Surface Reconstruction of Transparent Objects from Gaussian Radiance Fields

Peter Kravár
Supervised by: Martin Madaras, Lukáš Gajdošech

Faculty of Informatics Masaryk University
Brno / Czechia

Abstract

Methods based on radiance fields have continually pushed the boundaries of novel view synthesis since their first appearance in 2020. The capabilities of radiance fields have since been extended from pure view synthesis to high quality surface reconstruction. However, these surface reconstruction methods generally display poor performance in scenes with transparent objects, producing holes and artifacts in the resulting geometry. Surface reconstruction fails even for radiance fields, which are able to represent the transparent objects accurately in view synthesis. Our work investigates the causes of such failure cases in scenes represented by 3D Gaussians and proposes an improvement to the mesh extraction process. We build upon the method of Gaussian Opacity Fields and utilize it for both scene optimization and geometry extraction. By supplying their geometry extraction pipeline with Gaussians of multiple training steps reached during scene optimization, we achieve a significant uplift in extracted mesh quality for transparent objects. In order to avoid the global reconstruction quality loss incurred by this approach, we segment transparent objects in the scene and selectively apply our pipeline modifications only to those objects.

Keywords: 3D Gaussian, radiance field, transparent object, surface reconstruction, mesh extraction

1 Introduction

Recent advancements in view synthesis based on radiance fields have facilitated new approaches to 3D mesh reconstruction from multi-view images. Among these, methods leveraging Neural Radiance Fields (NeRFs) [26] have demonstrated significant progress [10, 36, 24].

The field of view synthesis has since seen the rise of radiance fields with explicit Gaussian point cloud representations after the release of Gaussian Splatting [18]. GS and its successors preserve and often exceed the rendering quality of NeRFs, while offering real-time rendering and optimization which is orders of magnitude faster. Several Gaussian-based techniques [13, 15, 43, 5] have matched the state-of-the-art view synthesis and surface reconstruction results achieved by NeRFs.

However, both NeRF and GS-based approaches struggle to reconstruct surfaces of transparent objects. This is despite their ability to represent transparent objects convincingly in view synthesis, suggesting that surface reconstruction failures stem from a limitation in the geometry extraction process rather than a complete inability of radiance fields to model transparent materials.

The challenge of accurate 3D reconstruction for transparent objects extends beyond methods based on radiance fields. Conventional approaches, such as LiDAR scanning and photogrammetry, also struggle with transparent materials due to their reliance on depth estimation techniques that assume consistent surface reflections [17].

Extending the capabilities of Gaussian-based radiance fields further to allow surface reconstruction of transparent objects is significant to the broader field of surface reconstruction. Even the more constrained problem of depth estimation for transparent materials has major impact for robotic applications such as object grasping [16].

We propose a simple approach for improving mesh quality and demonstrate our improvements in a qualitative evaluation, comparing the mesh outputs of our pipeline to the unmodified GOF.

2 Related Work

This chapter begins with an overview of early radiance fields methods invented for view synthesis. It summarizes subsequent developments which established the research field. The following subchapter explains how the capabilities of radiance fields were extended from pure view synthesis to surface reconstruction, and what improvements have been proposed since then. The last subchapter discusses existing techniques for segmentation of radiance fields.

2.1 Novel View Synthesis

Neural Radiance Fields (NeRF) [26], in which the authors utilized a multilayered perceptron (MLP) to learn a 3D scene by optimizing photometric loss through volume rendering, emerged as a critical innovation in the area of novel view synthesis.

The original NeRF implementation was computationally demanding for both training and rendering. Hence a major line of research focused on improving viewing performance by baking the NeRF into different representations after training [42, 14]. Poor training performance was addressed by later works [12, 32], which showed that the MLP in NeRF was being used as an inefficient volumetric data structure and could be replaced by a voxel grid. While voxel grids do achieve significantly faster inference than NeRF, they consume large amounts of memory and struggle to represent fine details due to their limited resolution. Nonetheless, their emergence sparked a broader search for more efficient and expressive data structures for radiance field representation. Later approaches attempted to replace the MLP with a wide variety of data structures, even large amounts of smaller MLPs [31], and 4D tensors [4]. Meanwhile, a major breakthrough in NeRF performance was achieved by multiscale hash encoding in InstantNGP [27], yet the search for an efficient data structure continued and trended towards increasingly more explicit representations. Such representations include neural point clouds [40], and finally sets of Gaussians.

