

Decoupling Appearance Variations with 3D Consistent Features in Gaussian Splatting

Jiaqi Lin¹, Zhihao Li², Binxiao Huang³, Xiao Tang², Jianzhuang Liu⁴,
Shiyong Liu², Xiaofei Wu², Fenglong Song², Wenming Yang^{1*}

¹Tsinghua University

²Huawei Noah's Ark Lab

³The University of Hong Kong

⁴Shenzhen Institute of Advanced Technology

Abstract

Gaussian Splatting has emerged as a prominent 3D representation in novel view synthesis, but it still suffers from appearance variations, which are caused by various factors, such as modern camera ISPs, different time of day, weather conditions, and local light changes. These variations can lead to floaters and color distortions in the rendered images/videos. Recent appearance modeling approaches in Gaussian Splatting are either tightly coupled with the rendering process, hindering real-time rendering, or they only account for mild global variations, performing poorly in scenes with local light changes. In this paper, we propose DAVIGS, a method that decouples appearance variations in a plug-and-play and efficient manner. By transforming the rendering results at the image level instead of the Gaussian level, our approach can model appearance variations with minimal optimization time and memory overhead. Furthermore, our method gathers appearance-related information in 3D space to transform the rendered images, thus building 3D consistency across views implicitly. We validate our method on several appearance-variant scenes, and demonstrate that it achieves state-of-the-art rendering quality with minimal training time and memory usage, without compromising rendering speeds. Additionally, it provides performance improvements for different Gaussian Splatting baselines in a plug-and-play manner.

Introduction

As a promising representation in novel view synthesis, Gaussian Splatting (GS) enables high-quality reconstruction and real-time rendering. Since the introduction of 3DGS (Kerbl et al. 2023), this representation has evolved into several variants, including Mip-Splatting (Yu et al. 2024b), 2DGS (Huang et al. 2024a), and GOF (Yu, Sattler, and Geiger 2024), among others (Dai et al. 2024; Yang et al. 2024; Yu et al. 2024a; Lu et al. 2024; Niemeyer et al. 2024; Zhao, Wang, and Liu 2024). These methods have been applied to various tasks, such as avatar creation (Li et al. 2024; Moreau et al. 2024; Kocabas et al. 2024; Abdal et al. 2024; Qian et al. 2024; Zheng et al. 2024; Jiang et al. 2024; Liu et al. 2024a), autonomous driving (Zhou et al. 2024c; Yan et al. 2024b; Zhao et al. 2024; Zhou et al. 2024a), 3D asset generation (Chen et al. 2024; Tang et al. 2023; Yi et al. 2024;

*Corresponding author.




Figure 1: (a) There are often global and local (orange box) appearance variations in real-world captures for 3D reconstruction. (b) Such appearance variations may lead to floaters. (c) Compared to existing SOTA methods, our DAVIGS achieves advanced reconstruction results with faster optimization and rendering.

Tang et al. 2024; Ling et al. 2024; Zhou et al. 2024b), and SLAM (Keetha et al. 2024; Matsuki et al. 2024; Yan et al. 2024a; Huang et al. 2024b). However, the impressive performance of GS relies on strict multi-view consistency, which is often compromised by appearance variations in real-world captures. Many factors, such as different camera image signal processings (ISPs), time of day, weather conditions, and local light changes can lead to noticeable alterations in the captured images, resulting in floaters and color distortions in the reconstructed scenes, as shown in Fig. 1.

Previous appearance modeling methods (Martin-Brualla et al. 2021; Chen et al. 2022; Lin et al. 2024; Zhang et al. 2024) typically optimize for an average scene appearance, followed by view-dependent transformations to match the actual appearances from different viewpoints. In NeRF-based methods (Martin-Brualla et al. 2021; Chen et al. 2022), this transformation occurs at the sampled points and is supervised in a pixel-wise manner. Appearance embeddings are attached to the point features and fed into a multi-layer perceptron (MLP) to generate view-dependent colors. GS-based methods (Lin et al. 2024; Zhang et al. 2024) use a discrete 3D representation and are optimized in a frame-wise manner. Their appearance modeling can be categorized

into two kinds: decoupled and coupled, as shown in Fig. 2.

The decoupled methods (Lin et al. 2024; Darmon et al. 2024; Niemeyer et al. 2024) perform transformations directly at the image level. These methods are advantageous due to their low computational and video memory requirements. Additionally, the decoupling nature allows them to be discarded after optimization, thereby not impacting rendering speed. However, their primary drawback is the limited expressiveness, as existing decoupled methods are only suitable for mild and global changes, and struggle with more complex variations.

The coupled methods (Zhang et al. 2024; Kulhanek et al. 2024) model appearance variations in the 3D space, and achieve stronger expressiveness by transforming Gaussian properties, such as color and opacity. This direct manipulation of 3D representations provides larger capacities to handle complex appearance variations. However, they are less efficient during rendering, as they cannot function without their appearance modeling. Moreover, they require changes of the Gaussian properties to store additional information.

