

Efficient Perspective-Correct 3D Gaussian Splatting Using Hybrid Transparency

Florian Hahlbohm¹, Fabian Friederichs¹, Tim Weyrich^{2,3}, Linus Franke²,
Moritz Kappel¹, Susana Castillo¹, Marc Stamminger², Martin Eisemann¹, and Marcus Magnor^{1,4}

¹Computer Graphics Lab, TU Braunschweig, Germany

²Visual Computing Erlangen, FAU Erlangen-Nürnberg, Germany

³University College London (UCL), United Kingdom

⁴Department of Physics and Astronomy, University of New Mexico, USA

<https://fhahlbohm.github.io/htgs/>



Figure 1: We present the first perspective and simultaneously geometrically accurate approach for real-time rendering of 3D Gaussian splats. Our temporally stable blending formulation based on hybrid transparency effectively removes the popping artifacts caused by the approximate, global sorting scheme of 3D Gaussian Splatting (3DGS) [KKLD23]. Additionally, we replace the approximate projection used in 3DGS with a novel, numerically stable formulation for evaluation of general 3D Gaussians along per-pixel rays to improve rendering quality. We also achieve a significant speedup compared to state-of-the-art approaches addressing these challenges [HYC*24, RSP*24], as our blending formulation only requires partial depth-ordering and our splat evaluation is well-suited for rasterization.

Abstract

3D Gaussian Splats (3DGS) have proven a versatile rendering primitive, both for inverse rendering as well as real-time exploration of scenes. In these applications, coherence across camera frames and multiple views is crucial, be it for robust convergence of a scene reconstruction or for artifact-free fly-throughs. Recent work started mitigating artifacts that break multi-view coherence, including popping artifacts due to inconsistent transparency sorting and perspective-correct outlines of (2D) splats. At the same time, real-time requirements forced such implementations to accept compromises in how transparency of large assemblies of 3D Gaussians is resolved, in turn breaking coherence in other ways. In our work, we aim at achieving maximum coherence, by rendering fully perspective-correct 3D Gaussians while using a high-quality approximation of accurate blending, hybrid transparency, on a per-pixel level, in order to retain real-time frame rates. Our fast and perspective-correct approach for evaluation of 3D Gaussians does not require matrix inversions, thereby ensuring numerical stability and eliminating the need for special handling of degenerate splats, and the hybrid transparency formulation for blending maintains similar quality as fully resolved per-pixel transparencies at a fraction of the rendering costs. We further show that each of these two components can be independently integrated into Gaussian splatting systems. In combination, they achieve up to 2× higher frame rates, 2× faster optimization, and equal or better image quality with fewer rendering artifacts compared to traditional 3DGS on common benchmarks.

CCS Concepts

• Computing methodologies → Rendering; Point-based models; Rasterization; Machine learning approaches;

1. Introduction

Novel view synthesis has undergone a significant transformation with the advent of Neural Radiance Fields (NeRF) [MST*20] and 3D Gaussian Splatting (3DGS) [KKLD23], both of which have established themselves as powerful techniques for rendering complex 3D scenes. Among these, 3DGS has emerged as the de-facto representation for user-facing radiance field applications due to its fast optimization and real-time rendering capabilities. Its speed and efficiency stem from an explicit point-based representation, where each point’s extent is modeled by an anisotropic 3D Gaussian. Using fast, tile-based rasterization, this can be efficiently implemented on modern GPUs.

However, the remarkable performance of 3DGS is achieved through a series of approximations that, while effective in many scenarios, introduce limitations in multi-view consistency. One significant issue arises from the affine approximation used in projecting 3D Gaussian splats onto the image plane, causing differing pixel contributions depending on camera placement. Another one is the per-primitive sorting, which results in patchy colors appearing and disappearing during movements (called popping [RSP*24]). In order to solve the issues, it is necessary to accurately sort and evaluate Gaussians per pixel, introducing a strong bottleneck on the sorting stage. In this paper, we introduce a pipeline which improves the original 3DGS framework in eliminating these artifacts while offering superior performance through a hybrid transparency approximation.



Figure 2: In 3DGS [KKLD23], Kerbl et al. use an affine approximation for projecting 3D Gaussians. Common benchmarks do not include viewpoints where this approximation matter is highlighted. Here, we demonstrate the benefits of our perspective-correct projection by comparing renderings from 3DGS and our method.

For projection artifacts, although the affine approximation performs well on benchmark datasets, it fails to model perspective distortion correctly, especially when parts of the scene are viewed at close distances. The result are visually disturbing artifacts, where the projected Gaussians take on extreme, distorted shapes, severely affecting the rendering quality (see Fig. 2).

Recent work on 2D Gaussian Surfels [HYC*24] has made strides in achieving perspective-accurate rendering by leveraging established techniques from Sigg et al. [SWBG06] and Weyrich et al. [WHA*07]. For 3D Gaussians, concurrent research explores evaluating splats by calculating intersection points with per-pixel viewing rays [YSG24]. This approach can be implemented

in ray-tracing frameworks [MLMP*24], whose rendering speeds are significantly lower. Alternatively, the rasterization-based approach faces numerical instability due to the matrix inversions required, making optimization via gradient descent challenging and prone to catastrophic failure if handled incorrectly.

In this paper, we propose a fast, differentiable method for perspective-accurate 3D Gaussian splat evaluation at the point of maximum contribution along per-pixel viewing rays that avoids matrix inversion entirely.

The second issue with 3DGS lies in depth ordering during rendering, where only the view-space depth of the mean is considered. This causes incorrect per-pixel blending order of fragments, leading to the so-called “popping” artifacts that disrupt multi-view consistency and especially the viewing experience in motion. Recent work, such as *StopThePop* [RSP*24], addresses this through hierarchical sorting, which is comprised of global presorting and then locally sorting within a sliding window, with progressively lower tile sizes.

We propose a similar idea, based on the established rendering paradigm of *Hybrid Transparency* [Wym16], which lets us skip the global presorting and provides high performance. By alpha-blending the first K fragments (called the *core*) in correct depth-order per pixel and accumulating remaining contributions (the *tail*) using an order-independent residual, our method mitigates the popping artifacts (see Fig. 3) while maintaining superior performance.




Figure 3: Radl et al. [RSP*24] discuss how 3DGS approximates depth-ordering for alpha blending, leading to popping artifacts during view rotation. To solve this without sacrificing performance, we propose a hybrid transparency approach, combining alpha blending with order-independent transparency, resulting in temporally stable rendering and an improved viewing experience.

These two enhancements – perspective-accurate splat evaluation

and improved depth-ordering – can be integrated into existing 3DGS systems independently. In our implementation, we combine both improvements, demonstrating that they not only eliminate the respective artifacts but also enhance rendering speed, crucial for real-time applications that rely on fast scene inspection. In summary, we make the following contributions:

- We present a novel, matrix inversion-free ray-splat intersection method within a differentiable renderer for 3D Gaussian splats, resulting in numerically stable optimization and fully perspective-correct rendering.
- We introduce hybrid transparency to 3D Gaussian splat rendering, resulting in both very stable results and improved performance.
- We evaluate our efficient implementation, which outperforms established methods significantly, while maintaining visual quality.

2. Related Work

This work is at the intersection of novel view synthesis, Gaussian splatting, and order-independent transparency.

2.1. Novel View Synthesis

Novel view synthesis aims to render images of a scene from arbitrary viewpoints given a fixed number of input images.

The field has recently been revolutionized by Neural Radiance Fields (NeRF) introduced by Mildenhall et al. [MST*20]. NeRF represents a scene by optimizing a large MLP that maps positions and viewing directions to volumetric density and color. The use of differentiable volume rendering enables optimization using gradient descent, which leads to high-quality results; however, NeRF’s slow training and rendering process hinders its use in interactive applications. To accelerate NeRF, voxel grids with trilinear interpolation were introduced, offering faster, continuous representations [FKYT*22, SSC22]. Memory-efficient methods, such as multi-resolution hash tables [MESK22] and tensor factorization [CXG*22, RSV*23], further reduce computational overhead. The current state-of-the-art method in terms of image quality, Zip-NeRF [BMV*23], combines these grid-based methods with solutions addressing aliasing issues [BMT*21, BMV*22].

Alternatively, point-based models use explicit point clouds for geometry, rendering images through fast rasterization. Due to the discrete nature of point clouds, recent point-based radiance field methods employ large convolutional neural networks (CNN) for hole-filling in image-space [ASK*20, RFS22, FRF*23, KPLD21, KLR*22, HFF*23], but these CNNs are computationally expensive. Addressing this limitation, Franke et al. showed that assigning a radius to each point in combination with trilinear interpolation into an image pyramid allows using a smaller CNN [FRFS24]. While enabling fast rendering, image quality of point-based models is often lacking behind implicit NeRF- and grid-based models. Very recently, Hahlbohm et al. showed that point-based models and grid-based models can be combined to improve image quality and robustness [HFK*24]. Remaining issues with point-based models are temporal instabilities and the reliance on dedicated GPUs, which are both due to the used CNN. Another approach involves converting a trained NeRF model into a faster representation for

inference [HSM*21, YLT*21]. Recent methods achieve impressive frame rates with only a small quality loss via triangle mesh baking [YHR*23, RGS*24] or distilling of state-of-the-art NeRF models into memory-efficient, hybrid representations [DHR*24]. While they render exceptionally fast even on low-end mobile devices, a major disadvantage is their reliance on the NeRF model that needs to be trained before the – often equally expensive – baking process. Other approaches that proved effective directly optimize a triangle meshes [CFHT23] or a set of tetrahedra [KS23].

