
NeuRodin: A Two-stage Framework for High-Fidelity Neural Surface Reconstruction

Yifan Wang^{1,2} Di Huang² Weicai Ye^{2,3,✉} Guofeng Zhang³ Wanli Ouyang² Tong He^{2,✉}
¹Shanghai Jiao Tong University ²Shanghai Artificial Intelligence Laboratory
³State Key Lab of CAD&CG, Zhejiang University
maikeyeweicai@gmail.com tonghe90@gmail.com



Figure 1: We present **NeuRodin**, a novel two-stage framework designed for high-fidelity neural surface reconstruction with intricate structures. Requiring only posed RGB captures as inputs, **NeuRodin** not only recovers large-scale areas but also accurately reconstructs fine-grained details.

Abstract

Signed Distance Function (SDF)-based volume rendering has demonstrated significant capabilities in surface reconstruction. Although promising, SDF-based methods often fail to capture detailed geometric structures, resulting in visible defects. By comparing SDF-based volume rendering to density-based volume rendering, we identify two main factors within the SDF-based approach that degrade surface quality: *SDF-to-density representation* and *geometric regularization*. These factors introduce challenges that hinder the optimization of the SDF field. To address these issues, we introduce **NeuRodin**, a novel two-stage neural surface reconstruction framework that not only achieves high-fidelity surface reconstruction but also retains the flexible optimization characteristics of density-based methods. **NeuRodin** incorporates innovative strategies that facilitate transformation of arbitrary topologies and reduce artifacts associated with density bias. Extensive evaluations on the Tanks and Temples and ScanNet++ datasets demonstrate the superiority of **NeuRodin**, showing strong reconstruction capabilities for both indoor and outdoor environments using solely posed RGB captures. Project website: <https://open3dvlab.github.io/NeuRodin/>

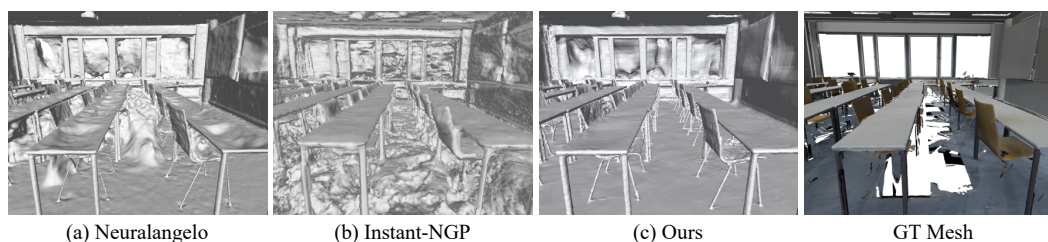


Figure 2: **Comparative analysis of SDF-based and density-based volume rendering methods.** (a) Neuralangelo [13] experiences difficulties with topological transformations, leading to incorrect surfaces. (b) Instant-NGP [18] approximates the correct surface positioning yet produces a noisy surface. (c) Our method achieves high-quality surfaces with fine details.

1 Introduction

3D surface reconstruction [34, 16, 35, 28, 8, 24, 26, 40, 39, 14, 12, 2, 38, 25] is a long-standing research topic in the field of computer vision. This process involves using images to recover the underlying 3D geometry, typically represented as meshes. These reconstructed meshes find diverse applications in various domains, including video games and augmented/virtual reality systems. In this paper, we specifically address the challenge of reconstructing 3D surfaces from posed RGB images.

Inspired by the density-based representation [15] for task of novel view synthesis, recent works for neural surface reconstruction commonly introduce signed distance functions (SDF) [28, 35] to recover high-quality geometry.

However, incorporating SDF to the density function is nontrivial and often fails to intricate geometric details. We illustrate this by comparing two methods for reconstructing the same scene: *Instant-NGP* [18], which employs density-based volume rendering, and *Neuralangelo* [13], which utilizes SDF-based volume rendering. Both methods use a similar multi-resolution hash table representation. The mesh is produced by is by TSDF-fusion [4] for *Instant-NGP*. As illustrated in Figure 2, *Instant-NGP* reconstructs surfaces with accurate localization albeit with a certain roughness, while *Neuralangelo* produces smoother surfaces yet encounters issues in correctly positioning portions of the surface. This disparity underscores the limitations and highlights the need for improved modeling capabilities in SDF-based surface reconstruction methods.

Why do SDF-based surface reconstruction methods face challenges in accurately capturing intricate geometric details, and how can these methods be improved? In this paper, we thoroughly analyze the reconstruction pipeline and identify two primary factors in current SDF-based pipelines that contribute to suboptimal surface reconstruction:

- *SDF-to-density conversion:* SDF-based volume rendering requires a conversion function to relate the SDF field with the density field. Existing methods use a conversion function that assigns uniform density across the same level sets, which restricts the representation of arbitrary non-negative density values. Additionally, there is no assurance that the geometric representation within a volume rendering framework will align perfectly with the implicit surface. This misalignment often results in accurate visual renderings on incorrect surfaces, due to inherent biases.
- *Geometric regularization:* Regularization constraints imposed on the implicit surface can limit topological changes during optimization. These constraints often introduce biases, complicating the convergence of the model and hindering its ability to accurately represent complex geometries.

To tackle these challenges, we introduce NeuRodin (Figure 1), a high-fidelity 3D surface reconstruction method that innovatively overcomes the limitations previously outlined. Firstly, we refine the SDF-to-density conversion by transitioning from a global scale parameter to a local adaptive parameter. Unlike previous methods that enforce the same densities for points with identical SDF values, our approach enhances the flexibility and effectiveness of the SDF function by allowing adaptive density values. Secondly, we implement a novel loss function designed to align the maximum probability distance with the zero-level set in volume rendering, improving the alignment of geometric representations. Additionally, We incorporates above innovations within a two-stage optimization framework to tackle the over-regularization imposed by geometric constraints. Initially, we employ a coarse optimization stage in which the SDF field operates similarly to a density field,

exhibiting minimal influence from topological transformations. Subsequently, a refinement stage is conducted to achieve a surface with enhanced smoothness. We also introduce the stochastic-step numerical gradient estimation technique to maintain a natural zero level set for the coarse stage. With the design described, our method enables high-fidelity surface reconstruction suited to both large-scale and intricate geometries.

We conducted extensive experiments on the Tanks and Temples dataset [11] and the ScanNet++ dataset [41], where our method demonstrated superior performance over the previous state-of-the-art in both indoor and outdoor environments. Notably, recognizing the lack of an established benchmark for ScanNet++, we executed comparative analyses using six baseline methods and established a new benchmark for ScanNet++ reconstruction, significantly enriching the community resources and setting a foundation for future research. In comparative tests, our model outperformed Neuralangelo on the Tanks and Temples training set, delivering superior results with only 1/8 the parameters of the comparative model. Our approach excels in optimizing complex topological structures and preserving intricate details, enabling high-fidelity, fine-grained surface reconstruction.

2 Related Work

Multi-view 3D reconstruction. In traditional 3D surface reconstruction, methods based on Multi-View Stereo (MVS) have long been prevalent, serving as a foundational approach for mining sparse geometric data from multiple views and generate detailed 3D models by comparing and analyzing the disparities across multiple camera perspectives. The traditional MVS methods [24], while effective in texture-rich domains, often stumble upon the challenge of processing ambiguous observations. The point clouds produced by traditional MVS methods suffer from noise, undermining the reliability of the surface triangle meshes reconstructed from these point clouds. For learning-based MVS methods [34, 44, 33, 32], the generated point clouds are still plagued by noise, leading to consistently incomplete reconstructions.

Neural surface reconstruction. NeRF [16, 37, 10, 17] pioneers the use of neural network to represent neural radiance fields for novel view synthesis and optimizes these scenes through differentiable volume rendering. Following NeRF, subsequent research has combined implicit surfaces with differentiable volume rendering [19, 28, 35]. These methods typically represent implicit surfaces as SDF and use the zero-level set of SDFs to describe geometry, achieving high-quality reconstruction on individual objects. Various improvements have been made based on this foundation, including the incorporation of different positional encoding to enhance representational capabilities [22, 29, 47, 48] and the introduction of additional priors to deal with surfaces that exhibit specular highlights or have low textures [43, 27]. Several studies refine the modeling of SDF-to-density conversion [46, 31] to address bias issues in density. Meanwhile, other works employ patch-match techniques to improve multi-view consistency [5, 7]. Neuralangelo [13] enhances the network’s representational capability by introducing hash encoding. Additionally, it proposes numerical gradients and coarse-to-fine optimization strategies to enhance the quality of surface reconstruction.

3 Study on SDF-based Volume Rendering

3.1 Preliminary

Density-based volume rendering Density-based volume rendering methods model a 3D scene as a volume density field. Given a camera position \mathbf{o} and view direction \mathbf{d} , the ray emitted from \mathbf{o} in direction \mathbf{d} is denoted as $\{\mathbf{r}(t) = \mathbf{o} + t\mathbf{d} | t > 0\}$. A set of n points is sampled along this ray. The predicted density $\sigma(\mathbf{r}(t))$ and geometry features $\mathbf{z}(\mathbf{r}(t))$ of the point $\mathbf{r}(t)$ are obtained from a geometry network ϕ_{geo} . the density is parameterized by an activation function, such as ReLU, softplus, or the exp function, prior to being output by the network.

A color network ϕ_{color} predicts the color $\mathbf{c}(\mathbf{r}(t), \mathbf{d})$, taking as inputs the geometry feature $\mathbf{z}(\mathbf{r}(t))$, and the viewing direction \mathbf{d} . The rendered color of this ray can be calculated as:

$$\hat{\mathbf{C}}(\mathbf{r}) = \int_0^{+\infty} T(t)\sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t)), \quad T(t) = \exp\left(-\int_0^t \sigma(\mathbf{r}(u))du\right). \quad (1)$$

The networks are trained to minimize a color loss $\mathcal{L}_{\text{color}}$ that quantifies the difference between the rendered colors and the ground truth colors:

$$\mathcal{L}_{\text{color}} = \frac{1}{m} \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{C}(\mathbf{r}) - C(\mathbf{r})\|_1, \quad (2)$$

with \mathcal{R} representing the set of m rays in each training batch and $C(\mathbf{r})$ being the ground-truth color of ray \mathbf{r} .

