




Rethinking Directional Parameterization in Neural Implicit Surface Reconstruction

Zijie Jiang^{1*}, Tianhan Xu^{2*}, and Hiroharu Kato²

¹ Tokyo Institute of Technology

² Preferred Networks, Inc.

zjiang@ok.sc.e.titech.ac.jp, {xutianhan,hkato}@preferred.jp

Abstract. Multi-view 3D surface reconstruction using neural implicit representations has made notable progress by modeling the geometry and view-dependent radiance fields within a unified framework. However, their effectiveness in reconstructing objects with specular or complex surfaces is typically biased by the directional parameterization used in their view-dependent radiance network. *Viewing direction* and *reflection direction* are the two most commonly used directional parameterizations but have their own limitations. Typically, utilizing the viewing direction usually struggles to correctly decouple the geometry and appearance of objects with highly specular surfaces, while using the reflection direction tends to yield overly smooth reconstructions for concave or complex structures. In this paper, we analyze their failed cases in detail and propose a novel hybrid directional parameterization to address their limitations in a unified form. Extensive experiments demonstrate the proposed hybrid directional parameterization consistently delivered satisfactory results in reconstructing objects with a wide variety of materials, geometry and appearance, whereas using other directional parameterizations faces challenges in reconstructing certain objects. Moreover, the proposed hybrid directional parameterization is nearly parameter-free and can be effortlessly applied in any existing neural surface reconstruction method.

Keywords: surface reconstruction · neural implicit representation · directional parameterization

1 Introduction

Recent advances in neural fields [32] have revolutionized the field of 3D geometry reconstruction from multi-view images [22,29,36]. Compared to traditional multi-view stereo methods [7], a key advantage of using neural fields is their ability to model and optimize complex light reflections on object surfaces, which improves robustness to specular or non-Lambertian surfaces. Since a radiance network is trained jointly with a geometry network via differentiable rendering, the radiance network is able to take full advantage of geometry and view-dependent information such as viewing direction to decouple the geometry and appearance.

* Equal contribution. Work done while Zijie Jiang interned at Preferred Networks, Inc.




Fig. 1: Comparison of reconstruction results using different directional parameterizations in [2]. Using the viewing directional parameterization fails to reconstruct objects with specular surfaces, while using the reflection directional parameterization tends to result in incorrect geometry for objects with concave structures. Using our proposed hybrid parameterization shows consistently satisfactory reconstruction results compared to the existing parameterizations.

Since the decomposition of geometry and appearance is inherently an ill-posed problem [40]³, it is essential to introduce an appropriate inductive bias into the radiance network to facilitate correct decomposition. For example, simply inputting the viewing direction to the radiance network sometimes leads to incorrect reconstructions of highly specular objects (Fig. 1 (a)). To address this issue, the recent study [28] suggests that providing the reflection direction of rays considering the learned local surface geometry (*i.e.*, normals) would result in a simpler radiance field to learn and facilitate correct decomposition, especially for specular objects. Some subsequent studies have shown that adopting this parameterization results in satisfactory reconstructions of highly specular objects [8, 13, 14] and also non-reflecting objects [8]. However, although this parameterization seems promising (Fig. 1 (b) top), we observed that it still does not work well when reconstructing objects with complex structures such as concave surfaces (Fig. 1 (b) bottom). Interestingly, while using viewing direction fails to recover specular surfaces, it succeeds in recovering such concave structures (Fig. 1 (a) bottom). Based on these, we pose the question: **Is there a unified directional parameterization that can effectively handle both specular surfaces and complex structures?**

In this paper, to answer the question, we start by studying the limitations of existing directional parameterizations and provide the first in-depth analysis of why using reflection direction may fail to reconstruct objects with concave or complex structures. As mentioned above, one important feature of the reflection directional parameterization is its dependency on the learnable geometry (*i.e.*,

³ If the capacity of the radiance network is sufficiently high, even for a fixed geometry such as a sphere, the image reconstruction loss can converge to zero [40].

normals). Our insight into this fact is that for sampled 3D points along a sampled ray during training, while utilizing reflection directions of sampled points near the intersecting surface is helpful in decoupling geometry and appearance, the use of reflection directions of sampled points elsewhere may hinder the accurate reconstruction of geometry. Specifically, we identified that the cause of this issue lies in the sampled points beyond a certain distance from the intersecting surface during the optimization. As we will discuss in Sec. 3.2, the normals of these points computed from learnable geometry may either be associated with surfaces unrelated to the intersection or introduce high-frequency variations in the input of the radiance network, which hinders the proper updating of geometry and radiance. This hindrance is particularly noticeable in the reconstruction of concave surfaces. The analysis also indirectly explains the advantage of using viewing directional parameterization in reconstructing concave or complex structures, as it is non-learnable and remains constant for all sampled points along a sampled ray.

Based on our insights and analysis, we propose a novel hybrid directional parameterization to address the limitations of existing directional parameterizations in a unified form. The main idea is to utilize the predicted distance to the surface for each sampled point: if the distance is small, indicating proximity to the intersecting surface, it is preferable to use its reflection direction as the directional parameterization, which facilitates the correct decomposition of geometry and appearance; in other cases, we instead utilize the non-learnable viewing direction to avoid the optimization issue mentioned above. The proposed hybrid directional parameterization smoothly transitions from the reflection direction to the viewing direction based on the predicted distance to the surface. The rate of this directional transition is controlled by a learnable scaling factor. We integrate the proposed hybrid directional parameterization into different methods and conduct extensive experiments to validate its effectiveness across various datasets. Compared to the existing viewing and reflection directional parameterizations, our proposed hybrid directional parameterization achieves consistently satisfactory results, as shown in Fig. 1 (c). Moreover, the proposed hybrid directional parameterization is nearly parameter-free and can be effortlessly applied in any existing neural implicit surface reconstruction method.

