

# RayletDF: Raylet Distance Fields for Generalizable 3D Surface Reconstruction from Point Clouds or Gaussians

Shenxing Wei<sup>†</sup> Jinxi Li<sup>†</sup> Yafei Yang Siyuan Zhou Bo Yang<sup>\*</sup>  
 vLAR Group, The Hong Kong Polytechnic University

{shenxing.wei, jinxi.li, ya-fei.yang, siyuan.zhou}@connect.polyu.hk, bo.yang@polyu.edu.hk




Figure 1. We train all methods on the large-scale ARKitScene dataset and then directly test them on the unseen ScanNet++ dataset. Our RayletDF shows superior generalizability compared to baselines, achieving significantly more accurate 3D scene surface reconstructions.

## Abstract

*In this paper, we present a generalizable method for 3D surface reconstruction from raw point clouds or pre-estimated 3D Gaussians by 3DGS from RGB images. Unlike existing coordinate-based methods which are often computationally intensive when rendering explicit surfaces, our proposed method, named **RayletDF**, introduces a new technique called raylet distance field, which aims to directly predict surface points from query rays. Our pipeline consists of three key modules: a raylet feature extractor, a raylet distance field predictor, and a multi-raylet blender. These components work together to extract fine-grained local geometric features, predict raylet distances, and aggregate multiple predictions to reconstruct precise surface points. We extensively evaluate our method on multiple public real-world datasets, demonstrating superior performance in surface reconstruction from point clouds or 3D Gaussians. Most notably, our method achieves exceptional generalization ability, successfully recovering 3D surfaces in a single-forward pass across unseen datasets in testing. Our code and datasets are available at <https://github.com/vLAR-group/RayletDF>.*

## 1. Introduction

Learning efficient, accurate, and generalizable 3D surface representations is crucial for many applications in mixed reality, embodied AI, and graphics. Given RGB/D images and/or point clouds, a series of 3D representations has

been developed to recover 3D geometry, including occupancy fields (OF) [38], un/signed distance fields (U/SDF) [13, 44], radiance fields (NeRF) [39], and vector fields (VF) [70]. While achieving excellent results in 3D reconstruction, these coordinate-based methods and their variants typically demand dense network evaluations to obtain explicit surfaces, thus being computationally heavy.

Recently, the point-based method 3D Gaussian Splatting (3DGS) [30] has emerged as an appealing alternative to those coordinate-based methods, thanks to its impressive real-time performance in synthesizing high-fidelity RGB images at high resolutions. However, it still falls short in rendering high-quality depth views, due to its failure in capturing fine-grained surface geometry, though various constraints such as depth priors [33, 62], local smoothness or planarity regularization [24, 63] can be added to enhance the quality of surfaces recovered.

On the other hand, a line of recent ray-based methods such as DRDF [32], PRIF [20], DDFs [1], Pointersect [8], and RayDF [36] have demonstrated excellent performance in representing intricate surfaces based on light rays. Nevertheless, due to the limitation of existing ray parametrizations such as Plucker and spherical coordinates, they are often limited to recovering object-level surfaces and require per-scene training, lacking the desired generalizability to infer diverse 3D scene structures in a single forward pass.

In this paper, we present a generalizable 3D surface representation pipeline to accurately recover 3D geometry. The core module of our pipeline is a new technique called raylet distance field, with inspiration from ray-based methods.

The term *raylet* is defined as a unit segment of a light ray

<sup>†</sup> Equal contribution <sup>\*</sup> Corresponding author




Figure 2. An illustration of raylets and raylet distances.

and its starting point is sampled (or located) near the surface of a shape. The *raylet distance* is defined as the signed distance between the surface hit point and the raylet starting point. As illustrated in Figure 2, given three light rays  $\{r_1, r_2, r_3\}$  emitting from light sources, we sample three raylets  $\{l_1, l_2, l_3\}$  near the surface  $S$ , where raylet distance  $d_1$  is assigned positive as its surface hit point  $p$  locates further away from the raylet starting point  $p_{l_1}$ , whereas  $\{d_2, d_3\}$  are assigned negative accordingly. Note that, given a single light ray, we can sample numerous raylets on both sides of the surface hit point. Intuitively, such definitions of raylets and raylet distances primarily focus on nuanced local patterns of surfaces, and these local geometric patterns are usually generalizable across diverse shapes.

With this merit of raylets, we simply formulate the problem of generalizable 3D surface reconstruction into learning raylet distance fields from visual observations. In particular, our pipeline comprises three modules: 1) a **raylet feature extractor** to extract geometry features from an input 3D scene for a query raylet; 2) a **raylet distance field** to predict the signed raylet-surface distance value for the input query raylet; and 3) a **multi-raylet blender** to aggregate multiple predicted raylet distances along every single ray, ultimately recovering the accurate surface point.

These three modules together allow us to learn generalizable ray-based surface representations from large-scale training data with ground truth depth scans as supervision signals. Once the pipeline is well trained, it can be used to infer high-quality 3D scene surfaces on unseen datasets from arbitrary query viewing angles in a single forward pass. Notably, our pipeline can not only train and evaluate on point clouds, but is also amenable to any scene represented by 3D Gaussians recovered by the popular 3DGS [30] from RGB images. Our pipeline is called **RayletDF** and Figure 1 shows qualitative results on diverse indoor scenes. Our contributions are:

- We propose a generic pipeline for explicit 3D surface reconstruction from either point clouds or 3D Gaussians.
- We introduce a new raylet distance field followed by a blender to learn generalizable surface patterns.
- We show superior accuracy in surface reconstruction across multiple datasets, clearly outperforming baselines, especially in terms of generalizability to new datasets.
- We benchmark generalizable surface reconstruction on 3D Gaussians and will release 7770 3D scene Gaussians trained on ScanNet/++, ARKit, and MultiScan datasets.

## 2. Related Works

Classical approaches to recovering 3D surfaces from images mainly include SfM [43] and SLAM [6] systems to obtain sparse point clouds, optionally followed by various global or local smoothness priors to recover continuous surfaces, but the reconstructed shapes usually lack fine details. Early learning methods to model explicit 3D structures mainly include voxel grids [14, 67, 68], point clouds [17], octree [54], meshes [29] based pipelines, but the fidelity of these discrete shape representations is often limited by the spatial resolutions and memory footprint. A comprehensive survey of these methods can be found in [4, 25].

**Coordinate-based Methods for Surface Reconstruction:** To avoid the discretization issue of classical explicit 3D representations, a series of implicit representations has been developed to use simple MLPs to recover continuous 3D shapes [21]. These representations, including OF [12, 38], SDF [44], UDF [13, 58], VF [47, 70], and NeRF [39, 56], typically take 3D coordinates as input and predict various properties of input coordinates. While demonstrating exceptional accuracy in recovering continuous 3D surfaces and/or rendering 2D views, these methods and their variants [55, 61, 65] demand dense 3D point sampling and network evaluations to regress explicit surface points given any query view, thus being computationally heavy. In this paper, our RayletDF learns efficient raylet fields and does not require dense sampling to regress surface points.