SDF-based volume rendering SDF-based volume rendering methods, such as NeuS [28] and VolSDF [35], combine volume rendering with an SDF representation. Unlike density-based methods, the geometry is represented as the zero level set of the SDF. The predicted SDF $f(\mathbf{r}(t))$ and geometry feature $\mathbf{z}(\mathbf{r}(t))$ at the point $\mathbf{r}(t)$ are obtained from the geometry network ϕ_{geo} . Then, the SDF $f(\mathbf{r}(t))$ is transformed into density $\sigma(\mathbf{r}(t))$ by a predefined function Ψ and a global scale s . For instance, VolSDF defines the density as the scaled cumulative distribution function of the negative SDF:

$$\sigma(\mathbf{r}(t)) = \Psi_s(f(\mathbf{r}(t))) = \begin{cases} \frac{1}{2s} \exp\left(\frac{-f(\mathbf{r}(t))}{s}\right) & \text{if } f(\mathbf{r}(t)) \geq 0, \\ \frac{1}{s} \left(1 - \frac{1}{2} \exp\left(\frac{f(\mathbf{r}(t))}{s}\right)\right) & \text{if } f(\mathbf{r}(t)) < 0. \end{cases} \quad (3)$$

Moreover, to encourage the SDF to have a unit norm gradient, the Eikonal loss [9] is often employed:

$$\mathcal{L}_{\text{eik}} = \frac{1}{mn} \sum_{\mathbf{r}, t} (\|\nabla f(\mathbf{r}(t))\| - 1)^2. \quad (4)$$

Incorporating this loss not only helps to avoid suboptimal solutions at the zero level set but also promotes smoothness. The networks are trained under the supervision of both the color loss $\mathcal{L}_{\text{color}}$ and the Eikonal loss \mathcal{L}_{eik} .

3.2 Challenges in Previous SDF-Based Volume Rendering

State-of-the-art SDF-based volume rendering techniques frequently fail to reconstruct surfaces with accuracy in scenarios where density-based methods manage to renders realistic novel views. This disparity highlights the inherent limitations of SDF-based volume rendering approaches. To further elucidate these issues, we explore the fundamental distinctions between SDF-based and density-based volume rendering. Please refer to the appendix for a more in-depth analysis.

Unsuitable assumption for SDF-to-density conversion. SDF-based volume rendering methods typically employ a predefined function Ψ and a global scale parameter s to convert SDF values into density, as described by Equation (3). These methods often result in uniform density values for points sharing identical SDF values. Such a global scaling mechanism restricts the representation capability of the density field derived from the SDF field. Intuitively, previous top-performing methods for novel view synthesis can generate arbitrary non-negative density values within $(0, +\infty)$. In contrast, incorporating SDF representation with a global scaling factor for surface reconstruction can only result in density values within $(0, \frac{1}{s}]$.

Bias of the density. When applying an Eikonal constraint or any form of smooth regularization to an SDF, the geometric representation within the rendering framework must align with that of the SDF. Unfortunately, current SDF-based methods often fail to ensure this alignment, particularly at larger-scale parameters s as explored mathematically in [46] to analyze this issue. Recent studies [31, 46, 3] have attempted to tackle this issue, proposing designs for SDF to density conversion that aim to minimize bias. Despite these advancements, these solutions still exhibit inherent biases. Additionally, the introduction of geometric regularization often exacerbates this bias, complicating model convergence and resulting in the creation of inaccurate surfaces. A more detailed analysis of this issue is provided in Section 4.2 and Appendix B to Appendix D.

Over-regularization of Geometry. To maintain a high-quality surface, previous methods often introduce geometric constraints, such as Eikonal loss or smoothness constraints. However, these global constraints result in excessive smoothing across all regions, both flat and intricate, leading to a loss of fine details. Moreover, in the framework of SDF-based volume rendering, the prediction of color typically necessitates being conditioned on normals following IDR [36], a characteristic that distinctly sets it apart from the density-based volume rendering approach. When optimizing the

color conditioned on the normal and explicitly constraining the SDF by geometric regularization, the optimization process restricts the topological structure. Please refer to the Appendix F for the impact on the normal condition.

4 Method

4.1 Uniform SDF, Diverse Densities

To deal with the representation limitations of SDF-transformed density, instead of using a global scale s for the transformation from SDF to density, we have employed a strategy akin to that of [30], which utilizes a non-linear mapping to obtain the unique scale s associated with a given point $\mathbf{r}(t)$. More precisely, it is defined as follows:

$$(f(\mathbf{r}(t)), s(\mathbf{r}(t)), \mathbf{z}(\mathbf{r}(t))) = \phi_{\text{geo}}(\mathbf{r}(t)), \quad \sigma(\mathbf{r}(t)) = \Psi_{s(\mathbf{r}(t))}(f(\mathbf{r}(t))). \quad (5)$$

With this particular design, the density is not identical within the same SDF level set and can achieve any non-negative value through the continuous representation that maps an input coordinate to its corresponding scale.

This approach ensures that densities within the same SDF level set are no longer uniformly identical. Instead, they can vary, achieving any non-negative value through a continuous representation that maps input coordinates to their corresponding scales. This design greatly enhances the flexibility and accuracy of our density modeling, enabling more realistic and detailed reconstructions. More detailed analysis regarding the local scale can be found in Appendix A.

4.2 Explicit Bias Correction

The issue of bias represents a critical concern frequently addressed within SDF-based volume rendering. As demonstrated in Figure 3, it is necessary to align the geometric representation under the volume rendering framework with that of the implicit surface. For the volume rendering framework, the most intuitive way to represent geometry is through the *rendered distance*:

$$\hat{D}_{\text{rendered}}(\mathbf{r}) = \int_0^{+\infty} T(t)\sigma(\mathbf{r}(t))t dt. \quad (6)$$

We can also consider the position where $w(t)$ is maximized — that is, the probability that the light ray arrives and collides is the greatest, or in other words, the location that contributes the most to the color — as the geometric representation within the volume rendering framework:

$$\hat{D}_{\text{prob}}(\mathbf{r}) = \arg \max_{t \in (0, +\infty)} w(t) = \arg \max_{t \in (0, +\infty)} T(t)\sigma(\mathbf{r}(t)). \quad (7)$$

We shall refer to $\hat{D}_{\text{prob}}(\mathbf{r})$ as the *maximum probability distance*. For an implicit surface, the *zero level set* offers a direct geometric representation. In an ideal scenario, irrespective of whether convergence has been achieved, the geometric representations of volume rendering (i.e., *rendered distance* and *maximum probability distance*) and the geometric representation of the implicit surface (i.e., *zero level set*) should be aligned, as illustrated in Figure 3 (a). However, in the practical optimization process, conflicts such as those depicted in Figure 3 (b) may arise, leading to misalignment between the two representations.

Multiple past methods have broached this topic, offering various solutions. For example, in TUVR [46], an unbiased model is proposed and it is mathematically proven that the zero-crossing point of the SDF is at a local maximum of rendering weight. However, bias still exists. The experimental analysis is presented in Section 5.3. In the work

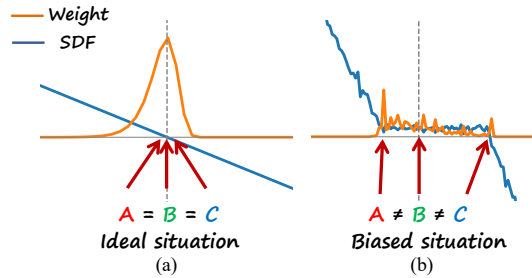


Figure 3: **Visualization of the bias of the density.** (a) An ideal scenario where the geometry of the volume rendering scheme (*A: maximum probability distance* and *B: rendered distance*) aligns precisely with the geometry of the implicit surface (*C: zero level set*). (b) A biased scenario showcasing misalignment.

of [3], a penalty is imposed on the rendered distance to ensure that its SDF value is equal to zero. However, it is important to note that the rendered distance is subject to significant distortion due to existing biases, particularly in the early stages of convergence as shown in Figure 3 (b). We offer further analysis from Appendix B and Appendix D.

We propose an explicit bias correction in which we opt to deliberately align the maximum probability distance with the zero level set. Specifically, we define

$$\mathcal{L}_{\text{bias}} = \frac{1}{m} \sum_{\mathbf{r} \in \mathcal{R}} \max(f(\mathbf{r}(t^* + \epsilon_{\text{bias}})), 0), \quad t^* = \arg \max_{t \in (0, +\infty)} T(t) \sigma(\mathbf{r}(t)), \quad (8)$$

where ϵ_{bias} is a bias correction factor. The loss function is designed to penalize the positive portion of $f(\mathbf{r}(t^* + \epsilon_{\text{bias}}))$, which encourages the SDF to take on negative values after the maximum probability distance. We have shown in Appendix C that manually scheduling the lower bound of the local scale, coupled with the penalty after the point of maximum probability distance, can effectively alleviate bias issues. During the experiment, we approximate t^* by directly using the sampled point with the largest $w(t)$, which, despite introducing a certain degree of deviation, does not affect the overall effectiveness. Please refer to Appendix C for details on the design.

4.3 Two-Stage Optimization to Tackle Geometry Over-Regularization

Previous methods often produce incorrect surfaces due to over-regularization of geometry as shown in Figure 2 (a). However, we have discovered that methods based on density are not constrained by changes in topology as shown in Figure 2 (b), prompting us to question whether the SDF field can be first optimized as freely as density field, then refine to a smooth surface by geometry regularization. We now propose a novel two-stage optimization approach. This approach allows the optimization process to initially mimic density-based behavior in the first stage and subsequently refines to a smooth surface in the second stage.

For the first stage, our objective is to tackle the over-regularization issue. An intuitive solution might be to eliminate or downweight any geometric constraints and avoid conditioning the color on the predicted normal, but this approach often results in an unnatural zero level set [9, 28, 35, 36]. We experimentally validated in the Appendix H.6.

We have identified a simple but effective method to preserve the natural level sets of large-scale structures while allowing the formation of complex structures to be unimpeded by geometric regularization. Instead of applying geometric regularization directly to the gradient $\nabla f(\mathbf{r}(t))$, we elect to impose them upon an estimated gradient $\hat{\nabla} f(\mathbf{r}(t))$, to which we introduce uncertainty through a specific design. Specifically, the x -component of the estimated gradient is

$$\hat{\nabla}_x f(\mathbf{r}(t)) = \frac{f(\mathbf{r}(t) + \epsilon_x) - f(\mathbf{r}(t) - \epsilon_x)}{2\epsilon}, \quad \text{where } \epsilon_x = (\epsilon, 0, 0) \text{ and } \epsilon \sim U(0, \epsilon_{\text{max}}). \quad (9)$$

The gradient is calculated through finite differences similar to those described in [13, 21]. However, the step size for gradient estimation at each iteration is stochastically sampled from a uniform distribution in the range of $(0, \epsilon_{\text{max}})$.

Using this technique, we observed that estimated larger-scale normals have smaller variance, while fine details exhibit larger variance, as depicted in Figure 4. This introduces uncertainty in geometric regularization, ensuring stability for large features and flexibility for complex details. We provide further explanation in Appendix E.

