

# Normal-NeRF: Ambiguity-Robust Normal Estimation for Highly Reflective Scenes

Ji Shi, Xianghua Ying\*, Ruohao Guo, Bowei Xing, Wenzhen Yue

National Key Laboratory of General Artificial Intelligence  
School of Intelligence Science and Technology, Peking University  
sji118@pku.edu.cn, xhying@pku.edu.cn

## Abstract

Neural Radiance Fields (NeRF) often struggle with reconstructing and rendering highly reflective scenes. Recent advancements have developed various reflection-aware appearance models to enhance NeRF’s capability to render specular reflections. However, the robust reconstruction of highly reflective scenes is still hindered by the inherent shape ambiguity on specular surfaces. Existing methods typically rely on additional geometry priors to regularize the shape prediction, but this can lead to oversmoothed geometry in complex scenes. Observing the critical role of surface normals in parameterizing reflections, we introduce a transmittance-gradient-based normal estimation technique that remains robust even under ambiguous shape conditions. Furthermore, we propose a dual activated densities module that effectively bridges the gap between smooth surface normals and sharp object boundaries. Combined with a reflection-aware appearance model, our proposed method achieves robust reconstruction and high-fidelity rendering of scenes featuring both highly specular reflections and intricate geometric structures. Extensive experiments demonstrate that our method outperforms existing state-of-the-art methods on various datasets.

**Code** — <https://github.com/sji118/Normal-NeRF>

## Introduction

Neural Radiance Fields (NeRF) (Mildenhall et al. 2021) has been extensively studied for its powerful 3D scene reconstruction and rendering capabilities. NeRF utilizes multi-layer perceptrons (MLPs) to encode a 3D scene into continuous fields of volume density and view-dependent radiance. Despite NeRF’s proficiency in capturing fine geometric structures and smoothly varying view-dependent appearance, it often fails to accurately represent highly specular reflections, as the high-frequency view-dependent appearance prevents NeRF from eliminating the shape-radiance ambiguity (Zhang et al. 2020).

Recent studies attempt to enhance NeRF’s capability in rendering reflective scenes by incorporating reflection-aware appearance models. Ref-NeRF (Verbin et al. 2022) reparameterizes the view-dependent radiance as a function of

the reflection direction about the surface normal, rather than the viewing direction itself. Later works further employ advanced reflection-aware appearance models, such as Microfacet Reflectance Model (Mai et al. 2023) and Whitted-Style Ray Tracing (Zeng et al. 2023). While these advanced reflection models are capable of representing complex reflections, achieving accurate reflection modeling remains a challenge due to the inherent shape ambiguity on highly reflective surfaces.

We observe that surface normals are critical in nearly all reflection-aware appearance models. Previous works typically derive surface normals based on the negative normalized gradient of the density field, which we term the “density gradient”. However, this approach becomes unreliable on highly specular surfaces, as NeRF tends to fake specular reflections with foggy artifacts embedded behind the surfaces (Verbin et al. 2022). This phenomenon prevents the density values from monotonically increasing towards the object’s interior, resulting in inconsistent gradient directions, as illustrated in Figure 1(a).

In this paper, we propose an ambiguity-robust normal estimation pipeline to enhance NeRF’s capability in reconstructing and rendering highly reflective scenes. To address the unreliability in conventional normal estimation approaches, we introduce the concept of “transmittance gradient”, which can provide accurate normal estimates even under conditions of ambiguous shape predictions. Instead of calculating the gradient of the density, we compute the gradient of the transmittance, which is naturally monotonic along any specific ray, as illustrated in Figure 1(b). Furthermore, we observe that reflection-aware appearance models generally prefer locally smooth surface normals, whereas NeRF tends to reconstruct a steep density field to achieve sharper object boundaries. To address this conflict, we propose a dual activated densities module that accommodates the distinct requirements of both the surface normals and the density field. Additionally, we design a stop-gradient warmup strategy for the predicted normal loss to prevent it from impeding the optimization of the density field.

By integrating a reflection-aware appearance model, our model achieves robust reconstruction and high-fidelity rendering in highly reflective scenes. Our key contributions can be summarized as follows:

- To the best of our knowledge, we are the first to identify

\*Xianghua Ying is the corresponding author.  
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

