

Radiant Foam: Real-Time Differentiable Ray Tracing

Shrisudhan Govindarajan^{*1}, Daniel Rebain^{*2}, Kwang Moo Yi², Andrea Tagliasacchi^{1,3,4}

¹Simon Fraser University, ²University of British Columbia,

³University of Toronto, ⁴Google DeepMind, ^{*}equal contributions

radfoam.github.io

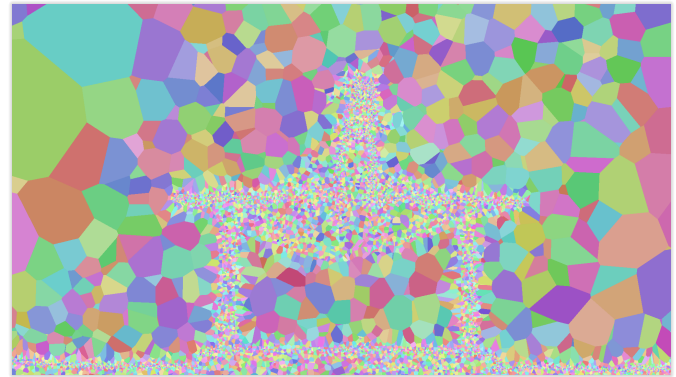


Figure 1. **Teaser** – We introduce the Radiant Foam differentiable 3D representation, which can be used to learn accurate radiance fields for any novel view synthesis applications (left). If we slice our foam along the plane highlighted by the red “laser”, we expose (right) the internal structure of our representation: a polyhedral mesh that provides an injective parameterization of the 3D domain. Our representation is a *foam*, as the polyhedral cell structure is analogous to a closed-cell foam which partitions space into regions physically separated by thin, flat walls. It is *radiant*, as each foam bubble emits a view-dependent radiance that can be used to model the plenoptic function.

Abstract

Research on differentiable scene representations is consistently moving towards more efficient, real-time models. Recently, this has led to the popularization of splatting methods, which eschew the traditional ray-based rendering of radiance fields in favor of rasterization. This has yielded a significant improvement in rendering speeds due to the efficiency of rasterization algorithms and hardware, but has come at a cost: the approximations that make rasterization efficient also make implementation of light transport phenomena like reflection and refraction much more difficult. We propose a novel scene representation which avoids these approximations, but keeps the efficiency and reconstruction quality of splatting by leveraging a decades-old efficient volumetric mesh ray tracing algorithm which has been largely overlooked in recent computer vision research. The resulting model, which we name Radiant Foam, achieves rendering speed and quality comparable to Gaussian Splatting, without the constraints of rasterization. Unlike ray traced Gaussian models that use hardware ray tracing acceleration, our method requires no special hardware or APIs beyond the standard features of a programmable GPU.

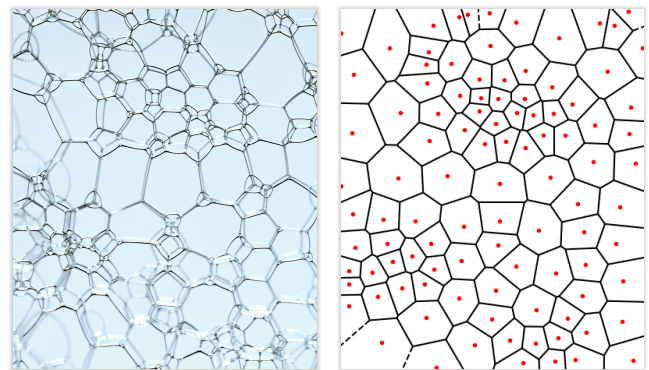


Figure 2. **Radiant Foam** – (left) In a stable foam, the pressure inside each bubble is roughly equal. The interfaces between bubbles settle into shapes that balance forces, leading to polygonal cells resembling Voronoi patterns (right). Our representation is nothing but a dense Voronoi tessellation of 3D space, where each point belongs to exactly one Voronoi cell. The position of Voronoi sites is differentiable, making it amenable to gradient-based optimization.



Figure 3. **Ray-based effects** – Ray tracing simplifies the implementation of many effects which are difficult to approximate with rasterization. To motivate our work, we show here examples of integrating reflections, refractions, and non-linear camera models into our rendering pipeline. Each would be complicated to achieve with rasterization, but requires only minor modification to our rendering code. S

1. Introduction

Neural radiance fields (NeRF) have revolutionized 3D computer vision by allowing the extraction of dense 3D representations from collections of 2D images [29]. In their early days, radiance fields suffered from low training and testing/rendering speed. Since then, researchers have developed more efficient models [31], as well as techniques for distilling models which enable highly efficient rendering/testing [5]. In particular, the real-time rendering performance was made possible by leveraging the rasterization pipeline readily available on GPU hardware. Soon thereafter, researchers proposed to also incorporate rasterization into the training process, leading to the development of the now wildly popular 3D Gaussian Splatting (3DGS) [19].

By comparing the development of radiance fields to that of classical graphics, one may argue that history is repeating itself: rasterization became ubiquitous in computer graphics because it is amenable to real-time rendering workloads. However, light effects that are trivial to implement in ray tracing (e.g. reflections, refractions, and transparency; see Fig. 3), are rather difficult to implement with rasterization.¹ Consider the “Graphics Gems” series of books² as a practical example, given how much of their content describes clever implementation tricks to express advanced light effects with rasterization engines. Nonetheless, since the introduction of dedicated hardware support (NVIDIA RTX) in 2018, real-time rendering engines have, at least partially, returned to simulate light transport via ray tracing.

Recently, research on 3DGS representations has flourished, and much of this research is seemingly trying to fix issues *caused* by rasterization, such as the removal of 3DGS popping artifacts [38], or the introduction of more complex camera projection models [16]. Others integrated ray tracing with 3DGS, so to accelerate rendering and enable more complex light behavior [30]. But these developments have not discouraged researchers from investigating new, better, 3D representations. The community craves a return to poly-

gonal meshes, as meshes are the unquestionable workhorse of modern computer graphics. And while we can find very interesting attempts at employing meshes for the modeling of radiance fields [11, 23, 49], as we will later discuss, none of these has realistically been able to oust 3DGS as *the* representation for learning radiance fields.