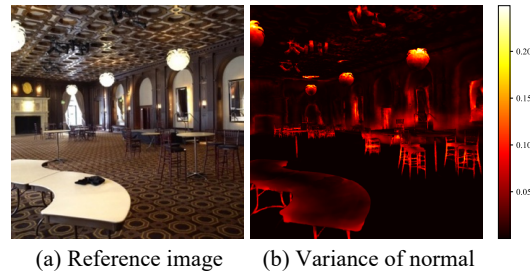


Figure 4: The heatmap for the variance of the normal predicted using random step.

During the initial stage, our goal is to reconstruct the approximate, coarse structure of the 3D content. This is primarily addressed by tackling the issues of over-regularization and the bias in density estimation that we previously mentioned. We employ the stochastic-step numerical gradient estimation along with the explicit bias correction to address the initial reconstruction of the coarse

Metric	F-Score \uparrow								
Scene	NeuralWarp	COLMAP	NeuS	Geo-NeuS	NeuS-NGP	MonoSDF	Neuralangelo*	Neuralangelo	Ours
Barn	0.22	0.55	0.29	0.33	0.46	0.49	<u>0.61</u>	0.70	0.70
Caterpillar	0.18	0.01	0.29	0.26	0.32	0.31	<u>0.34</u>	0.36	0.36
Courthouse	0.08	0.11	0.17	0.12	0.08	0.12	0.13	0.28	<u>0.21</u>
Ignatius	0.02	0.22	0.83	0.72	0.81	0.78	0.82	0.89	<u>0.87</u>
Meetingroom	0.08	0.19	0.24	0.20	0.08	0.23	0.22	<u>0.32</u>	0.43
Truck	0.35	0.19	0.45	0.45	0.44	0.42	0.45	0.48	<u>0.47</u>
Mean	0.15	0.21	0.38	0.35	0.37	0.39	0.43	<u>0.50</u>	0.51

Table 1: **Quantitative evaluation of our method on the Tanks and Temples training subset.** The **best** performance and the **second-best** outcomes are highlighted for easy reference. Note that the hash grid parameters used in our method is the same as Neuralangelo*, which possesses 2^{19} hash entries per resolution.

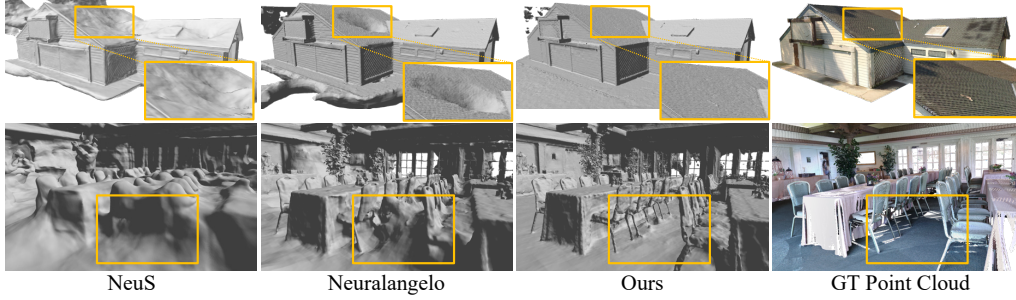


Figure 5: **Quantitative comparison on the training subset of Tanks and Temples dataset.**

structure of the 3D content. Additionally, we utilize VolSDF’s SDF-to-density conversion, with the local scale modeling delineated in Equation 5, to facilitate this primary formulation. Consequently, the training loss at this stage is formulated as:

$$\mathcal{L}_{\text{coarse}} = \mathcal{L}_{\text{color}} + \lambda_{\text{eik}} \mathcal{L}_{\text{eik}}(\hat{\nabla} f) + \lambda_{\text{bias}} \mathcal{L}_{\text{bias}}. \quad (10)$$

During the refinement stage, we discontinue the use of the estimated gradients given that the fundamental 3D content has been initially restored and the issue of over-regularization no longer presents a concern. In a similar vein, we move past the explicit bias correction, as the significant surface errors induced by bias were initially addressed in the first stage. We incorporate a standard Eikonal loss alongside a smoothness constraint from PermutoSDF [22] to enforce local smoothness:

$$\mathcal{L}_{\text{smooth}} = \frac{1}{mn} \sum_{\mathbf{r}, t} (\mathbf{n}(\mathbf{r}(t)) \cdot \mathbf{n}(\mathbf{r}(t) + \epsilon_s \boldsymbol{\eta}(\mathbf{r}(t))) - 1)^2, \quad (11)$$

where $\boldsymbol{\eta}(\mathbf{r}(t)) = \mathbf{n}(\mathbf{r}(t)) \times \boldsymbol{\tau}$ and $\boldsymbol{\tau}$ is a random unit vector. Furthermore, as the model nears convergence at this stage, we adopt the SDF-to-density conversion method proposed by TUVr [46], which ensures minimal bias and preserves fine object details. The loss function employed in the second phase is defined as follows:

$$\mathcal{L}_{\text{fine}} = \mathcal{L}_{\text{color}} + \lambda_{\text{eik}} \mathcal{L}_{\text{eik}}(\nabla f) + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}}. \quad (12)$$

5 Experiments

Experimental setup. We carry out experimental evaluations on two benchmark datasets: Tanks and Temples [11] and ScanNet++ [41]. We include several baselines for comparisons: VolSDF [35], NeuralWarp [5], COLMAP [23], NeuS [28], Geo-NeuS [7], Neuralangelo [13] and MonoSDF [43]. We extract mesh through marching cube algorithm with a resolution of 2048 applied across all scenes and report the F-score for surface evaluation. More details are provided in the supplementary materials. Please refer to Appendix H for additional experimental results.

5.1 Tanks and Temples

NeuRodin outperforms previous state-of-the-art methods in terms of the average F-score. Owing to our explicit bias correction technique, the barn’s roof maintains its structural integrity without

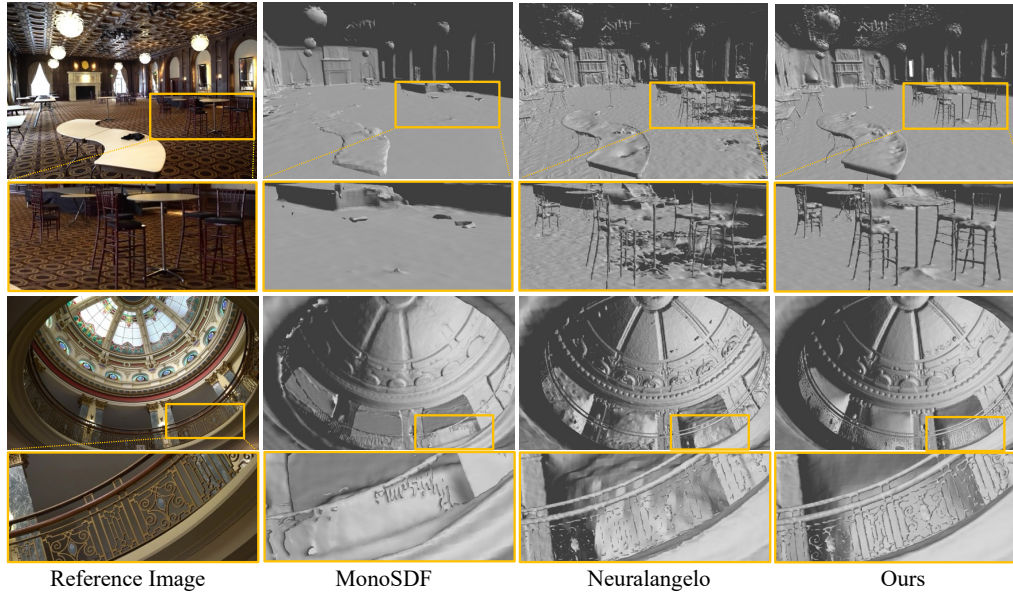


Figure 6: **Quantitative evaluation of our method on the Tanks and Temples advance subset.**

collapsing as shown at the top of Figure 5, a marked improvement over other methods which often fail to prevent such collapse. Thanks to our two-stage optimization approach, we effectively mitigate the issue of excessive geometric regularization as shown at the bottom of Figure 5. Consequently, NeuRodin achieves a more detailed surface representation than Neuralangelo with 1/8 fewer parameters.

It is evident that our method outperforms previous work on the Tanks and Temples advance subset as shown in Table 2. In the comparison in Figure 6, we demonstrate enhanced accuracy and completeness when reconstructing large-scale surfaces, alongside capturing more fine-grained details compared to Neuralangelo.

Benefiting from our explicit bias correction in the first stage and the TUVr modeling employed in the second stage, fine structures are restored to the zero level set in the initial phase and maintain a sufficiently small bias in the subsequent phase, thus enabling the refinement of high-quality surfaces.

Method	VolSDF	MonoSDF	Neuralangelo*	Neuralangelo	Ours
Mean	6.72	20.89	27.14	26.28	28.84

Table 2: **Quantitative evaluation of our method versus prior work on the Tanks and Temples advance subset.** The **best** performance and the second-best outcomes are highlighted for easy reference.

5.2 ScanNet++ Benchmark

Since no public results are available for the ScanNet++ dataset, we randomly selected 8 scenes to construct a benchmark. For more details and results on our ScanNet++ benchmark, please refer to the supplementary materials. Quantitative results are shown in Table 3. We surpassed the methods we compared against in most scenes and achieved comparable results to those with prior knowledge in terms of F-score. We provide more visual result on ScanNet++ dataset in the supplementary.

F-Score \uparrow	Without Prior					With Prior	
method	NeuS	VolSDF	Neuralangelo*	Neuralangelo	Ours	MonoSDF-MLP	MonoSDF-Grid
Mean	0.455	0.391	0.507	0.564	<u>0.638</u>	0.439	0.642

Table 3: **Quantitative evaluation of our method versus prior work on the ScanNet++ dataset.** The **best** performance and the second-best outcomes are highlighted for easy reference.

5.3 Analysis

Ablation Study. To validate the efficacy of the proposed techniques, we performed an ablation study on scene *Meetingroom* from Tanks and Temples dataset. As illustrated in Figure 7 (a), applying a global scale for SDF-to-density conversion results in inaccurate surfaces, primarily due to the

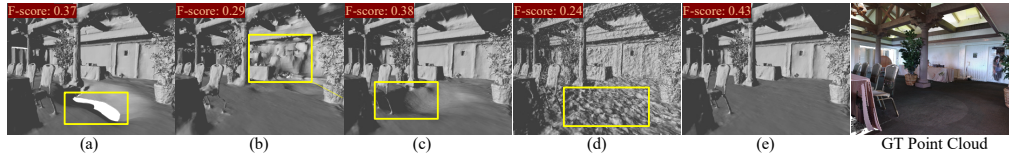


Figure 7: **Ablation results.** (a) Without local scale for SDF-to-density conversion. (b) Without stochastic-step numerical gradient estimation. (c) Without explicit bias correction. (d) Without stage-two refinement. (e) Full model.