2.2. Gaussian Splatting

In early work on point-based rendering, points were rasterized as opaque splats [GD98, RL00, PZVBG00] and thus suffered from strong aliasing. It was shown that quality can be significantly improved using semitransparent splats [ZPvBG01b], however this requires costly sorting and blending operations. To address performance issues, fast algorithms for then GPUs [BHZZK05] or dedicated hardware implementations [WHA*07] were proposed.

More recently, point-based rendering regained momentum as a powerful primitive for differentiable rendering in novel view synthesis [KPLD21, ASK*20, RFS22]. In 2023, Kerbl et al. [KKLD23] re-introduced rasterization-based rendering of point primitives with anisotropic Gaussian extent in their seminal work *3D Gaussian Splatting (3DGS)*. Utilizing fast, tile-based rasterization [LZ21], 3DGS optimizes a set of explicit 3D Gaussians via gradient descent and an adaptive density control mechanism. Due to exceptionally fast rendering during inference, 3DGS has inspired a plethora of follow-up research addressing, e.g., anti-aliasing [YCH*24], compression [BKL*24], dynamic scenes [LKL24, WYF*24], large scenes [KMK*24], and the reliance on an initial point cloud [KRS*24, NMR*24]. Especially relevant to this work is *StopThePop* by Radl et al. [RSP*24], which employs a hierarchical sorting approach to reduce popping artifacts originating from approximate sorting used in 3DGS. Notably, the per-pixel ordering resulting from this hierarchical sorting approach is not guaranteed to be fully correct. We also present a method for avoiding these popping artifacts, but achieve this by changing the blending procedure to a hybrid transparency-based approach so that our method only requires partial depth-ordering of the initial K splats contributing to each pixel color. Thus, we avoid the need for a sophisticated sorting implementation which, even when compared to the approximate solution used by 3DGS, improves performance. Equally important in the context of this work is the perspective-accurate *2D Gaussian Splatting (2DGS)* by Huang et al. [HYC*24] as well as concurrent work that introduces ray tracing into the 3DGS framework [YSG24, MLMP*24]. While the latter also provide a perspective-accurate rendering formulation by using ray tracing, they introduce a significant slowdown during both the optimization and subsequent inference. Furthermore, their evaluation of the 3D Gaussian primitives relies on matrix inversion, which is prone to numerical instabilities when Gaussians fall flat along one or more principal axes. The perspective-correct approach by Huang et al. [HYC*24] circumvents these numerical instabilities by utilizing an optimized approach for 2D, i.e., degenerate 3D Gaussians [SWBG06, WHA*07]. In this work, we present an approach that enables using perspective-correct rendering of 3D Gaussians through ray-splat intersection: we extend the perspective-correct

approach of Weyrich et al. [WHA*07] to work with non-degenerate 3D Gaussians, which enables the use of ray-splat intersection for rasterization-based rendering with negligible overhead and no issues with respect to numerical instabilities.

2.3. Order-Independent Transparency

The seminal Porter-Duff algorithm [PD84] blends two semi-transparent surfaces correctly, but extending it to more than two requires costly sorting of the input primitives, and even produces wrong results in case of intersecting primitives. Several techniques have been developed relying on even more costly per-pixel sorting to circumvent this problem, either by sorting explicitly [Car84] or implicitly using multiple render passes [Eve01, BM08]. These algorithms are in general invariant to the order of input primitives and produce correct results, and are therefore called *exact order-independent transparency (OIT)* methods. A good overview can be found in [Wym16]. A plethora of approaches try to approximate exact OIT by avoiding the costly sorting step. Some of them are not truly order-independent: the outcome still varies slightly depending on the order of input primitives [BCL*07, ESSL10, SML11, SV14]. True, *approximate OIT* algorithms generally modify the blending operations to achieve order-invariance. Ignoring the order-dependent parts of the alpha blending formula results in the “weighted sum” operator [Mes07], which works well with surfaces of small opacity but becomes increasingly inaccurate for more opaque surfaces, causing over-darkening or over-brightening. The weighted blended order-independent transparency (WBOIT) operator avoids the over-darkening by replacing the opacity and surface colors with opacity-weighted averages and introducing a monotonic decreasing depth weighting factor, heuristically reducing the influence of distant surfaces [MB13]. This approach works well for surfaces of low opacity and similar color, like smoke, but becomes problematic for almost opaque surfaces, and the weighting function needs to be tuned for every scene to achieve optimal results. Extending this approach to rendering the surfaces into layers and applying the WBOIT operator to each before blending the resulting colors in order provides an interesting trade-off between quality and efficiency [FEE21]. Moment-based transparency techniques remedy some of the shortcomings of WBOIT by computing moments of the transmittance which are used in a second pass to reconstruct an approximate transmittance function, that weighs the influence of each surface on the final color [MKKP18]. Sharp changes in the transmittance function are not well captured at a lower number of power moments, whereas higher numbers decrease the computational efficiency. As the influence of later surfaces is generally less than the first surfaces due to decreasing transmittance, another, truly order-independent approach is *hybrid transparency* that blends the first K surfaces correctly (including sorting), whereas the remaining surfaces are blended without sorting [MCTB13]. In this work, we integrate this hybrid transparency approach into the 3DGS framework and extend the formulation proposed by Maule et al. to improve its behavior within a gradient descent-based optimization scheme.

3. Method

We introduce the employed scene representation and then describe the two parts of our contribution.

3.1. Preliminaries

Following Kerbl et al. [KKLD23], we represent the scene using a set of 3D points with anisotropic Gaussian extent as a proxy for geometry. Cutting each Gaussian at 3σ results in ellipsoid-shaped splats. Each splat is defined by its 3D mean $\mu \in \mathbb{R}^3$, a scalar opacity value $o \in \mathbb{R}$, three principal tangential vectors $t_u, t_v, t_w \in \mathbb{R}^3$ modeling its orientation, and three scalars $s_u, s_v, s_w \in \mathbb{R}$ modeling its scale. Note that t_u, t_v, t_w represent a 3D rotation, which allows modeling these parameters using a quaternion $q \in \mathbb{R}^4$ that is easier to optimize via gradient descent. For rendering, we obtain a valid opacity value in $[0, 1]$ using a Sigmoid activation, represent s_u, s_v, s_w in logarithmic space, and ensure q is normalized. The view-dependent color representation based on spherical harmonics (SH) is identical to 3DGS, i.e., each point uses SH up to degree 3 resulting in 48 coefficients total.

For any point $x \in \mathbb{R}^3$, we can obtain an α value for a given 3D Gaussian by multiplying its opacity with the value of the Gaussian at that point:

$$\alpha = o \cdot e^{-\frac{1}{2}\rho(x)^2} \quad \text{using} \quad \rho(x)^2 = (x - \mu)^T \Sigma^{-1} (x - \mu), \quad (1)$$

where Σ^{-1} can easily be computed from (t_u, t_v, t_w) and (s_u, s_v, s_w) (cf. Kerbl et al. [KKLD23]). These α values are then used in conjunction with the view-dependent RGB color $c \in \mathbb{R}^3$ when computing per-pixel colors through standard depth-sorted alpha blending of all N contributing splats:

$$C = \sum_{i=1}^N \alpha_i T_i c_i, \quad \text{with} \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j). \quad (2)$$

3.2. Accurate Splat Bounding and Evaluation

In 3DGS, Kerbl et al. [KKLD23] project a 3D Gaussian onto the image plane using the Jacobian of the affine approximation of the projective transformation [ZPvBG01a]. The result is a 2D Gaussian that can easily be evaluated for each pixel and also allows for easy construction of an axis-aligned bounding box around the center (e.g., at 3σ) in screen space for accelerated computation. However, this approach is not fully perspective-accurate, meaning it introduces artifacts when splats are viewed from certain angles. For example, these artifacts become visible when placing the camera on top of a flat surface with the viewing direction perpendicular to the surface normal (see first row in Fig. 2).

