

Volumetrically Consistent 3D Gaussian Rasterization

Chinmay Talegaonkar Yash Belhe Ravi Ramamoorthi Nicholas Antipa
 University of California San Diego

{ctalegaonkar, ybelhe, ravir, nantipa}@ucsd.edu

Abstract

Recently, 3D Gaussian Splatting (3DGS) has enabled photorealistic view synthesis at high inference speeds. However, its splatting-based rendering model makes several approximations to the rendering equation, reducing physical accuracy. We show that the core approximations in splatting are unnecessary, even within a rasterizer; we instead volumetrically integrate 3D Gaussians directly to compute the transmittance across them analytically. We use this analytic transmittance to derive more physically-accurate alpha values than 3DGS, which can directly be used within their framework. The result is a method that more closely follows the volume rendering equation (similar to ray-tracing) while enjoying the speed benefits of rasterization. Our method represents opaque surfaces with higher accuracy and fewer points than 3DGS. This enables it to outperform 3DGS for view synthesis (measured in SSIM and LPIPS). Being volumetrically consistent also enables our method to work out of the box for tomography. We match the state-of-the-art 3DGS-based tomography method with fewer points. Our code is publicly available at:

<https://github.com/chinmay0301ucsd/Vol3DGS>

1. Introduction

Recently, there has been tremendous progress in view synthesis methods using differentiable volume rendering [5, 27, 37, 39, 51]. These techniques can be broadly categorized as rasterization or ray-tracing methods. Both produce photorealistic results but with distinct trade-offs.

The speed of rasterization-based methods like 3D Gaussian Splatting (3DGS) [27] comes at the cost of reduced physical accuracy compared to their ray-tracing-based counterparts. Rasterization itself makes some unavoidable approximations for faster rendering, like per-tile sorting and no overlap handling (though these can be mitigated at some computational expense [22, 35, 41], which 3DGS inherits. However, 3DGS makes further approximations due to its use of Ellipsoidal Weighted Averaging (EWA) splatting [60]. These approximations do not merely

have a theoretical impact — they have important practical consequences too. For example, as we show Fig. 3, they make it harder for 3D Gaussians to represent opaque surfaces prevalent in 3D scenes.

Recent works have taken inspiration from ray-tracing-based methods to make 3DGS more physically-accurate [23, 25, 53]. However, these methods don't volumetrically integrate 3D Gaussians for color computation. Instead, like 3DGS, they rely on *splatting*, i.e., projecting 3D Gaussians to 2D Gaussians in screen space.

We identify that *we do not need splatting*, including all its approximations, to have a fast rasterization-based technique. Instead of splatting, we analytically evaluate the volume rendering equation by integrating Gaussians directly in 3D space (assuming correct sorting and no overlap) — all within a rasterizer. See Fig. 1 for a qualitative illustration.

We first show the volume rendering equation (without splatting approximations) can be expressed as an alpha-blending operation over 3D Gaussians (Sec. 4.1). Next, we derive the corresponding alpha values which are more physically-accurate than 3DGS and use them within 3DGS's rasterization framework (Sec. 4.2). We then show that our alpha values enable representing opaque objects better than 3DGS (Sec. 4.3).

In practice, our method consistently matches or exceeds 3DGS's view synthesis quality (as measured by LPIPS and SSIM) over a wide variety of scenes (Sec. 6.1). Our method also works out-of-the-box for 3DGS-based tomographic reconstruction and is able to match a recent state-of-the-art method's quality with a lower memory footprint (Sec. 6.2).

2. Related Work

2.1. Volume rendering with ray marching

Image-based rendering (view synthesis) works have a long history in computer graphics [10, 18, 32, 45]. Recently, NeRF and its followups [3, 16, 37, 39, 51] pioneered using differentiable volume rendering [36] and have achieved photorealistic reconstruction. These methods use neural networks and grids to represent a scene's volumetric density and color. Colors along rays are computed evaluating

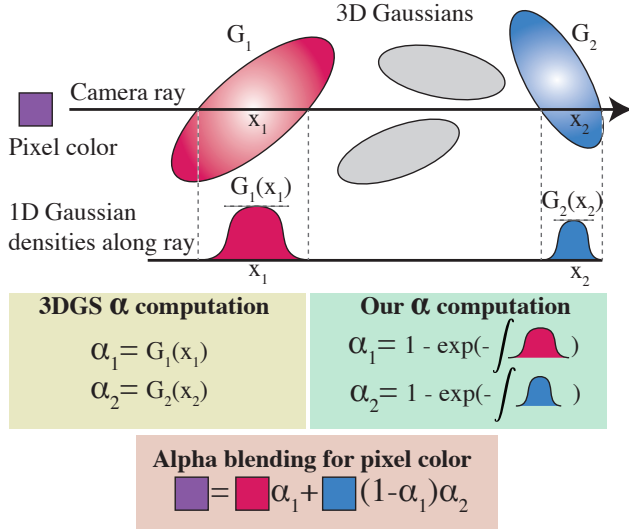


Figure 1. **Volumetrically consistent α Computation** Top row: 3D Gaussians that are not grey contribute to the pixel along the camera ray. Next, 1D Gaussian densities along the ray are used to compute α_i values required for color computation. 3DGS’s splatting approximates the volume rendering equation and sets $\alpha_i = G_i(x_i)$, that is the maximum density of the Gaussian long the ray. Our approach instead performs volumetrically consistent α computation by accumulating the density along the ray $\alpha_i = 1 - \exp(-\int G_i(x) dx)$ in accordance with the volume rendering equation, which we derive in Sec. 4. Both methods compute the final pixel color via α blending using their respective α values.

the volume rendering equation using ray marching. Ray marching typically requires several samples along the ray, each of which involves expensive neural network queries making the (differentiable) rendering process slow. There has been work on real-time NeRF rendering approaches [11, 14, 43, 44, 48, 49] but these are typically multi-stage approaches, and thus overall still slow.