To summarize, our contributions are as follows:

- We are the first to provide an in-depth analysis of why the use of reflection directional parameterization tends to reconstruct incorrect geometry of objects with concave or complex structures.
- We propose a novel hybrid directional parameterization to address the limitations of existing directional parameterizations, which can effectively handle both specular surfaces and complex structures in a unified form.
- We validate the effectiveness of the proposed directional parameterization across different methods on various datasets, and demonstrate an overall superiority compared to existing directional parameterizations.

2 Related Work

2.1 Neural Implicit Surface Reconstruction

With the significant advancement of neural implicit representations [3, 16–19, 23, 24, 37] and differentiable rendering [10, 11, 17, 20, 25, 27, 33] in recent years, neural implicit surface reconstruction methods [21, 29, 35, 36] have emerged as a promising alternative to traditional methods in the field of image-based 3D surface reconstruction. Compared to traditional multi-view stereo methods, neural implicit surface methods employ a more flexible and representative implicit shape representation, such as the signed distance function (SDF), and introduce the directional radiance field [17] to model the view-dependent appearance of the target object. Subsequent works on neural implicit surface reconstruction have mostly built upon the foundation of NeuS [29] or VolSDF [35], and made improvements for specific scenarios including fast rendering [30, 31], sparse-view inputs [15], large-scale scenes [12, 38] and more.

2.2 Reconstruction for Specular Objects

Despite neural implicit surface reconstruction methods consider the view-dependent appearance of objects by incorporating the directional radiance field, reconstruction for specular objects remains a challenging task due to its significant shape-radiance ambiguity. Several works [14, 26, 34, 39, 41] model the view-dependent appearance of specular objects by incorporating additional physical components (*e.g.*, materials and lighting), with the hope of jointly recovering both the geometric shape and physical parameters. However, as the number of unknown parameters increases, these works typically rely on a good initial geometry to ensure favorable results, leading to a chicken-and-egg problem. On the other hand, Ref-NeRF [28] was the first to propose the use of a radiance field conditioned on reflection direction to better decouple the geometry and radiance of specular objects. Subsequent neural implicit surface reconstruction methods targeting specular objects [8, 13, 14] have also adopted this directional parameterization. However, these methods often fail to correctly recover the geometry of objects with concave or complex structures.

3 Method

In this section, we introduce a novel directional parameterization for neural implicit surface reconstruction methods. We first give a brief review on neural implicit surface reconstruction and its commonly used directional parameterizations, including viewing and reflection directions (Sec. 3.1). In the next, we summarize the observed failed reconstructions when utilizing existing directional parameterizations and conduct a detailed analysis of their causes (Sec. 3.2). Finally, a hybrid directional parameterization is proposed based on our analysis, which can effectively address the limitations of existing directional parameterizations in a unified form (Sec. 3.3).

3.1 Preliminaries of Neural Implicit Surface Reconstruction

Neural implicit surface reconstruction methods [2, 6, 29, 36] typically represent the 3D object as the combination of a signed distance function (SDF) network Θ_f , which maps a spatial position $\mathbf{x} \in \mathbb{R}^3$ to its SDF value $f(\mathbf{x}) \in \mathbb{R}$, and a directional radiance network Θ_c , which maps a spatial position $\mathbf{x} \in \mathbb{R}^3$ and a directional vector $\mathbf{d} \in \mathbb{S}^2$ to the view-dependent color at \mathbf{x} . The surface \mathcal{S} of the object is extracted as the zero-level set of $f(\mathbf{x})$. To enable the joint optimization of the SDF network Θ_f and radiance network Θ_c using only 2D image supervisions, volume rendering is typically utilized [29, 35]. Given the sampled pixels in an image, camera rays $\{\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}_{\text{view}} | t > 0\}$ are cast through the camera origin \mathbf{o} and the sampled pixels, where \mathbf{d}_{view} is the unit directional vector of the ray. To obtain the rendered colors of the sampled rays, a number of points are sampled along each ray, and queried for their SDF values and view-dependent colors from Θ_f and Θ_c , where positional encoding (PE) [17] is typically incorporated. A relationship between the SDF value and opaque density is then formulated, such as using S-density [29], to translate the queried SDFs and colors to the ray color akin to NeRF [17]. Finally, the training loss is calculated from the difference between the rendered ray colors and the ground-truth colors of the sampled pixels.