Rather than revisiting history, and proposing clever engineering tricks that enable rasterization to work slightly better, we take a drastically different approach. In particular, we highlight how, over two decades ago, Weiler et al. [57] showed that fields represented by volumetric meshes admit a very efficient ray-tracing algorithm which requires no special hardware. This approach has been subsequently overlooked in the resurgence of volume rendering methods, and we hope to re-introduce volumetric mesh models which can benefit from it to the differentiable rendering community.³ In this paper, we make this representation differentiable, and carefully design refinement techniques for the underlying mesh. These refinements allow us to accurately represent the surface of objects, and yet render efficiently by skipping empty portions of volume.

In a nutshell, our solution, which we name *Radiant Foam* (Fig. 2) provides 3DGS-like rendering speed and quality, but has a training modality based on rays that resembles the one from NeRF [29]. This implies that many NeRF techniques can now be seamlessly adapted to our method, with the significant advantage that the underlying geometry is *explicitly* represented by a volumetric mesh. We parameterize this volumetric mesh as a 3D Voronoi diagram (Fig. 1), which enables training of a mesh structure with dynamic connectivity through gradient descent by avoiding the discontinuities associated with discrete changes in other representations. We also propose a coarse-to-fine training approach, which enables rapid construction of mesh models with adaptive resolution.

¹After all, is this really that surprising, given that the physics of light is taught as the propagation of light rays through an environment?

²<http://www.graphicsgems.org>

³It is also interesting to note that the volume rendering tutorial by Max and Chen [28] that is cited in conjunction with Mildenhall et al. [29] describes volumetric rendering within the context of volumetric meshes that *partition* space, which is more similar to [27] than to [29].

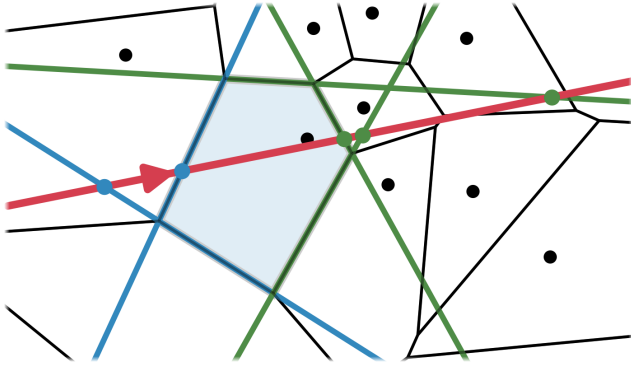


Figure 4. **Ray tracing foams** – When a ray (red) enters a cell, we iterate through all the planar cell faces to identify the face through which the ray exits. This exit intersection is the first intersection along the ray with a normal vector less than 90 degrees from the ray direction (green); other intersections are considered back-facing (blue) and ignored. As the faces each correspond to a neighboring cell, the tracing then proceeds by stepping into the cell associated with the exit intersection and repeating the process.

2. Related Work

Neural Radiance Fields (NeRFs) [29] represent 3D scenes as volumetric radiance fields encoded within coordinate-based neural networks. This representation allows querying the network at any spatial location to retrieve volumetric density and view-dependent color, facilitating the generation of photo-realistic novel views. The success of NeRFs has led to numerous follow-up works. For instance, significant effort has been devoted to enhancing training speed [31, 41], quality [1, 2], and to extract surfaces from the representation [54, 59]. Finally, several works investigated ways to speed up the inference by baking the neural fields to more performant representations [5, 8, 42, 48, 55]. While achieving high quality and fast rendering speeds, these methods often employ multi-stage training procedures.

Point-based rasterization. Point-based rendering with primitives such as circles, spheres, or ellipsoids [36, 44, 61] laid the foundations for point-based rasterization. Differentiable rendering using depth-based blending has also been extended to volumetric particles in Pulsar [24], which employs sphere-based differentiable rasterization to render scenes with millions of spheres in real time. While Pulsar provides a differentiable representation, Kerbl et al. [19] provides more expressive primitives in the form of soft, anisotropic Gaussians, which can be differentially optimized to fit the scene content. This approach has inspired several follow-up works aimed at reducing its reliance on strong initialization [21], reducing rendering time and memory footprints [9, 32, 35], enabling the reconstruction of surfaces [12, 15], and improving their ability to be trained for large spatial extents [20, 43].

Extensions of 3DGS. While significant progress has been

made, these approaches still inherit the inherent limitations of rasterization. They struggle to handle camera distortions, model secondary lighting effects, or simulate sensor-specific properties like rolling shutter or motion blur. Recent efforts have aimed to address these challenges. For example Niemeyer et al. [33] distilled a Zip-NeRF [3] into a 3DGS so to model distorted cameras and rolling shutter effects. To capture secondary lighting effects, recent works have explored incorporating occlusion information into spherical harmonics for each Gaussian [10, 25], or leveraging shading models and environment maps [17]. These methods either rely on rasterization during inference [10], or during training [17, 25] hence inheriting the limitations of rasterization. For complex lens effects, Seiskari et al. [47] modeled motion blur and rolling shutter by *approximating* them in screen space through rasterization and pixel velocities. Our approach introduces a principled framework for efficient ray tracing of volumetric primitives, overcoming the aforementioned limitations and enabling the simulation of effects like reflection, refraction, and camera distortion.

2.1. Ray tracing of volumetric primitives

Ray tracing has been a cornerstone of photo-realistic rendering since its introduction [58]. It enables the accurate simulation of light interactions with scene geometry, making it indispensable for applications requiring effects such as shadows, reflections, and refractions. Modern advancements in hardware have further accelerated ray tracing, allowing for real-time applications [18]. Recently, Moenne-Loccoz et al. [30] and Condor et al. [6] integrated the NVIDIA OptiX ray tracer with 3DGS [19] for fast ray tracing, while Mai et al. [26] employed it with constant ellipsoids. While these methods perform real-time ray tracing, they suffer from the tendency of 3DGS [19] to produce overlapping primitives which degrade the quality of the hierarchical acceleration structure and increase the number of intersections per ray. Our method represents the scene using a non-overlapping polyhedral mesh without the need for a secondary acceleration structure, thereby efficiently avoiding these limitations.