**Point-based Methods for Surface Reconstruction:** Point-based methods aim to recover surfaces by rendering disconnected geometry samples, including the simplest point sample [23] and circular/elliptic discs/ellipsoids/surfels sample rendering methods [45, 48, 77]. Recently, a new point-based method 3DGS [30] has been introduced to represent the scene by a set of 3D Gaussians with various properties such as position, covariance, and color. While achieving real-time rendering of 2D views thanks to the tile-based rasterizer, 3DGS falls short in synthesizing highly accurate depth views as it is hard to capture fine-grained surfaces, though advanced techniques [9, 11, 15, 18, 24, 26, 27, 33, 42, 50, 63, 64, 66, 73–76] can mitigate this issue but at the expense of generalization across various scenes. Meanwhile, researchers also introduce fast winding number to balance the efficiency and accuracy [2, 10]. These approaches achieve better geometry fidelity and convergence speed but still need per-scene optimization and point normals. In this paper, our introduced raylet distance field and the blender aim to learn nuanced local surface features, thus being accurate and generalizable.

**Ray-based Methods for Surface Reconstruction:** A line of works formulates 3D shapes as ray-based neural functions. Existing ray-based methods simply take individual rays as input and estimate either RGB values such as LFN [49], NeuLF [34] and others [7, 19, 41, 51], or surface




Figure 3. Overview of our proposed pipeline. The leftmost block shows the raylet feature extractor module, the middle block shows the raylet distance field module, and the rightmost block shows our multi-raylet blender module.

points including PRIF [20], DDFs [1], RayDF [36], and others [28, 32, 52, 53, 72]. Although these methods have shown excellent efficiency and accuracy in rendering surfaces, they can typically represent object-level shapes and lack generality due to the limitation of light ray representations. In this paper, instead of using global light rays to represent surfaces, we use raylets to recover geometry.

**Image-based Depth Estimation:** With the advancement of techniques for novel view RGB synthesis, existing methods [5, 46, 69] for image-based depth estimation can also be used to recover per-view depth scans, followed by depth fusion for 3D surface reconstruction. However, such a pipeline inherently lacks consistent depth alignment and geometry consistency priors. Therefore, the recovered 3D structures are inferior in quality.

### 3. RayletDF

#### 3.1. Overview

Our RayletDF models 3D surface as a neural network  $f$ . It takes a scene volumetric data  $P$  and a query raylet  $l$  as input, directly predicting the corresponding raylet distance  $d_l$ , *i.e.*, the distance between surface hit point and raylet starting point as shown in Figure 2. The scene volumetric data can be a sparse point cloud or on-the-shelf 3D Gaussians recovered by the popular 3DGS [30] from RGBs, both denoted by  $P \in \mathcal{R}^{N \times C_0}$  for simplicity, where  $N$  represents the total number of surface points or Gaussian centers and  $C_0$  represents the number of per-point features such as point/Gaussian center  $xyz$  optionally with other properties like colors or opacity. The query raylet  $l \in \mathcal{R}^{1 \times 6}$  is parameterized by its starting point  $xyz$  and a unit vector representing ray direction. Formally, it is defined as below:

$$d_l = f(P, l) \quad P \in \mathcal{R}^{N \times C_0}, \quad l \in \mathcal{R}^{1 \times 6} \quad (1)$$

Since our method aims at representing accurate and generalizable 3D surface representations, the pipeline is designed to be trained on a large number of 3D scenes with depth images as supervision signals. Once well-trained, it is expected to be applied to various unseen datasets, directly

inferring high-quality depth maps at arbitrary query viewing angles. As shown in Figure 3, our pipeline has three components clarified in the following Sections 3.2&3.3&3.4.

#### 3.2. Raylet Feature Extractor

This module is designed to learn local features from an input scene volumetric data point  $P$  for any specific query raylet  $l$ . As shown in the leftmost block of Figure 3, given an input scene  $P$ , we simply adopt the powerful SparseConv architecture [22] without any pretraining step as the backbone network, obtaining per-point features  $F \in \mathcal{R}^{N \times C}$  where the embedding length  $C$  is predefined as 32. Implementation details are in Appendix 5.5.

Having the input scene  $P$  and its per-point features  $F$ , we need to extract the local features for a specific input query raylet  $l$  whose starting point is denoted by  $p_l$  and ray direction by a unit vector  $u_l$ . In particular, as shown in Figure 3, for the query raylet  $l$ , we identify the top  $K$  nearest points  $\{p_l^1 \cdots p_l^k \cdots p_l^K\}$  in the input scene  $P$  with regard to the raylet starting point  $p_l$ . The simple K-nearest neighbors (KNN) algorithm is adopted based on the point-wise Euclidean distances. After that, we gather the information, denoted by  $\hat{f}_l^k$ , for the  $k^{th}$  neighboring point as follows:

$$\hat{f}_l^k = \left( p_l^k \oplus \frac{(p_l^k - p_l)}{\|p_l^k - p_l\|} \oplus \|p_l^k - p_l\| \right) \oplus f_l^k \quad (2)$$

where  $p_l$  and  $p_l^k$  are the  $xyz$  positions of points,  $\oplus$  is the concatenation operation,  $\|\cdot\|$  calculates the Euclidean distance between two points, and  $f_l^k$  is the feature vector extracted from per-point features  $F$ .

Lastly, we stack all information of the total  $K$  neighboring points into a local feature vector  $f_l$  for the query raylet  $l$  for simplicity as follows, though other methods such as max/mean and attentional pooling are also applicable.

$$f_l = \hat{f}_l^1 \oplus \cdots \hat{f}_l^k \cdots \oplus \hat{f}_l^K \quad (3)$$

Overall, our key insight of this module is that the features extracted for a query raylet should retain local geometrical patterns near surfaces, allowing the learned raylet distance representations to be generalizable across diverse scenes.

### 3.3. Raylet Distance Field

This module aims to learn accurate raylet surface distance value for any specific query raylet  $l$ . As illustrated in the middle block of Figure 3, we feed the query raylet, *i.e.*, the starting point  $\mathbf{p}_l$  and direction  $\mathbf{u}_l$ , and its feature vector  $\mathbf{f}_l$  (ref to Section 3.2) into MLPs, directly predicting the corresponding raylet surface distance value  $d_l$ .

To improve the generalizability and robustness of raylet distance predictions, we further predict a confidence score  $s_l$  alongside the distance value  $d_l$ . This score is basically designed to weight the prediction and blend multiple raylets sampled along a single ray which will be detailed in Section 3.4. Formally, the raylet distance field is defined as:

$$(d_l, s_l) = MLPs(\mathbf{p}_l \oplus \mathbf{u}_l \oplus \mathbf{f}_l) \quad (4)$$

For fast computation, the MLPs consist of 8 layers and each layer has 256 hidden neurons, though more sophisticated neural layers can be applied as well. Since the predicted distance value could be positive or negative, there is no non-linear activation function added at the output layer. More details are in Appendix 5.6.