3D Gaussian Splatting [18], unlike NeRF, does not rely on a slow ray marching approach for synthesizing images. 3D Gaussians are instead projected into 2D, sorted by distance from the camera, and overlaid to produce images in real-time. Due to their simple representation, Gaussian point clouds are easy to integrate into common 3D software. Their practical utility and superior performance made Gaussian representations the focus of subsequent research. However, research also continues with implicit representations. Numerous findings made with either representation have since been implemented using the other, making both lines of work complementary and mutually relevant.

2.2 Radiance Fields Surface Reconstruction

NeRF allows for depth maps to be obtained for free as a byproduct of rendering even in the original implementation [26]. The problem of surface reconstruction can therefore be reduced to mesh extraction from depth maps by applying a TSDF mesh extraction algorithm [28]. However, later methods commonly use signed distance functions and other implicit abstractions to represent the surface prior to mesh extraction in order to mitigate surface artifacts [14, 29, 36, 41, 30]. Despite some NeRF-based surface reconstruction methods achieving impressively high detail [24], they commonly require upwards of 12 hours of training per scene on expensive hardware.

3D Gaussian Radiance Fields make geometry extraction more challenging due to their explicit nature and unclear surface boundaries. SuGaR [13] mitigates the problem of unclear surface boundary by introducing a regularization term into the training pipeline. The regularization term constrains Gaussians to form more consistent surfaces, additionally allowing for estimation of a normal as the shortest axis for any given Gaussian. The normal information enables the use of Poisson mesh extraction, yielding meshes of higher quality. DN-Splatter [34] extends their approach further by integrating depth and normal cues into the optimization process, improving alignment with true scene geometry and enhancing surface smoothness.

The key differences in 3DGS-based surface reconstruction methods often lie in their approaches to estimation of normal data from Gaussians. PGSR [5] uses multiple regularizers and a mix of depth-estimation techniques to achieve view-consistent depth and normal estimations. Instead of approximating normals from 3D Gaussians, 2DGS [15] chooses to represent the scene with 2D Gaussians and obtain normals directly. The surface can then be approximated by optimizing a set of small 2D oriented disks to lie on the surface while matching its normals. Gaussian Opacity Fields (GOF) [43] proposes a direct level-set extraction approach, eliminating the need for Poisson reconstruction altogether and enabling more efficient surface reconstruction in unbounded scenes. Numerous other methods aim to guide the mesh extraction process using priors, more complex regularizers, and multimodal data [37, 35, 22].

Surface reconstruction of transparent objects is a largely underexplored problem in context of radiance fields. Several methods have attempted to leverage NeRFs [16, 19, 9, 11] and Gaussian Splatting [1, 20] for depth estimation and downstream tasks such as robotic grasping of transparent objects. However, these approaches are limited to producing imprecise depth maps and are not capable of generating explicit surface meshes.

True surface reconstruction was attempted by few methods, each of which introduces specific limitations. NeRRF [7] requires object silhouette as an additional input. Nu-NeRF [33] is computationally demanding. αSurf [38], despite its very impressive results, builds on a voxel-based representation [12] and thus inherits its associated limitations in scene resolution and memory efficiency. To the best of our knowledge, the only Gaussian-based method that has demonstrated success in reconstructing transparent surfaces under unconstrained conditions is Car-GS [23], which focuses specifically on car models. As of the time of writing, the implementation of Car-GS has not been publicly released.

2.3 3D Segmentation of Radiance Fields

Most existing radiance field segmentation methods work in image space by utilizing the SAM segmentation model [21] to obtain semantic maps from ground truth images. The semantic maps are then extrapolated into 3D in various ways. Certain Gaussian-based methods use the maps to train additional 3D Gaussian features [3]. Runtime segmentation can be then guided by user input in the form of one-shot click prompts [8, 3], or open-vocabulary prompts [25, 6, 39].