2.2. Delaunay triangulation and Voronoi diagrams

Voronoi diagrams have been used for representing information in various fields. In computer graphics, Delaunay triangulations [7] and Voronoi diagrams [53] have been extensively studied for meshing (surfaces and volumes), spatial partitioning [13, 50, 56], and differentiable simulation [34]. Recently, they have also found novel applications in the realm of differentiable rendering. DMTet [49] introduced a deformable tetrahedral grid, generated using Delaunay triangulation, as an underlying 3D representation, while Tet-Splatting [11] implements a differentiable rasterizer for this representation. Similarly, DMesh [51] proposes a differential representation to model Delaunay triangulations

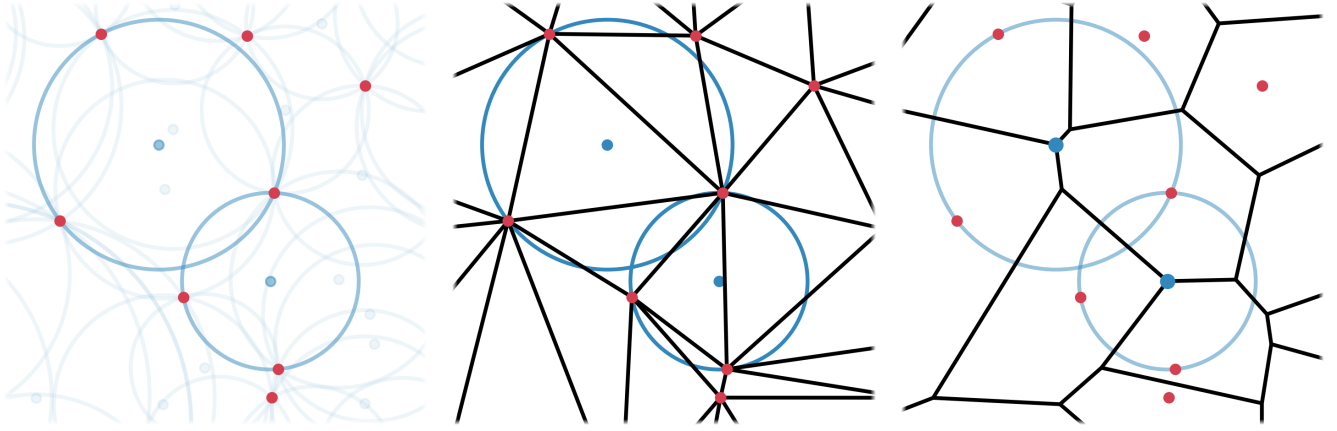


Figure 5. **Delaunay and its dual Voronoi** – (left) Given a set of points (red) in \mathbb{R}^N , we can find circumcircles (blue) which each pass through $N+1$ points. (center) The set of all circumcircles which contain no points on their interior defines the Delaunay triangulation, where the $N+1$ points tangent to each circumcircle form a simplex. (right) These circumcircles also describe the Delaunay triangulation’s dual, the Voronoi diagram: the centers of circumcircles tangent to a point become the vertices of the Voronoi cell containing that point.

to optimize for mesh connectivity. DeRF [39] leverages Voronoi diagrams to spatially decompose a scene, resulting in faster training and inference, while Tetra-NeRF [23] employs a tetrahedral Delaunay mesh to model scenes more effectively. While these methods inherit the ray-marching approach of [29] for volume rendering, they share a common limitation: slow rendering speeds caused by many MLP evaluations required per ray. In contrast, with our method we propose to *optimize* the mesh topology generated by the Voronoi diagram as a *differentiable* polygonal mesh. Our efficient ray tracing of volumetric particles also significantly accelerates the rendering speed.

3. Method

Our method addresses the now-familiar problem of constructing representations of scenes from image collections. As with NeRF [29] and its numerous successors, we achieve this reconstruction by gradient-based optimization of a differentiable scene representation. In the following sections we propose a volumetric mesh-based differentiable representation, and explain how we are able to effectively optimize it from image supervision.

3.1. Volume rendering

Volume rendering has become the workhorse of modern differentiable scene reconstruction methods. Volume rendering allows all points in space to make a continuously varying contribution to the observed color of viewing rays which pass through those points. The effect of this contribution is controlled by a density field, which creates occlusions, and a radiance field, which determines the brightness of light observed. Unlike many alternative rendering formulations, volume rendering of continuous fields is fully continuous with respect to all degrees of freedom, including both the

viewpoint, as well as the density and radiance field values. This property makes it very amenable to gradient-based optimization. Volume rendering is defined by an integral over the segment (t_{\min}, t_{\max}) of a viewing ray. Specifically the observed color $\mathbf{c}_{\mathbf{r}}$ for a ray \mathbf{r} is:

$$\mathbf{c}_{\mathbf{r}} = \int_{t_{\min}}^{t_{\max}} T(t) \cdot \sigma(\mathbf{r}(t)) \cdot \mathbf{c}(\mathbf{r}(t)) dt, \quad (1)$$

$$T(t) = \exp\left(-\int_{t_{\min}}^t \sigma(\mathbf{r}(u)) du\right), \quad (2)$$

where $\mathbf{r}(t)$ denotes the point in 3D space at distance t along ray \mathbf{r} , and $\sigma(\cdot)$ and $\mathbf{c}(\cdot)$ denote the density and radiance fields respectively; see [52] for more details.

Piecewise constant volumes. In the case of *piecewise constant* $\sigma(\cdot)$ and $\mathbf{c}(\cdot)$, the integral can be expressed as a sum over all N ray segments with constant field values:

$$\mathbf{c}_{\mathbf{r}} = \sum_{n=1}^N T_n \cdot (1 - \exp(-\sigma_n \delta_n)) \cdot \mathbf{c}_n dt, \quad (3)$$

$$T_n = \prod_{j=1}^n \exp(-\sigma_j \delta_j), \quad (4)$$

where δ_n is the width of segment n . This formulation admits a simple implementation as a loop over the segments in order of depth. NeRF-based methods typically use Eq. (3) as an *approximation* to Eq. (1), with segments sampled according to some importance sampling scheme, e.g. see [2].