In this paper, we present DAVIGS (Decoupling Appearance Variations In GS), a novel appearance modeling method that integrates the strengths of both decoupled and coupled approaches. Our method transforms the rendering results at the image level, retaining the low-cost advantage of decoupled appearance modeling approaches without compromising rendering speed. To address the weakness of the decoupled methods, we incorporate 3D consistent features into the image transformation process, ensuring strong expressiveness.

We conceptualize appearance variation information as two components: global appearance embeddings associated with each image, and local features for different responses of various 3D locations, stored in multi-resolution hash grids. These features are decoded by a lightweight MLP into an affine transformation matrix for each pixel in the image. To stabilize the optimization process and prevent unpleasant colors in the reconstructed scene, we introduce an identity regularization term that constrains the transformation to be close to the identity transformation. Additionally, we propose a cell-based query approach to minimize the additional optimization time introduced by our appearance modeling.

Our DAVIGS method can be adapted to a wide range of complex appearance variations and ensures 3D consistency across views. Our contributions are summarized as follows:

- We propose DAVIGS, a novel decoupled appearance modeling approach that effectively mitigates appearance variation-induced floaters and color distortions with low optimization time and video memory overhead, without slowing down rendering speeds.
- We introduce a transformation regularizer and a cell-based query method to ensure stable and efficient optimization of our decoupled appearance module.
- Experiments conducted on both our new dataset and existing datasets demonstrate that DAVIGS achieves state-of-the-art performance in rendering quality, optimization time, video memory, and rendering speeds. Besides, it is plug-and-play for various GS methods.

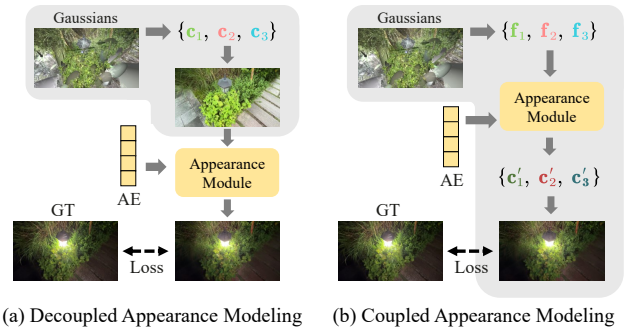


Figure 2: **Different appearance modelings.** The gray parts represent the rendering process in Gaussian Splatting, where f denotes the features of Gaussian primitives, and c denotes the colors. (a) Decoupled appearance modeling applies transformations to rendering results at the image level, which can be discarded after optimization. (b) Coupled appearance modeling addresses appearance variations by manipulating the features of Gaussian primitives, making it coupled with the rendering process.

Related Work

There are two categories of appearance modeling approaches: decoupled and coupled, as shown in Fig. 2.

Decoupled Appearance Modeling

Decoupled appearance modeling applies transformations to rendering results at the image level. VastGaussian (Lin et al. 2024) first proposes such an appearance modeling methodology. It optimizes an appearance embedding (AE) for each image, which is concatenated to the downsampled rendered image and fed into a convolutional neural network (CNN) to decode a transformation map on the image. Robust Gaussian Splatting (Darmon et al. 2024) adopts a similar approach, but it directly optimizes a set of affine transformation parameters for each image. RadSplat (Niemeyer et al. 2024) introduces an exposure processing module that adjusts the rendered image based on the image’s exposure duration and ISO setting.

The key advantage of these methods lies in their decoupling nature. Once optimized, the appearance modules can be discarded without slowing down the rendering speed. They are generally efficient during optimization, but their expressiveness is limited, as they can only adapt to relatively mild changes, such as camera auto-exposure and auto-white balance. Our proposed method, DAVIGS, belongs to this category but overcomes this limitation with 3D consistent features across views.

Coupled Appearance Modeling

Coupled appearance modeling operates directly on the 3D representations. In NeRF-based methods (Martin-Brualla et al. 2021; Chen et al. 2022), AEs are attached to 3D sampled points in ray-marching and processed through an MLP to determine the points’ colors. For GS-based methods, the manipulation is performed on Gaussian primitives. SWAG




Figure 3: **Overall pipeline of DAVIGS.** For each pixel p in the rendered image \mathcal{I}^r , we calculate its 3D spatial position \mathbf{x} by back-projecting with its depth $\mathcal{D}(p)$ in the depth map \mathcal{D} , and then look up the multi-resolution hash grids for its 3D consistent features $\{\mathbf{f}_i\}$. They are then concatenated with a view-dependent appearance embedding \mathbf{l} and fed into an MLP f to obtain a transformation matrix $\mathcal{M}(p) \in \mathbb{R}^{3 \times 4}$, which is used to perform affine transformation on the color $\mathcal{I}^r(p)$ to obtain $\mathcal{I}^t(p)$. The losses \mathcal{L}_1 and $\mathcal{L}_{D\text{-SSIM}}$ are calculated between the transformed image \mathcal{I}^t and the ground truth image \mathcal{I} . A regularization term \mathcal{L}_{ID} is applied to $\mathcal{M}(p)$ for constraining it close to the identity transformation matrix \mathcal{M}_{ID} .