Directional parameterization. The key to the ability of the radiance network to model view-dependent appearance lies in its incorporation of a directional vector \mathbf{d} as an input. Common directional parameterizations for \mathbf{d} includes the viewing direction \mathbf{d}_{view} and reflection direction \mathbf{d}_{ref} . The viewing directional parameterization is intuitively straightforward, as it models the view-dependent appearance from the perspective of the observer. The reflection directional parameterization is a little more complicated, which encodes the viewing direction \mathbf{d}_{view} and local normal direction \mathbf{n} into a new compact directional vector \mathbf{d}_{ref} , by considering the interaction between light and surface. Intuitively, it represents the direction of the incoming light, assuming that the surface is specular. \mathbf{d}_{ref} is formulated as follows:

$$\mathbf{d}_{\text{ref}} = 2(\mathbf{d}_{\text{view}} \cdot \mathbf{n})\mathbf{n} - \mathbf{d}_{\text{view}}, \quad (1)$$

where \mathbf{n} can be derived from the gradient of the SDF as $\mathbf{n}(\mathbf{x}) = \frac{\nabla f(\mathbf{x})}{\|\nabla f(\mathbf{x})\|}$.

3.2 Observations and Analysis

Observations. Although existing neural implicit surface reconstruction methods widely adopt either viewing direction or reflection direction as their directional parameterization, both types of parameterizations have their respective limitations when it comes to reconstructing certain objects. Using viewing direction often faces difficulty in correctly reconstructing objects with specular surfaces (*e.g.*, Fig. 1 (a) and Fig. 2 (c)). On the other hand, while using reflection direction can reconstruct specular objects with simple geometry (*e.g.*, Fig. 1 (b) top and Fig. 2 (d)), it often results in overly smooth reconstructions with




Fig. 2: Scenarios where existing directional parameterizations succeed and struggle. Reconstruction results are obtained by integrating the two directional parameterizations into NeuS [29].

missing concave surfaces when reconstructing objects with concave or complex structures (*e.g.*, Fig. 1 (b) bottom and Fig. 2 (b)). Interestingly, using viewing direction may yield better results in these parts with complex or concave surfaces (*e.g.*, Fig. 1 (a) bottom and Fig. 2 (a)).

Based on these observations, we raise the question: Is there a unified directional parameterization that can effectively handle both specular surfaces and complex geometry? To answer this, we start by analyzing the reasons behind their reconstruction failures:

Analysis of failures when using viewing direction. An analysis of why using viewing direction is unsuitable for modeling the view-dependent appearance of specular objects is provided in Verbin *et al.* [28], which points out that the radiance network conditioned on viewing direction is poorly suited for interpolation. We consider this characteristic leads to difficulties in correctly decoupling the geometry and appearance of specular objects when using viewing direction. On the other hand, reflection directional parameterization can effectively leverage the prior knowledge of the interaction between light and surface, which facilitates the correct decoupling of geometry and appearance, especially for specular objects.

Analysis of failures when using reflection direction. In contrast, no work has yet provided a detailed analysis of why using reflection directional parameterization fails in the reconstruction of concave or complex structures, making us the first to provide this analysis. We first highlight two key differences between viewing direction and reflection direction: the viewing direction is non-learnable throughout the optimization and remains constant along a sampled ray, whereas the reflection direction 1) depends on learnable geometry (*i.e.*, normals) and 2) may vary significantly during the optimization. We argue that these two properties of the reflection direction may pose challenges to the accurate recovery of geometry, especially for concave or complex structures.

A 2D toy example is first illustrated on Fig. 3 (left) to better demonstrate why the first property of reflection direction may hinder the recovery of geometry, which assumes the scenario of recovering an “L”-shaped concave structure. A sampled ray intersecting the surface and the sampled points on the ray along with the normals at these points are illustrated. In this case, we notice that the normals at sampled points that are slightly distant from the intersecting surface

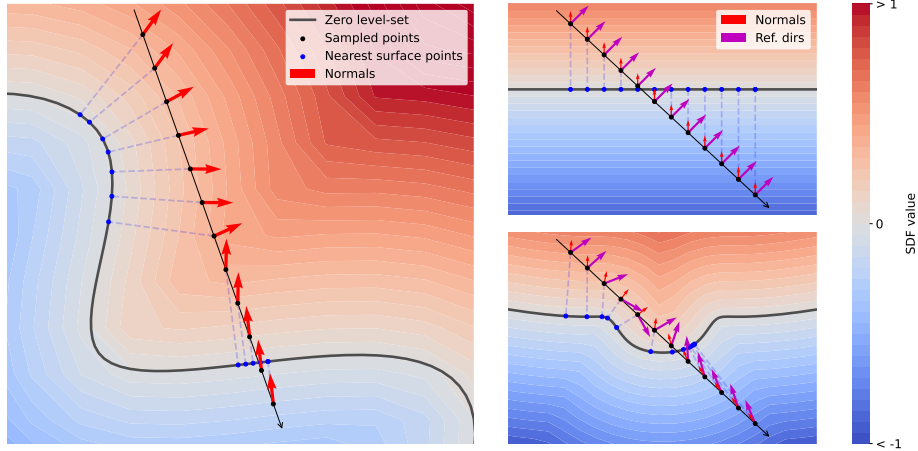


Fig. 3: Impact of incorporating normals on view-dependent radiance modeling. (Left) Normals at sampled points may be influenced by unrelated surfaces other than the intersection surface. (Top right) Smooth surfaces yield similarity in normals and reflection directions for the sampled points, while (Bottom right) surfaces with intricate local details, *e.g.*, concavities, induce a scattered distribution of normals and reflection directions (Ref. dirs), particularly pronounced slightly distant from the vicinity of the zero-level set, which adversely affects geometry optimization. The details are explained in Sec. 3.2.

do not necessarily relate to the normal of the intersecting surface (the bottom side of the zero level-set), but rather to the normals of surfaces elsewhere (the left side of the zero level-set). As a result, since the training loss of existing neural implicit surface reconstruction methods is usually ray-based, which is calculated using the local color of the intersecting surface, using this loss to optimize the normals of irrelevant surfaces elsewhere may introduce ambiguity in the optimization of geometry and result in incorrect geometric updates.

