

# Geometry Field Splatting with Gaussian Surfels

Kaiwen Jiang<sup>1</sup> Venkataram Sivaram<sup>1</sup> Cheng Peng<sup>2</sup> Ravi Ramamoorthi<sup>1</sup>

<sup>1</sup>University of California, San Diego <sup>2</sup>Johns Hopkins University

## Abstract

*Geometric reconstruction of opaque surfaces from images is a longstanding challenge in computer vision, with renewed interest from volumetric view synthesis algorithms using radiance fields. We leverage the geometry field proposed in recent work for stochastic opaque surfaces, which can then be converted to volume densities. We adapt Gaussian kernels or surfels to splat the geometry field rather than the volume, enabling precise reconstruction of opaque solids. Our first contribution is to derive an efficient and almost exact differentiable rendering algorithm for geometry fields parameterized by Gaussian surfels, while removing current approximations involving Taylor series and no self-attenuation. Next, we address the discontinuous loss landscape when surfels cluster near geometry, showing how to guarantee that the rendered color is a continuous function of the colors of the kernels, irrespective of ordering. Finally, we use latent representations with spherical harmonics encoded reflection vectors rather than spherical harmonics encoded colors to better address specular surfaces. We demonstrate significant improvement in the quality of reconstructed 3D surfaces on widely-used datasets.*

## 1. Introduction

The reconstruction of opaque surfaces from a collection of calibrated RGB images is a classic challenge in computer vision. This problem has received broader attention in recent years due to the success of Neural Radiance Field (NeRF) [30] and other volumetric novel view synthesis algorithms. Although methods like NeRF [30] are effective for view synthesis, they fall short in geometric reconstruction applications, since they do not explicitly reconstruct an accurate surface. Follow up methods, which are specialized towards surface reconstruction, have proposed instead that an explicit geometry representation should parameterize the volume density field.

The typical approach taken by previous methods [25, 27, 29, 31, 45, 48, 54, 55] is to parameterize signed or unsigned distance fields with a neural network, which can then be converted to a density field for volume rendering [31, 45]. Eikonal regularization [16] is often used to promote smooth-




Figure 1. We introduce a geometry reconstruction method from a set of calibrated RGB images by splatting a geometry field with Gaussian surfels. We enable almost exact and efficient differentiable rendering. Compared to 2DGS, we reach better geometry quality and achieve overall smooth and detailed geometry without having cracks or holes. The reader may wish to zoom into the electronic version in the figures. (Note: object colors used are for visualization, not original object color).

ness of the reconstructed surfaces. In particular, Miller et al. [31] provides a theoretical bridge between surface representations and volume rendering with stochastic geometry, and generalizes the use of distance fields to introduce a general stochastic geometry field, which we refer to as simply a *geometry field*. Chen et al. [5] shows that such geometry fields can be parameterized with a Poisson field.

Recent advances in Gaussian splatting based methods [23] have motivated a new class of surface modeling methods (e.g., [7, 18, 20, 49, 59, 60]). These methods utilize kernels, i.e., 3D Gaussians or 2D Gaussians, to parameterize a density or opacity field and then perform volume splatting [23, 64] for rendering. Geometry is then recovered from the converged density or opacity field through varying definitions of the surface. The benefit compared to the neural methods comes from the speed and ability to generalize to complex scenes (e.g., [4, 11, 61]). Despite the success achieved by these approaches, the surface is not theoretically clearly defined for a general density field, which incurs approximation that limits the reconstruction quality.

In this paper, we adapt a modern volumetric representation based on Gaussian kernels to splat the geometry field that provides precise geometry for reconstructing opaque solids. We show that we can utilize the kernels, i.e., 2D Gaussians or Gaussian surfels, to parameterize the geometry

field and still use the efficient volume splatting algorithm for differentiable rendering. We carefully remove the approximations in the original volume splatting algorithm to ensure almost exact rendering, except for the use of global sorting to approximate per-ray sorting. We also observe that, to make the represented stochastic geometry gradually converge to the deterministic case, kernels need to be clustered around the surface. Such clustering leads to an unstable ordering of kernels in volume splatting and the rendered color a discontinuous function of the properties of kernels. The view synthesis loss then creates a discontinuous loss landscape that is detrimental for optimization. We propose a solution to ensure that the rendered color is a continuous function of the properties of the kernels. Finally, inspired by [14, 22, 40, 41, 44, 50], we also explore different choices of color representation other than spherical harmonics [35] to tackle specular surfaces. Through extensive experiments, we show that such a design with geometry clearly defined can significantly improve the quality of the reconstructed surface (see Figs. 1 and 8).

Our main contributions are summarized as follows:

- We derive an efficient and differentiable rendering algorithm for geometry fields parameterized by Gaussian surfels (Sec. 4.1). Besides, we carefully analyze and reduce the approximations in the rendering algorithm.
- We expose and remedy the discontinuity of the loss landscape, which ultimately helps our method to converge from stochastic geometry to a deterministic shape (Sec. 4.2).

## 2. Related Work

**Multi-view Stereo Surface Reconstruction.** The computer vision community has studied 3D reconstruction from multi-view images for decades using multi-view stereo [1, 13, 15, 19, 38, 39, 42]. Despite their high accuracy, these methods do not provide complete geometry. Therefore, CNN-based approaches (e.g., [2, 9, 10, 17, 24, 43, 46, 47, 51, 52, 58, 63]) were introduced to overcome this limitation, but they cannot reconstruct reflective surfaces. Neural surface reconstruction approaches are later introduced to address more general surface types.

**Neural Surface Reconstruction.** Thanks to the advance of NeRF [30], some works (e.g., [14, 25, 28, 31, 44, 45, 48, 54, 55]) propose to utilize various representations, such as an MLP [30] or Hash-grid [33], to parameterize a SDF field and then convert it into a density field for rendering. The mesh is then extracted using marching cubes from the SDF field after convergence. Identifying the limitation of the SDF field for open-surface objects, some other works (e.g., [27, 29]) propose to leverage an UDF field instead of a SDF field. Inspired by multi-view stereo (MVS), another group of methods (e.g., [8, 12, 37]) add a multi-view consistency loss to further enhance the quality of reconstructed geometry.

**Gaussian Splatting based Surface Reconstruction.** Kerbl et al. [23] propose to utilize 3D Gaussians to param-




Figure 2. Illustration of a ray intersecting with non-overlapping kernels  $\mathcal{K}_1, \mathcal{K}_2, \dots, \mathcal{K}_N$ , which are sorted based on their intersection intervals  $[a_1, b_1], [a_2, b_2], \dots, [a_N, b_N]$ .

eterize the density field for differentiable view synthesis. Owing to its efficiency and high quality, it has become popular to extend Gaussian splatting (GS) for geometry reconstruction. SuGaR [18] constrains the 3D Gaussians to be flattened for extracting the geometry but still exhibits holes. 2DGS [20] and Dai et al. [7] then propose to directly use the 2D Gaussians to parameterize the density field and utilize depth-normal consistency loss or a monocular normal prior to ensure the smoothness of extracted geometry. In contrast, RaDe-GS [60] and GOF [59] still rely on 3D Gaussians, but propose custom surface definitions for the 3D Gaussian. Surfaces are extracted through TSDF fusion [6, 62] or tetrahedral grids [59]. Different from these methods, we parameterize the geometry field with 2D Gaussians and then convert it into the density field for rendering with refined volume splatting. We still use TSDF fusion to extract the geometry and achieve significant improvements in terms of geometry quality. Concurrently, PGSR [3] introduces a multi-view constraint into the optimization.

