On the Error Analysis of 3D Gaussian Splatting and an Optimal Projection Strategy

Letian Huang[®], Jiayang Bai[®], Jie Guo^{*}[®], Yuanqi Li[®], and Yanwen Guo[®]

Nanjing University, Nanjing, China

{lthuang, jybai, yuanqili}@smail.nju.edu.cn, {guojie, ywguo}@nju.edu.cn https://letianhuang.github.io/op43dgs/

Abstract. 3D Gaussian Splatting has garnered extensive attention and application in real-time neural rendering. Concurrently, concerns have been raised about the limitations of this technology in aspects such as point cloud storage, performance, and robustness in sparse viewpoints, leading to various improvements. However, there has been a notable lack of attention to the fundamental problem of projection errors introduced by the local affine approximation inherent in the splatting itself, and the consequential impact of these errors on the quality of photo-realistic rendering. This paper addresses the projection error function of 3D Gaussian Splatting, commencing with the residual error from the first-order Taylor expansion of the projection function. The analysis establishes a correlation between the error and the Gaussian mean position. Subsequently, leveraging function optimization theory, this paper analyzes the function's minima to provide an optimal projection strategy for Gaussian Splatting referred to Optimal Gaussian Splatting, which can accommodate a variety of camera models. Experimental validation further confirms that this projection methodology reduces artifacts, resulting in a more convincingly realistic rendering.

Keywords: 3D deep learning \cdot view synthesis \cdot radiance fields \cdot 3D gaussians \cdot real-time rendering \cdot error analysis

1 Introduction

Reconstructing 3D scenes from 2D images and synthesizing novel views has been a critical task in computer vision and graphics. Recently, the trend of novel view synthesis is spearheaded by Neural Radiance Fields (NeRFs) [25] and its variants [1, 2], which achieves photo-realistic rendering quality. NeRF represents scenes with Multi-Layer Perceptron (MLP) and renders novel views using volumetric ray-marching. However, these MLP-based methods contain a relatively large number of layers, leading to lengthy training and rendering times, which still fall significantly short of real-time rendering. To accelerate rendering,

^{*} Corresponding author.



Fig. 1: By minimizing the projection error through error analysis, we have achieved an improvement in the rendering image quality compared to the original 3D-GS [16]. These images are cropped from the complete images of the TRAIN scene, in order to be highlighted.

many methods resort to auxiliary data structures, such as hash grids [26], points or voxel [8, 36, 37, 39]. However, due to the dense sampling of points along rays, they still face challenges in achieving real-time performance.

With a substantial demand for high-speed and photo-realistic rendering effects, 3D Gaussian Splatting (3D-GS) [16] has emerged as an efficient representation that can be rendered at high speed on the GPU. It has departed from the implicit scene representation using MLPs and instead opted for an explicit representation using Gaussian functions. As a result, 3D-GS avoids the expensive cost of densely sampling points in volume rendering, thus achieving real-time performance. Given a set of images with the camera parameters calibrated by SfM [31], 3D-GS aims to optimize a set of 3D Gaussians as graphical primitives to explicitly represent the scene. To render 2D images, these 3D Gaussians are projected onto the image plane (z = 1 plane) via a projection function denoted as ϕ for differentiable rasterization.

Unfortunately, Gaussian functions are closed under affine transformations but not under projective transformations. Therefore, 3D-GS adopts a local affine approximation [44], specifically approximating the projection function with the first two terms of its Taylor expansion. Nevertheless, approximations introduce errors, and the local affine approximation similarly contributes to these errors which may lead to artifacts in the rendered images. This paper exploits the relationship between the error of 3D-GS and the Gaussian mean through the analysis of the Taylor remainder term. Additionally, it identifies the circumstances under which the error is minimized by determining the extremum of the error function.

Finally, based on the extremum analysis of the error function, we propose a novel optimal projection method to minimize the projection error in 3D-GS. This method requires only minor code modifications, does not affect real-time rendering performance, and yet achieves a significant improvement in rendering quality. Specifically, we project along the direction from the Gaussian mean to the camera center, where the projection plane is tangent to the line connecting the Gaussian mean and the camera center. Compared with the straightforward approach of projecting 3D Gaussians onto the image plane of the camera, the proposed projection reduces expected error while maintaining real-time and realistic rendering performance. Furthermore, we find that the projection errors in 3D-GS tend to escalate with the continuous expansion of the field of view and the corresponding reduction in focal length. This increasing error leads to more and larger elongated Gaussians, as well as cloud-like Gaussian artifacts, consequently degrading the overall image quality. In contrast, our method exhibits greater realism and robustness. As illustrated in Figure 1, our method significantly reduces blurriness and artifacts caused by the projection error. Through experiments on various datasets and scenes, our method consistently outperforms the original projection method, particularly in settings with short focal lengths. And the method is easily implemented and adaptable to various camera models. To summarize, we provide the following contributions:

- We provide thorough error analysis for 3D-GS that may lead to artifacts and consequently degrade rendering quality. We have identified the correlation between this error and the Gaussian position.
- We derive the mathematical expectation of this error function and analyze when this function takes extrema through methods of function optimization.
- We propose a novel optimal projection capable of adapting to diverse camera models that adopts different tangent plane projections based on the directions from each Gaussian mean to the camera center instead of naively projecting all Gaussians onto the same plane

2 Related Work

Novel view synthesis is a longstanding challenge in computer vision and graphics. From traditional techniques [3,5,10,11,19,31,33], to neural network-based scene representations [9, 15, 18, 24, 28, 29, 32], various approaches have struggled to address the problem of synthesizing a new view from captured images.