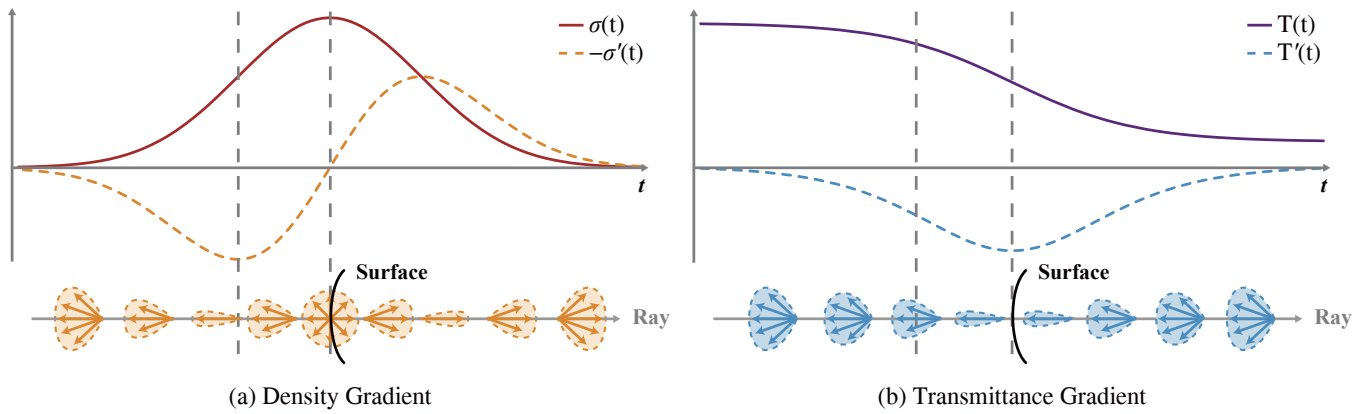


Figure 1: Illustration of our transmittance gradient compared to conventional density gradient for normal estimation. We plot the density  $\sigma(t)$  and the transmittance  $T(t)$  along a ray passing through a semi-transparent surface (top), and the distribution of the estimated normal vectors (bottom). (a) Since the derivatives of density approach zero near the surface, the directions of density gradients in nearby area can change rapidly. (b) In contrast, the derivatives of transmittance are large near the surface, thereby producing consistent normal estimates that align well with the true surface normals.

and analyze the inherent limitations of the density gradient traditionally used for NeRF normal estimation.

- We introduce the **Transmittance Gradient** to address inconsistencies and irregularities in normal estimation within NeRF.
- We propose a **Dual Activated Densities** module to effectively bridge the gap between smooth surface normals and sharp object boundaries.

### Related Work

**Neural Radiance Fields.** Neural Radiance Fields (NeRF) (Mildenhall et al. 2021) is a successful pipeline for novel view synthetic of complex scenes by optimizing the volumetric underlying function using a sparse set of input views. Later studies have improved the rendering quality (Barron et al. 2021, 2022, 2023) and accelerated the rendering speed (Müller et al. 2022; Sun, Sun, and Chen 2022; Chen et al. 2022) of NeRF using various techniques. NeRF has also inspired many subsequent works that extend its application, including few-shot rendering (Chen et al. 2021; Yu et al. 2021; Jain, Tancik, and Abbeel 2021; Yang, Pavone, and Wang 2023), dynamic scene rendering (Park et al. 2021a,b; Li et al. 2022), and 3D generation (Schwarz et al. 2020; Jain et al. 2022; Poole et al. 2023).

**Radiance Fields with Reflections.** Although the view-dependent radiance function of NeRF enables the modeling of non-Lambertian effects, it often encounters difficulties in accurately capturing the appearance of objects with high-frequency reflections due to the shape-radiance ambiguity (Zhang et al. 2020). NeRFReN (Guo et al. 2022) and MS-NeRF (Yin et al. 2023) treat specular reflections on planar mirrors as virtual images behind the surface, and model them with separate radiance fields to avoid inconsistency between front and back views of the mirror. Neural Catacaustics (Kopanas et al. 2022) further leverages a neural warp field, enabling the modeling of reflections on non-planar sur-

faces. However, due to the lack of physically-based modeling of interactions between light and surfaces, these methods face challenges in accurately representing reflections on complex surfaces.

Recent works incorporate NeRF with reflection-aware appearance models. Ref-NeRF (Verbin et al. 2022) conditions the view-dependent radiance on the reflected view direction instead of the camera view direction to make the radiance MLP easier to interpolate. Neural Microfacet Fields (Mai et al. 2023) employs a microfacet reflectance model for physically-accurate reflection rendering. Mirror-NeRF (Zeng et al. 2023) models multi-bounce reflections with Whitted-Style Ray Tracing. ENVIDR (Liang et al. 2023) leverages a pretrained neural renderer to enable high-fidelity scene relighting. SpecNeRF (Ma et al. 2024) proposes a Gaussian Directional Encoding to model near-field lighting in room-scale scenes. NeRF-Casting (Verbin et al. 2024) efficiently casts reflection rays to synthesize consistent reflections.

**Normal Estimation within Radiance Fields.** Estimating normals within NeRF poses a non-trivial challenge due to the absence of an explicit defined surface. Early research (Bi et al. 2020) employ an MLP to predict normal vectors directly without any regularization. Later works (Boss et al. 2021; Srinivasan et al. 2021) characterize the normal vectors as the negative normalized gradient of the density field, thereby enforcing the consistency between the surface normals and the density field. Subsequent studies (Zhang et al. 2021; Kuang et al. 2022; Verbin et al. 2022) integrate these approaches by tying the normals predicted by MLP to the normals estimated from the density field. However, these approaches become unreliable in highly reflective scenes, due to the non-monotonic nature of the density field under conditions of ambiguous shape predictions. More recent studies attempt to resolve this ambiguity by employing planar constraints (Zeng et al. 2023) or by enforcing the surfaces to be opaque (Mai et al. 2023). Nevertheless, these methods rely

on additional geometry priors, which may not be suitable for all scenes.

Another series of research (Yariv et al. 2021; Wang et al. 2021) substitutes the density field in NeRF with a signed distance function (SDF), providing an explicit definition of surfaces. However, this approach also encounters challenges when reconstructing surfaces within highly reflective scenes. Ref-NeuS (Ge et al. 2023) attempt to reduce the ambiguity by calculating a reflection score to identify specular regions. Despite this, errors in the reflection score can still lead to incorrect geometry reconstruction, particularly on concave surfaces.

