HDSDF: Hybrid Directional and Signed Distance Functions for Fast Inverse Rendering

Tarun Yenamandra[1], Ayush Tewari[2,3], Nan Yang[1,5], Florian Bernard[4], Christian Theobalt[2], and Daniel Cremers[1]

[1]TU Munich, [2]MPI for Informatics, SIC [3]MIT [4]University of Bonn [5]Artisense

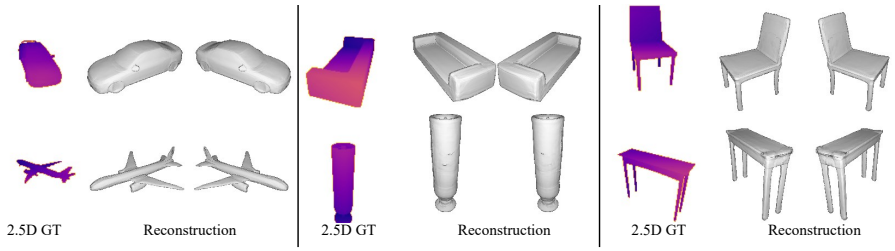

Fig. 1. We propose the first SDF-DDF hybrid model that allows for 3D shape optimization via rendering with a single evaluation. In each of the 3 parts of the figure, we show the given ground truth depth map (2.5D GT) on left, and two views of our reconstructed 3D shape on the right

Abstract. Implicit neural representations of 3D shapes form strong priors that are useful for various applications, such as single and multiple view 3D reconstruction. A downside of existing neural representations is that they require multiple network evaluations for rendering, which leads to high computational costs. This limitation forms a bottleneck particularly in the context of inverse problems, such as image-based 3D reconstruction. To address this issue, in this paper (i) we propose a novel hybrid 3D object representation based on a signed distance function (SDF) that we augment with a directional distance function (DDF), so that we can predict distances to the object surface from any point on a sphere enclosing the object. Moreover, (ii) using the proposed hybrid representation we address the multi-view consistency problem common in existing DDF representations. We evaluate our novel hybrid representation on the task of single-view depth reconstruction and show that our method is several times faster compared to competing methods, while at the same time achieving better reconstruction accuracy.

Keywords: Sphere tracing, neural implicit representations, directional distance function, signed distance function

1 Introduction

There are various alternatives for representing 3D shapes, among them being explicit and implicit representations. Explicit representations readily give us information about the surface of the represented 3D shape, e.g. triangular meshes, point clouds, and splines. In contrast, implicit 3D shape representations encode the surface of a shape as the level-set of a function, i.e. all points that have a specific function value are considered to form the surface of the shape. Recently, with the progress of deep learning, the use of neural networks to approximate functions for representing a class of shapes in an implicit manner have been widely adopted [27,22,7]. These representations have proven to be versatile in representing shapes of varied topology, and it has been demonstrated that respective parameterized neural representations generalize well to unseen objects.

Despite their great success, existing implicit neural representations are slow to render and optimize for inverse problems. The most common algorithms to render an implicit representation are ray marching and sphere tracing [13], where the latter is specifically designed for rendering SDFs. However, to utilise these algorithms in the context of neural implicit representations, it is necessary to perform multiple evaluations of the network per ray to obtain the distance to the surface along the ray. This is particularly prominent for objects with high geometric complexity, since in this case it would take many steps to march along the ray due to multiple local optima. The computation and time intensive step of rendering an SDF is aggravated when we use a neural SDF representation to solve an inverse problem [53,20], such as 3D reconstruction from images using optimization. Here, we need to render the neural SDF at each iteration, since we update the network parameters (or a latent code vector that represents a shape) during each optimization step.

To compensate for this shortcoming of neural implicit representations, in this work we introduce a novel hybrid 3D representation.

Specifically, we augment a neural signed distance function (SDF) representation with a neural directional distance function (DDF) that is defined on a unit sphere enclosing the 3D shape (see Fig. 2). Our main motivation for incorporating the DDF representation is to obviate the need for computationally expensive sphere tracing when solving inverse problems. By sphere-tracing an SDF, we get the distance to the surface along a ray – in contrast, this information is already explicitly encoded in a DDF representation. Thus, while sphere-tracing neural SDFs require multiple evaluations of the network along the ray, the DDF representation has the strong advantage that it just needs a single network evaluation per ray.

However, a downside of traditional DDF representations is that they suffer from multi-view inconsistencies. This is because a point on the surface of the object can be reached from multiple points (in different directions) on the unit sphere, see Fig. 2. As the DDF is represented by a general neural network function, we cannot guarantee that the rays will reach the same point on the surface with the predicted distance. In this work we present a solution to this problem by exploiting the fusion of neural SDF with DDF representations in terms of a

novel neural network architecture. Our main goal is to learn a neural network-

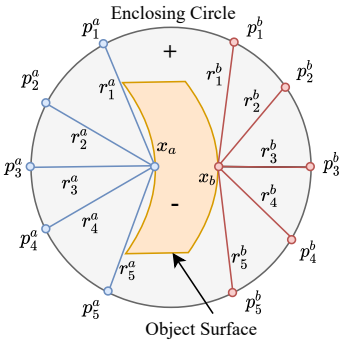

SDF: In the region in grey (+) inside the enclosing unit circle the SDF evaluates to a positive value, whereas in the region in orange (−) the SDF evaluates to a negative value. **DDF multi-view inconsistency:** The points p_i^a in blue represent a few possible points from which the point on the surface x_a is visible in the respective direction r_i^a.The points p_i^b in red represent a few possible points from which the point on the surface x_b is visible in the respective direction r_i^b. The number of possible points on the sphere and directions for the points x_a and x_b may be uncountable.

Fig. 2. A 2D illustration of a signed distance function and the directional distance function multi-view inconsistency.

based DDF model for representing a class of shapes, for which we consider an autodecoder framework. Overall, we demonstrate a substantial improvement in terms of runtime compared to the previous state of the art, as we do not need sphere tracing during the optimization process. Furthermore, as the DDF and SDF form strong priors, our approach leads to an improved accuracy compared to existing methods.

We note the following as the core contributions of our work.

1. A novel hybrid neural-3D representation that simultaneously models the signed distance function and directional distance function in order to represent a class of shapes.
2. A novel network architecture that predicts SDF when not conditioned on directions and predicts DDF when conditioned on directions.
3. Surface prediction with the DDF model with just one network evaluation per ray using the proposed hybrid representation for accelerating inverse rendering problems.

2 Related Work

In this section we present the works from topics that are most relevant to us: implicit representations, accelerating implicit representations, and 3D reconstruction methods.