2.1 Neural Radiance Fields

The Neural Radiance Field (NeRF) [25] stands out as a successful neural rendering method based on MLPs, primarily owing to its encoding of position and direction. This encoding allows for effective reconstruction of high-frequency information in scenes. Notable improvements on this encoding have been made by Mip-NeRF [1], NeRF-W [22], FreeNeRF [41], and Instant NGP [26]. These enhancements enable the handling of multi-resolution image inputs, multi-illumination with occlusion image inputs, sparse-view inputs, and achieve nearly real-time rendering capabilities, respectively. Barron et al. introduced MipNeRF360 [2] as an extension of MipNeRF [1] to address the issue of generating low-quality renderings for unbounded scenes in NeRF. Importantly, these methods thoroughly exploit the intrinsic capabilities of NeRF as an implicit scene representation without introducing additional model priors.

However, due to the implicit representation of scenes and the dense sampling of points along rays, they still face challenges in achieving real-time performance. 4 L. Huang et al.

2.2 3D Gaussian Splatting

3D Gaussian Splatting (3D-GS) [16] has demonstrated notable advancements in rendering performance by departing from MLPs and ray sampling, opting instead for Gaussian functions and Gaussian Splatting. This paradigm shift has attracted considerable attention within the industry, leading to various studies building upon this methodology [6].

3D-GS finds application in SLAM systems [23], scene editing [7] and segmentation [4]. Efforts [20] have also been made to integrate 3D-GS with popular diffusion models [14, 34, 35] and dynamic scenes [21]. Recent endeavors aim to enhance the robustness in sparse-view scenarios [38, 40], performance, storage efficiency [27] and mesh reconstruction [12] of 3D-GS. However, these improvements do not specifically address errors associated with Gaussian projection. Potential methods for enhancing Gaussian functions include [30, 42]. Nevertheless, they introduce a Gaussian filter and a novel Gaussian depth sorting algorithm, respectively, without investigating the projection function and associated errors. These approaches still rely on projecting onto the z = 1 plane.

This paper aims to explore the potential of 3D Gaussian Splatting by analyzing errors that may arise during the projection process. The analysis will delve into the factors contributing to these errors and propose methods to mitigate and reduce them.

3 Preliminaries

Similar to NeRF [25], the input to 3D-GS consists of a set of images, together with the corresponding cameras calibrated by SFM [31]. However, in contrast to NeRF, 3D-GS takes the sparse point cloud generated during the SFM calibration process as input. From these points, it constructs a set of 3D Gaussians as graphical primitives explicitly representing the scene. Each 3D Gaussian is characterized by its position (mean) μ , covariance matrix Σ , and opacity α , and carries anisotropic spherical harmonics (SH) representing directional appearance component (color) of the radiance field, following standard practice [8,26]. Subsequently, these 3D Gaussians are projected onto the image plane (z = 1 plane) via a projection function for differentiable rasterization. Next, we will provide a detailed explanation of the Gaussian projection process in 3D-GS.

In the world space, we define a 3D Gaussian function by G, characterized by its mean μ and covariance matrix Σ . To project the Gaussian onto the image plane, the initial step in 3D-GS involves an affine transformation of this Gaussian function from the world space to the camera space via a viewing transformation matrix \mathbf{W} , yielding a transformed Gaussian function:

$$G(\mathbf{x}) = \exp\left\{-1/2\left(\mathbf{W}\mathbf{x} - \mathbf{W}\boldsymbol{\mu}\right)^{\top}\left(\mathbf{W}\boldsymbol{\Sigma}\mathbf{W}^{\top}\right)^{-1}\left(\mathbf{W}\mathbf{x} - \mathbf{W}\boldsymbol{\mu}\right)\right\}.$$
 (1)

We denote G' as the 3D Gaussian in the camera space, with $\mathbf{x}' = \mathbf{W}\mathbf{x}, \boldsymbol{\mu}' = \mathbf{W}\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}' = \mathbf{W}\boldsymbol{\Sigma}\mathbf{W}^{\top}$ as the point, mean and covariance matrix in the camera space, respectively. Subsequently, traditional 3D-GS requires projecting the Gaussian function in the camera space onto the z = 1 plane, with the projection function φ defined as:

$$\varphi\left(\mathbf{x}'\right) = \mathbf{x}'\left(\mathbf{x_0}^{\top}\mathbf{x}'\right)^{-1}\left(\mathbf{x_0}^{\top}\mathbf{x_0}\right) = \mathbf{x}'\left(\mathbf{x_0}^{\top}\mathbf{x}'\right)^{-1}$$
(2)

where $\mathbf{x}_{\mathbf{0}} = \begin{bmatrix} 0, 0, 1 \end{bmatrix}^{\top}$ represents the projection of the camera space's origin onto this plane. Expanding this projection function to the first order using Taylor series yields:

$$\varphi\left(\mathbf{x}^{'}\right) = \varphi\left(\boldsymbol{\mu}^{'}\right) + \frac{\partial\varphi}{\partial\mathbf{x}^{'}}\left(\boldsymbol{\mu}^{'}\right)\left(\mathbf{x}^{'}-\boldsymbol{\mu}^{'}\right) + R_{1}\left(\mathbf{x}^{'}\right)$$
$$\approx \varphi\left(\boldsymbol{\mu}^{'}\right) + \frac{\partial\varphi}{\partial\mathbf{x}^{'}}\left(\boldsymbol{\mu}^{'}\right)\left(\mathbf{x}^{'}-\boldsymbol{\mu}^{'}\right)$$
(3)

where $\frac{\partial \varphi}{\partial \mathbf{x}'}(\boldsymbol{\mu}')$ is the Jacobian of the affine approximation of the projective transformation, denoted as **J**. Applying this local affine approximation by neglecting the Taylor remainder term allows to derive the 2D Gaussian function G_{2D} projected onto the z = 1 plane as:

$$G_{2D}\left(\mathbf{x}'\right) = \exp\left\{-1/2\left(\mathbf{J}\mathbf{x}' - \mathbf{J}\boldsymbol{\mu}'\right)^{\top} \left(\mathbf{J}\boldsymbol{\varSigma}'\mathbf{J}^{\top}\right)^{-1} \left(\mathbf{J}\mathbf{x}' - \mathbf{J}\boldsymbol{\mu}'\right)\right\}$$
$$\approx \exp\left\{-1/2\left(\varphi\left(\mathbf{x}'\right) - \varphi\left(\boldsymbol{\mu}'\right)\right)^{\top} \left(\mathbf{J}\boldsymbol{\varSigma}'\mathbf{J}^{\top}\right)^{-1} \left(\varphi\left(\mathbf{x}'\right) - \varphi\left(\boldsymbol{\mu}'\right)\right)\right\}.$$
(4)

