

Volumetric Rendering with Baked Quadrature Fields

Gopal Sharma¹, Daniel Rebain¹, Kwang Moo Yi¹, Andrea Tagliasacchi^{2, 3, 4}
¹ University of British Columbia, ² Google DeepMind,
³ Simon Fraser University, ⁴ University of Toronto

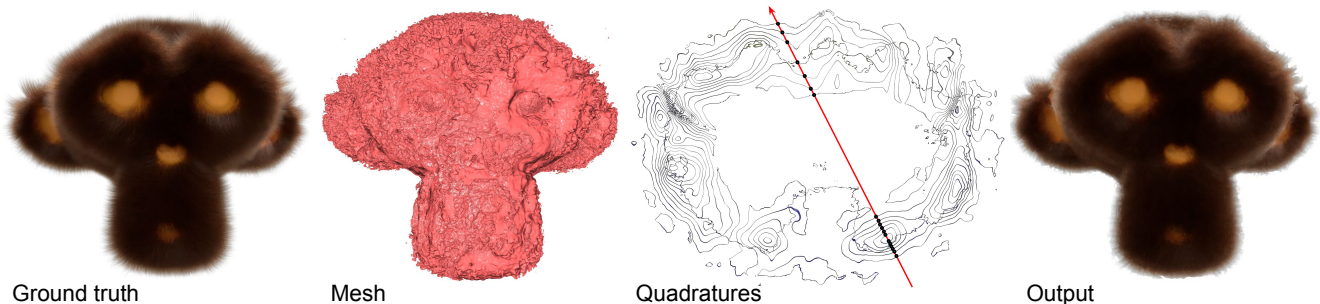


Figure 1. **Teaser** – we propose using textured polygons with NeRF to efficiently render non-opaque scenes, combining high-quality rendering with modern graphics hardware. To model a scene, we produce a mesh that gives quadrature points along a ray (visualized as points on the intersection with the cross-section of the mesh) required in volumetric rendering to render an image.

Abstract

We propose a novel Neural Radiance Field (NeRF) representation for non-opaque scenes that allows fast inference by utilizing textured polygons. Despite the high-quality novel view rendering that NeRF provides, a critical limitation is that it relies on volume rendering that can be computationally expensive and does not utilize the advancements in modern graphics hardware. Existing methods for this problem fall short when it comes to modelling volumetric effects as they rely purely on surface rendering. We thus propose to model the scene with polygons, which can then be used to obtain the quadrature points required to model volumetric effects, and also their opacity and colour from the texture. To obtain such polygonal mesh, we train a specialized field whose zero-crossings would correspond to the quadrature points when volume rendering, and perform marching cubes on this field. We then rasterize the polygons and utilize the fragment shaders to obtain the final colour image. Our method allows rendering on various devices and easy integration with existing graphics frameworks while keeping the benefits of volume rendering alive.

1. Introduction

Neural Radiance Fields (NeRFs) [29] have gained popularity by demonstrating impressive capabilities in generating photo-

realistic novel views. They use a continuous volumetric function to represent a scene with a 5D implicit function that estimates the density and radiance for any position and direction. NeRF representations are trained to achieve multi-view colour consistency for a set of posed images. One of the main challenges to the widespread adoption of NeRF is the need for specialized rendering algorithms that are not well-suited for commonly available hardware. For example, the traditional implementation of NeRF involves a volumetric rendering algorithm that calculates the density and radiance by evaluating a large MLP at hundreds of sample positions along the ray for each pixel. This rendering process is too slow for interactive visualization without powerful GPUs.

Researchers have thus been exploring real-time rendering methods using voxel grids [21] and polygonal meshes [7, 52] to address the challenge of representing scene geometry in neural volumetric rendering. While MobileNeRF [7] and BakedSDF [52] have made progress in using *binary* opacities to restrict volumetric content to polygonal meshes they cannot represent *transparent* surfaces such as glass or clouds as they rely on a *single* point to render each ray, unable to represent anything other than hard surfaces. To overcome this limitation, multiple quadrature points need to be sampled along a ray within the neural volumetric rendering setup.¹

¹Note that simply using multiple quadrature points does not enable modeling complicated physics-based rendering such as refraction, but we limit ourselves in this work to what is possible strictly with volume rendering.

However, precisely and efficiently storing such quadrature points is difficult and still an open problem. Most relevantly, SNERG [21] circumvents this problem by, instead of storing quadrature points precisely, using a volumetric representation to store opacity, diffuse colour, and feature vectors. They then march rays through a regular sparse voxel grid combined with deferred rendering to generate pixel colour. This method, however, requires large GPU memory to store volumetric data, and is incompatible with standard rendering pipelines when it comes to representing transparent objects.

In this work, we recall the traditional painter’s algorithm. The painter’s algorithm works by sorting triangles back-to-front and drawing the farthest triangles first. This naturally allows for achieving rendering of transparent surfaces, as the front triangle needs to be combined in order with the back ones using alpha blending. Thus, transparency can be implemented as a set of rasterization steps. Naturally, we aim to “bake” an existing NeRF model capable of approximating both solid and non-solid objects, while still being capable of running at interactive speed on a variety of devices. To do so, we need to design an approach that gives triangular meshes with continuous opacities, such that for solid objects we need to do a single step of rasterization and for non-solid objects, we perform multiple rasterization steps. The multi-step rasterization, when seen at the pixel level, is analogous to volumetric rendering in NeRFs. While the insight is straightforward, implementing it is not.

To implement this insight we propose learning an auxiliary neural field whose zero crossing surfaces induce a set of quadrature points for NeRF volumetric rendering – we name this field the *quadrature field*. We train the quadrature field so that it aligns with the *surface-field* [15] of the scene being represented; see Fig. 2.² With this field, we use marching cubes to extract the polygonal mesh, and for each ray, we use the intersection points with the mesh as quadrature points for volume rendering. To train this field, we use its gradients to encourage quadrature points to occur near the surfaces.

We evaluate our method on several synthetic object-centric datasets and open-scene datasets. To summarize, our contributions are as follows:

- Our approach is the first to use rasterization for rendering volumetric effects and successfully reconstructs non-solid objects.
- We introduce a quadrature field and use it to train a neural field to extract quadrature points for both solid and transparent objects.
- Our approach produces a compact representation of the scene and produces comparable results to the originally trained NeRF models.
- We create a small dataset of shapes with non-opaque surfaces to encourage research in this direction.

²Surface fields are applicable to both solid and transparent objects.

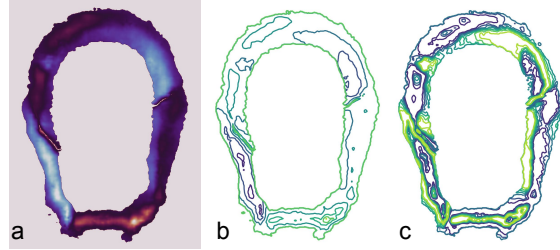


Figure 2. **Learned quadrature field from Fig. 1.** a) quadrature field along a cross-section, b) zero crossings at $\omega = 1$ and c) zero crossings at $\omega = 50$.