Implicit Representations. Implicit surface representations, e.g., signed distance functions (SDF), have been studied for decades due to the flexibility to represent arbitrary topologies of 3D shapes [19,10,18,42,33,24,41,16]. However, to represent real-world objects, analytical expressions rarely exist, so that in practice the SDF values are usually discretized and stored in a voxelized space, which makes it non-trivial to represent detailed shapes. Recently, deep implicit

representations are proposed to overcome the aforementioned limitations. Park *et al.* [27] proposed DeepSDF that learns the SDF values using an autodecoder architecture in which latent codes of different shape priors can be encoded and optimized. The concurrent works of Mescheder *et al.* [22] and Chen *et al.* [7] learn object surfaces as decision boundaries between the interior and exterior of shapes. Instead of encoding only global priors, local priors [5,15,8,30] have been explored to handle large-scale scenes and more detailed representations. To relax the requirements of the availability of 3D ground-truth labels, differentiable formulations have been proposed that allow to learn from only 2D images or depth maps [26,40,20]. Pixel features [56,35,36] have been used to condition implicit representations on local image features leading to high-quality multiview RGB reconstructions from a single image. Such representations also allow to render multiple views of an object given a single view RGB image. However, they can also not model the geometry of the object satisfactorily. Applications on human [36,35,46], face [55], and hair [34,55] modeling are also explored and have achieved superior results to classical methods.

Accelerating Neural Implicit Representation Rendering. Recent works [43,2,32] have focused on rendering an image from a neural network representing a single shape in real time. However, representing a class of shapes with these kinds of networks is an open problem. There have been two recent work which predict the the occupancy density distribution along the rays [31], or, alternatively, a region along the ray instead of distance [25]. They can however only model single objects and they still need to do local sampling as they work in a volumetric setting. We note PDDF [27] and SDDF [57] as our contemporary works. However, they suffer from multi-view inconsistencies as they represent shapes only with a directional distance function. In contrast, we address multi-view inconsistencies via our hybrid SDF-DDF representation.

3D Reconstruction. Reconstructing 3D shapes from images is a fundamental computer vision task. Depending on the number of input images, one can categorize the 3D reconstruction methods into single-view based methods and multi-view based methods. Although single-view reconstruction is generally an ill-posed problem, a lot of advances have been made due to its practicality and simplicity of inference. To constrain the problem, one can introduce explicit priors as statistical 3D models of certain classes [3,21,28] and optimize using such a model as prior [45,4,48]. When reconstructing dynamic objects from videos, temporal information and consistency constraints are considered to deliver better results [44]. To extend the methods to more general objects, lower-level geometric or photometric properties [12,47,50], as well as joint 3D reconstruction and correspondence matching [54] have also been explored. Multi-view methods leverage established 3D geometric constraints like epipolar constraints and can more reliably reconstruct the scene. Multi-view stereo estimates the dense depth map of the reference image using multiple images and their corresponding camera parameters [38,37,1]. Deep neural networks are also utilized to improve the reconstruction quality [14,52,49,17]. Recent novel neural representations encode scenes as neural radiance fields [23,29] and reconstruct dynamic and static ob-

jects using self-supervised learning. More recently, LFN [39] has been proposed to exploit light field principles to reconstruct multiple views given a single view. While it is possible to obtain sparse depth, the method does not model dense geometry.

Encoder-based single-view reconstruction methods [8,51,26] can obtain the latent code with a query to an encoder. However, they cannot work with arbitrary image sizes without architectural changes to the learned-encoder. An optimisation-based latent code prediction, such as the one that we propose in this work, on the other hand can not only work with any arbitrary image sizes, but also work with unstructured data such as point clouds, and open meshes.

3 Method

We learn a neural network-based novel hybrid representation which represents both the signed distance function and the directional distance function. The directional distance function represents the distance to the surface of the object from a point on the unit sphere along a given direction. This is particularly desirable as it helps avoid the computationally intensive sphere tracing step for rendering images, especially for solving image-based inverse problems such as 3D reconstruction. In the following, we first introduce our representation, followed by our network architecture, and our proposed algorithm based on the novel representation and architecture for single-view reconstruction from depth maps.

3.1 Proposed Hybrid Shape Representation

Our main objective is to learn a neural 3D representation that simultaneously contains information about two aspects: (i) a distance d to the surface of an object along a given direction r from a point p on the surface of the unit sphere, and (ii) the signed distance s at every point inside the unit sphere enclosing the object, as shown in Fig. 2.

The directional distance d and signed distance s are related as follows: Outside the surface of the object the signed distance is positive and the value is the minimum distance between a given point p and the object surface (in any given direction), i.e.

$$\mathrm{SDF}(p) = \min_r \mathrm{DDF}(p, r), \tag{1}$$

where $s = \mathrm{SDF}(p)$ is the signed distance at the point p, and $d = \mathrm{DDF}(p, r)$ is the directional distance to the surface of the object from point p in a direction r. In our work we learn to predict the DDF only on the surface of the sphere, and use the SDF to predict whether the ray hits the surface. In other words, the point $x = p + dr = p + \mathrm{DDF}(p, r) \cdot r$ is on the object surface if $s = \mathrm{SDF}(p) = 0$. In practice, we threshold s to predict if the ray hits. In the following, we elucidate how we approximate the distances for a class of objects using neural networks.

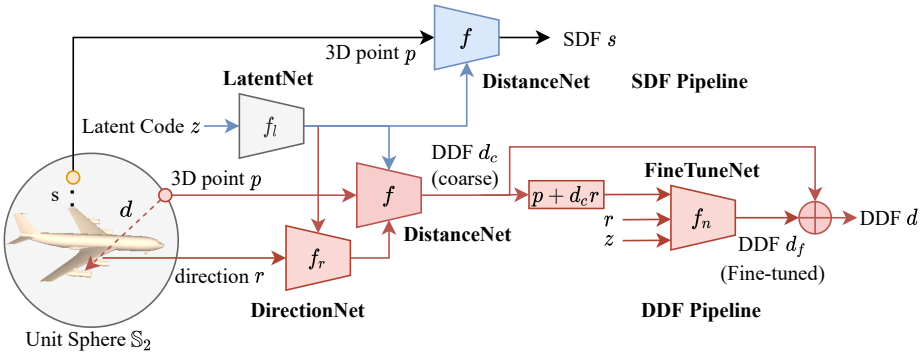

Fig. 3. Our novel network predicts SDF and DDF simultaneously. Our network structure for DDF prediction is shown in red and grey. The network takes as input a point p on the unit sphere, a direction r, and a latent code z. The LatentNet (grey) encodes the latent code z and the DirectionNet encodes the direction r, which are used to modulate the outputs of different layers [56] of the DistanceNet. The predicted distance from the DistanceNet is coarse, therefore we finetune it using the FineTuneNet. The latter has the task to predict the DDF at points close to the surface $p + d_c r$, where d_c is the predicted coarse distance. Finally, DDF $d = d_c + d_f$, where d_f is the fine-tuned DDF prediction near the surface. Our network structure for SDF prediction is shown in blue and grey. Given an encoded latent code $f_l(z)$, and a point inside the unit sphere p, the DistanceNet predicts signed distance s. Note that there is no DirectionNet conditioning during SDF prediction

3.2 Network

Motivated by the relationship between the SDF and DDF as illustrated in Eq. (1), we propose a neural network model that can predict both SDF and DDF, as shown in Fig. 3. Towards that, we use four main MLPs to predict DDF and SDF functions: the distance prediction network (*DistanceNet*) f, the latent conditioning network (*LatentNet*) f_l, the directional conditioning network (*DirectionNet*) f_r, and the DDF fine-tuning network f_n. LatentNet encodes the latent code and produces a conditioning for the DistanceNet to represent a class of objects. DirectionNet encodes directions given a LatentNet conditioning to condition the DistanceNet to predict DDF. DistanceNet is the main component of our network that predicts DDF given directional conditioning and predicts SDF without the directional conditioning. FineTuneNet predicts fine-tuned directional distances for accurate DDF prediction.