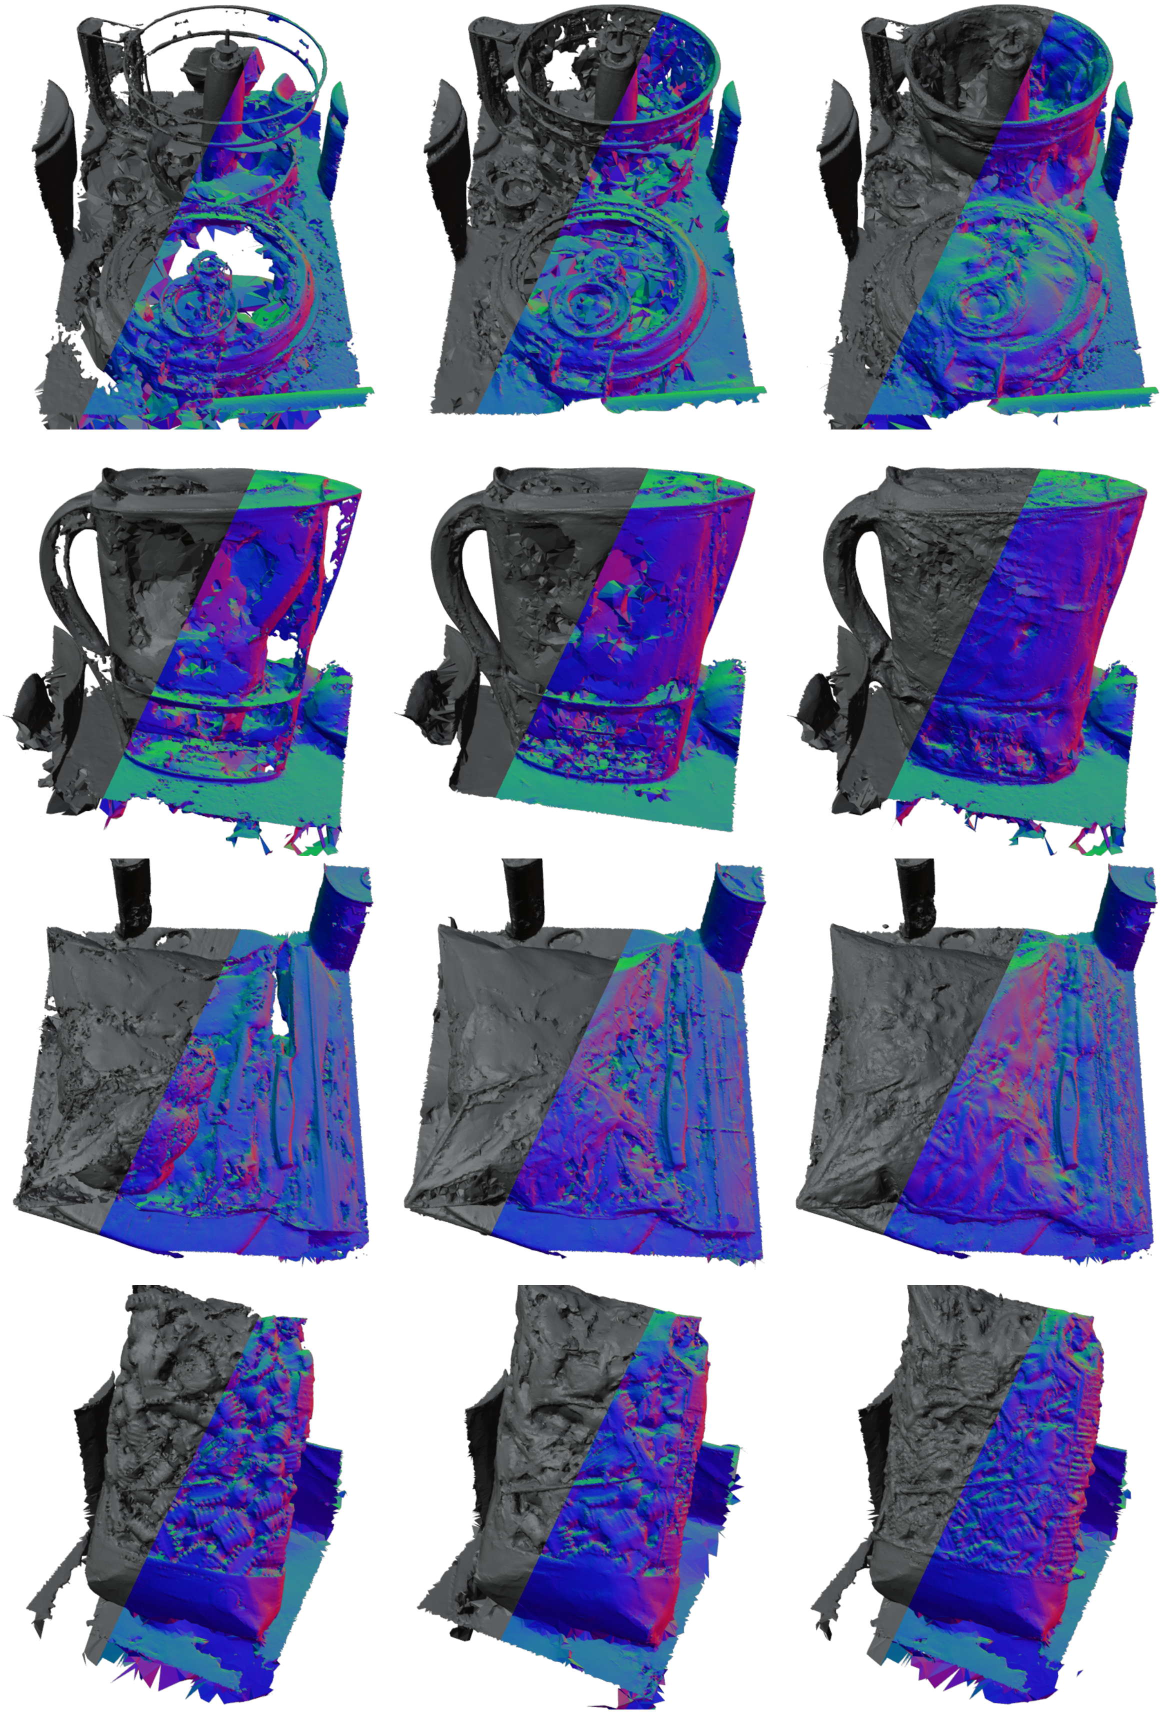

(a) 30k, SDF=0.5

(b) 30k, SDF=0.1

(c) Ours

Figure 1: Visual comparison of transparent object surface reconstruction on the "Counter" scene of mip-NeRF 360 dataset.

3 Method

In this section, we examine the GOF mesh extraction method. We observe its failure on transparent objects, make an argument about the cause, and perform experiments to support our claims. Based on the results of our experiments, we devise a modification to the mesh extraction pipeline, which extracts the surfaces of transparent objects more accurately.

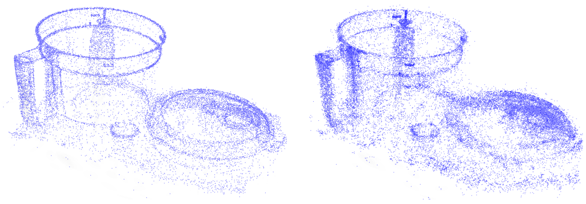

(a) 30k (b) 1k2k3k4k

Figure 2: Point cloud of Gaussian origins

3.1 Initial Findings

To assess the limitations of existing approaches in mesh reconstruction for transparent objects, we conduct experiments using Gaussian Opacity Fields (GOF) [43]. This method was selected as our baseline due to its high quality, performance in training 3D Gaussian scene representations, and its robust mesh extraction pipeline.

Firstly, we replicate the mesh extraction experiment conducted in the original GOF study, using the mip-NeRF 360 dataset [2]. Our experimental results confirm that GOF produces high quality meshes of the scenes, but fails to reconstruct transparent objects (see Fig. 1a). The issue persists when using a relaxed signed distance function (SDF) level of 0.1 (see Fig. 1b).

The reconstruction failure stems from the way in which transparent materials are represented in a Gaussian radiance field. During training, 3D Gaussians change to match ground truth image observations. Because transparent objects produce low contribution to photometric loss, they are modeled by Gaussians with low opacities. Note that this trait was inherited by most GS-based methods [18]. The behavior is expected and poses no issue in view synthesis. However, the assumption that visual opacity can be used as a measure for physical density is incorrect for surface reconstruction of transparent objects. We argue that this is the main cause of reconstruction failures in currently available state of the art methods for general surface reconstruction [43, 15, 5]. Novel works on surface reconstruction of transparent objects align with and support our reasoning [38, 23].

