

UNIS: A Unified Framework for Achieving Unbiased Neural Implicit Surfaces in Volume Rendering

Supplementary Material

This supplementary material is organized as follows: Appendix A presents two challenging applications: reconstruction under sparse views and reconstruction from reflective surfaces. These scenarios often pose significant difficulties for existing neural implicit surface techniques. We demonstrate that UNIS effectively addresses these challenges using a typical multi-stage training strategy with tailored component choices for each stage. Appendix B explores various types of prior knowledge that can be leveraged to enhance neural implicit surfaces. Appendix E presents the details of using alternative implicit function f_{GWN} within UNIS. Appendix F discusses the assumption of UNIS. Appendix G presents visual results for all test models listed in Table 2 of the main paper. Finally, Appendix H discusses the potential of UNIS.

A. Applications

Sparse-view reconstruction Reconstructing 3D models from sparse-view settings poses significant technical challenges, primarily due to the increased risk of overfitting caused by the limited number of views. As shown in Fig. 7, HF-NeuS [23], which decomposes the SDF into base and displacement components, struggles in this scenario as it tends to overfit to colors rather than accurately learning geometry. Existing methods often incorporate 2D or 3D priors to enhance training stability and robustness [18, 26, 29]. Here, we demonstrate that a two-stage UNIS training strategy, without relying on any priors, can achieve satisfactory reconstruction results. Specifically, we utilize NeuS in the first stage to learn a coarse geometry. Subsequently, we employ a new UNIS case study, referred to as UNIS-C6, which utilizes an exponent $k = 2$, the kernel function h_{laplace} , and the implicit function f_{SDF} . Computational results show that this configuration generates a stable SDF that effectively captures both the geometry and appearance of the target object. See Fig. 7 for the visual results.

Reflective surfaces Reflective surfaces also pose challenges in neural rendering due to their complex light interactions, which often lead to ambiguous geometry and appearance information that traditional rendering techniques struggle to disentangle. To address these challenges, we adopt a two-stage training strategy similar to the one used for sparse-view reconstruction. Specifically, we utilize UNIS-C6 in the first stage to stabilize geometry learning, followed by NeuS in the second stage to refine both geometry and appearance. See Fig. 8 for two examples from the

ShinyBlender dataset [20].

B. Supervision with priors

Beyond kernel and implicit functions, prior knowledge is essential in volume rendering, helping to reduce ambiguity in this inherently ill-posed problem and enhancing surface reconstruction quality. This section discusses commonly used priors that improve reconstruction fidelity.

B.1. Common priors

Mask supervision. Though often overlooked, the original NeuS [22] training framework supports mask supervision for neural implicit surface representations. This mask supervision serves as a prior by effectively distinguishing foreground objects from the background, simplifying the training process.

3D supervision. Geo-NeuS [6] and PG-NeuS [30] adopt point-based supervision by identifying multiple surface points via SfM techniques [8]. They then apply loss functions to encourage the implicit field values at these points to approach zero, guiding the network toward accurate surface representation.

Normal and depth information. MonoSDF [28] leverages a pretrained model to predict normals and depths from monocular input images, supervising the learned distance field to align with these predictions. Similarly, NeuRIS [21] integrates predicted normal information into the reconstruction framework, assigning confidence scores to predictions and using only those deemed reliable for supervision.

Environmental priors. For indoor scenes, the Manhattan World assumption [4] is commonly applied. This assumption posits that walls and floors are planar and aligned with the coordinate axes. Manhattan-SDF [7], for example, employs image segmentation to identify walls and floors, introducing a planar constraint to enforce normal consistency across pixels within the same planar surface.

Patch warping. Patch warping leverages color information across images to improve reconstruction fidelity. For instance, NeuralWarp [5] incorporates color from existing image patches alongside the color predicted by the color networks, refining volume rendering by guiding the geometry network to learn a representation closer to the ground



Figure 7. Sparse-view reconstruction of a human head. The input consists of five images with views covering the human face. We adopt a two-stage training strategy, using NeuS in the first stage followed by UNIS-C6 in the second stage, to effectively learn both geometry and appearance. This figure presents the reconstructed geometry alongside six rendered images: three from training views and three from novel views, with PSNR values reported to quantify rendering quality for each method. Due to the limited number of views, HF-NeuS tends to overfit to colors in the training views, producing high-quality renderings for these views. However, its inability to accurately reconstruct geometry leads to a significant drop in performance for novel views. In contrast, our method, while slightly underperforming in PSNR for training views, effectively reconstructs both geometry and appearance, achieving much better results on novel views.




