

Mirror-3DGS: Incorporating Mirror Reflections into 3D Gaussian Splatting

Jiarui Meng^{1*}, Haijie Li^{1*}, Yanmin Wu¹, Qiankun Gao¹, Shuzhou Yang¹, Jian Zhang^{1†}, Siwei Ma^{2†}

¹ School of Electronic and Computer Engineering, Peking University, China

² School of Computer Science, Peking University, China

{mengjiarui, 2301212733, wuyanmin, gqk, szyang}@stu.pku.edu.cn, {zhangjian.sz, swma}@pku.edu.cn

Abstract—3D Gaussian Splatting (3DGS) has significantly advanced 3D scene reconstruction and novel view synthesis. However, like Neural Radiance Fields (NeRF), 3DGS struggles with accurately modeling physical reflections, particularly in mirrors, leading to incorrect reconstructions and inconsistent reflective properties. To address this challenge, we introduce Mirror-3DGS, a novel framework designed to accurately handle mirror geometries and reflections, thereby generating realistic mirror reflections. By incorporating mirror attributes into 3DGS and leveraging plane mirror imaging principles, Mirror-3DGS simulates a mirrored viewpoint from behind the mirror, enhancing the realism of scene renderings. Extensive evaluations on both synthetic and real-world scenes demonstrate that our method can render novel views with improved fidelity in real-time, surpassing the state-of-the-art Mirror-NeRF, especially in mirror regions.

Index Terms—Gaussian Splatting, Mirror Scene.

I. INTRODUCTION

The 3D reconstruction and novel view synthesis are crucial in numerous applications, including the movie industry, computer gaming, virtual reality, and autonomous navigation. However, accurately capturing and rendering reflections from mirrors, which are prevalent in real-world scenes and significantly increase the visual complexity, remains a significant challenge for these technologies.

Traditional methods [1]–[5] simulate surface reflections by decomposing material properties and lighting conditions to handle scenes with mirrors. However, they often struggle to accurately infer surfaces and mirror components, especially in managing pure mirror reflections. With advancements in Neural Radiance Fields (NeRF), new methods [2], [6], [7] have emerged. NeRFReN [6] enhances rendering quality by using independent radiance fields for reflective and transmissive parts. Mirror-NeRF [7] employs Whitted Ray Tracing [8] for physically accurate mirror reflections, achieving high realism in novel view synthesis. Despite these improvements, NeRF-based methods face challenges such as long training times and slow rendering speeds, limiting their real-time application.

Recently, 3D Gaussian Splatting (3DGS) [9] has emerged as a promising approach for real-time, high-fidelity 3D scene rendering. Instead of the slow volumetric rendering of NeRF,

* means equal contribution and † means corresponding authors. This work was supported in part by the National Natural Science Foundation of China no. 61931014, U21B2012, in part by the Beijing Natural Science Foundation no. L242014, and in part by the Fundamental Research Funds for the Central Universities.




Fig. 1. The synthesized novel views and depth maps. 3DGS confuses mirror reflection, resulting in erroneous mirror depth. Ours properly handles content in the mirror with depth map almost identical to GT.

3DGS explicitly represents scenes using 3D Gaussians and rasterizes 3D Gaussians onto the image plane using splatting techniques, achieving significant improvements in training and rendering speeds [10]. However, 3DGS tends to treat mirror reflections as separate scenes, constructing Gaussian points for them as if they physically exist, leading to inaccurate capture of these virtual images within mirrors during reconstruction and rendering, affecting the realism and visual coherence of the scene, as shown in Fig. 1.