2.2. Rasterizing point-based volumes

Rasterization-based methods [27, 31, 60] instead represent the scene using explicit primitives, like 3D Gaussians, which are projected onto screen space and composited together. Tile-based rasterizers like 3DGS are fast, because they exploit the coherence across pixels within the tile by making some assumptions: a) primitives don’t overlap; b) sorting order for all primitives is consistent across all pixels in a tile. 3DGS also assumes that blending weights (volumetric transmittance) can be approximated using 2D Gaussians, a result of splatting 3D Gaussians onto the screen [59].

Very recently, there have been several follow-ups to 3DGS to address these approximations within the rasterization framework. 2DGS [24] uses 2D (instead of 3D) Gaussian primitives; they avoid the affine approximation in splat-

ting by computing ray-2D-Gaussian intersection in closed form. Concurrent work, Gaussian Opacity Fields [52] estimates the volume rendering equation in world-space by placing a single sample per Gaussian at its maximum value along a ray. Others [13, 19, 26, 58] also compute ray-primitive intersections for more accurate depth.

While these methods eliminate the affine approximation, none of them compute physically-accurate transmittance. Our method, by analytically integrating 3D Gaussians (instead of splatting them onto the screen) is able to do so, resulting in a more physically-based rendering model.

Hamdi et al. [20] replace 3D Gaussians with generalized exponentials which are then splat onto the screen; their approach enables primitives to be more opaque, resulting in a more compressed representation. Radl et al. [41] reduce popping artifacts due to inconsistent primitive sorting and Hou et al. [22] sidestep primitive sorting altogether by using order-independent transparency. Many recent works [7, 28, 29, 34, 50] have improved the densification heuristics in 3DGS and reduced its dependency on initialization. Our method is a drop-in replacement for the alpha computation in 3DGS, making it compatible with all these recent works as well.

2.3. Ray-tracing point-based volumes

Recently, there have been several concurrent works that ray-trace point-based volumes for view synthesis [12, 33, 38]. These methods vary in their physical-correctness: a) Moenne-Loccoz et al. [38] approximate the volume rendering equation by placing samples at locations of maximum contribution for each primitive along the ray; b) Condor et al. [12], like us, derive a closed-form expression for transmittance across 3D Gaussian and Epanechnikov kernels; c) Mai et al. [33] compute the volume rendering equation exactly (including sorting and overlap handling) for constant density ellipsoids. These methods, however, are much slower to train and render than rasterization-based ones, and require specialized hardware for fast ray-tracing.

2.4. Computed Tomography

Tomographic reconstruction aims to estimate the underlying 3D density of an object from 2D projections. Classical analytic methods like FDK [15] invert the radon transform [42] and are fast but struggle with sparse measurements. Recently, NeRF-based approaches [9, 46, 54–56] have shown promising results, but are slow due to ray marching. Several 3DGS-based tomography approaches have been proposed [8, 17, 57]. Zha et al. [57] are the state-of-the art and achieve high quality by fixing integration bias in 3DGS’s Splatting routine.

3. Background and Motivation

In this section we review the volume rendering equation in Sec. 3.1 and describe all the approximations 3DGS makes to estimate it in Sec. 3.2.

3.1. Volume rendering

NeRF-based methods represent a 3D scene as an emissive volume with density (attenuation coefficient) $\sigma(\mathbf{x})$ and view-dependent color $c(\mathbf{x}, \mathbf{d})$. The resulting color along a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ (origin \mathbf{o} and direction \mathbf{d}) is given by the volume rendering equation

$$C(\mathbf{r}) = \int_0^\infty T(0, t) \sigma(\mathbf{r}(t)) c(\mathbf{r}(t), \mathbf{d}) dt, \quad (1)$$

where $T(a, b)$ is the (exponential) transmittance across the ray segment beginning at $\mathbf{r}(a)$ and ending at $\mathbf{r}(b)$

$$T(a, b) = \exp \left(- \int_a^b \sigma(\mathbf{r}(s)) ds \right). \quad (2)$$

3.2. Volume rendering of 3D Gaussians

We briefly recap the Elliptical Weighted Average (EWA) splatting algorithm [60] used by 3DGS [27], and point out the different approximations it makes by underlining them. The density field

$$\sigma(\mathbf{x}) = \sum_{i=1}^N \kappa_i G_i(\mathbf{x}) \quad (3)$$

is a weighted (κ_i) sum of unnormalized 3D Gaussians

$$G_i(\mathbf{x}) = \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}_i)^T \boldsymbol{\Sigma}_i^{-1} (\mathbf{x} - \boldsymbol{\mu}_i) \right\} \quad (4)$$

which are assumed to not overlap, have view-dependent color c_i , and be sorted front-to-back [60]. We also adopt these two assumptions. To compute pixel color, both EWA and 3DGS, linearize the exponential transmittance ($e^{-x} \approx 1 - x$), an assumption not strictly necessary for splatting. Our method relaxes this assumption and subsequent assumptions required for splatting. For splatting, both EWA and 3DGS assume Gaussians are not self-occluding, resulting in the alpha blending equation describing the color at pixel p

$$C(p) = \sum_{i=1}^N c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j). \quad (5)$$