Our design of this raylet distance field is particularly simple in concept, but has a clear advantage over the prevailing coordinate-based surface representations such as OF, SDF, and UDF: our method takes raylets as input, without requiring sampling dense coordinates along a single ray, thus being efficient to regress surface points.

### 3.4. Multi-raylet Blender

The ultimate goal of our pipeline is to infer the surface point for any given light ray  $r$ , *i.e.*, the distance  $D$  between the query camera center  $\mathbf{p}_{cam}$  and the surface hit point along the ray. Since our raylet distance field learns representations for raylets instead of entire rays, this allows us to sample multiple raylets along each single ray followed by blending multiple estimations, potentially improving the generalizability and robustness of our method.

To this end, as shown in the rightmost block of Figure 3, we opt for sampling  $T$  raylets  $\{l_1 \cdots l_t \cdots l_T\}$  near the surface along a single ray  $r$ . These  $T$  raylets have the same unit vector of ray direction, but with different starting point positions, denoted by  $\{\mathbf{p}_{l_1} \cdots \mathbf{p}_{l_t} \cdots \mathbf{p}_{l_T}\}$ . Naturally, we feed these  $T$  raylets into our raylet distance field in parallel, obtaining the corresponding distances and scores:  $\{(d_{l_1}, s_{l_1}) \cdots (d_{l_t}, s_{l_t}) \cdots (d_{l_T}, s_{l_T})\}$ . After that, we blend all these predictions and estimate the distance between camera center  $\mathbf{p}_{cam}$  and surface hit point as follows:

$$D = \sum_{t=1}^T \hat{s}_{l_t} \left( \|\mathbf{p}_{cam} - \mathbf{p}_{l_t}\| + d_{l_t} \right), \hat{s}_{l_t} = \frac{e^{s_{l_t}}}{\sum_{t=1}^T e^{s_{l_t}}} \quad (5)$$

where  $\hat{s}_{l_t}$  is a normalized score.  $\mathbf{p}_{cam}$  and  $\mathbf{p}_{l_t}$  are always known during sampling in training and test.




Figure 4. Raylet sampling on virtual balls of point clouds.

The prediction  $D$  is fully supervised using  $\ell_1$  loss, and the ground truth distance values are converted from the available depth images in training set.

### 3.5. Sampling Raylets for Training and Test

Given a specific 3D scene  $\mathbf{P}$  as input, if it is a raw point cloud, for a specific query ray  $r$ , we sample multiple raylets for both training or test in the following steps:

- Step 1: For every point  $\mathbf{p}$  in the input scene point cloud  $\mathbf{P}$ , we search its nearest point  $\mathbf{p}'$  and identify the Euclidean distance  $\|\mathbf{p} - \mathbf{p}'\|$  as a radius of the ball centering at point  $\mathbf{p}$ . In this way, every point in  $\mathbf{P}$  is assigned a virtual ball. Intuitively, the underlying continuous 3D scene surface, which is unknown and yet to be estimated, should be bounded by the union space of all virtual spheres, as illustrated in the left block of Figure 4.
- Step 2: as illustrated in the right block of Figure 4, given the query ray  $r$  with a camera center  $\mathbf{p}_{cam}$ , it shoots through multiple virtual balls. We then project the center point of each intersected ball onto the ray, obtaining a perpendicular distance and the foot of the perpendicular. After that, we select  $T$  foot points, belonging to intersected virtual balls with top  $T$  shortest perpendicular distances, as the starting points of  $T$  raylets. If the number of intersected balls is fewer than  $T$ , we just keep the actual number of raylets. Note that, if there is no ball intersected, meaning that the ray shoots outside the target 3D surface, the ray is discarded in both training and test.

If the input 3D scene  $\mathbf{P}$  is a set of 3D Gaussians recovered by 3DGS [30] from RGBs, we follow the technique [31, 74] to calculate an intersection point between the query ray and each Gaussian, and then we select  $T$  intersection points, belonging to the Gaussians with top  $T$  contributions based on alpha blending, as the starting points of  $T$  raylets.

For efficiency, we follow the tile-based splitting strategy in 3DGS [30] to calculate intersection points on both point clouds and Gaussians. More details are in Appendix 5.7.

### 3.6. Surface Normal Derivation

Similar to RayDF, our method also has a nice property of deriving a closed-form expression for a normal vector, thanks to the nature of our ray-based formulation. In particular, given an input raylet  $l = (\mathbf{p}_l, \mathbf{u}_l)$  and its estimated raylet distance value  $d_l$  from the network, the corresponding normal vector  $\mathbf{n}_l$  is provided in Appendix 5.8.

This normal vector can be used in an additional loss to regularize surface points, or acts as a threshold to remove outlier predictions, which are left for future exploration.

Table 1. Quantitative results of all methods for ray surface distance estimation. † indicates the scale is aligned with ground truth depth.

		test on → ARKitScenes					test on → ScanNet/ScanNet++					test on → MultiScan				
		ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta$ ↑	ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta$ ↑	ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta$ ↑
DepthAnythingV2† [69]	3DGS [30]	0.268	0.379	0.173	0.098	0.707	0.321	0.461	0.152	0.149	0.761	0.431	0.532	0.322	0.225	0.402
	GOF [74]	0.318	0.469	0.213	0.159	0.659	0.320	0.507	0.155	0.214	0.775	0.571	0.700	0.404	0.373	0.363
	PGSR [9]	0.219	0.376	0.168	0.183	0.793	0.202	0.536	0.099	0.379	0.877	0.315	0.430	0.285	0.516	0.567
	Depth-Prof† [5]	0.206	0.294	0.144	0.093	0.860	0.168	0.248	0.118	0.066	0.887	0.228	0.303	0.205	0.117	0.775
	Depth-Prof† [5]	0.294	0.403	0.202	0.242	0.750	0.220	0.319	0.120	0.106	0.861	0.280	0.363	0.234	0.157	0.659
train on ARKitScenes	MVSGaussian [35]	0.584	0.755	0.392	0.416	0.425	0.592	0.797	0.395	0.376	0.538	0.543	0.676	0.444	0.434	0.363
	Pointersect [8]	0.286	0.397	0.209	0.135	0.749	0.366	0.513	0.190	0.151	0.733	0.266	0.351	0.217	0.117	0.671
	RayDF [36]	0.183	0.303	0.139	0.088	0.846	0.227	0.363	0.131	0.123	0.838	0.326	0.425	0.287	0.202	0.583
	PFGS [59]	0.264	0.386	0.171	0.098	0.719	0.303	0.455	0.145	0.145	0.779	0.536	0.653	0.442	0.375	0.316
	<b>RayletDF (Ours)</b>	<b>0.115</b>	<b>0.218</b>	<b>0.082</b>	<b>0.047</b>	<b>0.928</b>	<b>0.175</b>	<b>0.320</b>	<b>0.085</b>	<b>0.098</b>	<b>0.894</b>	<b>0.216</b>	<b>0.311</b>	<b>0.182</b>	<b>0.107</b>	<b>0.739</b>
train on ScanNet++	MVSGaussian [35]	0.600	0.774	0.404	0.448	0.413	0.591	0.787	0.295	0.367	0.528	0.548	0.681	0.448	0.445	0.363
	Pointersect [8]	0.328	0.462	0.196	0.142	0.665	0.433	0.604	0.198	0.199	0.658	0.404	0.504	0.296	0.198	0.449
	RayDF [36]	0.587	0.704	0.348	0.302	0.317	0.202	0.337	0.114	0.818	0.871	0.604	0.690	0.447	0.467	0.233
	PFGS [59]	0.274	0.399	0.177	0.104	0.704	0.286	0.433	0.137	0.142	0.798	0.534	0.669	0.448	2.449	0.327
	<b>RayletDF (Ours)</b>	<b>0.132</b>	<b>0.241</b>	<b>0.094</b>	<b>0.053</b>	<b>0.906</b>	<b>0.145</b>	<b>0.276</b>	<b>0.072</b>	<b>0.079</b>	<b>0.922</b>	<b>0.259</b>	<b>0.353</b>	<b>0.236</b>	<b>0.148</b>	<b>0.662</b>