3.2 Reconstructing Transparent Surfaces

Instead of making changes to the training pipeline, we attempt to improve surface reconstruction by modifying the Gaussian point cloud after training. Increasing the opacity of 3D Gaussians before surface reconstruction is insufficient for substantial improvements in mesh quality (see Fig. 7d). A closer inspection of the trained Gaussian point cloud reveals that the Gaussians of transparent objects are generated too sparsely to allow for meaningful surface estimation (see Fig. 2a), likely due to the low contribution of transparent objects to photometric loss, as explained earlier.

We obtain a denser point cloud by sampling the Gaussians from multiple Gaussian Radiance Fields trained on the same set of images. Instead of fully training the same scene with different seeds and sampling across training runs, we sample 3D Gaussians across iterations of the same training run in order not to increase training time.

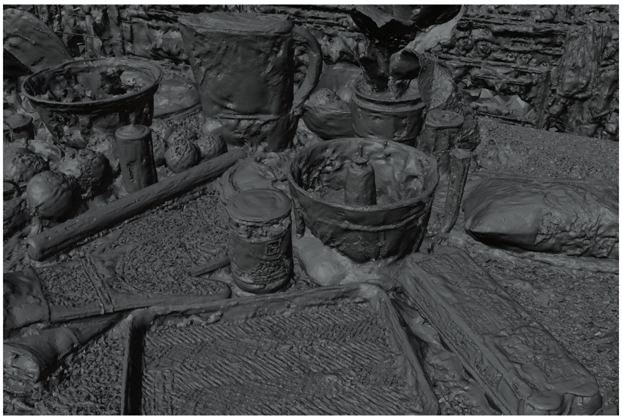

Figure 3: Surface reconstruction from merged Gaussians of training iterations 1k2k3k4k.

The process of training a Gaussian scene representation takes 30,000 iterations, in which Gaussians randomly shift their positions by small increments until they converge. Our assumption is that after a sufficiently high number of steps, Gaussians lie in close proximity to the surface and can be sampled to approximate it. Since the features of 3D Gaussians need to be mature enough to approximate the surface reliably, yet are most volatile in early training stages, we choose iterations 1k, 2k, 3k, and 4k for sampling. We merge the Gaussians of training iterations 1k, 2k, 3k, and 4k into a single point cloud, and reexamine the positions of their origins (see Fig. 2b). The merged point cloud is dense even in areas, which are sparse in the completely trained representation.

With this simple modification to the radiance field, the extraction pipeline yields greatly superior mesh quality for transparent objects (see Fig. 3). However, overlaying multiple training steps and using low training iterations decreases the mesh quality for other objects in the scene.

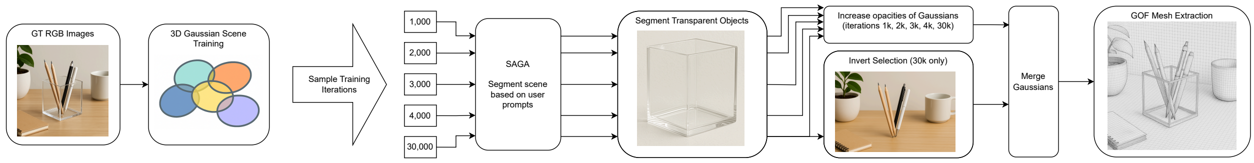

Figure 4: An overview of our surface reconstruction pipeline.

3.3 Complete Method

In order to avoid the incurred reconstruction errors, we segment the scene and extract Gaussians belonging to transparent objects. To perform the segmentation, we use SAGA [3], a SAM-based 3D Gaussian segmentation tool to segment transparent objects. The method requires the user to manually click once on each transparent object which is to be segmented.

Figure 5: Segmentation process using SAGA.