To address these challenges, we propose Mirror-3DGS, a novel rendering framework based on 3DGS, aimed at producing high-fidelity novel views in mirror-containing scenes. *Our Mirror-3DGS is centered on the principle that the virtual image presented on the mirror is equivalent to observing the real world from behind the mirror.* In general, We first identify 3D Gaussians representing the mirror, then observe from both the current and mirrored viewpoints, integrating these perspectives to render the final image. We introduce an attribute to signify mirror-like properties to 3D Gaussians and estimate the mirror plane equation, allowing us to derive the mirror transformation matrix and mirrored camera parameters. Our method involves a two-stage training process: 1) During the first stage, We learn a rough 3D Gaussian representation of the scene without mirror content and estimate the mirror plane equation. The mirror content is replaced with a fixed color based on the ground truth mirror mask to better learn mirror attributes and avoid interference from virtual images. 2) In the second stage, We filter Gaussians with high mirror properties

for precise estimation of the mirror equation. By integrating views from both the current and mirrored viewpoints, we optimize the rendering quality of the entire scene. To ensure stability and improve the accuracy of plane estimation, we introduce planar consistency and depth constraints throughout the training process. These constraints help achieve more realistic and refined rendering results. Experiments on three synthetic and three real-world scenes demonstrate that Mirror-3DGS delivers rendering quality comparable to the state-of-the-art Mirror-NeRF, while significantly reducing training time and enabling real-time rendering. Our contributions can be summarized as follows:

- **Innovative Rendering Framework for Mirror Scenes.** We present Mirror-3DGS, a novel rendering framework that addresses the limitations of 3DGS in handling mirror scenes, enabling realistic synthesis of novel views.
- **Precise Physical Modeling of Mirror Reflection.** We leverage the principle of plane mirror imaging to enhance 3DGS, enabling accurate modeling of physical space and producing high-fidelity rendering results.
- **Extensive Experiments.** We conduct comprehensive experiments across synthetic and real scenes to validate our proposed method. The results demonstrate our superiority in real-time rendering of high-quality novel views, particularly surpassing the state-of-the-art in mirror regions.

II. METHODOLOGY

A. Preliminaries

3D Gaussian Splatting (3DGS) is a state-of-the-art approach in novel view synthesis, utilizing the splatting technique [11] for efficient real-time rendering. In 3DGS, the scene is represented using anisotropic Gaussians characterized by center positions $\boldsymbol{\mu} \in \mathbb{R}^3$, covariances $\boldsymbol{\Sigma} \in \mathbb{R}^{3 \times 3}$, color defined by spherical harmonic (SH) coefficients $\boldsymbol{c} \in \mathbb{R}^{3 \times (k+1)^2}$ (with k denoting the SH order), and opacity $\alpha \in \mathbb{R}^1$. The 3D Gaussian can be queried as follows:

$$G(\mathbf{x}) = e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}, \quad (1)$$

where \mathbf{x} represents the position of the query point. Subsequently, an efficient 3D to 2D Gaussian mapping [12] is employed to project the Gaussian onto the image plane:

$$\boldsymbol{\mu}' = \mathbf{P}\mathbf{W}\boldsymbol{\mu}, \quad \boldsymbol{\Sigma}' = \mathbf{J}\mathbf{W}\boldsymbol{\Sigma}\mathbf{W}^\top \mathbf{J}^\top, \quad (2)$$

where $\boldsymbol{\mu}'$ and $\boldsymbol{\Sigma}'$ separately represent the 2D mean position and covariance of the projected 3D Gaussian. \mathbf{P} , \mathbf{W} and \mathbf{J} denote the projective transformation, viewing transformation, and Jacobian of the affine approximation of \mathbf{P} , respectively. The color of the pixel on the image plane, denoted by \mathbf{p} , uses a typical neural point-based rendering [13], [14]. Let $\mathbf{C} \in \mathbb{R}^{H \times W \times 3}$ represent the color of rendered image where H and W represents the height and width of images, the rendering process outlined as follows:

$$\begin{aligned} \sigma_i &= \alpha_i e^{-\frac{1}{2}(\mathbf{p}-\boldsymbol{\mu}')^\top \boldsymbol{\Sigma}'^{-1}(\mathbf{p}-\boldsymbol{\mu}')}, \\ \mathbf{C}(\mathbf{p}) &= \sum_{i=1}^N c_i \sigma_i \prod_{j=1}^{i-1} (1 - \sigma_j), \end{aligned} \quad (3)$$