In this work, we propose a solution for this problem that builds upon an established approach for perspective accurate rendering of ellipsoidal surfaces by Sigg et al. [SWBG06] and the respective optimizations by Weyrich et al. [WHA*07] for 2D Gaussian splats. It should be noted that these approaches elegantly avoid the matrix inversions required by similar approaches (e.g., Eq. (1)), which makes them numerically stable even for degenerate splats, i.e., those where one or more main axes vanish. Notably, Huang et al. [HYC*24] also base their perspective accurate ray-splat intersection and evaluation on the aforementioned work. However, their approach only works for 2D Gaussians, i.e., degenerate splats. In contrast, we now introduce an approach that enables perspective accurate 3D splat evaluation along per-pixel viewing rays. Thus, our approach is more general




Figure 4: The perspective-correct screen-space bounding box of a splat (a) is given by the projection of its bounding frustum in view space (b). When transformed into local splat coordinates, the frustum planes align with tangential planes of the unit sphere (c). Our approach for splat evaluation along viewing rays makes use of the Plücker coordinate representation $(\mathbf{d} : \mathbf{m})$. In local splat coordinates, the point along the ray that maximizes the Gaussian’s value corresponds to the point \mathbf{x} that minimizes the perpendicular distance $\|\mathbf{x}\|$ to the origin (d). Parts (a-c) courtesy of Weyrich et al. [WHA*07]; used with permission.

compared to Huang et al., as the 3D Gaussians we use may still degenerate to 2D Gaussians without impacting the optimization or rendering due to the numerical stability of our approach.

We start by describing the approach for computing tight, axis-aligned screen-space bounding boxes for each 3D Gaussian. We define the model-view matrix $M \in \mathbb{R}^{4 \times 4}$, the projection to clip-space $P \in \mathbb{R}^{4 \times 4}$, and the viewport transform $V \in \mathbb{R}^{4 \times 4}$ mapping to window coordinates. Using the per-splat attributes described in Sec. 3.1, the transformation $T' \in \mathbb{R}^{4 \times 4}$ from the normalized splat space (where the Gaussian’s ellipsoid becomes the unit sphere) to screen space is:

$$T' = VPMT, \quad \text{where} \quad T = \begin{pmatrix} | & | & | & | \\ s_u t_u & s_v t_v & s_w t_w & \mu \\ | & | & | & | \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (3)$$

is the transformation from normalized splat space to world space. Using Weyrich et al.’s optimization of Sigg et al.’s bounding box computation [SWBG06], extended to non-degenerate (full 3D) Gaussians, the desired tight 3D bounding box in screen space is $[b_1, t_1] \times [b_2, t_2] \times [b_3, t_3]$ (bottom and top values, b_i and t_i , for $i \in \{1, 2, 3\}$ the x , y , and z coordinate, respectively) is:

$$b_i = p_i - h_i \quad (4)$$

$$t_i = p_i + h_i, \quad (5)$$

with

$$s = \langle (\rho_c, \rho_c, \rho_c, -1), T'_4 \odot T'_4 \rangle, \quad \rho_c = 2 \ln \frac{\rho}{\tau_\alpha}, \quad (6)$$

$$\mathbf{f} = \frac{1}{s} (\rho_c, \rho_c, \rho_c, -1), \quad (7)$$

$$p_i = \langle \mathbf{f}, T'_i \odot T'_i \rangle, \quad (8)$$

$$h_i = \sqrt{p_i^2 - \langle \mathbf{f}, T'_i \odot T'_i \rangle}, \quad (9)$$

and T'_i the i -th row of T' , \odot denoting component-wise multiplication, and the dot product denoted by $\langle \cdot, \cdot \rangle$. A key advantage of this calculation, for which we provide a visual explanation in Fig. 4, is that it offers a closed-form solution without the need for matrix inversion and regardless of whether the ellipsoid is degenerate in any direction. Following Radl et al. [RSP*24], we also compute an individual cutoff value for each splat to obtain smaller bounding

boxes when their opacity value is less than one based on τ_α , a hyperparameter defining the minimum α value, Eq. (1), of a fragment for it not to be skipped during blending.

Next, we seek to compute the value of the 3D Gaussian for a given viewing ray. In computer graphics, ray intersections with deformed primitives, such as our 3D splats, are an established practice since the dawn of ray tracing, generally following the approach of inversely transforming a viewing ray into the undeformed space. In our case, re-using quantities of the bounding-box setup above, this would mean transforming the ray by $(VPMT)^{-1}$ into splat space, where the distance of the transformed ray to the origin yields ρ , Eq. (1). In order to avoid the matrix inversion that we successfully sidestepped during bounding-box calculation, we once again follow Weyrich et al. and represent a viewing ray through a pixel at (x_s, y_s) as two perpendicular planes[†] $\pi_x = (1, 0, 0, -x_s)^\top$ and $\pi_y = (0, 1, 0, -y_s)^\top$ in screen space. In contrast to 3D points, transforming those planes by $(VPMT)^{-1}$ into planes $\pi_{x/y}^s$ in splat space requires the inverse-transposed mapping, that is:

$$\pi_{x/y}^s = ((VPMT)^{-1})^{-\top} \pi_{x/y} = (VPMT)^\top \pi_{x/y}, \quad (10)$$

once again sidestepping inversion. Next up, intersection of π_x^s and π_y^s yields the viewing ray in splat space, and its distance to the origin is ρ . In contrast to Weyrich et al., however, ρ cannot quite as easily be computed, due to the non-degeneracy of our 3D Gaussians. Instead, we employ the dual Plücker coordinate representation that derives the mapped viewing ray \mathcal{L}_s^* from the intersection of $\pi_{x/y}^s$ from their coefficients $\pi_x^s = (a^1 : a^2 : a^3 : a^0)$, and $\pi_y^s = (b^1 : b^2 : b^3 : b^0)$, respectively:

$$\mathcal{L}_s^* = (p^{23} : p^{31} : p^{12} : p^{01} : p^{02} : p^{03}), \quad (11)$$

$$\text{with } p^{ij} = \begin{vmatrix} a^i & a^j \\ b^i & b^j \end{vmatrix} = a^i b^j - a^j b^i. \quad (12)$$

The line’s dual Plücker coordinates \mathcal{L}_s^* are numerically equivalent to its primal Plücker coordinates \mathcal{L}_s up to some common scaling

[†] We follow the common convention of representing planes as homogeneous vectors $\mathbf{p} = (a, b, c, d)^\top$, so that a point \mathbf{x} lies on the plane iff $\mathbf{p}^\top \mathbf{x} = 0$.

factor λ :

$$\mathcal{L}_s =: (\mathbf{d} : \mathbf{m}) = (p_{01} : p_{02} : p_{03} : p_{23} : p_{31} : p_{12}) \quad (13)$$

$$= (\lambda p^{23} : \lambda p^{31} : \lambda p^{12} : \lambda p^{01} : \lambda p^{02} : \lambda p^{03}) \quad (14)$$

where \mathbf{d} is the ray direction in splat space and \mathbf{m} is its moment around the origin (see Fig. 4 for a visualization). From the known equality $\text{distance}(\mathcal{L}_s, \text{origin}) = \frac{\|\mathbf{m}\|}{\|\mathbf{d}\|}$, it follows that:

$$\rho(\mathbf{x})^2 = \frac{\|\lambda(p^{01}, p^{02}, p^{03})^\top\|^2}{\|\lambda(p^{23}, p^{31}, p^{12})^\top\|^2} = \frac{\begin{vmatrix} a^0 b^1 - a^1 b^0 \\ a^0 b^2 - a^2 b^0 \\ a^0 b^3 - a^3 b^0 \end{vmatrix}^2}{\begin{vmatrix} a^2 b^3 - a^3 b^2 \\ a^3 b^1 - a^1 b^3 \\ a^1 b^2 - a^2 b^1 \end{vmatrix}^2}. \quad (15)$$

Note that we directly compute $\rho(\mathbf{x})^2$ for Eq. (1). This formulation is numerically stable, as the only potential issue – a vanishing denominator in Eq. (15) – can easily be detected and corresponds to the ray missing the splat. Conveniently, the obtained value corresponds to the point along the pixel’s viewing ray where the Gaussian has the highest value, essentially making it equivalent to the numerically unstable and slower approaches used in concurrent works [YSG24, MLMP*24]. Compared to Kerbl et al. [KKLD23], our calculations incur reduced setup costs for each Gaussian but slightly higher costs per pixel, which is in fact a desirable tradeoff whenever graphics primitives cover only few pixels [MBDM97]. Moreover, the added computations are exclusively additions and multiplications, except for the division in Eq. (15), meaning they are both fast and trivial to differentiate.

3.3. Temporally-Stable Rendering via Hybrid Transparency

In 3DGS, Kerbl et al. [KKLD23] use a tile-based rasterization approach to efficiently render images from a set of 3D Gaussian primitives. To achieve robust optimization via gradient descent, they compute per-pixel color by applying standard α -blending, Eq. (2), in depth-sorted order. In this work, we propose to integrate a hybrid-transparency approach for blending into the 3DGS framework. This is motivated by three observations: (1) The global sorting scheme proposed by Kerbl et al. [KKLD23] sorts splats as a whole, leading to incorrectly resolved splat-splat intersections and temporal incoherence (“popping”) during rendering and optimization; resolving these artifacts through a full sort of all 3D Gaussian pixel fragments, however, is prohibitively slow and thus approximate sorting is required (cf. Radl et al. [RSP*24]). (2) As we will demonstrate in our experiments, using “full” OIT, which forgoes sorting completely, works well during optimization and even improves background reconstruction but results in semi-transparent foreground rendering, which is clearly undesirable for novel-view synthesis and reconstruction methods. (3) Blending only the first K contributions for each pixel has been shown to work well in differentiable rendering frameworks while also much faster than approaches that require a complete sort [LZ21, FRFS24].