(Dahmani et al. 2024) optimizes hash grids and an MLP along with the attributes of 3D Gaussians. For each 3D Gaussian, the features of its position, its color and an AE are fed into the MLP to compute the color and an opacity residual. Wild-GS (Xu, Mei, and Patel 2024) takes a more complex approach by feeding a global appearance embedding extracted from the image, local features obtained from the back-projection of the image, and intrinsic features of 3D Gaussians into an MLP to decode spherical harmonics (SH) coefficients. GS-W (Zhang et al. 2024) uses adaptive sampling for each 3D Gaussian to sample appearance features from the feature maps extracted from the input image. The features are blended and processed through an MLP along with a view direction to obtain the color. WildGaussians (Kulhanek et al. 2024) extends the attributes of 3D Gaussians by adding an internal embedding, which is combined with the image’s AE and processed through an appearance MLP to obtain an affine mapping for determining view-dependent colors.

These coupled methods are capable of handling complex appearance variations, such as day-night transitions and different weather conditions. However, they require additional attributes within Gaussian primitives and coupled MLP inference during rendering, which significantly increases optimization time and video memory, and slows down rendering speeds.

Method

Preliminaries

In this section, we present our novel DAVIGS method for efficient, plug-and-play appearance modeling in Gaussian Splatting. Since our primary experiments are conducted on 3D Gaussian Splatting (3DGS), which serves as the foundation for other GS variants, we begin with its basic preliminaries. 3DGS represents both geometry and appearance using a set of 3D Gaussians, each characterized by its position, anisotropic covariance, opacity, and spherical harmonic coefficients for view-dependent colors. During rendering, each

3D Gaussian is projected onto the image plane as a 2D Gaussian. These projected 2D Gaussians are then assigned to different tiles, sorted, and alpha-blended into the final image through a point-based volume rendering process.

The dataset used for scene optimization consists of a sparse point cloud and a set of training views. During the optimization process of 3DGS, one image \mathcal{I}^r is rendered at each iteration using a differentiable rasterizer. This rendered image is then compared to the corresponding ground truth image \mathcal{I} to compute the L1 loss and the D-SSIM loss. The overall loss function is given by:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_1(\mathcal{I}^r, \mathcal{I}) + \lambda\mathcal{L}_{D\text{-SSIM}}(\mathcal{I}^r, \mathcal{I}), \quad (1)$$

where λ is a hyperparameter. The gradients obtained through backpropagation of this loss are used to optimize the properties of the 3D Gaussians. This optimization process is interleaved with adaptive point densification, which is triggered when the cumulative gradient of a point reaches a predefined threshold.

DAVIGS

We introduce DAVIGS to solve the appearance variation problem in novel view synthesis, which includes global changes (e.g., image ISPs and daylight) and local changes (e.g., street lamps at night). To achieve decoupling, we model appearance variations at the image level rather than at the Gaussian level. We use globally shared appearance embeddings (AEs) for each view, and learn distinct transformations for local changes of individual pixels. To query 3D consistent local features, we compute the 3D position of each pixel, and then calculate the transformation based on both the global AE and the local spatial features at that pixel.

Architecture As illustrated in Fig. 3, we encode the global appearance of each view using an AE vector $\mathbf{l} \in \mathbb{R}^A$, store the local spatial features using multi-resolution hash grids $\{\mathbf{G}_i\}_{i=0}^{L-1}$ with L resolution levels (Müller et al. 2022), and

employ a lightweight MLP f to learn affine transformation matrices. We transform the rendered image \mathcal{I}^r at the image level. Specifically, during the differentiable rasterization process, we also render a depth map \mathcal{D} . For each pixel $\mathcal{I}^r(p)$ in the rendered image, we compute its corresponding 3D coordinate $\mathbf{x} \in \mathbb{R}^3$ via back-projection using its depth value $\mathcal{D}(p)$ and the camera parameters. Similar to grid-based methods in NeRF (Müller et al. 2022), we look up the surrounding voxels of coordinate \mathbf{x} in the hash grids $\{\mathbf{G}_i\}_{i=0}^{L-1}$ and linearly interpolate the corner features to obtain the positional features of \mathbf{x} , denoted as $\{\mathbf{f}_i \in \mathbb{R}^F\}_{i=0}^{L-1}$. We then concatenate these positional features $\{\mathbf{f}_i\}_{i=0}^{L-1}$ with the corresponding appearance embedding $\mathbf{l} \in \mathbb{R}^A$ of the current view:

$$\mathbf{f} = \mathbf{f}_0 \oplus \mathbf{f}_1 \oplus \dots \oplus \mathbf{f}_{L-1} \oplus \mathbf{l}. \quad (2)$$

