

Fisheye-GS: Lightweight and Extensible Gaussian Splatting Module for Fisheye Cameras

Zimu Liao¹, Siyan Chen¹, Rong Fu^{1,†}, Yi Wang¹, Zhongling Su^{1,‡}, Hao Luo¹,
Linning Xu², Bo Dai¹, Hengjie Li¹, Zhilin Pei¹, and Xingcheng Zhang¹

¹ Shanghai Artificial Intelligence Laboratory

² The Chinese University of Hong Kong

[†]furong@pjlab.org.cn [‡]suzhongling@pjlab.org.cn

Abstract. Recently, 3D Gaussian Splatting (3DGS) has garnered attention for its high fidelity and real-time rendering. However, adapting 3DGS to different camera models, particularly fisheye lenses, poses challenges due to the unique 3D to 2D projection calculation. Additionally, there are inefficiencies in the tile-based splatting, especially for the extreme curvature and wide field of view of fisheye lenses, which are crucial for its broader real-life applications. To tackle these challenges, we introduce Fisheye-GS*. This innovative method recalculates the projection transformation and its gradients for fisheye cameras. Our approach can be seamlessly integrated as a module into other efficient 3D rendering methods, emphasizing its extensibility, lightweight nature, and modular design. Since we only modified the projection component, it can also be easily adapted for use with different camera models. Compared to methods that train after undistortion, our approach demonstrates a clear improvement in visual quality.

Keywords: 3D Gaussian Splatting · Fisheye Camera Model

1 Introduction

The advent of high-resolution fisheye cameras has transformed how we capture and interact with the world, providing a 180° or greater field of view that offers unique perspectives. These perspectives are essential for various real-world applications. In immersive virtual reality (VR), fisheye cameras enable the creation of panoramic environments that enhance the user experience by providing a more comprehensive and realistic field of vision [29]. This is crucial for applications such as VR gaming [4], virtual tourism [39], and training simulations [11], where a wide and uninterrupted view can significantly improve immersion and usability. In surveillance systems, fisheye cameras allow for broader area coverage with fewer cameras, reducing costs and complexity. They are particularly valuable in security and monitoring scenarios, such as in large public spaces [14], retail stores [24, 32], and traffic monitoring [7], where capturing every angle with great detail is vital for effective oversight and incident management. However,

* Available at <https://github.com/zmliao/Fisheye-GS>



Fig. 1: Our Fisheye-GS. We have directly trained our 3DGS [13] model from images captured from fisheye cameras without undistortion to pinhole cameras. We then integrate our Fisheye-GS as a lightweight module within FlashGS [8], an efficient rendering technique for 3DGS, to evaluate its visual quality and performance.

integrating fisheye imagery into novel view synthesis poses several challenges due to the inherent curvature and distortion. While Neural Radiance Fields (NeRF) [19] handles the distortions of fisheye lenses through its ray sampling approach, which can easily accommodate various camera models, the state-of-the-art 3D Gaussian Splatting (3DGS) [13] faces difficulties. The original 3DGS approach relies on perspective projections. This is less adept at managing the extreme curvature and wide field of view of fisheye imagery, leading to a lack of accuracy and efficiency in rendering.

In this work, we aim to modularize the lightweight integration of fisheye cameras with 3D Gaussian Splatting (3DGS) to fully leverage the potential of fisheye lenses in real-world applications requiring real-time rendering performance. By combining the wide field of view of fisheye lenses with the efficient rendering capabilities of 3DGS, we seek to address the challenges posed by the curvature and distortion inherent in fisheye imagery.

Our proposed method, Fisheye-GS, is the first work that seamlessly integrates fisheye camera data into 3D Gaussian Splatting, unlocking new possibilities for applications in virtual reality, gaming, security, and more. We achieve this by:

- We first directly calculate the equidistant projection transformation and derive the means and covariance of Gaussians for fisheye lenses. It is crucial for maintaining image quality and precision across the extreme curvature and wide field of view.
- We then formulate the gradients required for optimization in the training phase and develop corresponding CUDA implementations, ensuring that the 3DGS model learns efficiently and produces distortion-free imagery.
- Finally, we only modified the projection model. This lightweight approach can be effectively modularized for integration into high-performance 3DGS methods and can be extended to accommodate other camera models.

Extensive experiments and comparisons with existing methods demonstrate the effectiveness of Fisheye-GS in producing high-quality, distortion-free imagery,

making it ideal for various 3D graphics applications. Up to our best knowledge, Fisheye-GS is the first open-source project that supports both training and rendering of 3DGS on fisheye cameras with efficient speed and performance.