## Preliminary

Neural Radiance Fields (NeRF) (Mildenhall et al. 2021) encodes a scene as continuous volumetric fields, where the density  $\sigma(\mathbf{x}) \in \mathbb{R}$  at any 3D position  $\mathbf{x} \in \mathbb{R}^3$  and the color  $\mathbf{c}(\mathbf{x}, \mathbf{d}) \in \mathbb{R}^3$  at any 3D position  $\mathbf{x} \in \mathbb{R}^3$  under any viewing direction  $\mathbf{d} \in \mathbb{R}^2$  can be queried from MLPs. The color of a ray  $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$  is rendered as:

$$\hat{\mathbf{C}}(\mathbf{r}) = \int_0^{+\infty} T(t)\sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt, \quad (1)$$

where

$$T(t) = \exp\left(-\int_0^t \sigma(\mathbf{r}(s))ds\right) \quad (2)$$

is the transmittance along the ray, which indicates the probability of light traveling along the ray over the interval  $[0, t)$  without being absorbed or scattered. To approximate the integral  $\hat{\mathbf{C}}(\mathbf{r})$ , NeRF samples a set of points  $\{\mathbf{x}_i = \mathbf{o} + t_i\mathbf{d}\}$  and denotes the distance between adjacent samples by  $\delta_i = t_{i+1} - t_i$ . The density  $\sigma_i$  and the color  $\mathbf{c}_i$  at each point  $\mathbf{x}_i$  under the direction  $\mathbf{d}$  are then queried to approximate the color of the ray as:

$$\mathbf{C}(\mathbf{r}) = \sum_i T_i(1 - \exp(-\sigma_i\delta_i))\mathbf{c}_i, \quad (3)$$

where

$$T_i = \exp\left(-\sum_{j<i} \sigma_j\delta_j\right). \quad (4)$$

NeRF is optimized by minimizing the L2 difference between the ground truth color  $\mathbf{C}_{\text{gt}}(\mathbf{r})$  of each pixel taken from input images and the predicted color  $\mathbf{C}(\mathbf{r})$  of the ray corresponding to this pixel:

$$\mathcal{L}_c(\mathbf{r}) = \|\mathbf{C}(\mathbf{r}) - \mathbf{C}_{\text{gt}}(\mathbf{r})\|^2. \quad (5)$$

## Irregularity in Normal Estimation

Since volume density and surface normals both characterize object shape, it naturally follows to estimate normal vectors from the density field. Observing that volume density increases drastically at the boundary between non-opaque air and opaque objects, NeRD (Boss et al. 2021) characterizes normal vectors directly as the negative normalized gradients of the density field, which we refer to as the ‘‘density gradients’’:

$$\mathbf{n}^d(\mathbf{x}) = -\frac{\nabla\sigma(\mathbf{x})}{\|\nabla\sigma(\mathbf{x})\|}. \quad (6)$$

However, this approach is only effective on opaque surfaces, where the volume density can strictly increase towards the interior of the object. When encountering highly reflective surfaces, NeRF tends to fake reflections by positioning them underneath the surfaces, resulting in semi-transparent surface predictions during the optimization process. Local maxima in the density field will inevitably occur near a semi-transparent surface. Since the gradient at the maximum point is zero, the directions and magnitudes of gradients in nearby areas can change rapidly, leading to an irregular distribution of estimated normal vectors, as illustrated in Figure 1(a). Additionally, the gradients on either side of the surface may point in opposite directions.

Later studies (Verbin et al. 2022; Mai et al. 2023) introduce an orientation loss to penalize normal vectors that face away from the camera. However, this regularization encourages opaque surface predictions, compelling the model to select a specific surface from potential candidates and discard the rest without substantial evidence. Therefore, these methods still struggle with the accurate reconstruction of highly reflective scenes, as demonstrated in Figure 4.

## Method

Our goal is to enhance NeRF’s robustness and fidelity in reconstructing and rendering highly reflective scenes. We begin by introducing the concept of the transmittance gradient to address the irregularity in normal estimation under ambiguous shape conditions. Subsequently, we employ the dual activated densities to meet the distinct requirements of surface normals and object boundaries. Finally, we present details of our training process, including a stop-gradient warmup strategy for the predicted normal loss, and a reflection-aware appearance model.

### Transmittance Gradient

Despite the failure of density gradients to provide robust normal estimates on highly reflective surfaces, the density field remains a valuable source of geometric information for estimating normal vectors. Additionally, we find that the source of the irregularity in density gradients can be attributed to the non-monotonic nature of the volume density. Building on these insights, we identify the transmittance, as defined in Eq. (2), as a promising mediator for linking the density field and surface normals. The accumulated transmittance along any specific ray is a monotonically decreasing function, and its derivative at any position is equal to the rendering weight at that position (Tagliasacchi and Mildenhall 2022). Therefore, any sample point that contributes significantly to the final rendering will possess a correspondingly large derivative value, thereby avoiding the irregularity encountered by density gradients, as illustrated in Figure 1(b).

However, Both the transmittance and its derivative are defined along a specific 1D ray. To extract 3D directional information from them, we introduce the concept of the ‘‘transmittance gradient’’ as the normalized gradient of the trans-

mittance over a segment of ray with respect to the translation of this ray. More specifically, consider a ray originating at point  $\mathbf{o}$  and directed along  $\mathbf{d}$ . For any point  $\mathbf{x} = \mathbf{o} + t\mathbf{d}$  on this ray, the transmittance over the segment between  $\mathbf{o}$  and  $\mathbf{x}$  can be reparameterized by substituting  $\mathbf{o}$  with  $\mathbf{x} - t\mathbf{d}$ :

$$\hat{T}(\mathbf{x}; \mathbf{d}, t) = \exp\left(-\int_0^t \sigma(\mathbf{x} - s\mathbf{d}) ds\right). \quad (7)$$

Then the transmittance gradient at point  $\mathbf{x}$  is formulated as:

$$\mathbf{n}^t(\mathbf{x}; \mathbf{d}, t) = \frac{\nabla_{\mathbf{x}} \hat{T}(\mathbf{x}; \mathbf{d}, t)}{\left\| \nabla_{\mathbf{x}} \hat{T}(\mathbf{x}; \mathbf{d}, t) \right\|}. \quad (8)$$

If the ray is originated from a camera and the point  $\mathbf{x}$  is visible to the camera, the transmittance gradient  $\mathbf{n}^t(\mathbf{x}; \mathbf{d}, t)$  serves as an estimate of the normal vector at the position  $\mathbf{x}$ .