Since the rank of the matrix \mathbf{J} is 2, the inverse of $\mathbf{J}\boldsymbol{\Sigma}'\mathbf{J}^{\top}$ is, in fact, the inverse of the covariance matrix of the 2D Gaussian obtained by skipping the third row and column [44]. We denote $\mathbf{x}_{2D} = \varphi\left(\mathbf{x}'\right)$, $\boldsymbol{\mu}_{2D} = \varphi\left(\boldsymbol{\mu}'\right)$ and $\boldsymbol{\Sigma}_{2D} = \mathbf{J}\boldsymbol{\Sigma}'\mathbf{J}^{\top}$ as the point, mean and 2 × 2 covariance matrix of G_{2D} in the image space, respectively. For a pixel (u, v), we obtain the value of the 2D Gaussian function by querying $\mathbf{x}_{2D} = [u, v]^{\top}$, and then perform alpha blending to derive the color of the pixel. The blue box in Figure 3 illustrates the projection process.

4 Local Affine Approximation Error

From Equation 3 and Equation 4 we observe that 3D Gaussian Splatting introduces an approximation during the projection transformation, i.e. the local affine approximation. Consequently, 3D-GS actually adopts a 2D Gaussian function that is not the true projection of the initial 3D Gaussian. The error introduced by this approximation can be characterized by the Taylor remainder term in Equation 3:

$$R_{1}\left(\mathbf{x}'\right) = \varphi\left(\mathbf{x}'\right) - \varphi\left(\boldsymbol{\mu}'\right) - \frac{\partial\varphi}{\partial\mathbf{x}'}\left(\boldsymbol{\mu}'\right)\left(\mathbf{x}' - \boldsymbol{\mu}'\right)$$
(5)

6 L. Huang et al.

where

$$\frac{\partial \varphi}{\partial \mathbf{x}'} \left(\boldsymbol{\mu}' \right) = \mathbb{I} \otimes \left(\mathbf{x_0}^\top \boldsymbol{\mu}' \right)^{-1} - \boldsymbol{\mu}' \left(\mathbf{x_0}^\top \boldsymbol{\mu}' \right)^{-1} \left(\boldsymbol{\mu}'^\top \mathbf{x_0} \right)^{-1} \mathbf{x_0}^\top.$$
(6)

I represents the identity matrix and \otimes denotes the multiplication of a matrix and a scalar. Clearly, this Taylor remainder term is a three-dimensional vector related to the random variable \mathbf{x}' and the mean $\boldsymbol{\mu}'$ of the Gaussian function. Therefore, by computing the square of the Frobenius norm of this vector $||R_1(\mathbf{x}')||_F^2$ and taking the mathematical expectation of this norm function with respect to the random variable \mathbf{x}' , we obtain an error function that depends solely on the Gaussian mean $\boldsymbol{\mu}'$:

$$\epsilon\left(\boldsymbol{\mu}'\right) = \int_{\mathbf{x}'\in\mathcal{X}'} \|R_1\left(\mathbf{x}'\right)\|_F^2 \mathrm{d}\mathbf{x}'.$$
(7)

The simplified error function ϵ involves three unit vectors: \mathbf{x}' , $\mathbf{x_0}$ and $\boldsymbol{\mu}'$. Their spherical coordinates are given as follows:

$$\mathbf{x_0} = \begin{bmatrix} \sin(\phi_0)\cos(\theta_0) \\ -\sin(\theta_0) \\ \cos(\phi_0)\cos(\theta_0) \end{bmatrix}, \mathbf{x}' = \begin{bmatrix} \sin(\phi)\cos(\theta) \\ -\sin(\theta) \\ \cos(\phi)\cos(\theta) \end{bmatrix}, \boldsymbol{\mu}' = \begin{bmatrix} \sin(\phi_{\mu})\cos(\theta_{\mu}) \\ -\sin(\theta_{\mu}) \\ \cos(\phi_{\mu})\cos(\theta_{\mu}) \end{bmatrix}$$
(8)

where $\phi_0 = 0$ and $\theta_0 = 0$. Please refer to the detailed proofs in the supplementary materials.

The size of the integration region is related to the covariance of the Gaussian function. Since the current discussion specifically focuses on the impact of the mean on the error rather than the variance, let's assume the integration region to be $\left\{ \mathbf{x}' \mid \theta \in \left[-\pi/4 + \theta_{\mu}, \pi/4 + \theta_{\mu}\right] \land \phi \in \left[-\pi/4 + \phi_{\mu}, \pi/4 + \phi_{\mu}\right] \right\}$. Then, based on Equation 5-8, we obtain:

$$\epsilon\left(\theta_{\mu},\phi_{\mu}\right) = \int_{-\pi/4+\theta_{\mu}}^{\pi/4+\theta_{\mu}} \int_{-\pi/4+\phi_{\mu}}^{\pi/4+\phi_{\mu}} \left(-\frac{\sin\left(\phi-\phi_{\mu}\right)\cos\left(\theta\right)}{\cos^{2}\left(\phi_{\mu}\right)\cos\left(\theta_{\mu}\right)} + \tan\left(\phi\right) - \tan\left(\phi_{\mu}\right)\right)^{2} + \left(\frac{\sin\left(\theta\right)}{\cos\left(\phi_{\mu}\right)\cos\left(\theta_{\mu}\right)} - \frac{\sin\left(\theta_{\mu}\right)\cos\left(\phi\right)\cos\left(\theta\right)}{\cos^{2}\left(\phi_{\mu}\right)\cos^{2}\left(\theta_{\mu}\right)} + \frac{\tan\left(\theta_{\mu}\right)}{\cos\left(\phi_{\mu}\right)} - \frac{\tan\left(\theta\right)}{\cos\left(\phi\right)}\right)^{2} d\theta d\phi$$
(9)

Fortunately, this integral has a close-form expression, but its form is very complicated. We report the close-form expression in the supplemental materials and visually depict it in Figure 2.