Fig. 2. **The underlying principles in handling mirrored contents.** **a)** Mirror-NeRF employs a technique involving ray tracing and sampling to achieve mirror reflection. **b)** 3DGS mistakenly considers objects reflected by the mirror to be placed behind it, resulting the floaters behind the mirror. **c)** Our proposed Mirror-3DGS, captures the reflected objects by simulating a novel viewpoint situated behind the mirror.

where N represents the number of sample Gaussians that overlap the pixel \mathbf{p} , c_i and α_i denote the color and opacity of the i -th Gaussian, respectively.

B. Mirror-aware 3D Gaussian Representation

In the real world, a point’s color varies with different viewpoints due to its material and environmental lighting. While NeRF and 3DGS use view-dependent MLPs and spherical harmonics to simulate this, mirrors’ extreme specularity amplifies view-dependent inconsistencies. We review existing solutions and introduce our reflection-agnostic Mirror-3DGS.

- **Mirror-NeRF:** Mirror-NeRF achieves mirror modeling by learning a unified probabilistic reflectance field and simulates mirror reflection through physics-based ray tracing, enabling accurate mirror rendering, as illustrated in Fig. 2(a). However, its limitations lie in the intensive computational cost of ray tracing and the prolonged NeRF training time.

- **3DGS:** 3DGS is based on explicit gaussian spheres representation and alpha-blending along ray directions, tending to generate points behind mirrors as shown in Fig. 2(b). This violates physical laws and undermines the rendering quality of mirror back-side scenes.

- **Mirror-3DGS (Ours):** We innovatively introduce a mirror imaging method – “**moving behind the mirror for a view of the world**”, where the virtual image presented on the mirror is equivalent to observing from the real world from behind the mirror, as illustrated in Fig. 2(c). Therefore, we avoid the time-consuming ray tracing required by Mirror-NeRF.

More specifically, **1)** we render an image from the original viewpoint, *i.e.* observing from in front of the mirror, and retain the parts outside the mirror; **2)** we mirror the original viewpoint to behind the mirror, rendering an image from the virtual viewpoint, and only keep the parts of the mirror (see Sec. II-C); **3)** finally, we merge these two images to form the final result (see Sec. II-C).

To realize the first step above, we need to enable mirror detection. We add a learnable mirror attribute $m \in [0, 1]$ for each Gaussian in the original 3DGS, representing the probability of it being a mirror. The rendering process is similar to color rendering:

$$\mathbf{M}(\mathbf{p}) = \sum_{i=0}^N m_i \sigma_i \prod_{j=1}^{i-1} (1 - \sigma_j), \quad (4)$$

where $\mathbf{M} \in \mathbb{R}^{H \times W \times 1}$ represents the rendered mirror mask, and the rest of the symbols are defined as in Eq. (3). We utilize

the GT mirror mask \mathbf{M}_{gt} provided by the dataset to supervise the learning of the mirror attribute:

$$\mathcal{L}_{mask} = \mathcal{L}_1(\mathbf{M}, \mathbf{M}_{gt}). \quad (5)$$

C. Virtual Mirrored Viewpoint Construction and Image Fusion

Mirror Parameterization. In the previous section, we acquire the mirror attribute for each Gaussian. Using this attribute and opacity, we filter out Gaussians belonging to the mirror to construct a plane in 3D space. We adopt the parameterization of an infinite plane [15]: $\pi = (\mathbf{n}_\pi^\top, d)$, where $\mathbf{n}_\pi = (a, b, c)^\top$ represents the normal of the plane, and $\frac{d}{\|\mathbf{n}_\pi\|}$ represents the distance from the origin of the plane to the world coordinate origin. Thereby, any point $\mathbf{p} = (x, y, z)^\top$ on the plane satisfies:

$$\mathbf{n}_\pi^\top \mathbf{p} + d = 0. \quad (6)$$

The optimization of the plane π is presented in [Sec. II-D](#).