Figure 8. Reconstruction from reflective surfaces. UNIS-C6, combined with NeuS, achieves the best results compared to VolSDF, HF-NeuS, and NeuS alone. In this scenario, the zero-level sets of HF-NeuS include extraneous surfaces encapsulating the object. Its visual results shown here were obtained after removing these unnecessary components.

truth. Similarly, Geo-NeuS [6] applies a patch-based loss to enforce multi-view geometry consistency.

Color prediction with reflection models. For reflective objects, material-based color functions are introduced to improve color prediction. Rather than predicting color as a single component, many methods decompose it into diffuse and specular components or BRDF parameters [2, 17, 19, 24, 31], enhancing the accuracy of reconstructed colors

for objects with reflective surfaces.

C. Proof of Theorem 1

Theorem 1. *Let $f(t)$ be a monotonically decreasing implicit function along the ray \mathbf{r} , with its zero level set defining the surface \mathcal{S} . Modeling the accumulated transmittance $T(t)$ as a function derived from $f(t)$ is equivalent to modeling the volume density $\sigma(t)$ using $f(t)$. These two repre-*

sentations are inter-convertible as follows:

$$T(t) = \exp \left\{ - \int_0^t \sigma(u) du \right\},$$

and

$$\sigma(t) = -\frac{T'(t)}{T(t)}.$$

Similarly, modeling color weights $w(t)$ is equivalent to modeling accumulated transmittance $T(t)$ as given by:

$$w(t) = -T'(t),$$

and

$$T(t) = 1 - \int_0^t w(u) du.$$

Proof. We present the proof by showing the transformations for $w \Leftrightarrow T$ and $T \Leftrightarrow \sigma$, respectively.

(1) Between weights w and accumulated transmittance T

$T \Rightarrow w$: Straightforward calculation shows

$$\begin{aligned} T'(t) &= \frac{d}{dt} \exp \left\{ - \int_0^t \sigma(u) du \right\} \\ &= \exp \left\{ - \int_0^t \sigma(u) du \right\} \cdot \frac{d}{dt} \left(- \int_0^t \sigma(u) du \right) \\ &= \exp \left\{ - \int_0^t \sigma(u) du \right\} \cdot (-\sigma(t)) \\ &= -T(t)\sigma(t) \\ &= -w(t). \end{aligned}$$

$w \Rightarrow T$: Given that $T'(t) = -w(t)$, we can integrate both hand sides to get:

$$\begin{aligned} \int_0^t T'(u) du &= - \int_0^t w(u) du \\ \Rightarrow T(t) - T(0) &= - \int_0^t w(u) du \\ \Rightarrow T(t) &= 1 - \int_0^t w(u) du. \end{aligned}$$

(2) Between accumulated transmittance T and volume density σ

$T \Rightarrow \sigma$: From $w(t) = T(t)\sigma(t)$, we obtain

$$\sigma(t) = \frac{w(t)}{T(t)} = \frac{-T'(t)}{T(t)}.$$

$\sigma \Rightarrow T$: The accumulated transmittance $T(t)$ can be derived directly from the density σ by the definition of accumulated transmittance.

We have shown that w and T are inter-convertible. T and σ are also inter-convertible. Therefore, modeling any one of these three quantities is equivalent to modeling the other two. \square

Remark. The weights $w(t)$ and the volume density $\sigma(t)$ can also be directly converted into each other.

$\sigma \Rightarrow w$: By the definition of weights, we have

$$w(t) = T(t)\sigma(t) = \exp \left\{ - \int_0^t \sigma(u) du \right\} \cdot \sigma(t).$$

$w \Rightarrow \sigma$: Conversely, to convert weights $w(t)$ back into volume density $\sigma(t)$, we need

$$\sigma(t) = \frac{w(t)}{T(t)} = \frac{w(t)}{1 - \int_0^t w(u) d(u)}.$$

D. Sufficient conditions

In the main paper, we introduced the following necessary condition for unbiased volume density mapping:

$$\sigma'(t^*) = \sigma^2(t^*). \quad (25)$$

To establish a sufficient condition, we additionally require that $w'(t) > 0$ in the vicinity of t^* for $t \leq t^*$ and $w'(t) < 0$ in the vicinity of t^* for $t > t^*$. This ensures that $w(t)$ reaches its maximum value at t^* . Below, we reformulate these requirements on the color weights w in terms of the density function σ .