2. Related works

Novel view synthesis has been studied extensively in the literature [44]. In this section, we review previous works with a focus on real-time rendering.

Light fields. When viewpoints are densely sampled, novel-view synthesis can be achieved through light field rendering [25]. Lumigraph [17] performs interpolation between observed rays for rendering novel views, but this approach demands significant memory and restricts camera movements. These challenges can be mitigated by utilizing optical flow [4] for image interpolation or by employing neural networks to represent light fields [2]. Multi-plane [10, 28, 35] and multi-sphere image representations [1] have demonstrated usefulness, although they still limit camera movement. However, in practical settings where observed viewpoints are not densely captured, reconstructing a 3D representation of the scene is crucial for rendering convincing novel views.

Mesh rendering (classical). Traditional approaches to generating novel views utilize triangle meshes, typically reconstructed from point clouds via a multi-step process involving multi-view stereo [12], Poisson surface reconstruction [23], and marching cubes [27]. To create novel views, observed images are re-projected into each desired viewpoint and merged using either predetermined [5] or learned blending weights [20, 40, 41]. Although mesh-based representations are suitable for real-time rendering, they often exhibit inaccurate geometry in regions with intricate details or complex materials, resulting in visible imperfections.

Mesh rendering (differentiable). It is also possible to compute explicit triangle meshes through differentiable inverse rendering. For instance, DefTet [13] differentially renders a tetrahedral grid, considering occupancy and colour at each vertex, and then composes the interpolated values along a ray. NVDiffRec [31] combines differentiable marching tetrahedra [43] with differentiable rasterization to perform full inverse rendering, allowing extraction of triangle meshes, materials, and lighting from images. While these approaches enable scene editing and relighting, they tend to compromise

view synthesis quality.

Neural radiance fields (NeRF). Neural radiance fields [29] learn 3D consistent scene representation using continuous opacities and radiance with the help of MLP. This approach has produced excellent results for the novel-view synthesis of 3D objects with reflections [46], outdoor bounded scenes [3] and unbounded scenes [39]. However, as volumetric rendering produces pixel colour by evaluating the MLP over hundreds of points per ray, rendering speed is limited.

Efficient rendering of NeRFs. There have been several ways to speed up training and inference of NeRF. Early attempts include AutoInt [49], which models the radiance and opacity of a segment of ray instead of individual points and alpha composite the values over segments to get pixel colour, and DeRF [36], which combines multiple small NeRFs trained to represent disjoint spaces. Garbin et al. [14] introduced caching a factorized representation of neural radiance field for fast inference albeit at higher memory requirement for the cache. More recently, this problem has most commonly been addressed by trading off compute vs. storage by storing features into grids. These feature grids can be dense voxel grids [11], sparse voxel grids [21], multi-resolution hash grids [30], small MLPs distributed spatially [37] and low-rank tensor approximations of dense grid [6]. While in these methods features are converted into radiance/density by a small MLP, diffuse colours can also be stored on the grid and view-dependent radiance be represented by spherical harmonics [11, 22, 54]. Recently, Kerbl et al. [24] proposed learning a sparse set of 3D Gaussians to represent the scene, which allows skipping the empty regions of the scene easily and gives real-time rendering. Their representation is orthogonal to ours, which is purely surface-based.

Baking neural features. Rather than accelerating the NeRF directly, one can also “bake” the neural features into polygonal meshes or volumetric textures. SNERG [21] proposed to store features in sparse volumetric textures, and volumetric ray marching combined with deferred rendering to generate pixel colours. However, this approach cannot take advantage of the GPU rasterization pipeline, and requires a large GPU amount of memory to store volumetric data. A concurrent work MERF [38] allows for fast rendering of large-scale scenes while utilizing smaller memory in comparison to SNERG by utilizing a sparse feature grid and high-resolution 2D feature planes. They use a hybrid of planar and volumetric representation to model the scene, which is orthogonal to our approach which is purely surface-based. In contrast, MobileNeRF [7] proposed using a classical triangular mesh for baking the neural features into a texture map, and used *binary* opacities to optimize for a rasterized representation via volumetric rendering. BakedSDF [52] starts with VolSDF to learn a surface-based neural radiance field and use learned features baked at mesh vertices for

real-time rendering. This produces smoother and cleaner meshes in comparison to MobileNeRF owing to SDF priors via Eikonal loss and leads to better performance. However, both MobileNeRF and BakedSDF are unable to represent transparent objects faithfully as both representations assume rays to terminate at the first intersection with a surface.

Modeling transparent scenes. α Surf [50] proposed to reconstruct 3D geometry of semi-transparent objects such as glass. Their approach is based on a field that is initialized using normalized values of volumetric density, in comparison to our approach which is based on the surface field derived from volumetric weights. Their approach extracts faithful surfaces for transparent scenes but they do not explore the accurate placement of these surfaces for novel-view synthesis applications. Recently, NEMTO [47] models transparent objects by extracting the geometry of the object using DeepSDF [33] based approach and using another network to model bending of the ray through transparent media. Their reliance on SDF-based representation prevents them from modeling volumetric effects like smoke and fur, which our models can model easily as shown in Sec. 4.

3. Method

Given a set of posed images, our goal is to create a compact 3D representation of the scene that allows fast rendering. Similarly to Chen et al. [7], the representation consists of a triangular mesh and a texture map consisting of neural features and continuous opacities. Our rendering process consists of two steps:

- We use depth peeling [9] to render a transparent scene (i.e. continuous opacities) by sequential rasterization of non-transparent scenes (i.e. binarized opacities); see 3.4.
- We render view-dependent effects via spherical Gaussian lobes stored in the texture-map;

This representation is created in four-stages:

- *Training the NeRF.* We train a NeRF model based on instant-NGP [30] with continuous opacities in which quadrature points are sampled using importance sampling (Sec. 3.1).
- *Training the quadrature-field.* We train the quadrature field network with the help of NeRF. We use the trained quadrature field to extract a mesh (Sec. 3.2).
- *Fine-tuning.* We fine-tune the mesh vertices and NeRF with a network that produces the deformation field (3.3).
- *Baking.* We extract the neural features on the surface of the mesh and bake these features into a texture-map (Sec. 3.4).