The resulting vector $\mathbf{f} \in \mathbb{R}^{LF+A}$ is then fed into the MLP f :

$$f(\mathbf{f}) = \mathbf{v}. \quad (3)$$

The output vector $\mathbf{v} \in \mathbb{R}^{12}$ is reshaped into an affine transformation matrix $\mathcal{M}(p) \in \mathbb{R}^{3 \times 4}$. Finally, we apply this affine transformation to the RGB color of the pixel:

$$\mathcal{I}^t(p) = \mathcal{I}^r(p)\mathcal{M}(p), \quad (4)$$

where \mathcal{I}^t is the transformed image. This appearance modeling method encodes the global view-dependent appearance information in the appearance embeddings, while the local responses of different locations on the scene’s surface to the appearance features are sparsely stored in the hash grids. This approach ensures 3D consistency in both global and local appearances in the reconstructed scene. Since our approach transforms the rendered results at the image level, the module is plug-and-play for various GS methods and can be discarded after optimization without slowing down rendering speeds.

Optimization Without regularization of the appearance learning process, the Gaussian primitives tend to adjust their colors instead of allowing the appearance model to optimize its parameters to fit the appearance variations. Although the final optimized scene may also converge to an average appearance, it might not be visually pleasing; e.g., a homogeneous surface could become blotchy. To address this, we introduce a regularization term \mathcal{L}_{ID} on the affine transformation matrix $\mathcal{M}(p)$ to ensure that it remains close to the identity transformation matrix \mathcal{M}_{ID} :

$$\mathcal{L}_{\text{ID}} = \mathcal{L}_1(\mathcal{M}_{\text{ID}}, \mathcal{M}(p)), \quad \mathcal{M}_{\text{ID}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \quad (5)$$

The final loss function is:

$$\mathcal{L} = (1 - \lambda_1)\mathcal{L}_1(\mathcal{I}^t, \mathcal{I}) + \lambda_1\mathcal{L}_{\text{D-SSIM}}(\mathcal{I}^t, \mathcal{I}) + \lambda_2\mathcal{L}_{\text{ID}}, \quad (6)$$

where λ_1 and λ_2 are hyperparameters. Although \mathcal{I}^r is not directly supervised by the ground truth image \mathcal{I} , it is still properly constrained in an indirect manner, thanks to the regularization term \mathcal{L}_{ID} which penalizes the difference between \mathcal{I}^t and \mathcal{I}^r .




Figure 4: **Cell-based query.** (a) We divide the depth map \mathcal{D} into non-overlapping cells. For cell $\mathcal{C}_{i,j}$ with mean depth $\mathcal{D}_{i,j}$, we perform back-projection on its center $\mathbf{o}_{i,j}$ to obtain spatial coordinate $\mathbf{x}_{i,j}$. We use $\mathbf{x}_{i,j}$ to query the appearance module to obtain the transformation matrix $\mathcal{M}_{i,j}$. (b) For pixel p , we transform its color with the bilinearly interpolated affine transformation matrix.

Cell-Based Query Querying the appearance model for every pixel is time-consuming, so we implement a cell-based query to compute the affine transformation matrix for one cell at a time. Specifically, as shown in Fig. 4, let the resolution of the rendered image be $H \times W$; we divide it into non-overlapping cells in 2D with step size s , each of size $\frac{H}{s} \times \frac{H}{s}$. We then take the average of the depth map \mathcal{D} within each cell indexed by i, j (denoted as $\mathcal{D}_{i,j}$) and use the cell’s central position $\mathbf{o}_{i,j} = (\frac{H}{s}i + \frac{H}{2s}, \frac{W}{s}j + \frac{W}{2s})$ to compute its 3D position via back-projection, denoted as $\mathbf{x}_{i,j}$. We use $\mathbf{x}_{i,j}$ to carry out the query in our appearance module to obtain the affine transformation matrix for the cell, denoted as $\mathcal{M}_{i,j}$. For any pixel p in the rendered image, we use the affine transformation matrices of its neighboring cells to perform bilinear interpolation to obtain its affine transformation $\mathcal{M}(p)$. This cell-based query effectively reduces the optimization time. Given that spatially adjacent pixels tend to have similar appearance responses, this approach allows for faster optimization with minimal compromise in rendering quality.

Experiments

Experimental Setups

Implementation Details In our main experiments, DAVIGS is implemented based on 3DGS. The MLP f contains two hidden layers with widths of 128 and 64, respectively. The hyperparameter λ_1 is set to 0.2, while λ_2 follows a cosine annealing function, warming up linearly from 0 to 0.3 in the first 5k iterations and eventually decaying to 0.2. For the depth used in back-projection, we employ the mean depth (Huang et al. 2024a) calculated during differentiable rasterization. For cell-based query, we use a cell size of 8.

Datasets Our DAVIGS method focuses on decoupling appearance variations to alleviate floaters and color distortions. Existing datasets for evaluating this task are limited: some contain only global and mild appearance variations, such as the Mill-19 dataset (Turki, Ramanan, and Satyanarayanan



Figure 5: Qualitative comparison between DAVIGS and previous work. Floaters and other artifacts are pointed out by arrows.