Virtual Mirrored Viewpoint Construction. Based on the above plane equation, given an original observation viewpoint $\mathbf{P}_o \in \mathbb{R}^{4 \times 4}$, our goal is to obtain the mirrored viewpoint $\mathbf{P}_m \in \mathbb{R}^{4 \times 4}$ of this viewpoint relative to the plane. This transformation can be formulated as:

$$\mathbf{P}_m = \mathbf{T}_m \mathbf{P}_o, \quad (7)$$

where we define \mathbf{T}_m as the mirror transformation matrix, represented as:

$$\mathbf{T}_m = \begin{bmatrix} 1 - 2a^2 & -2ab & -2ac & -2ad \\ -2ab & 1 - 2b^2 & -2bc & -2bd \\ -2ac & -2bc & 1 - 2c^2 & -2cd \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (8)$$

Image Fusion. The image from the mirror viewpoint is obtained using the same rendering equation. We acquire the rendered image $\mathbf{C}_{\mathbf{P}_o} \in \mathbb{R}^{H \times W \times 3}$ from the original viewpoint, the rendered image $\mathbf{C}_{\mathbf{P}_m} \in \mathbb{R}^{H \times W \times 3}$ from the mirror viewpoint, and the mirror mask \mathbf{M} . The final composite image is fused as follows:

$$\mathbf{C}_{fuse} = \mathbf{C}_{\mathbf{P}_o} \odot (\mathbb{1} - \mathbf{M}) + \mathbf{C}_{\mathbf{P}_m} \odot \mathbf{M}, \quad (9)$$

where \odot means element-wise multiplication.

D. Two-Stage Training Strategy

We use a two-stage training strategy, as illustrated in [Fig. 3](#).

Stage 1: The training objective of the first stage is to obtain a rough Gaussian representation without mirror content and to optimize the mirror equation. We use the GT mirror mask provided by the dataset to erase the mirror content in the GT image \mathbf{C}_{gt} and fill it with red, denoted as \mathbf{C}_{wo}^m . This is done to optimize the mirror attribute of Gaussians and to avoid generating false-Gaussians inside the mirror. The loss for rendering RGB images is consistent with 3DGS [9], *i.e.*, L1 loss and D-SSIM term:

$$\mathcal{L}_{rgb}^1 = (1-\gamma)\mathcal{L}_1(\mathbf{C}_{\mathbf{P}_o}, \mathbf{C}_{wo}^m) + \gamma\mathcal{L}_{D-SSIM}(\mathbf{C}_{\mathbf{P}_o}, \mathbf{C}_{wo}^m) \quad (10)$$

where γ is the weight, and $\mathbf{C}_{\mathbf{P}_o}$ is the rendered image from the original viewpoint, as defined in [Sec. II-C](#).

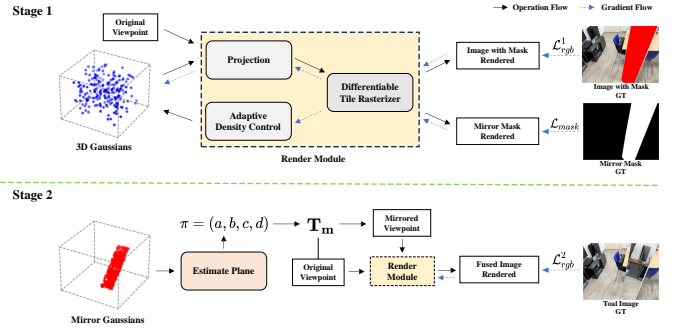


Fig. 3. **The illustration of our two-stage training pipeline.** In the **Stage 1**, we use the mirror mask and view image with mirror content replaced by red to learn the mirror plane and a coarse 3D Gaussian representation of the scene. In the **Stage 2**, based on the estimated mirror plane, we fuse views from the original and mirrored viewpoints to form the final rendered result, further optimizing the 3D Gaussians.