Given that: $w'(t) = T(t) [\sigma'(t) - \sigma^2(t)]$, the sign of $w'(t)$ is determined by the sign of $\sigma'(t) - \sigma^2(t)$. With the proposed mappings, we have:

$$\begin{aligned} \sigma'(t) &= g_1(f(t), f'(t))f'(t) \\ &= cs^2(f'(t))^{2k+2}h'(sf(t)(f'(t))^k) \end{aligned}$$

and:

$$\sigma^2(t) = c^2s^2(f'(t))^{2k+2}h^2(sf(t)(f'(t))^k).$$

To facilitate discussion, we introduce the notation

$$l(t) \triangleq sf(t)(f'(t))^k.$$

Then, the difference between σ' and σ^2 is given by:

$$\sigma'(t) - \sigma^2(t) = s^2(f'(t))^{2k+2} [ch'(l(t)) - c^2h^2(l(t))].$$

Thus, the sign of $w'(t)$ is determined by:

$$ch'(l(t)) - c^2 h^2(l(t)).$$

For $w(t^*)$ being a maximum value, we need to ensure that $w''(t^*) < 0$, which means:

$$\frac{d}{dt}ch'(l(t^*)) - \frac{d}{dt}c^2 h^2(l(t^*)) < 0. \quad (26)$$

Under the assumption of ray-plane intersection, $f'(t^*)$ is a constant. Therefore, we have:

$$\frac{d}{dt}ch'(l(t)) = cs(f'(t))^{k+1} h''(l(t)),$$

and:

$$\frac{d}{dt}c^2 h^2(l(t)) = 2c^2 s(f'(t))^{k+1} h(l(t))h'(l(t)).$$

Inequality (26) is equivalent to

$$s(f'(t^*))^{k+1} [ch''(l(t^*)) - 2c^2 h(l(t^*))h'(l(t^*))] < 0.$$

Since $f(t^*) = 0$, it follows that $l(t^*) = 0$. Consequently, the above inequality simplifies to:

$$s(f'(t^*))^{k+1} [ch''(0) - 2c^2 h(0)h'(0)] < 0. \quad (27)$$

Substituting $c = h'(0)/h^2(0)$ into (27), and noting that the learnable parameter $s > 0$, we arrive at the following condition:

$$(f'(t^*))^{k+1} \frac{h'(0)}{h^3(0)} [h(0)h''(0) - 2h'(0)^2] < 0. \quad (28)$$

Inequality (28), together with the condition $\sigma'(t^*) = \sigma^2(t^*)$, forms the sufficient conditions for achieving first-order unbiasedness. In Tab. 3, we verify that UNIS-C0 as well as the proposed UNIS-C1 through UNIS-C7 satisfy these conditions.

Mapping	$(f'(t^*))^{k+1}$	$h(0)$	$h'(0)$	$h''(0)$
UNIS-C0	< 0	0.5	-0.25	0
UNIS-C1	< 0	0.5	-0.5	0
UNIS-C2	< 0	0.5	-0.5	0
UNIS-C3	< 0	$\ln 2$	-0.5	0.25
UNIS-C4	> 0	0.5	0.25	0
UNIS-C5	< 0	0.5	-0.25	0
UNIS-C6	< 0	0.5	-0.5	0
UNIS-C7	< 0	0.5	-0.25	0

Table 3. Verification of the sufficient conditions for the volume density mappings proposed in the paper.

E. Implicit function f_{GWN}

In this section, we explore yet another implicit function, which is defined as the generalized winding number field [1, 10].

As mentioned in the main paper, we require the implicit function f to be positive outside and negative inside, also not increasing through the first intersection. In practical scenarios, the ray would generally have more than one intersection. Our proposed framework allows for multiple color weight local maximum at each ray-plane intersection point, but the volume rendering formula ensures that later intersections have much fewer color weights than the first intersection, so that the effect is negligible. In general, we do not impose more restrictions on the implicit function.