## 3. Preliminaries

We first introduce and analyze the volume splatting algorithm and volume representation of opaque solids.

**Volume Splatting.** Consider a ray  $\mathbf{x}(t) = \mathbf{o} + t\boldsymbol{\omega}$ , where  $\mathbf{o}$  denotes the camera origin,  $\boldsymbol{\omega}$  denotes the ray direction and  $t > 0$  denotes the depth. Volume splatting [64] aims to evaluate the following volume rendering equation:

$$\mathbf{C} = \int_0^\infty \mathbf{c}(t)\sigma(t) \exp\left(-\int_0^t \sigma(t') dt'\right) dt, \quad (1)$$

where  $\mathbf{c}(t), \sigma(t)$  denote the color and density values, or the extinction coefficient of the volume, at  $\mathbf{x}(t)$ .

As shown in Fig. 2, Zwicker et al. [64] propose to decompose the density field as a set of independent kernels, each of which is chosen to have finite support range by cutoff, i.e.,  $\sigma(t) = \sum_{i=1}^N w_i \mathcal{K}_i(t)$ , where  $N$  denotes the number of kernels, and  $w_i$  and  $\mathcal{K}_i(t)$  respectively denote the weight and evaluated value at  $t$  for the  $i^{\text{th}}$  kernel. The intersections between these kernels and the ray are assumed to be non-overlapping and, therefore, these kernels are sorted along the ray based on the intersection interval. The intersection interval of the  $i^{\text{th}}$  kernel is denoted as  $[a_i, b_i]$ , i.e.,  $\mathcal{K}_i(t) \neq 0$  if and only if  $a_i \leq t \leq b_i$ . The color values inside the  $i^{\text{th}}$  kernel are defined to be constant as  $\mathbf{c}_i$ .

By ignoring the self-occlusion and expanding the transmittance term with the Taylor Series, Eqn. (1) can be rewrit-

ten as:

$$\mathbf{C} = \sum_{i=1}^N \mathbf{c}_i \rho_i \prod_{j=1}^{i-1} (1 - \rho_j), \quad (2)$$

where  $\rho_i$  is the integration of density value along the ray, when considering only the  $i^{\text{th}}$  kernel. Namely,  $\rho_i = \int_{a_i}^{b_i} \sigma(t) dt = \int_{a_i}^{b_i} w_i \mathcal{K}_i(t) dt$ , and it is called the footprint of the  $i^{\text{th}}$  kernel with respect to the current ray. Notice that due to the expansion,  $\rho_i \in [0, 1]$ .

Therefore, it is important that the footprint function be easy to evaluate efficiently, such that the evaluation of Eqn. (2) can be efficient. Volume splatting chooses to use the 3D Gaussian as the kernel, and analytically solve the footprint function by approximating the perspective transform.

In summary, volume splatting makes the following assumption and approximation:

- Assume there is no overlapping of any two intersections between kernels and the ray.
- Self-occlusion of each kernel is ignored, and transmittance terms and footprint functions are approximated.

Even though remarkable view synthesis from a set of calibrated images is achieved by making the volume splatting differentiable [23], the non-overlapping assumption leads to a discontinuous view synthesis loss function and the constraint that the footprint function is within the range of  $[0, 1]$ . We will show that these problems are detrimental for the surface reconstruction and propose solutions. Besides, Kerbl et al. [23] uses a global sorting to approximate per-ray sorting.

**Stochastic Geometry Representation for Opaque Solid.** Miller et al. [31] propose that for an opaque solid existing in the  $\mathbb{R}^3$  space, each point  $\mathbf{x} \in \mathbb{R}^3$  is associated with a random variable  $G(\mathbf{x}) \in \mathbb{R}$ . Without loss of generality, the random variable is chosen to obey the normal distribution, i.e.,  $G(\mathbf{x}) \sim \mathcal{N}(\mu(\mathbf{x}), 1/s(\mathbf{x}))$ , where  $\mu(\mathbf{x})$  and  $s(\mathbf{x})$  denote the mean and inverse of standard deviation. Each point  $\mathbf{x}$  is then associated with an occupancy value  $o(\mathbf{x})$  and vacancy value  $v(\mathbf{x})$ , defined as:

$$\begin{aligned} o(\mathbf{x}) &= \Pr\{G(\mathbf{x}) \geq 0\} = \Psi(\mu(\mathbf{x})s(\mathbf{x})) \\ v(\mathbf{x}) &= \Pr\{G(\mathbf{x}) < 0\} = \Psi(-\mu(\mathbf{x})s(\mathbf{x})), \end{aligned} \quad (3)$$

where  $\Psi(\cdot)$  denotes the cdf of the standard normal distribution. The solid is defined at  $\mathbb{I}_{\{G(\mathbf{x}) \geq 0\}}$ .

From [31], the density at  $\mathbf{x}$  is defined as:

$$\begin{aligned} \sigma(\mathbf{x}) &= \frac{\|\nabla v(\mathbf{x})\|}{v(\mathbf{x})} \cdot \left( \alpha(\mathbf{x})|\boldsymbol{\omega} \cdot \mathbf{n}(\mathbf{x})| + \frac{1 - \alpha(\mathbf{x})}{2} \right) \\ &= \frac{\psi(-\mu(\mathbf{x})s(\mathbf{x}))}{\Psi(-\mu(\mathbf{x})s(\mathbf{x}))} \|\nabla(\mu(\mathbf{x})s(\mathbf{x}))\| \cdot \\ &\quad \left( \alpha(\mathbf{x})|\boldsymbol{\omega} \cdot \mathbf{n}(\mathbf{x})| + \frac{1 - \alpha(\mathbf{x})}{2} \right), \end{aligned} \quad (4)$$

where  $\psi(\cdot)$  denotes the pdf of the standard normal distribution,  $\mathbf{n}(\mathbf{x}) = -\nabla(\mu(\mathbf{x})s(\mathbf{x}))/\|\nabla(\mu(\mathbf{x})s(\mathbf{x}))\|$  and

$\alpha(\mathbf{x}) \in [0, 1]$  denotes an anisotropic parameter which is close to 1 near the surface, and close to 0 in the interior of the opaque solid. As  $s(\mathbf{x}) \rightarrow \infty$ , the stochastic geometry converges to the deterministic case.

Notice that, in other works (e.g., [31, 45]), the  $\mu(\mathbf{x})$  is defined to be a SDF field. Defining it as a SDF field allows the usage of Eikonal loss to promote the smoothness, but there is no constraint over the meaning of  $\mu(\mathbf{x})$ .

Since  $\mu(\mathbf{x})$  and  $s(\mathbf{x})$  are always together in Eqn. (4), we define that

$$F(\mathbf{x}) := \mu(\mathbf{x})s(\mathbf{x}), \quad (5)$$

which is called the *geometry field* as it indicates the underlying geometry. When  $\forall \mathbf{x} \in \mathbb{R}^3, |F(\mathbf{x})| \rightarrow \infty$ , the represented geometry is assumed to be converged to the deterministic case. Eqn. (4) can be then written as:

$$\sigma(\mathbf{x}) = \frac{\psi(-F(\mathbf{x}))}{\Psi(-F(\mathbf{x}))} \|\nabla F(\mathbf{x})\| \cdot \left( \alpha(\mathbf{x})|\boldsymbol{\omega} \cdot \mathbf{n}(\mathbf{x})| + \frac{1 - \alpha(\mathbf{x})}{2} \right), \quad (6)$$

where  $\mathbf{n}(\mathbf{x}) = -\nabla F(\mathbf{x})/\|\nabla F(\mathbf{x})\|$ .