2 Related Work

2.1 Novel View Synthesis

Neural Radiance Fields (NeRF). NeRF [19] is a groundbreaking technique in novel view synthesis, enabling the generation of realistic images of 3D scenes from any viewpoint. NeRF uses a deep neural network to encode a continuous volumetric scene function, which maps 3D points and viewing directions to corresponding RGB color and volume density values. To improve rendering quality, some methods refine the point sampling strategy during ray marching, leading to more accurate modeling of the volume rendering process [30]. Others enhance rendering by reparameterizing the scene representation, creating a more compact encoding, and simplifying the learning process [2, 3]. Despite these optimizations, NeRF remains computationally intensive during rendering. Typically, NeRF is represented as a global MLP encoding the entire scene space, which can be inefficient and costly for complex or large-scale scenes. To address this, some works have explored alternative scene representations, such as voxel grids and octree-based approaches [9, 36]. While these methods improve rendering speed, achieving real-time performance is still challenging due to the inherent ray marching strategy in volume rendering.

3D Gaussian Splatting 3D Gaussian Splatting (3DGS) [13] models the scene with a set of Gaussians, allowing efficient rendering through rasterizing these Gaussians into images. While 3DGS produces high-quality reconstruction results, there is room for improvement. [5, 6, 16, 28, 40] are designed to improve rendering performance under challenging inputs like sparse views. These methods enhance the robustness and accuracy of the rendered images in difficult scenarios. For anti-aliasing, approaches like [17, 31, 37] are utilized. Anchor-based methods, such as [18, 20] focus on improving efficiency and quality using hierarchical structures and anchor points. [8, 33] has proposed a suite of optimization strategies to amplify the performance of the rasterizing process in 3DGS. These methods enhance the computational efficiency and scalability of the rendering process.

2.2 Various Camera Models for Radiance Fields

Since NeRF cast rays from the camera center and samples points in 3D space for volume rendering, it is flexible in handling various camera models, as also integrated in [23], enables the reconstruction of a wide range of scenes from different types of input images. A few striking examples of models can be SC-NeRF [12] for reconstruct 3D scenes and camera parameters without traditional calibration objects, 360FusionNeRF [15] for 360-degree panorama rendering, OmniNeRF [22]

for enhancing 3D reconstruction at scene edges, PERF [25] for demonstrating the superiority of extensive 3D roaming, and Neuro Lens [27] for end-to-end optimization of the image formation process.

As for 3DGS, it primarily works in pinhole cameras. Adapting it to cameras with distortions is challenging because the projection of covariance is more difficult to compute. Some approaches have tried to extend 3DGS for specific cameras with distortion. For example, [1] projects 3D Gaussians onto the tangent plane of the unit sphere, addressing challenges in spherical projections and enhancing rendering quality by leveraging structural priors from panoramic layouts. [10] introduces an optimal projection strategy for various cameras by analyzing and minimizing projection errors by projecting Gaussians into a unit sphere.

3 Preliminaries

3.1 3D Gaussian Splatting

Definition. 3DGS represents a 3D model as a set of anisotropic 3D Gaussian functions and renders images via a differentiable tile-based rasterizer. Each Gaussian function is designated as the mean $\boldsymbol{\mu} \in \mathbb{R}^3$, anisotropic covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{3 \times 3}$,

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

where \mathbf{x} is an arbitrary position in the world coordinate system. To constrain the positive semi-definite form, the covariance is defined as $\boldsymbol{\Sigma} = \mathbf{R}\mathbf{S}\mathbf{S}^T\mathbf{R}^T$, where $\mathbf{R} \in SO(3)$ is the rotation matrix represented by quaternions and $\mathbf{S} = \text{diag}(s_1, s_2, s_3)$ is the scaling matrix.

Transformation. To project the Gaussian into the 2D plane, 3DGS applies the viewing transformation from the camera coordinate system $\varphi \in \mathbb{R}^3 \rightarrow \mathbb{R}^3$ and the projective transformation $\phi \in \mathbb{R}^3 \rightarrow \mathbb{R}^2$. We can model φ as the linear transformation $\varphi(\mathbf{x}) = \mathbf{W}\mathbf{x} + \mathbf{b}$ where $\mathbf{W} \in \mathbb{R}^{3 \times 3}$ and $\mathbf{b} \in \mathbb{R}^3$. Then the projected mean $\boldsymbol{\mu}_p \in \mathbb{R}^2$ and covariance $\boldsymbol{\Sigma}_p \in \mathbb{R}^{2 \times 2}$ are given as follows:

$$\boldsymbol{\mu}_p = \phi(\boldsymbol{\mu}_c) = \phi(\mathbf{W}\boldsymbol{\mu} + \mathbf{b}), \quad \boldsymbol{\Sigma}_p = \mathbf{J}\mathbf{W}\boldsymbol{\Sigma}\mathbf{W}^T\mathbf{J}^T, \quad (2)$$

where $\boldsymbol{\mu}_c \in \mathbb{R}^3$ is the mean in the camera coordinate system and $\mathbf{J} \in \mathbb{R}^{2 \times 3}$ is the Jacobian of ϕ . Notice that the $\boldsymbol{\Sigma}_p$ is the affine approximation. Then the projected 2D Gaussian can be modeled as an ellipse on the image plane.