The error function is partially differentiated with respect to θ_{μ} and ϕ_{μ} . It is easy to check that at $\theta_{\mu} = \theta_0 = 0$ and $\phi_{\mu} = \phi_0 = 0$, the following holds:

$$\frac{\partial \epsilon}{\partial \theta_{\mu}} (0,0) = 0, \frac{\partial \epsilon}{\partial \phi_{\mu}} (0,0) = 0.$$
(10)



Fig. 2: Visualization of the 3D Gaussian Splatting error function $\epsilon(\theta_{\mu}, \phi_{\mu})$ under two distinct domains. The domain of the function is $\{(\theta_{\mu}, \phi_{\mu}) \mid \theta_{\mu} \in [-\lambda \pi/4, \lambda \pi/4] \land \phi_{\mu} \in [-\lambda \pi/4, \lambda \pi/4]\}$ where $(\theta_{\mu}, \phi_{\mu})$ represents the polar coordinate of the 3D Gaussian mean and λ represents the scaling factor of the domain.

Therefore, this point is actually an extremum point of the function. We also confirm that this point is the minimum of the function. Indeed, from Figure 2a, it is also apparent that the function is concave, with its minimum value occurring at (0,0). Additionally, it can be observed that the minimum value of the error function is greater than zero, due to the fact that Gaussian functions are not closed under projection transformations. Figure 2b shows that in the majority of the region close to the origin, the function values vary very gently. However, as the independent variable of the function approaches the boundaries, the function values increase rapidly, resulting in a substantial difference between the maximum and minimum values.

The above analysis indicates that the error is small and does not significantly affect the quality of rendered images. Therefore, traditional 3D-GS [16] utilizes local affine approximation but still successfully reconstructs the scene to obtain a high-quality novel view image. Nevertheless, the naive projection of all Gaussians onto the same plane z = 1 in 3D-GS may lead to larger projection errors for Gaussians farther from the plane center (the projection point of the camera center to the plane), causing artifacts. These artifacts are severe when using wide-angle lenses, significantly degrading the quality of the renderings.

5 Optimal Gaussian Splatting

Based on the analysis of the error function, we introduce Optimal Gaussian Splatting shown in Figure 3. This helps us to achieve smaller projection errors, resulting in higher-quality rendering. The first step, similar to the original 3D-GS [16], involves transforming world coordinates into the camera coordinate system. Instead of projecting Gaussians onto the z = 1 plane, each Gaussian is radially projected based on its mean along the line connecting it to the camera

7



Fig. 3: Illustration of the rendering pipeline for our Optimal Gaussian Splatting and the projection of 3D-GS [16]. The blue box depicts the projection process of the original 3D-GS, which straightforwardly projects all Gaussians onto the same projection plane. In contrast, the red box illustrates our approach, where we project individual Gaussians onto corresponding tangent planes.

center, projecting onto a plane tangent to the unit sphere and perpendicular to this line. After projection, colors of points on the unit sphere are computed by alpha blending the 2D Gaussians on the tangent plane. Finally, for each pixel in the image, rays are cast onto the unit sphere to retrieve the color of that pixel, resulting in the rendered image.

5.1 Optimal Projection

We discover that the error function attains its minimum value when the projection of the Gaussian mean on the plane coincides with the projection from the camera center to the plane. Therefore, in Optimal Gaussian Splatting, we employ an optimal projection method. Specifically, instead of naively projecting different Gaussians onto the same plane z = 1, we adopt distinct projection planes for each Gaussian. These projection planes are determined by the tangent planes formed by the Gaussian mean and the line connecting it to the camera center. The specific projection plane is formulated as:

$$\mathbf{x}_{\mathbf{p}}^{\top} \cdot \left(\mathbf{x}^{'} - \mathbf{x}_{\mathbf{p}} \right) = 0 \tag{11}$$

where

$$\mathbf{x}_{\mathbf{p}} = \varpi \left(\boldsymbol{\mu}' \right) = \boldsymbol{\mu}' \left(\boldsymbol{\mu}'^{\top} \boldsymbol{\mu}' \right)^{-1/2}$$
(12)

represents the projection of the camera space's origin onto this plane and ϖ function represents the transformation that projects points onto the unit sphere. According to Equation 11, the optimal projection function $\varphi_{\rm p}$ is obtained as:

$$\varphi_{\mathbf{p}}\left(\mathbf{x}'\right) = \mathbf{x}'\left(\mathbf{x}_{\mathbf{p}}^{\top}\mathbf{x}'\right)^{-1}\left(\mathbf{x}_{\mathbf{p}}^{\top}\mathbf{x}_{\mathbf{p}}\right) = \mathbf{x}'\left(\mathbf{x}_{\mathbf{p}}^{\top}\mathbf{x}'\right)^{-1}.$$
 (13)