## 4. Method

In Sec. 4.1, we describe how we can reduce the approximations taken by the splatting algorithm. We then derive a refined splatting method which enables differentiable rendering of the opaque solid represented by the geometry field. Next, in Sec. 4.2, we propose a novel solution to address the issue of discontinuous loss landscapes which arises as the stochastic geometry gradually becomes deterministic. Finally, Sec. 4.3 introduces an improvement in color representation to account for specular surfaces.

### 4.1. Geometry Field Splatting

**Revisit Splatting Algorithm.** As the volume splatting is an *approximate* rendering algorithm, we aim to reduce the approximations, and investigate the situation under which the refined volume splatting algorithm is *exact*.

Specifically, we do not ignore the self-occlusion and do not expand the transmittance term with the Taylor series. Eqn. (2) is then replaced with:

$$\mathbf{C} = \sum_{i=1}^N \mathbf{c}_i (1 - \exp(-\rho_i)) \prod_{j=1}^{i-1} \exp(-\rho_j). \quad (7)$$

Please find the derivation in A.1 of the supplementary document. Notice that then  $\rho_i \in \mathbb{R}_+$  with no further restrictions.

We are now left with the non-overlapping assumption and footprint function approximation. We notice the following property:

**Property 1.** *When the intersection intervals of two and only two kernels  $\mathcal{K}_i, \mathcal{K}_j$  fully overlap, i.e.,  $a_i = a_j \wedge b_i = b_j$ , and they have the same color, it is equivalent to having one kernel  $\mathcal{K}_{ij}$  there with  $\rho_{ij} = \rho_i + \rho_j$  and the same color.*

*Proof.* As these two kernels are the only two kernels whose intersection intervals fully overlap at  $[a_i, b_i]$ , we can conclude that  $j = i + 1$  without losing generality. Therefore,  $(1 - \exp(-\rho_i)) + (1 - \exp(-\rho_j)) \exp(-\rho_i) = 1 - \exp(-(\rho_i + \rho_j))$ . Notice that the alpha-blended value of these two kernels is equivalent to having a single kernel there.  $\square$

This property can be extended into the following form through mathematical induction:

**Lemma 1.** *When the intersection intervals of  $m$  kernels  $\mathcal{K}_{i_1}, \mathcal{K}_{i_2}, \dots, \mathcal{K}_{i_m}$  fully overlap, i.e.,  $a_{i_1} = a_{i_2} = \dots = a_{i_m} \wedge b_{i_1} = b_{i_2} = \dots = b_{i_m}$ , and they have the same color, it is equivalent to having one kernel  $\mathcal{K}$  there with  $\rho = \rho_{i_1} + \rho_{i_2} + \dots + \rho_{i_m}$  and the same color.*

Therefore, we can define an operator  $\oplus$ , which satisfies the commutative, associative and identity properties, such that  $\mathcal{K}_{i_j} = w_i \mathcal{K}_i \oplus w_j \mathcal{K}_j$  where  $\rho_{i_j} = \rho_i + \rho_j$ . When the kernels on the ray have either non-overlapping or fully overlapped intersection intervals, the refined splatting algorithm defined by Eqn. (7) is an *exact* rendering algorithm of density field defined as:  $\sigma(\mathbf{x}) = \bigoplus_{i=1}^N w_i \mathcal{K}_i(\mathbf{x})$ , because we remove the approximations when evaluating Eqn. (1).

By considering the remaining approximation, we can then reach the following theorem:

**Theorem 1.** *When 1) the intersection intervals of any two kernels are non-overlapping, or fully overlapped in which case these two kernels have the exact same color value, and, 2) the footprint function is accurately calculated, then the sorting of kernels is well-defined and Eqn. (7) gives an exact rendering of a density field which is defined as:  $\sigma(\mathbf{x}) = \bigoplus_{i=1}^N w_i \mathcal{K}_i(\mathbf{x})$ .*

**Parameterize and Render Geometry Field with Gaussian Surfels.** Our overall idea is illustrated in Fig. 3. We propose to parameterize the geometry field in Eqn. (5) with kernels, then demonstrate how to convert it into the density field for volume splatting, and reveal under which conditions it is an exact rendering algorithm. We will also show that choosing Gaussian surfels as our kernels helps build a both efficient and almost exact rendering algorithm.

Specifically, we decompose the  $F(\mathbf{x})$  as:

$$F(\mathbf{x}) = \bigoplus_{i=1}^N w_i \mathcal{K}_i(\mathbf{x}) - c, \quad (8)$$

where  $\bigoplus$  is the custom plus operator which we have just discussed and will elaborate on later,  $w_i > 0$ , and  $c$  is a manually chosen constant. Since we expect the kernels to cluster around the surface, the anisotropic parameter  $\alpha(\mathbf{x})$  is then set to 1 as discussed in [31].

We first discuss the property of such a formulation of  $F(\mathbf{x})$ , when the represented geometry converges to the deterministic case. At the place where it is vacant, i.e., there are

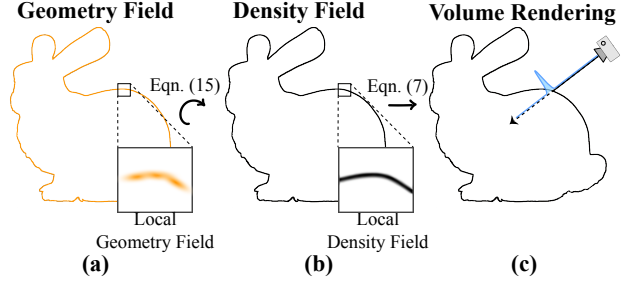


Figure 3. Overview of our algorithm. (a) We first use 2D Gaussians to parameterize the geometry field,  $F$ . These kernels are expected to cluster around the surface. (b) We then convert the geometry field into the density field  $\sigma$  and lastly, (c) we leverage our refined volume splatting algorithm for differentiable rendering.

no kernels,  $F(\mathbf{x}) = -c$ , such that  $v(\mathbf{x}) \approx 1$ . Therefore,  $c$  is a manually chosen large positive value. At the place where kernels cluster,  $F(\mathbf{x})$  is expected to be a large positive value, such that  $o(\mathbf{x}) \approx 1$ . Therefore, when a ray travels in the space, the transmittance is expected to almost immediately fall from 1 to 0 at the place where it first intersects with kernels as shown in Fig. 3 (c).

We still start from the assumption that there is no overlapping of any two intersections between kernels and the ray and then remove this assumption. From Eqn. (6), the corresponding density field is defined as:

$$\sigma(\mathbf{x}) = \frac{\psi(-F(\mathbf{x}))}{\Psi(-F(\mathbf{x}))} \|\nabla F(\mathbf{x})\| \cdot |\boldsymbol{\omega} \cdot \mathbf{n}(\mathbf{x})|. \quad (9)$$

As Eqn. (7) indicates, given a ray, our goal is to evaluate the footprint function for the  $i^{\text{th}}$  kernel:

$$\rho_i = \int_{a_i}^{b_i} \sigma(\mathbf{x}(t)) dt. \quad (10)$$

We now make our choice of kernels. Aiming at exact rendering, we choose to use 2D Gaussians or Gaussian surfels [7, 20, 34] as our kernels, instead of 3D Gaussians [23]. In general, the intersections between Gaussian surfels and a ray are points. Thus, any two of the intersections are either not overlapping at all or fully coincide, satisfying the conditions of Theorem. 1.