While signed distance fields are a prime example of the implicit functions discussed back in Sec. 5, where $f_{\text{SDF}}(\mathbf{x}) = 0$ for $\mathbf{x} \in \mathcal{S}$ and $\|\nabla f_{\text{SDF}}(\mathbf{x})\| = 1$ for $\mathbf{x} \notin \mathcal{S}$. Other implicit functions, such as the indicator function in Poisson surface reconstruction (PSR) [9, 12, 13] and the generalized winding number field [10] used in recent 3D reconstruction methods [14–16, 27], can also be integrated within UNIS. As shown in [3], the indicator function in PSR is equivalent to a regularized GWN field. Therefore, we consider the GWN field f_{GWN} , which is governed by Poisson’s equation with either Dirichlet or jump boundary conditions. For points off the surface $\mathbf{x} \notin \mathcal{S}$, they can be regularized using Laplacian constraints, $\Delta f_{\text{GWN}}(\mathbf{x}) = 0$. Since GWN takes a value of 1 for interior points and 0 for exterior points, with the $\frac{1}{2}$ -level set defining the target surface, we use a shifted and inverted function $f_{\text{GWN}}(\mathbf{x}) = \frac{1}{2} - \text{GWN}(\mathbf{x})$. Fig. 9 illustrates the implicit functions f_{SDF} and f_{GWN} of the Stanford Bunny model.




Figure 9. Illustrations of f_{SDF} and f_{GWN} on the Stanford Bunny model. We show cross-sectional views of the fields, visualized with a heat color map where warm colors represent high values and cold colors represent low values.

To illustrate training with GWN, we introduce UNIS-C7. This method utilizes GWN as the implicit function f . Its density mapping is as follows:

$$\sigma_{\text{UNIS-C7}}(t) = -sf'_{\text{GWN}}(t)h_{\text{logistic}}(-sf_{\text{GWN}}(t)). \quad (29)$$

For the GWN-based case study C7, we first obtain a dense point cloud, then initialize the GWN field using

DWG [16], which is an efficient method that computes the GWN field from unoriented points in just a few seconds. Following this, we pretrain a lightweight MLP network to overfit the GWN field. During the main training process, we adopt a similar approach to [3], using a network to predict the GWN of the space for rendering, with the pretrained network providing supervision.

We observe that when GWN is used as implicit functions, its performance on Chamfer Distances is still comparable to other UNIS case studies. In some cases, it significantly improves the reconstruction quality.

F. Assumption

Following the approach in NeuS [22], we assume throughout this paper that the target surface is locally approximated by a plane and analyze the ray-plane intersection scenario. This assumption is reasonable, as smooth surfaces can be well approximated by piecewise linear functions, achieving a quadratic convergence rate. This explains why polygonal meshes remain the dominant representation for 3D surfaces of arbitrary geometry.

When representing surfaces using polygonal meshes, the local planar approximation is exact within each individual face. The only discrepancy arises at mesh edges, where the surface normal is discontinuous. However, these edges form a set of measure zero in the continuous domain, meaning that the assumption holds almost everywhere except at these singularities. In practice, this suggests that ray-plane assumptions remains valid for most rays, with only a negligible subset of cases violating the assumption.

G. Visual results

Figure 10 and Figure 11 present the visual results for all test models from the DTU dataset [11] and the OmniObject3D dataset [25]. For each model, we include results from the two baseline methods, VolSDF and NeuS, alongside our proposed UNIS case studies.

H. Final remarks

UNIS decomposes the volume rendering process for neural implicit surfaces into distinct modules: kernel function h , exponent k , implicit function f , and the incorporation of priors. This modular design offers flexibility, enabling the development of multi-stage training strategies for handling challenging scenarios, with tailored component choices for each stage. We believe that this framework holds significant potential for guiding the advancement of new neural implicit surface techniques and fostering further innovation in the field.