The corresponding local affine approximation Jacobian matrix $\mathbf{J}_{\mathbf{p}}$ is:

$$\mathbf{J}_{\mathbf{p}} = \frac{\partial \varphi_{\mathbf{p}}}{\partial \mathbf{x}'} \left(\boldsymbol{\mu}' \right) = \mathbb{I} \otimes \left(\mathbf{x}_{\mathbf{p}}^{\top} \boldsymbol{\mu}' \right)^{-1} - \boldsymbol{\mu}' \left(\mathbf{x}_{\mathbf{p}}^{\top} \boldsymbol{\mu}' \right)^{-1} \left(\boldsymbol{\mu}'^{\top} \mathbf{x}_{\mathbf{p}} \right)^{-1} \mathbf{x}_{\mathbf{p}}^{\top}.$$
 (14)

When $\boldsymbol{\mu}' = [\mu_x, \mu_y, \mu_z]^{\top}$, the specific form of this Jacobian matrix is:

$$\mathbf{J}_{\mathbf{p}} = \begin{bmatrix} \frac{\mu_y^2 + \mu_z^2}{(\mu_x^2 + \mu_y^2 + \mu_z^2)^{\frac{3}{2}}} & -\frac{\mu_x \mu_y}{(\mu_x^2 + \mu_y^2 + \mu_z^2)^{\frac{3}{2}}} & -\frac{\mu_x \mu_z}{(\mu_x^2 + \mu_y^2 + \mu_z^2)^{\frac{3}{2}}} \\ -\frac{\mu_x \mu_y}{(\mu_x^2 + \mu_y^2 + \mu_z^2)^{\frac{3}{2}}} & \frac{\mu_x^2 + \mu_z^2}{(\mu_x^2 + \mu_y^2 + \mu_z^2)^{\frac{3}{2}}} & -\frac{\mu_y \mu_z}{(\mu_x^2 + \mu_y^2 + \mu_z^2)^{\frac{3}{2}}} \\ -\frac{\mu_x \mu_z}{(\mu_x^2 + \mu_y^2 + \mu_z^2)^{\frac{3}{2}}} & -\frac{\mu_y \mu_z}{(\mu_x^2 + \mu_y^2 + \mu_z^2)^{\frac{3}{2}}} & \frac{\mu_x^2 + \mu_y^2}{(\mu_x^2 + \mu_y^2 + \mu_z^2)^{\frac{3}{2}}} \end{bmatrix}.$$
(15)

It's very easy to implement modifications to the Jacobian matrix in the forward process.

5.2 Unit Sphere Based Rasterizer

For the proposed Optimal Projection, we have proposed a Unit Sphere Based Rasterizer to generate images. Through Optimal Projection, we obtain the projection of the three-dimensional Gaussians on the tangent plane of the unit sphere instead of obtaining the Gaussians in the image space. Therefore, we need to rasterize based on this unit sphere.

Specifically, for a pixel (u, v) on the image, similar to NeRF, we cast a ray. However, unlike NeRF, we do not involve the extensive sampling of points along the ray, which would significantly degrade performance. Instead, our focus is solely on determining which tangent-plane Gaussians the ray intersects on the unit sphere:

$$\mathbf{x}_{2\mathrm{D}} = \varphi_{\mathrm{p}} \left(\begin{bmatrix} \left(u - c_x \right) / f_x \\ \left(v - c_y \right) / f_y \\ 1 \end{bmatrix} \right)$$
(16)

where c_x , c_y , f_x , f_y denote the intrinsic parameters of the pinhole camera. We then query the function values of these Gaussians for alpha blending to obtain the color

$$G_{2D}\left(\mathbf{x}_{2D}\right) = \exp\left\{-1/2\left(\mathbf{x}_{2D} - \varphi_{p}\left(\boldsymbol{\mu}'\right)\right)^{\top} \left(\mathbf{J}_{p}\boldsymbol{\Sigma}'\mathbf{J}_{p}^{\top}\right)^{-1} \left(\mathbf{x}_{2D} - \varphi_{p}\left(\boldsymbol{\mu}'\right)\right)\right\}.$$
(17)

Note that both $\mathbf{J}_{\mathbf{p}}$ in Equation 14 and \mathbf{J} in Equation 6 are matrices with a rank of 2. Consequently, the inverse of the covariance matrix of the 2D Gaussian is also obtained by skipping the third row and column. However, the third row and third column of $\mathbf{J}_{\mathbf{p}}\boldsymbol{\Sigma}'\mathbf{J}_{\mathbf{p}}^{\top}$ are not entirely zeros, as opposed to $\mathbf{J}\boldsymbol{\Sigma}'\mathbf{J}^{\top}$. To ensure the equation still holds, an invertible matrix \mathbf{Q} needs to be used for left multiplication with \mathbf{x}_{2D} , $\varphi_{p}\left(\boldsymbol{\mu}'\right)$ and $\mathbf{J}_{\mathbf{p}}$:

$$\mathbf{Q} = \begin{bmatrix} \frac{\mu_z}{\sqrt{\mu_x^2 + \mu_z^2}} & 0 & -\frac{\mu_x}{\sqrt{\mu_x^2 + \mu_z^2}} \\ -\frac{\mu_x \mu_y}{\sqrt{\mu_x^2 + \mu_z^2} \sqrt{\mu_x^2 + \mu_z^2}} & \frac{\sqrt{\mu_x^2 + \mu_z^2}}{\sqrt{\mu_x^2 + \mu_y^2 + \mu_z^2}} & -\frac{\mu_y \mu_z}{\sqrt{\mu_x^2 + \mu_z^2} \sqrt{\mu_x^2 + \mu_y^2 + \mu_z^2}} \\ \frac{\mu_x}{\sqrt{\mu_x^2 + \mu_y^2 + \mu_z^2}} & \frac{\mu_y}{\sqrt{\mu_x^2 + \mu_y^2 + \mu_z^2}} & \frac{\mu_z}{\sqrt{\mu_x^2 + \mu_y^2 + \mu_z^2}} \end{bmatrix}.$$
(18)

10 L. Huang et al.

Since our projection does not rely on the perspective image plane z = 1 as in the original 3D-GS [16], we can adapt to different camera models such as fisheye cameras and panoramic images by simply modifying the transformation from image space to camera space, as described in Equation 16.

