

Micro-macro Wavelet-based Gaussian Splatting for 3D Reconstruction from Unconstrained Images

Yihui Li^{1,2}, Chengxin Lv^{1,2}, Hongyu Yang^{3,4*}, Di Huang^{1,2}

¹ State Key Laboratory of Complex and Critical Software Environment, Beijing, China

² School of Computer Science and Engineering, Beihang University, China

³ School of Artificial Intelligence, Beihang University, China

⁴ Shanghai Artificial Intelligence Laboratory, Shanghai, China
{kidleyh, 20231195, hongyuyang, dhuang}@buaa.edu.cn

Abstract

3D reconstruction from unconstrained image collections presents substantial challenges due to varying appearances and transient occlusions. In this paper, we introduce Micro-macro Wavelet-based Gaussian Splatting (MW-GS), a novel approach designed to enhance 3D reconstruction by disentangling scene representations into global, refined, and intrinsic components. The proposed method features two key innovations: **Micro-macro Projection**, which allows Gaussian points to capture details from feature maps across multiple scales with enhanced diversity; and **Wavelet-based Sampling**, which leverages frequency domain information to refine feature representations and significantly improve the modeling of scene appearances. Additionally, we incorporate a Hierarchical Residual Fusion Network to seamlessly integrate these features. Extensive experiments demonstrate that MW-GS delivers state-of-the-art rendering performance, surpassing existing methods.

Introduction

3D reconstruction from images is a longstanding and critical task in computer vision, with applications ranging from immersive virtual reality to 3D content creation. Recent advancements in this field have been driven by both implicit representations, such as Neural Radiance Fields (NeRF) (Mildenhall et al. 2021) and their subsequent developments (Barron et al. 2023; Tancik et al. 2022), and explicit representations, such as 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) and related techniques (Yu et al. 2024; Lu et al. 2024). These methods have demonstrated significant progress, particularly in reconstructing static scenes under controlled conditions, where stable lighting and minimal occlusions prevail. However, real-world applications often involve images captured in unconstrained and dynamic environments, where traditional methods struggle to maintain consistent reconstruction quality, leading to issues such as blurriness, artifacts, and overall performance degradation (Yang et al. 2023).

To address the challenges of dynamic appearance variations in real-world scenes, NeRF-W (Martin-Brualla et al. 2021) introduced per-image appearance embeddings, later

refined by (Chen et al. 2022b; Yang et al. 2023) to better handle these variations across different views. Despite these advancements, global embeddings often fall short in representing fine-grained details, as they fail to fully account for the significant appearance variations across different scene locations influenced by object properties and environmental factors. Furthermore, the high computational costs associated with implicit representations hinder their applicability in real-time rendering scenarios.

Some Gaussian-related methods (Dahmani et al. 2024; Kulhanek et al. 2024) have attempted to improve dynamic appearance modeling by combining intrinsic Gaussian features with global appearance embeddings. However, these approaches still face limitations similar to those of NeRF-based methods, primarily due to the less expressive nature of global embeddings. The state-of-the-art GS-W method (Zhang et al. 2024) has made notable progress by allowing Gaussian points to adaptively sample detailed dynamic appearance information from 2D feature maps. This technique provides more flexibility and captures richer details, yet challenges like blurriness persist, particularly when images are closely inspected.

In this paper, we introduce Micro-macro Wavelet-based Gaussian Splatting (MW-GS), a novel approach designed to address the limitations of existing 3D reconstruction techniques. Our method decomposes Gaussian features into three distinct components: global appearance, refined appearance, and intrinsic features, which together provide a comprehensive representation of dynamically changing scenes. Global features capture the overall scene characteristics, such as color tone and lighting conditions, while refined features model detailed textures and region-specific effects like highlights and shadows. Intrinsic features represent consistent properties of the scene, ensuring robustness across varying conditions.