The hybrid-transparency approach [MCTB13] splits up the per-pixel color computation into two parts: The *core* that uses standard α -blending to blend the first K contributions for each pixel, including

sorting. All remaining contributions are combined into a *tail*, that is blended without sorting. To decide whether a splat is one of the first K , we need to compute a pixel-specific depth value for each of them, i.e., the z -value of the point of evaluation \mathbf{x} (see Fig. 4) in view space. Re-using the intermediate values of our splat evaluation (see Sec. 3.2), \mathbf{x}_{view} can be computed efficiently by

$$\mathbf{x}_{\text{view}} = MT \frac{\mathbf{d} \times \mathbf{m}}{\|\mathbf{d}\|^2}. \quad (16)$$

As in Sec. 3.2, \mathbf{x}_{view} is the splat’s point of maximum contribution along the pixel ray, which again makes our computation compatible with the numerically unstable and slower approaches used in recent work [RSP*24]. In our implementation, we keep track of the first K contributions and accumulate all of the $(N - K)$ contributions for each pixel to then compute its color as

$$C = \sum_{i=1}^K \alpha_i T_i \mathbf{c}_i + T_{K+1} \left((1 - T_{\text{tail}}) \mathbf{c}_{\text{tail}} + T_{\text{tail}} \mathbf{c}_{\text{bg}} \right), \quad (17)$$

with

$$T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad \mathbf{c}_{\text{tail}} = \frac{\sum_{i=K+1}^N \alpha_i \mathbf{c}_i}{\sum_{i=K+1}^N \alpha_i}, \quad T_{\text{tail}} = \prod_{i=K+1}^N (1 - \alpha_i). \quad (18)$$

Notably, our computation of T_{tail} is more accurate compared to the average-based formulation of Maule et al. [MCTB13]. As the partial derivative of the blending function is the same for all splats inside the tail, our approach allows pre-computing all required partial derivatives for each pixel during the forward pass. Thus, our backward pass can be implemented very efficiently.

4. Experiments

We conduct multiple experiments on established datasets to validate our approach.

4.1. Setup

Following recent work, we evaluate on the established Mip-NeRF360 [BMV*22] and Tanks and Temples [KPZK17] datasets, which provide a diverse set of 17 real scenes with varying challenges. We use all nine scenes from the Mip-NeRF360 dataset as well as all eight scenes from the *intermediate* set of the Tanks and Temples dataset resulting in a total of 17 scenes. This allows us to evaluate our method on a broad spectrum of challenges regarding both geometric and photometric aspects. As is common practice, we use the $4 \times 2 \times$ downsampled images for the outdoor/indoor scenes of the Mip-NeRF360 dataset. For the Tanks and Temples scenes, we use the full-resolution images to evaluate performance at full HD resolution. To compute meaningful quality metrics, we use the established 7:1 train/test split [BMV*22] for all scenes. Central to our evaluation is the comparison against the original 3D Gaussian Splatting [KKLD23]. Furthermore, we compare against the recent StopThePop [RSP*24] and 2D Gaussian Splatting [HYC*24], as they also provide solutions for the problems addressed in this work. We also compare against Gaussian Opacity Fields (GOF) [YSG24], a very recent work that uses ray tracing-based volume rendering. For all baselines, we use the official implementation to compute the results for all 17 scenes and use the same script for all quality metric

Method	Mip-NeRF360 [BMV*22] (Img. Res. \approx 1MP-1.5MP)						Tanks&Temples [KPZK17] (Img. Res. \approx 2MP)					
	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	Training	FPS	#Splats	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	Training	FPS	#Splats
Zip-NeRF [BMV*23]	0.828	28.56	0.219	5h	0.2	N/a	0.878	26.75	0.233	5h	0.1	N/a
3DGS [KKLD23]	0.814	27.20	0.254	18m40s	316.6	3.31M	0.866	25.27	0.276	20m18s	216.2	1.50M
StopThePop [RSP*24]	0.814	27.30	0.252	12m03s	212.5	3.28M	0.866	24.99	0.270	14m27s	206.2	1.49M
2DGS [HYC*24]	0.795	26.81	0.297	21m08s	134.9	2.01M	0.851	24.55	0.314	28m45s	95.8	0.95M
GOF [YSG24]	0.821	27.27	0.238	1h16m	16.7	2.91M	0.866	25.01	0.274	1h35m	12.3	1.26M
Ours	0.822	27.17	0.234	10m07s	457.3	2.19M	0.866	24.62	0.272	8m33s	380.7	0.81M

Table 1: Quantitative comparisons on the Mip-NeRF360 and Tanks and Temples datasets. Our approach of using perspective correct splat evaluation in combination with hybrid transparency significantly reduces training and rendering times with image quality being similar to the baselines’. Excluding Zip-NeRF, the three best results are highlighted in green in descending order of saturation.

computations. We use a single RTX 4090 for all experiments and outline our approach for measuring frame rates in Appendix A.

4.2. Implementation and Optimization

We implement our method from scratch in PyTorch and CUDA using the implementation of Kerbl et al. [KKLD23] as a reference. While the optimization pipeline and hyperparameters are mostly the same as in 3DGS, we make multiple modifications. First, we replace the opacity reset with a decay strategy [RSP*24] that multiplies splat opacity by $\lambda_o = 0.9995$ every 50 iterations during densification. We tackle aliasing problems by employing our own fast CUDA implementation of the 3D filter proposed by Yu et al. [YCH*24]. As the screen-space gradients 3DGS uses for densification are not available with our splat evaluation scheme, we follow Moenne-Loccoz et al. and instead directly use the gradient of the splat means scaled by half the respective distance to the camera in view space [MLMP*24]. Furthermore, we adapt the densification changes from GOF to our approach [YSG24] and apply visibility score-based pruning as proposed by Niemeyer et al. [NMR*24]. For our blending formulation, we use a core with $K = 16$ and employ a second threshold $\tau_K = 0.05$ specifying the minimum α of a fragment to be considered for the core. As the insertion sort that we employ for keeping track of per-pixel core causes significant divergence, we reduce the tile size to 8×8 in our implementation. Lastly, we make use of an optimized SSIM implementation [MKVC*24] to speed up loss computation.

4.3. Results

The quantitative results in Tab. 1 show that our approach provides significantly faster training and rendering. Compared to the next fastest baseline, we achieve a training speedup of $1.4\times$ and render $1.6\times$ faster on average. In terms of image quality, we observe that our method achieves about equal results across the board. However, all 3DGS-based approaches are clearly outperformed by Zip-NeRF [BMV*23]. Taking our contributions into consideration, we observe that our approach outperforms approaches with perspective correct splat evaluation: Compared to 2DGS [HYC*24], we achieve better results across the board reassuring that 3D Gaussian splats are more expressive than the 2D Gaussian surfels used by 2DGS. GOF [YSG24] achieves slightly higher image quality but requires longer training and fails to render in real-time. It should be noted that 2DGS and GOF were fine-tuned for mesh extraction after

optimization, which is an interesting topic but not a focus of this work. Seeing that we achieve similar image quality but faster training and rendering compared to StopThePop [RSP*24], the results confirm that our approach is a valid alternative to prevent popping artifacts for rasterization-based 3D Gaussian splat rendering.

We also show visual comparisons for multiple scenes in Fig. 5. Beyond those comparisons, we often observe more accurate reconstruction of the background when inspecting trained models. This is likely caused by our perspective correct splat evaluation that allows better use of the multi-view signal provided by the training images where the background is generally located near image borders where perspective distortion has its highest influence.

Method	Preprocess	Tiling	Blending	Total
3DGS [KKLD23]	0.451 ms	2.212 ms	2.134 ms	4.797 ms
StopThePop [RSP*24]	0.649 ms	0.738 ms	3.599 ms	4.986 ms
Ours	0.556 ms	0.754 ms	2.007 ms	3.317 ms

Table 2: Breakdown of render timings at a resolution of 1920×1080 pixels measured on an RTX 4090 GPU. For our approach, we render the trained bicycle model (4.35M splats) for all test-set viewpoints and average across 100 runs. Values for 3DGS and StopThePop are copied from Radl et al. [RSP*24], who use a slightly different approach for their performance benchmarks but average in a similar manner, which makes numbers comparable.

Performance Breakdown. In Tab. 2, we compare how parts of the rendering influence the total runtime. We denote the creation of per-tile primitive lists (cf. Kerbl et al. [KKLD23]) as *tiling* and per-tile color computations as *blending*. The speed of 3DGS is equally limited by tiling and blending. StopThePop [RSP*24] implements significant optimizations for the tiling step in an attempt to reduce the overhead of their hierarchical sorting during blending, which is the main bottleneck. In contrast, we use a 16-bit unsigned integer as the key for each Gaussian/tile instance, as we do not require any global depth-sorting. This speeds up tiling by a similar amount as the revised culling strategy of StopThePop. Regarding the blending step, our hybrid transparency approach requires only partial depth-ordering making it significantly faster than the approach of StopThePop.