Additionally, to improve the quality of rendering and geometric structure, we add depth supervision. First, a pretrained depth estimation model [16] customized for mirror scenes is utilized to extract the depth map \mathbf{D}_{gt} from image \mathbf{C}_{wo}^m . Then a modified rasterization library [17] is used to render the predicted depth map \mathbf{D}_{pred} , supervised with a similar loss:

$$\mathcal{L}_{depth} = \mathcal{L}_1(\mathbf{D}_{pred}, \mathbf{D}_{gt}). \quad (11)$$

The first stage focuses on optimizing the mirror equation $\pi = (\mathbf{n}_\pi^\top, d)$ to accurately generate virtual mirrored viewpoints. We filter the Gaussians with high mirror probability and use an improved RANSAC algorithm [18] to remove noise Gaussians. Then we supervise that the remaining Gaussians fall on the mirror plane:

$$\mathcal{L}_{plane} = \frac{1}{N_f} \sum_{i=1}^{N_f} |\mathbf{n}_\pi^\top \mathbf{p}_i + d|, \quad (12)$$

where \mathbf{p}_i is the coordinate of the Gaussian to be optimized, N_f is the number of filtered mirror Gaussians. The overall loss for the first stage is:

$$\mathcal{L}_{stage}^1 = \lambda_{mask}\mathcal{L}_{mask} + \mathcal{L}_{rgb}^1 + \lambda_{depth}\mathcal{L}_{depth} + \mathcal{L}_{plane}, \quad (13)$$

where $\lambda_{mask}, \lambda_{depth}$ are the weight of mask loss and depth loss, respectively.

Stage 2: The objective of the second stage is to refine the rendering quality of the fused images, during which the mirror equation is fixed. As described in [Sec. II-C](#), we fuse images from the current viewpoint and the virtual mirrored viewpoint, supervising with the complete, unmasked GT image \mathbf{C}_{gt} :

$$\mathcal{L}_{rgb}^2 = (1-\gamma)\mathcal{L}_1(\mathbf{C}_{fuse}, \mathbf{C}_{gt}) + \gamma\mathcal{L}_{D-SSIM}(\mathbf{C}_{fuse}, \mathbf{C}_{gt}) \quad (14)$$

where \mathbf{C}_{fuse} is the fused rendered image (see [Eq. \(9\)](#)). The overall loss of the second stage is expressed as:

$$\mathcal{L}_{stage}^2 = \lambda_{mask}\mathcal{L}_{mask} + \mathcal{L}_{rgb}^2. \quad (15)$$



Fig. 4. **Qualitative comparison of novel view synthesis.** Two rows display washroom and discussion room scenes. The top-right shows magnified mirror regions, and the bottom-right shows predicted mirror masks. PSNR is calculated only for the mirror regions. Both Mirror-NeRF and Mirror-3DGS accurately predict mirror masks, with Mirror-3DGS achieving the best performance.

TABLE I
QUANTITATIVE COMPARISON ON THE MIRROR-NeRF DATASET. BEST RESULTS ARE **BOLDED** AND SECOND BEST ARE UNDERLINED. METHODS MARKED WITH ‘*’ ARE EVALUATED ONLY ON MIRROR REGIONS.

Method	Synthetic				Real			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FPS \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FPS \uparrow
InstantNGP [19]	23.54	0.71	0.42	34.89	10.51	0.20	0.71	34.36
DVGO [20]	28.05	0.82	0.29	4.35	22.18	0.67	0.33	4.46
Mirror-NeRF [7]	38.07	0.99	0.01	0.75	<u>25.04</u>	0.86	0.06	0.07
3DGS [9]	37.00	<u>0.97</u>	0.04	341.13	22.55	0.73	0.24	209.62
Mirror-3DGS (Ours)	<u>37.89</u>	<u>0.97</u>	0.01	<u>171.98</u>	25.93	<u>0.81</u>	0.06	<u>95.97</u>
3DGS* [9]	28.35	-	-	-	21.99	-	-	-
Mirror-NeRF* [7]	29.47	-	-	-	22.39	-	-	-
Mirror-3DGS* (Ours)	32.40	-	-	-	24.53	-	-	-

III. EXPERIMENTS

A. Experimental Setup

Dataset. We utilize the publicly released dataset by Mirror-NeRF [7], which includes common scenarios with prevalent mirrors. The dataset comprises three synthetic scenes (living room, office, and washroom) and three real-world scenes (discussion room, lounge, and market). Each scene has 200 to 300 training views with accompanying mirror masks.