Conversely, in our model these forms are exactly equivalent. We propose to leverage the algorithm by Weiler et al. [57] along with a model of constant field values within the cells of a volumetric mesh (see Fig. 4), arriving at a representation for which Eq. (3) gives the *exact* volume rendering

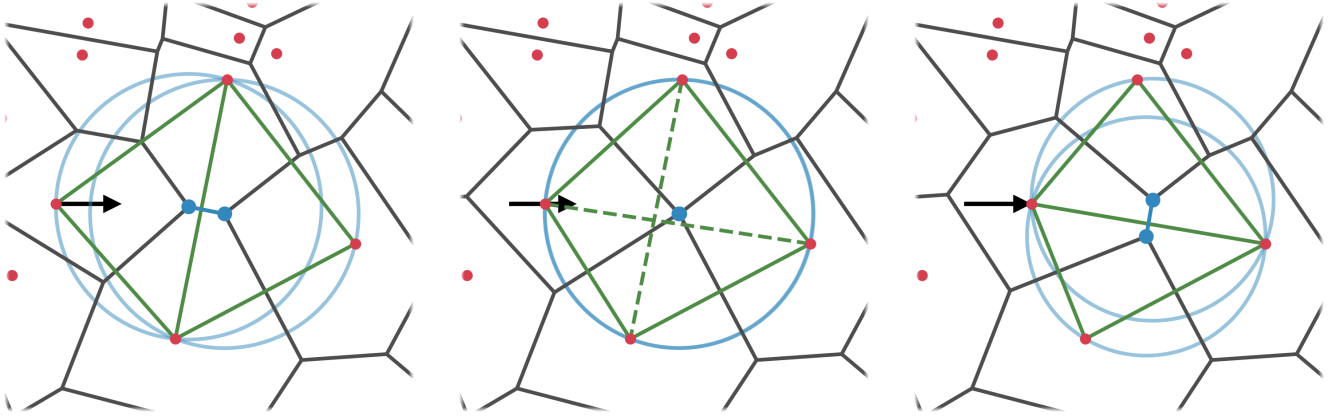


Figure 6. **Edge flips** – The connectivity of the Delaunay graph (green) is sensitive to small positional perturbations, leading to “edge flips” in the triangulation. These discrete changes occur at configurations where the circumcenters (blue) of two neighbouring simplices become identical. In the dual Voronoi diagram, this configuration also corresponds to a discrete change, but the cell boundary which changes has an area of zero at the moment of the flip (center). Consequently, while there are still discrete changes in the adjacency structure of the Voronoi diagram, the *shapes of the cells* vary continuously with the positions of the points (red), which enables gradient-based optimization.

result. In the volume rendering literature, this is also known as regular tracking [37, Chapter 15.2.2]. While this choice is highly advantageous in avoiding any complicated or expensive sampling schemes, we must take great care to not interfere with the continuity of the representation, which is critical for gradient-based optimizers. In particular, the volumetric mesh itself receives gradients only through the segment widths δ_n , which are determined by the locations of ray intersections with cell boundaries. It is therefore critical that these boundary intersections *vary continuously* with the optimizable parameters of the model.

3.2. Differentiable mesh representation

Constructing mesh representations of volumes or surfaces through gradient-based optimization is hardly a new problem, and many methods tackle it using a variety of strategies; see Shen et al. [49] and citations therein. Our method, however, encounters and must solve a challenge that most mesh optimization methods either avoid, or solve with discrete optimization techniques. Generally, meshes are determined by degrees of freedom in two groups: vertex locations \mathcal{V} and cells \mathcal{C} (i.e. connectivity, or mesh topology). Assuming all other parameters are fixed, optimizing for vertex locations is typically straightforward, as the intersections between rays and cell boundaries vary smoothly with vertex locations. The true challenge arises if one desires to also optimize the connectivity \mathcal{C} of the representation.

The Delaunay triangulation. Because mesh connectivity is inherently discrete, we can not optimize it directly with gradient-based methods. To avoid this problem, we can define connectivity in terms of the vertex locations, such that each possible configuration of vertices corresponds to a *unique* set of cells. The most obvious choice for this mapping is the *Delaunay triangulation* [7], which in 3D

comprises the set of all tetrahedral cells $c_i \in \mathcal{C}$ which may be formed from four vertices such that their circumcenters contain no additional vertices (the so-called “Delaunay criterion”). This construction is unique for point sets in general position (those lacking groups of 5 or more co-spherical points), and is easily computed using well-known and efficient algorithms; see Section 3.4. This strategy encounters two significant problems if we wish to use it as a basis for differentiable volume rendering: (1) the discrete nature of the mesh connectivity is not entirely avoided, as the boundaries of Delaunay cells undergo discrete “flips” whenever a vertex enters the circumsphere of another cell. These flips introduce *discontinuities* into the optimization landscape, which interfere with the convergence of gradient descent; (2) the number of tetrahedra in this model is not fixed in this model, and therefore it is not straightforward to associate each cell with optimizable σ and c values as required for volume rendering. The latter issue could be solved by associating field values with vertices (vs. tets), and interpolating those values within the cell, but this complicates volume rendering.

The Voronoi diagram. Rather than utilize the Delaunay mesh directly and suffer its limitations, we instead look to the dual graph of the Delaunay triangulation: the *Voronoi diagram* [53]. As shown in Figure 5, this is constructed by placing a vertex at the circumcenter of each Delaunay tetrahedron; it partitions space into convex polyhedral cells $c_i \in \mathcal{C}$ consisting of points which share a nearest neighbour among the primal vertices⁴ \mathbf{p}_i of the Delaunay triangulation:

$$c_i = \{x \in \mathbb{R}^3 : \arg \min_j \|x - \mathbf{p}_j\| = i\}. \quad (5)$$

This may at first seem like a strange choice... after all, *if the*

⁴Also known as Voronoi “sites” or “seeds”.