References

- [1] Gavin Barill, Neil G. Dickson, Ryan Schmidt, David I. W. Levin, and Alec Jacobson. Fast winding numbers for soups and clouds. *ACM Transactions on Graphics*, 37(4), 2018. 4
- [2] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P.A. Lensch. NeRD: Neural reflectance decomposition from image collections. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12664–12674, 2021. 2
- [3] Hanyu Chen, Bailey Miller, and Ioannis Gkioulekas. 3D reconstruction with fast dipole sums. *ACM Transactions on Graphics*, 43(6), 2024. 4, 5
- [4] J.M. Coughlan and A.L. Yuille. Manhattan world: compass direction from a single image by bayesian inference. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, pages 941–947, 1999. 1
- [5] François Darmon, Bénédicte Bascle, Jean-Clément Devaux, Pascal Monasse, and Mathieu Aubry. Improving neural implicit surfaces geometry with patch warping. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6250–6259, 2022. 1
- [6] Qiancheng Fu, Qingshan Xu, Yew Soon Ong, and Wenbing Tao. Geo-Neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. In *Advances in Neural Information Processing Systems*, pages 3403–3416. Curran Associates, Inc., 2022. 1, 2
- [7] Haoyu Guo, Sida Peng, Haotong Lin, Qianqian Wang, Guofeng Zhang, Hujun Bao, and Xiaowei Zhou. Neural 3D scene reconstruction with the manhattan-world assumption. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5501–5510, 2022. 1
- [8] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2 edition, 2004. 1
- [9] Fei Hou, Chiyu Wang, Wencheng Wang, Hong Qin, Chen Qian, and Ying He. Iterative poisson surface reconstruction (iPSR) for unoriented points. *ACM Transactions on Graphics*, 41(4), 2022. 4
- [10] Alec Jacobson, Ladislav Kavan, and Olga Sorkine-Hornung. Robust inside-outside segmentation using generalized winding numbers. *ACM Transactions on Graphics*, 32(4), 2013. 4
- [11] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engil Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 406–413, 2014. 5, 7
- [12] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics*, 32(3), 2013. 4
- [13] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Symposium on Geometry Processing*. The Eurographics Association, 2006. 4
- [14] Siyou Lin, Dong Xiao, Zuoqiang Shi, and Bin Wang. Surface reconstruction from point clouds without normals by parametrizing the gauss formula. *ACM Transactions on Graphics*, 42(2), 2022. 4