Splatting i -th 3D Gaussian onto the image plane results in 2D opacity

$$\alpha_i = o_i \hat{G}_i(p), \quad (6)$$

where o_i is the peak opacity, restricted to be in $[0, 1]$, of 2D Gaussian $\hat{G}_i(p)$. The 2D mean μ'_i is the projection of 3D mean $\boldsymbol{\mu}_i$ onto the image plane and the 2D covariance is given by

$$\boldsymbol{\Sigma}' = \mathbf{J} \mathbf{W} \boldsymbol{\Sigma} \mathbf{W}^T \mathbf{J}^T. \quad (7)$$

\mathbf{J} is the affine approximation (Jacobian) to the camera projection matrix. For more details please see Zwicker et al. [60]. We stress that the values α_i are approximate, even when the sorting and non-overlapping assumptions are met, due to the three additional approximations made by splatting. We do not perform splatting, thereby relaxing the extra approximations 3DGS entails. Instead we compute α_i by analytically evaluating Eq. (2), which accumulates the density $\sigma(\mathbf{x})$ along ray $\mathbf{r}(t)$. This results in an efficient rendering method that matches or exceeds state-of-the-art results in both view synthesis and sparse-view tomography.

4. Method

Our method analytically computes the volume rendering equation along a ray for a mixture of 3D Gaussians. We first describe how this analytic integral (without splatting approximations) can be expressed as an alpha-blending operation in Sec. 4.1. We derive the corresponding alpha values in Sec. 4.2. Next, by swapping 3DGS's alpha computation with ours, we show how our method can produce more accurate renderings of opaque objects in Sec. 4.3.

4.1. Alpha blending without splatting

In this section we describe how our method, which analytically integrates transmittance, can also be written as alpha blending due to the properties of exponential transport. Substituting Eq. (3) into Eq. (1) results in the color for a pixel p (with the ray $\mathbf{r}(t)$ passing through its center) being

$$C(p) = \sum_{i=1}^N \int_{t_{in}}^{t_{if}} T(0, t) c_i \kappa_i G_i(\mathbf{r}(t)) dt. \quad (8)$$

Here $[t_{in}, t_{if}]$ are the limits of Gaussian i along the ray. Assuming the Gaussians are sorted (front to back) and non-overlapping, we can separate out the transmittance due to the previous Gaussians

$$C(p) = \sum_{i=1}^N \left(\prod_{j=1}^{i-1} \bar{T}_j \right) c_i \int_{t_{in}}^{t_{if}} T(t_{in}, t) \kappa_i G_i(\mathbf{r}(t)) dt. \quad (9)$$

The innermost integral weights the color contribution from Gaussian i , accounting for self-occlusion. The accumulated transmittance across previous Gaussians $j < i$ is

$$\bar{T}_j = T(t_{jn}, t_{jf}) = \exp \left(- \int_{t_{jn}}^{t_{jf}} \kappa_j G_j(\mathbf{r}(s)) ds \right). \quad (10)$$

The second equality holds because we assume the Gaussians are non-overlapping. Note that from the definition of exponential transmittance in Eq. (2) and differentiating, $dT(t_{in}, t) = -T(t_{in}, t)\sigma(\mathbf{r}(t))$, which corresponds directly to the integrand in Eq. (9). Therefore, this integral results in the i -th Gaussian’s accumulated opacity, $\alpha_i = 1 - \bar{T}_i$:

$$\int_{t_{in}}^{t_{if}} T(t_{in}, t) \kappa_i G_i(\mathbf{r}(t)) dt = - \int_{t_{in}}^{t_{if}} dT(t_{in}, t) \quad (11)$$

$$= T(t_{in}, t_{in}) - T(t_{in}, t_{if}) \quad (12)$$

$$= 1 - \bar{T}_i = \alpha_i. \quad (13)$$

Substituting this result into Eq. (9) yields the final alpha-blending equation

$$C(p) = \sum_{i=1}^N c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j). \quad (14)$$

We observe that this equation is identical to the alpha-blending equation, Eq. (5), adopted by 3DGS. The difference is that Eq. (14) requires accurately evaluating α_i , the accumulated density of the i -th primitive. Note that thus far, we do not rely on any specific property of 3D Gaussians; we only require the primitives G_i to be non-overlapping. This opens up a design space for other compact kernels which we leave for future work; we use 3D Gaussians. In the next section, we show that the α values in Eq. (14) can be computed analytically for 3D Gaussians. The result closely matches ray tracing results (assuming no overlap and correct sorting), but uses rasterization for efficient rendering, see Supp.A.

4.2. Analytic transmittance computation

The final step required to compute the color $C(p)$ is to estimate the transmittance Eq. (10), which requires integrating a 3D Gaussian along a 1D ray. The corresponding (scaled) 1D Gaussian $g(t)$ in ray-coordinates (i.e. parameterized by the distance t , ray origin \mathbf{o} and direction \mathbf{d}) is

$$g_j(t) = G_j(\gamma_j \mathbf{d}) \exp \left\{ \frac{-(t - \gamma_j)^2}{2\beta_j^2} \right\}, \quad (15)$$

where $G_j(\gamma_j \mathbf{d})$ is the maximum value of the Gaussian along the ray, and the 1D Gaussian has mean γ_j and variance β_j (proof in Supp.B)

$$\gamma_j = \frac{(\boldsymbol{\mu}_j - \mathbf{o})^T \Sigma_j^{-1} \mathbf{d}}{\mathbf{d}^T \Sigma_j^{-1} \mathbf{d}}, \quad \beta_j = \frac{1}{\sqrt{\mathbf{d}^T \Sigma_j^{-1} \mathbf{d}}}. \quad (16)$$