The  $i^{\text{th}}$  2D Gaussian kernel  $\mathcal{G}(\mathbf{x})$  is defined as:

$$\mathcal{G}_i(\mathbf{x}) = \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{m}_i)^T \Sigma_i^{-1}(\mathbf{x} - \mathbf{m}_i)\right), \quad (11)$$

where  $\mathbf{m}_i$  denotes its center and  $\Sigma_i$  denotes its covariance matrix. As shown in Fig. 4 (a), given a ray  $\mathbf{x}(t) = \mathbf{o} + t\boldsymbol{\omega}$ , we are able to calculate the intersected depth  $t_i$  and evaluate its kernel value at  $\mathbf{x}(t_i)$  as  $f_i = w_i \mathcal{G}_i(\mathbf{x}(t_i))$  [20]. We then have  $F(\mathbf{x}(t))$  near  $t_i$ :

$$F(\mathbf{x}(t)) = \begin{cases} f_i - c & t = t_i \\ -c & |t - t_i| < \epsilon, \text{ where } \epsilon \rightarrow 0. \end{cases} \quad (12)$$



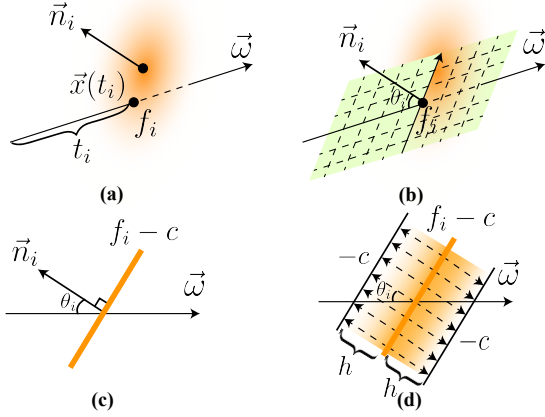


Figure 4. **(a)** An illustration of intersecting the  $i^{\text{th}}$  Gaussian surfel, whose normal vector is denoted as  $\mathbf{n}_i$  that is perpendicular to the plane containing the surfel, with a ray, whose direction is denoted as  $\omega$ . The depth of the intersected point is denoted as  $t_i$ , and the intersected point is denoted as  $\mathbf{x}(t_i)$ . The value of the Gaussian surfel at  $\mathbf{x}(t_i)$  is denoted as  $f_i$ . **(b)** We analyze the intersection by creating a 2D coordinate plane which is parallel to the  $\mathbf{n}_i$  and  $\omega$ , and passes through  $\mathbf{x}(t_i)$ . The angle between  $\mathbf{n}_i$  and  $\omega$  is denoted as  $\theta_i$ . **(c)** On the 2D coordinate plane, within an infinitesimally small range, the intersected line, drawn as an orange line, between the Gaussian surfel and the plane can be seen as having constant geometry field value  $F = f_i - c$ . **(d)** We expand the intersected line by linearly decaying it into  $-c$  with length  $h$  and direction  $\mathbf{n}_i$ . It gives the surfel a 3D width, which does not follow the Gaussian distribution. As  $h \rightarrow 0$ , it can be seen as equivalent to (c).

Such a function is not even a continuous function. It seems that we cannot calculate Eqn. (9). However, since we are interested in the footprint function instead of the density field, we can reach a closed form solution by slightly modifying this function, which makes the splatting algorithm a perfect fit. Specifically, we start analyzing the intersection by creating a plane as a 2D coordinate system in Fig. 4 (b). We look at an infinitesimally small range around  $\mathbf{x}(t_i)$  as shown in Fig. 4 (c), such that the geometry field at the intersected line, drawn as an orange line, between the Gaussian surfel and the plane is constant, i.e.,  $f_i - c$ . We enable the calculation of Eqn. (9) and Eqn. (10) by extruding the intersected line with a linear decay. The geometry field value decays linearly from  $f_i - c$  to  $-c$  with the direction of  $\mathbf{n}_i$ , which is perpendicular to the plane containing the surfel, and length of  $h$ , as shown in Fig. 4 (d). It is equivalent to the original case as  $h \rightarrow 0$ . The  $F(\mathbf{x}(t))$  near  $t_i$  is modified as:

$$F(\mathbf{x}(t)) = f_i \times \left(1 - \frac{|t - t_i|}{h / \cos \theta_i}\right) - c, \text{ when } |t - t_i| < h / \cos \theta_i, \quad (13)$$

where  $\cos \theta_i = |\omega \cdot \mathbf{n}_i|$ . The  $\|\nabla F(\mathbf{x})\|$  near  $t_i$  is then calculated as:

$$\|\nabla F(\mathbf{x}(t))\| = f_i / h, \text{ when } |t - t_i| < h / \cos \theta_i \quad (14)$$

and it is parallel to  $\mathbf{n}_i$ . Therefore, we can evaluate Eqn. (10)




Figure 5. Illustration of opacity value on different Gaussian surfels. **(a)** 2DGS only allows the opacity to reach 1 at the center. **(b)** Due to the enforced transformation from geometry field into density as in Eqn. (15), the opacity on a Gaussian surfel in our case does not follow a Gaussian distribution, thus differentiating it from Gaussian splatting. Furthermore, we also allow a larger central area to reach 1, thus making the Gaussian surfel more opaque which benefits the surface reconstruction.

with Eqn. (9) as:

$$\begin{aligned} \rho_i &= \int_{t_i - h / \cos \theta_i}^{t_i + h / \cos \theta_i} \frac{\psi(-F(\mathbf{x}(t)))}{\Psi(-F(\mathbf{x}(t)))} \|\nabla F(\mathbf{x}(t))\| \cdot |\omega \cdot \mathbf{n}(\mathbf{x}(t))| dt \\ &= -2 \ln \Psi(c - f_i). \end{aligned} \quad (15)$$

Please find the derivation in A.2 of the supplementary. We define  $g(u) := -2 \ln \Psi(c - u)$  to simplify the notation. As the expansion goes to 0, i.e., without extrusion,  $\lim_{h \rightarrow 0} \rho_i = g(f_i)$ .

Notice that the value range of  $g(u)$  exceeds  $[0, 1]$ , making approximated volume splatting in Eqn. (2) not applicable and our refined volume splatting in Eqn. (7) suitable. We can then use it for rendering in Eqn. (7) as:

$$\mathbf{C} = \sum_{i=1}^N \mathbf{c}_i (1 - \exp(-g(f_i))) \prod_{j=1}^{i-1} \exp(-g(f_j)). \quad (16)$$

We now remove the non-overlapping assumption. We define  $S(u) := -2 \ln \Psi(c - u)$ . Notice that  $S(0) = 0$ , and  $S(u)$  is a monotonically increasing function with respect to  $u$ . Therefore, there exists a  $S^{-1}(v), \forall v \geq 0$ . We then define  $\oplus(\cdot)$ , such that  $a \oplus b = S^{-1}(S(a) + S(b))$ . Besides, when  $m$  Gaussian surfels  $\mathcal{G}_{j+1}, \mathcal{G}_{j+2}, \dots, \mathcal{G}_{j+m}, j \in \mathbb{N}_+, m \in \mathbb{N}_+$  intersect the ray at the same place  $\mathbf{x}(t_{j+1})$ , from Lemma. 1, Eqn. (7) is equivalent to having a kernel whose footprint function equals to  $\sum_{k=1}^m g(w_{j+k} \mathcal{G}_{j+k}(\mathbf{x}(t_{j+1})))$ .