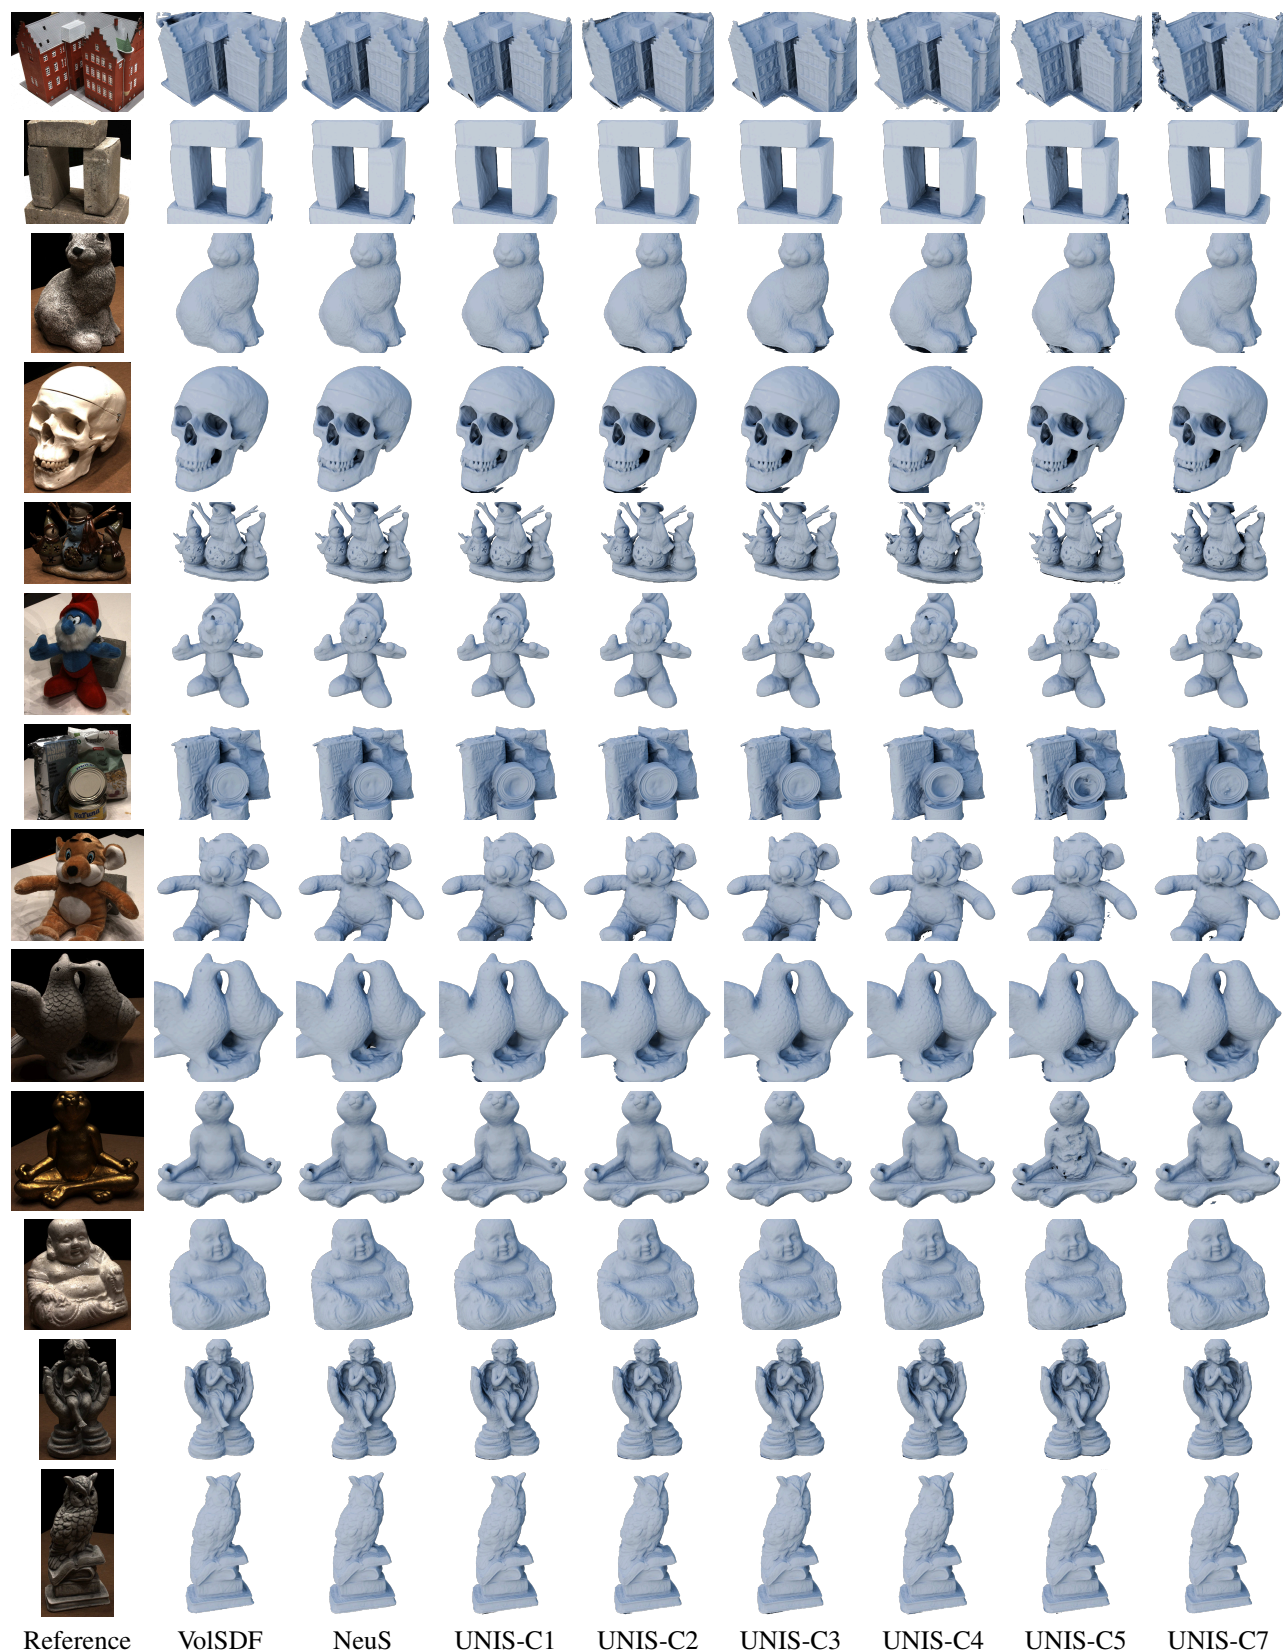


Figure 10. Visual results on the DTU dataset.

Method	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean
UNIS-C7	0.90	0.94	0.47	0.41	0.98	0.65	0.64	1.47	0.85	0.97	0.53	1.25	0.40	0.48	0.52	0.764

Method	Antique	Light-1	Light-2	Light-3	Orans	Ornament-1	Ornament-2	Rice	Sofa	Mean
UNIS-C7	1.96	2.73	8.87	3.02	2.22	4.52	5.04	2.50	2.98	3.760

Table 4. Chamfer distances ($\times 10^{-3}$) on the DTU dataset [11] (top) and the OmniObjects3D dataset [25] (bottom) under UNIS-C7, which uses GWN as implicit functions.