Perspective Correct Splat Evaluation. As previously discussed, our contributions can be integrated into 3D Gaussian splatting implementations independently. Our hybrid transparency approach

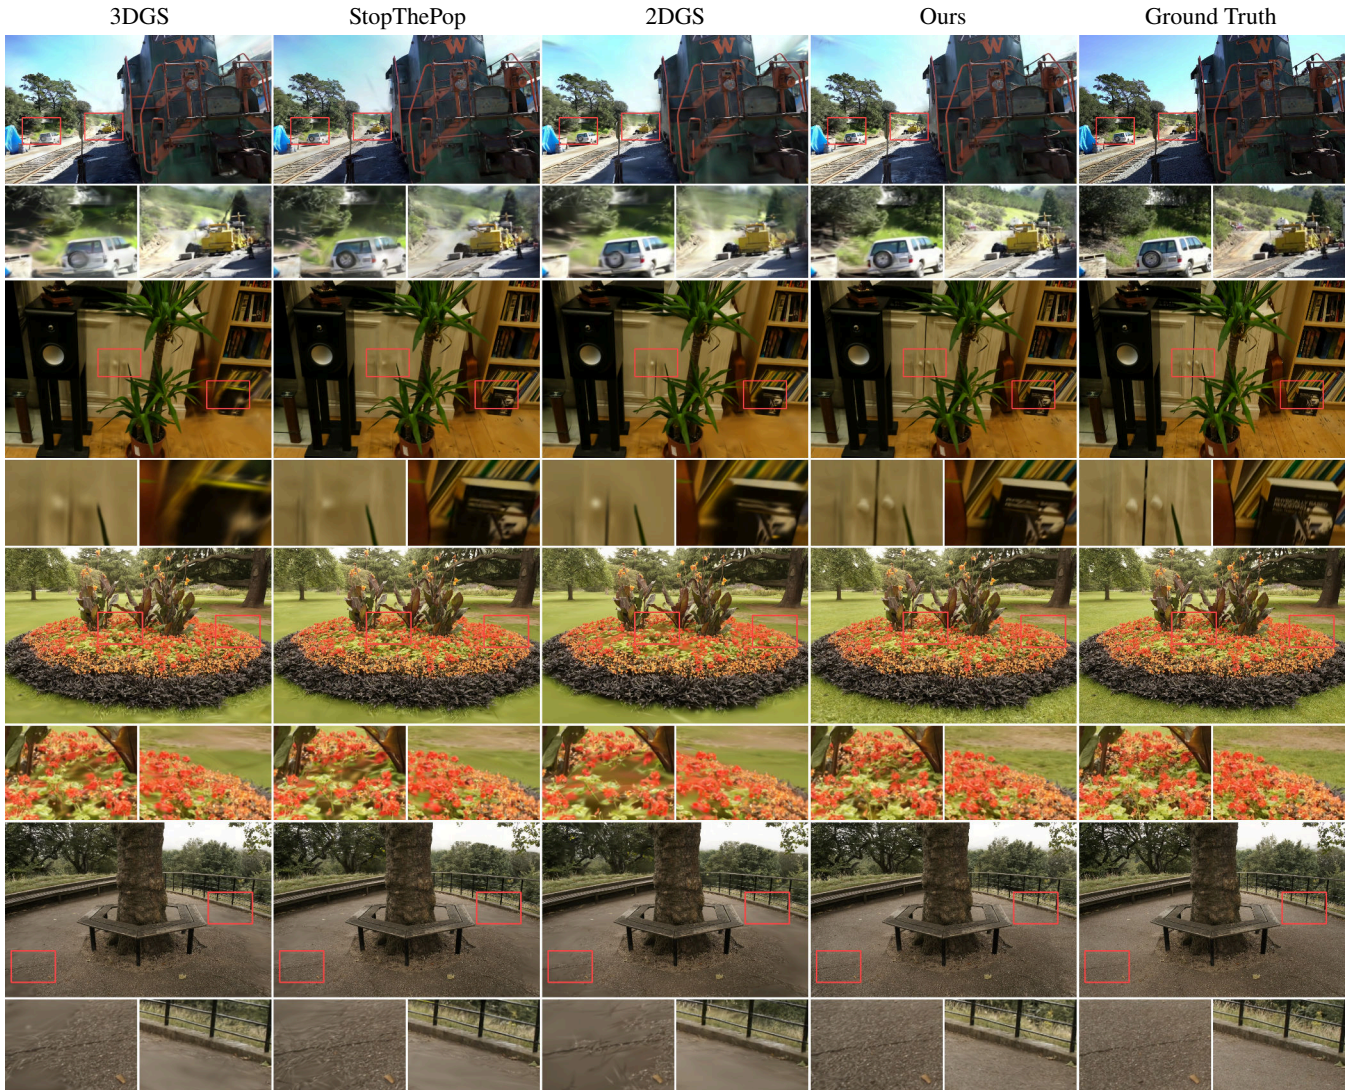


Figure 5: Visual comparisons of baselines that allow for real-time rendering.

(Sec. 3.3) is mainly a method for addressing popping artifacts without sacrificing computational efficiency. However, our perspective-accurate splat evaluation (Sec. 3.2) should prove beneficial across a much broader range of applications due to its accuracy. As can be seen in Tab. 3, using it alongside the approximately depth-ordered alpha blending from 3DGS [KKLD23] improves results for all image quality metrics. Note that the increase in speed compared to 3DGS is mostly due to improved global as well as per-pixel culling, which reduces the number of required computations.

Hybrid Transparency Ablations. To analyze the behavior of our hybrid transparency-based blending, we perform ablations on the nine scenes from the Mip-NeRF360 dataset [BMV*22]. The quantitative results shown in Tab. 4 clearly indicate that $K = 16$ is the optimal choice for our approach. We also show visual comparisons in Fig. 6 showing that blending without any sorting is not a valid option as it causes the foreground to become transparent. However, we do

find that background reconstruction seems to improve in this case. With our default configuration of $K = 16$, omitting tail computation during inference causes a small drop in image quality, which is most noticeable in the sky. Importantly, not using the tail during training causes catastrophic failure.

5. Discussion

The experiments validate that our approach improves robustness and accuracy of the optimization due to our perspective-accurate and numerically stable splat evaluation. Similarly, they show that hybrid transparency can effectively be integrated into differentiable rasterizers for 3D Gaussian splats to accelerate training and rendering without sacrificing image quality. Importantly, both of our contributions improve the rendering quality beyond what shows in the quantitative results for standardized benchmarks. To address this,



Figure 6: Visual comparisons for different model configurations regarding our hybrid transparency approach. Using a smaller core size K causes issues for reflective surfaces, as radiance fields commonly model these using semi-transparency. Disabling the order-independent tail only slightly reduces quality, especially in the sky, whereas not using it during optimization results in catastrophic failure.

Method	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	Training	FPS
3DGS	0.730	24.64	0.265	21m17s	293.2
Ours	0.742	24.62	0.237	12m09s	480.6
Ours w/o per-pixel HT	0.748	24.81	0.230	17m26s	448.7

Table 3: We showcase the potential of our perspective-correct splat evaluation by omitting our hybrid transparency-based blending in favor of the approximately depth-ordered alpha blending from 3DGS [KKLD23]. As it forgoes per-pixel ordering, this version of our model is subject to popping artifacts, and it is slightly less computationally efficient than our default configuration, but outperforms 3DGS across all metrics. The reported values are averaged results for the five outdoor scenes from the Mip-NeRF360 dataset [BMV*22].

Configuration	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	Training	FPS
$K = 8$	0.811	26.70	0.245	10m02s	437.9
$K = 32$	0.821	27.12	0.236	31m11s	372.4
Ours ($K = 16$)	0.822	27.17	0.234	10m07s	457.3
Inference w/o Tail	0.817	26.92	0.240	10m07s	467.0

Table 4: Ablations for our hybrid transparency approach computed on the nine scenes from the Mip-NeRF360 dataset [BMV*22].

our supplemental includes two videos that clearly demonstrate the advantage of our two-part contribution, which we encourage the reader to view. The perspective-correct splat evaluation allows our rendering to avoid artifacts that arise from the approximate projection used in 3DGS [KKLD23]. We find that the commonly used benchmark datasets largely fail to reveal this artifact making it invisible in quantitative evaluation. However, when inspecting trained 3D Gaussian splat models in a real-time viewer, it is very easy to find scenarios where the perspective-inaccurate projection causes disturbing artifacts.

Equally as noticeable are the popping artifacts caused by the approximate depth-ordering in 3DGS. StopThePop [RSP*24] uses a hierarchical sorting approach to resolve this issue. They sort splats globally based on their means [KKLD23] and then use a sliding window approach to sort the per-tile lists approximately. In contrast, our approach guarantees correct sorting order for the first K elements in each pixel and then uses fast blending without sorting for the rest.

Our results show that this strategy results in a good tradeoff between quality and performance: The correct sorting of the front-most fragments leads to high rendering quality and temporal stability, with little impact on performance. While the remaining elements from the tail are essential for robust optimization as demonstrated in Fig. 6, our results show that the much cheaper unsorted blending is sufficient to maintain high performance, very good rendering quality, and robust optimization.

