesa Choudhuri, Terrence Chen, and Ziyan Wu. 7dgs: Unified spatial-temporal-angular gaussian splatting. In *ICCV*, 2025b. 2, 3, 6, 7, 9
- Abdullah Hamdi, Luke Melas-Kyriazi, Jinjie Mai, Guocheng Qian, Ruoshi Liu, Carl Vondrick, Bernard Ghanem, and Andrea Vedaldi. Ges: Generalized exponential splatting for efficient radiance field rendering. In *CVPR*, 2024. 3, 12
- Jan Held, Renaud Vandeghen, Abdullah Hamdi, Adrien Deliege, Anthony Cioppa, Silvio Giancola, Andrea Vedaldi, Bernard Ghanem, and Marc Van Droogenbroeck. 3d convex splatting: Radiance field rendering with 3d smooth convexes. In *CVPR*, 2025. 3
- Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *SIGGRAPH*, 2024. 12
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. In *SIGGRAPH*, 2023. 1, 2, 3, 7, 12
- Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Jeff Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3d gaussian splatting as markov chain monte carlo. In *NeurIPS*, 2024. 12
- Joo Chan Lee, Daniel Rho, Xiangyu Sun, Jong Hwan Ko, and Eunbyung Park. Compact 3d gaussian representation for radiance field. In *CVPR*, 2024. 2
- Rong Liu, Dylan Sun, Meida Chen, Yue Wang, and Andrew Feng. Deformable beta splatting. In *SIGGRAPH*, 2025. 1, 3, 6, 12
- Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *3DV*, 2024. 1, 3
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2, 6, 7, 12
- Wieland Morgenstern, Florian Barthel, Anna Hilsmann, and Peter Eisert. Compact 3d scene representation via self-organizing gaussian grids. In *ECCV*, 2024. 2
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. In *SIGGRAPH*, 2022. 2, 12
- Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, 2021a. 3

- Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. In *SIGGRAPH*, 2021b. 3
  - Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 3, 6, 7
  - Jin-Chuan Shi, Miao Wang, Hao-Bin Duan, and Shao-Hua Guan. Language embedded 3d gaussians for open-vocabulary scene understanding. In *CVPR*, 2024. 2
  - Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. In *TIP*, 2004. 6
  - Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *CVPR*, 2024.
  - Qi Wu, Janick Martinez Esturo, Ashkan Mirzaei, Nicolas Moenne-Loccoz, and Zan Gojcic. 3dgut: Enabling distorted cameras and secondary rays in gaussian splatting. In *CVPR*, 2025. 6
  - Zhiwen Yan, Weng Fei Low, Yu Chen, and Gim Hee Lee. Multi-scale 3d gaussian splatting for anti-aliased rendering. In *CVPR*, 2024. 2
  - Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. In *ICLR*, 2024a. 1, 3, 7
  - Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *CVPR*, 2024b. 1, 3
  - Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *CVPR*, 2024. 2, 12
  - Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 6
  - Junyu Zhou, Yuyang Huang, Wenrui Dai, Junni Zou, Ziyang Zheng, Nuowen Kan, Chenglin Li, and Hongkai Xiong. 3dgabsplat: 3d gabor splatting for frequency-adaptive radiance field rendering. In *ACM MM*, 2025. 3
  - Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suya You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *CVPR*, 2024. 2
  - Jialin Zhu, Jiangbei Yue, Feixiang He, and He Wang. 3d student splatting and scooping. In *CVPR*, 2025. 3

## A APPENDIX

#### A.1 COMPARISON WITH STATE-OF-THE-ART METHODS

Table 3 presents a comprehensive comparison of UBS-6D with state-of-the-art neural rendering methods on the Mip-NeRF360 (Barron et al., 2022) and NeRF Synthetic (Mildenhall et al., 2020) benchmarks. Our method achieves the highest PSNR on both datasets, demonstrating the effectiveness of adaptive Beta kernels over fixed Gaussian profiles. The consistent improvements across both synthetic and real-world scenes validate the universal applicability of our Beta kernel framework.

Table 3: Comparison with SOTA methods on Mip-NeRF360 (Barron et al., 2022) and NeRF Synthetic (Mildenhall et al., 2020). We adopt the original number from corresponding papers when available. A dash ("-") indicates results that were not reported.