Rendering. To render the color $C \in \mathbb{R}^3$ of the pixel, 3DGS uses differentiable alpha blending from ellipses of N Gaussians intersecting the pixel according to the depth order:

$$\mathbf{C} = \sum_{i=1}^n \mathbf{c}_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (3)$$

where α is evaluated by projected 2D Gaussians. To determine intersections between pixels and 2D Gaussians, 3DGS applies the technique of dividing the

image plane into several 16×16 tiles. Firstly each ellipse determines its axis-aligned bounding box. Then 3DGS traverses each bounding box to store the intersections and sorts them by tile IDs and depths. Finally, each pixel in a certain tile performs alpha blending on intersected Gaussians to render its color.

3.2 Equidistant Projection Model for Fisheye Cameras

Fisheye cameras use wide-angle lenses with a large field of view, creating images with significant radial distortion to capture extensive scenes. Unlike pin-hole cameras, which form images on a plane, fisheye cameras form images on a spherical surface. Their projection models include equidistant, equisolid angle, orthogonal, and stereographic projections.

We denote θ as the angle of incidence, and r_d as the distance from the optical center to the projection point. In equidistant projection, this distance equals the arc length on the projection plane, given by:

$$r_d = f\theta, \quad (4)$$

where f is the camera’s focal length. This model ensures equal distances for equal angles of incidence.

For the point $\mathbf{p} \in \mathbb{R}^3 = (x_c, y_c, z_c)^T$ in camera space, the distance to the z-axis is:

$$l_z = \sqrt{x_c^2 + y_c^2}. \quad (5)$$

The angle of incidence is $\theta = \arctan\left(\frac{\sqrt{x_c^2 + y_c^2}}{z_c}\right)$. Using equidistant projection,

the pixel coordinates are:

$$x_p = c_x + \frac{f_x \theta x_c}{l_z}, \quad y_p = c_y + \frac{f_y \theta y_c}{l_z}, \quad (6)$$

where f_x and f_y are the focal lengths in the x and y directions, and (c_x, c_y) is the optical center’s pixel coordinate.

4 Methods

4.1 Analysis of the Rendering Process of 3DGS

At first, we identify the projection-related components by analyzing the rendering process of 3D Gaussian Splatting (3DGS). As is shown in Figure 2, the rendering process is mainly composed of 4 processes, geometry preprocessing, tile binning, sorting, and rasterizing. The preprocessing is implemented by `preprocessCUDA`, which projects the Gaussians into a 2D surface via the projection transform φ . The tile binning implemented by `duplicateWithKeys` binds the tiles overlapping the corresponding Gaussians. Then the sorting process uses

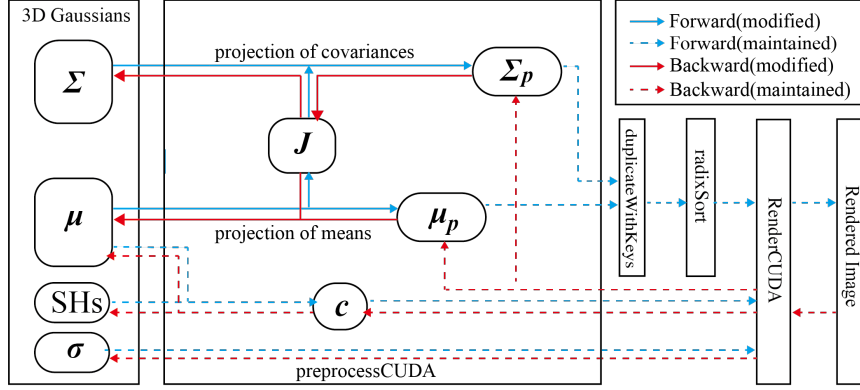


Fig. 2: Analysis and Modification of 3DGS pipeline. To apply a fisheye camera for 3DGS, we derive the equidistant projection and its Jacobian matrix. Then we implement the projection as the red arrows. We have also adjusted the back-propagation illustrated in the purple arrows to align the modified projection of the fisheye cameras. Our module enables both training and rendering of 3DGS for fisheye cameras.

radixSort to sort Gaussians with tiles according to tile IDs and depths. Finally, renderCUDA renders the Gaussians to image as the rasterizing process.

To apply fisheye cameras for 3DGS, we can only modify the projection transform according to equidistant projection ϕ in the kernel **preprocessing**. This approach constrains the scope of modifications, facilitating the subsequent implementation of lightweight modules and the extension to other projection types. In detail, we can only adjust the projection of means and covariances as the red arrows and the back-propagation as purple arrows in Figure 2.

4.2 The Modified Projection and Gradients of the Means

Forward (blue arrows in Fig 2). Suppose the mean in camera space is $\mu_c = (x_c, y_c, z_c)^T$ and the mean in pixel coordinates is $\mu_p = (x_p, y_p)^T$, the projection of means is simply applying the equidistant projection model ϕ as Section 3.2 shows:

$$\mu_p = \phi(\mu_c). \quad (7)$$

Backward (red arrows in Fig 2). Suppose the loss is \mathcal{L} , then the gradients of projected mean $\frac{\partial \mathcal{L}}{\partial \mu_p}$ can be calculated by vanilla 3DGS. To compute the mean in camera space $\frac{\partial \mathcal{L}}{\partial \mu_c}$, we can apply the chain rule:

$$\frac{\partial \mathcal{L}}{\partial \mu_c} = \frac{\partial \mu_p}{\partial \mu_c} \cdot \frac{\partial \mathcal{L}}{\partial \mu_p}. \quad (8)$$