The artifacts addressed by our approach are most apparent while the camera is moving (see our supplemental videos), which is an issue for VR/AR applications or fast-paced games where the user rarely keeps their head or the camera perfectly still. As VR also requires high frame rates to mitigate cybersickness, we believe that the combined advantages of our approach – fast, perspective-correct rendering without popping artifacts – make it a great fit for such applications. Another area where we think our approach could prove beneficial are 3DGS implementations for low-end devices, which currently require expensive CPU sorting. As our approach naturally supports depth peeling due to the fixed size of the core, it should be possible to implement a hardware-accelerated viewer for our approach that does not require a dedicated sorting routine.

Regarding reconstruction and image quality, we observe that the adaptive density control mechanism by Kerbl et al. [KKLD23] introduced for 3DGS has a significant influence while being very sensitive to changes. As outlined in Sec. 4.2, approaches evaluating splats along per-pixel viewing rays do not naturally provide the information used by the densification approach of 3DGS, i.e., the gradient of the splat’s position in view-space. After experimenting with the solutions presented in recent works [KRS*24, HYC*24, MLMP*24], we find scaling the gradients of the 3D positions by two times the distance to the camera [MLMP*24] works best for our approach. Nonetheless, we find that the direct connection between local splat attributes and the global densification scheme introduces issues regarding the robustness of all 3DGS-based approaches we investigated. Additional challenges, such as the photometric variation in images from the Tanks and Temples dataset [KPZK17] further amplify this issue.

Limitations and Future Work. Naturally, our approach is not without limitations. While sole use of alpha blending allows for early stopping based on a preset transmittance threshold [KKLD23], this is not an option for hybrid transparency by default. In our current implementation, we accumulate the RGB color and alpha of all

per-pixel contributions that are not part of the core. Intuitively, this is a problem when scenes become larger due to increase in depth complexity. However, we are confident that this issue can be resolved, e.g., by combining a depth prepass with a depth-weighting function as in WBOIT [MB13]. Additionally, we find that our approach achieves slightly lower PSNR than that of 3DGS, which we think is due to the absence of a screen-space aliasing filter or multi-sampling approach. As suggested by Huang et al. [HYC*24], we tried using the approach of Botsch et al. [BHZK05], but did not find it beneficial. We think it is possible to extend our splat evaluation to efficiently handle lens distortion thus enabling training on images without applying lossy undistortion operations during data pre-processing. Lastly, we would like to highlight that our hybrid transparency approach unveils an avenue for integrating otherwise too expensive ideas into real-time rendering of 3D Gaussians splats. While correct volume rendering of all contributing splats is not computationally feasible, it should be possible for the first $K = 16$ Gaussians. Similarly, we think that the fixed size of the core in our hybrid transparency approach allows to use more powerful appearance models [DHR*24] to improve image quality.

6. Conclusion

In this paper, we addressed two significant challenges prevalent in 3D Gaussian splat rendering. First, we introduced a novel approach for fast and perspective-accurate splat evaluation, which eliminates artifacts that arise from the affine approximation used by current 3D Gaussian Splatting implementations. Our approach elegantly avoids matrix inversion thus ensuring robust optimization and accurate rendering without numerical instabilities, even when splats degenerate. Second, we proposed an approach for view-consistent rendering through hybrid transparency, which uses fast, approximate per-pixel depth sorting that guarantees correct order for the first K contributions and thus prevents popping artifacts. Importantly, these improvements have a largely positive impact on performance, as evidenced by our analysis of training time, rendering speed, image quality metrics, and visual fidelity. We believe our two-part contribution can both independently and jointly improve accuracy and efficiency of 3D Gaussian splat rendering in the future.

Acknowledgements

We would like to thank Timon Scholz and Carlotta Harms for their help with comparisons and the supplemental material. The authors gratefully acknowledge financial support from the German Research Foundation (DFG) for the projects “Real-Action VR” (ID 523421583) and “Increasing Realism of Omnidirectional Videos in Virtual Reality” (ID 491805996), as well as from the L3S Research Center, Hanover, Germany. Linus Franke was supported by the 5G innovation program of the German Federal Ministry for Digital and Transport under the funding code 165GU103B.

References

- [ASK*20] ALIEV K.-A., SEVASTOPOLSKY A., KOLOS M., ULYANOV D., LEMPIŃSKI V.: Neural point-based graphics. In *ECCV* (2020). doi:10.1007/978-3-030-58542-6_42. 3
- [BCL*07] BAVOIL L., CALLAHAN S. P., LEFOHN A., COMBA J. L., SILVA C. T.: Multi-fragment effects on the gpu using the k-buffer. In *Proceedings of the 2007 symposium on Interactive 3D graphics and games* (2007), pp. 97–104. 4
- [BHZK05] BOTSCH M., HORNING A., ZWICKER M., KOBELT L.: High-quality surface splatting on today’s gpus. In *Proceedings Eurographics/IEEE VGTC Symposium Point-Based Graphics, 2005*. (2005), IEEE, pp. 17–141. 3, 10
- [BKL*24] BAGDASARIAN M. T., KNOLL P., LI Y.-H., BARTHEL F., HILSMANN A., EISERT P., MORGENSTERN W.: 3dgs.zip: A survey on 3d gaussian splatting compression methods, 2024. arXiv:2407.09510. 3
- [BM08] BAVOIL L., MYERS K.: Order independent transparency with dual depth peeling. *NVIDIA OpenGL SDK 1*, 12 (2008), 2–4. 4
- [BMT*21] BARRON J. T., MILDENHALL B., TANCIC M., HEDMAN P., MARTIN-BRUALLA R., SRINIVASAN P. P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. *ICCV* (2021). doi:10.1109/ICCV48922.2021.00580. 3
- [BMV*22] BARRON J. T., MILDENHALL B., VERBIN D., SRINIVASAN P. P., HEDMAN P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR* (2022). doi:10.1109/CVPR52688.2022.00539. 3, 6, 7, 8, 9, 12, 13
- [BMV*23] BARRON J. T., MILDENHALL B., VERBIN D., SRINIVASAN P. P., HEDMAN P.: Zip-nerf: Anti-aliased grid-based neural radiance fields. In *ICCV* (October 2023), pp. 19640–19648. doi:10.1109/ICCV51070.2023.01804. 3, 7, 13
- [Car84] CARPENTER L.: The a-buffer, an antialiased hidden surface method. In *Proceedings of the 11th annual conference on Computer graphics and interactive techniques* (1984), pp. 103–108. 4
- [CFHT23] CHEN Z., FUNKHOUSER T., HEDMAN P., TAGLIASACCHI A.: Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. In *CVPR* (2023). doi:10.1109/CVPR52729.2023.01590. 3
- [CXG*22] CHEN A., XU Z., GEIGER A., YU J., SU H.: TensorRF: Tensorial radiance fields. In *ECCV* (2022). doi:10.1007/978-3-031-19824-3_20. 3
- [DHR*24] DUCKWORTH D., HEDMAN P., REISER C., ZHIZHIN P., THIBERT J.-F., LUČIĆ M., SZELISKI R., BARRON J. T.: Smerf: Streamable memory efficient radiance fields for real-time large-scene exploration. *ACM Trans. Graph.* 43, 4 (July 2024). doi:10.1145/3658193. 3, 10, 12
- [ESSL10] ENDERTON E., SINTORN E., SHIRLEY P., LUEBKE D.: Stochastic transparency. In *3D '10: Proceedings of the 2010 symposium on Interactive 3D graphics and games* (New York, NY, USA, 2010), pp. 157–164. 4
- [Eve01] EVERITT C.: Interactive order-independent transparency. *White paper, nVIDIA* 2, 6 (2001), 7. 4
- [FEE21] FRIEDERICH S., EISEMANN E., EISEMANN M.: Layered weighted blended order-independent transparency. In *Proc. Graphics Interface (GI)* (May 2021), pp. 1–6. 4
- [FKYT*22] FRIDOVICH-KEIL S., YU A., TANCIC M., CHEN Q., RECHT B., KANAZAWA A.: Plenoxels: Radiance fields without neural networks. In *CVPR* (2022). doi:10.1109/CVPR52688.2022.00542. 3
- [FRF*23] FRANKE L., RÜCKERT D., FINK L., INNMANN M., STAMMINGER M.: VET: visual error tomography for point cloud completion and high-quality neural rendering. In *SIGGRAPH Asia Conference Papers* (New York, NY, USA, Dec. 2023), ACM. 3
- [FRFS24] FRANKE L., RÜCKERT D., FINK L., STAMMINGER M.: TRIPS: Trilinear Point Splatting for Real-Time Radiance Field Rendering. *Comput. Graph.* 43, 2 (2024). doi:https://doi.org/10.1111/cg.15012. 3, 6
- [GD98] GROSSMAN J. P., DALLY W. J.: Point sample rendering. In *Rendering Techniques '98: Proceedings of the Eurographics Workshop in Vienna, Austria, June 29–July 1, 1998* 9 (1998), Springer, pp. 181–192. 3