3.1. Training the NeRF

Mildenhall et al. [29] introduced a 3D scene representation consisting of an MLP with trainable parameters θ that takes a position $\mathbf{x} \in \mathbb{R}^3$ and a direction vector \mathbf{d} and outputs radiance $\mathbf{c}(\mathbf{x}, \mathbf{d})$ and density $\sigma(\mathbf{x})$. Given a camera pose,

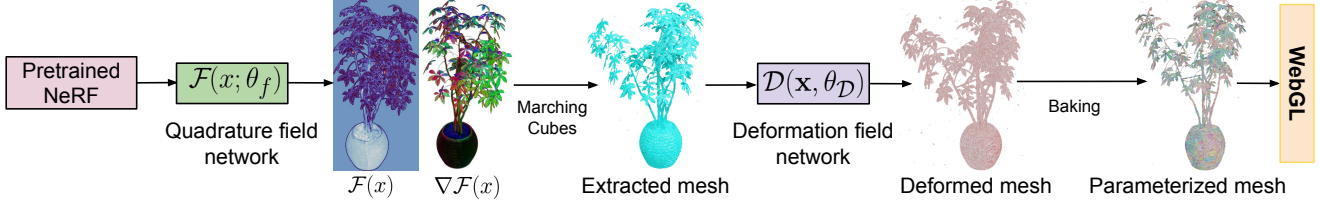


Figure 3. **Overview of our pipeline.** We start with a pre-trained network to train a quadrature field that learns the placement of quadrature points. The extracted mesh from the quadrature field is fine-tuned using a deformation field (deformation is shown using red colour on the deformed mesh). Lastly, the neural features are baked into a texture map and the mesh, which can be rendered with WebGL.

the pixel colour is computed using volumetric integration along a ray $\mathbf{r} = (\mathbf{o}, \mathbf{d})$ which is sampled at quadratures $\{t_i\}$ inducing a set of spatial samples $\mathbf{x}_i = \mathbf{o} + t_i \mathbf{d}$:

$$\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) = \sum_i w_i \cdot c_i \quad (1)$$

where $w_i = \alpha_i \prod_{j < i} \exp(1 - \alpha_j)$ is the weight given to a sampled point, $\alpha_i = 1 - \exp(-\sigma(\mathbf{x}_i) \delta_i)$ is the opacity and $\delta_i = t_{i+1} - t_i$. The MLP parameters are optimized by minimizing the difference between the predicted ray colour and the ground truth ray colour, specifically:

$$\mathcal{L}(\mathbf{r}; \boldsymbol{\theta}) = \|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_{gt}(\mathbf{r})\|^2 \quad (2)$$

In this work, we employ the instant-NGP [30] variant of NeRF for efficient training. We further use Spherical Gaussians (SG) to represent the colour which allows faster computation of view-dependent effects.

3.2. Training the quadrature field

The value w_i in (1) represents the weight given to the quadrature point i ; the higher the weight, the more likely light traveling along the direction d hits the point \mathbf{x}_i . In this sense, w_i can be seen as (view-dependent) *surfacedness*, as described by Goli et al. [15]. While for solid surfaces a single quadrature point should be sufficient to approximate the rendering equation, for non-solid objects more than one quadrature point is needed. Traditional NeRF models sample a fixed number of quadrature points from the probability density function w to approximate integration via (1). However, it is not clear how to determine these quadrature points deterministically. In this work, we propose a deterministic way to find quadrature points for all surface types. We emphasize here that we are interested in *multiple* quadrature points along the ray to enable volumetric effects, unlike those that only allow a single point [7, 52].

Defining the quadrature field. We seek to find a field that *concentrates* quadrature points in regions where surfaces are more likely to occur.³ We take inspiration from parameterization literature [19], as well as from methods that exploit the

³Note that this is not restricted to solids, but also transparent surfaces that require many quadrature points per ray to be accurately represented.

zero crossing of a signed distance field to define quadrature points [32, 48, 51], and define our *quadrature field* as:

$$\mathcal{Q}(\mathbf{x}) = \sin(\omega \mathcal{F}(\mathbf{x}; \boldsymbol{\theta}_{\mathcal{F}})) \quad (3)$$

$$\text{where } \mathcal{F}(\mathbf{x}; \boldsymbol{\theta}_{\mathcal{F}}) : \mathbb{R}^3 \rightarrow \mathbb{R} \quad (4)$$

and ω is a hyperparameter that controls the frequency of zero-crossings as shown in Fig. 2. Our quadrature points are then defined by the intersection of a ray and the zero crossings of \mathcal{Q} . Note the field \mathcal{Q} is only a function of position, as the quadrature points will be represented as a surface mesh, whose geometry does not change according to a viewpoint. To train the quadrature field, we optimize the parameters $\boldsymbol{\theta}_{\mathcal{F}}$ of the function \mathcal{F} , and create quadrature points that approximate the volume-rendering integral along a given ray.

Training the quadrature field. For quality rendering, the field q should have more zero-crossings in the region where the weight function w attains higher values. To fulfill this objective, we make two simple observations:

- Assuming local linearity, the number of zero-crossings of (3) will be proportional to the gradient of (4);
- As the weight function w in (1) is a *view-dependent* quantity, we can only supervise the *directional* derivative of (4).

Putting these two observations together, the constraint $\nabla f(\mathbf{x}; \boldsymbol{\theta}_{\mathcal{F}}) \cdot \mathbf{d} \approx w$ emerges, which we approximately satisfy via the following loss:

$$\mathcal{L}_f(\mathbf{x}; \boldsymbol{\theta}_{\mathcal{F}}) = \|\nabla \mathcal{F}(\mathbf{x}; \boldsymbol{\theta}_{\mathcal{F}}) \cdot \mathbf{d} - \max(w(\mathbf{x}, \mathbf{d}), w(\mathbf{x}, -\mathbf{d}))\| \quad (5)$$

where note the function w is non-optimized and instead kept fixed. As $w(\mathbf{x}, d)$ is a function of direction, whereas $\nabla f(\mathbf{x})$ is direction independent, we force the field f to vary most in the direction for which $w(\mathbf{x}, \pm d)$ is the largest; see Fig. 4. We use an MLP on top of a hash-grid [30] to model \mathcal{F} .

An alternative approach would be to take maximum along all directions as is suggested in nerf2nerf [16, Eq.10], but would require an impractical amount of compute in a training loop. Alternatively, MobileNeRF initializes a pruning grid to zero in [7, Eq.10] and use surface-field as a lower-bound. In our case, this entails initializing the network such that $\nabla f(\mathbf{x}) = 0 \forall \mathbf{x}$, which is harder to achieve. Our formulation, (5), provides simple and stable optimization.

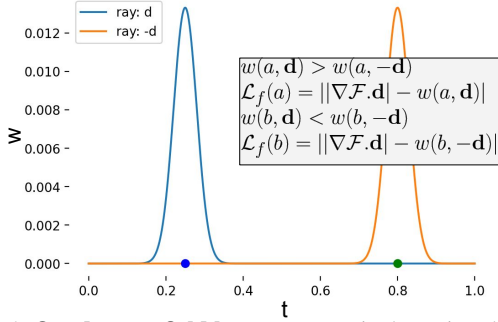


Figure 4. **Quadrature field loss.** For a particular point, the quadrature field is supervised to predict the directional gradient to be equal to the maximum weight between the bi-directions ((5)).