Delaunay Triangulation is unavoidably discontinuous in optimization, will not its dual also suffer this issue? The answer to this question lies in the fact that any discrete change in the connectivity of the Voronoi Diagram exactly coincides with the point that the affected cell faces attain zero surface area. As shown in Figure 6, the discrete flips are effectively *hidden* within these zero-volume regions of space, and the resulting field representation remains completely continuous with respect to the primal vertex locations. For a rigorous treatment of this property of the Voronoi diagram, see the proof by Reem [40] that Voronoi cells in \mathbb{R}^d are stable with respect to primal vertex positions. The number of cells in the Voronoi diagram is also constant regardless of configuration, which makes the association of σ and c values to cells... trivial. The resulting model is then effectively a *learnable point cloud*, not dissimilar to the formulation of 3D Gaussian Splatting [19], though lacking the per-point covariance matrix. The only remaining issue is that our ray tracing algorithm [27] expects *tetrahedral* cells. We therefore modify the algorithm to handle the more general case of convex cells. This modified tracing process pre-fetches the neighboring vertices of each vertex, thereby allowing an efficient iteration over the cell faces to find ray intersections. For more detail on this algorithm, see the supplementary material.

3.3. Optimization

Similarly to Kerbl et al. [19], the (mostly) local nature of Voronoi cells renders the optimization landscape more prone to local minima. We follow a similar strategy, by first carefully initializing the optimization, and then adaptively *densifying* and *pruning* Voronoi sites. Additionally, to promote the formation of surface-like densities, we also employ a regularization objective similar to the *distortion* loss [2] commonly used by NeRF methods.

Densification. Similarly to Kerbl et al. [19], to initialize training we start with a sparse point cloud obtained from COLMAP [46]. Over training, we perform densification and pruning operations to control the number of Voronoi sites and their density, allowing the model to adaptively re-allocate representational capacity to areas of space with more geometric and/or photometric detail. We observe that gradients of the reconstruction loss with respect to Voronoi site locations can be used to identify cells which are underfitting the training signal. We therefore use the norm of this gradient multiplied by the approximate radius of the cell as a measure of which cells require further densification. Inspired by Kheradmand et al. [21], we select the candidates for densification by sampling a Multinomial distribution with probability mass function proportional to this measure.

Pruning. Towards building a parsimonious representation, we remove cells from the Voronoi diagram that do not con-

tribute to rendering. However, it is not sufficient to simply delete “empty” cells (i.e. zero density), as the geometry of Voronoi cells is determined by the positions of adjacent sites, even when the density in those cells is zero. To accurately represent object boundaries, it is therefore essential to retain cells with near-zero density that define the boundary. For this reason, our pruning strategy removes Voronoi sites that have very low density, *and* are surrounded by neighbors with very low density. This pruning ensures that we eliminate sites that neither contribute to nor define the surfaces, thereby maintaining the accuracy of the object boundaries.

Training objectives. Similarly to the *distortion* loss of Mip-NeRF 360 [2], we apply a regularization on the distribution of contribution to the volume rendering integral along the ray. This additional loss function encourages the density to concentrate at surfaces and reduces visible “floater” artifacts. This objective is computed as

$$\mathcal{L}_{\text{quantile}} = \mathbb{E}_{t_1, t_2 \sim \mathcal{U}[0,1]} [|W^{-1}(t_1) - W^{-1}(t_2)|], \quad (6)$$

where $W^{-1}(\cdot)$ denotes the quantile function (inverse CDF) of the volume rendering weight distribution along the ray. This form has the same effect as the distortion loss [2], but avoids the need for a quadratic nested sum which would increase the computational cost and memory footprint of training. Denoting with $\mathcal{L}_{\text{rgba}}$ the typical L2 photometric reconstruction loss (including ground truth opacity equal to one for all rays), our overall training objective is

$$\mathcal{L} = \mathcal{L}_{\text{rgba}} + \lambda \mathcal{L}_{\text{quantile}} \quad (7)$$

3.4. Implementation details

We implemented our method in Python and C++ within the PyTorch framework, with custom CUDA kernels for ray tracing, Delaunay triangulation, and other operations requiring high efficiency. This implementation includes an interactive viewer and a (very) low overhead renderer, which we used to measure frame rates. Importantly, nothing in our implementation is dependent on dedicated ray tracing hardware, or the OptiX library, which is required by methods like [30]. Therefore, with some engineering effort, our entire rendering loop could easily be implemented in a portable rendering framework like WebGL. For more details on training and Voronoi optimization please refer to Section 7 in the supplementary material.

4. Experiments

We evaluated our algorithm on a total of 9 real-world scenes sourced from two publicly available datasets. Specifically, we utilized the complete set of scenes from the Mip-NeRF 360 dataset [2] except for two private scenes (flowers and treehill) and two scenes from the Deep Blending dataset [14]. These datasets contain scenes with a diverse range of capture styles, including bounded indoor scenes and expansive,

	MipNeRF 360 [2]				Deep Blending [14]			
	PSNR↑	SSIM↑	LPIPS↓	FPS↑	PSNR↑	SSIM↑	LPIPS↓	FPS↑
Rasterization								
3DGS* [19]	28.69	0.87	0.22	293	29.41	0.90	0.32	319
Mip-Splatting [60]	29.39	0.88	0.20	241	29.47	0.90	0.32	260
3DGS-MCMC [21]	29.72	0.89	0.19	302	29.71	0.90	0.32	662
Ray Tracing / Marching								
Plenoxels [45]	23.63	0.67	0.44	< 30	23.06	0.80	0.51	< 30
iNGP-Big [31]	26.75	0.75	0.30	< 30	24.96	0.82	0.39	< 30
MipNeRF360 [2]	29.23	0.84	0.21	< 1	29.40	0.90	0.25	< 1
3DGRT** [30]	28.71	0.85	0.25	78	29.23	0.90	0.32	119
RadiantFoam (v1)	28.47	0.83	0.21	200	28.95	0.89	0.26	301

Table 1. **Novel view synthesis** – We evaluate our method’s accuracy in reconstructing held-out views for two standard datasets. Our method has similar performance to 3DGS and 3DGRT, while providing *significantly* higher frame rates than the latter. FPS is measured as an average on the test set for each scene at their native resolutions. *For 3DGS, we report corrected LPIPS scores provided by Bulò et al. [4] **Note that 3DGRT FPS results were measured with an RTX 6000 Ada GPU. We report the results from the original publication, as the code is not open source.

unbounded outdoor environments. For Mip-NeRF 360, to make our results compatible with [19, 30], we downsample images for the indoor scenes by a factor of two, and the outdoor scenes by four. For Deep Blending scenes, we use the original image resolutions. All frame rates were measured on a consumer-grade RTX 4090 GPU.