Figure 11. Visual results on the OmniObjects3D dataset.

- [15] Weizhou Liu, Xingce Wang, Haichuan Zhao, Xingfei Xue, Zhongke Wu, Xuequan Lu, and Ying He. Consistent point orientation for manifold surfaces via boundary integration. In *ACM SIGGRAPH 2024 Conference Papers*, New York, NY, USA, 2024. Association for Computing Machinery.
- [16] Weizhou Liu, Jiaze Li, Xuhui Chen, Fei Hou, Shiqing Xin,

- Xingce Wang, Zhongke Wu, Chen Qian, and Ying He. Diffusing winding gradients (DWG): A parallel and scalable method for 3D reconstruction from unoriented point clouds. *ACM Transactions on Graphics*, 44(2), 2025. 4, 5
- [17] Yuan Liu, Peng Wang, Cheng Lin, Xiaoxiao Long, Jiepeng Wang, Lingjie Liu, Taku Komura, and Wenping Wang.

- NeRO: Neural geometry and BRDF reconstruction of reflective objects from multiview images. *ACM Transactions on Graphics*, 42(4), 2023. [2](#)
- [18] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. SparseNeuS: Fast generalizable neural surface reconstruction from sparse views. In *Computer Vision – ECCV 2022*, pages 210–227, Cham, 2022. Springer Nature Switzerland. [1](#)
- [19] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. NeRV: Neural reflectance and visibility fields for relighting and view synthesis. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7491–7500, 2021. [2](#)
- [20] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5481–5490, 2022. [1](#)
- [21] 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 *Computer Vision – ECCV 2022*, pages 139–155, Cham, 2022. Springer Nature Switzerland. [1](#)
- [22] 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. In *Advances in Neural Information Processing Systems*, pages 27171–27183. Curran Associates, Inc., 2021. [1](#), [5](#)
- [23] Yiqun Wang, Ivan Skorokhodov, and Peter Wonka. HF-NeuS: Improved surface reconstruction using high-frequency details. In *Advances in Neural Information Processing Systems*, pages 1966–1978. Curran Associates, Inc., 2022. [1](#)
- [24] Tong Wu, Jia-Mu Sun, Yu-Kun Lai, and Lin Gao. DE-NeRF: Decoupled neural radiance fields for view-consistent appearance editing and high-frequency environmental relighting. In *ACM SIGGRAPH 2023 Conference Proceedings*, New York, NY, USA, 2023. Association for Computing Machinery. [2](#)
- [25] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, Dahua Lin, and Ziwei Liu. OmniObject3D: Large-vocabulary 3D object dataset for realistic perception, reconstruction and generation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 803–814, 2023. [5](#), [7](#)
- [26] Baixin Xu, Jiarui Zhang, Kwan-Yee Lin, Chen Qian, and Ying He. Deformable model-driven neural rendering for high-fidelity 3D reconstruction of human heads under low-view settings. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 17878–17888, 2023. [1](#)
- [27] Rui Xu, Zhiyang Dou, Ningna Wang, Shiqing Xin, Shuangmin Chen, Mingyan Jiang, Xiaohu Guo, Wenping Wang, and Changhe Tu. Globally consistent normal orientation for point clouds by regularizing the winding-number field. *ACM Transactions on Graphics*, 42(4), 2023. [4](#)
- [28] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. MonoSDF: Exploring monocular geometric cues for neural implicit surface reconstruction. In *Advances in Neural Information Processing Systems*, pages 25018–25032. Curran Associates, Inc., 2022. [1](#)
- [29] Yu-Jie Yuan, Yu-Kun Lai, Yi-Hua Huang, Leif Kobbelt, and Lin Gao. Neural radiance fields from sparse RGB-D images for high-quality view synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7):8713–8728, 2023. [1](#)
- [30] Chen Zhang, Wanjun Su, Qingshan Xu, Xinyao Liao, and Wenbing Tao. PG-NeuS: Robust and efficient point guidance for multi-view neural surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, to appear. [1](#)
- [31] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. PhysSG: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5449–5458, 2021. [2](#)