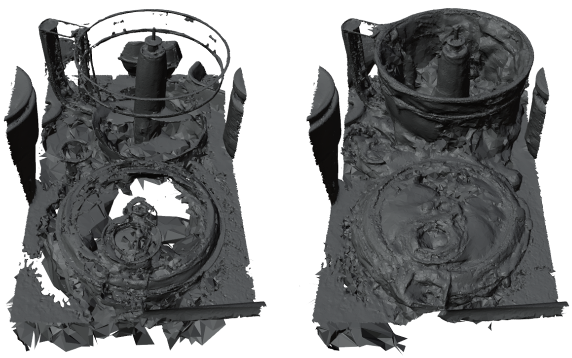

(a) GOF default settings (b) Ours

Figure 6: Cropped regions of extracted meshes.

To obtain the final Gaussian representation ready for mesh extraction, we sample the training process at iterations 1k, 2k, 3k, 4k, 30k. We then segment the Gaussians of transparent objects (see Fig. 5) from all of the selected training steps, increase their opacities to 1, merge them together, and add the non-transparent background from iteration 30k as shown in Fig.4. The resulting mesh quality is greatly improved (see Fig. 6).

We evaluate our results qualitatively on the mip-NeRF 360 "Counter" scene. Our experiments were carried out on an NVIDIA A100 80GB VRAM GPU. SAGA models were trained for 5k iterations with default settings. Mesh cutouts shown in Fig. 1 and Fig. 6 were obtained by removing all vertices outside manually established bounding boxes which span areas of interest.

4 Limitations & Future Work

One major limitation of our approach is the amount of necessary user interaction. Hence a possible improvement to our method would be to use an open-vocabulary segmentation approach and use predefined prompts for transparent objects.

Another limitation is performance. Despite the segmentation happening in under 10 milliseconds, SAGA requires upwards of an hour of total training time for all sampled scenes to learn semantic Gaussian features used when segmenting user prompts.

Lastly, our fixed choice of iterations which are to be segmented is merely a proof of concept. It can potentially be extended to an automated system, which identifies sparsely populated areas and samples Gaussians adaptively across all training iterations.

5 Conclusions

We inspected Gaussian Opacity Fields [43], an existing method of high quality surface reconstruction for 3D Gaussian radiance fields. We made the argument that erroneously representing geometric density with visual opacity is the primary cause of its inability to capture surfaces of transparent objects. We supported this claim with an experiment, in which we improved the extracted mesh quality of transparent objects by densifying their Gaussian representations and increasing their opacity. Additionally, we showed that such improvements can be obtained by simply sampling the 3D Gaussians across multiple training iterations of the same training run, rather than training the scene multiple times with different seeds. Lastly, we demonstrated the ability of an existing 3D Gaussian segmentation technique, SAGA [3], to obtain 3D Gaussians of transparent objects, which we leveraged to perform mesh extraction of higher quality for the entire scene.