Now, the final integral necessary to compute Eq. (10) is

$$\int_{t_{jn}}^{t_{jf}} g_j(t) dt = \int_{t_{jn}}^{t_{jf}} G_j(\gamma_j \mathbf{d}) \exp \left\{ \frac{-(t - \gamma_j)^2}{2\beta_j^2} \right\} dt, \quad (17)$$

$$= G_j(\gamma_j \mathbf{d}) \sqrt{\frac{\pi}{2}} \beta_j \operatorname{erf} \left(\frac{t - \gamma_j}{\sqrt{2} \beta_j} \right) \Big|_{t_{jn}}^{t_{jf}}. \quad (18)$$

Assuming infinite support, the transmittance is given by

$$\bar{T}_j = \exp \left(-\kappa_j G_j(\gamma_j \mathbf{d}) \sqrt{2\pi} \beta_j \right). \quad (19)$$

Concurrent works [38, 52] also estimate the volume rendering equation by sampling the maximum value along the ray (instead of analytic integration), which approximates transmittance; Condor et al. [12] derive the same transmittance expression as us and use it for ray-tracing 3D Gaussians.

4.3. Benefits of accurate alpha computation

Comparing our alpha derivation to 3DGS’s, we note that 3DGS makes three extra approximations (see Sec. 3.2): it assumes a) linearization of the exponential; b) no-self occlusion; c) linearization of the covariance. Each of these make renderings less physically-based and hurt view synthesis quality.

The result of the first two approximations is that for 3DGS α_i is given by a 2D Gaussian, see Eq. (6). This 2D Gaussian only reaches its maximum opacity of 1 at its center (Fig. 2 (b), dotted), limiting 3DGS’s ability to represent opaque surfaces ($\alpha = 1$). As a consequence, our method is better at fitting piecewise constant textures than 3DGS for a fixed budget of Gaussians. We demonstrate this through examples in Fig. 2 and Fig. 3. In Fig. 2 we plot the opacity, α_i , of a volumetrically rendered Gaussian with varying density and scale. We observe that when increasing the density κ_i , our method results in opacity close to 1 over a much larger region (Fig. 2 (b), solid), thereby enabling it to better represent opaque surfaces (Fig. 2, bottom row).

Another consequence of splatting a 3D primitive to 2D is that it discards the z-scale (along the camera ray) of the Gaussian, which makes renderings invariant to it (Fig. 2 (a), dotted). This is physically inaccurate; larger z-scales should increase a Gaussian’s contribution, making it more opaque. Our method captures this effect, as shown in Fig. 2 (a) (solid).

Finally, the linearization of the covariance results in artifacts away from the point of linearization (the center of the Gaussian), as has been shown by several prior works [13, 19, 24, 26, 52, 58]. Since our method directly integrates Gaussians in ray-space, it doesn’t have this issue and the artifacts it causes.

The overall effect is that our method can better fit opaque objects. We demonstrate this in Fig. 3 where we have fit




Figure 2. **Opacity as a function of density and scale.** *Top row:* 1D cross sections of renderings for our method (solid) and 3DGS (dotted). Our method represents opaque objects better by two mechanisms: 1) increasing the scale along the camera ray (a); 2) increasing the volume density (b); both increase the flat region where $\alpha = 1$. Since 3DGS splats 3D Gaussians, irrespective of the scale (a) or the opacity (b), the cross-sections are all Gaussian with $\alpha = 1$ possible only at the center. *Bottom row:* We optimize the parameters of a single 3D Gaussian to fit a circle using 3DGS and our method. As we saw in the top row (a) and (b), regardless of scale or opacity, 3DGS can only render a 2D Gaussian in image-space, resulting in a blurry fit of the opaque object. On the other hand, our method adjusts density and scale to produce a more opaque rendering.

3D Gaussians to a single image (Fig. 3(a)). We observe that our method (Fig. 3(b)) better approximates sharp edges and constant regions versus 3DGS (Fig. 3(c)), which blurs edges and shows bleeding artifacts. We also validate this observation in a synthetic multiview setting in Supp.C.

5. Implementation details

Our method is implemented in the 3DGS framework using their SLANG.D rasterizer [2]. We found that this configuration produces better results for 3DGS (and is slightly slower) than the implementation used in their paper; for fair comparison, we use this configuration for both 3DGS and our method. In the rasterizer, we simply swap out 3DGS’s α computation with ours.

Reparameterizing density: Our method is able to make the density κ high to fit opaque shapes well. However, opaque shapes are close to piecewise-constant, resulting in very small gradients, which can hamper optimization. To




Figure 3. **Fitting piecewise constant shapes with equal numbers of 3D Gaussians.** (Please zoom in.) Our method, by increasing the volume density κ can make rendered 3D Gaussians close to opaque (b), which is a closer match to the ground truth (a), with LPIPS 0.005 (lower is better). 3DGS splats 3D Gaussians onto the screen as 2D Gaussians, which are only fully opaque at their center; this leads to artifacts especially near edges (see inset in second row) and much worse LPIPS 0.027.

handle this, similar to concurrent work [33], we reparameterize the density

$$\kappa = -\log(1 - 0.99\theta) \frac{1}{3} \left(\frac{1}{s_x} + \frac{1}{s_y} + \frac{1}{s_z} \right), \quad (20)$$

where $\theta \in [0, 1]$. The reparameterization promotes high density for small Gaussians and discourages it for larger ones, which we have found to benefit view synthesis quality.