2022), which primarily includes camera ISP variations; others contain other distractors, like the PhotoTourism dataset (Jin et al. 2021), which includes both appearance changes and transient objects, alongside with varying image resolutions and qualities. To facilitate a fair comparison, we create a high-quality *GLAV dataset* comprising seven scenes that exhibit both **G**lobal and **L**ocal **A**ppearance **V**ariations. The images are captured with an iPhone 15 Pro camera at 1080p resolution and do not have any distractors other than appearance variations. The camera poses and sparse point clouds are estimated using COLMAP (Schönberger and Frahm 2016). Additionally, we also evaluate our method on the PhotoTourism dataset, as it is widely used by other methods.

Metrics We evaluate rendering quality using PSNR, SSIM, and LPIPS, and assess efficiency using optimiza-

tion time, VRAM usage, and rendering speed. Following the common practice, we use the left half of the test images during training to optimize the appearance-related parameters and reserve the right half for testing. All the metrics are tested on a single Tesla V100 GPU.

Compared Methods We compare DAVIGS with several other methods, including NeRF-W (Martin-Brualla et al. 2021), Ha-NeRF (Chen et al. 2022), VastGaussian (Lin et al. 2024), Gaussian-W (Zhang et al. 2024), and WildGaussians (Kulhanek et al. 2024). For NeRF-based methods, we follow the default parameter settings of Ha-NeRF, training for 20 epochs with a batch size of 1024. For GS-based methods, we optimize the models for 30k iterations, keeping other parameters the same as those in the original papers.

Dataset	GLAV					
Metrics	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Training \downarrow	VRAM \downarrow	FPS \uparrow
NeRF-W	14.13	0.470	0.496	69h48min	8.3G	<1
Ha-NeRF	21.09	0.613	0.381	82h49min	7.7G	<1
3DGS	22.80	0.753	0.221	16min	<u>3.8G</u>	134
GS-W	27.23	0.839	0.141	1h47min	8.1G	24
GS-W-c	27.23	0.839	0.141	1h47min	8.1G	84
WildGS	24.68	0.764	0.244	1h18min	22.7G	28
VastGS	23.21	0.782	0.186	28min	4.7G	<u>148</u>
DAVIGS	28.31	0.861	0.134	<u>19min</u>	3.4G	151

Table 1: Effectiveness (PSNR, SSIM, and LPIPS) and efficiency (training time, training VRAM usage, and rendering speed) evaluations on the GLAV dataset. Our method achieves superior performance in all the metrics.

Dataset	PhotoTourism					
Metrics	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Training \downarrow	VRAM \downarrow	FPS \uparrow
NeRF-W	20.78	0.799	0.208	114h20min	8.3G	<1
Ha-NeRF	21.41	0.793	0.178	271h13min	7.7G	<1
3DGS	22.39	0.789	0.232	20min	6.7G	130
GS-W	<u>23.81</u>	0.841	0.135	2h28min	13.9G	18
GS-W-c	23.81	0.841	0.135	2h28min	13.9G	49
WildGS	22.81	0.824	0.158	1h40min	18.8G	22
VastGS	23.08	0.827	0.153	43min	6.9G	111
DAVIGS	23.95	<u>0.838</u>	<u>0.148</u>	<u>26min</u>	6.3G	<u>123</u>

Table 2: Despite some other methods apply additional modules (e.g., sky handling) for the PhotoTourism dataset, DAVIGS achieves comparable high quality with minimal training time and VRAM usage.

Result Analysis

As shown in Tab. 1 and Fig. 5, DAVIGS demonstrates superior rendering quality in comparison with other methods. 1) Compared to *vanilla 3DGS*, it is slightly slower during training, but has significant improvement in rendering quality. Since it reduces the floaters caused by appearance variances, it is slightly faster in rendering. Despite incorporating additional modules, DAVIGS requires less VRAM than 3DGS. 2) Compared to the *coupled appearance modeling methods* (GS-W, GS-W with cached features denoted as GS-W-c, and WildGS), DAVIGS has substantial advantages in terms of optimization time and VRAM usage since we transform the rendering results at the image level. The decoupled nature of DAVIGS also results in a much faster rendering process. 3) Compared to another *decoupled appearance modeling method* (VastGS), ours provides better 3D consistency, effectively resolving floaters and color distortions, leading to significant improvements in rendering quality. On the PhotoTourism dataset used by other methods, DAVIGS also achieves comparable quality with minimal training time and VRAM usage (see Tab. 2).

Plug-and-Play on Different Baselines

To demonstrate the plug-and-play nature of our method, we implement DAVIGS on some other representative GS-based baselines, including Mip-Splatting (Yu et al. 2024b), 2DGS (Huang et al. 2024a), and GOF (Yu, Sattler, and Geiger




Figure 6: DAVIGS can be plugged into different baselines to suppress floaters caused by appearance variations.