Unlike the density, which is encoded by a differentiable MLP, the transmittance in NeRF is approximated through numerical integration. Consequently, the transmittance gradient cannot be directly obtained via automatic differentiation. Following the quadrature rule used in NeRF, we estimate the transmittance gradient with the same discrete set of samples:

$$\mathbf{n}_i^t = -\frac{\sum_{j < i} \nabla \sigma(\mathbf{x}_j) \delta_j}{\left\| \sum_{j < i} \nabla \sigma(\mathbf{x}_j) \delta_j \right\|}. \quad (9)$$

### Dual Activated Densities

While our transmittance gradient effectively mitigates irregularities in normal estimation under conditions of ambiguous shape prediction, the disparity between the density and surface normal can still induce instability. Sharp object boundaries, which are crucial for high-fidelity renderings and geometric details (Sun, Sun, and Chen 2022), necessitate a steep density field. In contrast, reflection-aware appearance models generally prefer locally smooth surface normals to accurately reconstruct high-frequency reflections. To bridge this gap, we propose a dual activated densities module that simultaneously supports smooth surface normals and maintains sharp object boundaries.

In our design, we apply two distinct activation functions, `softplus` and `exp`, to the output of the density MLP. Specifically, we activate a sharp density,  $\hat{\sigma} = \exp(b)$ , and a smooth density,  $\tilde{\sigma} = \text{softplus}(b)$ , where  $b$  represents the pre-activation output of the density MLP. The sharp density  $\hat{\sigma}$  will be used to calculate the rendering weights in Eqs. (3) and (4), while the smooth density  $\tilde{\sigma}$  will be used to compute the transmittance gradient according to Eq. (9), by substituting  $\sigma$  in these formulas. Shared learnable MLP parameters ensure consistency between the two densities, while dual activation functions allow for varying degrees of steepness.

Moreover, the dual activated densities also reduces numerical instability when calculating the transmittance gradient. As our transmittance gradient is approximated through numerical integration, a steep density field can induce artifacts caused by numerical instability, as demonstrated in Figure 2. Conversely, the smooth density  $\tilde{\sigma}$  activated by `softplus` significantly enhances the stability of the numerical integration process.

### Training

While the transmittance gradient provides an accurate normal estimate, it lacks precision. Specifically, the transmittance gradient value at a given spatial position can vary slightly with changes in viewing direction and sampling strategy. To address this, we employ a spatial MLP to refine the transmittance gradient, ensuring both accuracy and precision in normal prediction, which is crucial for reflection modeling. For any position  $\mathbf{x}$  within the scene, we predict the normal  $\mathbf{n}^p(\mathbf{x}) \in \mathbb{R}^3$  by normalizing a three-dimensional vector output from the spatial MLP.

### Predicted Normal Loss with Stop-Gradient Warmup.

Ref-NeRF (Verbin et al. 2022) ties the normals predicted by MLP and the normals estimated from the density field using a simple predicted normal loss. However, we observe that this loss function may impede the optimization of the density field. To address this issue, we extend the predicted normal loss with a stop-gradient warmup strategy. Specifically, we apply the stop-gradient operator `sg` to allow for the adjustment of the proportion of gradients flowing from the predicted normals  $\mathbf{n}^p$  to the density field (including the rendering weights  $w$  and the transmittance gradients  $\mathbf{n}^t$ ):

$$\mathcal{L}_n = \lambda_n \overleftarrow{\mathcal{L}}_n + (1 - \lambda_n) \overrightarrow{\mathcal{L}}_n, \quad (10)$$

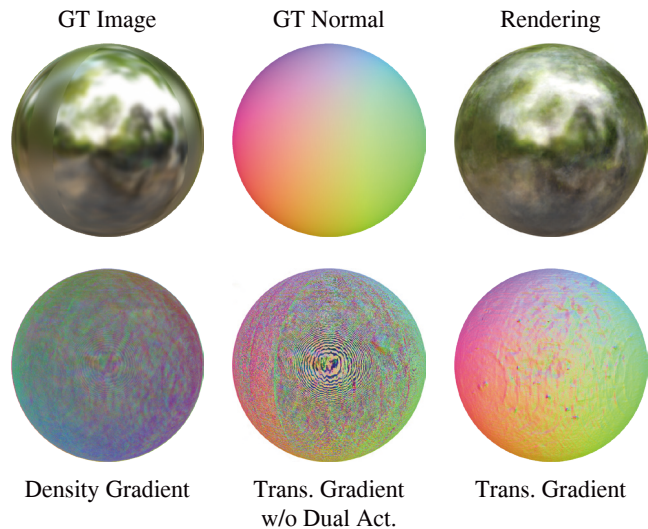


Figure 2: We intentionally reconstruct a highly reflective scene using a baseline model that excludes any reflection-aware appearance model, which is unable to eliminate the ambiguity in shape prediction. Under this ambiguous shape condition, we visually compare different normal estimation methods. The density gradient method completely fails to produce reasonable normal estimates. Omitting dual activated densities in our transmittance gradient method leads to artifacts from numerical instability, while our full pipeline produces normal estimates that align well with the ground truth.