Heuristics and hyperparameters: We make a few small changes to 3DGS’s adaptive densification process for our method: a) we halve density κ after splitting and cloning; b) we also split points with high density κ ; c) we prune points with low density (based on θ i.e. density before reparameterization); d) we disable opacity reset since we found it didn’t help; We list other hyperparameters and implementation details in Supp.D.

6. Applications

We developed our method primarily for view synthesis. However, since we analytically compute transmittance, our proposed rendering model also works out-of-the-box (with no modifications) for 3DGS-based tomography pipelines, where we match or exceed the state-of-the-art in performance, with fewer points. Our method is compatible with other extensions of 3DGS, for both view synthesis [28] and downstream applications like sonar reconstruction [40]; we leave these for future work.

Table 1. **Quantitative Comparison.** We compare our approach against prior work on the MipNerf-360, Tanks and Temples [30] and DeepBlending [21] datasets. Metrics for 3DGS-Slang and our method are obtained from our experiments using an Nvidia 3090-Ti GPU. We trained and evaluated ZipNeRF using an Nvidia A-40 GPU. We obtained the numbers directly from the paper authors for concurrent work EVER. All the other baseline numbers are taken from their respective papers. The results are colored as **best**, **second-best**, and **third-best** among point-based methods (see Sec. 6.1). We use the same train-test split as Kerbl et al. [27].

Dataset	Mip-NeRF360				Tanks&Temples				DeepBlending			
Method—Metric	PSNR↑	SSIM↑	LPIPS↓	FPS	PSNR↑	SSIM↑	LPIPS↓	FPS	PSNR↑	SSIM↑	LPIPS↓	FPS
Plenoxels [51]	23.08	0.626	0.463	6.79	21.08	0.719	0.379	13.0	23.06	0.795	0.510	11.2
INGP-Big [39]	25.59	0.699	0.331	9.43	21.92	0.745	0.305	14.4	24.96	0.817	0.390	2.79
M-NeRF360 [5]	27.69	0.792	0.237	0.06	22.22	0.759	0.257	0.14	29.40	0.901	0.245	0.09
Zip-NeRF [6]	29.16	0.830	0.179	0.03	28.45	0.921	0.189	0.09	30.54	0.929	0.201	0.045
Don’t Splat [12]	22.07	0.536	-	28-39	23.93	0.850	-	28-39	-	-	-	-
EVER [33]	27.51	0.825	0.194	36	23.60	0.870	0.160	37	29.58	0.908	0.308	44
3DGS-Slang [27]	27.52	0.813	0.222	159	23.64	0.850	0.176	225	29.78	0.906	0.248	163
GES [20]	26.91	0.794	0.250	186	23.35	0.836	0.198	210	29.68	0.901	0.252	160
Ours	27.30	0.813	0.209	136	23.74	0.854	0.167	185	29.72	0.908	0.247	172

6.1. View synthesis

We evaluate our method on MipNeRF-360, DeepBlending, and Tanks&Temples Datasets. For DeepBlending and Tanks&Temples we use the same scenes as 3DGS. We compare with two kinds of baselines. NeRF and voxel-based methods [4, 6, 16, 39], and point-based methods (rows below Zip-NeRF in Tab. 1) that include ray-tracing [12, 33], and rasterization-based 3DGS follow-ups that change the alpha computation similar to what we do [20]. Many works improve 3DGS along various axes [28, 29] which are orthogonal to our contribution. We only show comparisons with methods that directly change the alpha computation and are rasterization-based.

Reconstruction quality: We consistently match or outperform rasterization-based methods such as 3DGS and GES on SSIM and LPIPS. Since SSIM and LPIPS are more sensitive to edge quality and perceptual similarity, our method performs well on these metrics compared to PSNR, where the ability of 3DGS to render more diffuse textures skews results in its favor. Our method allows 3D Gaussian primitives to represent opaque textures better than 3DGS, leading to comparable or better quality. While GES produces fewer primitives, ours has better reconstruction quality and is compatible with the primitives proposed in their work. For point-based ray-tracing methods, we have better reconstruction performance than Condor et al. [12], which is closest to our work algorithmically. EVER [33] has better reconstruction quality, but since it uses ray tracing, it has a much lower FPS and higher training times than ours. Zip-NeRF [6] achieves state-of-the-art performance but is prohibitively slow to train and render. We report per-scene metrics from our method on all three datasets in Supp.E.

Runtimes and storage: Since our method is rasterization-based, we are faster than all the ray tracing-

based methods for rendering images (and similar to 3DGS / GES). While GES has the lowest storage requirements and highest FPS (though its reconstruction quality is worse than ours). Our training times are much lower than point-based ray-tracing and NeRF-based methods, but slightly slower than 3DGS, see Supp.F. The training slowdown is due to more complex computations in the forward and backward pass, and extra computations in the vertex shading phase for our method compared to 3DGS; more details are in Supp.F. We also include details on our ZipNeRF training setup in Supp.F. We show that our method outperforms 3DGS across a wide range of memory budgets in Supp.G. We ablate our method and present the results in Supp.H.

6.2. Tomography

We briefly review the goal and setup of computed tomography, see [57] for an in-depth review. The goal is to recover the underlying 3D density field $\sigma(\mathbf{x})$ of a scene from 2D images. The (logarithmic) intensity of a pixel, governed by Beer-Lamberts law is given by

$$I(p) = \log(I_0) - \log(I') = \int_{t_n}^{t_f} \sigma(\mathbf{r}(t)) dt. \quad (21)$$

where I_0 is the source intensity and $I'(\mathbf{r})$ is the detected intensity.