Metrics. We assess each method using three widely recognized image quality metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS). In the supplementary web page we include rendered video paths for selected scenes, showcasing views significantly different from the input images.

Quantitative results. We report our quantitative results for the Mip-NeRF 360 [2] and Deep Blending [14] in Table 1. In terms of quality, our method achieves results comparable to, or slightly below, those of 3DGS [19] and 3DGRT [30] (the state of the art differentiable ray tracing method). However, our method excels in rendering speed. As shown in Table 1, our efficient ray-tracing implementation achieves in some cases over 300 FPS, more than twice as fast as 3DGRT (119 FPS), while maintaining a similar rendering speed to rasterization methods.

Qualitative results. We include in the supplementary material additional qualitative comparison showing the differences between our reconstruction and that of 3D Gaussian splatting. These results highlight how 3DGS tends to over-smooth some details and miss some view dependence, but also more gracefully extrapolates to under-observed regions.

SfM	Densify	Prune	Quantile	Point opt	Bonsai	Garden	Playroom	Mean
✗	✓	✓	✓	✓	29.65	25.83	26.34	27.00
✗	✗	✓	✓	✓	20.23	18.88	19.55	19.36
✓	✓	✗	✓	✓	32.25	26.58	29.46	29.15
✓	✓	✓	✗	✓	29.62	25.35	29.59	27.90
✓	✓	✓	✓	✗	30.51	25.03	28.35	27.96
✓	✓	✓	✓	✓	32.15	26.58	29.59	29.15

Table 2. **Ablation table** – We evaluate the impact of various components in our method by systematically excluding them and analyzing the reconstruction quality (PSNR↑) on the Bonsai and Garden scenes from MipNeRF 360 [2], as well as the Playroom scene from Deep Blending [14].

4.1. Ablation

We isolated the different contributions and choices we made and constructed a set of experiments to measure their effect. Specifically, we test the following aspects of our algorithm: initialization from SfM, our densification and pruning strategies, and regularization loss. The quantitative effect of each choice is summarized in Table 2.

Initialization from SfM. We assess the importance of initializing our Voronoi sites from the SfM point cloud. For this, we initialize our representation with 2^{17} points from a normal distribution with a standard deviation of 25. We observe that our method performs relatively well even without the SfM points. Instead, it degrades mainly in the background and tends to have more floaters in regions that are sparsely covered in the training views, see Figure 7.

Densification and pruning. Due to numerical approximations, our triangulation algorithm can fail when processing very close (or identical) points. This limitation prevents us from initializing our representation directly with the final set of points obtained by duplicating and perturbing the Structure from Motion (SfM) points, as SfM typically produces numerous closely spaced points. Consequently, for the ablation study on densification, we utilize the random initialization strategy described earlier to initialize our representation with the intended number of points. We observe that our method significantly underperforms without densification, resulting in an under-represented scene where resources are not adequately allocated to regions with complex geometry or texture. In the ablation study of our pruning strategy, we cease pruning those Voronoi sites that are neither surface nor boundary points. We notice that this pruning approach has minimal impact on the quality of the renderings because the number of prunable points is very low. By starting with a sparse point set and progressively densifying only in under-represented areas, we ensure that points are allocated exclusively to necessary regions, resulting in few prunable points.

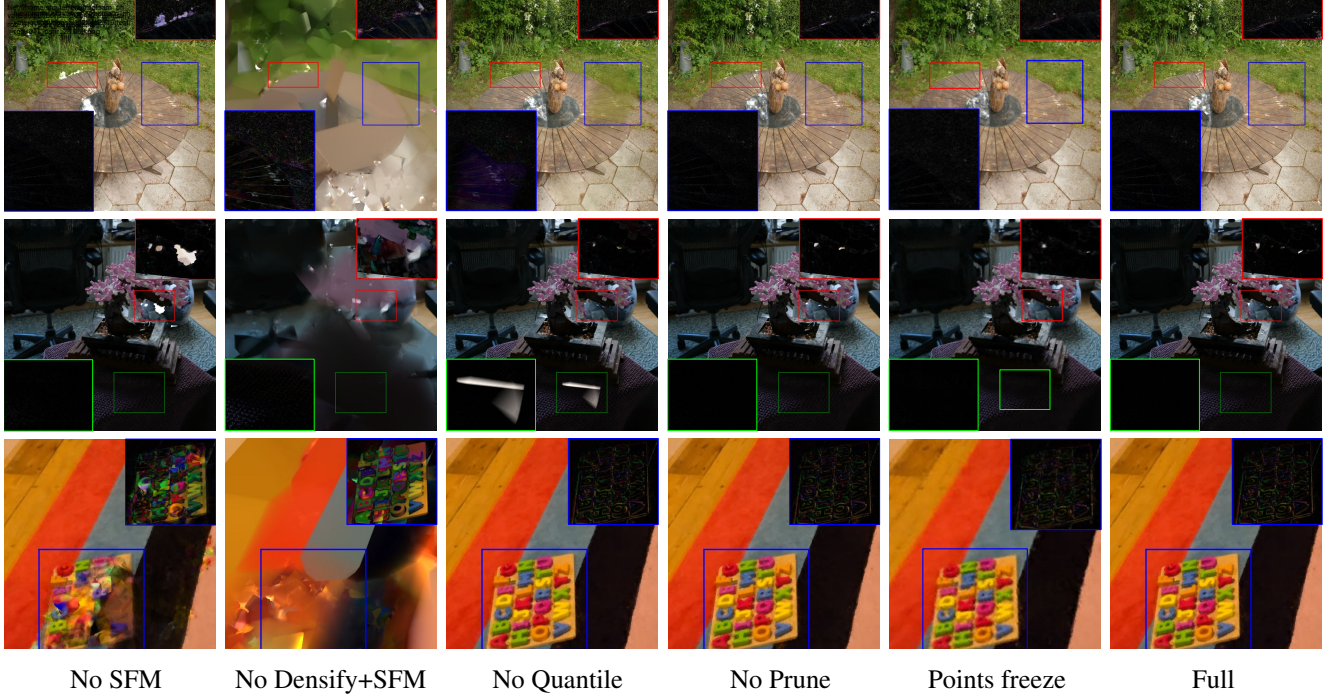


Figure 7. **Ablations** – We qualitatively analyze the impact of excluding various components from our method. For better visualization, we zoom in on specific image regions and present their corresponding error maps. *No SFM*: While the method performs relatively well without initialization, it degrades in sparsely covered regions of the training views. *No Densify+SFM*: Without densification, the method significantly underperforms, leading to under-represented scenes. *No Quantile*: Disabling the quantile loss introduces floaters, degrading rendering quality. *No Pruning*: Ceasing point pruning has minimal impact due to the low number of prunable points. *Points freeze*: Fixing points position leads to sub-optimal mesh topology, resulting in a lack of fine details.