6 Experiments

To validate the effectiveness of the Optimal Gaussian Splatting derived from theoretical error analysis, we conduct a series of experiments, comparing it with the original 3D-GS [16] and some current state-of-the-art methods.

6.1 Implementation

We implemented Optimal Gaussian Splatting based on the PyTorch framework in 3D-GS [16] and wrote custom CUDA kernels for rasterization. We used the default parameters of 3D-GS to maintain consistency with the original 3D-GS and prevent other factors from introducing interference into the results.

6.2 Datasets

We test our algorithm on a total of 13 real scenes which are the same as those used in the original 3D-GS [16]. In particular, we evaluate our approach on the complete set of scenes featured in Mip-NeRF360 [2], which currently represents the state-of-the-art in NeRF rendering quality. Additionally, we test our method on two scenes from the Tanks & Temples dataset [17] and two scenes provided by Hedman et al. [13]. The selected scenes exhibit diverse capture styles, encompassing both confined indoor environments and expansive, unbounded outdoor settings.

6.3 Results

We compare our algorithm with 3D-GS [16] and three state-of-the-art NeRFbased approaches: Mip-NeRF360 [2], InstantNGP [26] and Plenoxels [8]. All methods are configured using the same settings as outlined in the 3D-GS [16] to control variables, preventing the introduction of variables other than our Optimal Splatting method.

Quantitative comparisons We adopt a train/test split for datasets following the methodology proposed by Mip-NeRF360 [2]. Specifically, every 8th photo is reserved for testing, ensuring a consistent and meaningful basis for comparisons to generate error metrics. Standard metrics such as PSNR, LPIPS [43], and SSIM are employed for evaluation. We report quantiative results in Table 1. While the 3D-GS significantly outperforms MipNeRF360 in rendering performance, its rendering quality slightly falls behind MipNeRF360, with a lower average PSNR across the 13 scenes. This might be attributed to MipNeRF360 not employing

Table 1: Quantitative evaluation of our method compared to previous work across three datasets. The Mip-NeRF360 [2] dataset includes both indoor and outdoor scenarios. Tanks. and Deep. indicates the average results over the Tanks&Templates [17] and Deep Blending [13] datasets. Average represents the mean value across all scenarios in these datasets.

Dataset	Average			Mip-NeRF360			Tanks. and Deep.		
Method Metric	$PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$
Plenoxels [8]	22.77	0.666	0.457	23.08	0.625	0.463	22.07	0.757	0.444
INGP-Base [26]	24.49	0.699	0.373	25.30	0.671	0.371	22.67	0.760	0.377
INGP-Big [26]	24.93	0.724	0.336	25.59	0.699	0.331	23.44	0.781	0.347
M-NeRF360 [2]	27.11	0.803	0.241	27.69	0.792	0.237	25.81	0.830	0.251
3D-GS [16]	26.92	0.832	0.214	27.21	0.815	0.214	26.27	0.872	0.213
Ours	27.17	0.836	0.210	27.48	0.821	0.209	26.44	0.872	0.214



Fig. 4: We show comparisons of our method to previous methods and the corresponding ground truth images from held-out test views. The scenes are, from the top down: TRUCK from Tanks&Temples [17]; PLAYROOM from the Deep Blending dataset [13] and BONSAI, COUNTER from Mip-NeRF360 dataset [2]. Differences in quality highlighted by arrows/insets.

significant approximations throughout its entire pipeline. However, by replacing the original projection method of 3D-GS with our optimal projection, which incurs smaller errors, we not only maintain rendering performance far surpassing MipNeRF360 but also achieve a quality improvement.

Qualitative comparisons For our qualitative experiments, we select only two methods, 3D-GS and MipNeRF360, as our baselines. It can be observed that



Fig. 5: We derive the projection error function with respect to image coordinates (u, v) and focal length through the transformation between image and polar coordinates $(\theta_{\mu}, \phi_{\mu})$. We visualize this function, showing 3D-GS rendered images with long and short focal lengths, followed by the corresponding error functions.

our method is capable of generating more realistic details, with fewer artifacts compared to 3D-GS [16]. And this is precisely the superiority brought about by our projection method, which results in smaller errors. Furthermore, We see that in some cases Mip-NeRF360 has remaining artifacts that our method avoids (e.g., black lines - in PLAYROOM).

6.4 Discussions

Impacts of Decreasing Focal Length Based on Equation 9, we further derive the projection error function with respect to image coordinates (u, v) and focal length f_x , f_y through the coordinate transformation between image coordinates and polar coordinates (θ_μ, ϕ_μ) . As illustrated in Figure 5, the peak of projection error mainly occurs at the image edges under a long focal length, resulting in artifacts at the image edges. As the focal length decreases, the field of view expands, leading to more Gaussians deviating from the projection center and an overall increase in projection error. In such cases, 3D-GS exhibit more artifacts, such as needle-like structures or cloud-like large Gaussians, which obscure parts that perform well under long focal length, thereby significantly degrading the overall image quality. Since our method employs a central radial projection for the projection plane, theoretically, we should not encounter such defects. To validate this perspective, we conduct experiments at various focal length reduction ratios (reduction ratios of $\times 0.2$ and $\times 0.3$).

For quantitative experiments, considering that there is no ground truth for large field-of-view in the 13 scenes from the 3D-GS [16], we apply a focal length mask. This mask selects central parts of the rendered images with corresponding ground truth based on the focal length scaling factor for metric calculation. The results are presented in the Table 2. Whether it is at $\times 0.2$ or $\times 0.3$ reduction ratios, our method outperforms the original 3D-GS significantly across various scenes in terms of all metrics. As the focal length decreases from $\times 0.3$ to $\times 0.2$, 3D-GS [16] experiences a significant degradation due to increased projection errors, while our method demonstrates greater robustness.