In addition, the second property of the reflection direction is more closely related to the phenomenon that using reflection directional parameterization tends to reconstruct overly smooth surfaces with missing concave or complex structures. Assuming that the geometry reconstruction begins with a smooth surface exhibiting local flatness (Fig. 3 top-right), which aligns with the geometric initialization practices commonly employed by existing neural implicit surface reconstruction methods [6, 29], and the geometry of the target object includes concave structures (Fig. 3 bottom-right). We observed that for the same sampled ray, when the surface geometry is updated from a smooth surface to a concave surface, the normals and reflection directions of the same sampled points also change. This change is particularly pronounced for points that are slightly distant from the surface, and could be further pronounced due to the presence of Positional Encoding (PE). The noticeable changes in reflection directions at these points within each surface update introduce high-frequency variances to

the directional inputs of the radiance network. This makes learning a stable radiance field challenging. To be specific, assuming that the radiance network conditioned on reflection direction has already learned to minimize the color loss to some extent before the surface update, the substantial change in directional inputs due to the update could likely lead to an increase in color loss, which hinders the optimization and is susceptible to falling into local minima.

In contrast, we notice that using the viewing direction often leads to concave or overly intricate surface details (although some may be incorrect), as shown in Fig. 1 (a) and Fig. 2 (a) (c). We consider that this is because the viewing direction does not rely on the surface geometry, which means that changes in geometry do not affect the optimization of the radiance network conditioned on viewing direction. This allows the SDF can undergo more “aggressive” updates without considering the impact on radiance.

3.3 Hybrid Directional Parameterization

Building upon the analysis from the previous section, we discovered that viewing direction and reflection direction seem to have a spatially complementary relationship: On one hand, for sampled points near the intersecting surface, the normals here accurately reflect the normals of the intersecting surface, thereby enabling the reflection direction to better model the interaction between light and surface, which in turn promotes the correct decomposition of geometry and appearance. On the other hand, for other sampled points along the ray, using the reflection direction might introduce irrelevant surface geometry into the optimization, or high-frequency variations in the directional inputs of the radiance network, both potentially hindering the optimization of the geometry and radiance networks. By contrast, using the viewing direction at these sampled points can avoid these issues. Motivated by these thoughts, we propose a spatial-aware directional representation, which transitions from reflection direction \mathbf{d}_{ref} to viewing direction \mathbf{d}_{view} based on the distance to the surface (*i.e.*, the absolute value of SDF). We refer to this fused direction as **hybrid direction** \mathbf{d}_{hyb} , formulated as:

$$\mathbf{d}_{\text{hyb}} = \text{normalize}(\alpha \cdot \mathbf{d}_{\text{ref}} + (1 - \alpha) \cdot \mathbf{d}_{\text{view}}). \quad (2)$$

Here, $\alpha \in [0, 1]$ represents the blend weight. We normalize the resulting linear blend to a unit vector thereby representing a valid direction. The blend weight α is implemented as

$$\alpha = \exp(-\gamma \cdot \text{detach}(|f(\mathbf{x})|)), \quad (3)$$

where $\gamma \in \mathbb{R}^+$ is an optimizable parameter that controls the transition speed between reflection/viewing direction. The gradient of the SDF is detached to prevent interference with geometric updates. At the surface of the object, *i.e.*, when SDF equals 0, α simply becomes 1, which means that the hybrid direction aligns with the reflection direction. As \mathbf{x} moves further away from the surface, α gradually approaches 0, indicating that the direction tends towards the viewing

direction. Here, α can be any bell-shaped function w.r.t. the SDF value. In this case, we utilize the PDF of the Laplace distribution for simplicity.

In the next section, extensive experimental results validate the superiority of the proposed hybrid directional parameterization compared to existing parameterizations.

4 Experiments

4.1 Experimental Settings

Datasets. We conduct experiments on objects from different datasets, including the DTU dataset [1], the Shiny Blender dataset [28], and the real captured dataset [9]. The DTU dataset [1] contains indoor multi-view captures of diverse objects, including challenging conditions for geometry reconstruction such as non-Lambertian surfaces and complex structures. We use the same 15 objects as previous works [2, 29, 31, 36] for evaluation. In addition, we utilize the foreground masks provided by Yariv *et al.* [36]. The Shiny Blender dataset [28] is a synthetic dataset rendered in Blender [4], which contains different objects with specular surfaces. We follow the experimental settings used in Ref-NeuS [8] on this dataset, which selects four objects. The real captured dataset [9] contains a set of three outdoor captures of real objects with complex geometry and specular surfaces. We select the “toycar” scene for evaluation.