Finally, to encourage *sparse* creation of surfaces, thus fewer quadratures to rasterize, we have explored the use of an additional ℓ_1 regularization. We observed that it is sufficient to simply rely on the pruning data structure in instant-NGP [30] to remove surfaces from low-density regions. After learning the quadrature field, we extract the quadrature mesh \mathcal{M} via marching cubes on the function \mathcal{Q} .

3.3. Fine-tuning

While (5) can supervise the density of quadrature points, their precise location should be derived by directly optimizing for photometric reconstruction losses. One way to achieve this would be to fine-tune the vertices of the mesh and the original NeRF directly, as in Chen et al. [7]. However this saturates GPU memory and, in our experience, leads to non-smooth optimization. Instead, we employ a vector field to represent deformations continuously in space:

$$\mathcal{D}(\mathbf{x}; \boldsymbol{\theta}_{\mathcal{D}}) : \mathbb{R}^3 \rightarrow \mathbb{R}^3 \quad (6)$$

and use a perturbed set of quadrature positions $\{\tilde{\mathbf{x}}_i\}$ to evaluate photometric reconstruction via (1), where perturbation is restricted to happen *along* the ray direction \mathbf{d} :

$$\tilde{\mathbf{x}}_i = \mathbf{x}_i + \delta(\mathbf{x}_i) \quad \delta(\mathbf{x}_i) = \kappa \mathbf{d} \cdot \mathcal{D}(\mathbf{x}_i; \boldsymbol{\theta}_{\mathcal{D}}) \mathbf{d} \quad (7)$$

where a hyperbolic tangent is used as the final activation function in the \mathcal{D} , together with κ , to limit the perturbations within the spatial support of the marching cube mesh and thus stabilize training. We implement the deformation field using an MLP on top of a hash-grid [30]. Given the deformation field, a mesh vertex \mathcal{V}_i is updated as:

$$\mathcal{V}_i \leftarrow \mathcal{V}_i + \mathbb{E}_c[w_{i,c}(\kappa \mathbf{d}_c \cdot \mathcal{D}(\mathcal{V}_i; \boldsymbol{\theta}_{\mathcal{D}})) \mathbf{d}_c] / \mathbb{E}_c[w_{i,c}] \quad (8)$$

where c indices over training views, and the perturbation is weighted by the weight $w_{i,c}$ on the basis of how much a view c affects the given vertex \mathcal{V}_i .

Training loss and regularization. Training is done by defining photometric loss with perturbed quadrature points as:

$$\mathcal{L}_{\text{def}}(\mathbf{r}, \boldsymbol{\theta}, \boldsymbol{\theta}_{\mathcal{D}}) = \left\| \sum_i w(\tilde{\mathbf{x}}_i, \mathbf{r}) c(\tilde{\mathbf{x}}_i) - \mathbf{C}_{gt}(\mathbf{r}) \right\|^2 \quad (9)$$

We jointly optimize the parameters of NeRF ($\boldsymbol{\theta}$) and the deformation field ($\boldsymbol{\theta}_{\mathcal{D}}$). We also encourage $\boldsymbol{\theta}_{\mathcal{D}}$ to be smooth by minimizing the norm of deformation, and by encouraging the deformation to be smooth for each triangle \mathcal{T} :

$$\mathcal{L}_{\text{reg}}(\boldsymbol{\theta}_{\mathcal{D}}) = \mathbb{E}_{\mathcal{T} \in \mathcal{M}} \mathbb{E}_{(\mathbf{x}_a, \mathbf{x}_b) \in \mathcal{T}} \left(\|\mathcal{D}(\mathbf{x}_a; \boldsymbol{\theta}_{\mathcal{D}})\|_2^2 + \|\mathcal{D}(\mathbf{x}_a; \boldsymbol{\theta}_{\mathcal{D}}) - \mathcal{D}(\mathbf{x}_b; \boldsymbol{\theta}_{\mathcal{D}})\|_2^2 \right) \quad (10)$$

Training implementation. We employ block coordinate descent, where we first optimize the deformation field till convergence, and then we update the mesh vertices. In order to update vertices in (8), we perform a sweep over all training views to compute w_i . Note that we do not change the topology of the mesh during the above process, as we assume that the extracted mesh from the quadrature field is a good approximation.

3.4. Baking and rendering neural features

After fine-tuning, we now prepare the triangular mesh and the texture maps that we ultimately use to render. We start by post-processing the mesh to remove surfaces that are not visible from training views – these are likely artifacts as they were never “seen”. We further remove the surfaces for which the maximum volumetric weight w across all training views is below a threshold. We then construct the texture map by first parameterizing the mesh using a publicly available library [53]. We provide more detail on parameterization in the Supplementary material.

For compatibility with the WebGL rasterization pipeline, as well as general efficiency, we compress the representation of colour, alpha, and spherical Gaussian parameters to 8-bit texture maps. Specifically, we use a sigmoid transformation to bring unbounded RGB coefficients to a $[0, 1]$ range, represent spherical Gaussian lobe axes as 8-bit azimuth and elevation angles, and compress lambda values using a logarithmic mapping. These quantization results in minimal loss in performance as shown in Tab. 4. We implement the depth peeling algorithm [9] to efficiently render the textured, transparent meshes that result from our baking process while taking full advantage of hardware-accelerated rasterization.

3.5. Implementation details

We implement our NeRF using instant-NGP, where we use 10 spherical Gaussians for view-dependent effects. We use the Nerfacc [26] library based on Pytorch [34], as it provides stable training at mixed precision with comparable performance as the original instant-NGP paper. For real scenes, we use contraction mapping [3] to train the NeRF. For extracting meshes, we use a 1024 sized voxel grid for synthetic scenes and a 2048 sized voxel grid for real scenes. We further add meshes extracted from the density field of the NeRF with the mesh extracted from our quadrature field, which helps in avoiding holes like artifacts. We ablate this choice in the