LatentNet The latent conditioning network $f_l : \mathbb{R}^{256} \rightarrow \mathbb{R}^{512}$ is an MLP that takes a latent code z and encodes the latent code to the size of the hidden layer dimensions of DistanceNet and DirectionNet, i.e.

$$z_e = f_l(z; \Theta_l)\,, \tag{2}$$

where z_e is the encoded latent code and Θ_l are the network parameters of the LatentNet.

DirectionNet The directional conditioning network $f_r : \mathbb{R}^3 \times \mathbb{R}^{512} \rightarrow \mathbb{R}^{512}$ conditions the DistanceNet f to predict directional distances on the surface of the unit sphere. Without the directional conditioning, the DistanceNet predicts signed distances everywhere in the unit sphere. The DirectionNet is an MLP with encoded latent z_e modulating the layers' outputs similar to pixelNeRF [56]. This conditioning helps the network predict different features for different objects based on a given a direction. The inputs to this network are encoded with a positional encoding [23]. Given a direction r, we obtain the encoded directional vector as

$$r_e = f_r(r, z_e; \Theta_r)\,, \tag{3}$$

where $r_e \in \mathbb{R}^{512}$ is the encoded directional vector for a given encoded latent code z_e in Eq. (2).

DistanceNet The distance prediction network f predicts both the signed distance $f : \mathbb{R}^3 \times \mathbb{R}^{512} \rightarrow \mathbb{R}$ and the directional distance $f : \mathbb{R}^3 \times \mathbb{R}^{512} \times \mathbb{R}^{512} \rightarrow \mathbb{R}$. When the network is conditioned with f_r, it predicts the directional distance on the unit sphere. Otherwise, i.e. without conditioning with f_r, it predicts the signed distances everywhere inside the unit sphere. As with the DirectionNet, we add the output of the LatentNet to a few output layers of the DistanceNet to condition the network to predict directional distances to the surface of different shapes.

SDF Prediction The signed distance at a point $p \in \mathbb{R}^3$ inside the unit sphere is given by

$$s = f(p, z_e; \Theta_D), \tag{4}$$

where z_e is the encoded latent code from Eq. (2), and $p \in \mathbb{R}^3$ is a point inside the unit sphere, and $s \in \mathbb{R}$ is the signed distance value. We truncate the predicted signed distances [27].

Coarse DDF Prediction The coarse directional distance $d_c \in \mathbb{R}_+$ (as DDF is always positive on outside the shape) at a point $p \in \mathbb{S}_2$ on the unit sphere in a direction $r \in \mathbb{S}_2$ is given by

$$d_c = f(p, r_e, z_e; \Theta_D), \tag{5}$$

where z_e is the encoded latent code from Eq. (2), r_e is the encoded direction code, and d_c is the coarse directional distance.

FineTuneNet The DDF fine-tuning network $f_n : \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^{256} \rightarrow \mathbb{R}$ fine-tunes the coarse DDF d_c predicted by the DistanceNet by predicting a local directional distance d_l from points close to the surface, i.e.

$$d_f = f_n(p + d_c r, r, z; \Theta_n)\,, \tag{6}$$

where $p + d_c r$ is a point close to the surface obtained using the coarse directional distance d_c, point p on the unit sphere and direction r, and z is the latent code. The final directional distance is given by

$$d = d_l + d_c. \tag{7}$$

3.3 Training

We train the model per object class of the ShapeNet dataset [6].

Data Preprocessing For training the network, we use two supervision signals, the signed distance supervision, and the directional distance supervision. We use the preprocessing pipeline from DeepSDF [27] to sample about 1 million points and signed distances inside and near the surface of each shape. Further, we randomly sample 1 million points on the unit sphere and random directions that point to the surface of an object. We sample directions with the help of sampled points and the shape's point cloud. We use the rendering function from Trimesh [11] to find the distance to the surface for each pair of the sampled points and directions.

Losses We train the network with the following losses:

- **SDF loss $\mathcal{L}_s$.** We utilize the SDF loss to train DistanceNet to predict signed distances. Note that the DistanceNet does not have directional conditioning for predicting signed distances in the unit sphere. Specifically, we use

$$\mathcal{L}_s(s) = \|s - s_{GT}\|_1 \,, \tag{8}$$

  where s is the predicted signed distance value from Eq. (4), and s_{GT} is the ground truth signed distance value at the point p.
- **Coarse DDF loss $\mathcal{L}_{dc}$.** We enforce the coarse DDF loss to predict the coarse DDF value at a point on the unit sphere as

$$\mathcal{L}_{dc}(d_c) = \|d_c - d_{GT}\|_1 \,, \tag{9}$$

  where d_{GT} is the ground truth distance value, and d_c is the coarse predicted distance from Eq. (5).
- **Fine-tuned DDF loss $\mathcal{L}_{df}$.** We enforce the fine-tuned DDF loss to train the FineTuneNet to predict DDF values near the surface as

$$\mathcal{L}_{df}(d_f, d_c) = \|d_c + d_f - d_{GT}\|_1 \,, \tag{10}$$

  where d_f is the local directional distance inside the unit sphere from Eq. (6). Note that we enforce both d_c and $d = d_c + d_f$ to predict the ground truth directional distance d_{GT} as a means to regularize that the FineTuneNet only needs to predict fine DDF values close to the object surface.

- **Latent code regularizer** $\mathcal{L}_l$. As we use an Autodecoder framework [27], we enforce that the latent codes for different shapes are close to each other, so that the model can generalize to new shapes. This can be achieved by penalizing latent codes with large norms, so that latent codes are close to zero, i.e.

$$\mathcal{L}_l(z) = \|z\|_2 \,, \tag{11}$$

  where z is the latent code for a given shape. Note that we use the same latent code for both SDF and DDF data paths for a given shape.
- **Total training loss**. The complete training loss is given as

$$\mathcal{L} = w_s\mathcal{L}_s + w_{dc}\mathcal{L}_{dc} + w_{df}\mathcal{L}_{df} + w_l\mathcal{L}_l \,, \tag{12}$$

  where w_s, w_{dc}, w_{df}, and w_l are the weights for the SDF loss, coarse DDF loss, fine-tuned DDF loss, and latent code regularizer, respectively.

Optimization We optimize the loss in Eq. (12) for the network parameters and the optimization parameters. We optimize for the loss in Eq. (12) to obtain the parameters Θ_D, Θ_r, Θ_l, Θ_n, and Z, where $Z = \{z_i | i = 1 \dots N\}$ is the set of latent codes representing all the N training shapes, Θ_D are the learnable network parameters of the DistanceNet f, Θ_r are the learnable network parameters of the DirectionNet f_r, Θ_l are the learnable network parameters of the LatentNet f_z, and Θ_n are the learnable network parameters of the FineTuneNet f_n.