Implementation details. We integrate the proposed hybrid directional parameterization into two types of neural surface reconstruction methods: NeuS [29] and NeuDA [2]. The former utilizes a deep MLP network as the scene representation, while the latter leverages a shallow MLP decoder combined with multi-level voxel grids. Both methods adopt a radiance network conditioned on viewing direction in their original versions. The modification of the directional input from the viewing direction to the proposed hybrid direction can be accomplished with less than 10 lines of code. We additionally implement the radiance network that depends on reflection direction for both methods, to facilitate a comprehensive comparison of different directional parameterizations. Regarding the learnable parameter γ in Eqn. 3, we implement it as $\gamma = \exp(10 \cdot \gamma_b)$. We assign an initial value of 0.3 to γ_b except for the “toaster” scene from the Shiny Blender dataset. For the “toaster” scene, an initial value of 0.1 is assigned since it results in better reconstruction. The training losses of the modified methods are identical to those of their original versions. The number of training iterations is set to 300,000 to ensure convergence.

Evaluation metrics. For the DTU dataset, we follow the evaluation protocol of previous works [2, 29] and evaluate the Chamfer Distance (CD) between the estimated and ground-truth geometry. For the Shiny Blender dataset, we follow the evaluation protocol proposed in Ref-NeuS [8]. It evaluates the accuracy (Acc.) instead of the CD for unbiased evaluation, and also evaluates the mean angular error (MAE) between the ground-truth normal maps and the rendered normal maps. For the real captured dataset [9], we only provide a qualitative evaluation, as ground-truth geometry is not provided in this dataset.

Table 1: Quantitative evaluation in terms of Chamfer distance of integrating different directional parameterizations in NeuS [29] and NeuDA [2] on the DTU dataset [1]. The experimental objects are categorized into diffuse objects and specular objects. The numbers in the headers denote the scene IDs. **Bold** results have the best score.

Methods	Scan of Diffuse Objects													Scan of Specular Objects					Total Avg.
	24	37	40	55	65	83	105	106	114	118	122	Avg.	63	69	97	110	Avg.		
NeuS [29]	0.83	0.98	0.56	0.37	0.59	1.45	0.78	0.52	0.36	0.45	0.45	0.67	1.13	0.60	0.95	1.43	1.03	0.77	
NeuS, w/ reflection dir.	1.39	4.39	0.73	0.39	1.34	1.49	0.92	0.59	0.36	0.46	0.54	1.14	1.11	0.63	1.46	0.92	1.03	1.11	
NeuS, w/ hybrid dir. (Ours)	0.81	0.97	0.54	0.38	0.61	1.47	0.77	0.51	0.35	0.45	0.47	0.67	1.10	0.60	0.85	0.90	0.86	0.72	
NeuDA [2]	0.51	0.76	0.39	0.37	0.56	1.37	0.79	0.50	0.34	0.42	0.46	0.59	1.08	0.57	1.13	0.80	0.90	0.67	
NeuDA, w/ reflection dir.	0.58	1.03	0.63	0.37	0.79	1.45	0.89	0.51	0.37	0.43	0.50	0.69	1.03	0.55	0.90	0.76	0.81	0.72	
NeuDA, w/ hybrid dir. (Ours)	0.53	0.77	0.40	0.36	0.57	1.38	0.76	0.47	0.33	0.41	0.47	0.59	1.02	0.54	0.76	0.75	0.77	0.63	

4.2 Comparison of Different Directional Parameterizations

We first evaluate the performance of integrating different directional parameterizations into the two baseline methods NeuS [29] and NeuDA [2] on the DTU dataset [1]. The experimental objects from the DTU dataset are categorized into diffuse objects and specular objects for better demonstration. Tab. 1 summarizes the quantitative results. For both baseline methods, we notice that substituting the original viewing directional parameterization with the reflection direction only yields better results for a minority of specular objects. For diffuse objects and the remaining specular objects, employing reflection directional parameterization demonstrates significant degradation in performance. Considering the relatively complex geometry of DTU objects, we attribute this to the optimization issue analyzed in Sec. 3.2. Conversely, the incorporation of the proposed hybrid directional parameterization enables both modified methods to yield consistently superior performance for specular objects while achieving results on par with the original version for diffuse objects. The first two rows in Fig. 4 showcase the visual comparison of adopting different directional parameterizations for the reconstruction of two specular objects from the DTU dataset. Both baseline methods fail to reconstruct the specular surfaces of objects. On the other hand, although integrating the reflection direction yields better reconstruction for specular objects with simple geometry (*e.g.*, Scan 63), they fail to reconstruct the specular surfaces at the concave structure in Scan 97. In contrast, integrating the proposed hybrid direction achieves the best reconstruction results in both scenarios. More qualitative comparisons on the DTU dataset are available in our supplementary.

We further validate the advantage of the proposed hybrid directional parameterization on the Shiny Blender dataset [28], as demonstrated in Tab. 2. Both baseline methods exhibit large reconstruction errors on the surfaces of highly specular objects in this dataset. By contrast, due to the relatively simple geometry of objects, replacing the viewing directional parameterization with the reflection direction results in a significant improvement in reconstruction quality on most objects in this dataset, except for the “coffee” object. For the reconstruction of the “coffee” object, adopting the reflection directional parameterization fails to recover the concave structure formed by the inner wall of the coffee cup

Table 2: Quantitative evaluation of integrating different directional parameterizations in NeuS [29] and NeuDA [2] on the Shiny Blender dataset [28]. **Bold** results have the best score.