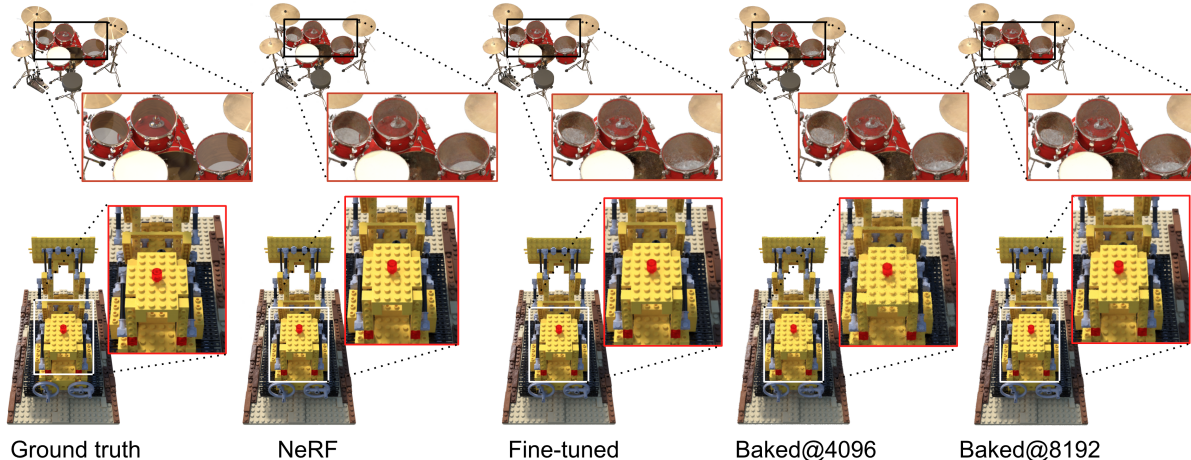


Figure 5. **Visualization of nerf-synthetic dataset.** a) Ground truth, b) instant-ngp with spherical gaussians, c) fine-tuning the meshes produced by our quadrature field (3.3), d) baked neural features using 4096 texture map, and e) using 8192 texture map. A larger texture size enables more detailed reconstruction. Notice that our approach is able to represent transparency in the drums.

Methods	Mean	Lego	Chair	Ship	Mic	Drums	Mat.	Ficus	Hotd.
instant-NGP	33.12	35.74	35.72	30.30	36.59	25.68	29.60	33.97	37.37
instant-NGP (SG)	33.00	35.50	35.52	29.99	36.67	25.76	29.50	33.90	37.18
MobileNeRF [7]	30.90	34.18	34.09	29.06	32.48	25.02	26.72	30.20	35.46
SNeRG [21]	30.38	33.82	33.24	27.97	32.60	24.57	27.21	29.32	34.33
VMesh [18]	30.70	-	-	-	-	-	-	-	-
Fine-tuned (ours)	31.78	33.90	34.38	29.11	34.50	25.48	27.79	33.08	35.98
Baked (ours)	31.02	32.25	33.62	27.96	34.16	25.20	26.70	32.85	35.41

Table 1. **Quantitative performance on NeRF synthetic dataset.** We compare our approach with various real-time approaches that are based on baking neural features.

Tab. 4. We use $\omega = 100$ for synthetic scenes and $\omega = 10$ for real scenes. We use super-sampling at twice the resolution to achieve anti-aliasing. **We will release our code upon acceptance of this manuscript.**

4. Experiments

We design our experiments to showcase superior reconstruction quality in real-time on a variety of scenes. We use the NeRF-Synthetic dataset consisting of synthetic 360-degree scenes [29]. We also experiment with a more challenging real-world dataset consisting of 7 scenes from MipNeRF 360⁴ [3]. We focus on comparing previous works like Mobile-NeRF and Baked-SDF that bake neural features into a triangular mesh. We provide ablations showing the effect of our design choices. Finally, we show results on a small dataset we created of four scenes with non-opaque objects, two of which are synthetic, and the other two are captured in indoor environments. We evaluate the performance of our method using Peak Signal to Noise Ratio (PSNR). Other metrics such as Learned Perceptual Image Patch Similarity (LPIPS) and Structural Similarity Index (SSIM) are provided in the Supp. material.

⁴Two scenes have been excluded because of license limitation.

4.1. Novel-view synthesis

NeRF synthetic dataset. Our work produces triangular meshes using the quadrature field that is based on the surface field (Sec. 3.2). An alternative approach to generate meshes would be to use Delaunay Triangulation on the surface field. This approach produces more vertices of tetrahedra in regions where the surface is likely to occur. We further compare with other methods that propose baking neural features into a volumetric representation such as voxels (SNeRG [21]), and meshes (MobileNeRF [7]) and hybrid volumetric and mesh representation (VMesh [18]).

We report the results on the NeRF synthetic dataset in the Tab. 1 and qualitative visualization in Fig. 5. Our baseline mesh reconstructed using Delaunay triangulation results in bad representation of the underlying surface as is shown in Fig. 8. Our approach is able to reconstruct transparent objects such as the drums scene and produces better reconstruction than the baselines as shown in Fig. 7.

MipNeRF 360 dataset. We further evaluate our approach to a real-world dataset and compare it with surface-based representation (MobileNeRF and BakedSDF) and hybrid volumetric representation (MeRF [38]). Tab. 2 quantitative performance of our approach and comparison with the base-



Figure 6. Visualization of mip-NeRF 360 dataset. Top: Ground truth and Bottom: Ours.

Method	Mean-indoor	Kitchen	Room	Bonsai	Counter	Mean-outdoor	Garden	Bicycle	Stump
Instant-NGP	30.07	30.30	30.96	31.46	27.57	24.39	25.47	23.08	24.61
Instant-NGP-SG	29.55	29.75	30.64	30.74	27.07	24.25	25.30	23.02	24.44
Mobile-NeRF	-	-	-	-	-	23.06	23.54	21.70	23.95
Baked-SDF	27.06	26.72	28.68	27.17	25.69	23.52	24.94	22.04	23.59
Ours	27.32	28.18	28.67	26.62	25.83	22.86	25.17	21.02	22.38
Ours (2×)	27.76	28.71	29.03	27.05	26.25	23.06	25.43	21.15	22.60

Table 2. Qualitative evaluation on mip-NeRF 360 dataset. We compare our approach with various real-time approaches that are based on baking neural features. MobileNeRF does not report results on indoor scenes. (2×) is super-sampling at twice the resolution.

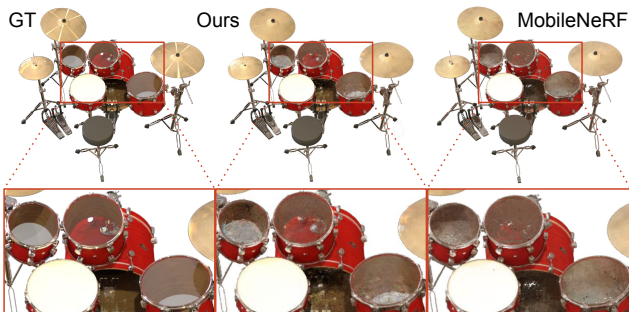


Figure 7. Comparison with MobileNeRF [7]. Our approach is able to represent transparency whereas MobileNeRF fails.

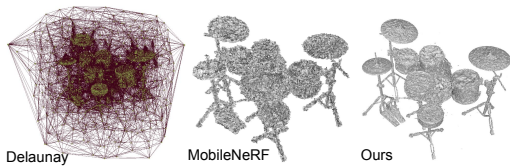


Figure 8. Mesh reconstruction. Our reconstructed mesh provides a good representation of the underlying object.

lines. Fig. 6 shows qualitative performance at different stages of training. Our approach outperforms the surface based baselines in indoor scenes. Finally, Fig. 9 shows a comparison with a surface-based method, *e.g.*, BakedSDF. Previous works that are based on binary opacities like BakedSDF struggle to represent transparent objects.