3.4 3D Reconstruction from Single View Depth Maps

Our autodecoder framework allows us to work with any type of data without having to learn a new encoder for each type of data. Hence, during test time we merely need to optimize for the latent code z. One of the main advantages of our representation is that, unlike DIST [20], we obviate the need for sphere tracing at every iteration of the optimization.

For 3D reconstruction we assume a depth map with an object mask and a given camera pose as input. We obtain the points of intersection of the rays r from the camera with the unit sphere as p. For some initial latent code z, we evaluate our DDF pipeline for the directional distance from p along r using Eq. (7). We evaluate the SDF pipeline for the SDF at the points $p + dr$ as s.

We optimize for the latent code z_r of the object in the given depth map using the loss function

$$\mathcal{L}_{rec} = w_m\mathcal{L}_S + w_D\mathcal{L}_D + w_l\mathcal{L}_l + w_{DDF}\mathcal{L}_{DDF} \,, \tag{13}$$

where $\mathcal{L}_S$ is the silhouette loss, $\mathcal{L}_D$ is the depth loss, $\mathcal{L}_{latent}$ is the regularizer for learning the latent code, and $\mathcal{L}_{DDF}$ is the regularizer for DDF-SDF consistency, with w_m, w_D, w_l, and w_{DDF} as their respective weights. In the following we explain the individual loss terms:

– **Silhouette Loss** $\mathcal{L}_S$. The silhouette loss is enforced as

$$\mathcal{L}_S(s) = \|s_{m_{=1}}\|_1 + \||s_{m_{=0}}| - \tau\|_1 \,, \tag{14}$$

where m is the given mask, s is the predicted signed distance from Eq. (4), and τ is the truncation distance. $m_{=0}$ is an indicator function where mask is 0, similarly $m_{=1}$ where mask is 1. The first term in the loss enforces that along the rays that must hit the surface, the signed distance must be close to 0, and the second term enforces that the signed distance value must be close to the truncation distance where the rays must not hit the surface.

– **Depth Loss** $\mathcal{L}_D$. The depth loss is simply the error between the given depth and the predicted depth, i.e.

$$\mathcal{L}_D(\delta) = \|\delta - \delta_{GT}\|_1 \,, \tag{15}$$

where δ is the predicted depth and δ_{GT} is the given depth.

– **Latent Regularizer** $\mathcal{L}_l$. We enforce the same loss as in Eq. (11) for learning the latent code during inference.

– **Regularizer for DDF-SDF consistency**. As the predicted directional distance and the signed distance for an object need not agree, we enforce a DDF-SDF consistency regularizer. Towards that, we randomly sample points on the unit sphere and on the currently predicted shape. We enforce that the signed distance of the points (using $p + dr$) predicted by the directional distance must be close to 0 by using

$$\mathcal{L}_{DDF}(s) = \|s\|_1 \,. \tag{16}$$

4 Experiments

In this section we first provide the training details for our model followed by experiments that evaluate our design choices and our method.

4.1 Training Details

We train a model for each class of the ShapeNet dataset with the splits from [26]. We use a batch size of 64 and 16384 samples per scene, train for 3000 iterations for classes with less than 3000 shapes, and 2000 otherwise. Each batch takes about 3s on two Nvidia A100 GPUs. Depending on the number of shapes in a class it takes 1-7 days to train a per-class model.

4.2 Inference Setup

We use an Nvidia RTX8000 GPUs for all the inference tasks based-on optimization both for ours and DIST [20] to ensure that the optimization performance we report are consistent.

4.3 3D Reconstruction from Single View Depth Maps

Our method predicts distance to the surface of the shape given a latent code representing the shape, the ray origin and the ray direction. This makes our method well-suited for solving inverse 3D reconstructions problems from images. In this experiment, we evaluate our method on reconstructing 3D shape given a depth map and camera pose.

Given a depth image, and camera parameters, we render a depth image from our network. We optimize for the latent code to minimize the distance between the predicted and given depth images as discussed in Sec 3.4. The main advantage of using our method compared to other methods, such as DIST [20], is that others need to perform *multiple* forward passes to sphere trace. Specifically, for finding the object surface our network only requires a single evaluation, whereas others may require multiple passes due to their involved line search.

We test our trained models on different classes of shapes – airplanes, cars, chairs, lamps, sofas, and tables. We test our method on the first 200 instances of the test splits of each class (except for the cars class we use the first 140 instances). We obtain the first image of the rendered ShapeNet dataset from 3D-R2N2 [9]. The images are of resolution 137×137, and we re-render the images with the same camera matrices to obtain the depth maps and use these depth maps with camera poses to reconstruct the 3D shape. We use the official shape completion code from IF-Net [8] to complete the sparse point clouds. As such, we obtain the sparse point clouds for completion using IF-Net by un-projecting the depth maps using the given camera poses.

Qualitative Results We show qualitative results in Fig. 4. It can clearly be seen that our method can complete the shape given a sparse depth of about 3k pixels while resulting in the most plausible 3D shapes. Whilst DIST [20] is close in terms of the overall quality, our method is about $8.6\times$ faster per iteration (see Tab. 1). Further, given that we have a prior on the distances, the DDF, we can better reconstruct finer details such as the airplane's tail in 6th column. We also compare with IF-NET [8], a state-of-the-art encoder-based neural implicit representation. While IF-NET leads to plausible reconstructions in the observed locations, where there are depth maps, it does not complete the unobserved shapes, as shown in the last two rows of Fig. 4.

Quantitative Results We extract meshes at 256^3 resolution to evaluate our experiments quantitatively. We use the symmetric L2 Chamfers distance as the metric for our quantitative experiments, i.e.

$$\text{CD} = \frac{1}{N}\Sigma_{i=1}^{N} \min_{y \in Y} \|x_i - y\|_2^2 + \frac{1}{N}\Sigma_{i=1}^{N} \min_{x \in X} \|x - y_i\|_2^2 \,, \tag{17}$$

where $X = \{x_i \in \mathbb{R}^3 | i = 1...N\}$ and $Y = \{y_i \in \mathbb{R}^3 | i = 1...N\}$ are points on the surfaces of two shapes, and N is the number of points sampled on the two shapes. We show the quantitative results in Tab. 1. The quantitative results