Notice that $\frac{\partial \boldsymbol{\mu}_p}{\partial \boldsymbol{\mu}_c}$ is the Jacobian matrix of equidistant projection ϕ , which also needs to be modified:

$$\mathbf{J}_\phi = \frac{\partial \boldsymbol{\mu}_p}{\partial \boldsymbol{\mu}_c} = \begin{bmatrix} \frac{f_x x_c^2 z_c}{l_z^2 l_2} + \frac{y_c^2 \theta}{l_z^3} & x_c y_c \left(\frac{f_x z_c}{l_z^2 l_2} - \frac{\theta}{l_z^3} \right) & -\frac{f_x x_c}{l_2} \\ x_c y_c \left(\frac{f_y z_c}{l_z^2 l_2} - \frac{\theta}{l_z^3} \right) & \frac{f_y y_c^2 z_c}{l_z^2 l_2} + \frac{x_c^2 \theta}{l_z^3} & -\frac{f_y y_c}{l_2} \end{bmatrix}, \quad (9)$$

where l_2 is the square length of the vector $\boldsymbol{\mu}_p$, i.e., $l_2 = x_c^2 + y_c^2 + z_c^2$.

4.3 The Modified Projection and Gradients of the Covariance

Forward. Using the modified Jacobian matrix above, we can compute the projected covariance via equidistant projection. Let $\mathbf{T} = \mathbf{J}\mathbf{W} \in \mathbb{R}^{2 \times 3}$, then we have the projected covariance:

$$\boldsymbol{\Sigma}_p = \mathbf{T}\boldsymbol{\Sigma}\mathbf{T}^T. \quad (10)$$

Backward. As the matrix \mathbf{T} is the function of $\boldsymbol{\mu}_c$, the projection above propagates gradients to both transformation matrix \mathbf{T} and covariance $\boldsymbol{\Sigma}$. GSplat [34] and Vanilla 3DGS implementation [13] have derived the gradients above from the covariance projection by matrix calculus. Since the view transformation matrix \mathbf{W} is not a function of $\boldsymbol{\mu}$, we can directly modify the partial derivative of the Jacobian matrix from Vanilla 3DGS to propagate the gradients:

$$\frac{\partial \mathbf{J}}{\partial x_c} = \begin{bmatrix} \frac{f_x S_1}{l_z^2 l_2^5} & \frac{f_x S_2}{l_z^2 l_2^5} & \frac{f_x S_5}{l_2^2} \\ \frac{f_y S_2}{l_z^2 l_2^5} & \frac{f_y S_3}{l_z^2 l_2^5} & 2f_y x_c y_c \\ \frac{f_x S_5}{l_z^2 l_2^5} & \frac{f_y S_3}{l_z^2 l_2^5} & \frac{l_2^2}{l_2^2} \end{bmatrix} \quad (11)$$

$$\frac{\partial \mathbf{J}}{\partial y_c} = \begin{bmatrix} \frac{f_x S_2}{l_z^2 l_2^5} & \frac{f_x S_3}{l_z^2 l_2^5} & 2f_y x_c y_c \\ \frac{f_y S_3}{l_z^2 l_2^5} & \frac{f_y S_4}{l_z^2 l_2^5} & \frac{l_2^2}{l_2^2} \\ \frac{f_x S_2}{l_z^2 l_2^5} & \frac{f_y S_3}{l_z^2 l_2^5} & \frac{l_2^2}{l_2^2} \end{bmatrix}, \quad (12)$$

$$\frac{\partial \mathbf{J}}{\partial z_c} = \begin{bmatrix} \frac{f_x S_5}{l_2^2} & \frac{2f_x x_c y_c}{l_2^2} & \frac{2f_x x_c z_c}{l_2^2} \\ \frac{2f_y x_c y_c}{l_2^2} & \frac{f_y S_6}{l_2^2} & \frac{2f_y y_c z_c}{l_2^2} \\ \frac{f_x S_5}{l_2^2} & \frac{2f_x x_c y_c}{l_2^2} & \frac{2f_x x_c z_c}{l_2^2} \end{bmatrix}, \quad (13)$$

where

$$\begin{aligned} S_1 &= x_c(3S_7 + x_c^2 S_8), & S_2 &= y_c(S_7 + x_c^2 S_8), \\ S_3 &= x_c(S_7 + y_c^2 S_8), & S_4 &= y_c(3S_7 + y_c^2 S_8), \\ S_5 &= x_c^2 - y_c^2 - z_c^2, & S_6 &= y_c^2 - x_c^2 - z_c^2 \\ S_7 &= -l_z^2 l_2^2 \theta + l_2 l_z^3 z_c, & S_8 &= 3l_2^2 \theta - 3l_2 l_z z_c - 2l_z^3 z_c \end{aligned} \quad (14)$$

Finally, the mean $\boldsymbol{\mu}'$ accumulates the gradients from the backpropagation of the projection of the mean itself and its covariance to optimize the Gaussian's position. The projection of covariance also propagates gradients to the covariance itself.