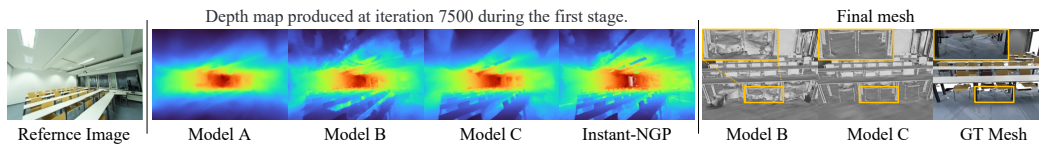


Figure 8: **Analysis on stochastic-step numerical gradient estimation.** We visualize the produced depth maps at iteration 7500 of the first stage. Model A: Change stochastic-step numerical gradient estimation to analytical gradient. Model B: Change stochastic-step numerical gradient estimation to progressive numerical gradient estimation from Neuralangelo. Model C: Ours.

assumption of uniform density across the same level sets. This assumption leads to the convergence of surfaces with subtler textures to incorrect locations. Figure 7 (b) demonstrates that the absence of stochastic-step numerical gradient estimation hinders the model’s ability to form arbitrary topologies, leading to incorrect surfaces. The incorrect ground collapsing depicted in Figure 7 (c) is due to the bias in density, creating inconsistencies as shown in Figure 3(c). Without the application of our proposed explicit bias correction, this bias issue causes visibly incorrect surfaces. Figure 7 (d) presents the outcomes of optimization conducted in a single phase; under stochastic-step numerical gradient estimation, the Eikonal loss’s preference for smoothness is somewhat compromised, resulting in a rougher surface finish. Finally, Figure 7 (e) showcases our full model, which achieves high-fidelity smooth surfaces while preserving most details. We conducted additional ablation experiments in Appendix H.6.

Analysis on Stochastic-step Numerical Gradient Estimation. We further demonstrate experimentally the impact of our stochastic-step numerical gradient estimation on the optimization process, as illustrated in Figure 8. The depth maps produced indicate that Model A is severely constrained by geometric regularization, making it difficult to alter its topological structure during optimization. Model B employs a progressive numerical gradient estimation technique that, even after 7500 steps, does not yield an accurate depth map. However, utilizing our stochastic-step numerical gradient estimation, we achieve an approximately accurate scene geometry in as few as 7500 steps.

At this stage, the depth map of Model C is similar to that of Instant-NGP, yet ours displays a more natural and smooth depth profile. This suggests that our approach, similar to instant-NGP, is capable of freely altering topological shapes for optimization, while also maintaining a natural zero-level set surface.

Furthermore, final mesh of Model B still manifests inaccurate floor collapses, whereas mesh of model C, with our stochastic-step numerical gradient estimation, maintains correct and smooth floors. This is attributed to the fact that the floor, being a large scale, allows our stochasticity to continuously apply the Eikonal constraint on large-scale areas. This results in a natural zero level set in these vast regions. However, for progressive steps, once the step size is reduced beyond a certain point, it no longer imposes the Eikonal constraint on large-scale surfaces, resulting in unnatural zero level sets.

Analysis on Explicit Bias Correction. To substantiate the versatility and effectiveness of our explicit bias correction approach, we further conducted experiments to verify its potential as a plug-and-play correction method independent of our full model. We implemented this correction technique on our coarse stage and tested it across various renderers, including NeuS, VolSDF, and TUVr. These experiments were designed to evaluate the adaptability and efficacy of our bias correction in diverse SDF-to-density modelings.

As depicted in Figure 9, the ceiling of the room displays a pronounced bias issue that leads to a collapse. However, with the application of our proposed explicit bias correction approach, the issue of ceiling collapse is significantly ameliorated.

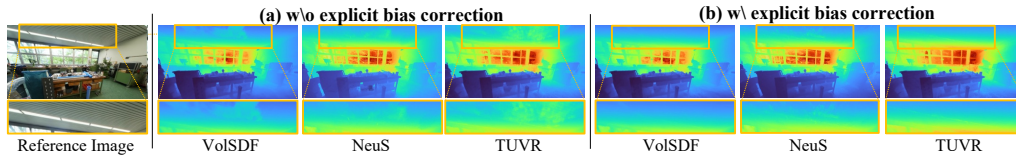


Figure 9: **Analysis on explicit bias correction.** We visualize the produced depth maps with and without our explicit bias correction for different SDF-to-density modelings at the first stage.

6 Conclusion

This paper proposes NeuRodin, a two-stage framework for high-fidelity neural surface reconstruction with intricate details. It introduces key designs to tackle SDF-based rendering challenges, notably the local scale adjustment for SDF-to-density conversion, which enables any non-negative value to be achieved, facilitating accurate density derivation from SDF. Additionally, an explicit bias correction method is employed to ensure the geometry of the volume rendering scheme coherently aligns with that of the implicit surface, thereby preventing the emergence of incorrect surfaces. Finally, a two-stage optimization strategy effectively resolves the issue of over-regularization imposed by geometric constraints. Comprehensive experiments demonstrate that NeuRodin simultaneously delivers superior quality.

Acknowledgments and Disclosure of Funding

This work is funded in part by the National Key R&D Program of China No.2022ZD0160102, and Shanghai Artificial Intelligence Laboratory.

References

- [1] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022.
- [2] Danpeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin Liu, Hujun Bao, and Guofeng Zhang. Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *arXiv preprint arXiv:2406.06521*, 2024.
- [3] Decai Chen, Peng Zhang, Ingo Feldmann, Oliver Schreer, and Peter Eisert. Recovering fine details for neural implicit surface reconstruction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4330–4339, 2023.
- [4] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996.
- [5] François Darmon, Bénédicte Bascle, Jean-Clément Devaux, Pascal Monasse, and Mathieu Aubry. Improving neural implicit surfaces geometry with patch warping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6260–6269, 2022.
- [6] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10786–10796, 2021.
- [7] Qiancheng Fu, Qingshan Xu, Yew Soon Ong, and Wenbing Tao. Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 35:3403–3416, 2022.
- [8] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *Proceedings of the IEEE international conference on computer vision*, pages 873–881, 2015.

- [9] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of the 37th International Conference on Machine Learning*, pages 3789–3799, 2020.
- [10] Chenxi Huang, Yuenan Hou, Weicai Ye, Di Huang, Xiaoshui Huang, Binbin Lin, Deng Cai, and Wanli Ouyang. Nerf-det++: Incorporating semantic cues and perspective-aware depth supervision for indoor multi-view 3d detection. *arXiv preprint arXiv:2402.14464*, 2024.
- [11] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017.
- [12] Hai Li, Weicai Ye, Guofeng Zhang, Sanyuan Zhang, and Hujun Bao. Saliency guided subdivision for single-view mesh reconstruction. In *2020 International Conference on 3D Vision (3DV)*, pages 1098–1107. IEEE, 2020.
- [13] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2023.
- [14] Xiangyu Liu, Weicai Ye, Chaoran Tian, Zhaopeng Cui, Hujun Bao, and Guofeng Zhang. Coxgraph: multi-robot collaborative, globally consistent, online dense reconstruction system. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8722–8728. IEEE, 2021.
- [15] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021.
- [16] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [17] Yuhang Ming, Weicai Ye, and Andrew Calway. idf-slam: End-to-end rgb-d slam with neural implicit mapping and deep feature tracking. *arXiv preprint arXiv:2209.07919*, 2022.
- [18] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022.
- [19] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5589–5599, 2021.
- [20] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [21] Albert Pumarola, Artsiom Sanakoyeu, Lior Yariv, Ali Thabet, and Yaron Lipman. Visco grids: Surface reconstruction with viscosity and coarea grids. *Advances in Neural Information Processing Systems*, 35:18060–18071, 2022.
- [22] Radu Alexandru Rosu and Sven Behnke. Permutosdf: Fast multi-view reconstruction with implicit surfaces using permutohedral lattices. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8466–8475, 2023.
- [23] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016.

- [24] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, pages 501–518. Springer, 2016.
- [25] Ziyu Tang, Weicai Ye, Yifan Wang, Di Huang, Hujun Bao, Tong He, and Guofeng Zhang. Nd-sdf: Learning normal deflection fields for high-fidelity indoor reconstruction. *arXiv preprint arXiv:2408.12598*, 2024.
- [26] Hoang-Hiep Vu, Patrick Labatut, Jean-Philippe Pons, and Renaud Keriven. High accuracy and visibility-consistent dense multiview stereo. *IEEE transactions on pattern analysis and machine intelligence*, 34(5):889–901, 2011.
- [27] Jiepeng Wang, Peng Wang, Xiaoxiao Long, Christian Theobalt, Taku Komura, Lingjie Liu, and Wenping Wang. Neuris: Neural reconstruction of indoor scenes using normal priors. In *European Conference on Computer Vision*, pages 139–155. Springer, 2022.
- [28] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021.
- [29] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3295–3306, 2023.
- [30] Zian Wang, Tianchang Shen, Merlin Nimier-David, Nicholas Sharp, Jun Gao, Alexander Keller, Sanja Fidler, Thomas Müller, and Zan Gojcic. Adaptive shells for efficient neural radiance field rendering. *ACM Transactions on Graphics (TOG)*, 42(6):1–15, 2023.
- [31] Yuting Xiao, Jingwei Xu, Zehao Yu, and Shenghua Gao. Debsdf: Delving into the details and bias of neural indoor scene reconstruction. *arXiv preprint arXiv:2308.15536*, 2023.
- [32] Qingshan Xu, Weihang Kong, Wenbing Tao, and Marc Pollefeys. Multi-scale geometric consistency guided and planar prior assisted multi-view stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4945–4963, 2022.
- [33] Qingshan Xu and Wenbing Tao. Multi-scale geometric consistency guided multi-view stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5483–5492, 2019.
- [34] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mysnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, pages 767–783, 2018.
- [35] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021.
- [36] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33:2492–2502, 2020.
- [37] Weicai Ye, Shuo Chen, Chong Bao, Hujun Bao, Marc Pollefeys, Zhaopeng Cui, and Guofeng Zhang. IntrinsicNeRF: Learning Intrinsic Neural Radiance Fields for Editable Novel View Synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [38] Weicai Ye, Xinyu Chen, Ruohao Zhan, Di Huang, Xiaoshui Huang, Haoyi Zhu, Hujun Bao, Wanli Ouyang, Tong He, and Guofeng Zhang. Datap-sfm: Dynamic-aware tracking any point for robust dense structure from motion in the wild. *arXiv preprint arXiv:2411.13291*, 2024.
- [39] Weicai Ye, Xinyue Lan, Shuo Chen, Yuhang Ming, Xingyuan Yu, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. Pvo: Panoptic visual odometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9579–9589, June 2023.