The key innovation of our work lies in **Micro-macro Projection**, which significantly enhances refined appearance modeling—a critical aspect where previous methods often fall short. Specifically, our method employs adaptive sampling within narrow and broad conical frustums on the 2D feature map, optimizing the 3D Gaussian points to capture both fine-grained textures and relatively long-range features that influence overall scene fidelity, such as highlights. Inspired by traditional MipMap operations, we further introduce a

*Corresponding author

simple yet effective jitter to the projected position for each Gaussian point at the **micro scale**, rather than relying on a fixed position as in previous methods. This jitter introduces variability into the sampling process, enabling capturing a richer set of features. Additionally, we incorporate **Wavelet-based Sampling**, leveraging frequency domain information to further enhance refined appearance modeling and reconstruction accuracy. This multi-scale approach ensures that each Gaussian point accurately captures fine-grained features while maintaining feature diversity. Finally, we propose a Hierarchical Residual Fusion Network (HRFN) to effectively integrate features from different levels, ensuring a cohesive and precise 3D reconstruction.

In summary, the key contributions of this paper are:

- **Micro-macro Wavelet-based Gaussian Splatting (MW-GS):** A novel approach for 3D reconstruction that effectively decomposes and models global, refined, and intrinsic scene features.
- **Micro-macro projection:** An innovative sampling strategy that introduces jitter and adaptive conical frustums to capture a rich set of features, significantly enhancing refined appearance modeling.
- **Wavelet-based Sampling:** This component enhances multi-resolution sampling, leveraging frequency domain information to refine feature representation and improve reconstruction accuracy.

Our experiments across multiple scenes demonstrate that MW-GS achieves superior reconstruction results and rendering quality, surpassing existing methods.

Related Work

Scene Representations

Existing traditional methods include meshes (Groueix et al. 2018; Kanazawa et al. 2018; Kato, Ushiku, and Harada 2018; Wen et al. 2019), point clouds (Qi et al. 2017a,b; Shi et al. 2020), and voxels (Schwarz et al. 2022; Wu et al. 2015). Recently, Neural Radiance Fields (NeRF) (Mildenhall et al. 2021) have revolutionized the synthesis of novel, photo-realistic views from images. Extensions to NeRF enhance visual quality (Barron et al. 2022; Hu et al. 2023), rendering speed (Chen et al. 2023; Reiser et al. 2023), and convergence (Müller et al. 2022; Chen et al. 2022a), though limitations persist in speed and detail. More recently, 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023), an explicit representation method, offers real-time rendering with high-resolution quality. Recent advances in 3DGS include improvements in efficiency (Lee et al. 2024), surface reconstruction (Huang et al. 2024), and incorporating semantic attributes for multimodal applications (Xie et al. 2024; Shi et al. 2024; Qin et al. 2024). 3DGS has also been extended to various tasks, including autonomous driving (Zhou et al. 2024c; Yan et al. 2024; Zhou et al. 2024a), 3D generation (Chen et al. 2024b; Chung et al. 2023), and controllable 3D scene editing (Chen et al. 2024a; Wang et al. 2024; Zhou et al. 2024b).

Novel View Synthesis from Unconstrained Photo Collections

Traditional novel view synthesis methods assume static geometry, materials, and lighting conditions. However, unconstrained photo collections, such as those from the internet (Snavely, Seitz, and Szeliski 2006), often feature variable lighting and dynamic occlusions, which challenge these assumptions. NeRF-W (Martin-Brualla et al. 2021) pioneered addressing these challenges by incorporating learnable appearance and transient embeddings for each image. Subsequent NeRF-based methods, such as Ha-NeRF (Chen et al. 2022b) and CR-NeRF (Yang et al. 2023), extended this approach with CNN-based appearance encoders to handle dynamic variations more effectively.

Despite these advancements, methods based on implicit representations suffer from slow rendering speeds, prompting research into the use of 3D Gaussian Splatting (3DGS) as a more efficient alternative to NeRF. Approaches such as SWAG (Dahmani et al. 2024) and WildGaussians (Kulhanek et al. 2024) incorporate learnable appearance embeddings to modulate the color of 3D Gaussians using a multi-layer perceptron (MLP). WE-GS (Wang, Wang, and Qi 2024) enhances CNN-based appearance representations through a spatial attention module, while GS-W (Zhang et al. 2024) and Wild-GS (Xu, Mei, and Patel 2024) leverage CNNs to generate feature maps for dynamic appearance modulation. GS-W projects each Gaussian onto feature maps and performs adaptive sampling, whereas Wild-GS uses depth information to construct triplane maps for local appearance embedding.