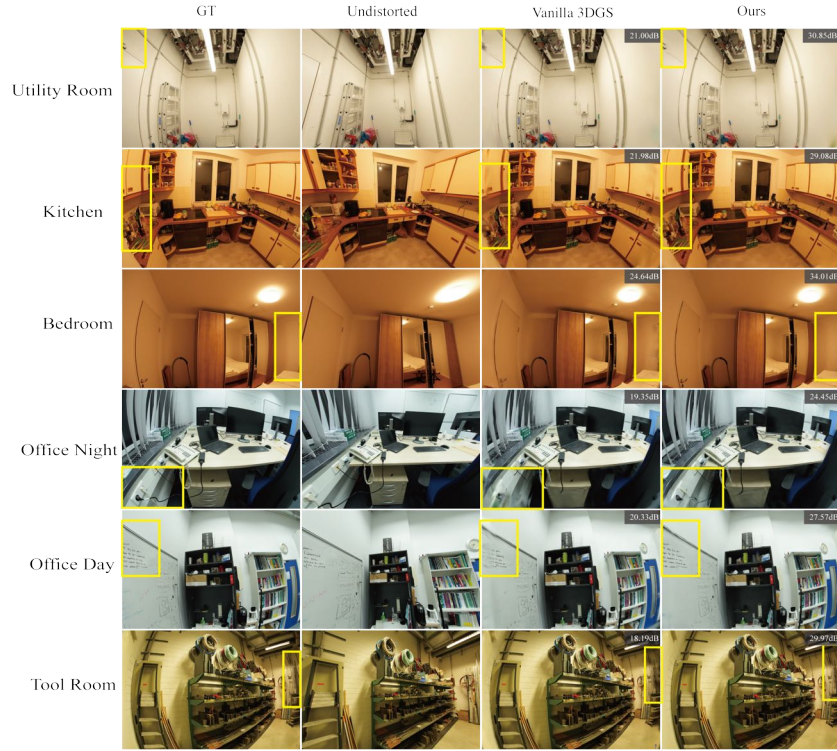


Fig. 3: Qualitative comparison between Fisheye-GS and the baseline. The baseline struggles to render on edges and corners due to the clipping and interpolation from undistortion.

5 Experiments

5.1 Experimental Setup

Dataset. We have adopted 6 indoor scenes from Scannet++ [35] dataset as our real-world dataset, where the images are captured by fisheye cameras and resized to the resolution of 1752×1168 . The dataset also provides point cloud, camera intrinsic, and poses of each view as COLMAP [21] SfM format. We have also generated a synthetic dataset from [19] with Blender 3.6 by equidistant fisheye camera models. We generate 200 frames with a resolution of 800×800 by rotating around the model for the Lego scene and 100 frames for other scenes. For each scene, we render the dataset separately for the field of view (FOV) settings of 120° and 180° . For every 8 frames of the scene, the first frame is used as the test set and the others are used for training.

Optimization. We have initialized Gaussians from point cloud on Scannet++ dataset as COLMAP format. For each, we train our models for 30000 iterations and use the same hyperparameters, training strategy, and loss function as [13].

Scene	Method	#Gaussians	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Utility Room	Vallina 3DGS	730320	21.69	0.821	0.215
	Ours	722639	26.42	0.894	0.177
Big Office	Vanilla 3DGS	1800937	20.21	0.72	0.264
	Ours	1442433	24.95	0.849	0.223
Small Office	Vanilla 3DGS	712107	24.02	0.87	0.190
	Ours	685902	26.15	0.905	0.177
Kitchen	Vanilla 3DGS	704898	25.54	0.867	0.211
	Ours	594699	30.63	0.933	0.179
Bedroom	Vanilla 3DGS	355658	25.91	0.882	0.223
	Ours	1442433	31.89	0.946	0.191
Tool Room	Vanilla 3DGS	3558209	19.19	0.63	0.316
	Ours	2605005	27.12	0.855	0.223

Table 1: Quantitative Comparison on Scannet++ dataset. Our Fisheye-GS achieves better quality than the baseline in all scenes and is evaluated by all metrics.

For the Scannet++ dataset, we randomize the background color for each iteration to model the walls of indoor scenes. Black backgrounds are used for training on synthetic datasets.

Baseline and Evaluation. We compare our method against original 3DGS [13] as our baseline. Because 3DGS is applied for the pinhole camera model, we have undistorted the images from the Scannet++ dataset for training. The baseline shares the same training technique as our method. For a fair comparison, we evaluate the baseline from ground-truth images before undistortion, using our Fisheye-GS on both training results. The evaluation resolution is the same as the training dataset. For the synthetic dataset, we have tested the rendering quality in different FOVs.