Figure 3: Qualitative comparisons on test views of synthetic scenes.

where

$$\begin{aligned} \overleftarrow{\mathcal{L}}_n &= \sum_i w_i \|\mathbf{n}_i^p - \mathbf{n}_i^t\|^2, \\ \overrightarrow{\mathcal{L}}_n &= \sum_i \text{sg}(w_i) \|\mathbf{n}_i^p - \text{sg}(\mathbf{n}_i^t)\|^2. \end{aligned} \quad (11)$$

In all of our experiments, the parameter  $\lambda_n$  follows an exponential warmup, increasing from 0.01 to 1 over 20k iterations.

Our key insight regarding this design is that the density field is more reliable than the predicted normals at the beginning of training. Although both the density field and predicted normals are initialized randomly, the density field quickly converges to reasonable shape predictions. Conversely, in the absence of the predicted normal loss, the predicted normals may even degenerate into a piecewise constant function (Verbin et al. 2022). Consequently, at the beginning of training, the gradients from the predicted normals could disrupt the density field. In contrast, the gradients from the density field can help in regularizing the predicted normals. Therefore, we restrict the majority of gradients flowing from the predicted normals to the density field at the beginning of training.

**Reflection-Aware Appearance.** To demonstrate the efficacy of our normal estimation technique, we develop our reflection-aware appearance model as follows. For any sample point  $\mathbf{x}$  observed under viewing direction  $\mathbf{d}$ , we compute the reflection direction using the predicted normal  $\mathbf{n}^p$  at this position:

$$\mathbf{d}^r = 2(\mathbf{d} \cdot \mathbf{n}^p)\mathbf{n}^p - \mathbf{d}, \quad (12)$$

We then feed the reflection direction into an environment MLP  $\mathcal{F}_{\text{env}}$  to obtain an environment feature:

$$\mathbf{f}_{\text{env}} = \mathcal{F}_{\text{env}}(\mathbf{d}^r). \quad (13)$$

Incorporating a material feature  $\mathbf{f}_{\text{mat}}$  conditioned exclusively on spatial position, we decompose the color into its

diffuse and specular components:

$$\begin{aligned} \mathbf{c}_s &= \mathcal{F}_s(\mathbf{f}_{\text{mat}}, \mathbf{f}_{\text{env}}), \\ \mathbf{c}_d &= \mathcal{F}_d(\mathbf{f}_{\text{mat}}). \end{aligned} \quad (14)$$

Finally, we combine the diffuse component and the specular component in linear space and then convert it to sRGB space with gamma tone mapping (Anderson et al. 1996):

$$\mathbf{c} = \gamma(\mathbf{c}_d + \mathbf{c}_s). \quad (15)$$

## Experiments

**Datasets.** To comprehensively validate the effectiveness and robustness of our proposed method, we conduct evaluation on several datasets, including widely-used NeRF Synthetic dataset (Mildenhall et al. 2021), two reflective objects datasets: Shiny Blender (Verbin et al. 2022) and Glossy Synthetic (Liu et al. 2023), and one real captured dataset from Ref-NeRF (Verbin et al. 2022).

**Baselines and Metrics.** We compare our method against the following baselines: **Zip-NeRF** (Barron et al. 2023), a state-of-the-art grid-based NeRF variant with no special treatment for reflection; **Ref-NeRF** (Verbin et al. 2022), a NeRF-based method focusing on reflective objects rendering; **NMF** (Mai et al. 2023), a NeRF-based method for inverse rendering with microfacet reflectance model; **Ref-NeuS** (Ge et al. 2023), a SDF-based method for reflective surface reconstruction with a reflection-aware photometric loss. **ENVIDR** (Liang et al. 2023), a SDF-based method for scene relighting with a pretrained neural renderer. We evaluate rendering quality using PSNR, SSIM and LPIPS (Zhang et al. 2018), and assess normal accuracy with mean angular error (MAE) (Verbin et al. 2022).

**Implementation Details.** All experiments are conducted on an NVIDIA RTX 4090 GPU. We implement our model within Nerfstudio (Tancik et al. 2023), building upon the Instant-NGP (Müller et al. 2022) framework. We train our

	NeRF Synthetic				Shiny Blender				Glossy Synthetic		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Zip-NeRF	<i>33.69</i>	<b>0.974</b>	<b>0.028</b>	-	30.12	0.948	0.089	-	29.88	0.944	<i>0.065</i>
Ref-NeRF	<u>33.99</u>	<i>0.966</i>	<i>0.038</i>	24.077	<u>35.96</u>	0.967	<i>0.059</i>	18.384	<i>30.39</i>	0.941	0.074
NMF	30.71	0.940	0.053	<i>20.952</i>	<u>34.56</u>	0.963	<u>0.053</u>	6.061	28.76	0.932	0.086
Ref-NeuS	25.83	0.916	0.099	22.973	33.26	<i>0.971</i>	<u>0.053</u>	<i>4.852</i>	<u>30.49</u>	<u>0.956</u>	<u>0.064</u>
ENVIDR	29.33	0.942	0.064	<u>20.549</u>	<i>35.02</i>	<u>0.972</u>	<u>0.053</u>	<u>4.602</u>	28.47	<i>0.948</i>	0.066
Ours	<b>34.53</b>	<u>0.972</u>	<u>0.030</u>	<b>14.486</b>	<b>39.24</b>	<b>0.982</b>	<b>0.040</b>	<b>4.241</b>	<b>33.24</b>	<b>0.971</b>	<b>0.043</b>

Table 1: Comparison with SOTAs on NeRF Synthetic (Mildenhall et al. 2021), Shiny Blender (Verbin et al. 2022) and Glossy Synthetic (Liu et al. 2023). The best results are **bold**, the second best results are underlined, and the third best results are *italics*.