Using 3DGS out-of-the-box for tomography reconstruction produces poor results. This is because the splatting process isn’t physically accurate, as shown by R^2 -Gaussian Zha et al. [57]. They identify that unnormalized 3D Gaussians (as used by 3DGS), when splatted onto the screen, are view inconsistent. To fix this, they devise a specialized rendering model for tomography by normalizing the 3D Gaussians, which achieves state-of-the-art reconstruction results. Tomography reconstruction is slightly simpler than view synthesis since the forward rendering model does not depend on the ordering of the Gaussians,




Figure 4. **Qualitative Results on View Synthesis** We qualitatively show the strengths of our approach on several scenes from real datasets. Each row from top to bottom are for the scenes FLOWERS, ROOM, TRAINS,PLAYROOM, COUNTER and STUMP. Our method has less blurry regions compared to 3DGS in FLOWERS. Our ability to represent opaque regions accurately shows up prominently in ROOM. Our method is able to better capture fine details in the tiles in PLAYROOM. In COUNTER 3DGS misses details on the floor that our method is able to capture. Finally, in STUMP, we are able to capture sharper details in the grass that 3DGS misses.

and there is no emission or color involved, only the transmittance along a ray is measured. While Zha et al. [57] remove the integration bias from the alpha computation process by normalizing Gaussians, their approach still incurs the affine approximation in 3DGS.

Our method requires no changes to the alpha computa-

tion process and still works out of the box for tomography. We implement our method for tomography by integrating our rasterizer in Zha et al.’s codebase and evaluate on their real and synthetic datasets. Since our method provides a more accurate forward rendering model for tomography, we produce fewer points compared to Zha et al., while still




Figure 5. **Tomography visual results:** For our method (top row) and the splatting-based SOTA R^2 -Gaussian (Zha et al. [57]), we show slices of the reconstructed 3D density along the axial plane, perpendicular to the frontal views used during training. Each column corresponds to one of the distinct scenes from the datasets in [57]. Both methods produce visually similar reconstructions, with an average 2D PSNR of 29.94 (ours) and 30.06 (3DGS) for the displayed images.

Methods	Synthetic Dataset				Real Dataset			
	PSNR \uparrow	SSIM \uparrow	Time \downarrow	Pts	PSNR \uparrow	SSIM \uparrow	Time \downarrow	Pts
FDK [15]	22.99	0.317	-	-	23.30	0.335	-	-
SART [1]	31.14	0.825	1m47s	-	31.52	0.790	1m47s	-
ASD-POCS [47]	30.48	0.847	56s	-	31.32	0.810	56s	-
IntraTomo [55]	34.68	0.914	2h19m	-	35.85	0.835	2h18m	-
NAF [56]	33.91	0.893	31m1s	-	32.92	0.772	51m24s	-
SAX-NeRF [8]	34.33	0.905	13h3m	-	33.49	0.793	13h25m	-
R^2 -Gaussian [57]	34.97	0.921	18m55s	66.4K	35.01	0.835	15m57s	71K
Ours	35.08	0.922	17m39s	58.2K	34.98	0.835	30m23s	67.4K

Table 2. **Sparse-view (25 views) tomography.** We report 3D-PSNR and 3D-SSIM computed on the estimated density for synthetic and real data. The metrics for other baselines are taken from Zha et al. [57]. For both Zha et al. and ours, we report results at 30k iterations from our experiments.

matching them in performance on 3D PSNR and 3D SSIM on both real and synthetic datasets. We show results on their most challenging sparse-view (25 views) reconstruction in Tab. 2 and a subset of visual results in Fig. 5.

7. Limitations

Sorting and overlap: Since our method, like 3DGS, is rasterization-based, it inherits some of its issues. Per-tile (instead of per-pixel) primitive sorting and no overlap handling introduces popping artifacts. Several recent works [22, 35, 41] ameliorate this issue for 3DGS (their methods are readily applicable to us), but they cannot fully eliminate it like ray-tracing-based methods can [33].

Sensitivity to poor calibration: In cases of large errors in camera calibration or distortion, our method produces opaque artifacts while 3DGS degrades more gracefully by producing blurry reconstructions, as shown in Fig. 6. Refer




Figure 6. **Opaque artifacts** Our method is more sensitive than 3DGS to poor camera calibration and distorted images (Treehill from MipNerf-360). We produce opaque artifacts whereas 3DGS prefers a blurry solution with lower error (see red arrows).

to supp.I for a thorough analysis of the failure cases in the MipNerf-360 dataset, which exemplifies this effect.

8. Future work

Compact primitives: Our method applies to all primitives for which alpha in Eq. (10) can be evaluated analytically. Of particular interest are primitives with compact support, like the Epanechnikov kernel, which concurrent work [12] shows has analytic alpha values, since its compactness minimizes errors due to overlap.

Gaussian tile size: To compute overlap between primitives and tiles, we use the same heuristics as 3DGS, though this should ideally be adjusted to account for our method’s different rendering model.

9. Conclusion

We introduce a more physically-accurate version of 3DGS by computing alpha values via analytical 3D Gaussian integration instead of splatting. Our approach integrates the strengths of physically-based ray tracing with the efficiency of rasterization. It enables higher quality reconstruction than 3DGS (as measured by SSIM and LPIPS). We also show that since our method is a plug-and-play replacement for alpha computation in 3DGS, it can be directly used for other 3DGS-based applications such as tomography.