	Methods	M	ip-NeRF3	60	NeRF Synthetic			
	- Treate as	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	
cit	Instant-NGP (Müller et al., 2022)	25.51	0.684	0.398	33.18	0.959	0.055	
implicit	Mip-NeRF360 (Barron et al., 2022)	27.69	0.792	0.237	33.25	0.962	0.039	
Ē.	Zip-NeRF (Barron et al., 2023)	28.54	0.828	0.189	33.10	0.971	0.031	
	3DGS (Kerbl et al., 2023)	27.20	0.815	0.214	33.31	0.969	0.037	
	GES (Hamdi et al., 2024)	26.91	0.794	0.250	-	-	-	
÷	2DGS (Huang et al., 2024)	27.04	0.805	0.297	33.07	-	-	
ij	Mip-Splatting (Yu et al., 2024)	27.79	0.827	0.203	33.33	0.969	0.039	
explicit	3DGS-MCMC (Kheradmand et al., 2024)	28.29	0.840	0.210	33.80	0.970	0.040	
	DBS (Liu et al., 2025)	28.60	0.844	0.182	34.64	0.973	0.028	
	UBS-6D ( <b>Ours</b> )	28.66	0.840	0.184	34.92	0.975	0.026	

### A.2 PER-PRIMITIVE PARAMETER ANALYSIS

Table 4 provides a detailed breakdown of parameter counts per primitive across different splatting methods. UBS achieves significant parameter efficiency by eliminating the need for auxiliary color encodings through its native high-dimensional representation.

For static scenes, 3DGS requires 59 parameters per primitive: 11 for geometry (pos: 3, scale: 3, rotation: 4, opacity: 1) and 48 for spherical harmonic color encoding (16 coefficients  $\times$  3 RGB channels). 6DGS increases this to 76 parameters by expanding the covariance representation to 28 parameters while maintaining the same 48-parameter SH encoding. In contrast, UBS-6D requires only 35 parameters: 32 for N-dimensional geometry (pos: 3, dir: 3, spatial-orthogonal covariance: 21, beta shapes: 4, opacity: 1) and just 3 for direct RGB color, achieving a 41% parameter reduction compared to 3DGS.

For dynamic scenes, the parameter efficiency becomes even more pronounced. 4DGS requires 161 parameters per primitive: 17 for 4D geometry (pos: 3, time: 1, 4D-scale: 4, 4D-rotation: 8, opacity: 1) and 144 for temporal spherical harmonics (48 coefficients × 3 RGB channels). 7DGS reduces this to 84 parameters: 36 for geometry (pos: 3, time: 1, dir: 3, 7D covariance: 28, opacity: 1) and 48 for standard SH encoding. UBS-7D achieves the most compact representation with only 44 parameters: 41 for unified 7D geometry (pos: 3, time: 1, dir: 3, spatial-orthogonal covariance: 28, beta shapes: 5, opacity: 1) and 3 for direct RGB color, representing a 73% parameter reduction compared to 4DGS.

This dramatic parameter efficiency stems from UBS's ability to natively encode view-dependent and temporal variations within the kernel itself, eliminating the need for separate spherical harmonic or

Table 4: Per-primitive parameter count comparison across static and dynamic splatting methods. UBS significantly reduces parameter counts compared to existing NDGS methods: UBS-6D is 59% of 3DGS and UBS-7D is 27% of 4DGS.

Comparison		Static Solution	s	Dynamic Solutions				
,	3DGS	6DGS	UBS-6D (Ours)	4DGS	7DGS	UBS-7D (Ours)		
Parameter (Geom. + Color) Ratio	$11 + 48 = 59$ $1 \times$	28 + 48 = 76 $1.29 \times$	$\begin{array}{ c c c c } 32 + 3 = 35 \\ 0.59 \times \end{array}$	17 + 144 = 161 1×	$\begin{vmatrix} 36 + 48 = 84 \\ 0.52 \times \end{vmatrix}$	$ \begin{array}{ c c c c } \hline 41 + 3 = 44 \\ 0.27 \times \end{array} $		

Table 5: Comparison of rotation parameterizations for covariance matrix construction on the bicycle scene for UBS-3D. Our spatial-orthogonal approach balances performance, efficiency, and parameter count.