## 4. Experiments

**Datasets:** Our method is evaluated on four real-world datasets based on the available train/test splits: **1) ScanNet** [16] consisting of 1201 and 100 scenes for training and test; **2) ScanNet++** [71] comprising 865 and 50 scenes for training and test; **3) ARKitScenes** [3] with 4498 and 549 scenes for training and test, which is the largest real-world indoor scene dataset captured by mobile RGBD sensor from Apple LiDAR scanner; **4) MultiScan** [37] only consisting of about 200 indoor scenes in total captured by smartphones and tablets, where all scenes are used as a test set in our experiments. Due to the domain overlap, we merge ScanNet and ScanNet++, getting a joint train set and a joint test set.

**Baselines:** We choose 5 representative groups of methods as our baselines: 1) the state-of-the-art per-scene optimization based 3D Gaussians splatting methods **GOF** [74] and **PGSR** [9] particularly designed for high-fidelity surface reconstruction; 2) the latest feed-forward generalizable 3D Gaussian splatting method **MVSGaussian** [35]; 3) the state-of-the-art point cloud rendering method **PFGS** [59]; 4) the state-of-the-art ray-based methods **Pointersect** [8] and **RayDF** [36] which is adapted as a generalizable version where the ray feature extractor is exactly the same as ours; 5) the latest image-based depth estimation methods **DepthAnythingV2** [69] and **DepthPro** [5].

**Metrics:** Primarily, we evaluate the surface reconstruction by measuring the error between predicted ray-surface distance and ground truth distance. We report the per ray-surface **absolute distance error (ADE)** in meters across all test images, and other four commonly used metrics [57] including **Root Mean Squared Error (RMSE)**, the **Absolute Relative Distance (Abs Rel)**, the **Squared Relative Distance (Sq Rel)**, and the **Threshold accuracy ( $\delta \leq 1.25$ )** which measures the percentage of predicted value within a certain threshold  $\delta$  of ground truth. In addition, we also evaluate the reconstructed 3D meshes, reporting **Ac-**

**curacy, Completion, Precision, Recall, Chamfer-L1 distance, Normal Consistency, and F-scores** with a threshold of 5cm. All details are in Appendix 5.1&5.2&5.3&5.4.

### 4.1. Evaluation on 3D Gaussians

In this setting, our method takes 3D Gaussians as input and predicts raylet distances for query rays. To prepare for this experiment, for every 3D scene in ScanNet/ScanNet++, ARKitScenes, and MultiScan, we first use the vanilla 3DGS to learn per-scene 3D Gaussians from RGB images in an offline fashion, and then evaluate our method in the following two groups of experiments.

- Group 1: We train our method from scratch on the joint train set of ScanNet/ScanNet++. After that, we evaluate the trained model on the test sets of ScanNet/ScanNet++, ARKitScenes, and MultiScan.
- Group 2: We train our method from scratch on the train set of ARKitScenes, and then test it on the test sets of ScanNet/ScanNet++, ARKitScenes, and MultiScan.

For a fair comparison, the four feed-forward generalizable baselines MVSGaussian [35], Pointersect [8], RayDF [36], and PFGS [59] are all trained with the same amount of training data and supervision signals as our method. For a reference, the per-scene methods 3DGS [30], GOF [74] and PGSR [9] directly render depth images from their optimized 3D Gaussians on test set respectively. For the image-based depth estimation methods, we first render test view RGBs from 3D Gaussians and then feed them into a depth estimator, obtaining depth results. Lastly, we use ground truth depth values to align the scale of estimated depths, which is strongly in favor of these baselines. More details of experiments are in Appendix 5.4.

**Results & Analysis:** Table 1 compares the quantitative results of all methods for estimating distance values of all query rays from test views. Table 2 compares the quantitative results of 3D meshes reconstructed by TSDF or Marching Tetrahedra (MT) from the rendered depth views for the

Table 2. Quantitative results of 3D meshes reconstructed from estimated ray surface distances. Our model is trained on Scannet/++ dataset.

		test on → ARKitScenes							test on → ScanNet/ScanNet++						
		Acc.↓	Comp.↓	Pre.↑	Rec.↑	C-L1↓	NC↑	F1↑	Acc.↓	Comp.↓	Pre.↑	Rec.↑	C-L1↓	NC↑	F1↑
3DGS [30]	TSDF	0.111	0.219	0.305	0.247	0.165	0.622	0.268	0.175	0.160	0.274	0.290	0.168	0.650	0.270
GOF [74]	MT	0.135	0.575	0.335	0.294	0.355	0.571	0.314	0.165	0.173	0.372	0.385	0.169	0.632	0.368
PGSR [9]	TSDF	0.153	0.188	0.437	0.344	0.171	<b>0.691</b>	0.378	<b>0.113</b>	0.088	<b>0.527</b>	0.543	<b>0.100</b>	<b>0.746</b>	0.532
<b>RayletDF (Ours)</b>	TSDF	<b>0.051</b>	<b>0.083</b>	<b>0.671</b>	<b>0.609</b>	<b>0.067</b>	0.689	<b>0.633</b>	0.151	<b>0.066</b>	<b>0.527</b>	<b>0.633</b>	0.109	0.705	<b>0.566</b>

Table 3. Quantitative results of estimated ray-surface distances from 3D point clouds.