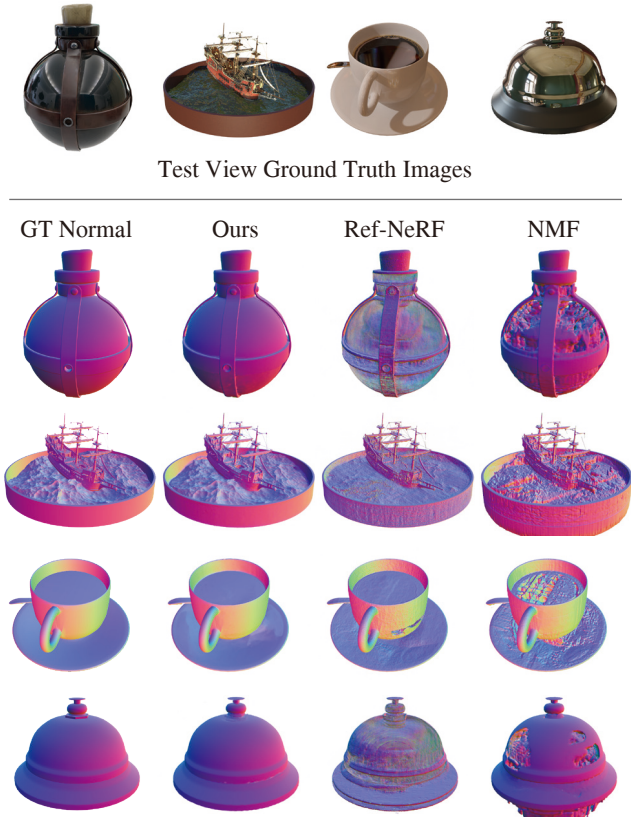


Figure 4: Normal map visualizations of NeRF-based methods.

model for 50k iterations with a batch size of  $2^{19}$  sample points. Please refer to supplementary material for more details.

### Comparison with State-of-the-Arts

Table 1 reports quantitative results on synthetic datasets. Our method demonstrates superior performance on Shiny Blender dataset and Glossy Synthetic dataset, and achieves performance on par with Zip-NeRF (Barron et al. 2023) on non-specular NeRF Synthetic dataset. Visual comparisons of rendering quality are demonstrated in Figure 3. Other NeRF-based baselines (Ref-NeRF, NMF and Zip-NeRF) struggle with the highly specular reflections (second and third rows),




Figure 5: Visualization on a reflective yet semi-transparent surface.

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Zip-NeRF	23.44	<u>0.626</u>	<b>0.257</b>
Ref-NeRF	24.06	0.589	0.355
Ours	<b>24.09</b>	<b>0.630</b>	<u>0.293</u>

Table 2: Comparison with SOTAs on real captured scenes from Ref-NeRF (Verbin et al. 2022).

while SDF-based baselines (ENVIDR and Ref-NeuS) fails to capture the intricate geometric details (first row). Our method consistently recovers high-fidelity rendering across all these scenes, indicating the robustness and effectiveness of our method. Per-scene metrics and additional visualizations are presented in supplementary material.

We further visualize and compare the recovered normal maps of NeRF-based methods in Figure 4. We can see that Ref-NeRF (Verbin et al. 2022) tends to generate semi-transparent surfaces with noisy normal maps, due to its inability to sufficiently resolve ambiguities in highly reflective surfaces. In contrast, NMF (Mai et al. 2023), which employs the orientation loss directly to density gradients, recovers surfaces that are opaque but often irregular. Our model demonstrates a robust capability to produce accurate surface normals, effectively mitigating such ambiguities.

Figure 5 highlights a reflective yet semi-transparent surface, which SDF-based baselines reconstruct as opaque. In comparison, our model effectively preserves the semi-transparency of reconstructed surface while still generating a plausible normal map.

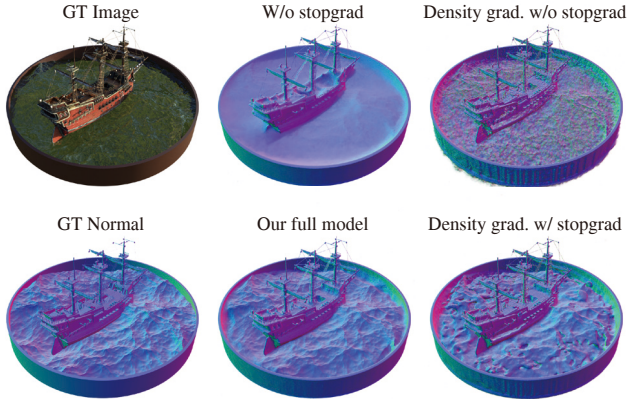


Figure 6: Ablation on normal estimation techniques, including the choice between transmittance gradient and density gradient, and the application of the stop-gradient warmup strategy.

Normal		Density		Metrics		
Trans.	Stopgrad	Softplus	Exp	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	<b>33.24</b>	<b>0.971</b>	<b>0.043</b>
$\checkmark$		$\checkmark$	$\checkmark$	31.64	0.950	0.072
	$\checkmark$	$\checkmark$	$\checkmark$	30.67	0.957	0.062
		$\checkmark$	$\checkmark$	27.82	0.925	0.108
$\checkmark$	$\checkmark$	$\checkmark$		<i>32.20</i>	<i>0.965</i>	<i>0.052</i>
$\checkmark$	$\checkmark$		$\checkmark$	<u>32.79</u>	<u>0.968</u>	<u>0.048</u>

Table 3: Quantitative comparisons for ablation runs on Glossy Synthetic dataset (Liu et al. 2023).