For quantitative experiments, we randomly select three scenes, each with two different focal lengths for demonstration. Additional experimental results can be found in the supplementary materials. It is evident from the Figure 6 **Table 2:** We quantitatively compare our method with the original 3D-GS [16] using focal lengths of 0.2 and 0.3 times the original. Six scenes were selected from three datasets: Mip-NeRF360 [2], Tank & Templates [17], and Deep Blending [13]. The scenes include TRAIN and TRUCK from Tanks & Templates, DRJOHNSON and PLAYROOM from Deep Blending, and ROOM and TREEHILL from Mip-NeRF360.

Metric	Focal	Method	Train	Truck	Dr Johnson	Playroom	Room	Treehill	Average
$\mathrm{PSNR}\uparrow$	$\times 0.2$	3D-GS [16]	10.12	12.39	19.24	16.38	20.02	15.36	15.58
		Ours	16.60	16.70	23.54	23.29	23.67	18.98	20.46
	$\times 0.3$	3D-GS [16]	12.85	14.89	23.36	22.51	25.00	18.35	19.49
		Ours	17.52	18.07	24.89	26.40	25.87	19.87	22.10
$\mathrm{SSIM}\uparrow$	$\times 0.2$	3D-GS [16]	0.325	0.380	0.616	0.619	0.680	0.371	0.499
		Ours	0.565	0.537	0.701	0.752	0.777	0.435	0.628
	$\times 0.3$	3D-GS [16]	0.471	0.529	0.740	0.791	0.812	0.437	0.630
		Ours	0.588	0.605	0.762	0.844	0.834	0.472	0.684
LPIPS↓	$\times 0.2$	3D-GS [16]	0.569	0.479	0.296	0.355	0.207	0.474	0.397
		Ours	0.277	0.250	0.216	0.203	0.178	0.385	0.251
	$\times 0.3$	3D-GS [16]	0.331	0.276	0.207	0.213	0.154	0.397	0.263
		Ours	0.243	0.217	0.201	0.182	0.158	0.385	0.231



Fig. 6: We compare our method to the original 3D-GS [16] with different focal length reductions. Two focal lengths, 0.2 and 0.3 times the original, are tested. Ground truth $\times 1$ in the left images indicates no focal length reduction. The red-boxed areas in the right images correspond to the same regions in the left images.

that our method exhibits greater robustness across various focal length settings, mitigating the artifacts introduced by projection in the original 3D-GS.

Adaptability to Various Camera Models By straightforwardly modifying the transformation from image space to camera space, we have achieved adaptation to various camera models, with the caveat that the original 3D-GS's projection struggles to support these camera models. Figure 7 illustrates that our projection, in addition to conforming to the original pinhole camera according to Equation 16, can also accommodate fisheye cameras and generate panoramas.



Fig. 7: Renderings of three scenes using our projection under three camera models. From left to right: the ground truth of the original perspective, wide-angle rendering with a pinhole camera, fisheye camera rendering, and the panorama.

Limitation Due to our method's independence from the image plane, it exhibits adaptability to various camera models. However, the slight increase in training time is a consequence of the subsequent transformation to the image plane after projection onto the tangent plane, rather than a direct projection onto the image plane. Fortunately, this method avoids the need for dense points sampling, ensuring the algorithm's time complexity remains unchanged and rendering performance stays unaffected. Moreover, it may be possible to mitigate the impact on training time, particularly concerning the corresponding backward propagation, through the implementation of optimized CUDA code.

7 Conclusion

Through theoretical analysis, we establish the correlation between the projection error and the Gaussian mean in 3D-GS. By optimizing the error function, we discover that the projection error is minimized when projecting on the plane perpendicular to the line connecting the Gaussian mean and the camera center. Leveraging this insight, we propose a novel projection that yields high-quality images without compromising rendering performance and can be easily adapted to various camera models through simple modifications.

In this work, our primary focus is on analyzing the relationship between the projection error and the mean of the 3D Gaussian under the assumption of a constant covariance. The influence of Gaussian covariance on projection requires further discussion in future work. It may be possible to further explore the potential of Gaussian Splatting as an explicit representation technique for scenes.

References

- Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5835–5844. IEEE (2021)
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mipnerf 360: Unbounded anti-aliased neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5470– 5479 (2022)
- Buehler, C., Bosse, M., McMillan, L., Gortler, S.J., Cohen, M.F.: Unstructured lumigraph rendering. In: Pocock, L. (ed.) Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2001, Los Angeles, California, USA, August 12-17, 2001. pp. 425–432. ACM (2001). https: //doi.org/10.1145/383259.383309, https://doi.org/10.1145/383259.383309
- Cen, J., Fang, J., Yang, C., Xie, L., Zhang, X., Shen, W., Tian, Q.: Segment any 3d gaussians. arXiv preprint arXiv:2312.00860 (2023)
- Chaurasia, G., Duchêne, S., Sorkine-Hornung, O., Drettakis, G.: Depth synthesis and local warps for plausible image-based navigation. ACM Trans. Graph. 32(3), 30:1-30:12 (2013). https://doi.org/10.1145/2487228.2487238, https://doi. org/10.1145/2487228.2487238
- Chen, G., Wang, W.: A survey on 3d gaussian splatting. arXiv preprint arXiv:2401.03890 (2024)
- Chen, Y., Chen, Z., Zhang, C., Wang, F., Yang, X., Wang, Y., Cai, Z., Yang, L., Liu, H., Lin, G.: Gaussianeditor: Swift and controllable 3d editing with gaussian splatting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21476–21485 (2024)
- Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5501–5510 (2022)
- 9. Genova, K., Cole, F., Sud, A., Sarna, A., Funkhouser, T.: Local deep implicit functions for 3d shape. In: CVPR (2020)
- Goesele, M., Snavely, N., Curless, B., Hoppe, H., Seitz, S.M.: Multi-view stereo for community photo collections. In: IEEE 11th International Conference on Computer Vision, ICCV 2007, Rio de Janeiro, Brazil, October 14-20, 2007. pp. 1–8. IEEE Computer Society (2007). https://doi.org/10.1109/ICCV.2007.4408933, https://doi.org/10.1109/ICCV.2007.4408933
- Gortler, S.J., Grzeszczuk, R., Szeliski, R., Cohen, M.F.: The lumigraph. In: Fujii, J. (ed.) Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1996, New Orleans, LA, USA, August 4-9, 1996. pp. 43–54. ACM (1996). https://doi.org/10.1145/237170.237200, https: //doi.org/10.1145/237170.237200
- Guédon, A., Lepetit, V.: Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5354– 5363 (2024)
- Hedman, P., Philip, J., Price, T., Frahm, J.M., Drettakis, G., Brostow, G.: Deep blending for free-viewpoint image-based rendering. ACM Transactions on Graphics (ToG) 37(6), 1–15 (2018)