- [40] Weicai Ye, Xingyuan Yu, Xinyue Lan, Yuhang Ming, Jinyu Li, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. Deflow slam: Self-supervised scene motion decomposition for dynamic dense slam. *arXiv preprint arXiv:2207.08794*, 2022.
- [41] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22, 2023.
- [42] Zehao Yu, Anpei Chen, Bozidar Antic, Songyou Peng Peng, Apratim Bhattacharyya, Michael Niemeyer, Siyu Tang, Torsten Sattler, and Andreas Geiger. Sdfstudio: A unified framework for surface reconstruction, 2022.
- [43] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in neural information processing systems*, 35:25018–25032, 2022.
- [44] Jingyang Zhang, Yao Yao, Shiwei Li, Zixin Luo, and Tian Fang. Visibility-aware multi-view stereo network. *arXiv preprint arXiv:2008.07928*, 2020.
- [45] Kai Zhang, Gernot Riegler, Noah Snaveley, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020.
- [46] Yongqiang Zhang, Zhipeng Hu, Haoqian Wu, Minda Zhao, Lincheng Li, Zhengxia Zou, and Changjie Fan. Towards unbiased volume rendering of neural implicit surfaces with geometry priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4359–4368, 2023.
- [47] Fuqiang Zhao, Yuheng Jiang, Kaixin Yao, Jiakai Zhang, Liao Wang, Haizhao Dai, Yuhui Zhong, Yingliang Zhang, Minye Wu, Lan Xu, et al. Human performance modeling and rendering via neural animated mesh. *ACM Transactions on Graphics (TOG)*, 41(6):1–17, 2022.
- [48] Yiyu Zhuang, Qi Zhang, Ying Feng, Hao Zhu, Yao Yao, Xiaoyu Li, Yan-Pei Cao, Ying Shan, and Xun Cao. Anti-aliased neural implicit surfaces with encoding level of detail. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–10, 2023.

A Analysis of Local Scale in SDF-to-Density Transformation

We further illustrate the importance of the local scale in Figure A, using a simple single plane scenario for explanation. This plane has a low-texture region on the left and a rich-texture region on the right.

Without special constraints, the rendering weight should converge to a Dirac delta function at the surface in the richly textured region and form a scattered distribution in the weakly textured region.

However, under the assumption of a global scale factor, all areas on the plane follow the same distribution which will be derived in the following sections. This means that both richly textured and weakly textured regions share the same density bias, preventing the surface from converging correctly to the richly textured surface with higher certainty.

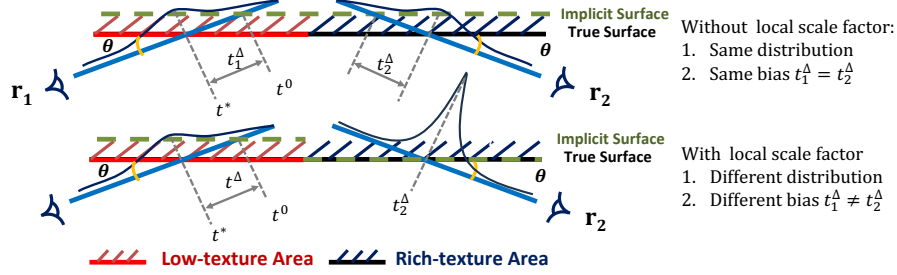


Figure 10: Explanation of the motivation behind the use of the local scale factor.

By introducing the local scale factor, the most significant difference is that the distribution of rendering weights along the ray is no longer uniform. The network can adaptively converge in the richly textured regions, and the density bias in these areas is no longer affected by the low-texture regions.

In Adaptive shells, their focus is primarily on rendering quality and they did not evaluate surface reconstruction metrics. A straightforward application to surface reconstruction can lead to issues such as increased density bias (as $w(t)$ is also a function of the scale factor $s(t)$). We tackle this by implementing special designs to ensure effectiveness, such as gradually scheduling the lower bound of the scale factor to correct relatively small biases and the explicit bias correction.

B Analysis of Density Bias

In this section, we analyze the density bias and discuss the main disadvantage of previous work (NeuS [28], VolSDF [35] and TUVr [46]). According to Equation (8), the rendering weight maximum point t^* should satisfy

$$t^* = \arg \max_{t \in (0, +\infty)} T(t) \sigma(\mathbf{r}(t)). \quad (13)$$

The derivation of rendering weight $w(t) = T(t) \sigma(\mathbf{r}(t))$ respected to t^* should equal to zero:

$$\begin{aligned} \left. \frac{\partial w(t)}{\partial t} \right|_{t=t^*} &= \left. \frac{\partial (T(t) \sigma(\mathbf{r}(t)))}{\partial t} \right|_{t=t^*} \\ &= \left. \frac{\partial T(t)}{\partial t} \right|_{t=t^*} \sigma(\mathbf{r}(t^*)) + \left. \frac{\partial \sigma(\mathbf{r}(t))}{\partial t} \right|_{t=t^*} T(t^*) \\ &= \left(\sigma^2(\mathbf{r}(t^*)) - \left. \frac{\partial \sigma(\mathbf{r}(t))}{\partial t} \right|_{t=t^*} \right) \exp \left(- \int_0^{t^*} \sigma(u) du \right) \\ &= 0. \end{aligned} \quad (14)$$

Then, we have

$$\sigma^2(\mathbf{r}(t^*)) = \left. \frac{\partial \sigma(\mathbf{r}(t))}{\partial t} \right|_{t=t^*}. \quad (15)$$

NeuS [28] and TUVr [46] seek a modeling approach for $\sigma(\mathbf{r}(t))$ such that it also satisfies the above equation at the point where the SDF value is zero $f(r(t^0)) = 0$. NeuS only satisfies this under the

first-order approximation of the distribution of the SDF along the ray, while TUVR extends this to arbitrary distributions.

Although TUVR's modeling ensures that the derivative of the rendering weight is zero at the point t^0 where the SDF value is zero, this does not mean that t^0 is the location of the global maximum of the weight.

During the optimization process, the distribution of rendering weight along the ray is a complex non-convex function. Therefore, TUVR only guarantees that t^0 is at a local maxima. Therefore, bias persists throughout the optimization process. We have shown more analysis on TUVR in Section D.

Here we take VolSDF's SDF-to-density modeling (3) as an example. And considering the simplest case, where the ray intersects with a single plane. In other words, the implicit surface at the current iteration is a single plane. In this case, if the angle between the ray and the plane is θ , and the distance from the ray's origin to the plane is d , then the distribution of the SDF along the ray can be expressed as $f(\mathbf{r}(t)) = -t \sin \theta + d$. In Equation (15), the expression on the right-hand side is

$$\frac{\partial \sigma(\mathbf{r}(t))}{\partial t} = -\frac{1}{2s^2} \frac{\partial f(\mathbf{r}(t))}{\partial t} \exp\left(\frac{|f(\mathbf{r}(t))|}{s}\right). \quad (16)$$

And the left-hand side is

$$\sigma^2(\mathbf{r}(t)) = \begin{cases} \frac{1}{4s^2} \exp\left(\frac{-2f(\mathbf{r}(t))}{s}\right) & \text{if } f(\mathbf{r}(t)) \geq 0, \\ \frac{1}{s^2} \left(1 - \frac{1}{2} \exp\left(\frac{f(\mathbf{r}(t))}{s}\right)\right)^2 & \text{if } f(\mathbf{r}(t)) < 0. \end{cases} \quad (17)$$

If $f(t^*) \geq 0$, then we have

$$\begin{aligned} \frac{1}{4s^2} \exp\left(\frac{-2f(\mathbf{r}(t^*))}{s}\right) &= -\frac{1}{2s^2} \frac{\partial f(\mathbf{r}(t))}{\partial t} \bigg|_{t=t^*} \exp\left(\frac{-f(\mathbf{r}(t^*))}{s}\right) \\ \exp\left(\frac{-f(\mathbf{r}(t^*))}{s}\right) &= -2 \frac{\partial f(\mathbf{r}(t))}{\partial t} \bigg|_{t=t^*} \\ \exp\left(\frac{t^* \sin \theta - d}{s}\right) &= 2 \sin \theta. \end{aligned} \quad (18)$$

We can directly obtain the closed-form solution for t^* only if $\sin \theta \leq 0.5$

$$t^* = \frac{s \ln(2 \sin \theta) + d}{\sin \theta}. \quad (19)$$

If $f(t^*) < 0$, then we have

$$\exp\left(\frac{-t \sin \theta + d}{s}\right) \sin \theta = 2 \left(1 - \frac{1}{2} \exp\left(\frac{-t \sin \theta + d}{s}\right)\right)^2 \quad (20)$$

Let $\exp((-t \sin \theta + d)/s)$ be denoted as m . It's easy to see that the above equation has a solution $m^* = 2 + \sin \theta - \sqrt{\sin^2 \theta + 4 \sin \theta}$ which can be obtained using the quadratic formula only if $\sin \theta > 0.5$, therefore

$$t^* = \frac{s \ln(1/m^*) + d}{\sin \theta}. \quad (21)$$

After organizing the two situations, can obtain the results:

$$t^* = \frac{s \ln(k) + d}{\sin \theta}, \quad k = \begin{cases} 2 \sin \theta & \text{if } \sin \theta \leq 0.5, \\ 1/m^* & \text{if } \sin \theta > 0.5. \end{cases} \quad (22)$$

In this case, the closed-form solution for t^0 can be directly obtained as $t^0 = d/\sin \theta$. Then the distance between t^* and t^0 is

$$t^\Delta = t^* - t^0 = \frac{\ln k}{\sin \theta} s. \quad (23)$$

It can be observed that t^Δ is a linear function of the scale factor s . If the scale factor s is relatively large, the implicit surface is far from the regions that contribute most to the color. At that time, the geometry regularization on the implicit surface acts on the incorrect surface, resulting in visible defects in the final mesh. Our explicit bias correction aims to minimize t^Δ during the coarse stage, so that the geometric regularization can be applied to the correct surface.

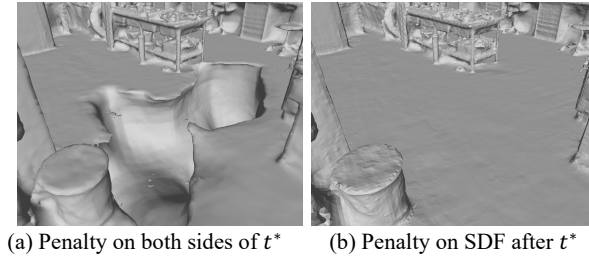


Figure 11: **Visual results on ScanNet++ with different designs of our explicit bias correction.** Penalty on both sides of t^* will lead to erroneous surface

C Design of Explicit Bias Correction

Although aligning t^* with the SDF zero-crossing is intuitive, our experiments showed that this approach requires special design considerations. We tried simply constraining $f(t^*)$ to approach zero but found that visible surface defects persisted. We also attempted to constrain the SDF values both before and after t^* , which led to very strange optimized surfaces, likely due to overly strong constraints on the SDF. We have shown the visual result in Figure 11.