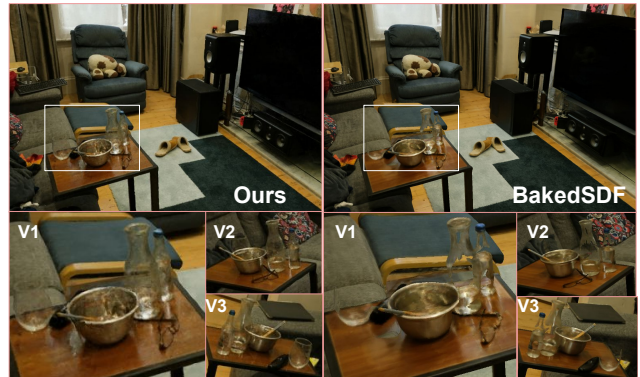


Figure 9. Comparison with BakedSDF [52]. Left: our approach, right: BakedSDF, and bottom three images focus on glassy objects from multiple views. In comparison to BakedSDF, our approach consistently reconstructs non-solid objects from different views.

4.2. Run-time performance

The rendering speed is evaluated on the synthetic dataset and is reported in the Tab. 3. The rendering performance can further be improved by pre-sorting the triangles, however, this is not the main focus of our work [42]. Run-time performance on different devices for MipNeRF 360-degree scenes is reported in the Supp. material.

	FPS @4096	FPS @8192
Macbook pro (M2)	15	8
Desktop Nvidia-3090	20	9
Desktop Nvidia A-6000	30	10

Table 3. Evaluation of rendering speed is done with different texture sizes on the mic model from the synthetic dataset with 780k faces.

Stage	Experimental Settings	PSNR
Stage 1	NGP	36.59
	NGP + SG + 4 lobes	36.24
	NGP + SG + 10 lobes (ours)	36.67
Stage 2	Finetune (ours)	34.50
	Finetune w/ density mesh	29.39
	Finetune w/ quad-field mesh	34.10
	Mesh with $\omega = 10$	31.89
	No finetune	30.58
	Finetune w/o continuous loss	33.7
Stage 3	Baked (ours)	34.15
	Baked w/o quantization	34.18
	Baked w/ 4096 texture map	33.39
	Baked w/ 8192 texture map	34.15

Table 4. **Ablations.** We ablate our proposed approach using 1) different number of spherical gaussian lobes, 2) whether we use density mesh along with mesh extracted from the quadrature field, 3) the effect of fine-tuning, 4) use of continuous rendering loss, 5) effect of ω used for mesh extraction, 6) effect of quantization and 7) finally the different sizes of texture map used for baking neural features. We use the Mic scene from the synthetic dataset to ablate.

4.3. Ablations

We ablate our algorithms on the following parameters in the Tab. 4 and observe the following:

- **Number of spherical-gaussian lobes:** Increasing spherical lobes improves reconstruction of view-dependent effects, albeit at the cost of rendering.
- **Include mesh from density field:** Often including coarse mesh extracted from the density fields complements the mesh extracted from the quadrature fields and helps fill the holes. We provide visualization in the Supp. material.
- **Effect of omega:** Larger omega leads to more quadrature points, which better captures the volumetric effects.
- **Effect of fine-tuning:** Fine-tuning aligns the mesh with the NeRF density field improving the reconstruction.
- **Size of texture map:** Increasing texels per triangle improves reconstruction but at the cost of speed and memory.

4.4. Experiments on Furry dataset

In order to test the capability of our approach in producing volumetric effects, we create a dataset with furry and smokey objects. The dataset consists of two scenes taken indoors using a Sony RX-0 II camera, capturing furry toy objects. The other two scenes are created via blender, depicting artificial smoke and fur. The resulting dataset and our final reconstructions are compared in Fig. 10. Our approach faithfully reconstructs volumetric effects in real-world datasets, however, also shows limits, when it comes to modeling pure



Figure 10. **Visualization of results on our dataset.** **Top:** our reconstruction on real scenes. **Bottom:** our reconstruction on synthetic scenes and comparison with surface rendering using NeuS [48].

volumetric objects such as the smoke, which requires too many quadrature points. While this could be alleviated by introducing more quadrature points, it would come at the cost of more computation. Still, our renderings as shown in Fig. 10, are of higher quality than surface renderings produced by NeuS [48] which is similar to BakedSDF in terms of representation power. More details and evaluations on this dataset are shown in Supp. material.

5. Conclusion

Our research addresses a critical limitation of the NeRF representation by introducing a novel approach that leverages textured polygons with continuous opacity and encodes feature vectors, enabling rapid rendering and integration into standard graphics pipelines. By training a specialized field to identify quadrature points and utilizing a novel gradient-based loss function, we achieve a quality mesh suitable for interactive rendering on desktops. Moreover, our method retains the ability to handle scenes featuring transparent objects, enhancing its practical applicability and potential impact in computer vision and graphics domains.

Limitations. Our method is bound by the limitations of NeRF; extending quadrature fields to more general rendering techniques would be interesting. Dealing with thin surfaces poses a challenge, as rarely, our method may miss thin surfaces when limited in capacity; recent developments for better quadrature for NeRFs [45] might be helpful. For large-scenes, reducing the memory footprint could be interesting to enable extremely low-end devices.

6. Acknowledgments

The authors would like to thank Sara Sabour for helping to create a dataset. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, NSERC Collaborative Research and Development Grant, Google, Digital Research Alliance of Canada, and Advanced Research Computing at the University of British Columbia.