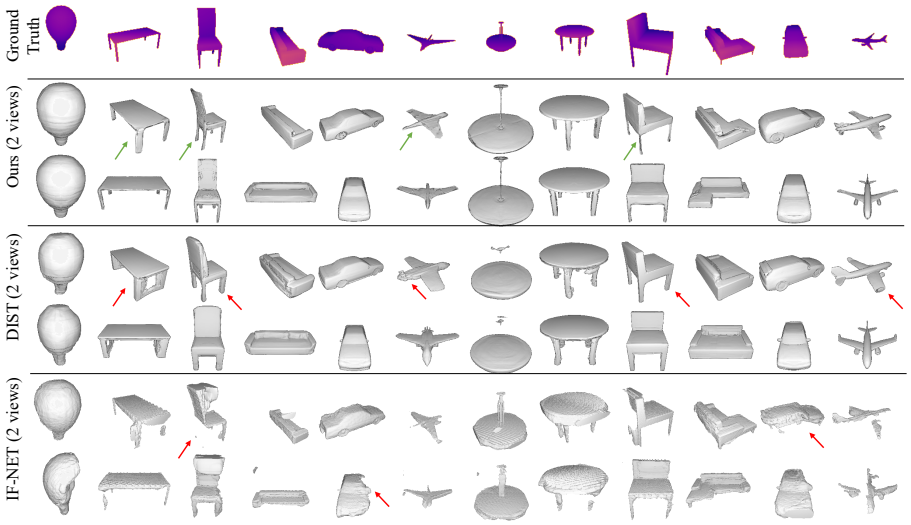

Fig. 4. 3D shapes reconstructed from a given depth map. Each column shows reconstruction results from the methods in our evaluation given the same depth image. Top row: Given ground truth depth map. Second and third rows: 3D shape reconstructed with our method shown in two views. Fourth and fifth rows: 3D shape reconstructed using DIST [20]. Last two rows: 3D shape reconstructed using IF-NET [8], a state-of-the-art encoder-based implicit representation network, using point clouds obtained by unprojecting depth maps using the given camera matrices. Our method outperforms existing methods, as it can for example better model fine-scale details (see e.g. the legs of the tables or chairs, or the geometry of the airplanes)

are consistent with the qualitative results, as our method can fit better to the given depth maps and thus obtain more plausible shapes compared to DIST [20]. Moreover, we outperform DIST by a small margin in most cases, while being 8.6× faster. Further, as IF-NET [8] does not complete the shape in unobserved areas, we significantly outperform IF-NET both quantitatively and qualitatively.

For computing the metrics, we use DeepSDF's [27] surface sampling algorithm to sample 30000 points on the visible surface of the object. We compute the symmetric Chamfer's distance, see Eq. (17), in ground truth scale of ShapeNetV2 [6].

4.4 Ablations

In this experiment we evaluate the design choices of our final model. We start with a baseline method which is a vanilla DeepSDF [27] for predicting the SDF and a similar network with larger capacity for DDF prediction. Next, we refine the DDF prediction in the baseline using FineTuneNet (see Sec. 3.2). Afterwards, we evaluate the our model without FineTuneNet. Finally, we evaluate our model with all the components.

Table 1. Quantitative comparisons of our method with DIST [20] and IF-NET [8]. Our method outperforms DIST in almost all the classes, showing that our DDF prior and our depth-fitting algorithm lead to more plausible 3D shape reconstructions. We also outperform the encoder-based implicit reconstruction method IF-NET by a large margin as IF-NET does not lead to plausible shapes in unobserved areas. Note that we report 1000× the L2 Chamfer distance (CD) (see Eq. (17)). We further compare the time it takes to perform one step in the optimization with DIST, as shown in the last two columns. As we do not need to perform sphere tracing to obtain an image each iteration, our approach is about 8.6× faster than DIST

Method	**Ours**	**DIST** [20]	**IF-NET** [8]	**Ours**	**DIST** [20]
Metric	CD ↓			seconds/iteration ↓	
Plane	**0.68**	0.94	2.08	**0.017**	0.227
Sofa	**1.42**	1.81	9.43	**0.0368**	0.275
Table	3.47	**2.79**	4.67	**0.0348**	0.263
Lamp	**5.55**	7.34	6.05	**0.0227**	0.247
Chair	**1.84**	1.92	5.45	**0.0335**	0.257
Car	**0.55**	0.613	4.09	**0.0359**	0.246

For the ablations experiments, we train the models with the first 256 shapes from the training split of the sofas class of ShapeNet dataset. We test our model on the first 64 test shapes from the test split. We use the first image from the renders of 3D-R2N2 [9] along with the camera matrices to obtain a depth map. We optimize for the latent code of our model representing a 3D shape, from the depth map as explained in Sec. 3.4.

Table 2. Ablations. We evaluate the design choices of our model by training a model on the sofas class with 256 shapes and testing on a set of 64 shapes. We report 1000× the L2 Chamfer distance (CD). We ablate with 4 models, (i) (Baseline WO FT) a baseline model with a DeepSDF [27] network for SDF prediction and a similar larger capacity network for DDF prediction, (ii) (Baseline FT) the DDF predictions of the baseline are fine-tuned with the FineTuneNet, (iii) (Ours WO FT) our model without the FineTuneNet, and (iv) (Ours) our final model with FineTuneNet. The quantitative results confirm that our final model performs better than our baselines and our FineTuneNet indeed helps improve the reconstruction accuracy

Method	**Baseline WO FT**	**Baseline W FT**	**Ours WO FT**	**Ours**
CD ↓	1.30	1.29	1.26	**1.24**

Quantitative Results We show the quantitative results in Tab. 2 We compute the chamfers distance between 30000 samples of the reconstructed shape and the ground truth shape. We use DeepSDF's [27] preprocessing step to obtain the surface samples. As can be seen in Tab. 2, our model with FineTuneNet outperforms our model without the fine tuning layers. This could be a result of the fact that an accurate prediction of DDF enables the model to infer a surface close to the iso-surface obtained from the given depth maps. The quality of DDF

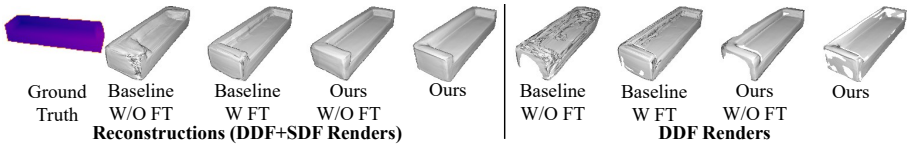

Fig. 5. Qualitative comparison of ablations. Given a ground truth depth map (left), we show renders of reconstructed 3D shapes (left-part). From left to right, baseline without FineTuneNet, baseline with FineTuneNet, ours without FineTuneNet and ours. As can be seen our models leads to the most plausible reconstructions. We show one pass forward renders with DDF in the right part, we probe the network for DDF and find if the point at the predicted distance is close to the surface (SDF≈ 0). This need not be the case for all predictions due to multi-view inconsistency. As can be seen, our model with FineTuneNet results in the most consistent prediction among all the others

prediction is important for us as we reconstruct with just a single evaluation through the network during optimization. The baselines with two models do not perform as well as our model as they do not implicity exploit the relation between DDF and SDF (which our models do).

Qualitative Results As can be seen in Fig. 5, our network performs better than the baselines. Further, in Fig. 5 (right) we observe that the DDF prediction is more consistent with the SDF prediction when there are fine-tuning layers.

5 Future Work

Overall, our method has shown significant improvement in terms of speed and quality of reconstructions from depth maps, and we believe that our proposed model opens up many exciting future possibilities. One of the most challenging problems is to make the DDF and SDF consistent. While we provide a solution to this problem based on test time optimization, we believe in better neural models that can couple DDF and SDF more tightly. Achieving this would mean that we marry the inherently implicit SDF representations with an explicit counterpart, thereby enabling even faster and consistent rendering. Another open problem with much wider applications that could help solve this limitation is to make the predictions of the neural network even more accurate.