The key to our method’s design lies in correcting large relative bias with explicit loss functions, while smaller biases are corrected by gradually scheduling the lower bound of the scale factor. We empirically found that visible surface errors are always caused by t^* is being before SDF zero-crossing points t^0 . Therefore, we penalize the SDF value just after t^* to trend towards negative values, which prevents t^* from being before t^0 . For the case where t^* is after t^0 , we can directly address it by gradually scheduling the lower bound of the scale factor, as the bias in this situation is relatively small.

We provided a mathematical explanation under the assumption of a sufficiently small local surface. According to Section B the distance between t^* and t^0 is as Equation (23).

When $t^\Delta < 0$, t^* is before t^0 , and when $t^\Delta > 0$, t^* is after t^0 . We visualized the values of t^Δ under different θ in Figure 12. We found that when t^* is before t^0 ($\sin \theta$ less than 0.5), the relative bias is significantly greater than when t^* is after t^0 . Therefore, we penalize the SDF value just after t^* to prevent t^* being before t^0 .

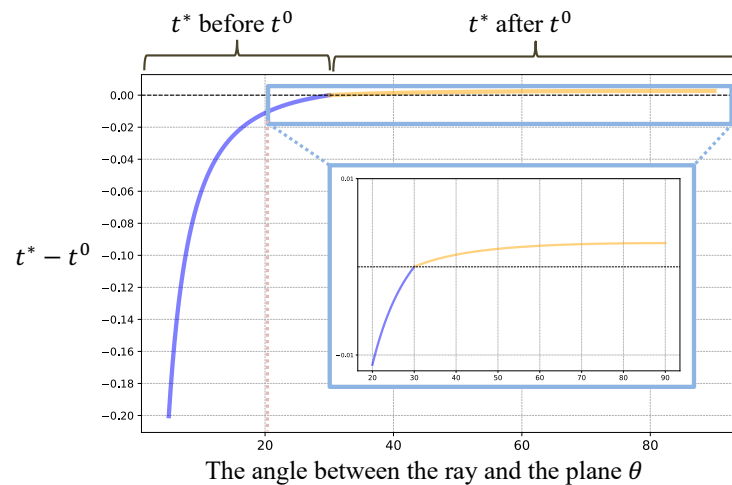


Figure 12: **The distance between t^* and t^0 as the angle between the ray and the single plane varies.** We find that when t^* precedes t^0 , a significantly larger relative bias is observed.

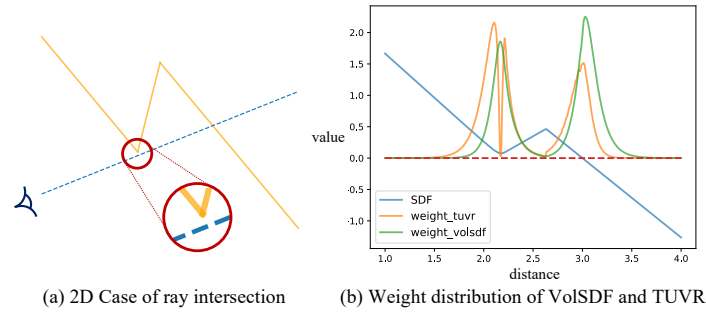


Figure 13: **A particular scenario for density bias.** In this situation TUVR exhibits a more pronounced bias compared to VolSDF.

D Analysis of TUVR

Previously, we mentioned that TUVR only proves t^0 as a local maximum for rendering weights, not a global one. Here, we present a scenario where bias exists in the modeling of TUVR. Consider a wall composed of three planes in space, with light passing through it, as shown in Figure 13 (a). In this situation, the rendering weights, as illustrated in Figure 13 (b), exhibit a bias in TUVR; although TUVR ensures a local peak at the SDF zero crossing point, the rendering weight is greater at a previous location. In this scenario, the bias in TUVR may even be greater than that in VolSDF.

Next, we analyze the unbiasedness of TUVR under the assumption of a local scale factor. To simplify the equations, we will henceforth abbreviate all $\mathbf{r}(t)$ as t and all $\frac{\partial}{\partial t}$ as $\cdot'(t)$. When using a local scale factor, the density represented under TUVR modeling is:

$$\sigma(t) = \begin{cases} \frac{1}{s(t)} \exp\left(\frac{-f(t)}{s(t)|f'(t)|}\right) & \text{if } f(t) \geq 0, \\ \frac{2}{s(t)} \left(1 - \frac{1}{2} \exp\left(\frac{f(t)}{s(t)|f'(t)|}\right)\right) & \text{if } f(t) < 0. \end{cases} \quad (24)$$

When $f(t) \geq 0$, the left-hand side of Equation (15) is:

$$\sigma'(t) = -\frac{s'(t)}{s^2(t)} \exp\left(-\frac{f(t)}{s(t)|f'(t)|}\right) + \frac{1}{s(t)} \left(-\frac{\left|\frac{f(t)}{|f'(t)|}\right|' s(t) - \frac{f(t)}{|f'(t)|} s'(t)}{s^2(t)}\right) \exp\left(-\frac{f(t)}{s(t)|f'(t)|}\right) \quad (25)$$

When $f(t) = 0$, the above expression can be simplified to:

$$\sigma'(t) = -\frac{s'(t)}{s^2(t)} - \frac{f'(t)}{s^2(t)|f'(t)|} \quad (26)$$

And we only consider the scenario where the light ray enters the plane, namely $f'(t) < 0$. So we have

$$\sigma'(t) = \frac{-s'(t) + 1}{s^2(t)} \quad (27)$$

At this point, $\sigma^2(t) = 1/s^2(t)$, so Equation (15) is only satisfied when $s'(t) = 0$, meaning the scale factor is a constant. Therefore, under the assumption of a local scale factor, TUVR's local unbiasedness (ensuring the SDF zero crossing point is a local maximum in rendering weights) cannot be achieved. When $f(t) < 0$, we can also arrive at a similar conclusion.

E Explanation of Stochastic Gradients

Our stochastic gradient estimation introduces some uncertainty into the true normals. For large-scale features, the Eikonal loss with varying step sizes can still be successfully minimized (since SDF near



Figure 14: **Visual results on Tanks and Temples from iteration 25,000 at the first stage.**

a large-scale plane should satisfy the Eikonal equation for different step sizes). In other words, the variance of stochastic gradients is small for large-scale features, making it easier for the model to minimize the Eikonal loss. However, for fine details, the random step sizes lead to high variance in the estimated normals, which reduces the impact of the Eikonal loss and makes it more challenging for the model to minimize samples with high variance.

F Impact of Color Conditioning on Normal

In Figure 14, we experimentally observed that in indoor scenes, when color conditioning on normals is applied, the optimization process becomes very slow and adversely affects the optimization results (both surface and rendering quality). However, this issue is nearly absent in outdoor scenes. We believe this is primarily due to the convergence behavior of the scale factor. Indoor scenes often have many weakly textured areas, leading to larger scale factors and a greater scatter distribution. Additionally, color conditioning on normals, which is intended for geometry disentanglement as mentioned in IDR and resonable for surface points, results in many details being optimized to incorrect positions before convergence is achieved.

However, our use of stochastic normals helps mitigate these erroneous surfaces. For detailed regions, the estimated normals have greater variance, which alleviates the impact of incorrect surfaces. This is also demonstrated in our ablation experiments, as shown in the Table 4. Cases D and E are as follows: Case D: Excludes the stage 1 Eikonal loss but incorporates color conditioning based on the estimated normal. Case E: Excludes the stage 1 Eikonal loss but includes color conditioning based on the analytical normal.

F-score \uparrow	D	E	Full Model
Courthouse (outdoor)	0.11	0.11	0.21
Meetingroom (indoor)	0.35	0.21	0.43

Table 4: **Impact of color conditioning on normal in different scenes.**

G Experimental Details

G.1 Datasets

We carry out experimental evaluations on two benchmark datasets: Tanks and Temples [11] and ScanNet++ [41]. The Tanks and Temples dataset is characterized by its large-scale, diverse real-world scenes, both indoors and outdoors. For our experiments, we utilize six scenes from the training subset, consistent with the scenes employed in Neuralangelo, to maintain comparability. Additionally, we extend our validation to four expansive indoor scenes from the advanced subset to further assess the robustness of our method. Turning to the ScanNet++ dataset, it is distinguished by its high-quality indoor scenes, supplemented with DSLR-quality images. From this dataset, we have selected eight scenes for our analysis.

G.2 Baselines

For the Tanks and Temples dataset, our methodology is compared against several prominent methods, including: NeuralWarp [5], COLMAP [23], NeuS [28], Geo-NeuS [7], and Neuralangelo [13]. In the context of the ScanNet++ dataset, our approach is contrasted with methods lacking prior knowledge, such as VolSDF [35], NeuS [28], and Neuralangelo [13]. Additionally, we evaluate our approach against methods incorporating pretrained prior information, notably MonoSDF [43]. It should be noted that our efforts to reproduce Neuralangelo for indoor scenes were unsuccessful. Instead, we employed the implementation of Bakedangelo from [42], which serves as an enhanced version of Neuralangelo. Bakedangelo utilizes the same proposal network as our setup. We discovered that manually adjusting the global scale of SDF-to-density conversion in Bakedangelo significantly