Given that

$$g\left(\bigoplus_{k=1}^m w_{j+k} \mathcal{G}_{j+k}(\mathbf{x}(t_{j+1}))\right) = \sum_{k=1}^m g(w_{j+k} \mathcal{G}_{j+k}(\mathbf{x}(t_{j+1}))),$$

it then corresponds to the geometry field value:

$$\begin{aligned} F(\mathbf{x}(t_{j+1})) &= \bigoplus_{k=1}^m w_{j+k} \mathcal{G}_{j+k}(\mathbf{x}(t_{j+1})) - c \\ &= \bigoplus_{i=1}^N w_i \mathcal{G}_i(\mathbf{x}(t_{j+1})) - c, \end{aligned} \quad (17)$$

as other kernels are 0 here. It follows our definition Eqn. (8).




Figure 6. We illustrate the discontinuous change of rendered color in the splatting algorithm with respect to the positions of kernels. As shown in (a) and (b), when two kernels continuously change their positions and are eventually swapped, the ordering could be unstable and the rendered color  $C$  undergoes a discontinuous change when the ordering changes. The opacity is denoted on top of the kernel.

Another benefit that comes with using Gaussian surfels is that we can have the expected depth  $D$  efficiently calculated with volume splatting as:

$$D = \sum_{i=1}^N t_i (1 - \exp(-g(f_i))) \prod_{j=1}^{i-1} \exp(-g(f_j)). \quad (18)$$

In contrast, with 3D Gaussians, the assumption for volume splatting that the property, i.e., depth in this case, which we want to render, is constant inside the kernel does not hold.

**Discussion.** From Theorem 1, if Gaussian surfels have the same color when they have the same intersected point, the rendering is then exact, but we make the only approximation to use the global sorting to approximate the per-ray sorting as in [20, 23] for efficiency. We will discuss in Sec. 4.2 that this color constraint is also important for having a continuous loss landscape for optimization.

In practice, as in Fig. 5, the distributions of opacity on the Gaussian surfels do not follow the Gaussian distribution anymore, thus differentiating our algorithm from Gaussian splatting. We also allow larger areas on the Gaussian surfels to become fully opaque. These benefit the surface reconstruction as shown in Sec. 5.

## 4.2. Remedying Loss Landscape Defects

With the rendering algorithm we introduced before, during training, we apply the view synthesis loss as  $\mathcal{L}_{\text{rgb}}$ , depth distortion loss as  $\mathcal{L}_d$  and depth-normal consistency loss as  $\mathcal{L}_n$ , following [20]. However, we use the expected depth defined in Eqn. (18) instead of the median depth in [20] to calculate  $\mathcal{L}_n$ . The final loss is defined as:

$$\mathcal{L} = \mathcal{L}_{\text{rgb}} + \lambda_1 \mathcal{L}_d + \lambda_2 \mathcal{L}_n, \quad (19)$$

where  $\lambda_1$  is set specific to the dataset, and  $\lambda_2 = 0.05$ .  $\mathcal{L}_d$  helps encourage the Gaussian surfels to cluster together, such that the represented stochastic geometry becomes deterministic,  $\mathcal{L}_d$  helps smooth the geometry, and  $\mathcal{L}_{\text{rgb}}$  is the main loss for driving the surface reconstruction. However,  $\mathcal{L}_{\text{rgb}}$  is not a continuous function of properties, including center and covariance, of the Gaussian surfel. As shown in Fig. 6, when two Gaussian surfels continuously change their intersected




Figure 7. Demonstration of continuously changing the colors of kernels when two kernels gradually swap their positions.

points and swap their position, the rendered color undergoes a discontinuous change. In the case of surface reconstruction, Gaussian surfels are encouraged to cluster together, which makes the ordering unstable and optimization harder.

We notice that, as shown in Fig. 7, if the color also continuously changes when the Gaussian surfel moves in such a way that whenever two Gaussian surfels intersect the ray at the exact same position, they have the same color, this discontinuity is solved. Therefore, given a ray which intersects with Gaussian surfels  $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_N$ , we propose to replace the color  $c_i$  of  $\mathcal{G}_i$  with  $\hat{c}_i$  as:

$$\hat{c}_i = \frac{\sum_{j=1}^N (1 - \exp(-w_j)) \exp(-\tau |t_j - t_i|) c_j}{\sum_{j=1}^N (1 - \exp(-w_j)) \exp(-\tau |t_j - t_i|)}, \quad (20)$$

where  $\tau = 100$ , and we blend the color based on the weight and distance.

**Theorem 2.** With the color defined in Eqn. (20), Eqn. (16) gives a continuous function of properties of Gaussian surfels.

*Proof.*  $\hat{c}_i$  is a continuous function of  $t_1, t_2, \dots, t_N$ . And when there exists an  $i' \in \mathbb{N}$ , such that  $t_i = t_{i'}$ ,  $\hat{c}_i = \hat{c}_{i'}$ .  $\square$

With the color defined in Eqn. (20), we also reach an almost exact rendering algorithm as discussed in Sec. 4.1. However, it is computationally inefficient to blend the color per ray. We instead propose an efficient approximation that blends the color in the  $\mathbb{R}^3$  space. Specifically, we have:

$$\hat{c}_i = \frac{\sum_{j=1}^N (1 - \exp(-w_j)) \exp(-\tau \|\mathbf{m}_i - \mathbf{m}_j\|_2) c_j}{\sum_{j=1}^N (1 - \exp(-w_j)) \exp(-\tau \|\mathbf{m}_i - \mathbf{m}_j\|_2)}, \quad (21)$$

which can be efficiently implemented as propagating the color before rendering.

## 4.3. Improve Color Representation

The method described above significantly improves surface reconstruction quality using spherical harmonics as the color representation. However, we find it struggles to reconstruct the specular surfaces, as observed in [5, 22, 26, 40, 56].

We instead assign a latent  $\mathbf{l}_i \in \mathbb{R}^{32}$  to the  $i^{\text{th}}$  Gaussian surfel and create a shallow MLP  $\Phi$ , such that:

$$\mathbf{c}_i = \Phi(\mathbf{l}_i, \text{SE}(\boldsymbol{\omega}), \text{SE}(\boldsymbol{\omega}_o)), \quad (22)$$

where  $\text{SE}(\cdot)$  denotes the spherical harmonics encoding [35],  $\boldsymbol{\omega}$  denotes the ray direction, and  $\boldsymbol{\omega}_o$  denotes the reflected ray direction with respect to the local normal vector, i.e.,  $\boldsymbol{\omega} = (\mathbf{m}_i - \mathbf{o}) / \|\mathbf{m}_i - \mathbf{o}\|_2$ , and  $\boldsymbol{\omega}_o = 2(\mathbf{n}_i \cdot \boldsymbol{\omega})\mathbf{n}_i - \boldsymbol{\omega}$ . We find it beneficial to input both  $\boldsymbol{\omega}$  and  $\boldsymbol{\omega}_o$  for dealing with both diffuse and specular surfaces, as in [41, 44].