$\lambda_n$ Scheduling	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MAE $\downarrow$
Exp Warmup	<b>37.02</b>	<b>0.979</b>	<b>0.048</b>	<b>4.534</b>
Constant 0	36.40	0.976	0.054	6.143
Constant 0.01	<u>36.82</u>	<u>0.978</u>	<i>0.051</i>	<u>5.199</u>
Constant 0.1	<i>36.78</i>	<u>0.978</u>	<u>0.049</u>	<i>5.353</i>
Constant 1	36.35	0.976	0.052	5.767

Table 4: Ablation on the scheduling of  $\lambda_n$ , which controls the ratio of stop-gradient, conducted on Shiny Blender dataset (Verbin et al. 2022).

To explore our method’s robustness in real world environments, we conduct experiments using the real captured scenes from Ref-NeRF (Verbin et al. 2022). The quantitative results presented in Table 2 show that our method performs on par with existing methods.

### Ablation Studies

We conduct a series of ablation studies to evaluate the effect of our key components.

**Normal Estimation.** We evaluate both transmittance gradient and the conventional density gradient for normal estimation, each with and without the stop-gradient warmup strategy. Quantitative results in Table 3 demonstrate that the absence of either the transmittance gradient or the stop-

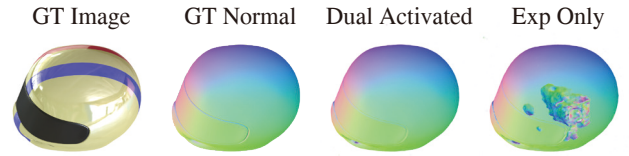


Figure 7: Normal map comparison between dual activated densities and single exp activated density at 10K iterations.

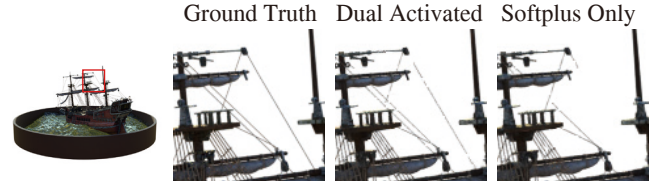


Figure 8: Visual comparison between dual activated densities and single softplus activated density.

gradient warmup strategy leads to a significant degradation in performance. Without the implementation of the stop-gradient warmup strategy, the randomly initialized predicted normals may oversmooth the surface reconstruction of water waves, as illustrated in Figure 6. Normal estimation based on the density gradients faces challenges in accurately reconstructing the water waves, even when employing the stop-gradient warmup strategy.

**Density Activation.** In addition to the dual activated densities, we also evaluate the performance using only an exp or softplus density activation. As presented in Table 3, employing a single activated density uniformly diminishes the metrics. As demonstrated in Figure 7, using a single exp activated density undermines the robustness of normal estimation. Additionally, employing a single softplus activated density hampers the reconstruction of thin geometric structures, as illustrated in Figure 8.

**Stop-Gradient Warmup.** The parameter  $\lambda_n$ , which controls the stop-gradient ratio in the predicted normal loss, follows an exponential warmup schedule in our design. We compare its performance with that of a constant  $\lambda_n$  applied throughout training. Table 4 shows that our warmup strategy consistently outperforms all tested constant values of  $\lambda_n$ .

## Conclusion

In this paper, we present a pipeline to enhance NeRF’s capability in reconstructing and rendering highly reflective scenes. The core of our approach is a transmittance-gradient-based normal estimation technique to improve the robustness and accuracy of surface normal estimation under conditions of ambiguous shape prediction. We also introduce dual activated densities to model objects with both smooth surfaces and sharp boundaries. Extensive experiments demonstrate that our approach quantitatively and qualitatively outperforms existing methods.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (NSFC) under Grant No. 62371009, and Beijing Natural Science Foundation under Grant No. L247029.