Scene	Metric	NeuS	VolSDF	Neuralangelo*	Neuralangelo	Ours	MonoSDF-MLP	MonoSDF-Grid
0e75f3c4d9	Acc	0.053	0.053	0.326	0.121	0.166	0.041	0.064
	Comp	0.065	0.078	0.039	0.053	0.029	0.038	0.027
	Prec	0.433	0.395	0.381	0.525	0.584	0.611	0.574
	Recal	0.437	0.363	0.734	0.719	0.790	0.665	0.705
	F-score	0.435	0.378	0.502	0.607	0.671	0.637	0.633
036bce3393	Acc	0.050	0.053	0.184	0.168	0.028	0.062	0.037
	Comp	0.066	0.098	0.028	0.023	0.029	0.097	0.037
	Prec	0.507	0.475	0.509	0.520	0.688	0.383	0.606
	Recal	0.492	0.372	0.754	0.780	0.733	0.333	0.656
	F-score	0.499	0.417	0.608	0.624	0.710	0.356	0.630
108ec0b806	Acc	0.047	0.049	0.121	0.092	0.044	0.059	0.041
	Comp	0.088	0.097	0.060	0.050	0.053	0.113	0.045
	Prec	0.519	0.475	0.389	0.498	0.597	0.395	0.575
	Recal	0.432	0.372	0.542	0.624	0.587	0.303	0.576
	F-score	0.472	0.417	0.453	0.554	0.592	0.343	0.576
21d970d8de	Acc	0.046	0.070	0.247	0.261	0.075	0.054	0.042
	Comp	0.050	0.060	0.071	0.049	0.031	0.062	0.031
	Prec	0.526	0.383	0.403	0.403	0.588	0.365	0.580
	Recal	0.565	0.410	0.595	0.656	0.726	0.371	0.676
	F-score	0.545	0.396	0.481	0.499	0.650	0.368	0.624
355e5e32db	Acc	0.042	0.043	0.073	0.079	0.034	0.043	0.032
	Comp	0.075	0.071	0.038	0.034	0.038	0.060	0.031
	Prec	0.534	0.501	0.575	0.575	0.672	0.439	0.657
	Recal	0.465	0.447	0.701	0.730	0.689	0.404	0.683
	F-score	0.497	0.475	0.632	0.643	0.681	0.421	0.669
578511c8a9	Acc	0.094	0.080	0.254	0.222	0.103	0.063	0.039
	Comp	0.174	0.212	0.057	0.030	0.044	0.171	0.044
	Prec	0.373	0.354	0.378	0.465	0.520	0.322	0.608
	Recal	0.328	0.271	0.599	0.737	0.623	0.264	0.657
	F-score	0.349	0.307	0.463	0.570	0.567	0.290	0.631
7f4d173c9c	Acc	0.080	0.097	0.414	0.420	0.155	0.094	0.138
	Comp	0.055	0.075	0.026	0.024	0.028	0.030	0.022
	Prec	0.503	0.457	0.437	0.457	0.657	0.704	0.708
	Recal	0.496	0.436	0.771	0.799	0.732	0.713	0.813
	F-score	0.500	0.446	0.557	0.582	0.693	0.709	0.757
09c1414f1b	Acc	0.170	0.085	0.348	0.106	0.081	0.070	0.050
	Comp	0.116	0.285	0.076	0.242	0.056	0.105	0.080
	Prec	0.348	0.347	0.289	0.475	0.545	0.418	0.651
	Recal	0.331	0.256	0.470	0.403	0.530	0.363	0.589
	F-score	0.339	0.294	0.358	0.436	0.537	0.389	0.618
Mean	Acc	0.073	0.066	0.246	0.184	0.086	0.061	0.055
	Comp	0.086	0.122	0.049	0.063	0.039	0.085	0.040
	Prec	0.468	0.423	0.420	0.490	0.606	0.455	0.620
	Recal	0.443	0.366	0.646	0.681	0.676	0.427	0.669
	F-score	0.455	0.391	0.507	0.564	0.638	0.439	0.642

Table 5: **ScanNet++ Benchmark** We established benchmarks for eight different scenes within the ScanNet++ dataset.

mitigates the optimization issues encountered with Neuralangelo in indoor scenes. Additionally, we eliminated the unnecessary 360 unbound setting and background modeling for indoor scenes. Therefore, in the context of the Tanks and Temples advanced subset and ScanNet++, we report the results obtained with Bakedangelo as a substitute for those of Neuralangelo.

G.3 Evaluation Metrics

Mesh is extracted through marching cube algorithm with a resolution of 2048 applied across all scenes. For the Tanks and Temples dataset, the evaluation metrics on the training subset are computed using the official Python script provided by the dataset’s maintainers¹. Meanwhile, for the four scenes from the advanced subset, our reconstructed results are submitted to the evaluation server², which calculates the evaluation metrics. For the ScanNet++ dataset, we calculate the F-score with a threshold of 0.025 to compare the resultant meshes with the ground truth mesh, which is derived from the point cloud captured by a laser scanner.

¹https://github.com/isl-org/TanksAndTemples/tree/master/python_toolbox/evaluation

²<https://www.tanksandtemples.org/>

G.4 Implementation Details

NeuRodin utilizes a multi-resolution hash grid for encoding, spanning from 2^5 to 2^{11} across 16 levels. Each hash entry possesses a channel size of 2 for room-level scenes such as ScanNet++, which is adjusted to 8 for larger-scale scenes like Tanks and Temples. The maximum number of hash entries for each resolution is set at 2^{19} . We incorporate per-image appearance encoding in the style of NeRF-W [15] while employing a proposal network [1] based on a compact hash grid. For the outdoor scene, we model the background using an additional network [45] with the hash grid. For the ScanNet++ datasets, we sample 512 pixels per iteration. In the case of the Tanks and Temples dataset, we sample 1024 pixels during the first stage and escalate to 8192 pixels for the second stage. We set the weights λ_{eik} and λ_{smooth} to be 0.01 and 0.005, respectively. For outdoor scenes, we set λ_{bias} to 0.1. For indoor scenes, we progressively increase λ_{bias} from 0.001 to 0.05 over the first 10,000 iterations through an exponential adjustment. Bakedangelo samples 8192 pixels per iteration in large indoor scenes of Tanks and Temples, aligning with the settings employed by Neuralangelo. In the context of room-level scenarios within ScanNet++, the batch size is adjusted to 1024 pixels per iteration. Our implementation of the method employs PyTorch [20] and utilizes the Adam optimizer, with a learning rate of 0.001 applied to both the hash grid and the network, alongside a weight decay set at 0.01. For the background model, we set the learning rate to 0.01. The total training steps amount to 300k, with the learning rate for the foreground model being decayed by a factor of 10 at 160k and 240k steps. For the background model, we employ an exponential schedule for the learning rate, reducing it to 0.0001. For the proposal network, the learning rate is decayed by a factor of 3 at steps 150k, 225k, and 270k. All our experiments were conducted on an A100 40G GPU. We roughly require 7 GPU hours to complete the reconstruction of an indoor scene from the ScanNet++ dataset. For large-scale scenes, the reconstruction takes approximately 18 GPU hours.

G.5 Implementation of the Explicit Bias Correction

In the actual implementation of explicit bias correction, when inferring the SDF at ϵ_{bias} after the point $\mathbf{r}(t^*)$ along the ray in Equation (8), we also infer the SDF at $\epsilon_{\text{bias mask}}$ beyond $\mathbf{r}(t^*)$ along the ray as a mask. If $f(\mathbf{r}(t^* + \epsilon_{\text{bias mask}}))$ is less than zero, we do not apply the bias loss to this particular ray. We have found that this approach effectively prevents incorrect alignments that may be caused by our approximate estimation of t^* . For outdoor scenes, we simply determine whether a ray is cast towards the background by checking if there exists a negative value of SDF at any sampled point along the ray. If so, we similarly refrain from performing bias correction. We have configured $\epsilon_{\text{bias mask}}$ to 0.001 for large-scale scenes, such as the Tanks and Temples dataset. For room-level scenes, such as the ScanNet++ dataset, we have set it to 0.01.

G.6 Implementation Details of the Two-Stage Optimization

The local scale modeling in Equation (5) can impede the model’s convergence to the surface to some extent. Therefore, we manually adjust the lower bound s_{coarse} for the scale to prevent the ambiguity that arises from excessively small scales. Analogous to the manual adjustments made in the first phase, our objective in this stage is to facilitate convergence from volume rendering to surface rendering, thereby aligning the implicit surface with volume rendering completely. To achieve this, we exponentially increase the lower bound of the scale s_{fine} to a substantial value. In all experiments, we set the value of s_{coarse} to 100 and s_{fine} to 3000.

Metric	F-Score (%) \uparrow				
Scene	VolSDF	MonoSDF	Neuralangelo*	Neuralangelo	Ours
Auditorium	3.16	10.97	<u>14.32</u>	14.09	16.03
Ballroom	11.61	29.30	<u>32.21</u>	28.93	33.10
Courtroom	7.71	21.58	36.53	32.81	33.53
Museum	4.41	21.71	25.49	<u>29.28</u>	32.71
Mean	6.72	20.89	<u>27.14</u>	26.28	28.84

Table 6: **Quantitative evaluation of our method versus prior work on the Tanks and Temples advance subset.** The best performance and the second-best outcomes are highlighted for easy reference.

		24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean
CD (mm) ↓	NeRF	1.90	1.60	1.85	0.58	2.28	1.27	1.47	1.67	2.05	1.07	0.88	2.53	1.06	1.15	0.96	1.49
	VolSDF	1.14	1.26	0.81	0.49	1.25	0.70	0.72	1.29	1.18	0.70	0.66	1.08	0.42	0.61	0.55	0.86
	NeuS	1.00	1.37	0.93	0.43	1.10	0.65	0.57	1.48	1.09	0.83	0.52	1.20	0.35	0.49	0.54	0.84
	Neuralangelo	0.37	0.72	0.35	0.35	0.87	0.54	0.53	1.29	0.97	0.73	0.47	0.74	0.32	0.41	0.43	0.61
	Ours	0.37	0.65	0.31	0.36	0.93	0.54	0.63	1.30	1.07	0.68	0.52	0.61	0.32	0.41	0.37	0.60

Table 7: **Quantitative results on DTU Benchmark.**

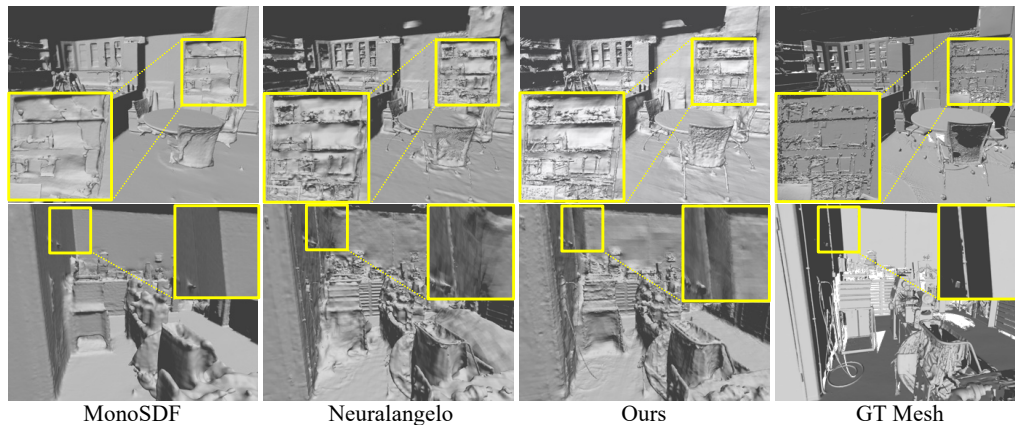


Figure 15: **Quantitative evaluation of our method on the ScanNet++ dataset.**

H More Experimental Results

H.1 Experimental Results on the Scannet++ Benchmark

For every scene, we use high-quality DSLR camera images from all frames for our experiments. For methods that do not use prior knowledge, we downsample the images from 1752×1168 to 876×584 for training. For methods that do use prior knowledge, we do something similar to MonoSDF. We first crop the images to 1152×1152 , then downsample them to 384×384 before feeding them into the Omnidata model [6] to predict geometry cues. First, we scale the pose to be centered within a bounding sphere of radius equal to 1. Subsequently, we scale the camera pose by 0.8 to ensure that all scene boundaries are contained within the bounding sphere. We apply the marching cubes algorithm to a cube with dimensions -1 to 1 at a resolution of 2048. Then, we evaluate the metrics: Accuracy (Acc), Completeness (Comp), Precision (Pred), Recall (Rec), and F-score. The final result is shown in Table 5. Here, we also present some visual results in Figure 15.