		test on → ARKitScenes					test on → ScanNet/ScanNet++					test on → MultiScan				
		ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta$ ↑	ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta$ ↑	ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta$ ↑
Pointersect [8]	train on ARKitScenes	0.335	0.456	0.198	0.132	0.643	0.389	0.553	0.177	0.148	0.680	0.254	0.330	0.201	0.093	0.637
RayDF [36]		0.166	0.281	0.112	0.074	0.876	0.186	0.321	0.094	0.055	0.883	0.154	0.234	0.134	0.062	0.822
<b>RayletDF (Ours)</b>		<b>0.088</b>	<b>0.205</b>	<b>0.066</b>	<b>0.060</b>	<b>0.949</b>	<b>0.107</b>	<b>0.269</b>	<b>0.051</b>	<b>0.049</b>	<b>0.943</b>	<b>0.067</b>	<b>0.149</b>	<b>0.064</b>	<b>0.037</b>	<b>0.943</b>
Pointersect [8]	train on ScanNet/++	0.299	0.398	0.209	0.141	0.691	0.344	0.477	0.181	0.133	0.708	0.233	0.289	0.225	0.103	0.656
RayDF [36]		0.289	0.407	0.189	0.118	0.699	0.161	0.291	0.086	0.049	0.893	0.349	0.429	0.313	0.259	0.561
<b>RayletDF (Ours)</b>		<b>0.096</b>	<b>0.217</b>	<b>0.074</b>	<b>0.077</b>	<b>0.946</b>	<b>0.093</b>	<b>0.234</b>	<b>0.047</b>	<b>0.040</b>	<b>0.949</b>	<b>0.130</b>	<b>0.229</b>	<b>0.131</b>	<b>0.216</b>	<b>0.914</b>

three excellent per-scene Gaussian based methods and our method. Note that, in the original paper, GOF uses MT for higher quality surface reconstruction. Figure 5 shows qualitative results. From the results, we can see that:

- When training/testing on ARKitScenes, ScanNet/ ScanNet++ datasets in domain, our method achieves the best accuracy, outperforming the second best method RayDF by large margins, with 0.115 *vs* 0.183 meters on ARKitScenes and 0.145 *vs* 0.202 meters on ScanNet/ScanNet++ on the key metric ADE. This means that our newly introduced raylet distance field has its clear advantage over existing ray-based representations.
- When evaluating all methods across new datasets, our method demonstrates superior generalizability on unseen datasets, clearly surpassing all other baselines usually by more than 0.10 meters in accuracy. This performance gap is much greater than the in-domain reconstruction results. Essentially, this is because our learned raylet distance representations capture the local surface geometric patterns which tend to be generalizable at various scenes.
- The per-scene optimization based 3D Gaussian methods and the depth estimation methods can achieve satisfactory results, but still being inferior than our method.

## 4.2. Evaluation on Point Clouds

In this setting, our method takes 3D point clouds as input and predicts raylet distances for query rays. To prepare for this experiment, for every 3D scene in ScanNet/ScanNet++, ARKitScenes, and MultiScan, we uniformly sample 10k points from the provided 3D scene mesh. Same as Section 4.1, we also conduct two groups of experiments. For a fair comparison, we also evaluate the top-performing ray-based methods Pointersect and RayDF with exactly the same training data and supervision signals as ours.

**Results & Analysis:** Table 3 compares the quantitative results and Figure 6 shows the qualitative results. We can see that, similar to the analysis in Section 4.1, our method demonstrates clearly better reconstruction accuracy on the in-domain datasets. Most notably, we achieve an outstanding generalization ability when evaluating on unseen datasets, clearly surpassing two strong baselines. For example, when trained on ARKitScenes or ScanNet/ScanNet++, our method achieves {0.067, 0.130} meters in ADE on the novel MultiScan dataset respectively, while the baselines often have more than 0.20 meters errors. Basically, our strong generalization ability on point clouds is also attributed to the effectiveness of our virtual ball based raylet sampling strategy designed in Section 3.5.

## 4.3. Evaluation on Raylet Sampling in Testing

Regarding our designed multi-raylet blender in Section 3.4, the single hyperparameter  $T$  of this module can be different in training and test phase, allowing flexibility during evaluation on various datasets. Intuitively, once our model is well-trained on a dataset, in testing, the more number of raylet samples along every single ray, we may obtain more accurate surface estimations.

We further conduct experiments to validate the flexibility of raylet sampling in test phase. In particular, we use our model well-trained on ScanNet/ScanNet++ where  $T$  is chosen as 5 during training, and then directly evaluate on the test splits of ScanNet/ScanNet++, ARKitScenes, and the whole MultiScan dataset, where  $T$  is chosen as {1, 5, 10, 20} respectively.

Tables 4&5&6 show the quantitative results. It can be seen that, the more raylets sampled along every single ray during testing, we generally obtain more accurate surface reconstruction. This is particularly obvious when

the trained model is evaluated on unseen datasets as shown in Tables 5&6. Primarily, this is because if more raylets are sampled along every single light ray, our multi-raylet blender tends to vote a more accurate surface point, instead of being affected by outlier predictions among all estimations. This validates the generalizability and robustness of our simple design. More results are in Appendix 5.10.

Table 4. Qualitative results for different number of raylet samples per ray when evaluating on the test split of ScanNet/ScanNet++.

(train samples: $T = 5$ )	ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta \uparrow$
test samples: $T = 1$	0.177	0.337	0.087	0.190	0.896
test samples: $T = 5$	0.145	0.276	0.072	0.079	0.922
test samples: $T = 10$	<b>0.140</b>	<b>0.268</b>	<b>0.070</b>	<b>0.070</b>	<b>0.927</b>
test samples: $T = 20$	0.141	0.271	0.071	0.073	0.925

Table 5. Qualitative results for different number of raylet samples per ray when evaluating on the test split of ARKitScenes.

(train samples: $T = 5$ )	ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta \uparrow$
test samples: $T = 1$	0.161	0.287	0.113	0.076	0.876
test samples: $T = 5$	0.132	0.241	0.094	0.053	0.906
test samples: $T = 10$	0.127	0.234	0.092	<b>0.052</b>	0.912
test samples: $T = 20$	<b>0.124</b>	<b>0.230</b>	<b>0.090</b>	<b>0.052</b>	<b>0.917</b>

Table 6. Qualitative results for different number of raylet samples per ray when evaluating on the whole MultiScan dataset.

(train samples: $T = 5$ )	ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta \uparrow$
test samples: $T = 1$	0.291	0.393	0.261	0.177	0.619
test samples: $T = 5$	0.259	0.353	0.236	<b>0.148</b>	0.662
test samples: $T = 10$	0.252	<b>0.346</b>	<b>0.234</b>	<b>0.148</b>	0.674
test samples: $T = 20$	<b>0.247</b>	0.347	<b>0.234</b>	0.153	<b>0.687</b>

#### 4.4. Ablations

Our pipeline has three major modules, 1) Raylet Feature Extractor, 2) Raylet Distance Field, and 3) Multi-raylet Blender. To evaluate the effectiveness of each module and the sensitivity of hyperparameters, we conduct the following ablations on the merged ScanNet/ScanNet++ dataset, and the input to our method is chosen as 3D Gaussians.