Methods	helmet		toaster		coffee		car		Avg.	
	Acc. ↓	MAE ↓	Acc. ↓	MAE ↓	Acc. ↓	MAE ↓	Acc. ↓	MAE ↓	Acc. ↓	MAE ↓
NeuS [29]	4.88	3.20	3.31	2.85	1.97	1.06	0.86	0.95	2.76	2.02
NeuS, w/ reflection dir.	0.31	0.36	0.39	0.94	3.30	1.37	0.48	0.73	1.12	0.85
NeuS, w/ hybrid dir. (Ours)	0.29	0.34	0.45	0.92	0.64	0.55	0.43	0.72	0.45	0.63
NeuDA [2]	5.73	3.77	5.98	3.39	1.16	0.81	0.87	1.03	3.44	2.25
NeuDA, w/ reflection dir.	0.34	0.37	1.03	1.71	3.66	1.53	0.62	0.87	1.41	1.12
NeuDA, w/ hybrid dir. (Ours)	0.32	0.36	0.95	1.54	0.55	0.57	0.60	0.84	0.61	0.83




Fig. 4: Qualitative comparison of integrating different directional parameterizations in NeuS [29] and NeuDA [2] on the DTU dataset [1] and Shiny Blender dataset [28].

and the surface of the water, as shown in the third row in Fig. 4. On the other hand, the integration of our proposed hybrid directional parameterization enables the successful reconstruction of both the specular surfaces and concave structures, obtaining consistently superior reconstructions of all objects in this dataset, compared to the other two directional parameterizations.

Lastly, we qualitatively present the reconstruction results on the “toycar” scene from the real captured scene in Fig. 5, which contains challenging reflective glass and complex geometry. Both modified methods employing the proposed hybrid directional parameterization succeed in recovering the complex geometry and the smooth surface of the reflective glass. In contrast, using either the reflection direction or the original viewing direction results in reconstructions that are either overly smooth, missing concave surface structures, or provide incorrect surface reconstructions at the reflective glass areas.

As the summary of this section, compared to the viewing directional parameterization which struggles with reconstructing specular surfaces, and the reflection directional parameterization which faces difficulties in reconstructing

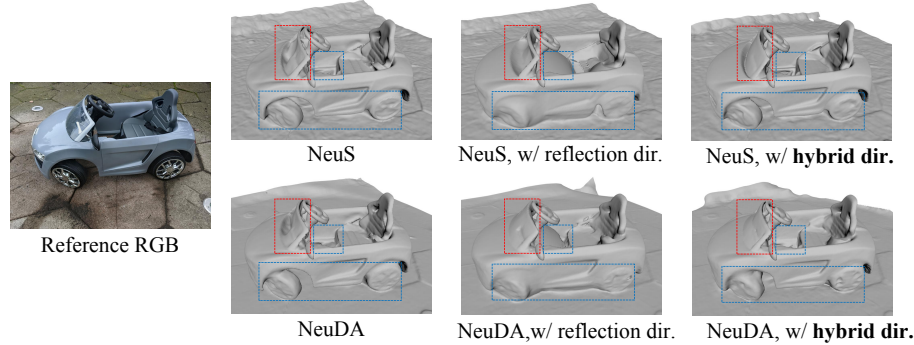


Fig. 5: Qualitative comparison of integrating different directional parameterizations in NeuS [29] and NeuDA [2] on the “toycar” scene from the real captured dataset [9].

Table 3: Quantitative evaluation in terms of Chamfer distance on the DTU dataset [1]. The numbers in the header denote the scene IDs. **Bold** results have the best score.

Methods	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Avg.
IDR [36]	1.63	1.87	0.63	0.48	1.04	0.79	0.77	1.33	1.16	0.76	0.67	0.90	0.42	0.51	0.53	0.90
NeuS [29]	0.83	0.98	0.56	0.37	1.13	0.59	0.60	1.45	0.95	0.78	0.52	1.43	0.36	0.45	0.45	0.77
Voxurf [31]	0.65	0.74	0.39	0.35	0.96	0.64	0.85	1.58	1.01	0.68	0.60	1.11	0.37	0.45	0.47	0.72
NeuDA [2]	0.51	0.76	0.39	0.37	1.08	0.56	0.57	1.37	1.13	0.79	0.50	0.80	0.34	0.42	0.46	0.67
Ours (NeuS)	0.81	0.97	0.54	0.38	1.10	0.61	0.60	1.47	0.85	0.77	0.51	0.90	0.35	0.45	0.47	0.72
Ours (NeuDA)	0.53	0.77	0.40	0.36	1.02	0.57	0.54	1.38	0.76	0.76	0.47	0.76	0.33	0.41	0.47	0.63

complex structures, our proposed hybrid directional parameterization effectively handles both scenarios in a unified form. This benefit is attributed to our hybrid directional parameterization effectively combining the advantages of both viewing and reflection directional parameterizations while avoiding their drawbacks.


4.3 Comparison to State-of-the-Art

Although our main purpose is not to develop a best-performing surface reconstruction system, we demonstrate that we can still achieve state-of-the-art performance by integrating the proposed hybrid directional parameterization into existing methods. For brevity, we refer to NeuS [29] and NeuDA [2] incorporating the proposed hybrid directional parameterization as Ours (NeuS) and Ours (NeuDA) in the following text. Tab. 3 reports the Chamfer distance (CD) metrics of different methods on the DTU dataset [1]. For a fair comparison, all methods in Tab. 3 utilize foreground masks provided by Yariv *et al.* [36] during training. Both Ours (NeuS) and Ours (NeuDA) improve their baselines by a certain margin in the average CD scores, and Ours (NeuDA) achieves the best performance compared to previous SOTA methods in the average CD metric.