Metrics	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Training \downarrow	VRAM \downarrow
3DGS	22.80	0.753	0.221	16min8s	3.76G
3DGS+Ours	28.31	0.861	0.134	18min51s	3.39G
Mip-Splatting	22.46	0.748	0.295	22min17s	13.74G
Mip-Splatting+Ours	28.31	0.870	0.130	24min52s	14.17G
GOF	22.38	0.737	0.231	27min52s	18.11G
GOF+Ours	28.21	0.861	0.139	28min9s	12.43G
2DGS	22.70	0.746	0.229	18min7s	6.16G
2DGS+Ours	27.76	0.842	0.158	20min	3.16G

Table 3: Our DAVIGS can be implemented into different baselines efficiently in a plug-and-play manner.

2024). The parameters of our hash grids and MLP remain the same as those in the previous subsection, while the other parameters are kept as in the original papers. We compare these baselines with and without our appearance module on the GLAV dataset. As shown in Tab. 3 and Fig. 6, integrating our module significantly improves the performance of these baselines. The integration requires minimal additional training time and even reduces VRAM usage due to the suppression of floaters, demonstrating the efficiency and plug-and-play capability of DAVIGS.

Ablation Studies

To assess the effectiveness of our introduced modules, we conduct ablation experiments on the GLAV dataset. Following Mip-NeRF 360 (Barron et al. 2022) and VastGaussian (Lin et al. 2024), we perform color correction before calculating the metrics to evaluate the quality of the average appearance (see the supplementary material).

3D-Aware Information To illustrate the importance of incorporating 3D-aware information in decoupled appearance modeling, we replace the 3D consistent features with four types of degenerated information: the uv -coordinates of each pixel in the rendered image, the depth, a combination of the previous two, and the RGB color. Each of these is embedded with positional encoding to match the length of




Figure 7: The 3D features (xyz) in DAVIGS are necessary for suppressing floaters.




Figure 8: Replacing the 3D features (xyz) with pixel colors or removing the transformation regularizer causes color distortions in the reconstructed scene.

our 3D features. As shown in Tab. 4 and Fig. 7, using pixel coordinates in the image or/and depth information significantly degrades the rendering quality compared to using the 3D information. When using only color information, the final optimized scene suffers from severe color distortions, as seen in Fig. 8.

Model Setting	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
$xyz \rightarrow uv$ coordinate	21.76	0.798	0.252
$xyz \rightarrow depth$	20.04	0.788	0.268
$xyz \rightarrow uv + depth$	22.12	0.802	0.248
$xyz \rightarrow color$	21.58	0.765	0.317
w/o \mathcal{L}_{ID}	21.33	0.767	0.298
Full Model	22.75	0.803	0.243

Table 4: Ablation on the introduction of the 3D features (xyz) and transformation regularization term \mathcal{L}_{ID} .

Transformation Regularization As shown in Tab. 4 and Fig. 8, removing the transformation regularizer allows the appearance module more freedom in learning appearances, which results in rendering quality degradation and color distortions in the optimized scene.




Figure 9: Visual quality comparison of transformed images with different query cell sizes.

Cell Size	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Training \downarrow
1	28.24	0.858	0.134	28min35s
2	28.35	0.861	0.134	20min36s
4	28.33	0.861	0.135	18min59s
8	28.31	0.861	0.134	18min51s
16	28.18	0.862	0.136	19min17s
32	27.76	0.858	0.140	22min23s

Table 5: Effect of different query cell sizes.

Different Cell Sizes We test the effect of different cell sizes on rendering quality and optimization speed (Tab. 5 and Fig. 9). The visual quality of the transformed images decreases as the cell size increases, because a smaller cell size allows for more fine-grained transformations. When the cell size is greater than 8, floaters start to appear in the scene. For cell sizes of 8 or smaller, the optimization time decreases as the cell size increases due to fewer queries in the appearance model. However, for cell sizes larger than 8, the optimization time increases instead. This is because when the cell size increases, the gradients of the appearance model’s parameters are back-propagated to more Gaussian primitives, which causes an increase in time that outweighs the time saved by the cell-based query.

Conclusions and Limitations

In this paper, we introduce DAVIGS, an efficient and plug-and-play decoupled appearance modeling method that effectively alleviates floaters and color distortions caused by image appearance variations in novel view synthesis. Our approach achieves state-of-the-art rendering quality with minimal optimization time and video memory overhead. Besides, since the introduced module can be discarded after optimization, it does not slow down rendering speeds.

However, when scaling up to larger scenes, our method may require larger multi-resolution hash grids and deeper MLPs. This can result in increased video memory usage, storage demand, and longer optimization time. This scaling problem can be solved by some divide-and-conquer strategies (Lin et al. 2024; Liu et al. 2024b), and we leave it for future work.

Acknowledgments

This work was partly supported by the National Natural Science Foundation of China (No. 62171251) and the Special Foundations for the Development of Strategic Emerging Industries of Shenzhen (No. KJZD20231023094700001).