5.2 Visual Quality Evaluation

Evaluation on Real-world Dataset. We use visual PSNR, SSIM [26] and LPIPS [38] to test our method and the baseline on 6 real indoor scenes captured by fisheye cameras. As is shown in table 1, our method outperforms the baseline on each metric and scene on average. Fig. 3 shows that the baseline performs well on the central area of the image. However, artifacts frequently appear on the edge. Besides the interpolation of the undistortion process leading to the loss of information, the difference between the pinhole and fisheye camera models plays a key role. The fisheye camera has a larger field of view (FOV) compared to the pinhole camera. This wide FOV leads to an undistortion process that clips the scenes at the edges of the image. Then the 3DGS struggles to learn the edge area of the scene. In other words, the corners of the scene can be more frequently viewed by our Fisheye-GS, add more geometry constraints, and show fewer artifacts in novel views.

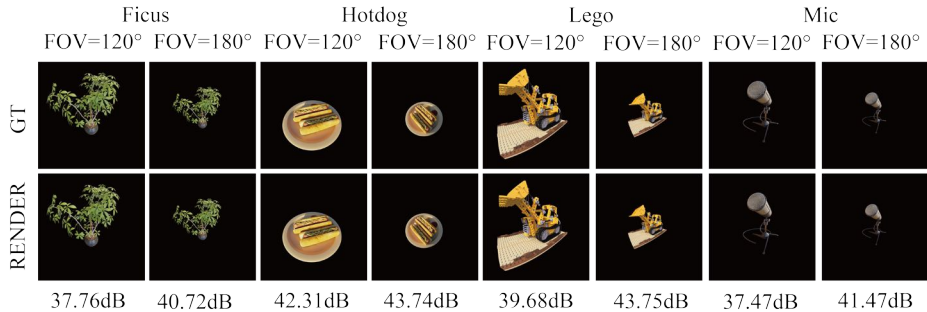


Fig. 4: Result on the synthetic dataset, varying from scenes and FOVs.

Evaluation on Synthetic Dataset. As is shown in table 2 and fig 4, our Fisheye-GS can successfully model the scene from random initialization, and adapt different views. The results show our method can adapt to larger distortions like Lego in fig 4.

Scene	FOV=120°				FOV=180°			
	#Gaussians	PSNR	SSIM	LPIPS	#Gaussians	PSNR	SSIM	LPIPS
Lego	144794	33.06	0.981	0.026	50357	30.62	0.9900	0.0141
Mic	47773	43.16	0.997	0.002	20371	47.77	0.9986	0.0013
Hotdog	62145	41.50	0.991	0.0128	37519	43.47	0.9956	0.0055
Ficus	201956	36.90	0.992	0.0084	85475	40.01	0.9962	0.0041

Table 2: Evaluation of the synthetic dataset from different scenes and FOVs.

5.3 Discussion and Limitation

Adaptation to Efficient Algorithms as Module Compared by op43dgs [10], Our method has only modified the geometry preprocessing and maintains other components as the vanilla 3DGS implements. This characteristic enables us to easily adapt Fisheye-GS as lightweight geometry preprocessing module to efficient rendering techniques like gsplat [33], FlashGS [8], etc. We have applied our method to FlashGS, the efficient rendering pipeline for 3DGS which reduces the computation redundancies and optimizes the rasterizing process. As Table 3 shows, we have evaluated FPS, maximum rendering time per frame, the number of intersections between tiles and projected Gaussians, and memory consumption via different GPUs, rendering techniques, and camera models. The results show that the speedup in fisheye cameras is close to that in pinhole cameras on different settings.

Extension to Other Cameras Models. We have also extended our method to other projection transformation. For panorama cameras, we have modified the projection transformation ϕ , the Jacobian matrix \mathbf{J} and their gradients according to the ideal panorama cameras. Fig 5 shows our method can both render the

Scene	GPU	Camera	Kernel	FPS↑	MaxTime	Intersects	Memory
					(ms)↓		
Small Office	A100	Fisheye	FlashGS	405	4.761	3034430	103.42
			Vanilla	121	17.155	13948372	1062.69
		Pinhole	FlashGS	510	4.399	2415742	89.26
			Vanilla	125	17.711	13145841	1035.06
	RTX 4090	Fisheye	FlashGS	286	4.614	3034430	103.42
			Vanilla	213	17.101	13948372	1062.69
		Pinhole	FlashGS	304	4.455	2415742	89.26
			Vanilla	189	15.000	13145841	1035.06
Kitchen	A100	Fisheye	FlashGS	570	4.038	2746724	93.36
			Vanilla	166	11.801	9834303	853.38
		Pinhole	FlashGS	534	3.793	2801239	94.61
			Vanilla	161	10.785	10945069	891.80
	RTX 4090	Fisheye	FlashGS	339	4.095	2746724	93.36
			Vanilla	282	12.002	9834303	853.38
		Pinhole	FlashGS	342	4.137	2801239	94.61
			Vanilla	264	13.999	10945069	891.80

Table 3: Performance Comparison on different GPUs w/o using FlashGS. The "Intersects" in the table is the number of intersections between the Gaussians and the overlapped tiles. The result shows that our method can adapt to Flash-GS in different GPUs, camera models, and scenes.