**(1)~(3) Including different information of neighboring points.** Regarding the raylet feature extractor, in Equation 2, for each neighboring point, we choose to include: (1) only the  $xyz$  (i.e.,  $\mathbf{p}_i^k$ ) of a neighboring point, (2) only the relative position information (i.e.,  $(\frac{\mathbf{p}_i^k - \mathbf{p}_l}{\|\mathbf{p}_i^k - \mathbf{p}_l\|} \oplus \|\mathbf{p}_i^k - \mathbf{p}_l\|)$ ), (3) both  $xyz$  and the relative position information.

**(4)~(7) Choosing different number of neighboring points  $K$ .** For the raylet feature extractor, in Equation 3, the hyperparameter  $K$  is chosen as  $\{1, 5, 10, 20\}$ . This aims to evaluate the influence of local neighborhood on raylet features. In all main experiments, we choose  $K$  as 5.

**(8)~(11) Sampling different number of raylets.** For the multi-raylet blender, along each light ray, a total of  $T$  raylets are sampled near the surface. In this ablation, the hyperparameter  $T$  is chosen as  $\{1, 5, 10, 20\}$ . In all main experiments, we choose  $T$  as 5.

**(12) Removing the prediction of confidence score  $s_l$ .** Regarding the raylet distance field in Equation 4, we remove the confidence score, and the subsequent multi-raylet blender turns to simply average out multiple raylet distance predictions, i.e.,  $\hat{s}_l$  is  $1/T$  in Equation 5.

**(13) Using alpha blending to aggregate multi-raylet predictions.** For the multi-raylet blender, we choose to regard the confidence score  $s_l$  in Equation 4 as opacity, and then use alpha blending to get the final prediction.

**(14) Using sigmoid to normalize the confidence score  $s_l$ .** For the multi-raylet blender in Equation 5, we turn to use sigmoid function to independently normalize confidence scores, meaning that multiple raylets along a single ray will be trained independently.

Table 7 shows the ablation results. We can see that: 1) The greatest impact is caused by the decrease of raylet sample  $T$  to be just 1 in training. Basically, this means that our raylet distance field degenerates to the existing ray-based method such as RayDF, which is ineffective to learn local surface patterns. 2) The second greatest impact is caused by the removal of confidence score  $s_l$  in Equation 4. Basically, this is closely related to our design of multi-raylet blender as ablated in (12)~(14). Without a suitable multi-raylet blending strategy like ours, the accuracy of estimated surface drops. 3) Different number of neighboring points  $K$  and choices of their features included could be helpful as shown in (1)~(7), but they are less important overall.

Table 7. Results of all ablated models on ScanNet/ScanNet++. The main settings are bolded.

	ADE↓	RMSE↓	Abs-Rel↓	Sq-Rel↓	$\delta \uparrow$
(1) Only $xyz$	0.148	0.278	0.073	0.074	0.921
(2) Only $(\frac{\mathbf{p}_i^k - \mathbf{p}_l}{\ \mathbf{p}_i^k - \mathbf{p}_l\ } \oplus \ \mathbf{p}_i^k - \mathbf{p}_l\ )$	0.146	0.275	0.072	0.078	0.921
<b>(3) Both</b>	<b>0.145</b>	<b>0.276</b>	<b>0.072</b>	<b>0.079</b>	<b>0.922</b>
(4) $K = 1$	0.151	0.279	0.077	0.077	0.919
<b>(5) <math>K = 5</math></b>	<b>0.145</b>	<b>0.276</b>	<b>0.072</b>	<b>0.079</b>	<b>0.922</b>
(6) $K = 10$	0.150	0.280	0.074	0.074	0.920
(7) $K = 20$	0.152	0.285	0.075	0.073	0.915
(8) $T = 1$	0.174	0.326	0.085	0.144	0.897
<b>(9) <math>T = 5</math></b>	<b>0.145</b>	<b>0.276</b>	<b>0.072</b>	<b>0.079</b>	<b>0.922</b>
(10) $T = 10$	0.141	0.274	0.070	0.066	0.925
(11) $T = 20$	0.141	0.271	0.070	0.070	0.924
(12) Removing score $s_l$	0.170	0.302	0.084	0.097	0.907
(13) Using alpha blending	0.156	0.290	0.075	0.073	0.914
(14) Using sigmoid	0.164	0.287	0.075	0.078	0.915
<b>(15) The Full Pipeline (RayletDF)</b>	<b>0.145</b>	<b>0.276</b>	<b>0.072</b>	<b>0.079</b>	<b>0.922</b>

#### 5. Conclusion

In this paper, we present a generalizable 3D surface reconstruction method. By introducing a new technique of raylet distance field, our pipeline accurately captures intricate local surface patterns from both point clouds and 3D Gaussians, demonstrating superior performance across diverse real-world 3D scene datasets. Remarkably, thanks to the learned local raylet features, it exhibits excellent generalizability to new and unseen scenes in testing, while all baselines fail to do so.




Figure 5. Qualitative results of our method and baselines for 3D surface reconstruction on multiple datasets. All methods are trained on pre-estimated 3D Gaussians.




Figure 6. Qualitative results of our method and baselines for 3D surface reconstruction on multiple datasets. All methods are trained on point clouds.

**Acknowledgment:** This work was supported in part by Research Grants Council of Hong Kong under Grants 15225522 & 25207822 & 15219125, in part by Otto Poon Charitable Foundation Smart Cities Research Institute (8-CDCQ), in part by Research Centre for Unmanned Autonomous Systems (1-CE3D), The Hong Kong Polytechnic University.