- [HFF*23] HARRER M., FRANKE L., FINK L., STAMMINGER M., WEYRICH T.: Inovis: Instant novel-view synthesis. In *SIGGRAPH Asia Conference Papers* (New York, NY, USA, Dec. 2023), ACM. doi:10.1145/3610548.3618216. 3
- [HFK*24] HALHBOHM F., FRANKE L., KAPPEL M., CASTILLO S., STAMMINGER M., MAGNOR M.: Inpc: Implicit neural point clouds for radiance field rendering. *arXiv preprint arXiv:2403.16862* (2024). 3
- [HSM*21] HEDMAN P., SRINIVASAN P. P., MILDENHALL B., BARRON J. T., DEBEVEC P.: Baking neural radiance fields for real-time view synthesis. *ICCV* (2021). doi:10.1109/ICCV48922.2021.00582. 3
- [HYC*24] HUANG B., YU Z., CHEN A., GEIGER A., GAO S.: 2D Gaussian Splatting for Geometrically Accurate Radiance Fields. In *SIGGRAPH 2024 Conference Papers* (2024), ACM. doi:10.1145/3641519.3657428. 1, 2, 3, 4, 6, 7, 9, 10, 12, 13
- [KKLD23] KERBL B., KOPANAS G., LEIMKÜHLER T., DRETTAKIS G.: 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM TOG* 42, 4 (July 2023). doi:10.1145/3592433. 1, 2, 3, 4, 6, 7, 8, 9, 12, 13
- [KLR*22] KOPANAS G., LEIMKÜHLER T., RAINER G., JAMBON C., DRETTAKIS G.: Neural point catacaustics for novel-view synthesis of reflections. *ACM TOG* (2022). doi:10.1145/3550454.3555497. 3
- [KMK*24] KERBL B., MEULEMAN A., KOPANAS G., WIMMER M., LANVIN A., DRETTAKIS G.: A hierarchical 3d gaussian representation for real-time rendering of very large datasets. *ACM Trans. Graph.* 43, 4 (July 2024). doi:10.1145/3658160. 3
- [KPLD21] KOPANAS G., PHILIP J., LEIMKÜHLER T., DRETTAKIS G.: Point-based neural rendering with per-view optimization. *CGF* (2021). doi:10.1111/cgf.14339. 3
- [KPZK17] KNAPITSCH A., PARK J., ZHOU Q.-Y., KOLTUN V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics* 36, 4 (2017). 6, 7, 9, 12, 13
- [KRS*24] KHERADMAND S., REBAIN D., SHARMA G., SUN W., TSENG J., ISACK H., KAR A., TAGLIASACCHI A., YI K. M.: 3d gaussian splatting as markov chain monte carlo. *arXiv preprint arXiv:2404.09591* (2024). 3, 9
- [KS23] KULHANEK J., SÄTTLER T.: Tetra-nerf: Representing neural radiance fields using tetrahedra. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)* (2023), pp. 18412–18423. doi:10.1109/ICCV51070.2023.01692. 3
- [LKLR24] LUITEN J., KOPANAS G., LEIBE B., RAMANAN D.: Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *3DV* (2024). 3
- [LZ21] LASSNER C., ZOLLHOFER M.: Pulsar: Efficient sphere-based neural rendering. In *CVPR* (2021). doi:10.1109/CVPR46437.2021.00149. 3, 6
- [MB13] MCGUIRE M., BAVOIL L.: Weighted Blended Order-Independent Transparency. *Journal of Computer Graphics Techniques (JCGT)* 2, 2 (December 2013), 122–141. URL: <http://jcgt.org/published/0002/02/09/>. 4, 10
- [MBDM97] MONTRYM J. S., BAUM D. R., DIGNAM D. L., MIGDAL C. J.: InfiniteReality: a real-time graphics system. In *Computer Graphics (Proc. ACM SIGGRAPH '97)* (1997), ACM Press, pp. 293–302. 6
- [MCTB13] MAULE M., COMBA J., TORCHELSEN R., BASTOS R.: Hybrid Transparency. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (2013), pp. 103–118. 4, 6
- [Mes07] MESHKIN H.: Sort-independent alpha blending. *GDC Talk* (2007). 4
- [MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG* 41, 4 (jul 2022). doi:10.1145/3528223.3530127. 3
- [MKKP18] MÜNSTERMANN C., KRUMPEN S., KLEIN R., PETERS C.: Moment-based order-independent transparency. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 1 (2018). doi:10.1145/3203206. 4
- [MKVC*24] MALLICK AND GOEL, KERBL B., VICENTE CARRASCO F., STEINBERGER M., DE LA TORRE F.: Taming 3dgs: High-quality radiance fields with limited resources. In *SIGGRAPH Asia 2024 Conference Papers* (2024). doi:10.1145/3680528.3687694. 7
- [MLMP*24] MOENNE-LOCOCOZ N., MIRZAEI A., PEREL O., DE LUTIO R., ESTURO J. M., STATE G., FIDLER S., SHARP N., GOJCIC Z.: 3d gaussian ray tracing: Fast tracing of particle scenes. *arXiv preprint arXiv:2407.07090* (2024). 2, 3, 6, 7, 9
- [MST*20] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHY R., NG R.: NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV* (2020). doi:10.1145/3503250. 2, 3
- [NMR*24] NIEMEYER M., MANHARDT F., RAKOTOSAONA M.-J., OECHSLE M., DUCKWORTH D., GOSULA R., TATENO K., BATES J., KAESER D., TOMBARI F.: Radsplat: Radiance field-informed gaussian splatting for robust real-time rendering with 900+ fps. *arXiv preprint arXiv:2403.13806* (2024). 3, 7
- [PD84] PORTER T., DUFF T.: Compositing digital images. *SIGGRAPH Comput. Graph.* 18, 3 (Jan. 1984), 253–259. doi:10.1145/964965.808606. 4
- [PZVBG00] PFISTER H., ZWICKER M., VAN BAAR J., GROSS M.: Surfels: Surface elements as rendering primitives. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (2000), pp. 335–342. 3
- [RFS22] RÜCKERT D., FRANKE L., STAMMINGER M.: Adop: Approximate differentiable one-pixel point rendering. *ACM TOG* (2022). doi:10.1145/3528223.3530122. 3
- [RGS*24] REISER C., GARBIN S., SRINIVASAN P. P., VERBIN D., SZELISKI R., MILDENHALL B., BARRON J. T., HEDMAN P., GEIGER A.: Binary opacity grids: Capturing fine geometric detail for mesh-based view synthesis. *arXiv preprint arXiv:2402.12377* (2024). 3
- [RL00] RUSINKIEWICZ S., LEVOY M.: Qsplat: A multiresolution point rendering system for large meshes. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (2000), pp. 343–352. 3
- [RSP*24] RADL L., STEINER M., PARGER M., WEINRAUCH A., KERBL B., STEINBERGER M.: StopThePop: Sorted Gaussian Splatting for View-Consistent Real-time Rendering. *ACM TOG* 4, 43 (2024). 1, 2, 3, 5, 6, 7, 9, 12, 13
- [RSV*23] REISER C., SZELISKI R., VERBIN D., SRINIVASAN P. P., MILDENHALL B., GEIGER A., BARRON J. T., HEDMAN P.: Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *SIGGRAPH* (2023). doi:10.1145/3592426. 3
- [SML11] SALVI M., MONTGOMERY J., LEFOHN A.: Adaptive transparency. In *Proceedings of the ACM SIGGRAPH Symposium on High Performance Graphics* (2011), pp. 119–126. 4
- [SSC22] SUN C., SUN M., CHEN H.: Direct voxel grid optimization: Superfast convergence for radiance fields reconstruction. In *CVPR* (2022). doi:10.1109/CVPR52688.2022.00538. 3
- [SV14] SALVI M., VAIDYANATHAN K.: Multi-layer alpha blending. In *Proceedings of the 18th meeting of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (2014), pp. 151–158. 4
- [SWBG06] SIGG C., WEYRICH T., BOTSCH M., GROSS M.: Gpu-based ray-casting of quadratic surfaces. In *Proceedings of the 3rd Eurographics / IEEE VGTC Conference on Point-Based Graphics* (Goslar, DEU, 2006), SPBG'06, Eurographics Association, p. 59–65. 2, 3, 4, 5
- [WHA*07] WEYRICH T., HEINZLE S., AILA T., FASNACHT D. B., OETIKER S., BOTSCH M., FLAIG C., MALL S., ROHRER K., FELBER N., KAESLIN H., GROSS M.: A hardware architecture for surface splatting. *ACM Trans. Graph.* 26, 3 (July 2007), 90–es. doi:10.1145/1276377.1276490. 2, 3, 4, 5
- [WYF*24] WU G., YI T., FANG J., XIE L., ZHANG X., WEI W., LIU W., TIAN Q., WANG X.: 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 20310–20320. 3