scenes with fisheye and panorama cameras. To extend our method to panorama cameras, we have to modify the projection as follows:

$$x_p = \frac{w(\operatorname{atan}_2(x_c, z_c) + \pi)}{2\pi} \quad (15)$$

$$y_p = \frac{h(\operatorname{atan}_2(y_c, \sqrt{x_c^2 + z_c^2}) + \frac{\pi}{2})}{\pi} \quad (16)$$

$$\mathbf{J}_\phi = \begin{bmatrix} \frac{Wz}{2\pi(x^2+z^2)} & 0 & -\frac{Wx}{2\pi(x^2+z^2)} \\ -\frac{Hxy}{\pi\sqrt{x^2+z^2}(x^2+y^2+z^2)} & \frac{H\sqrt{x^2+z^2}}{\pi(x^2+y^2+z^2)} & -\frac{Hy}{\pi\sqrt{x^2+z^2}(x^2+y^2+z^2)} \end{bmatrix} \quad (17)$$

where h, w is the height and width of the image, usually $w = 2h$.

Limitations. Our method derives and implements the training and rendering process of 3DGS for fisheye camera models. We also evaluate the rendering quality of our method. However, our method ignores the approximation error of the projection for covariances. And because we explicitly compute the gradients, our method struggles to be generalized to other distortion models. Besides, our method is only based on ideal camera models but has not provided a projection model for generic cameras and a calibration method for real large-FOV cameras. Due to the lack of fisheye camera datasets, our method has not been tested by various types of scenes, like unbounded scenes, city-scale scenes, etc.



Fig. 5: Comparison between rendered images in fisheye cameras and panorama cameras from the scene "Office Night" trained from Scannet++ dataset.

6 Conclusion

This paper introduces a new approach, Fisheye-GS, addressing the adaptability issues of 3D Gaussian Splatting (3DGS) across different camera models, particularly fisheye lenses. Fisheye-GS not only enables the reconstruction of 3DGS scenes captured by fisheye lenses but also functions as a lightweight module that can seamlessly integrate into high-performance 3DGS rendering algorithms. Additionally, it can be adapted for use with other camera models. This paper demonstrates that Fisheye-GS can effectively expand the application range of 3DGS technology to various camera perspectives, which is crucial for its broader real-life applications. We anticipate that the outcomes of this study will provide robust support and new perspectives for research and development in related fields.

Acknowledgment

This work was supported by the National Key R&D Program of China (Grant No. 2022ZD0160201).

References

1. Bai, J., Huang, L., Guo, J., Gong, W., Li, Y., Guo, Y.: 360-GS: Layout-guided panoramic gaussian splatting for indoor roaming. arXiv preprint arXiv:2402.00763 (2024) [4](#)
2. 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: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5855–5864 (2021) [3](#)
3. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5470–5479 (2022) [3](#)
4. Bourke, P.: idome: Immersive gaming with the Unity3D game engine. *Computer Games and Allied Technology* **9** (2009) [1](#)

5. Charatan, D., Li, S., Tagliasacchi, A., Sitzmann, V.: pixelSplat: 3D gaussian splats from image pairs for scalable generalizable 3D reconstruction. arXiv preprint arXiv:2312.12337 (2023) **3**
6. Chen, Y., Xu, H., Zheng, C., Zhuang, B., Pollefeys, M., Geiger, A., Cham, T.J., Cai, J.: MVSpLat: Efficient 3D gaussian splatting from sparse multi-view images. arXiv preprint arXiv:2403.14627 (2024) **3**
7. Deng, L., Yang, M., Qian, Y., Wang, C., Wang, B.: CNN based semantic segmentation for urban traffic scenes using fisheye camera. In: 2017 IEEE Intelligent Vehicles Symposium (IV). pp. 231–236. IEEE (2017) **1**
8. Feng, G., Chen, S., Fu, R., Liao, Z., Wang, Y., Liu, T., Pei, Z., Li, H., Zhang, X., Dai, B.: Flashgs: Efficient 3d gaussian splatting for large-scale and high-resolution rendering (2024), <https://arxiv.org/abs/2408.07967> **2, 3, 10**
9. 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) **3**
10. Huang, L., Bai, J., Guo, J., Li, Y., Guo, Y.: On the error analysis of 3d gaussian splatting and an optimal projection strategy (2024), <https://arxiv.org/abs/2402.00752> **4, 10**
11. Jakab, D., Deegan, B.M., Sharma, S., Grua, E.M., Horgan, J., Ward, E., Ven, P.v.d., Scanlan, A., Eising, C.: Surround-view fisheye optics in computer vision and simulation: Survey and challenges. IEEE Transactions on Intelligent Transportation Systems p. 1–22 (2024) **1**
12. Jeong, Y., Ahn, S., Choy, C., Anandkumar, A., Cho, M., Park, J.: Self-calibrating neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5846–5854 (2021) **3**
13. Kerbl, B., Kopanas, G., Leimkuehler, T., Drettakis, G.: 3D gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics (TOG) **42**, 1–14 (2023) **2, 3, 7, 8, 9**
14. Konrad, J., Cokbas, M., Ishwar, P., Little, T.D., Gevelber, M.: High-accuracy people counting in large spaces using overhead fisheye cameras. Energy and Buildings **307**, 113936 (2024) **1**
15. Kulkarni, S., Yin, P., Scherer, S.: 360FusionNeRF: Panoramic neural radiance fields with joint guidance. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 7202–7209. IEEE (2023) **3**
16. Li, J., Zhang, J., Bai, X., Zheng, J., Ning, X., Zhou, J., Gu, L.: DNGaussian: Optimizing sparse-view 3D gaussian radiance fields with global-local depth normalization. arXiv preprint arXiv:2403.06912 (2024) **3**
17. Liang, Z., Zhang, Q., Hu, W., Feng, Y., Zhu, L., Jia, K.: Analytic-splatting: Anti-aliased 3D gaussian splatting via analytic integration. arXiv preprint arXiv:2403.11056 (2024) **3**
18. Lu, T., Yu, M., Xu, L., Xiangli, Y., Wang, L., Lin, D., Dai, B.: Scaffold-GS: Structured 3D gaussians for view-adaptive rendering. arXiv preprint arXiv:2312.00109 (2023) **3**
19. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021) **2, 3, 8**
20. Ren, K., Jiang, L., Lu, T., Yu, M., Xu, L., Ni, Z., Dai, B.: Octree-GS: Towards consistent real-time rendering with lod-structured 3D gaussians. arXiv preprint arXiv:2403.17898 (2024) **3**