H.2 Experimental Results on the DTU Benchmark

Although our design is not specifically tailored for single-object datasets, we validated our method on the DTU Benchmark. The results are shown in the Table 7. We found that, even without special parameter tuning, our method achieves results comparable to Neuralangelo and surpasses other baseline methods.

H.3 Experimental Results on the Tanks and Temples Advance Subset

In this section, we present the individual results for each scene within the advanced subset of the Tanks and Temples dataset, as shown in Table 6. Except for the Courtroom scene, our method significantly outperforms comparative approaches in the other three large-scale indoor scenes, demonstrating the effectiveness of our approach. We present a portion of the mesh for the museum data in Figure 16 and more results in Figure 17, highlighting our method’s significant capability in capturing details.

H.4 Comparison with TUVR

Since TUVR is not open-source and only reports metrics on the DTU dataset, we reproduced its unbiased SDF-to-density technique and combined it with a hash grid to create the TUVR-Grid method for comparison. Additionally, we compared our results on the DTU dataset with the reported results in the paper that did not use MVS priors (TUVR-MLP). The results are shown in the table below.

The performance of TUVR is not ideal on all datasets because its unbiased nature is not fully guaranteed, and it is also somewhat affected by over-regularization.

F-score \uparrow	TUVR-Grid	Ours
Barn	0.57	0.70
Caterpillar	0.25	0.36
Courthouse	0.11	0.21
Meetingroom	0.03	0.43
Truck	0.33	0.47
Ignatius	0.66	0.86
Mean	0.40	0.51

Table 8: **Comparison with TUVR on the Tanks and Temples training set.**

		24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean
\rightarrow	TUVR-MLP	0.72	0.77	0.67	0.37	0.93	0.58	0.61	1.23	1.15	0.65	0.56	1.08	0.34	0.45	0.47	0.71
	TUVR-Grid	0.84	0.81	1.44	0.37	1.23	0.69	0.78	1.16	1.23	0.65	0.54	1.32	0.34	0.43	0.54	0.82
	Ours	0.37	0.65	0.31	0.36	0.93	0.54	0.63	1.30	1.07	0.68	0.52	0.61	0.32	0.41	0.37	0.60

Table 9: **Comparison with TUVR on the DTU Benchmark.**

H.5 Image Reconstruction Comparison

While our design is primarily intended for surface reconstruction tasks, we have also compared the results of NeuS and Neuralangelo using the image quality evaluation methods from Neuralangelo, across different parameter scales.

As shown in Table 10, with fewer parameters, our method (Ours-19) achieves image reconstruction quality close to that of Neuralangelo. With a larger number of parameters, our method (Ours-22) surpasses Neuralangelo in terms of image reconstruction quality.

PSNR	NeuS	Neuralangelo-22	Ours-19	Ours-22
Mean	24.58	27.24	26.90	27.67

Table 10: **Image Reconstruction Results on the Tanks and Temples training set.**

H.6 More Ablation Study

We conducted additional ablation experiments on the Tanks and Temples training set and we included two more cases to verify the role of Eikonal loss in maintaining the natural zero level set during the first stage, as well as the impact of color condition on normals.

The five cases are: A: Without the local scale factor. B: Changing the stage 1 estimated gradient to the analytical gradient. C: Without explicit bias correction. D: Without stage 1 Eikonal loss, but with color conditioning on the estimated normal. E: Without stage 1 Eikonal loss, but with color conditioning on the analytical normal.

F-score \uparrow	A	B	C	D	E	Full Model
Barn	0.70	0.58	0.68	0.54	0.56	0.70
Caterpillar	0.33	0.33	0.36	0.32	0.33	0.36
Courthouse	0.24	0.12	0.19	0.11	0.11	0.21
Meetingroom	0.37	0.29	0.38	0.35	0.21	0.43
Truck	0.46	0.39	0.47	0.42	0.43	0.47
Ignatius	0.81	0.79	0.86	0.77	0.78	0.86
Mean	0.49	0.42	0.49	0.42	0.40	0.51

Table 11: **More ablation study on Tanks and Temples training set.**

The quantitative results are in Table 11. The results from cases A, B, and C demonstrate the effectiveness of the techniques we proposed. The results from cases D and E indicate the necessity of applying Eikonal loss during the first stage.

I Limitation

Though NeuRodin has capabilities in reconstruction, it falls short in certain areas. Specifically, it struggles to faithfully reconstruct the correct surface in areas that are textureless and less observed. NeuRodin is also incapable of handling situations with strong ambiguity. Additionally, density is not guaranteed to be unbiased, hence, bias will always exist though the SDF-to-density conversion.

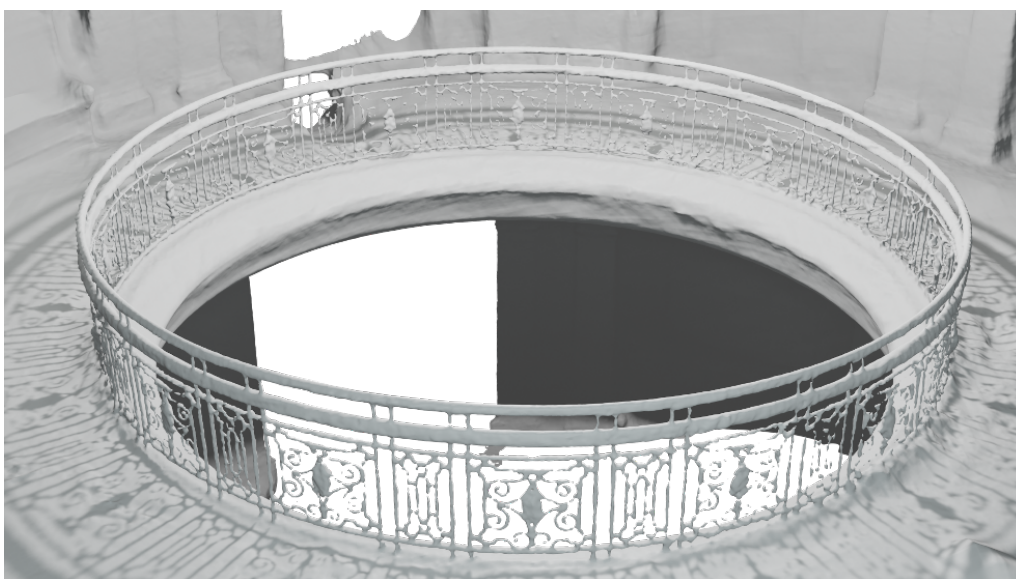


Figure 16: **Recovered mesh of the *Museum* data from the Tanks and Temples dataset.**

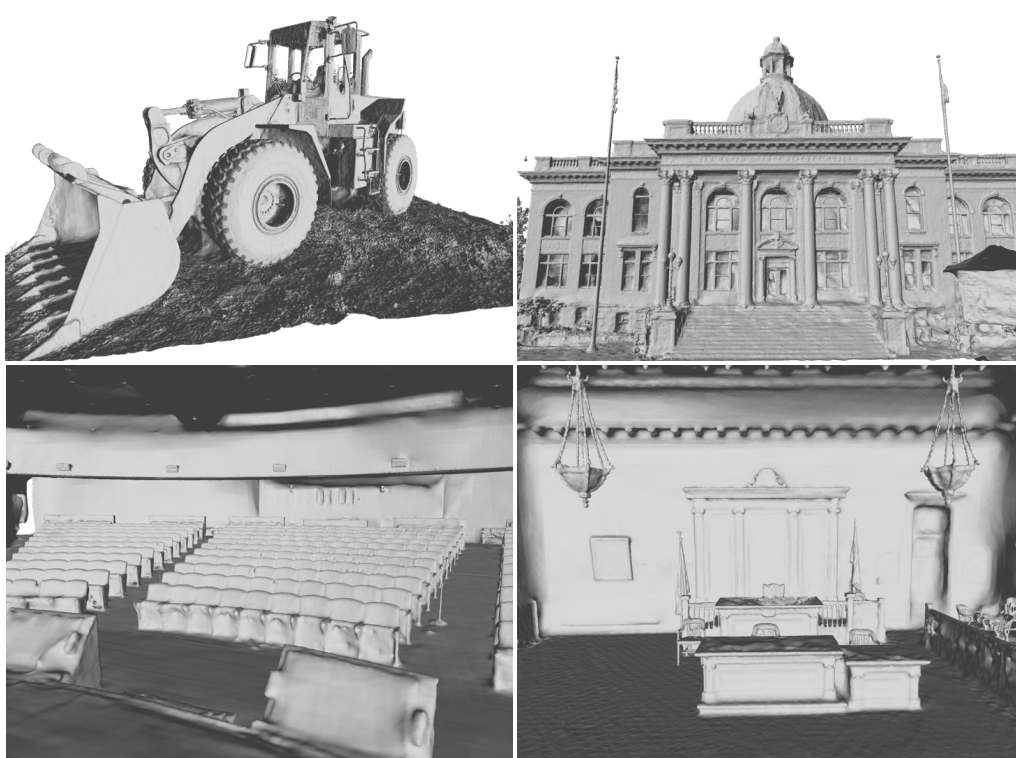


Figure 17: **More results from the Tanks and Temples dataset.**

Lastly, the time taken for the reconstruction of large scenes is considerable, often requiring hours to complete. This could prove to be inefficient in scenarios where time is of the essence.

J Societal Impact

Our model achieves high-fidelity 3D reconstruction. The societal impact of this development is multifaceted. On one hand, it enables significant advancements in fields such as architecture and augmented reality, improving professional practices and potentially benefiting the public by enhancing the precision and interactivity of digital models. On the other hand, the increase in computational power demands may lead to greater energy consumption, which poses environmental considerations. Moreover, there could be privacy concerns if such technology is applied to reconstruct environments from personal data without consent. Overall, while this technology presents opportunities for progress and innovation, it also requires careful consideration of ethical and environmental implications.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Yes, the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have discussed the limitations in the supplementary material.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We present all implementation details in the supplementary material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We have released the code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All the details of training and testing have been outlined in the supplementary material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Completing the experiments required substantial utilization of graphics card resources.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We have provided sufficient information on the computer resources in the supplementary material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research conducted in the paper conforms fully with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss potential negative societal impacts in our supplementary material.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have cited all the datasets we have used.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: Our paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.