References

- Abdal, R.; Yifan, W.; Shi, Z.; Xu, Y.; Po, R.; Kuang, Z.; Chen, Q.; Yeung, D.-Y.; and Wetzstein, G. 2024. Gaussian shell maps for efficient 3d human generation. In *CVPR*.
- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2022. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *CVPR*.
- Chen, X.; Zhang, Q.; Li, X.; Chen, Y.; Feng, Y.; Wang, X.; and Wang, J. 2022. Hallucinated neural radiance fields in the wild. In *CVPR*.
- Chen, Z.; Wang, F.; Wang, Y.; and Liu, H. 2024. Text-to-3d using gaussian splatting. In *CVPR*.
- Dahmani, H.; Bennehar, M.; Piasco, N.; Roldao, L.; and Tsishkou, D. 2024. SWAG: Splatting in the Wild images with Appearance-conditioned Gaussians. *arXiv preprint arXiv:2403.10427*.
- Dai, P.; Xu, J.; Xie, W.; Liu, X.; Wang, H.; and Xu, W. 2024. High-quality surface reconstruction using gaussian surfels. In *ACM SIGGRAPH*.
- Darmon, F.; Porzi, L.; Rota-Bulò, S.; and Kotschieder, P. 2024. Robust gaussian splatting. *arXiv preprint arXiv:2404.04211*.
- Huang, B.; Yu, Z.; Chen, A.; Geiger, A.; and Gao, S. 2024a. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH*.
- Huang, H.; Li, L.; Hui, C.; and Yeung, S.-K. 2024b. PhotoSLAM: Real-time Simultaneous Localization and Photorealistic Mapping for Monocular, Stereo, and RGB-D Cameras. In *CVPR*.
- Jiang, Y.; Shen, Z.; Wang, P.; Su, Z.; Hong, Y.; Zhang, Y.; Yu, J.; and Xu, L. 2024. Hifi4g: High-fidelity human performance rendering via compact gaussian splatting. In *CVPR*.
- Jin, Y.; Mishkin, D.; Mishchuk, A.; Matas, J.; Fua, P.; Yi, K. M.; and Trulls, E. 2021. Image matching across wide baselines: From paper to practice. *IJCV*.
- Keetha, N.; Karhade, J.; Jatavallabhula, K. M.; Yang, G.; Scherer, S.; Ramanan, D.; and Luiten, J. 2024. SplatAM: Splat, Track & Map 3D Gaussians for Dense RGB-D SLAM. In *CVPR*.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM ToG*.
- Kocabas, M.; Chang, J.-H. R.; Gabriel, J.; Tuzel, O.; and Ranjan, A. 2024. HUGS: Human Gaussian Splatting. In *CVPR*.
- Kulhanek, J.; Peng, S.; Kukulova, Z.; Pollefeys, M.; and Sattler, T. 2024. WildGaussians: 3D Gaussian Splatting in the Wild. *arXiv*.
- Lee, J. C.; Rho, D.; Sun, X.; Ko, J. H.; and Park, E. 2024. Compact 3D Gaussian Representation for Radiance Field. In *CVPR*.
- Li, Z.; Zheng, Z.; Wang, L.; and Liu, Y. 2024. Animatable Gaussians: Learning Pose-dependent Gaussian Maps for High-fidelity Human Avatar Modeling. In *CVPR*.
- Lin, J.; Li, Z.; Tang, X.; Liu, J.; Liu, S.; Liu, J.; Lu, Y.; Wu, X.; Xu, S.; Yan, Y.; and Yang, W. 2024. VastGaussian: Vast 3D Gaussians for Large Scene Reconstruction. In *CVPR*.
- Ling, H.; Kim, S. W.; Torralba, A.; Fidler, S.; and Kreis, K. 2024. Align Your Gaussians: Text-to-4D with Dynamic 3D Gaussians and Composed Diffusion Models. In *CVPR*.
- Liu, X.; Zhan, X.; Tang, J.; Shan, Y.; Zeng, G.; Lin, D.; Liu, X.; and Liu, Z. 2024a. Humangaussian: Text-driven 3d human generation with gaussian splatting. In *CVPR*.
- Liu, Y.; Guan, H.; Luo, C.; Fan, L.; Peng, J.; and Zhang, Z. 2024b. Citygaussian: Real-time high-quality large-scale scene rendering with gaussians.
- Lu, T.; Yu, M.; Xu, L.; Xiangli, Y.; Wang, L.; Lin, D.; and Dai, B. 2024. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *CVPR*.
- Martin-Brualla, R.; Radwan, N.; Sajjadi, M. S. M.; Barron, J. T.; Dosovitskiy, A.; and Duckworth, D. 2021. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *CVPR*.
- Matsuki, H.; Murai, R.; Kelly, P. H. J.; and Davison, A. J. 2024. Gaussian Splatting SLAM.
- Moreau, A.; Song, J.; Dharmo, H.; Shaw, R.; Zhou, Y.; and Pérez-Pellitero, E. 2024. Human gaussian splatting: Real-time rendering of animatable avatars. In *CVPR*.
- Morgenstern, W.; Barthel, F.; Hilsmann, A.; and Eisert, P. 2024. Compact 3D Scene Representation via Self-Organizing Gaussian Grids. In *ECCV*.
- Müller, T.; Evans, A.; Schied, C.; and Keller, A. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM ToG*.
- Niemeyer, M.; Manhardt, F.; Rakotosaona, M.-J.; Oechsle, M.; Duckworth, D.; Gosula, R.; Tateno, K.; Bates, J.; Kaeser, D.; and Tombari, F. 2024. RadSplat: Radiance Field-Informed Gaussian Splatting for Robust Real-Time Rendering with 900+ FPS. *arXiv.org*.
- Qian, S.; Kirschstein, T.; Schoneveld, L.; Davoli, D.; Giebenhain, S.; and Nießner, M. 2024. Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In *CVPR*.
- Schönberger, J. L.; and Frahm, J.-M. 2016. Structure-from-Motion Revisited. In *CVPR*.
- Tang, J.; Chen, Z.; Chen, X.; Wang, T.; Zeng, G.; and Liu, Z. 2024. LGM: Large Multi-View Gaussian Model for High-Resolution 3D Content Creation. In *ECCV*.
- Tang, J.; Ren, J.; Zhou, H.; Liu, Z.; and Zeng, G. 2023. DreamGaussian: Generative Gaussian Splatting for Efficient 3D Content Creation. *arXiv preprint arXiv:2309.16653*.
- Turki, H.; Ramanan, D.; and Satyanarayanan, M. 2022. Mega-NeRF: Scalable Construction of Large-Scale NeRFs for Virtual Fly-Throughs. In *CVPR*.