Tab. 4 reports the quantitative evaluation of different methods on the Shiny Blender dataset [28]. The results of [5, 6, 8, 22, 28, 35] are taken from the results

Table 4: Quantitative evaluation on the Shiny Blender dataset [28]. “-” denotes that the results are not available. **Bold** results have the best score.

Methods	helmet		toaster		coffee		car		Avg.	
	Acc. ↓	MAE ↓	Acc. ↓	MAE ↓	Acc. ↓	MAE ↓	Acc. ↓	MAE ↓	Acc. ↓	MAE ↓
Ref-NeRF [28]	-	29.48	-	42.87	-	12.24	-	14.93	-	24.88
NeuS [29]	4.88	3.20	3.31	2.85	1.97	1.06	0.86	0.95	2.76	2.02
UNISURF [22]	1.69	1.78	3.75	5.51	2.88	3.15	2.34	1.98	2.67	3.11
VolSDF [35]	1.55	1.37	2.02	2.53	2.23	2.28	0.58	1.04	1.60	1.81
NeuralWarp [5]	2.25	1.94	5.90	3.51	1.54	2.04	0.65	1.07	2.59	2.14
Geo-NeuS [6]	0.74	2.36	6.35	3.76	3.85	6.36	2.88	5.67	3.46	4.54
NeuDA [2]	5.73	3.77	5.98	3.39	1.16	0.81	0.87	1.03	3.44	2.25
Ref-NeuS [8]	0.29	0.38	0.42	1.47	0.77	0.99	0.37	0.80	0.46	0.91
Ours (NeuS)	0.29	0.34	0.45	0.92	0.64	0.55	0.43	0.72	0.45	0.63
Ours (NeuDA)	0.32	0.36	0.95	1.54	0.55	0.57	0.60	0.84	0.61	0.83

**Fig. 6:** Reconstruction results using different initial values of γ_b . The main differences are highlighted in red boxes.

reported in Ge *et al.* [8]. Both Ours (NeuS) and Ours (NeuDA) achieve better average MAE metrics than the previous SOTA method Ref-NeuS [8], and Ours (NeuS) achieve the best overall performance in terms of both Acc. and MAE metrics on this dataset. We refer to our supplementary for the qualitative comparisons between Ours (NeuS) and Ref-NeuS [8].

4.4 Ablation Studies

In this section, we conduct ablation studies on several key design choices of the proposed hybrid directional parameterization by using Ours (NeuDA).

The initial value of γ_b . We first validate the impact of setting different initial values for γ_b on the reconstruction results. Fig. 6 compares the obtained reconstruction results using five different initial values for γ_b . We find that when the initial value of γ_b is set low, our method tends to recover incorrect geometries similar to those caused by using the reflection directional parameterization. This phenomenon remains consistent with our analysis in Sec 3.2, as a smaller initial value for γ_b increases the likelihood of involving irrelevant surfaces and high-frequency variations in the directional inputs of the radiance network during the optimization. We empirically find the best initial value of γ_b to be 0.3 for most experimental objects. More discussions regarding the parameter γ_b and blend weight α can be found in our supplementary material.

Detach SDF gradient in Eqn. 3 or not? Fig. 7 compares the reconstruction results with and without detaching the gradient of the estimated SDF in Eqn. 3. We observe that our method struggles to reconstruct the fine-grained geometry




Fig. 7: Reconstruction results with and without detaching the gradient of the estimated SDF in Eqn. 3.

Fig. 8: Reconstruction results fusing directions before and after positional encoding (PE).

when the gradient of the estimated SDF is not detached in Eqn. 3. We consider that introducing the gradient of the estimated SDF places an additional burden on the learning of both the geometry and radiance networks, making it easier to fall into local minima. Thereby, we choose to detach the gradient of SDF in Eqn. 3.

Fuse the viewing direction and reflection direction before or after positional encoding? We introduce the fusion of viewing direction and reflection direction in Eqn. 2 before the operation of position encoding (PE). We compare the reconstruction results with different fusion orders, specifically the proposed pre-PE fusion: $\text{PE}(\mathbf{d}_{\text{hyb}})$ (ours) and alternative approach using post-PE fusion: $\alpha \cdot \text{PE}(\mathbf{d}_{\text{ref}}) + (1 - \alpha) \cdot \text{PE}(\mathbf{d}_{\text{view}})$. The comparison is shown in Fig. 8. We observe that fusing directions after the PE operation may result in degraded reconstruction results. We consider that fusing the directional information in the high-dimension space may increase the complexity of the learning space of the radiance network, which brings challenges in finding optimal reconstruction results. Thereby, we choose to fuse the viewing direction and reflection direction before the operation of positional encoding.

5 Conclusion and Limitation

In this paper, we conduct an in-depth investigation into the topic of directional parameterization in neural implicit surface reconstruction. We reveal that the commonly used directional parameterizations in modeling the view-dependent radiance field, such as viewing and reflection directions, both have their limitations and none of them can effectively handle both specular surfaces and complex structures. A novel hybrid directional parameterization is introduced to achieve this goal, which softly fuses the viewing and reflection directions based on the distance of the sampled points from the surface, uniting the advantages of each representation. Extensive experiments demonstrate that our hybrid directional parameterization outperforms the existing directional parameterizations in reconstructing objects with a wide variety of materials, geometry, and appearance.