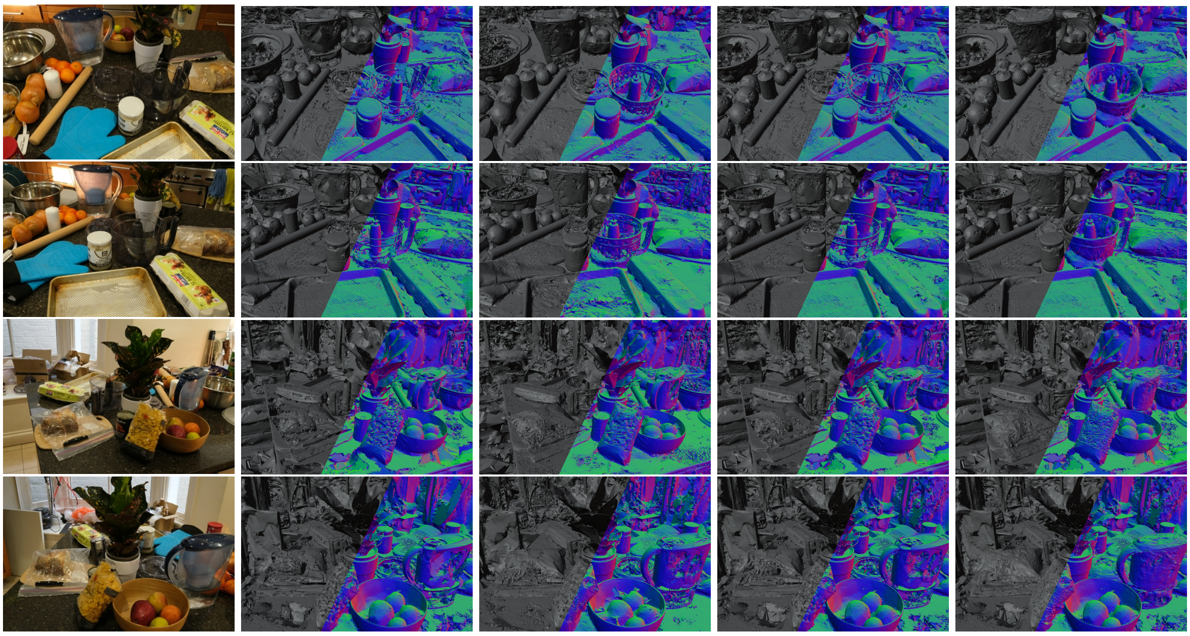

(a) GT (b) 30k SDF=0.5 (c) 30k SDF=0.1 (d) 30k alpha=1, SDF=0.5 (e) Ours

Figure 7: Comparison of GT and different training configurations.

References

[1] Aviral Agrawal, Ritaban Roy, Bardienus Pieter Duisterhof, Keerthan Bhat Hekkadka, Hongyi Chen, and Jeffrey Ichnowski. Clear-splatting: Learning residual gaussian splats for transparent object manipulation. In *RoboNerF: 1st Workshop On Neural Fields In Robotics at ICRA 2024*, 2024.

[2] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5460–5469, 2022.

[3] Jiazhong Cen, Jiemin Fang, Chen Yang, Lingxi Xie, Xiaopeng Zhang, Wei Shen, and Qi Tian. Segment any 3d gaussians, 2025.

[4] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 333–350, Cham, 2022. Springer Nature Switzerland.

[5] 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. *IEEE Transactions on Visualization and Computer Graphics*, page 1–12, 2024.

[6] Runnan Chen, Xiangyu Sun, Zhaoqing Wang, Youquan Liu, Jiepeng Wang, Lingdong Kong, Jiankang Deng, Mingming Gong, Liang Pan, Wenping Wang, and Tongliang Liu. Ovgaussian: Generalizable 3d gaussian segmentation with open vocabularies, 2024.

[7] Xiaoxue Chen, Junchen Liu, Hao Zhao, Guyue Zhou, and Ya-Qin Zhang. Nerrf: 3d reconstruction and view synthesis for transparent and specular objects with neural refractive-reflective fields, 2023.

[8] Seokhun Choi, Hyeonseop Song, Jaechul Kim, Taehyeong Kim, and Hoseok Do. Click-gaussian: Interactive segmentation to any 3d gaussians. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, editors, *Computer Vision – ECCV 2024*, pages 289–305, Cham, 2025. Springer Nature Switzerland.

[9] Qiyu Dai, Yan Zhu, Yiran Geng, Ciyu Ruan, Jiazhao Zhang, and He Wang. Graspnerf: Multiview-based 6-dof grasp detection for transparent and specular objects using generalizable nerf. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1757–1763, 2023.

[10] 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.

[11] Bardienus P. Duisterhof, Yuemin Mao, Si Heng Teng, and Jeffrey Ichnowski. Residual-nerf: Learning residual nerfs for transparent object manipulation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13918–13924, 2024.

[12] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5491–5500, 2022.

[13] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5354–5363, 2024.

[14] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5855–5864, 2021.

[15] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, SIGGRAPH '24, New York, NY, USA, 2024. Association for Computing Machinery.

[16] Jeffrey Ichnowski, Yahav Avigal, Justin Kerr, and Ken Goldberg. Dex-nerf: Using a neural radiance field to grasp transparent objects. In Aleksandra Faust, David Hsu, and Gerhard Neumann, editors, *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pages 526–536. PMLR, 08–11 Nov 2022.

[17] A. Karami, R. Battisti, F. Menna, and F. Remondino. 3d digitization of transparent and glass surfaces: State of the art and analysis of some methods. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLIII-B2-2022:695–702, 2022.

[18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4), July 2023.

[19] Justin Kerr, Letian Fu, Huang Huang, Yahav Avigal, Matthew Tancik, Jeffrey Ichnowski, Angjoo Kanazawa, and Ken Goldberg. Evo-nerf: Evolving nerf for sequential robot grasping of transparent objects. In Karen Liu, Dana Kulic, and Jeff Ichnowski, editors, *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pages 353–367. PMLR, 14–18 Dec 2023.