21. 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) [8](#)
22. Shen, J., Song, B., Wu, Z., Xu, Y.: OmniNeRF: Hybridizing omnidirectional distance and radiance fields for neural surface reconstruction. In: 2022 2nd International Conference on Computational Modeling, Simulation and Data Analysis (CMSDA). pp. 281–285. IEEE (2022) [3](#)
23. Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A., et al.: NeRFstudio: A modular framework for neural radiance field development. In: ACM SIGGRAPH 2023 Conference Proceedings. pp. 1–12 (2023) [3](#)
24. Vandewiele, F., Motamed, C., Yahiaoui, T.: Visibility management for object tracking in the context of a fisheye camera network. In: 2012 Sixth International Conference on Distributed Smart Cameras (ICDSC). pp. 1–6. IEEE (2012) [1](#)
25. Wang, G., Wang, P., Chen, Z., Wang, W., Loy, C.C., Liu, Z.: PERF: Panoramic neural radiance field from a single panorama. arXiv preprint arXiv:2310.16831 (2023) [4](#)
26. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing **13**(4), 600–612 (2004) [9](#)
27. Xian, W., Božič, A., Snavely, N., Lassner, C.: Neural lens modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8435–8445 (2023) [4](#)
28. Xiong, H., Muttukuru, S., Upadhyay, R., Chari, P., Kadambi, A.: SparseGS: Real-time 360° sparse view synthesis using gaussian splatting. arXiv e-prints pp. arXiv–2312 (2023) [3](#)
29. Xiong, Y., Turkowski, K.: Creating image-based VR using a self-calibrating fisheye lens. In: Proceedings of IEEE computer society conference on computer vision and pattern recognition. pp. 237–243. IEEE (1997) [1](#)
30. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Pointerf: Point-based neural radiance fields. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 5428–5438 (2022), <https://api.semanticscholar.org/CorpusID:246210101> [3](#)
31. Yan, Z., Low, W.F., Chen, Y., Lee, G.H.: Multi-scale 3D gaussian splatting for anti-aliased rendering. arXiv preprint arXiv:2311.17089 (2023) [3](#)
32. Yang, Y., Deng, H.: EPformer: an efficient transformer-based approach for retail product detection in fisheye images. Journal of Electronic Imaging **32**(1), 013017–013017 (2023) [1](#)
33. Ye, V., Kanazawa, A.: gsplat, <https://github.com/nerfstudio-project/gsplat> [3](#), [10](#)
34. Ye, V., Kanazawa, A.: Mathematical supplement for the gsplat library. arXiv preprint arXiv:2312.02121 (2023) [7](#)
35. Yeshwanth, C., Liu, Y.C., Nießner, M., Dai, A.: ScanNet++: A high-fidelity dataset of 3D indoor scenes. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12–22 (2023) [8](#)
36. Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenotrees for real-time rendering of neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5752–5761 (2021) [3](#)
37. Yu, Z., Chen, A., Huang, B., Sattler, T., Geiger, A.: Mip-Splatting: Alias-free 3D gaussian splatting. arXiv preprint arXiv:2311.16493 (2023) [3](#)

38. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018) [9](#)
39. Zhang, S., Zhao, W., Wang, J., Luo, H., Feng, X., Peng, J.: A mixed-reality museum tourism framework based on HMD and fisheye camera. In: Proceedings of the 15th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its Applications in Industry-Volume 1. pp. 47–50 (2016) [1](#)
40. Zhu, Z., Fan, Z., Jiang, Y., Wang, Z.: FSGS: Real-time few-shot view synthesis using gaussian splatting. arXiv preprint arXiv:2312.00451 (2023) [3](#)