## References

- [1] Tristan Aumentado-Armstrong, Stavros Tsogkas, Sven Dickinson, and Allan Jepson. Representing 3D Shapes with Probabilistic Directed Distance Fields. *CVPR*, 2022. 1, 3
- [2] Gavin Barill, Neil G Dickson, Ryan Schmidt, David IW Levin, and Alec Jacobson. Fast winding numbers for soups and clouds. *ACM Transactions on Graphics (TOG)*, 2018. 2
- [3] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, and Elad Shulman. ARK-itScenes: A Diverse Real-World Dataset For 3D Indoor Scene Understanding Using Mobile RGB-D Data. *NeurIPS*, 2021. 5
- [4] M Berger, A Tagliasacchi, L M Seversky, P Alliez, J A Levine, A Sharf, and C T Silva. A Survey of Surface Reconstruction from Point Clouds. *Computer Graphics Forum*, 2017. 2
- [5] Aleksei Bochkovskii, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R. Richter, and Vladlen Koltun. Depth Pro: Sharp Monocular Metric Depth in Less Than a Second. *ICLR*, 2025. 3, 5
- [6] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, Jose Neira, Ian D. Reid, and John J. Leonard. Past, Present, and Future of Simultaneous Localization and Mapping: Towards the Robust-Perception Age. *TRO*, 2016. 2
- [7] Junli Cao, Huan Wang, Pavlo Chemerys, Vladislav Shakhrai, Ju Hu, Yun Fu, Denys Makoviichuk, Sergey Tulyakov, and Jian Ren. Real-Time Neural Light Field on Mobile Devices. *CVPR*, 2023. 2
- [8] Jen-Hao Rick Chang, Wei-Yu Chen, Anurag Ranjan, Kwang Moo Yi, and Oncel Tuzel. Pointersect: Neural Rendering with Cloud-Ray Intersection. *CVPR*, 2023. 1, 5, 6, 4, 7
- [9] 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. *TVCG*, 2024. 2, 5, 6
- [10] Hanyu Chen, Bailey Miller, and Ioannis Gkioulekas. 3d reconstruction with fast dipole sums. *ACM Transactions on Graphics (TOG)*, 2024. 2
- [11] Hanlin Chen, Fangyin Wei, Chen Li, Tianxin Huang, Yunsong Wang, and Gim Hee Lee. VCR-GauS: View Consistent Depth-Normal Regularizer for Gaussian Surface Reconstruction. *NeurIPS*, 2024. 2
- [12] Zhiqin Chen and Hao Zhang. Learning Implicit Fields for Generative Shape Modeling. *CVPR*, 2019. 2
- [13] Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural Unsigned Distance Fields for Implicit Function Learning. *NeurIPS*, 2020. 1, 2
- [14] Christopher B. Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3D-R2N2: A Unified Approach for Single and Multi-view 3D Object Reconstruction. *ECCV*, 2016. 2
- [15] Jaeyoung Chung, Jeongtaek Oh, and Kyoung Mu Lee. Depth-Regularized Optimization for 3D Gaussian Splatting in Few-Shot Images. *CVPRW*, 2024. 2
- [16] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. *CVPR*, 2017. 5
- [17] Haoqiang Fan, Hao Su, and Leonidas Guibas. A Point Set Generation Network for 3D Object Reconstruction from a Single Image. *CVPR*, 2017. 2
- [18] Lue Fan, Yuxue Yang, Minxing Li, Hongsheng Li, and Zhaoxiang Zhang. Trim 3D Gaussian Splatting for Accurate Geometry Representation. *arXiv:2406.07499*, 2024. 2
- [19] Brandon Yushan Feng and Amitabh Varshney. SIGNET: Efficient Neural Representation for Light Fields. *ICCV*, 2021. 2
- [20] Brandon Y Feng, Yinda Zhang, Danhang Tang, and Ruofei Du. PRIF: Primary Ray-based Implicit Function. *ECCV*, 2022. 1, 3
- [21] A Tewari O Fried, J Thies V Sitzmann, S Lombardi K Sunkavalli, and R Martin-brualla T Simon J Saragih. State of the Art on Neural Rendering. *Computer Graphics Forum*, 2020. 2
- [22] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3D Semantic Segmentation with Submanifold Sparse Convolutional Networks. *CVPR*, 2018. 3
- [23] Jeff P. Grossman and William J. Dally. Point Sample Rendering. *Rendering Techniques*, 1998. 2
- [24] Antoine Guédon and Vincent Lepetit. SuGaR: Surface-Aligned Gaussian Splatting for Efficient 3D Mesh Reconstruction and High-Quality Mesh Rendering. *CVPR*, 2024. 1, 2
- [25] Xian-feng Han, Hamid Laga, and Mohammed Bannamoun. Image-based 3D Object Reconstruction: State-of-the-Art and Trends in the Deep Learning Era. *TPAMI*, 2019. 2
- [26] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D Gaussian Splatting for Geometrically Accurate Radiance Fields. *SIGGRAPH*, 2024. 2
- [27] Han Huang, Yulun Wu, Chao Deng, Ge Gao, Ming Gu, and Yu-Shen Liu. FatesGS: Fast and Accurate Sparse-View Surface Reconstruction using Gaussian Splatting with Depth-Feature Consistency. *AAAI*, 2025. 2
- [28] Linyi Jin, Nilesh Kulkarni, and David Fouhey. 3DFIRES: Few Image 3D REconstruction for Scenes with Hidden Surface. *CVPR*, 2024. 3
- [29] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3D Mesh Renderer. *CVPR*, 2018. 2
- [30] Bernhard Kerbl, Université Côte, Georgios Kopanas, Université Côte, Thomas Leimkühler, Max-planck-institut Informatik, and G R Aug. 3D Gaussian Splatting for Real-Time