6 Conclusion

We have presented a novel hybrid 3D representation that models both directional and signed distances. To this end, we have introduced a unique neural network architecture that can predict SDF and also DDF given an additional directional conditioning. We have demonstrated the effectiveness of our model by evaluating its performance on the task of 3D reconstruction from depth maps. For the first time, our model enables optimizing for a shape given a depth map

observation with just a single evaluation through the DDF and SDF networks, thereby making it not only 8.6× faster than existing methods, but also resulting in better reconstruction accuracy. Overall, we believe that our novel hybrid neural 3D shape representation can be used to address diverse other challenging inverse rendering tasks.

References

1. Agarwal, S., Furukawa, Y., Snavely, N., Simon, I., Curless, B., Seitz, S.M., Szeliski, R.: Building rome in a day. Communications of the ACM (Oct 2011). https://doi.org/10.1145/2001269.2001293, https://doi.org/10.1145/2001269.2001293
2. Alex Yu and Sara Fridovich-Keil, Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks (2021)
3. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (1999)
4. Bogo, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In: European Conference on Computer Vision (2016)
5. Chabra, R., Lenssen, J.E., Ilg, E., Schmidt, T., Straub, J., Lovegrove, S., Newcombe, R.: Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In: European Conference on Computer Vision. pp. 608–625. Springer (2020)
6. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: Shapenet: An information-rich 3d model repository (2015)
7. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
8. Chibane, J., Alldieck, T., Pons-Moll, G.: Implicit functions in feature space for 3d shape reconstruction and completion. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (jun 2020)
9. Choy, C.B., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: Proceedings of the European Conference on Computer Vision (ECCV) (2016)
10. Cremers, D.: Dynamical statistical shape priors for level set-based tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence **28**(8), 1262–1273 (2006)
11. Dawson-Haggerty et al.: trimesh, https://trimsh.org/
12. Georgia Gkioxari, Jitendra Malik, J.J.: Mesh r-cnn. Proceedings of the IEEE International Conference on Computer Vision (2019)
13. Hart, J.C.: Sphere tracing: a geometric method for the antialiased ray tracing of implicit surfaces. The Visual Computer **12**(10), 527–545 (Dec 1996). https://doi.org/10.1007/s003710050084
14. Huang, P.H., Matzen, K., Kopf, J., Ahuja, N., Huang, J.B.: Deepmvs: Learning multi-view stereopsis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2821–2830 (2018)
15. Jiang, C., Sud, A., Makadia, A., Huang, J., Nießner, M., Funkhouser, T., et al.: Local implicit grid representations for 3d scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6001–6010 (2020)

16. Kobbelt, L.P., Botsch, M., Schwanecke, U., Seidel, H.P.: Feature sensitive surface extraction from volume data. In: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques. p. 57–66. SIGGRAPH '01, Association for Computing Machinery, New York, NY, USA (2001). https://doi.org/10.1145/383259.383265, https://doi.org/10.1145/383259.383265
17. Koestler, L., Yang, N., Zeller, N., Cremers, D.: Tandem: Tracking and dense mapping in real-time using deep multi-view stereo. In: Conference on Robot Learning. pp. 34–45. PMLR (2022)
18. Kohlberger, T., Cremers, D., Rousson, M., Ramaraj, R., Funka-Lea, G.: 4d shape priors for a level set segmentation of the left myocardium in spect sequences. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 92–100. Springer (2006)
19. Leventon, M.E., Grimson, W.E.L., Faugeras, O.: Statistical shape influence in geodesic active contours. In: 5th IEEE EMBS International Summer School on Biomedical Imaging, 2002. pp. 8–pp. IEEE (2002)
20. Liu, S., Zhang, Y., Peng, S., Shi, B., Pollefeys, M., Cui, Z.: Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
21. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM Transactions on Graphics (Oct 2015)
22. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2019)
23. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
24. Muraki, S.: Volumetric shape description of range data using "blobby model". In: Proceedings of the 18th annual conference on Computer graphics and interactive techniques. pp. 227–235 (1991)
25. Neff, T., Stadlbauer, P., Parger, M., Kurz, A., Mueller, J.H., Chaitanya, C.R.A., Kaplanyan, A.S., Steinberger, M.: DONeRF: Towards Real-Time Rendering of Compact Neural Radiance Fields using Depth Oracle Networks. Computer Graphics Forum **40**(4) (2021). https://doi.org/10.1111/cgf.14340, https://doi.org/10.1111/cgf.14340
26. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2020)
27. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
28. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019)
29. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Deformable neural radiance fields. arXiv preprint arXiv:2011.12948 (2020)
30. Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: European Conference on Computer Vision. pp. 523–540. Springer (2020)

31. Piala, M., Clark, R.: Terminerf: Ray termination prediction for efficient neural rendering. In: 2021 International Conference on 3D Vision (3DV). pp. 1106–1114. IEEE (2021)
32. Rebain, D., Li, K., Sitzmann, V., Yazdani, S., Yi, K.M., Tagliasacchi, A.: Deep medial fields. arXiv preprint arXiv:2106.03804 (2021)
33. Ricci, A.: A constructive geometry for computer graphics. The Computer Journal **16**(2), 157–160 (1973)
34. Saito, S., Hu, L., Ma, C., Ibayashi, H., Luo, L., Li, H.: 3d hair synthesis using volumetric variational autoencoders. ACM Transactions on Graphics (TOG) **37**(6), 1–12 (2018)
35. Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2304–2314 (2019)
36. Saito, S., Simon, T., Saragih, J., Joo, H.: Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 84–93 (2020)
37. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2016)
38. Schönberger, J.L., Zheng, E., Pollefeys, M., Frahm, J.M.: Pixelwise view selection for unstructured multi-view stereo. In: European Conference on Computer Vision (2016)
39. Sitzmann, V., Rezchikov, S., Freeman, W.T., Tenenbaum, J.B., Durand, F.: Light field networks: Neural scene representations with single-evaluation rendering. In: Proc. NeurIPS (2021)
40. Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. Advances in Neural Information Processing Systems **32** (2019)
41. Sommer, C., Sang, L., Schubert, D., Cremers, D.: Gradient-sdf: A semi-implicit surface representation for 3d reconstruction. arXiv preprint (2021), `https://arxiv.org/abs/2111.13652`
42. Sturm, J., Bylow, E., Kahl, F., Cremers, D.: Copyme3d: Scanning and printing persons in 3d. In: German Conference on Pattern Recognition. pp. 405–414. Springer (2013)
43. Takikawa, T., Litalien, J., Yin, K., Kreis, K., Loop, C., Nowrouzezahrai, D., Jacobson, A., McGuire, M., Fidler, S.: Neural geometric level of detail: Real-time rendering with implicit 3D shapes (2021)
44. Tewari, A., Bernard, F., Garrido, P., Bharaj, G., Elgharib, M., Seidel, H.P., Pérez, P., Zöllhofer, M., Theobalt, C.: Fml: Face model learning from videos. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019)
45. Tewari, A., Zollöfer, M., Kim, H., Garrido, P., Bernard, F., Perez, P., Christian, T.: MoFA: Model-based Deep Convolutional Face Autoencoder for Unsupervised Monocular Reconstruction. In: Proceedings of the IEEE International Conference on Computer Vision (2017)
46. Tiwari, G., Sarafianos, N., Tung, T., Pons-Moll, G.: Neural-gif: Neural generalized implicit functions for animating people in clothing. In: International Conference on Computer Vision (ICCV) (October 2021)