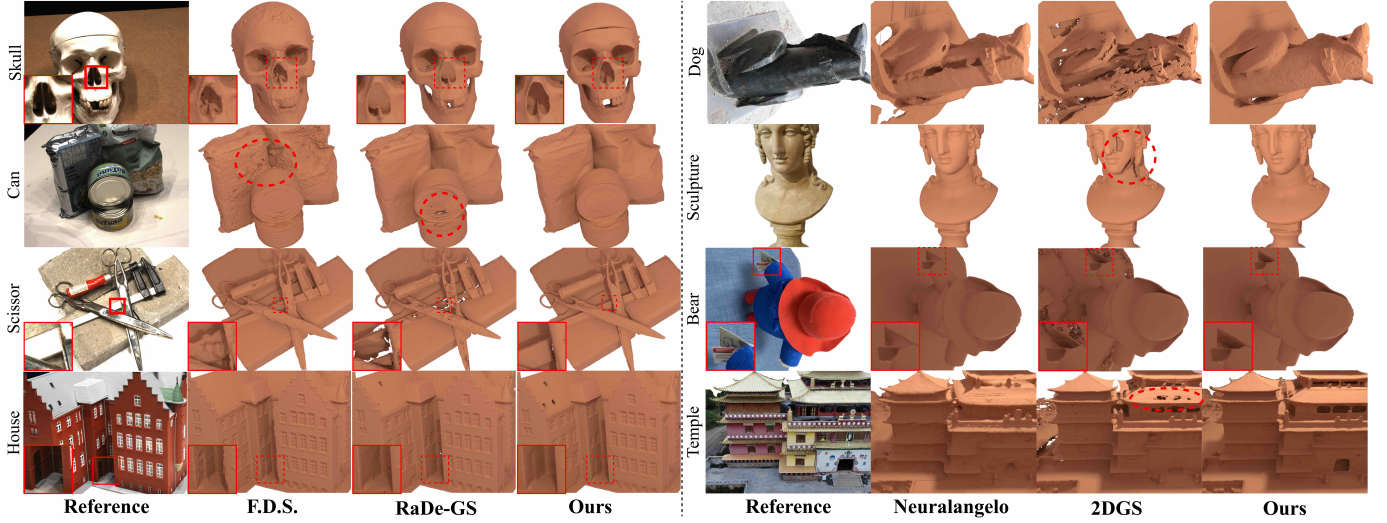


Figure 8. Qualitative comparison on DTU and BlendedMVS datasets. Non-obvious differences are highlighted by insets and red circles. “F.D.S.” denotes Chen et al. [5].

Methods	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Avg. ↓	Time
NeuS2 [48]	0.59	0.79	0.74	0.40	0.93	0.74	0.85	1.37	1.37	0.82	0.74	0.86	0.44	0.60	0.58	0.79	4m
N.A. [25]	0.39	0.79	0.38	0.30	0.82	0.72	1.91	1.20	1.94	0.67	0.38	1.15	0.34	0.84	0.95	0.85	> 12h
O.V. [31]	1.53	1.21	1.06	0.74	0.72	1.14	0.75	0.98	1.47	0.86	0.69	0.80	0.56	0.62	0.73	0.92	> 6h
NeUDF [27]	2.01	0.75	1.35	0.57	0.95	0.75	0.89	0.91	1.29	0.56	0.86	1.41	0.58	0.77	0.59	0.95	> 6h
F.D.S. [5]	0.44	0.66	0.33	0.32	0.90	0.67	0.54	1.28	1.10	0.70	0.65	0.57	0.41	0.40	0.39	0.62	1h
SuGaR [18]	1.47	1.33	1.13	0.61	2.25	1.71	1.15	1.63	1.62	1.07	0.79	2.45	0.98	0.88	0.79	1.33	42m
2DGS [20]	0.41	0.67	0.33	0.36	0.95	0.81	0.77	1.24	1.19	0.68	0.65	1.28	0.35	0.65	0.46	0.72	6m
G.S. [7]	0.64	0.85	0.58	0.43	0.99	1.18	0.90	1.14	1.11	0.81	0.92	1.46	0.53	0.61	0.58	0.85	3m
GOF [59]	0.46	0.68	0.39	0.38	1.15	0.82	0.70	1.14	1.21	0.65	0.68	1.04	0.48	0.68	0.49	0.73	33m
RaDe-GS [60]	0.46	0.66	0.35	0.41	0.79	0.73	0.67	1.14	1.18	0.61	0.64	0.87	0.37	0.69	0.48	0.67	10m
Ours (SH)	0.38	0.63	0.30	0.35	0.79	0.63	0.65	1.10	1.21	0.62	0.48	1.18	0.33	0.42	0.38	0.63	10m
Ours (Latent)	0.40	0.59	0.39	0.38	0.72	0.59	0.65	1.08	0.93	0.59	0.50	0.67	0.34	0.47	0.40	0.58	11m

Table 1. Quantitative evaluation on the DTU dataset based on the Chamfer Distance. Throughout the paper, the best metric is highlighted in red, the second best metric is highlighted in orange, and the third best metric is highlighted in yellow.

Methods	EvaUni	Temple	Excava	Museum	Avg. ↓
F.D.S. [5]	N/A	1.87	1.08	1.92	1.62
N.A. [25]	1.49	2.20	0.88	2.07	1.66
2DGS [20]	1.56	2.10	0.98	2.21	1.71
G.S. [7]	2.18	2.26	N/A	2.88	2.44
GOF [59]	1.55	1.99	0.81	2.23	1.65
RaDe-GS [60]	1.67	2.01	0.88	2.19	1.69
Ours (SH)	1.48	1.89	0.72	1.98	1.52
Ours (Latent)	1.58	2.13	0.87	2.26	1.71

Table 2. Quantitative comparison on the BlendedMVS dataset for scene-level cases based on the Chamfer Distance ( $\times 10^{-1}$ ). “N/A” denotes that the method fails to converge.

Methods	Bea	Clo	Dog	Dur	Jad	Man	Scu	Sto	Avg. ↓
F.D.S. [5]	0.33	0.10	0.21	2.63	0.19	0.80	0.44	0.91	0.70
N.A. [25]	0.32	0.12	0.31	3.11	0.15	0.80	0.43	0.77	0.75
2DGS [20]	1.17	0.46	0.45	2.87	0.15	0.53	0.59	0.90	0.89
G.S. [7]	0.57	0.35	0.27	6.62	0.16	0.53	0.76	1.45	1.34
GOF [59]	0.59	0.26	0.34	2.72	0.12	0.53	0.48	1.03	0.76
RaDe-GS [60]	0.51	0.25	0.24	2.83	0.12	0.52	0.43	0.84	0.72
Ours (SH)	0.55	0.17	0.40	2.63	0.14	0.43	0.36	0.71	0.67
Ours (Latent)	0.60	0.13	0.25	2.58	0.14	0.44	0.36	0.80	0.66

Table 3. Quantitative comparison on the BlendedMVS dataset for object-centric cases based on the Chamfer Distance ( $\times 10^{-2}$ ).

k-points [36], where  $k = 10$ . All the experiments are conducted on an NVIDIA 6000 Ada.

## 5. Results

We discuss implementation details and evaluation here. For detailed proof, further implementation details, evaluation results, and discussions, please refer to the supplementary.

**Implementation Details.** We base our implementations and hyper-parameters on [20] and use [32] for implementing the MLP. We replace the original densification strategy [20, 23] with [57]. We use the TSDF fusion [62] to extract the mesh from rendered depth maps. The proposed color blending over all kernels is approximated through closest-

### 5.1. Qualitative Evaluation

As shown in Fig. 8, we compare with Chen et al. [5], Neuralangelo [25], RaDe-GS [60], and 2DGS [20] as representative methods on DTU [21], and BlendedMVS [53]. Our method better captures geometric details (e.g., nose of the skull, door of the house), while handling specular surfaces well (e.g., reflective cans and scissors), compared to Chen et al. [5] and RaDe-GS. Our method also ensures smooth geometry without holes (e.g., dog, sculpture and temple).