Acknowledgements: This work was supported in part by NSF grants 2341952 and 2105806. We also acknowledge ONR grant N0001423-1-2526, gifts from Adobe, Google, Qualcomm and Rembrand, support from the Ronald L. Graham Chair, and the UC San Diego Center for Visual Computing. We also thank Alexander Mai, Kaiwen Jiang, and Zachary Novack for helpful discussions and perspectives.

References

- [1] Anders H Andersen and Avinash C Kak. Simultaneous algebraic reconstruction technique (sart): a superior implementation of the art algorithm. *Ultrasonic imaging*, 6(1):81–94, 1984. 8
- [2] Sai Bangaru, Lifan Wu, Tzu-Mao Li, Jacob Munkberg, Gilbert Bernstein, Jonathan Ragan-Kelley, Fredo Durand, Aaron Lefohn, and Yong He. SLANG.D: Fast, Modular and Differentiable Shader Programming. *SIGGRAPH Asia*, 2023. 5
- [3] 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. *ICCV*, 2021. 1
- [4] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. *CVPR*, 2022. 6
- [5] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022. 1, 6
- [6] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. *ICCV*, 2023. 6
- [7] Samuel Rota Buló, Lorenzo Porzi, and Peter Kotschieder. Revising densification in gaussian splatting. *arXiv:2404.06109*, 2024. 2
- [8] Yuanhao Cai, Jiahao Wang, Alan Yuille, Zongwei Zhou, and Angtian Wang. Structure-aware sparse-view x-ray 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2, 8
- [9] Yuanhao Cai, Yixun Liang, Jiahao Wang, Angtian Wang, Yulun Zhang, Xiaokang Yang, Zongwei Zhou, and Alan Yuille. Radiative gaussian splatting for efficient x-ray novel view synthesis. In *European Conference on Computer Vision*, pages 283–299. Springer, 2025. 2
- [10] Shenchang Eric Chen and Lance Williams. *View Interpolation for Image Synthesis*. Association for Computing Machinery, New York, NY, USA, 1 edition, 1993. 1
- [11] Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. MobileNeRF: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. *CVPR*, 2023. 2
- [12] Jorge Condor, Sebastien Speierer, Lukas Bode, Aljaz Bozic, Simon Green, Piotr Didyk, and Adrian Jarabo. Don’t splat your gaussians: Volumetric ray-traced primitives for modeling and rendering scattering and emissive media, 2024. 2, 4, 6, 8
- [13] Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. High-quality surface reconstruction using gaussian surfels. In *ACM SIGGRAPH 2024 Conference Papers*. Association for Computing Machinery, 2024. 2, 4
- [14] Daniel Duckworth, Peter Hedman, Christian Reiser, Peter Zhizhin, Jean-François Thibert, Mario Lučić, Richard Szeliski, and Jonathan T Barron. SMERF: Streamable Memory Efficient Radiance Fields for Real-Time Large-Scene Exploration. *ACM Transactions on Graphics*, 2024. 2
- [15] Lee A Feldkamp, Lloyd C Davis, and James W Kress. Practical cone-beam algorithm. *Josa a*, 1(6):612–619, 1984. 2, 8
- [16] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 1, 6
- [17] Zhongpai Gao, Benjamin Planche, Meng Zheng, Xiao Chen, Terrence Chen, and Ziyang Wu. Ddgs-ct: Direction-disentangled gaussian splatting for realistic volume rendering. *arXiv preprint arXiv:2406.02518*, 2024. 2
- [18] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996. 1
- [19] Florian Hahlbohm, Fabian Friederichs, Tim Weyrich, Linus Franke, Moritz Kappel, Susana Castillo, Marc Stamminger, Martin Eisemann, and Marcus Magnor. Efficient perspective-correct 3d gaussian splatting using hybrid transparency. *arXiv*, 2024. 2, 4
- [20] 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. *CVPR*, 2024. 2, 6
- [21] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics*. 6
- [22] Qiqi Hou, Randall Rauwendaal, Zifeng Li, Hoang Le, Farzad Farhadzadeh, Fatih Porikli, Alexei Bourd, and Amir Said. Sort-free gaussian splatting via weighted sum rendering. *arXiv preprint arXiv:2410.18931*, 2024. 1, 2, 8
- [23] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *SIGGRAPH 2024 Conference Papers*. Association for Computing Machinery, 2024. 1
- [24] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 2, 4
- [25] Kaiwen Jiang, Yang Fu, Mukund Varma T, Yash Belhe, Xiaolong Wang, Hao Su, and Ravi Ramamoorthi. A construct-optimize approach to sparse view synthesis without camera pose. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 1
- [26] Kaiwen Jiang, Yang Fu, Mukund Varma T, Yash Belhe, Xiaolong Wang, Hao Su, and Ravi Ramamoorthi. A construct-optimize approach to sparse view synthesis without camera pose. *SIGGRAPH*, 2024. 2, 4
- [27] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023. 1, 2, 3, 6
- [28] 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. *arXiv:2404.09591*, 2024. 2, 5, 6