- 16 L. Huang et al.
- 14. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual (2020), https://proceedings.neurips.cc/paper/2020/hash/ 4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html
- Jiang, C., Sud, A., Makadia, A., Huang, J., Nießner, M., Funkhouser, T.: Local implicit grid representations for 3d scenes. In: CVPR (2020)
- Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics 42(4) (2023)
- Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. ACM Transactions on Graphics (ToG) 36(4), 1–13 (2017)
- Kopanas, G., Leimkühler, T., Rainer, G., Jambon, C., Drettakis, G.: Neural point catacaustics for novel-view synthesis of reflections. ACM Trans. Graph. 41(6), 201:1-201:15 (2022). https://doi.org/10.1145/3550454.3555497, https://doi. org/10.1145/3550454.3555497
- Levoy, M., Hanrahan, P.: Light field rendering. In: Fujii, J. (ed.) Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1996, New Orleans, LA, USA, August 4-9, 1996. pp. 31–42. ACM (1996). https://doi.org/10.1145/237170.237199, https://doi.org/10.1145/ 237170.237199
- Li, X., Wang, H., Tseng, K.K.: Gaussiandiffusion: 3d gaussian splatting for denoising diffusion probabilistic models with structured noise. arXiv preprint arXiv:2311.11221 (2023)
- Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In: 2024 International Conference on 3D Vision (3DV). pp. 800–809. IEEE (2024)
- Martin-Brualla, R., Radwan, N., Sajjadi, M.S.M., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 7210–7219. Computer Vision Foundation / IEEE (2021). https://doi.org/10.1109/CVPR46437.2021.00713
- Matsuki, H., Murai, R., Kelly, P.H.J., Davison, A.J.: Gaussian Splatting SLAM. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2024)
- 24. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3D reconstruction in function space. In: CVPR (2019)
- Mildenhall, B., Srinivasan, P., Tancik, M., Barron, J., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European conference on computer vision (2020)
- Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics (ToG) 41(4), 1– 15 (2022)
- Niedermayr, S., Stumpfegger, J., Westermann, R.: Compressed 3d gaussian splatting for accelerated novel view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10349–10358 (2024)
- Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In: CVPR (2019)

- 29. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: CVPR (2019)
- Radl, L., Steiner, M., Parger, M., Weinrauch, A., Kerbl, B., Steinberger, M.: StopThePop: Sorted Gaussian Splatting for View-Consistent Real-time Rendering. ACM Transactions on Graphics 4(43) (2024)
- Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4104–4113 (2016)
- 32. Sitzmann, V., Zollhoefer, M., Wetzstein, G.: Scene representation networks: Continuous 3D-structure-aware neural scene representations. In: NeurIPS (2019)
- 33. Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring photo collections in 3d. ACM Trans. Graph. 25(3), 835–846 (2006). https://doi.org/10.1145/ 1141911.1141964, https://doi.org/10.1145/1141911.1141964
- 34. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada. pp. 11895-11907 (2019), https://proceedings.neurips.cc/paper/2019/hash/3001ef257407d5a371a96dcd947c7d93-Abstract.html
- 35. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Scorebased generative modeling through stochastic differential equations. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net (2021), https://openreview.net/forum? id=PxTIG12RRHS
- 36. Sun, C., Sun, M., Chen, H.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 5449–5459. IEEE (2022). https://doi.org/10.1109/CVPR52688.2022.00538, https://doi.org/10.1109/CVPR52688.2022.00538
- Takikawa, T., Litalien, J., Yin, K., Kreis, K., Loop, C.T., Nowrouzezahrai, D., Jacobson, A., McGuire, M., Fidler, S.: Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 11358–11367. Computer Vision Foundation / IEEE (2021). https://doi.org/10.1109/CVPR46437. 2021.01120
- Xiong, H., Muttukuru, S., Upadhyay, R., Chari, P., Kadambi, A.: Sparsegs: Realtime 360° sparse view synthesis using gaussian splatting. arXiv e-prints pp. arXiv-2312 (2023)
- 39. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 5428-5438. IEEE (2022). https://doi.org/10.1109/CVPR52688.2022.00536, https://doi.org/10.1109/CVPR52688.2022.00536
- 40. Yang, C., Li, S., Fang, J., Liang, R., Xie, L., Zhang, X., Shen, W., Tian, Q.: Gaussianobject: Just taking four images to get a high-quality 3d object with gaussian splatting. arXiv preprint arXiv:2402.10259 (2024)
- Yang, J., Pavone, M., Wang, Y.: Freenerf: Improving few-shot neural rendering with free frequency regularization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8254–8263 (2023)

- 18 L. Huang et al.
- 42. Yu, Z., Chen, A., Huang, B., Sattler, T., Geiger, A.: Mip-splatting: Alias-free 3d gaussian splatting. Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- 43. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)
- 44. Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: Ewa splatting. IEEE Transactions on Visualization and Computer Graphics 8(3), 223–238 (2002)