Quantile loss. We disable the quantile regularization while training, which results in floater artifacts that degrade rendering quality for novel views. Some scenes are more affected by this than others, as the floaters result from data-dependent ambiguities, but on average we find that the quantile regularization improves reconstruction metrics.

Point optimization. For this ablation, we keep the positions of the points fixed during training, leading to a suboptimal mesh topology that lacks adaptability to scene geometry. This rigidity prevents the mesh from dynamically optimizing its structure, resulting in degraded novel view quality, as the fixed topology fails to accurately represent fine details.

5. Conclusions

We have introduced Radiant Foam, a novel representation that allows real-time differentiable ray tracing. The core of our method is a foam structure of polyhedral cells, which allows efficient volumetric mesh ray tracing algorithms to be applied *without* relying on dedicated hardware such as NVIDIA RT cores. We allow these cells to be continuously optimized by parameterizing them as a Voronoi diagram, which we show to be differentiable under volume rendering. By doing so, we have achieved similar modeling quality as

3D Gaussian Splatting without sacrificing either the benefits of a true ray tracing-based volume renderer, or the fast rendering speed of rasterization-based renderers.

Limitations and future work. While the Voronoi-based representation we have proposed is very effective at constructing foam models through continuous optimization, the space of possible foam models which could be used in our rendering pipeline is much larger than what is parameterized by Voronoi. Most notably, our current model always requires that cell boundaries be equidistant between neighbouring points, which leads to many small, empty cells being needed to define a surface. Future work could potentially relax this requirement by generalizing beyond Voronoi diagrams. Other open research questions include how to compose multiple foam models together efficiently and accounting for varying illumination, how to model dynamic content instead of static scenes, how to enable editing of scenes, and how to integrate generative modeling with our representation. Progress in these directions could make foam models relevant in real-time ray tracing applications currently dominated by triangle meshes, as we have already found that foam-based ray tracing can exceed the performance of dedicated ray tracing hardware.

Acknowledgements. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant [2023-05617], NSERC Collaborative Research and Development Grant, the SFU Visual Computing Research Chair, Google Research, Digital Research Alliance of Canada, and Advanced Research Computing at the University of British Columbia. We would like to thank George Kopanas, Lily Goli, Sherwin Bahmani, Alex Evans, Thomas Muller, Bernhard Kerbl, Vincent Sitzmann, Forrester Cole, Or Litany, David Fleet, and Delio Vicini for their feedback and/or early research discussions.

References

- [1] 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. 3
- [2] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022. 3, 4, 6, 7, 12, 13
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. *ICCV*, 2023. 3
- [4] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. Revising densification in gaussian splatting, 2024. 7, 12
- [5] Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. In *CVPR*, 2023. 2, 3
- [6] Jorge Condor, Sebastian 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. *ACM Trans. Graph.*, 44(1), 2025. 3
- [7] Boris Delaunay. Sur la sphere vide. *Bulletin of the Academy of Sciences of the USSR Classe des Sciences Mathematiques et Naturelles*, 1934. 3, 5
- [8] 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 (TOG)*, 43(4):1–13, 2024. 3
- [9] Zhiwen Fan, Kevin Wang, Kairun Wen, Zehao Zhu, DeJia Xu, and Zhangyang Wang. Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. *arXiv preprint arXiv:2311.17245*, 2023. 3
- [10] Jian Gao, Chun Gu, Youtian Lin, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. Relightable 3d gaussian: Real-time point cloud relighting with brdf decomposition and ray tracing. *arXiv preprint arXiv:2311.16043*, 2023. 3
- [11] Chun Gu, Zeyu Yang, Zijie Pan, Xiatian Zhu, and Li Zhang. Tetrahedron splatting for 3d generation. In *NeurIPS*, 2024. 2, 3
- [12] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5354–5363, 2024. 3
- [13] Leonidas Guibas and Jorge Stolfi. Primitives for the manipulation of general subdivisions and the computation of voronoi. *ACM Trans. Graph.*, 4(2):74–123, 1985. 3
- [14] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. 37(6):257:1–257:15, 2018. 6, 7, 12, 13
- [15] 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*, New York, NY, USA, 2024. Association for Computing Machinery. 3
- [16] Letian Huang, Jiayang Bai, Jie Guo, Yuanqi Li, and Yanwen Guo. On the error analysis of 3d gaussian splatting and an optimal projection strategy. *arXiv preprint arXiv:2402.00752*, 2024. 2
- [17] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuexin Ma. Gaussianshader: 3d gaussian splatting with shading functions for reflective surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5322–5332, 2024. 3
- [18] Pushkar Joshi, Mark Meyer, Tony DeRose, Brian Green, and Tom Sanocki. Harmonic coordinates for character articulation. *ACM transactions on graphics (TOG)*, 26(3):71–es, 2007. 3
- [19] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 2, 3, 6, 7, 12, 13
- [20] Bernhard Kerbl, Andreas Meuleman, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, and George Drettakis. A hierarchical 3d gaussian representation for real-time rendering of very large datasets. *ACM Transactions on Graphics*, 43(4), 2024. 3
- [21] 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. *NeurIPS*, 2024. 3, 6, 7, 12, 13
- [22] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 12
- [23] Jonas Kulhanek and Torsten Sattler. Tetra-nerf: Representing neural radiance fields using tetrahedra. In *ICCV*, 2023. 2, 4
- [24] Christoph Lassner and Michael Zollhofer. Pulsar: Efficient sphere-based neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1440–1449, 2021. 3
- [25] Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. Gs-ir: 3d gaussian splatting for inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21644–21653, 2024. 3
- [26] Alexander Mai, Peter Hedman, George Kopanas, Dor Verbin, David Futschik, Qiangeng Xu, Falko Kuester, Jonathan T. Barron, and Yinda Zhang. Ever: Exact volumetric ellipsoid rendering for real-time view synthesis, 2024. 3