Xu, J.; Mei, Y.; and Patel, V. M. 2024. Wild-GS: Real-Time Novel View Synthesis from Unconstrained Photo Collections. *arXiv preprint arXiv:2406.10373*.

Yan, C.; Qu, D.; Xu, D.; Zhao, B.; Wang, Z.; Wang, D.; and Li, X. 2024a. Gs-slam: Dense visual slam with 3d gaussian splatting. In *CVPR*.

Yan, Y.; Lin, H.; Zhou, C.; Wang, W.; Sun, H.; Zhan, K.; Lang, X.; Zhou, X.; and Peng, S. 2024b. Street gaussians for modeling dynamic urban scenes. *arXiv preprint arXiv:2401.01339*.

Yang, Z.; Gao, X.; Zhou, W.; Jiao, S.; Zhang, Y.; and Jin, X. 2024. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *CVPR*.

Yi, T.; Fang, J.; Wang, J.; Wu, G.; Xie, L.; Zhang, X.; Liu, W.; Tian, Q.; and Wang, X. 2024. GaussianDreamer: Fast Generation from Text to 3D Gaussians by Bridging 2D and 3D Diffusion Models. In *CVPR*.

Yu, H.; Julin, J.; Milacski, Z. Á.; Niinuma, K.; and Jeni, L. A. 2024a. Cogs: Controllable gaussian splatting. In *CVPR*.

Yu, Z.; Chen, A.; Huang, B.; Sattler, T.; and Geiger, A. 2024b. Mip-Splatting: Alias-free 3D Gaussian Splatting. In *CVPR*.

Yu, Z.; Sattler, T.; and Geiger, A. 2024. Gaussian Opacity Fields: Efficient High-quality Compact Surface Reconstruction in Unbounded Scenes. *arXiv:2404.10772*.

Zhang, D.; Wang, C.; Wang, W.; Li, P.; Qin, M.; and Wang, H. 2024. Gaussian in the Wild: 3D Gaussian Splatting for Unconstrained Image Collections. *arXiv preprint arXiv:2403.15704*.

Zhao, C.; Sun, S.; Wang, R.; Guo, Y.; Wan, J.-J.; Huang, Z.; Huang, X.; Chen, Y. V.; and Ren, L. 2024. Tlcl-gs: Tightly coupled lidar-camera gaussian splatting for surrounding autonomous driving scenes. *arXiv preprint arXiv:2404.02410*.

Zhao, L.; Wang, P.; and Liu, P. 2024. BAD-Gaussians: Bundle Adjusted Deblur Gaussian Splatting. In *ECCV*.

Zheng, S.; Zhou, B.; Shao, R.; Liu, B.; Zhang, S.; Nie, L.; and Liu, Y. 2024. GPS-Gaussian: Generalizable Pixel-wise 3D Gaussian Splatting for Real-time Human Novel View Synthesis. In *CVPR*.

Zhou, H.; Shao, J.; Xu, L.; Bai, D.; Qiu, W.; Liu, B.; Wang, Y.; Geiger, A.; and Liao, Y. 2024a. HUGS: Holistic Urban 3D Scene Understanding via Gaussian Splatting. In *CVPR*.

Zhou, S.; Fan, Z.; Xu, D.; Chang, H.; Chari, P.; Bhadravaj, S.; Tejas You; Wang, Z.; and Kadambi, A. 2024b. DreamScene360: Unconstrained Text-to-3D Scene Generation with Panoramic Gaussian Splatting. In *ECCV*.

Zhou, X.; Lin, Z.; Shan, X.; Wang, Y.; Sun, D.; and Yang, M.-H. 2024c. Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. In *CVPR*.