- [Wym16] WYMAN C.: Exploring and expanding the continuum of OIT algorithms. In *High Performance Graphics* (2016), pp. 1–11. [2](#), [4](#)
- [YCH*24] YU Z., CHEN A., HUANG B., SATTTLER T., GEIGER A.: Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 19447–19456. [3](#), [7](#)
- [YHR*23] YARIV L., HEDMAN P., REISER C., VERBIN D., SRINIVASAN P. P., SZELISKI R., BARRON J. T., MILDENHALL B.: Baked sdf: Meshing neural sdfs for real-time view synthesis. *arXiv* (2023). [3](#)
- [YLT*21] YU A., LI R., TANCIK M., LI H., NG R., KANAZAWA A.: PlenOc-trees for real-time rendering of neural radiance fields. In *ICCV* (2021). [doi:10.1109/ICCV48922.2021.00570](#). [3](#)
- [YSG24] YU Z., SATTTLER T., GEIGER A.: Gaussian opacity fields: Efficient and compact surface reconstruction in unbounded scenes. *arXiv preprint arXiv:2404.10772* (2024). [2](#), [3](#), [6](#), [7](#), [12](#), [13](#)
- [ZPVBG01a] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: Ewa volume splatting. In *Proceedings Visualization, 2001. VIS '01.* (2001), pp. 29–538. [doi:10.1109/VISUAL.2001.964490](#). [4](#)
- [ZPVBG01b] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: Surface splatting. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, Association for Computing Machinery, p. 371–378. [doi:10.1145/383259.383300](#). [3](#)

Appendix A: Frame Rate Benchmarking

Following Duckworth et al. [DHR*24], we compute frame rate by rendering the test-set images of each scene 100 times and averaging over the results. While the four compared baselines – 3DGS [KKLD23], StopThePop [RSP*24], 2DGS [HYC*24], and GOF [YSG24] – are all implemented on top of the original 3D Gaussian Splatting implementation by Kerbl et al. [KKLD23], not all of them implement optimized viewers for inference. As such, we intentionally avoid basing our benchmarks on these real-time viewers and instead came up with a way of removing most of the overhead caused by the PyTorch-based framework used for training and quality metric computation. Specifically, we apply a “baking” step before rendering, which integrates all activation functions into the splat parameters and ensures all spherical harmonics coefficients are stored in the same buffer. As shown in Tab. 5, this has a significant influence on the resulting frame rates of 3DGS and StopThePop, as they are highly optimized for real-time rendering, but a negligible influence for 2DGS and GOF. Seeing that the obtained numbers align with what was shown in the supplemental video from 3DGS, we think that this is a fair method for comparison, which does not rely on method-specific optimizations for inference. Note that in our implementation, the rendering is also controlled by a PyTorch-based framework.

Method	Mip-NeRF360		Tanks&Temples	
	Default	Optimized	Default	Optimized
3DGS [KKLD23]	197.2	316.6	178.4	216.2
StopThePop [RSP*24]	157.3	212.5	180.6	206.2
2DGS [HYC*24]	122.2	134.9	94.9	95.8
GOF [YSG24]	16.2	16.7	12.2	12.3

Table 5: For fair comparison of frame rates, we “bake” activation functions in splat parameters and ensure all spherical harmonics coefficients are stored in the same buffer before rendering.

Appendix B: Per-Scene Metrics

We report per-scene SSIM, PSNR, and LPIPS for all evaluated scenes from Mip-NeRF360 [BMV*22] as well as Tanks and Temples [KPZK17] in Table 6 and Table 7 respectively.

SSIM \uparrow on Mip-NeRF360 [BMV*22]										
Method	Bicycle	Flowers	Garden	Stump	Treehill	Bonsai	Counter	Kitchen	Room	Average
3DGS [KKLD23]	0.770	0.602	0.869	0.774	0.637	0.938	0.905	0.921	0.913	0.814
StopThePop [RSP*24]	0.767	0.600	0.865	0.774	0.634	0.941	0.905	0.925	0.918	0.814
2DGS [HYC*24]	0.732	0.570	0.841	0.756	0.615	0.931	0.891	0.915	0.906	0.795
GOF [YSG24]	0.787	0.634	0.868	0.793	0.640	0.937	0.901	0.916	0.911	0.821
Ours	0.784	0.629	0.867	0.792	0.637	0.941	0.902	0.923	0.920	0.822
Zip-NeRF [BMV*23]	0.772	0.637	0.863	0.788	0.674	0.952	0.905	0.929	0.929	0.828
PSNR \uparrow on Mip-NeRF360 [BMV*22]										
Method	Bicycle	Flowers	Garden	Stump	Treehill	Bonsai	Counter	Kitchen	Room	Average
3DGS [KKLD23]	25.25	21.52	27.41	26.55	22.49	31.98	28.69	30.32	30.63	27.20
StopThePop [RSP*24]	25.18	21.51	27.24	26.63	22.51	32.00	28.53	31.09	30.98	27.30
2DGS [HYC*24]	24.74	21.10	26.63	26.20	22.33	31.22	28.09	30.29	30.68	26.81
GOF [YSG24]	25.46	21.69	27.38	27.00	22.38	31.57	28.69	30.82	30.45	27.27
Ours	25.31	21.35	27.22	26.85	22.37	31.55	28.40	31.02	30.45	27.17
Zip-NeRF [BMV*23]	25.85	22.33	28.22	27.35	23.95	34.79	29.12	32.36	33.04	28.56
LPIPS \downarrow on Mip-NeRF360 [BMV*22]										
Method	Bicycle	Flowers	Garden	Stump	Treehill	Bonsai	Counter	Kitchen	Room	Average
3DGS [KKLD23]	0.229	0.366	0.118	0.244	0.367	0.253	0.262	0.158	0.289	0.254
StopThePop [RSP*24]	0.232	0.365	0.121	0.244	0.365	0.248	0.254	0.154	0.280	0.252
2DGS [HYC*24]	0.301	0.404	0.165	0.300	0.434	0.280	0.291	0.180	0.316	0.297
GOF [YSG24]	0.206	0.310	0.122	0.229	0.326	0.241	0.258	0.166	0.286	0.238
Ours	0.204	0.313	0.122	0.226	0.321	0.235	0.254	0.159	0.273	0.234
Zip-NeRF [BMV*23]	0.228	0.309	0.127	0.236	0.281	0.196	0.223	0.134	0.238	0.219

Table 6: Per-scene image quality metrics for the Mip-NeRF360 dataset [BMV*22] separated into outdoor and indoor scenes. Excluding Zip-NeRF, the three best results are highlighted in green in descending order of saturation.

SSIM \uparrow on Tanks and Temples [KPZK17]									
Method	Family	Francis	Horse	Lighthouse	M60	Panther	Playground	Train	Average
3DGS [KKLD23]	0.871	0.901	0.889	0.834	0.901	0.910	0.834	0.791	0.866
StopThePop [RSP*24]	0.872	0.899	0.889	0.834	0.901	0.910	0.835	0.791	0.866
2DGS [HYC*24]	0.860	0.884	0.883	0.820	0.887	0.902	0.806	0.768	0.851
GOF [YSG24]	0.873	0.898	0.891	0.825	0.901	0.909	0.840	0.793	0.866
Ours	0.872	0.896	0.893	0.832	0.901	0.909	0.839	0.786	0.866
Zip-NeRF [BMV*23]	0.893	0.918	0.909	0.835	0.905	0.908	0.846	0.813	0.878
PSNR \uparrow on Tanks and Temples [KPZK17]									
Method	Family	Francis	Horse	Lighthouse	M60	Panther	Playground	Train	Average
3DGS [KKLD23]	25.05	27.64	24.18	21.76	27.82	28.35	25.65	21.69	25.27
StopThePop [RSP*24]	24.52	26.91	23.88	21.65	27.70	28.38	25.49	21.37	24.99
2DGS [HYC*24]	24.72	26.29	23.86	21.49	26.89	27.94	24.29	20.89	24.55
GOF [YSG24]	25.05	26.85	24.30	21.17	27.84	28.47	25.09	21.28	25.01
Ours	24.19	25.32	23.84	21.96	27.82	28.40	24.69	20.78	24.62
Zip-NeRF [BMV*23]	28.05	29.55	27.67	22.31	28.86	28.84	26.62	22.10	26.75
LPIPS \downarrow on Tanks and Temples [KPZK17]									
Method	Family	Francis	Horse	Lighthouse	M60	Panther	Playground	Train	Average
3DGS [KKLD23]	0.236	0.344	0.239	0.291	0.244	0.241	0.291	0.320	0.276
StopThePop [RSP*24]	0.228	0.338	0.235	0.288	0.237	0.236	0.285	0.313	0.270
2DGS [HYC*24]	0.272	0.386	0.264	0.333	0.278	0.266	0.349	0.366	0.314
GOF [YSG24]	0.223	0.347	0.230	0.310	0.238	0.238	0.286	0.320	0.274
Ours	0.227	0.346	0.226	0.303	0.233	0.235	0.276	0.331	0.272
Zip-NeRF [BMV*23]	0.172	0.270	0.181	0.281	0.212	0.217	0.251	0.279	0.233

Table 7: Per-scene image quality metrics for the Tanks and Temples dataset [KPZK17]. Excluding Zip-NeRF, the three best results are highlighted in green in descending order of saturation.