References

- [1] Benjamin Attal, Selena Ling, Aaron Gokaslan, Christian Richardt, and James Tompkin. MatryODShka: Real-time 6DoF video view synthesis using multi-sphere images. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [2] Benjamin Attal, Jia-Bin Huang, Michael Zollhöfer, Johannes Kopf, and Changil Kim. Learning neural light fields with ray-space embedding networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [3] 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, 5, 6
- [4] Tobias Bertel, Mingze Yuan, Reuben Lindroos, and Christian Richardt. Omniphotos: Casual 360° vr photography. *ACM Trans. Graph.*, 2020. 2
- [5] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. *Unstructured Lumigraph Rendering*. 2023. 2
- [6] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision (ECCV)*, 2022. 3
- [7] Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. *arXiv preprint arXiv:2208.00277*, 2022. 1, 3, 4, 5, 6, 7, 13
- [8] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *Under Review of ICLR2016 (1997)*, 2015. 12
- [9] Cass Everitt. Interactive order-independent transparency, 2001. Accessed on November 12, 2023. 3, 5
- [10] John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2362–2371, 2019. 2
- [11] Fridovich-Keil and Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 3
- [12] Yasutaka Furukawa and Carlos Hernández. Multi-view stereo: A tutorial. *Found. Trends. Comput. Graph. Vis.*, page 1–148, 2015. 2
- [13] Jun Gao, Wenzheng Chen, Tommy Xiang, Clement Fuji Tsang, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Learning deformable tetrahedral meshes for 3d reconstruction. In *Advances In Neural Information Processing Systems*, 2020. 2
- [14] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien P. C. Valentin. Fastnerf: High-fidelity neural rendering at 200fps. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14326–14335, 2021. 3
- [15] Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, and Andrea Tagliasacchi. nerf2nerf: Pairwise registration of neural radiance fields. 2022. 2, 4
- [16] Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, and Andrea Tagliasacchi. nerf2nerf: Pairwise registration of neural radiance fields. In *International Conference on Robotics and Automation (ICRA)*, 2022. 4
- [17] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. 1996. 2
- [18] Yuan-Chen Guo, Yan-Pei Cao, Chen Wang, Yu He, Ying Shan, Xiaohu Qie, and Song-Hai Zhang. Vmesh: Hybrid volume-mesh representation for efficient view synthesis, 2023. 6
- [19] Lei He, Scott Schaefer, and Kai Hormann. Parameterizing subdivision surfaces. *ACM Trans. Graph.*, 2010. 4
- [20] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Trans. Graph.*, 2018. 2
- [21] P. Hedman, P. P. Srinivasan, B. Mildenhall, J. T. Barron, and P. Debevec. Baking neural radiance fields for real-time view synthesis. 2021. 1, 2, 3, 6
- [22] Animesh Karnewar, Tobias Ritschel, Oliver Wang, and Niloy Mitra. Relu fields: The little non-linearity that could. In *ACM SIGGRAPH 2022 Conference Proceedings*, New York, NY, USA, 2022. Association for Computing Machinery. 3
- [23] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson Surface Reconstruction. In *Symposium on Geometry Processing*. The Eurographics Association, 2006. 2
- [24] 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. 3
- [25] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, 1996. 2
- [26] Ruilong Li, Hang Gao, Matthew Tancik, and Angjoo Kanazawa. Nerfacc: Efficient sampling accelerates nerfs. *arXiv preprint arXiv:2305.04966*, 2023. 5
- [27] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *SIGGRAPH Comput. Graph.*, page 163–169, 1987. 2
- [28] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 2019. 2

- [29] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. [1](#), [3](#), [6](#)
- [30] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM TOG*, 2022. [3](#), [4](#), [5](#)
- [31] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting Triangular 3D Models, Materials, and Lighting From Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8280–8290, 2022. [2](#)
- [32] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *International Conference on Computer Vision (ICCV)*, 2021. [4](#)
- [33] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [3](#)
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*. 2019. [5](#)
- [35] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *36(6)*, 2017. [2](#)
- [36] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. Derf: Decomposed radiance fields. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14148–14156, 2021. [3](#)
- [37] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *International Conference on Computer Vision (ICCV)*, 2021. [3](#)
- [38] Christian Reiser, Richard Szeliski, Dor Verbin, Pratul P Srinivasan, Ben Mildenhall, Andreas Geiger, Jonathan T Barron, and Peter Hedman. Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *arXiv preprint arXiv:2302.12249*, 2023. [3](#), [6](#)
- [39] Konstantinos Rematas, Andrew Liu, Pratul P. Srinivasan, Jonathan T. Barron, Andrea Tagliasacchi, Tom Funkhouser, and Vittorio Ferrari. Urban radiance fields. *CVPR*, 2022. [3](#)
- [40] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *European Conference on Computer Vision*, 2020. [2](#)
- [41] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. [2](#)
- [42] Pedro V. Sander, Diego Nehab, and Joshua Barczak. Fast triangle reordering for vertex locality and reduced overdraw. *ACM Trans. Graph.*, 2007. [7](#)
- [43] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [2](#)
- [44] A. Tewari, Ohad Fried, J. Thies, V. Sitzmann, S. Lombardi, Kalyan Sunkavalli, Ricardo Martin Brualla, T. Simon, J. Saragih, M. Nießner, R. Pandey, Sean Fanello, G. Wetzstein, Jun-Yan Zhu, C. Theobalt, Maneesh Agrawala, E. Shechtman, Dan Goldman, and M. Zollhöfer. State of the art on neural rendering. *Computer Graphics Forum*, 2020. [2](#)
- [45] Mikaela Angelina Uy, George Kiyohiro Nakayama, Guandaoyang Yang, Rahul Krishna Thomas, Leonidas Guibas, and Ke Li. Nerf revisited: Fixing quadrature instability in volume rendering. In *NeurIPS*, 2023. [8](#)
- [46] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. *CVPR*, 2022. [3](#)
- [47] Dongqing Wang, Tong Zhang, and Sabine Süsstrunk. NEMTO: Neural Environment Matting for Novel View and Relighting Synthesis of Transparent Objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [3](#)
- [48] 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. [4](#), [8](#), [11](#)
- [49] Song Weiping, Shi Chence, Xiao Zhiping, Duan Zhijian, Xu Yewen, Zhang Ming, and Tang Jian. AutoInt: Automatic feature interaction learning via self-attentive neural networks. *arXiv preprint arXiv:1810.11921*, 2018. [3](#)
- [50] Tianhao Wu, Hanxue Liang, Fangcheng Zhong, Gernot Riegler, Shimon Vainer, and Cengiz Oztireli. α surf: Implicit surface reconstruction for semi-transparent and thin objects with decoupled geometry and opacity, 2023. [3](#)
- [51] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. [4](#), [11](#)
- [52] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P. Srinivasan, Richard Szeliski, Jonathan T. Barron, and Ben Mildenhall. BakedSDF: Meshing neural SDFs for real-time view synthesis. *arXiv*, 2023. [1](#), [3](#), [4](#), [7](#), [11](#)
- [53] Jonathan Young. Xatlas. Available at: <https://github.com/jpcy/xatlas>, 2023. [5](#)
- [54] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. [3](#)
- [55] Zehao Yu, Anpei Chen, Bozidar Antic, Songyou Peng, Apratim Bhattacharyya, Michael Niemeyer, Siyu Tang, Torsten Sattler, and Andreas Geiger. Sdfstudio: A unified framework for surface reconstruction, 2022. [11](#)



Figure 11. **Renderings from our web application.** Left: Rendering of mic model from the nerf-synthetic dataset and Right: rendering of cropped mesh from the garden scene from mipNeRF dataset.