Figure 9. We demonstrate the differences between SH and latent color representations on the “scan110” sample of the DTU dataset. When reconstructing specular surfaces (left), the SH representation (middle) leads to visible holes in the recovered mesh; these artifacts are resolved when using a latent representation (right).

Neuralangelo fails to properly represent open-boundary objects (e.g., card held by the bear). We further prepare a preliminary experiment on highly specular surfaces in C.2 of the supplementary.

## 5.2. Quantitative Evaluation

We evaluate on the datasets with ground-truth geometries and their corresponding camera parameters. We compare with methods, including NeuS2 [48], Neuralangelo [25], Miller et al. [31], NeUDF [27] as representative neural methods, SuGaR [18], 2DGS [20], Dai et al. [7], GOF [59], RaDe-GS [60] as representative splatting methods, and Chen et al. [5] on the DTU [21] dataset in Tab. 1. Besides, we compare with Chen et al. [5], Neuralangelo, 2DGS, Dai et al. [7], GOF and RaDe-GS on the BlendedMVS [53] dataset for scene-level cases in Tab. 2 and object-centric cases in Tab. 3. We unify the evaluation protocol and do not apply masking loss during training, nor iterative closest point during evaluation.

It is noteworthy that our method has the best averaged metrics. We achieve the best or second best results across almost all diffuse scenes while using the SH representation. For specular cases (i.e., “scan97”, “scan110” of DTU, and “Dog”, “Clock” of BlendedMVS), our method with improved color representation significantly improves the reconstruction quality and achieves one of the top three metrics. We are also significantly faster than other competing methods, including, Neuralangelo and Chen et al. [5]. The additional computational cost compared to 2DGS comes from our color blending which relies on the closest-k-points.

## 5.3. Ablation Study

We evaluate our proposed components on the DTU dataset [21] with both SH representation and latent representation for colors, and report the Chamfer Distance (denoted as CD) in Tab. 4. Specifically, we replace the proposed geometry field-based splatting algorithm with the original 2DGS splatting algorithm, which leads to a severe drop of geometry quality. Besides, we also remove the proposed color propagation used for remedying loss landscape defects, which also leads to a drop in geometry quality. Furthermore, we study the effects of the only approximation we make, i.e., using global sorting to approximate the per-ray sorting. We

Methods	Ours (SH)	Ours (Latent)
Full Model	0.63	0.58
- Geometry Field Splatting	0.70	0.72
- Remedy Loss Landscape Defects	0.67	0.59
+ Per-Ray Sorting	0.61	0.57

Table 4. Ablation study for the proposed components on the DTU dataset based on the Chamfer Distance.

implement a per-ray priority queue to sort the surfels based on the depths they intersect with the ray. Even though the per-ray sorting further improves the geometry quality, it consumes much more time ( $3 - 4\times$ ) and memory. We then do not use per-ray sorting by default and in other experiments.

Moreover, we also demonstrate the benefits of using the latent representation for colors on the specular surfaces in Fig. 9. With SH representation, there are holes on the specular surface, while latent representation is free of this problem.

## 6. Conclusion

In this paper, we propose to first define a stochastic geometry field in the space with Gaussian surfels, then convert it into the density field, and render it with an efficient and differentiable refined splatting algorithm for accurate geometry reconstruction. We also identify the discontinuity of view synthesis loss and propose an efficient remedy to address it such that the represented stochastic geometry can converge to the deterministic geometry well through the optimization. Our method is not without limitations. For example, our method cannot handle transparent or semi-transparent objects due to the opaque assumption, or fuzzy objects well due to the smoothness constraint.

In conclusion, we show that it is possible to use Gaussian surfels to accurately define the geometry and then perform efficient and differentiable almost exact rendering for geometry reconstruction. We achieve significant improvement over the geometry quality compared to other baselines, including both neural and splatting-based methods.

**Acknowledgements** This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DOI/IBC) contract number 140D0423C0076. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government. We also acknowledge ONR grant N00014-23-1-2526, and NSF grants 2100237 and 2120019 for the Nautilus cluster, gifts from Adobe, Google, Qualcomm and Rembrand, the Ronald L. Graham Chair and the UC San Diego Center for Visual Computing. We thank Yash Belhe for discussions, proofreading and feedback. We also thank the NVIDIA for GPU gifts.



## References

- [1] Michael Bleyer, Christoph Rhemann, and Carsten Rother. Patchmatch stereo - stereo matching with slanted support windows. In *British Machine Vision Conference*, 2011. 2
- [2] Changjiang Cai, Pan Ji, Qingan Yan, and Yi Xu. Riav-mvs: Recurrent-indexing an asymmetric volume for multi-view stereo. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 919–928, 2023. 2
- [3] Danpeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin Liu, Hujun Bao, and Guofeng Zhang. Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *arXiv preprint arXiv:2406.06521*, 2024. 2
- [4] Decai Chen, Brianne Oberson, Ingo Feldmann, Oliver Schreer, Anna Hilsmann, and Peter Eisert. Adaptive and temporally consistent gaussian surfels for multi-view dynamic reconstruction. *arXiv preprint arXiv:2411.06602*, 2024. 1
- [5] Hanyu Chen, Bailey Miller, and Ioannis Gkioulekas. 3d reconstruction with fast dipole sums. *arXiv preprint arXiv:2405.16788*, 2024. 1, 6, 7, 8
- [6] Brian Curless and Marc Levoy. *A Volumetric Method for Building Complex Models from Range Images*. Association for Computing Machinery, New York, NY, USA, 1 edition, 2023. 2
- [7] Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. High-quality surface reconstruction using gaussian surfels. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 1, 2, 4, 7, 8
- [8] François Darmon, Bénédicte Bascle, Jean-Clément Devaux, Pascal Monasse, and Mathieu Aubry. Improving neural implicit surfaces geometry with patch warping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6260–6269, 2022. 2
- [9] Yikang Ding, Wentao Yuan, Qingtian Zhu, Haotian Zhang, Xiangyue Liu, Yuanjiang Wang, and Xiao Liu. Transmvsnet: Global context-aware multi-view stereo network with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8585–8594, 2022. 2
- [10] Yikang Ding, Qingtian Zhu, Xiangyue Liu, Wentao Yuan, Haotian Zhang, and Chi Zhang. Kd-mvs: Knowledge distillation based self-supervised learning for multi-view stereo. In *European conference on computer vision*, pages 630–646. Springer, 2022. 2
- [11] Zheng Dong, Ke Xu, Yaoan Gao, Hujun Bao, Weiwei Xu, and Rynson W. H. Lau. Gaussian surfel splatting for live human performance capture. *ACM Trans. Graph.*, 43(6), 2024. 1
- [12] Qiancheng Fu, Qingshan Xu, Yew Soon Ong, and Wenbing Tao. Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 35:3403–3416, 2022. 2
- [13] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(8):1362–1376, 2010. 2
- [14] Wenhao Ge, Tao Hu, Haoyu Zhao, Shu Liu, and Ying-Cong Chen. Ref-neus: Ambiguity-reduced neural implicit surface learning for multi-view reconstruction with reflection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4251–4260, 2023. 2
- [15] Michael Goesele, Noah Snavely, Brian Curless, Hugues Hoppe, and Steven M. Seitz. Multi-view stereo for community photo collections. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8, 2007. 2
- [16] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020. 1
- [17] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuoqun Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2495–2504, 2020. 2
- [18] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. *CVPR*, 2024. 1, 2, 7, 8
- [19] Heiko Hirschmüller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):328–341, 2008. 2
- [20] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 1, 2, 4, 6, 7, 8
- [21] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engil Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 406–413, 2014. 7, 8
- [22] 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. 2, 6
- [23] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 1, 2, 3, 4, 6, 7
- [24] Hamid Laga, Laurent Valentin Jospin, Farid Boussaid, and Mohammed Bennamoun. A survey on deep learning techniques for stereo-based depth estimation. *IEEE transactions on pattern analysis and machine intelligence*, 44(4):1738–1764, 2020. 2
- [25] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 7, 8
- [26] Yuan Liu, Peng Wang, Cheng Lin, Xiaoxiao Long, Jiepeng Wang, Lingjie Liu, Taku Komura, and Wenping Wang. Nero: Neural geometry and brdf reconstruction of reflective objects from multiview images. *ACM Transactions on Graphics (TOG)*, 42(4):1–22, 2023. 6
- [27] Yu-Tao Liu, Li Wang, Jie Yang, Weikai Chen, Xiaoxu Meng, Bo Yang, and Lin Gao. Neudf: Learning neural unsigned distance fields with volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 237–247, 2023. 1, 2, 7, 8