[20] Jeongyun Kim, Jeongho Noh, Dong-Guw Lee, and Ayoung Kim. Transplat: Surface embedding-guided 3d gaussian splatting for transparent object manipulation, 2025.

[21] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3992–4003, 2023.

[22] Meenakshi Krishnan, Liam Fowl, and Ramani Duraiswami. 3d gaussian splatting with normal information for mesh extraction and improved rendering, 2025.

[23] Congcong Li, Jin Wang, Xiaomeng Wang, Xingchen Zhou, Wei Wu, Yuzhi Zhang, and Tongyi Cao. Cargs: Addressing reflective and transparent surface challenges in 3d car reconstruction, 2025.

[24] 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 *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8456–8465, 2023.

[25] Siyun Liang, Sen Wang, Kunyi Li, Michael Niemeyer, Stefano Gasperini, Nassir Navab, and Federico Tombari. Supergseg: Open-vocabulary 3d segmentation with structured super-gaussians, 2024.

[26] 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. *Commun. ACM*, 65(1):99–106, December 2021.

[27] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022.

[28] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE International Symposium on Mixed and Augmented Reality*, pages 127–136, 2011.

[29] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5569–5579, 2021.

[30] Marie-Julie Rakotosaona, Fabian Manhardt, Diego Martín Arroyo, Michael Niemeyer, Abhijit Kundu, and Federico Tombari. Nerfmeshing: Distilling neural radiance fields into geometrically-accurate 3d meshes. In *International Conference on 3D Vision, 3DV 2024, Davos, Switzerland, March 18-21, 2024*, pages 1156–1165. IEEE, 2024.

[31] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14315–14325, 2021.

[32] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5449–5459, 2022.

[33] Jia-Mu Sun, Tong Wu, Ling-Qi Yan, and Lin Gao. Nu-nerf: Neural reconstruction of nested transparent objects with uncontrolled capture environment. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2024)*, 43(6), 2024.

[34] Matias Turkulainen, Xuqian Ren, Iaroslav Melekhov, Otto Seiskari, Esa Rahtu, and Juho Kannala. Dn-splatter: Depth and normal priors for gaussian splatting and meshing. In *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*, pages 2421–2431, February 2025.

[35] Jiepeng Wang, Yuan Liu, Peng Wang, Cheng Lin, Junhui Hou, Xin Li, Taku Komura, and Wenping Wang. Gaussurf: Geometry-guided 3d gaussian splatting for surface reconstruction. *arXiv preprint arXiv:2411.19454*, 2024.

[36] 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 *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, Red Hook, NY, USA, 2021. Curran Associates Inc.

[37] Qianyi Wu, Jianmin Zheng, and Jianfei Cai. Surface reconstruction from 3d gaussian splatting via local structural hints. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, editors, *Computer Vision – ECCV 2024*, pages 441–458, Cham, 2025. Springer Nature Switzerland.

[38] Tianhao Walter Wu, Hanxue Liang, Fangcheng Zhong, Gernot Riegler, Shimon Vainer, Jiankang Deng, and Cengiz Oztireli. Alphasurf: Implicit surface reconstruction for semi-transparent and thin objects with decoupled geometry and opacity. In *3DV*, 2025.

[39] Yanmin Wu, Jiarui Meng, Haijie Li, Chenming Wu, Yahao Shi, Xinhua Cheng, Chen Zhao, Haocheng Feng, Errui Ding, Jingdong Wang, and Jian Zhang. Opengaussian: Towards point-level 3d gaussian-based open vocabulary understanding, 2024.

[40] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5428–5438, 2022.

[41] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P. Srinivasan, Richard Szeliski, Jonathan T. Barron, and Ben Mildenhall. Bakedsdf: Meshing neural sdfs for real-time view synthesis. In *ACM SIGGRAPH 2023 Conference Proceedings*, SIGGRAPH '23, New York, NY, USA, 2023. Association for Computing Machinery.

[42] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenoctrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5732–5741, Montreal, QC, Canada, October 2021. IEEE.

[43] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Trans. Graph.*, 43(6), November 2024.