7. Supplementary Material

In this section, we provide information about the following:

- Surface parameterization
- Run-time performance on mipNeRF 360 scenes
- Visualization of ablations (quad mesh)
- Evaluation using more metrics
- More evaluations on fur dataset

7.1. Surface parameterization

To parameterize the mesh generated using $Q(x)$, we first segment the mesh into smaller patches using graph cuts. Each segmented patch is parameterized separately and then all patches are packed into an atlas. For mipNeRF 360 scenes, we parameterize the mesh in contracted space such that the triangles that are farther away use fewer texels per unit.

7.2. Run-time performance on mipNeRF 360 scenes

In order to fit the texture map in memory, we only evaluate the rendering speed on a mesh cropped around the object of interest as is shown in Fig. 11. We report the run-time in the Tab. 6.

7.3. Effect of different mesh extraction techniques

Mesh extracted from our quadrature field can result in holes. This can be alleviated by adding a coarse mesh extracted from the density field as shown in Fig. 12. This comes at a little additional cost of more triangles per scene.

7.4. Experiments on fur dataset

Methods	Hairy monkey			Smokey monkey		
	PSNR	LPIPS	SSIM	PSNR	LPIPS	SSIM
NGP + SG	36.78	0.1585	0.9669	44.05	0.1432	0.9807
NeuS [48]	22.81	0.1504	0.9069	23.34	0.1132	0.9148
Ours	34.76	0.1626	0.9565	35.86	0.1687	0.9546

Table 5. **Evaluation on synthetic fur dataset.** We compare our approach with NeuS by surface rendering.



Figure 12. Mesh extracted from our quadrature field can result in holes. This can be alleviated by adding a coarse mesh extracted from the density field.



Figure 13. **Comparison with VolSDF.** Our approach produces better volumetric effects in comparison to surface rendering done on mesh extracted from VolSDF. Notice the absence of fur in the surface rendering using VolSDF.

In order to compare with our mesh-based rendering approach, We create a baseline using NeuS [48]. In this, we extract a mesh and perform surface rendering instead of the volumetric rendering proposed in the original paper. This results in renderings that lack volumetric effects as is shown in Figure 10 in the main paper. We compare with this baseline in the Tab. 5.

For our real dataset, we compare our approach with VolSDF [51] (on which BakedSDF [52] is based) since the official implementation of BakedSDF is not available. We use the publicly available source code from sdfstudio [55] for comparison. Note that this implementation does not fine-tune the neural features on the vertices of the mesh, as proposed by the original BakedSDF. Our approach, which produces multiple surfaces, thanks to the learned quadrature field, produces volumetric effects like fur as shown in Fig. 13.

7.5. Evaluation metrics

We provide more evaluation using SSIM and LPIPS metrics in Tabs. 7 and 9 and Tabs. 8 and 10 respectively on nerf-synthetic and mipNeRF 360 dataset.

	FPS @4096	FPS @8192
Macbook pro (M2)	10	5
Desktop Nvidia-3090	10	3
Desktop Nvidia A-6000	18	9

Table 6. Evaluation of rendering speed is done with different texture sizes on the garden model from the Garden scene with 690k faces and rendering at 800×800 resolution.

7.6. Architecture design

We implement a quadrature field using the hash grid with an MLP with two hidden layers, each of width 16. The input coordinates are not only input to the hash grid but also concatenated with the output of the hash grid. We use Exponential Linear Units [8] in the MLP which allows double derivatives used to supervise our quadrature field. For synthetic experiments, we use the hash grid with a maximum resolution of 512, minimum resolution of 16 and codebook size of 2^{30} . For the real dataset, we use the hash grid with a maximum resolution of 4096, a minimum resolution of 16 and a codebook size of 2^{25} . We use a similar architecture design to implement our deformation network, except we use relu non-linearity in the MLP. The choice of these parameters is dependent on available GPU memory and validation performance.

Method	Mean	Lego	Chair	Ship	Mic	Drums	Materials	Ficus	Hotdog
NGP	0.961	0.980	0.985	0.891	0.991	0.934	0.948	0.982	0.982
NGP SG	0.961	0.979	0.985	0.886	0.991	0.937	0.947	0.982	0.981
MobileNeRF [7]	0.947	0.975	0.978	0.867	0.979	0.927	0.913	0.965	0.973
Fine-tune (ours)	0.954	0.974	0.980	0.875	0.987	0.931	0.927	0.979	0.976
Baked (ours)	0.950	0.965	0.976	0.865	0.986	0.929	0.926	0.978	0.975

Table 7. Evaluation using SSIM metric on NeRF-synthetic dataset.

Method	Mean	Lego	Chair	Ship	Mic	Drums	Materials	Ficus	Hotdog
NGP	0.052	0.024	0.021	0.142	0.015	0.081	0.069	0.026	0.035
NGP SG	0.051	0.024	0.021	0.146	0.014	0.075	0.067	0.026	0.035
MobileNeRF	0.062	0.025	0.025	0.145	0.032	0.077	0.092	0.048	0.050
Fine-tune (ours)	0.060	0.031	0.029	0.147	0.022	0.083	0.086	0.032	0.052
Baked (ours)	0.072	0.065	0.038	0.172	0.025	0.091	0.097	0.034	0.058

Table 8. Evaluation using LPIPS metric on NeRF-synthetic dataset.

Method	Mean (indoor)	Kitchen	Room	Bonsai	Counter	Mean (outdoor)	Garden	Bicycle	Stump
Instant-NGP	0.884	0.893	0.902	0.914	0.826	0.617	0.710	0.529	0.612
Instant-NGP SG	0.877	0.881	0.902	0.908	0.818	0.612	0.706	0.529	0.601
Mobile-NeRF	-	-	-	-	-	0.527	0.599	0.426	0.556
Baked-SDF	0.837	0.817	0.870	0.851	0.808	0.639	0.751	0.570	0.595
Ours	0.844	0.866	0.876	0.851	0.784	0.560	0.725	0.449	0.505
Ours (2x)	0.856	0.878	0.883	0.862	0.801	0.579	0.732	0.485	0.519

Table 9. Evaluation using SSIM metric on mipNeRF 360 dataset.

Method	Mean (indoor)	Kitchen	Room	Bonsai	Counter	Mean (outdoor)	Garden	Bicycle	Stump
Instant-NGP	0.2051	0.1445	0.2278	0.1826	0.2655	0.4004	0.2892	0.4884	0.4235
Instant-NGP SG	0.2133	0.1591	0.2300	0.1939	0.2704	0.4078	0.2922	0.4887	0.4424
Mobile-NeRF	-	-	-	-	-	0.4337	0.3580	0.5130	0.4300
Baked-SDF	0.2583	0.2370	0.2510	0.2590	0.2860	0.3173	0.2130	0.3680	0.3710
Ours	0.2519	0.1818	0.2637	0.2642	0.2978	0.4308	0.2667	0.5301	0.4956
Ours (2x)	0.2379	0.1693	0.2556	0.2511	0.2758	0.4127	0.2599	0.5032	0.4750

Table 10. Evaluation using LPIPS metric on mipNeRF 360 dataset.