Covariance Matrix Parameterization	Training Time	PSNR	# Parameters (N>3)	# Parameters (N=3)
Quaternion (w/o CUDA)	3h26m	25.50	Not Applicable	7
Cholesky	1h14m	25.28	N(1+N)/2	6
Skew-Symmetric w/ Matrix Exp	3h6m	25.49	N(1+N)/2	6
Skew-Symmetric w/ Cayley Transform	1h16m	25.52	N(1+N)/2	6
Skew-Symmetric w/ 2nd order Taylor	1h27m	25.54	N(1+N)/2	6
Skew-Symmetric w/ 1st order Taylor	1h	25.54	N(1+N)/2	6

temporal color encodings. The Beta kernel's inherent anisotropy across dimensions captures complex appearance variations that would otherwise require dozens of harmonic coefficients, making UBS both more expressive and more parameter-efficient than existing methods.

### A.3 ROTATION PARAMETERIZATION ANALYSIS

To validate our spatial-orthogonal Cholesky parameterization design, we conduct a comprehensive comparison of different spatial rotation parameterizations for covariance matrix construction on the <code>bicycle</code> scene from the Mip-NeRF 360 dataset (Barron et al., 2022) using UBS-3D (pure spatial Beta splatting without view or temporal dimensions). Note that this analysis focuses solely on spatial rotation parameterization within the 3D spatial subspace, not the full *N*-dimensional covariance structure used in UBS-6D or UBS-7D.

Table 5 presents our findings across different parameterization strategies:

**Quaternion.** The traditional 3DGS approach using quaternion rotation and diagonal scaling achieves competitive PSNR (25.50 dB) and the lowest parameter count (7 parameters: 3 position + 3 scale + 1 quaternion magnitude after normalization). However, this approach does not generalize to higher dimensions (N > 3), making it unsuitable for our N-dimensional framework.

Cholesky Decomposition. Pure Cholesky factorization  $\Sigma = LL^{\top}$  offers faster training (1h14m) and generalizes to arbitrary dimensions with N(N+1)/2 parameters. However, it suffers from a 0.22 dB PSNR drop compared to quaternions, as it lacks explicit orthogonality constraints that help maintain geometric consistency in the spatial subspace.

Skew-Symmetric Matrix Methods. Constructing spatial rotation matrices through  $\mathbf{R} = \exp(\mathbf{A})$  where  $\mathbf{A}$  is skew-symmetric preserves orthogonality and generalizes to higher dimensions. The naive implementation requiring full matrix exponentiation is prohibitively slow (3h6m), but approximation methods offer significant improvements. The Cayley transform  $\mathbf{R} = (\mathbf{I} - \mathbf{A})^{-1}(\mathbf{I} + \mathbf{A})$  reduces training to 1h16m with 25.52 dB PSNR. Taylor approximations provide the best balance: the 1st order Taylor approximation  $\mathbf{R} \approx \mathbf{I} + \mathbf{A}$  achieves both the fastest training time (1h) and the highest PSNR (25.54 dB), making it optimal for spatial rotation parameterization.

Our Choice - Spatial-Orthogonal Cholesky. Based on these results, we select the skew-symmetric approach with 1st order Taylor approximation for spatial rotation parameterization, as it provides the best combination of training speed (1h) and reconstruction quality (25.54 dB). For our full UBS framework, we adopt a hybrid spatial-orthogonal approach: skew-symmetric 1st order Taylor for the 3D spatial rotation subspace (preserving geometric interpretability and optimal performance) combined with Cholesky factors for modeling cross-dimensional correlations between spatial and non-spatial dimensions (view, time). This design achieves the optimal trade-off between performance, generalizability, and computational efficiency across all UBS variants.

The analysis confirms that the skew-symmetric 1st order Taylor approximation provides the best spatial rotation parameterization, enabling our spatial-orthogonal design to maintain geometric consistency in the spatial subspace while allowing flexible correlations with additional dimensions in the full N-dimensional UBS framework.

**Fundamental Limitations of Gaussian Kernels** N-dimensional Gaussian kernels suffer from inherent dimensional coupling that prevents modeling of independent scene properties. In methods like 6DGS and 7DGS, all dimensions are coupled through the joint Gaussian distribution: when the

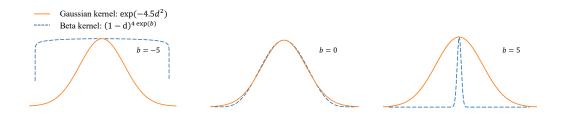


Figure 4: Adaptive Beta kernels.

input view direction or time changes, the primitive is inevitably affected across all dimensions—it shifts in space, changes shape, and modulates opacity simultaneously. This coupling stems from two fundamental constraints of the Gaussian kernel:

- 1. Fixed Bell-Shaped Profile: The Gaussian kernel maintains its characteristic bell-shaped exponential decay  $\exp(-r^2/2\sigma^2)$  across all dimensions. This fixed functional form cannot produce flat responses needed for view-independent or temporally-static elements. For example, in 6DGS, a primitive cannot remain view-independent because the Gaussian's bell-shaped profile in the angular dimensions will always modulate opacity based on viewing direction, even when the scene element should appear constant across views.
- 2. Symmetric Dimensional Coupling: The multivariate Gaussian distribution enforces symmetric relationships between dimensions through its covariance structure. When conditioned on viewing direction or time, the resulting 3D Gaussian inevitably shifts its spatial mean, changes its spatial covariance, and modulates its opacity according to the angular or temporal distance from the mean. This prevents primitives from exhibiting independence—a spatially-localized texture cannot remain fixed while only its angular response varies, or a static object cannot maintain constant appearance while dynamic elements change around it.

These limitations force inefficient representations where many coupled primitives must approximate what should be expressible by fewer independent primitives.

In contrast, UBS's Beta kernels enable per-dimension shape control: spatial dimensions can adopt flat profiles for view-independent geometry while angular dimensions use peaked responses for specular highlights, and temporal dimensions can remain nearly constant for static elements while varying sharply for dynamic components. This dimensional independence allows each primitive to model complex anisotropic phenomena efficiently—a single UBS primitive can simultaneously represent a static spatial surface with view-dependent reflectance, whereas Gaussian methods require multiple coupled primitives that cannot cleanly separate these properties.

Universal Beta Kernels: Mathematical Properties. UBS addresses these limitations through the N-dimensional Beta kernel:

$$B(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{b}) = \prod_{i=1}^{N} (1 - d_i(\boldsymbol{x}))^{\beta_i}, \quad d_i \in [0, 1),$$
(12)

where  $d_i(\boldsymbol{x})$  represents the bounded distance measure obtained through  $d_i = \max(\tanh(\frac{d_i^{\text{raw}}}{2}), 0)$  with  $d_i^{\text{raw}}$  being the Mahalanobis distance component along dimension i. The tanh and max transformations ensure  $d_i \in [0,1)$ , satisfying the Beta kernel input requirement. The shape parameter  $\beta_i = 4 \exp(b_i)$  controls the kernel profile.

## A.4 THEORETICAL FOUNDATION AND ANALYSIS

The Beta kernel exhibits three key properties:

**Property 1 (Shape Adaptivity):** The shape parameter  $b_i$  directly controls kernel behavior through  $\beta_i = 4 \exp(b_i)$ :

- $b_i < 0$ : Flatter kernel profile ( $\beta_i < 4$ , broader support)
- $b_i = 0$ : Gaussian-like profile ( $\beta_i = 4$ , standard bell curve)

•  $b_i > 0$ : More peaked kernel profile ( $\beta_i > 4$ , sharper localization)

When  $b_i=-5$ , we have  $\beta_i=4e^{-5}\approx 0.027$  and the kernel becomes nearly constant  $(1-d_i)^{0.027}\approx 1$ . When  $b_i=5$ ,  $\beta_i=4e^5\approx 594$  and the kernel becomes highly peaked. These practical ranges  $(b_i\in[-5,5])$  are sufficient to span from flat to sharp kernel profiles, as shown in Figure 4.

**Property 2 (Per-Dimension Shape Control):** The product form enables independent shape control for each dimension through its own  $b_i$ :

$$B(x) = \prod_{i=1}^{N} (1 - d_i(x))^{4 \exp(b_i)}, \qquad (13)$$

Unlike Gaussian kernels where all dimensions share the same bell-shaped functional form (only varying in scale), Beta kernels allow each dimension to adopt fundamentally different shapes: dimension i can be nearly flat ( $b_i = -3$ ) while dimension j is highly peaked ( $b_j = 3$ ). This perdimension shape control, combined with cross-dimensional correlations through  $\Sigma$ , enables efficient representation of complex anisotropic phenomena.

**Property 3 (Backward Compatibility):** When  $b_i = 0$ , we have  $\beta_i = 4 \exp(0) = 4$ , and the Beta kernel approximates a Gaussian:

$$(1-d)^4 \approx \exp(-4.5d^2)$$
 for  $d \in [0,1)$ , (14)

guaranteeing that UBS subsumes Gaussian methods as special cases when b = 0. This ensures seamless compatibility with existing Gaussian-based methods.