## References

- Anderson, M.; Motta, R.; Chandrasekar, S.; and Stokes, M. 1996. Proposal for a standard default color space for the internet—srgb. In *Color and imaging conference*, volume 4, 238–245. Society of Imaging Science and Technology.
- Barron, J. T.; Mildenhall, B.; Tancik, M.; Hedman, P.; Martin-Brualla, R.; and Srinivasan, P. P. 2021. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5855–5864.
- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2022. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5470–5479.
- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2023. Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 19697–19705.
- Bi, S.; Xu, Z.; Srinivasan, P.; Mildenhall, B.; Sunkavalli, K.; Hašan, M.; Hold-Geoffroy, Y.; Kriegman, D.; and Ramamoorthi, R. 2020. Neural Reflectance Fields for Appearance Acquisition. arXiv:2008.03824.
- Boss, M.; Braun, R.; Jampani, V.; Barron, J. T.; Liu, C.; and Lensch, H. 2021. Nerf: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12684–12694.
- Chen, A.; Xu, Z.; Geiger, A.; Yu, J.; and Su, H. 2022. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, 333–350. Springer.
- Chen, A.; Xu, Z.; Zhao, F.; Zhang, X.; Xiang, F.; Yu, J.; and Su, H. 2021. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14124–14133.
- Ge, W.; Hu, T.; Zhao, H.; Liu, S.; and Chen, Y.-C. 2023. Ref-NeuS: Ambiguity-Reduced Neural Implicit Surface Learning for Multi-View Reconstruction with Reflection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 4251–4260.
- Guo, Y.-C.; Kang, D.; Bao, L.; He, Y.; and Zhang, S.-H. 2022. Nerfren: Neural radiance fields with reflections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18409–18418.
- Jain, A.; Mildenhall, B.; Barron, J. T.; Abbeel, P.; and Poole, B. 2022. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 867–876.
- Jain, A.; Tancik, M.; and Abbeel, P. 2021. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5885–5894.
- Kopanas, G.; Leimkühler, T.; Rainer, G.; Jambon, C.; and Drettakis, G. 2022. Neural point catacaustics for novel-view synthesis of reflections. *ACM Transactions on Graphics (TOG)*, 41(6): 1–15.
- Kuang, Z.; Olszewski, K.; Chai, M.; Huang, Z.; Achlioptas, P.; and Tulyakov, S. 2022. Neroic: Neural rendering of objects from online image collections. *ACM Transactions on Graphics (TOG)*, 41(4): 1–12.
- Li, T.; Slavcheva, M.; Zollhoefer, M.; Green, S.; Lassner, C.; Kim, C.; Schmidt, T.; Lovegrove, S.; Goesele, M.; Newcombe, R.; et al. 2022. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5521–5531.
- Liang, R.; Chen, H.; Li, C.; Chen, F.; Panneer, S.; and Vijaykumar, N. 2023. ENVIDR: Implicit Differentiable Renderer with Neural Environment Lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 79–89.
- Liu, Y.; Wang, P.; Lin, C.; Long, X.; Wang, J.; Liu, L.; Komura, T.; and Wang, W. 2023. NeRO: Neural Geometry and BRDF Reconstruction of Reflective Objects from Multiview Images. In *SIGGRAPH*.
- Ma, L.; Agrawal, V.; Turki, H.; Kim, C.; Gao, C.; Sander, P.; Zollhöfer, M.; and Richardt, C. 2024. SpecNeRF: Gaussian Directional Encoding for Specular Reflections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 21188–21198.
- Mai, A.; Verbin, D.; Kuester, F.; and Fridovich-Keil, S. 2023. Neural Microfacet Fields for Inverse Rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 408–418.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Müller, T.; Evans, A.; Schied, C.; and Keller, A. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans. Graph.*, 41(4): 102:1–102:15.
- Park, K.; Sinha, U.; Barron, J. T.; Bouaziz, S.; Goldman, D. B.; Seitz, S. M.; and Martin-Brualla, R. 2021a. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5865–5874.
- Park, K.; Sinha, U.; Hedman, P.; Barron, J. T.; Bouaziz, S.; Goldman, D. B.; Martin-Brualla, R.; and Seitz, S. M. 2021b. HyperNeRF: a higher-dimensional representation for topologically varying neural radiance fields. *ACM Transactions on Graphics (TOG)*, 40(6): 1–12.
- Poole, B.; Jain, A.; Barron, J. T.; and Mildenhall, B. 2023. DreamFusion: Text-to-3D using 2D Diffusion. In *The Eleventh International Conference on Learning Representations*.

- Schwarz, K.; Liao, Y.; Niemeyer, M.; and Geiger, A. 2020. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems*, 33: 20154–20166.
- Srinivasan, P. P.; Deng, B.; Zhang, X.; Tancik, M.; Mildenhall, B.; and Barron, J. T. 2021. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7495–7504.
- Sun, C.; Sun, M.; and Chen, H.-T. 2022. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5459–5469.
- Tagliasacchi, A.; and Mildenhall, B. 2022. Volume Rendering Digest (for NeRF). arXiv:2209.02417.
- Tancik, M.; Weber, E.; Ng, E.; Li, R.; Yi, B.; Kerr, J.; Wang, T.; Kristoffersen, A.; Austin, J.; Salahi, K.; Ahuja, A.; McAllister, D.; and Kanazawa, A. 2023. Nerfstudio: A Modular Framework for Neural Radiance Field Development. In *ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH '23*.
- Verbin, D.; Hedman, P.; Mildenhall, B.; Zickler, T.; Barron, J. T.; and Srinivasan, P. P. 2022. Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5491–5500.
- Verbin, D.; Srinivasan, P. P.; Hedman, P.; Mildenhall, B.; Attal, B.; Szeliski, R.; and Barron, J. T. 2024. Nerf-casting: Improved view-dependent appearance with consistent reflections. In *SIGGRAPH Asia 2024 Conference Papers*, 1–10.
- Wang, P.; Liu, L.; Liu, Y.; Theobalt, C.; Komura, T.; and Wang, W. 2021. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. *Advances in Neural Information Processing Systems*, 34: 27171–27183.
- Yang, J.; Pavone, M.; and Wang, Y. 2023. FreeNeRF: Improving Few-shot Neural Rendering with Free Frequency Regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8254–8263.
- Yariv, L.; Gu, J.; Kasten, Y.; and Lipman, Y. 2021. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34: 4805–4815.
- Yin, Z.-X.; Qiu, J.; Cheng, M.-M.; and Ren, B. 2023. Multi-Space Neural Radiance Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12407–12416.
- Yu, A.; Ye, V.; Tancik, M.; and Kanazawa, A. 2021. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4578–4587.
- Zeng, J.; Bao, C.; Chen, R.; Dong, Z.; Zhang, G.; Bao, H.; and Cui, Z. 2023. Mirror-NeRF: Learning Neural Radiance Fields for Mirrors with Whitted-Style Ray Tracing. In *Proceedings of the 31th ACM International Conference on Multimedia*.
- Zhang, K.; Riegler, G.; Snavely, N.; and Koltun, V. 2020. NeRF++: Analyzing and Improving Neural Radiance Fields. arXiv:2010.07492.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhang, X.; Srinivasan, P. P.; Deng, B.; Debevec, P.; Freeman, W. T.; and Barron, J. T. 2021. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (ToG)*, 40(6): 1–18.