47. Toeppe, E., Nieuwenhuis, C., Cremers, D.: Volume constraints for single view reconstruction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2013)
48. Wang, R., Yang, N., Stueckler, J., Cremers, D.: Directshape: Photometric alignment of shape priors for visual vehicle pose and shape estimation. In: Proceedings of the IEEE International Conference on Robotics and Automation (2020)
49. Wimbauer, F., Yang, N., Von Stumberg, L., Zeller, N., Cremers, D.: Monorec: Semi-supervised dense reconstruction in dynamic environments from a single moving camera. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6112–6122 (2021)
50. Wu, S., Rupprecht, C., Vedaldi, A.: Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2020)
51. Xu, Q., Wang, W., Ceylan, D., Mech, R., Neumann, U.: Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems 32, pp. 492–502. Curran Associates, Inc. (2019), http://papers.nips.cc/paper/8340-disn-deep-implicit-surface-network-for-high-quality-single-view-3d-reconstruction.pdf
52. Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 767–783 (2018)
53. Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Ronen, B., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. Advances in Neural Information Processing Systems **33** (2020)
54. Ye, Z., Yenamandra, T., Bernard, F., Cremers, D.: Joint deep multi-graph matching and 3d geometry learning from inhomogeneous 2d image collections (2022)
55. Yenamandra, T., Tewari, A., Bernard, F., Seidel, H.P., Elgharib, M., Cremers, D., Theobalt, C.: i3dmm: Deep implicit 3d morphable model of human heads. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12803–12813 (2021)
56. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images (2020)
57. Zobeidi, E., Atanasov, N.: A deep signed directional distance function for object shape representation. arXiv preprint arXiv:2107.11024 (2021)

Appendix

In this supplementary material, we show additional experiments that further evaluate our model. Specifically, we show results of ablation studies on our algorithm for 3D reconstruction from depth maps (see Sec. 3.4 in the main paper), and 3D shape reconstruction from silhouettes.

A Implementation Details

In this section, we provide more details about our network architecture and hyperparameters.

A.1 Network Architecture

DistanceNet: The DistanceNet consists of 3 MLPs, as shown in Fig. 6, with 2 hidden layers each. The hidden layers have a dimension of 512. The output dimension of the first two MLPs is 512 and that of the last MLP is 1. Encoded latent code and encoded directional conditioning codes are added to the output of the MLPs before activation. We use ReLU activation function for all the layers except the output. We use Tanh activation when predicting SDF and no activation when predicting DDF.
Other Networks: For all the other networks, we use MLPs with two hidden layers and ReLU activations. The MLP of DirectionNet is conditioned similarly as DistanceNet with the encoded latent code.

A.2 Hyperparameters

During training, we use the weights $w_s = 1.0$, $w_{dc} = 1.0$, $w_{df} = 1.0$, and $w_l = 0.0001$ respectively for the SDF loss, coarse DDF loss, fine-tuned DDF loss, and latent code regularizer. For 3D shape reconstruction from depth maps, during inference, we set $w_m = 1.0$, $w_D = 1.0$, $w_{DDF} = 1.0$, and $w_l = 0.0001$ respectively for silhouette loss, depth loss, latent regularizer, and regularizer for SDF-DDF consistency. For 3D shape reconstruction from silhouettes, we set $w_m = 1.0$, $w_D = 0.0$, $w_{DDF} = 1.0$, and $w_l = 0.05$.

We set a threshold of 0.009 to mark the points obtained using the predicted distances of the DDF network as surface points. That is, if the SDF at the points evaluates to less than the threshold, they are marked as surface points and if the SDF evaluates to more than the threshold, the predictions are discarded as not surface points.

Similar to DeepSDF [27], we set a truncation distance $\tau = 0.1$. In other words, SDF values that are more than 0.1 are set to 0.1 and those less than -0.1 are set to -0.1.

We run the code and parameters as released by the authors of DIST [20], and IF-NET [8] for comparisons.

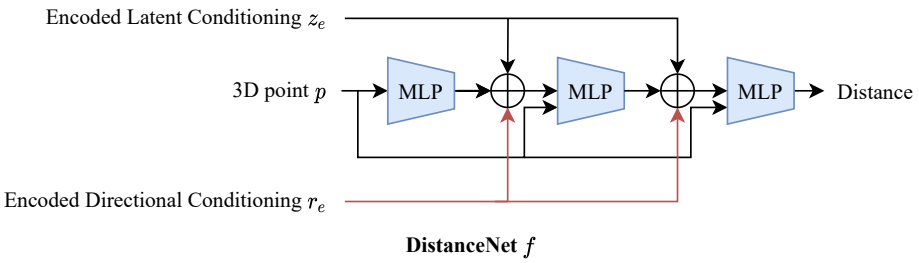

Fig. 6. Network architecture of our DistanceNet f. The encoded directional conditioning (red) is only applied while predicting DDF

B Shape from Silhouettes

We evaluate our method on the challenging task of reconstructing 3D shapes from silhouettes given a camera pose. The task is much more challenging compared with reconstructing from depth maps as we have no information about the shape of an object apart from a silhouette in an image.

We optimize for shape with the 3D reconstruction from depth maps algorithm, presented in Sec. 3.4 of the main paper, without the depth loss $w_D = 0.0$. We compare with DIST [20] for this task and also set the weight of the depth loss in its reconstruction algorithm to 0.0. We test on the same test split as in 3D reconstruction from depth maps, i.e., the first 200 shapes from test splits of each category and the first image from the test dataset of 3D-R2N2 [9].

Qualitative results: We show qualitative results of reconstructing 3D shape from a silhouette in Fig. 7. As can be seen, our method reconstructs more plausible shapes as per the given mask due to the strong prior formed by our hybrid representation, HDSDF. From a mask, our method is able to reproduce accurate finer details such as spoilers for cars, and wing geometry of airplanes.

Quantitative resuls: We show the quantitative results of our 3D shape from silhouettes experiment in Tab. 3. We report 1000× the L2 chamfer distance (see Sec. 4.3) between the reconstructed and ground truth meshes. Our method outperforms DIST [20] in almost all the classes by a large margin further corroborating that the prior learned by our hybrid representation is indeed strong and meaningful.

C Ablations - 3D Reconstruction from Depth Maps

We perform ablation studies on our 3D reconstruction from depth maps algorithm presented in Sec. 3.4 of the main paper. We study the main advantage of our model is the DDF prior used during reconstruction, therefore our ablation studies are focused on the DDF predictions.