Implementation. Training and testing are done at 480×360 resolution for real scenes and 400×400 for synthetic scenes. The training process is divided into two stages: we train 5000 steps in stage 1, followed by 65000 steps in stage 2. The loss weights λ_{mask} and λ_{depth} are set to 1 and 0.1, respectively.

Evaluation. The rendering quality of all methods is evaluated and compared using the widely recognized PSNR, SSIM [21], and LPIPS [22] metrics. Additionally, we compare the rendering speed using FPS (Frames Per Second). All experiments are conducted on a single RTX 3090 GPU for fair comparison.

B. Experimental Results

Quantitative Comparisons. We conduct quantitative comparisons between our method and several competitive counterparts, and the comparison results are detailed in Tab. I.

Among these competitors, Mirror-NeRF achieves high PSNR, SSIM, and LPIPS scores in synthetic and real scenes. However, its rendering speed is significantly lower due to

NeRF’s volumetric rendering inefficiencies. Even the fastest NeRF-based method, InstantNGP, is much slower than 3DGS and our approach in real scenes. Therefore, NeRF-based methods are challenging for real-time rendering scenarios.

Our method yields performance better than vanilla 3DGS. Specifically, our Mirror-3DGS demonstrates an average PSNR improvement of 0.89 and 3.38 dB over vanilla 3DGS on the synthetic and real scenes, respectively. Additionally, our SSIM and LPIPS metrics far exceed those of vanilla 3DGS, indicating that our synthesized views are significantly more realistic. These improvements are mainly attributed to our method’s specialized handling of mirrors and their reflections in the scene, which inevitably introduces some additional computation overhead. Therefore, while our rendering speed is slightly lower than that of vanilla 3DGS, it is significantly faster than NeRF-based methods, making it suitable for real-time rendering applications.

It is important to highlight that our Mirror-3DGS outperforms both Mirror-NeRF and vanilla 3DGS in terms of the PSNR metric when focusing solely on the mirror regions, as shown in the last three rows of Tab. I, demonstrating our method’s enhanced capability in handling mirror reflections. **Qualitative Comparisons.** The qualitative analysis presented in Fig. 4, focusing on the detailed rendering of mirrors. Our Mirror-3DGS, by incorporating the principles of plane mirror imaging and accurately modeling spatial geometric relationships, enables high-fidelity rendering of mirror contents. Across all six scenes, our Mirror-3DGS achieves rendering quality comparable to Mirror-NeRF, demonstrating its effectiveness in handling complex mirror reflection.

IV. CONCLUSION

In this paper, we introduce Mirror-3DGS, a novel rendering framework that addresses the challenge of accurately capturing and rendering physical reflections in mirror-containing scenes. Our Mirror-3DGS employs a two-stage training process that first filters mirror Gaussians and then uses this information to accurately estimate the mirror plane and derive mirrored camera parameters. The effectiveness of Mirror-3DGS in delivering high-quality renderings of scenes with mirrors is demonstrated by experiments on synthetic and real scenes.