- [29] Sieun Kim, Kyungjin Lee, and Youngki Lee. Color-cued efficient densification method for 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 775–783, 2024. 2, 6
- [30] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics*, 2017. 6
- [31] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Trans. Graph.*, 39(6), 2020. 2
- [32] Marc Levoy and Pat Hanrahan. *Light Field Rendering*. Association for Computing Machinery, New York, NY, USA, 1 edition, 2023. 1
- [33] Alexander Mai, Peter Hedman, George Kopanas, Dor Verbin, David Futschik, Qiangeng Xu, Falko Kuester, Jon Barron, and Yinda Zhang. Ever: Exact volumetric ellipsoid rendering for real-time view synthesis. *arXiv preprint arXiv:2410.01804*, 2024. 2, 5, 6, 8
- [34] Saswat Subhajiya Mallick, Rahul Goel, Bernhard Kerbl, Francisco Vicente Carrasco, Markus Steinberger, and Fernando De La Torre. Taming 3dgs: High-quality radiance fields with limited resources. *arXiv preprint arXiv:2406.15643*, 2024. 2
- [35] Marilena Maule, João Comba, Rafael Torchelsen, and Rui Bastos. Hybrid transparency. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pages 103–118, 2013. 1, 8
- [36] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. 1
- [37] 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
- [38] Nicolas Moenne-Loccoz, Ashkan Mirzaei, Or Perel, Riccardo de Lutio, Janick Martinez Esturo, Gavriel State, Sanja Fidler, Nicholas Sharp, and Zan Gojcic. 3d gaussian ray tracing: Fast tracing of particle scenes. *arXiv:2407.07090*, 2024. 2, 4
- [39] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022. 1, 6
- [40] Ziyuan Qu, Omkar Vengurlekar, Mohamad Qadri, Kevin Zhang, Michael Kaess, Christopher Metzler, Suren Jayasuriya, and Adithya Pediredla. Z-splat: Z-axis gaussian splatting for camera-sonar fusion. *arXiv preprint arXiv:2404.04687*, 2024. 5
- [41] Lukas Radl, Michael Steiner, Mathias Parger, Alexander Weinrauch, Bernhard Kerbl, and Markus Steinberger. Stopthepop: Sorted gaussian splatting for view-consistent real-time rendering. *arXiv:2402.00525*, 2024. 1, 2, 8
- [42] Johann Radon. On the determination of functions from their integral values along certain manifolds. *IEEE transactions on medical imaging*, 5(4):170–176, 1986. 2
- [43] Christian Reiser, Rick Szeliski, Dor Verbin, Pratul Srinivasan, Ben Mildenhall, Andreas Geiger, Jon Barron, and Peter Hedman. Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *ACM Transactions on Graphics (TOG)*, 42(4):1–12, 2023. 2
- [44] Christian Reiser, Stephan Garbin, Pratul P. Srinivasan, Dor Verbin, Richard Szeliski, Ben Mildenhall, Jonathan T. Barron, Peter Hedman, and Andreas Geiger. Binary opacity grids: Capturing fine geometric detail for mesh-based view synthesis. *SIGGRAPH*, 2024. 2
- [45] Steven M. Seitz and Charles R. Dyer. View morphing. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, page 21–30, New York, NY, USA, 1996. Association for Computing Machinery. 1
- [46] Liyue Shen, John Pauly, and Lei Xing. Nerp: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 2
- [47] Emil Y Sidky and Xiaochuan Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology*, 53(17):4777, 2008. 8
- [48] Zian Wang, Tianchang Shen, Merlin Nimier-David, Nicholas Sharp, Jun Gao, Alexander Keller, Sanja Fidler, Thomas Müller, and Zan Gojcic. Adaptive shells for efficient neural radiance field rendering. *SIGGRAPH Asia*, 2023. 2
- [49] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P Srinivasan, Richard Szeliski, Jonathan T Barron, and Ben Mildenhall. Baked sdf: Meshing neural sdfs for real-time view synthesis. *SIGGRAPH*, 2023. 2
- [50] Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. Absgs: Recovering fine details in 3d gaussian splatting. *ACM Multimedia 2024*, 2024. 2
- [51] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. 2021. 1, 6
- [52] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient high-quality compact surface reconstruction in unbounded scenes. *arXiv preprint arXiv:2404.10772*, 2024. 2, 4
- [53] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient and compact surface reconstruction in unbounded scenes. *arXiv:2404.10772*, 2024. 1
- [54] Guangming Zang, Ramzi Idoughi, Congli Wang, Anthony Bennett, Jianguo Du, Scott Skeen, William L Roberts, Peter Wonka, and Wolfgang Heidrich. Tomofluid: Reconstructing dynamic fluid from sparse view videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1870–1879, 2020. 2
- [55] Guangming Zang, Ramzi Idoughi, Rui Li, Peter Wonka, and Wolfgang Heidrich. Intratomo: self-supervised learning-based tomography via sinogram synthesis and prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1960–1970, 2021. 8
- [56] Ruyi Zha, Yanhao Zhang, and Hongdong Li. Naf: Neural attenuation fields for sparse-view cbct reconstruction. In *International Conference on Medical Image Computing and*

- Computer-Assisted Intervention*, pages 442–452. Springer, 2022. [2](#), [8](#)
- [57] Ruyi Zha, Tao Jun Lin, Yuanhao Cai, Jiwen Cao, Yanhao Zhang, and Hongdong Li. R²-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. *arXiv preprint arXiv:2405.20693*, 2024. [2](#), [6](#), [7](#), [8](#)
 - [58] Baowen Zhang, Chuan Fang, Rakesh Shrestha, Yixun Liang, Xiaoxiao Long, and Ping Tan. Rade-gs: Rasterizing depth in gaussian splatting. *ArXiv*, abs/2406.01467, 2024. [2](#), [4](#)
 - [59] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa volume splatting. *Visualization*, 2001. [2](#)
 - [60] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa splatting. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):223–238, 2002. [1](#), [2](#), [3](#)