- [27] Gerd Marmitt and Philipp Slusallek. Fast ray traversal of tetrahedral and hexahedral meshes for direct volume rendering. In *Proceedings of the Joint Eurographics/IEEE VGTC conference on Visualization*, pages 235–242, 2006. 2, 6
- [28] Nelson Max and Min Chen. Local and global illumination in the volume rendering integral. Technical report, Lawrence Livermore National Lab. (LLNL), Livermore, CA (United States), 2005. 2
- [29] 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. *ECCV*, 2020. 2, 3, 4
- [30] 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. *ACM Transactions on Graphics and SIGGRAPH Asia*, 2024. 2, 3, 6, 7, 12
- [31] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *SIGGRAPH*, 2022. 2, 3, 7, 13
- [32] Simon Niedermayr, Josef Stumpffegger, and Rüdiger Westermann. Compressed 3d gaussian splatting for accelerated novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10349–10358, 2024. 3
- [33] Michael Niemeyer, Fabian Manhardt, Marie-Julie Rakotsaona, Michael Oechsle, Daniel Duckworth, Rama Gosula, Keisuke Tateno, John Bates, Dominik Kaeser, and Federico Tombari. Radsplat: Radiance field-informed gaussian splatting for robust real-time rendering with 900+ fps. *arXiv.org*, 2024. 3
- [34] Logan Numerow, Yue Li, Stelian Coros, and Bernhard Thomaszewski. Differentiable voronoi diagrams for simulation of cell-based mechanical systems. *ACM Trans. Graph.*, 43(4), 2024. 3
- [35] Panagiotis Papantonakis, Georgios Kopanas, Bernhard Kerbl, Alexandre Lanvin, and George Drettakis. Reducing the memory footprint of 3d gaussian splatting. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 7(1):1–17, 2024. 3
- [36] Hanspeter Pfister, Matthias Zwicker, Jeroen Van Baar, and Markus Gross. Surfels: Surface elements as rendering primitives. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 335–342, 2000. 3
- [37] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically based rendering: From theory to implementation*. Third edition, 2023. 5
- [38] Lukas Radl, Michael Steiner, Mathias Parger, Alexander Weinrauch, Bernhard Kerbl, and Markus Steinberger. Stopthepop: Sorted gaussian splatting for view-consistent real-time rendering. *ACM Transactions on Graphics (TOG)*, 2024. 2
- [39] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. Derf: Decomposed radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14153–14161, 2021. 4
- [40] Daniel Reem. The geometric stability of voronoi diagrams with respect to small changes of the sites. In *Proceedings of the twenty-seventh annual symposium on Computational geometry*, pages 254–263, 2011. 6
- [41] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14335–14345, 2021. 3
- [42] Christian Reiser, Stephan Garbin, Pratul Srinivasan, Dor Verbin, Richard Szeliski, Ben Mildenhall, Jonathan Barron, Peter Hedman, and Andreas Geiger. Binary opacity grids: Capturing fine geometric detail for mesh-based view synthesis. *ACM Transactions on Graphics (TOG)*, 43(4):1–14, 2024. 3
- [43] Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Linning Xu, Zhangkai Ni, and Bo Dai. Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. *arXiv preprint arXiv:2403.17898*, 2024. 3
- [44] Liu Ren, Hanspeter Pfister, and Matthias Zwicker. Object space ewa surface splatting: A hardware accelerated approach to high quality point rendering. In *Computer Graphics Forum*, pages 461–470. Wiley Online Library, 2002. 3
- [45] Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 7, 13
- [46] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [47] Otto Seiskari, Jerry Ylilammi, Valtteri Kaatrasalo, Pekka Rantalankila, Matias Turkulainen, Juho Kannala, Esa Rahtu, and Arno Solin. Gaussian splatting on the move: Blur and rolling shutter compensation for natural camera motion. In *European Conference on Computer Vision*, pages 160–177. Springer, 2025. 3
- [48] Gopal Sharma, Daniel Rebain, Kwang Moo Yi, and Andrea Tagliasacchi. Volumetric rendering with baked quadrature fields. *ECCV*, 2024. 3
- [49] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *NeurIPS*, 2021. 2, 3, 5
- [50] Jonathan Richard Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry*, 22(1):21–74, 2002. 16th ACM Symposium on Computational Geometry. 3
- [51] Sanghyun Son, Matheus Gadelha, Yang Zhou, Zexiang Xu, Ming C. Lin, and Yi Zhou. Dmesh: A differentiable representation for general meshes, 2024. 3
- [52] Andrea Tagliasacchi and Ben Mildenhall. Volume rendering digest (for nerf), 2022. 4
- [53] Georges Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 1908. 3, 5
- [54] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit

- surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. [3](#)
- [55] 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. *arXiv preprint arXiv:2311.10091*, 2023. [3](#)
 - [56] D. F. Watson. Computing the n-dimensional delaunay tessellation with application to voronoi polytopes. *The Computer Journal*, 24(2):167–172, 1981. [3](#)
 - [57] M. Weiler, M. Kraus, M. Merz, and T. Ertl. Hardware-based view-independent cell projection. *IEEE Transactions on Visualization and Computer Graphics*, 9(2):163–175, 2003. [2](#), [4](#), [12](#)
 - [58] Turner Whitted. An improved illumination model for shaded display. In *ACM Siggraph 2005 Courses*, pages 4–es. 2005. [3](#)
 - [59] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. [3](#)
 - [60] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [7](#), [13](#)
 - [61] Matthias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross. Surface splatting. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, page 371–378, New York, NY, USA, 2001. Association for Computing Machinery. [3](#)