One limitation of our method is that the optimal choice of initial value for γ_b may vary depending on the target object, which would require additional efforts for parameter tuning. An automated method for selecting the optimal initial value of γ_b for target objects could be a potential future work for us.

References

1. Aanæs, H., Jensen, R.R., Vogiatzis, G., Tola, E., Dahl, A.B.: Large-scale data for multiple-view stereopsis. *IJCV* **120**, 153–168 (2016)
2. Cai, B., Huang, J., Jia, R., Lv, C., Fu, H.: Neuda: Neural deformable anchor for high-fidelity implicit surface reconstruction. In: *CVPR* (2023)
3. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: *CVPR* (2019)
4. Community, B.O.: Blender - a 3D modelling and rendering package. Blender Foundation, Stichting Blender Foundation, Amsterdam (2018)
5. Darmon, F., Bascle, B., Devaux, J., Monasse, P., Aubry, M.: Improving neural implicit surfaces geometry with patch warping. In: *CVPR* (2022)
6. Fu, Q., Xu, Q., Ong, Y.S., Tao, W.: Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. In: *NeurIPS* (2022)
7. Furukawa, Y., Hernández, C., et al.: Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision* **9**(1-2), 1–148 (2015)
8. Ge, W., Hu, T., Zhao, H., Liu, S., Chen, Y.C.: Ref-neus: Ambiguity-reduced neural implicit surface learning for multi-view reconstruction with reflection. In: *ICCV* (2023)
9. Hedman, P., Srinivasan, P.P., Mildenhall, B., Barron, J.T., Debevec, P.: Baking neural radiance fields for real-time view synthesis. In: *CVPR* (2021)
10. Kato, H., Beker, D., Morariu, M., Ando, T., Matsuoka, T., Kehl, W., Gaidon, A.: Differentiable rendering: A survey. *arXiv preprint arXiv:2006.12057* (2020)
11. Li, T.M., Aittala, M., Durand, F., Lehtinen, J.: Differentiable monte carlo ray tracing through edge sampling. *ACM Transactions on Graphics* **37**(6), 1–11 (2018)
12. Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: *CVPR* (2023)
13. Liang, R., Chen, H., Li, C., Chen, F., Panneer, S., Vijaykumar, N.: Envidr: Implicit differentiable renderer with neural environment lighting. *arXiv preprint arXiv:2303.13022* (2023)
14. Liu, Y., Wang, P., Lin, C., Long, X., Wang, J., Liu, L., Komura, T., Wang, W.: Nero: Neural geometry and brdf reconstruction of reflective objects from multiview images. *ACM Transactions on Graphics* (2023)
15. Long, X., Lin, C., Wang, P., Komura, T., Wang, W.: Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In: *ECCV* (2022)
16. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: *CVPR* (2019)
17. 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)
18. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics* **41**(4), 1–15 (2022)
19. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Occupancy flow: 4d reconstruction by learning particle dynamics. In: *ICCV* (2019)
20. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: *CVPR* (2020)
21. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: *CVPR* (2020)

22. Oechsle, M., Peng, S., Geiger, A.: Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In: ICCV (2021)
23. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: CVPR (2019)
24. Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: ECCV (2020)
25. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. arXiv preprint arXiv:2007.08501 (2020)
26. Srinivasan, P.P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B., Barron, J.T.: Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In: CVPR (2021)
27. Tulsiani, S., Zhou, T., Efros, A.A., Malik, J.: Multi-view supervision for single-view reconstruction via differentiable ray consistency. In: CVPR (2017)
28. Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J.T., Srinivasan, P.P.: Ref-nerf: Structured view-dependent appearance for neural radiance fields. In: CVPR (2022)
29. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In: NeurIPS (2021)
30. Wang, Y., Han, Q., Habermann, M., Daniilidis, K., Theobalt, C., Liu, L.: Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In: ICCV (2023)
31. Wu, T., Wang, J., Pan, X., Xu, X., Theobalt, C., Liu, Z., Lin, D.: Voxurf: Voxel-based efficient and accurate neural surface reconstruction. In: ICLR (2023)
32. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. *Computer Graphics Forum* **41**(2), 641–676 (2022)
33. Yan, X., Yang, J., Yumer, E., Guo, Y., Lee, H.: Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In: NeurIPS (2016)
34. Yao, Y., Zhang, J., Liu, J., Qu, Y., Fang, T., McKinnon, D., Tsin, Y., Quan, L.: Neilf: Neural incident light field for material and lighting estimation. In: ECCV (2022)
35. Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. In: NeurIPS (2021)
36. Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Ronen, B., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. In: NeurIPS (2020)
37. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: CVPR (2021)
38. Yu, Z., Peng, S., Niemeyer, M., Sattler, T., Geiger, A.: Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. In: NeurIPS (2022)
39. Zhang, J., Yao, Y., Li, S., Liu, J., Fang, T., McKinnon, D., Tsin, Y., Quan, L.: Neilf++: Inter-reflectable light fields for geometry and material estimation (2023)
40. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492 (2020)
41. Zhang, X., Srinivasan, P.P., Deng, B., Debevec, P., Freeman, W.T., Barron, J.T.: Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics* **40**(6), 1–18 (2021)