- Radiance Field Rendering. *SIGGRAPH*, 2023. 1, 2, 3, 4, 5, 6
- [31] Leonid Keselman and Martial Hebert. Approximate Differentiable Rendering with Algebraic Surfaces. *ECCV*, 2022. 4, 3
- [32] Nilesh Kulkarni, Justin Johnson, and David F Fouhey. Directed Ray Distance Functions for 3D Scene Reconstruction. *ECCV*, 2022. 1, 3
- [33] Jiahe Li, Jiawei Zhang, Xiao Bai, Jin Zheng, Xin Ning, Jun Zhou, and Lin Gu. DNGaussian: Optimizing Sparse-View 3D Gaussian Radiance Fields with Global-Local Depth Normalization. *CVPR*, 2024. 1, 2
- [34] Celong Liu, Zhong Li, Junsong Yuan, and Yi Xu. NeuLF: Efficient Novel View Synthesis with Neural 4D Light Field. *Eurographics Symposium on Rendering*, 2022. 2
- [35] Tianqi Liu, Guangcong Wang, Shoukang Hu, Liao Shen, Xinyi Ye, Yuhang Zang, Zhiguo Cao, Wei Li, and Ziwei Liu. MVSGaussian: Fast Generalizable Gaussian Splatting Reconstruction from Multi-View Stereo. *ECCV*, 2024. 5
- [36] Zhuoman Liu, Bo Yang, Yan Luximon, Ajay Kumar, and Jinxi Li. RayDF: Neural Ray-surface Distance Fields with Multi-view Consistency. *NeurIPS*, 2023. 1, 3, 5, 6, 4, 7
- [37] Yongsen Mao, Yiming Zhang, Hanxiao Jiang, Angel X Chang, and Manolis Savva. MultiScan: Scalable RGBD scanning for 3D environments with articulated objects. *NeurIPS*, 2022. 5
- [38] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy Networks: Learning 3D Reconstruction in Function Space. *CVPR*, 2019. 1, 2
- [39] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *ECCV*, 2020. 1, 2, 6
- [40] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 2022. 6
- [41] Julian Ost, Issam Laradji, Alejandro Newell, Yuval Bahat, and Felix Heide. Neural Point Light Fields. *CVPR*, 2022. 2
- [42] Wanli Ouyang, Guofeng Zhang, Yu Qiao, and Tong He. GigaGS: Scaling up Planar-Based 3D Gaussians for Large Scene Surface Reconstruction. *arXiv:2409.06685*, 2024. 2
- [43] Onur Ozyesil, Vladislav Voroninski, Ronen Basri, and Amit Singer. A Survey of Structure from Motion. *Acta Numerica*, 2017. 2
- [44] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. *CVPR*, 2019. 1, 2
- [45] Hanspeter Pfister, Matthias Zwicker, Jeroen van Baar, and Markus H. Gross. Surfels: surface elements as rendering primitives. *SIGGRAPH*, 2000. 2
- [46] Rene Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-Shot Cross-Dataset Transfer. *TPAMI*, 2022. 3
- [47] Edoardo Mello Rella, Ajad Chhatkuli, Ender Konukoglu, and Luc Van Gool. Neural Vector Fields for Surface Representation and Inference. *IJCV*, 2024. 2
- [48] Liu Ren, Hanspeter Pfister, and Matthias Zwicker. Object Space EWA Surface Splatting: A Hardware Accelerated Approach to High Quality Point Rendering. *Computer Graphics Forum*, 2003. 2
- [49] Vincent Sitzmann, Semon Rezkikov, William T Freeman, Joshua B. Tenenbaum, and Frédo Durand. Light Field Networks: Neural Scene Representations with Single-Evaluation Rendering. *NeurIPS*, 2021. 2
- [50] Gaochao Song, Chong Cheng, and Hao Wang. GVKF: Gaussian Voxel Kernel Functions for Highly Efficient Surface Reconstruction in Open Scenes. *NeurIPS*, 2024. 2
- [51] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Light Field Neural Rendering. *CVPR*, 2022. 2
- [52] Peder Bergebakken Sundt and Theoharis Theoharis. MARS: The Medial Atom Ray Field object representation. *Computers & Graphics*, 2023. 3
- [53] Peder Bergebakken Sundt and Theoharis Theoharis. Towards Multi-view Consistency in Neural Ray Fields using Parametric Medial Surfaces. *Computers & Graphics*, 2024. 3
- [54] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree Generating Networks: Efficient Convolutional Architectures for High-resolution 3D Outputs. *ICCV*, 2017. 2
- [55] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Yifan Wang, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, Tomas Simon, Christian Theobalt, Matthias Niessner, Jonathan T. Barron, Gordon Wetzstein, Michael Zollhoefer, and Vladislav Golyanik. Advances in Neural Rendering. *SIGGRAPH*, 2021. 2
- [56] Alex Trevithick and Bo Yang. GRF: Learning a General Radiance Field for 3D Representation and Rendering. *ICCV*, 2021. 2
- [57] Matias Turkulainen, Xuqian Ren, Iaroslav Melekhov, Otto Seiskari, Esa Rahtu, and Juho Kannala. DN-Splatter: Depth and Normal Priors for Gaussian Splatting and Meshing. *WACV*, 2025. 5
- [58] Bing Wang, Zhengdi Yu, Bo Yang, Jie Qin, Toby Breckon, Ling Shao, Niki Trigoni, and Andrew Markham. RangeUDF: Semantic Surface Reconstruction from 3D Point Clouds. *arXiv:2204.09138*, 2022. 2
- [59] Jiaxu Wang, Ziyi Zhang, Junhao He, and Renjing Xu. PFGS: High Fidelity Point Cloud Rendering via Feature Splatting. *ECCV*, 2024. 5
- [60] 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:2106.10689*, 2021. 6
- [61] Zhengren Wang. 3D Representation Methods: A Survey. *arXiv:2410.06475*, 2024. 2
- [62] Yaniv Wolf, Amit Bracha, and Ron Kimmel. GS2Mesh: Surface Reconstruction from Gaussian Splatting via Novel Stereo Views. *ECCV*, 2024. 1

- [63] Qianyi Wu, Jianmin Zheng, and Jianfei Cai. Surface Reconstruction from 3D Gaussian Splatting via Local Structural Hints. *ECCV*, 2024. 1, 2
- [64] Haodong Xiang, Xinghui Li, Xiansong Lai, Wanting Zhang, Zhichao Liao, Kai Cheng, and Xueping Liu. Gaussian-Room: Improving 3D Gaussian Splatting with SDF Guidance and Monocular Cues for Indoor Scene Reconstruction. *arXiv:2405.19671*, 2024. 2
- [65] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural Fields in Visual Computing and Beyond. *Computer Graphics Forum*, 2022. 2
- [66] Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. DepthSplat: Connecting Gaussian Splatting and Depth. *CVPR*, 2025. 2
- [67] Bo Yang, Stefano Rosa, Andrew Markham, Niki Trigoni, and Hongkai Wen. Dense 3D Object Reconstruction from a Single Depth View. *TPAMI*, 2019. 2
- [68] Bo Yang, Sen Wang, Andrew Markham, and Niki Trigoni. Robust Attentional Aggregation of Deep Feature Sets for Multi-view 3D Reconstruction. *IJCV*, 2020. 2
- [69] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth Anything V2. *NeurIPS*, 2024. 3, 5
- [70] Xianghui Yang, Guosheng Lin, Zhenghao Chen, and Luping Zhou. Neural Vector Fields: Implicit Representation by Explicit Learning. *CVPR*, 2023. 1, 2
- [71] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. ScanNet++: A High-Fidelity Dataset of 3D Indoor Scenes. *ICCV*, 2023. 5
- [72] Ruihong Yin, Yunlu Chen, Sezer Karaoglu, and Theo Gevers. Ray-Distance Volume Rendering for Neural Scene Reconstruction. *ECCV*, 2024. 3
- [73] Mulin Yu, Tao Lu, Linning Xu, Lihan Jiang, Yuanbo Xiangli, and Bo Dai. GSDF: 3DGS Meets SDF for Improved Rendering and Reconstruction. *NeurIPS*, 2024. 2
- [74] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian Opacity Fields: Efficient and Compact Surface Reconstruction in Unbounded Scenes. *SIGGRAPH Asia*, 2024. 4, 5, 6, 2
- [75] Baowen Zhang, Chuan Fang, Rakesh Shrestha, Yixun Liang, Xiaoxiao Long, and Ping Tan. RaDe-GS: Rasterizing Depth in Gaussian Splatting. *arXiv:2406.01467*, 2024.
- [76] Ziyu Zhang, Binbin Huang, Hanqing Jiang, Liyang Zhou, Xiaojun Xiang, and Shunhan Shen. Quadratic Gaussian Splatting for Efficient and Detailed Surface Reconstruction. *arXiv:2411.16392*, 2024. 2
- [77] Matthias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross. Surface splatting. *SIGGRAPH*, 2001. 2