REFERENCES

- [1] P. P. Srinivasan, B. Deng, X. Zhang, M. Tancik, B. Mildenhall, and J. T. Barron, “NeRV: Neural Reflectance and Visibility Fields for Relighting and View Synthesis,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [2] D. Verbin, P. Hedman, B. Mildenhall, T. Zickler, J. T. Barron, and P. P. Srinivasan, “Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [3] Y. Zhang, J. Sun, X. He, H. Fu, R. Jia, and X. Zhou, “Modeling Indirect Illumination for Inverse Rendering,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [4] Y. Shi, Y. Wu, C. Wu, X. Liu, C. Zhao, H. Feng, J. Liu, L. Zhang, J. Zhang, B. Zhou *et al.*, “GIR: 3D Gaussian Inverse Rendering for Relightable Scene Factorization,” *arXiv preprint arXiv:2312.05133*, 2023.
- [5] L. Ma, V. Agrawal, H. Turki, C. Kim, C. Gao, P. Sander, M. Zollhöfer, and C. Richardt, “SpecNeRF: Gaussian Directional Encoding for Spectular Reflections,” *arXiv preprint arXiv:2312.13102*, 2023.
- [6] Y.-C. Guo, D. Kang, L. Bao, Y. He, and S.-H. Zhang, “NeRFReN: Neural Radiance Fields With Reflections,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [7] J. Zeng, C. Bao, R. Chen, Z. Dong, G. Zhang, H. Bao, and Z. Cui, “Mirror-NeRF: Learning Neural Radiance Fields for Mirrors with Whitted-Style Ray Tracing,” in *Proceedings of the ACM International Conference on Multimedia (ACM MM)*, 2023.
- [8] T. Whitted, “An improved illumination model for shaded display,” in *ACM Siggraph 2005 Courses*, 2005.
- [9] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3D Gaussian Splatting for Real-Time Radiance Field Rendering,” *ACM Transactions on Graphics (TOG)*, 2023.
- [10] S. Yang, Y. Wang, H. Li, J. Meng, X. Meng, and J. Zhang, “Fourier123: One Image to High-Quality 3D Object Generation with Hybrid Fourier Score Distillation,” *arXiv preprint arXiv:2405.20669*, 2024.
- [11] W. Yifan, F. Serena, S. Wu, C. Öztireli, and O. Sorkine-Hornung, “Differentiable Surface Splatting for Point-based Geometry Processing,” *ACM Transactions on Graphics (TOG)*, 2019.
- [12] M. Zwicker, H. Pfister, J. Van Baar, and M. Gross, “EWA Volume Splatting,” in *IEEE Conference on Visualization (VIS)*, 2001.
- [13] G. Kopanas, T. Leimkühler, G. Rainer, C. Jambon, and G. Drettakis, “Neural Point Catacaustics for Novel-View Synthesis of Reflections,” *ACM Transactions on Graphics (TOG)*, 2022.
- [14] G. Kopanas, J. Philip, T. Leimkühler, and G. Drettakis, “Point-Based Neural Rendering with Per-View Optimization,” *Proceedings of the Eurographics Symposium on Rendering (EGSR)*, 2021.
- [15] M. Kaess, “Simultaneous Localization and Mapping with Infinite Planes,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [16] A. Costanzino, P. Z. Ramirez, M. Poggi, F. Tosi, S. Mattoccia, and L. Di Stefano, “Learning Depth Estimation for Transparent and Mirror Surfaces,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [17] J. Luiten, G. Kopanas, B. Leibe, and D. Ramanan, “Dynamic 3D Gaussians: Tracking by Persistent Dynamic View Synthesis,” in *Proceedings of the International Conference on 3D Vision (3DV)*, 2024.
- [18] M. A. Fischler and R. C. Bolles, “Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography,” *Communications of the ACM*, 1981.
- [19] T. Müller, A. Evans, C. Schied, and A. Keller, “Instant Neural Graphics Primitives with a Multiresolution Hash Encoding,” *ACM Transactions on Graphics (TOG)*, 2022.
- [20] C. Sun, M. Sun, and H.-T. Chen, “Direct Voxel Grid Optimization: Super-Fast Convergence for Radiance Fields Reconstruction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [21] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image Quality Assessment: From Error Visibility to Structural Similarity,” *IEEE Transactions on Image Processing (TIP)*, 2004.
- [22] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The Unreasonable Effectiveness of Deep Features as a Perceptual Metric,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.