- [28] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In *European Conference on Computer Vision*, pages 210–227. Springer, 2022. 2
- [29] Xiaoxiao Long, Cheng Lin, Lingjie Liu, Yuan Liu, Peng Wang, Christian Theobalt, Taku Komura, and Wenping Wang. Neuraludf: Learning unsigned distance fields for multi-view reconstruction of surfaces with arbitrary topologies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20834–20843, 2023. 1, 2
- [30] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2
- [31] Bailey Miller, Hanyu Chen, Alice Lai, and Ioannis Gkioulekas. Objects as volumes: A stochastic geometry view of opaque solids. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 87–97, 2024. 1, 2, 3, 4, 7, 8
- [32] Thomas Müller. tiny-cuda-nn, 2021. 7
- [33] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022. 2
- [34] 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*, page 335–342, USA, 2000. ACM Press/Addison-Wesley Publishing Co. 4
- [35] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, page 497–500, New York, NY, USA, 2001. Association for Computing Machinery. 2, 6
- [36] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 7
- [37] Xinlin Ren, Chenjie Cao, Yanwei Fu, and Xiangyang Xue. Improving neural surface reconstruction with feature priors from multi-view image. *arXiv preprint arXiv:2408.02079*, 2024. 2
- [38] Johannes L. Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *Computer Vision – ECCV 2016*, pages 501–518, Cham, 2016. Springer International Publishing. 2
- [39] Jian Sun, Nan-Ning Zheng, and Heung-Yeung Shum. Stereo matching using belief propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(7):787–800, 2003. 2
- [40] 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. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5481–5490. IEEE, 2022. 2, 6
- [41] Dor Verbin, Pratul P Srinivasan, Peter Hedman, Ben Mildenhall, Benjamin Attal, Richard Szeliski, and Jonathan T Barron. Nerf-casting: Improved view-dependent appearance with consistent reflections. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–10, 2024. 2, 6
- [42] Hoang-Hiep Vu, Patrick Labatut, Jean-Philippe Pons, and Renaud Keriven. High accuracy and visibility-consistent dense multiview stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(5):889–901, 2012. 2
- [43] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patchmatchnet: Learned multi-view patchmatch stereo. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14194–14203, 2021. 2
- [44] Fangjinhua Wang, Marie-Julie Rakotosaona, Michael Niemeyer, Richard Szeliski, Marc Pollefeys, and Federico Tombari. Unisdf: Unifying neural representations for high-fidelity 3d reconstruction of complex scenes with reflections. *Advances in Neural Information Processing Systems*, 37:3157–3184, 2025. 2, 6
- [45] 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. *NeurIPS*, 2021. 1, 2, 3
- [46] Shaoqian Wang, Bo Li, and Yuchao Dai. Efficient multi-view stereo by iterative dynamic cost volume. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8655–8664, 2022. 2
- [47] Xiang Wang, Chen Wang, Bing Liu, Xiaoqing Zhou, Liang Zhang, Jin Zheng, and Xiao Bai. Multi-view stereo in the deep learning era: A comprehensive review. *Displays*, 70: 102102, 2021. 2
- [48] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3295–3306, 2023. 1, 2, 7, 8
- [49] Yaniv Wolf, Amit Bracha, and Ron Kimmel. GS2Mesh: Surface reconstruction from Gaussian splatting via novel stereo views. In *European Conference on Computer Vision*, pages 207–224. Springer, 2025. 1
- [50] Tong Wu, Jiaqi Wang, Xingang Pan, Xudong Xu, Christian Theobalt, Ziwei Liu, and Dahua Lin. Voxurf: Voxel-based efficient and accurate neural surface reconstruction. *arXiv preprint arXiv:2208.12697*, 2022. 2
- [51] Kaiqiang Xiong, Rui Peng, Zhe Zhang, Tianxing Feng, Jianbo Jiao, Feng Gao, and Ronggang Wang. Cl-mvsnet: Unsupervised multi-view stereo with dual-level contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3769–3780, 2023. 2
- [52] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, pages 767–783, 2018. 2
- [53] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. *Computer Vision and Pattern Recognition (CVPR)*, 2020. 7, 8
- [54] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural

- surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33:2492–2502, 2020. 1, 2
- [55] 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. 1, 2
- [56] Keyang Ye, Qiming Hou, and Kun Zhou. 3d gaussian splatting with deferred reflection. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–10, 2024. 6
- [57] Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. Absgs: Recovering fine details in 3d gaussian splatting. In *ACM Multimedia 2024*, 2024. 7
- [58] Zehao Yu and Shenghua Gao. Fast-mvsnet: Sparse-to-dense multi-view stereo with learned propagation and gauss-newton refinement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1949–1958, 2020. 2
- [59] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient high-quality compact surface reconstruction in unbounded scenes. *arXiv:2404.10772*, 2024. 1, 2, 7, 8
- [60] Baowen Zhang, Chuan Fang, Rakesh Shrestha, Yixun Liang, Xiaoxiao Long, and Ping Tan. Rade-gs: Rasterizing depth in gaussian splatting. *arXiv preprint arXiv:2406.01467*, 2024. 1, 2, 7, 8
- [61] Yiqun Zhao, Chenming Wu, Binbin Huang, Yihao Zhi, Chen Zhao, Jingdong Wang, and Shenghua Gao. Surfel-based gaussian inverse rendering for fast and relightable dynamic human reconstruction from monocular video. *arXiv preprint arXiv:2407.15212*, 2024. 1
- [62] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3d: A modern library for 3d data processing. *arXiv preprint arXiv:1801.09847*, 2018. 2, 7
- [63] Qingtian Zhu, Chen Min, Zizhuang Wei, Yisong Chen, and Guoping Wang. Deep learning for multi-view stereo via plane sweep: A survey. *arXiv preprint arXiv:2106.15328*, 2021. 2
- [64] M. Zwicker, H. Pfister, J. van Baar, and M. Gross. Ewa volume splatting. In *Proceedings Visualization, 2001. VIS '01.*, pages 29–538, 2001. 1, 2