We compare three algorithms for this evaluation — (i) (SDF+DDF, ours) obtaining the surface with one evaluation of the DDF pipeline and one evaluation

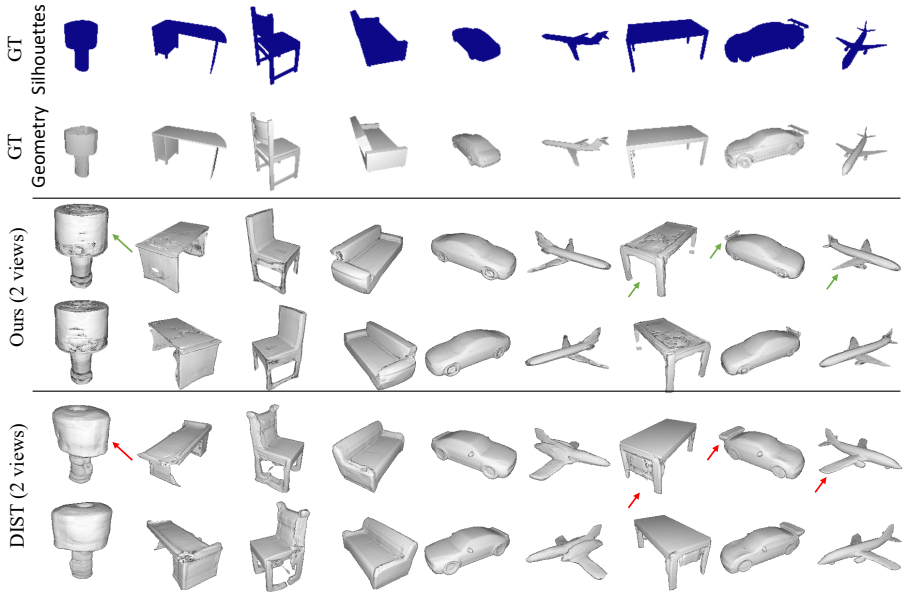

Fig. 7. Qualitative comparison of 3D shape reconstruction from silhouettes. In each column, given a silhouette (top) (ground truth geometry for reference in second row) we show renders of reconstructed 3D shapes with our method (middle rows) and shapes reconstructed with DIST [20] (bottom rows). As can be seen from images, our method better reconstructs shape from silhouettes proving that our hybrid representation, DDF along with SDF, is a stronger prior

of the SDF pipeline, (ii) (SDF as DDF) obtaining the surface with 1 evaluation of the SDF pipeline instead of the DDF pipeline to reach the surface, and (iii) (SDF+DIST) reconstructing 3D shape with our SDF prior using DIST's [20] reconstruction algorithm. For all the experiments we use the same split as in the ablation studies in the main paper (see Sec. 4.4).

SDF+DIST: Our reconstruction algorithm outperforms DIST's [20] both in terms of speed and accuracy. This could be due to two reasons, either our DDF model learns a strong prior of the surfaces and leads to better reconstructions or our SDF model is superior in comparison with DeepSDF [27] which DIST uses. Therefore, we study the performance of our SDF prior with the reconstruction algorithm from DIST [20]. **SDF as DDF**: We replace the DDF prediction with SDF predictions to study the influence of DDF in the reconstruction algorithm i.e., to ensure that the DDF predictions are indeed meaningful directional distances.

Qualitative results: We show the qualitative results in Fig. 8. As can be seen, our algorithm faithfully reconstructs finer details of the shape as compared with our SDF prior with DIST's reconstruction algorithm while being 8.6× per iteration faster than DIST. This shows that our DDF pipeline learns a strong

Table 3. Quantitative comparisons of our method with DIST [20] on 3D shape reconstruction given a silhouette. Our method outperforms DIST [20] in almost all the classes by a large margin, proving that our hybrid prior leads to more plausible 3D shape reconstructions given just a silhouette of the shape. Note that we report 1000× the L2 Chamfer distance (CD)

Method	**Ours**	**DIST** [20]
Metric	CD ↓	
Plane	**2.24**	5.74
Sofa	**2.84**	3.82
Table	9.98	**5.95**
Lamp	**9.54**	11.92
Chair	**4.79**	7.84
Car	**0.97**	1.53

prior on a class of shapes leading to better and faster reconstructions. Further, the SDF as DDF experiment shows that the DDF predictions are indeed meaningful directional distances to the surface and not just some direction-independent distance predictions.

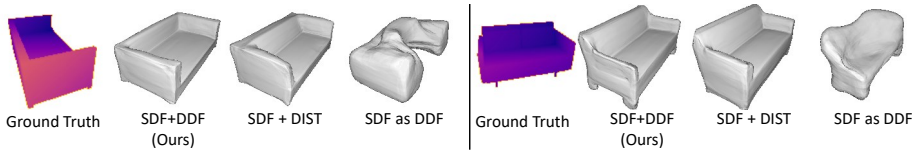

Fig. 8. Qualitative comparison of ablation studies on our 3D reconstruction algorithm. Given a ground truth depth map (left of each part of figure), we show renders of reconstructed 3D shapes (on the right part of each part of the figure). We show three experiments in the order — (i) (SDF+DDF) Our proposed algorithm where we only have one evaluation of the DDF and SDF pipelines, (ii) (SDF + DIST) Our SDF pipeline with DIST's [20] 3D reconstruction algorithm, and (iii) (SDF as DDF) Our SDF pipeline used instead of DDF. The 'SDF+DDF' ablation shows that the prior DDF learns is more meaningful leading to more plausible shapes. The 'SDF as DDF' ablation shows that the learned DDF is indeed meaningful and learns directional distances

Quantitative results: We show the quantitative results for the ablation studies in Tab. 4. In the table, we show 1000× the chamfer distance (see Sec. 4.3) between the reconstructed and ground truth meshes. As one would expect, when we use SDF instead of DDF prediction (SDF as DDF), the reconstructions are really bad therefore leading to bad results. Moreover, when we reconstruct with DIST's [20] algorithm with our SDF pipeline instead of DeepSDF's[27], we see a decrease in accuracy proving that our method performs better due to the strong prior learned by the DDF pipeline. Further, the experiment also shows that our

improvements, for the most part, are not just due to an improvement in the SDF predictions compared to DeepSDF [27].

Table 4. Ablation study on our algorithm to reconstruct 3D shape given a depth map. We evaluate the design choices of our reconstruction algorithm with our model trained on the sofas class with 256 shapes. We test the algorithm by reconstructing 64 shapes from depth maps. We report 1000× the L2 Chamfer distance (CD) between the ground truth shape and reconstructed shape. We ablate 3 design choices — (i) (SDF as DDF) use SDF predictions instead of DDF to show that our DDF predictions are meaningful directional distances, (ii) (SDF + DIST [20]) using the 3D reconstruction algorithm from DIST along with our SDF prior instead of our proposed algorithm to show that our hybrid representation and the presented reconstruction algorithm are indeed better, and (iii) ours

Method	**SDF as DDF**	**SDF + DIST [20]**	**Ours**
CD ↓	8.18	1.62	**1.24**

D Visualizations

Unless explicitly mentioned, we visualize the results of our method by one evaluation of DDF followed by local sphere tracing with the SDF network. We use sphere tracing to show the results of DIST [20], and we show reconstructed meshes for IF-NET [8]. For shading, we use an accelerated version of Blinn-Phong shader from i3DMM's code [55]. We render meshes of IF-NETs for visualizations, and 3D-R2N2's test data set for obtaining depth maps and masks with the help of Trimesh [11].