However, projecting 2D appearance reference images to 3D space can lead to sparse features and gaps. In this work, we address these limitations by introducing Micro-macro Wavelet-based Sampling, which enhances sampling accuracy and diversity. Our approach leverages frequency domain information to refine appearance features, representing the first integration of frequency domain data into 3DGS appearance representation. This innovation effectively improves both the appearance representation and reconstruction results for unstructured photo collections.

Preliminaries

3D Gaussian Splatting

3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) represents a scene using a collection of anisotropic 3D Gaussians. These Gaussians are differentiable and can be efficiently projected onto 2D splats, enabling rapid α -blending through tile-based rasterization. Each 3DGS is characterized by a complete 3D covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$, which is defined in world space and centered at a point (mean) $\mu \in \mathbb{R}^3$:

$$G(x) = \exp\left(-\frac{1}{2}(x - \mu)^\top \Sigma^{-1}(x - \mu)\right), \quad (1)$$

where x is an arbitrary position within the 3D scene. To maintain its positive semi-definiteness during optimization, Σ is formulated using a diagonal scaling matrix S and an orthogonal rotation matrix R :

$$\Sigma = RSS^\top R^\top. \quad (2)$$




Figure 1: Overview of Micro-macro Wavelet-based Gaussian Splatting (MW-GS). Beginning with a reference image, CNN-based sub-networks extract global appearance features, a visibility map, and a set of feature maps which then undergo wavelet transform to capture multi-resolution information for refined appearance modeling. 3D Gaussian points are projected onto these feature maps using Micro-macro Projection, which employs adaptive sampling within narrow and broad conical frustums to capture richer information. The global, refined, and intrinsic features are subsequently fused through a Hierarchical Residual Fusion Network (HRFN) to generate the final color, which is rendered using a tile rasterizer for high-quality 3D reconstruction.

In terms of implementation, the covariance matrix Σ is parameterized using a unit quaternion q and a 3D scaling vector s . Additionally, each Gaussian is associated with color \hat{c} and an opacity factor α , which is multiplied by $G(x)$ during the blending process. During rendering, these Gaussians are directly projected to the screen for high-speed rendering, called Splatting (Zwicker et al. 2001). After sorting the Gaussians by depth, alpha compositing is used to compute the final color \hat{C} for pixel \mathbf{p} . This process is expressed as:

$$\hat{C}(\mathbf{p}) = \sum_{i \in N} \hat{c}_i \alpha'_i \prod_{j=1}^{i-1} (1 - \alpha'_j), \quad (3)$$

where α'_i is the product of α_i and the splatted 2D Gaussian.

Discrete Wavelet Transform

Wavelet (Daubechies 1992; Strang and Nguyen 1996), has long been a widely used mathematical tool in image analysis (Chen et al. 2021; Duan et al. 2017; Li et al. 2022). The coefficients of the Discrete Wavelet Transform (DWT) describe signals across different frequency bands at various resolution levels, effectively capturing both local and global information. The 2D DWT decomposes an image into four different components in the frequency domain with low-pass (emphasizing relatively smooth parts) and high-pass (capturing high-frequency signals, including local texture information) filters, denoted by \mathbf{L} and \mathbf{H} , respectively. The combination of the two filters yields four distinct kernels, *i.e.*, \mathbf{LL} , \mathbf{LH} , \mathbf{HL} , and \mathbf{HH} , corresponding to different frequency and spatial information. Given a feature map $\mathbf{F} \in \mathbb{R}^{H \times W}$, where H and W indicate its height and width, respectively, we apply DWT to \mathbf{F} . The four sub-band features can be obtained by a one-level decomposition, formulated as:

$$\begin{aligned} \mathbf{F}_w^{LL} &= \mathbf{LFL}^\top, \mathbf{F}_w^{LH} = \mathbf{HFL}^\top, \\ \mathbf{F}_w^{HL} &= \mathbf{LFH}^\top, \mathbf{F}_w^{HH} = \mathbf{HFH}^\top. \end{aligned} \quad (4)$$

For a feature map with multi channels, the wavelet transform is applied individually to each channel. The resulting sub-bands from the same filter are then concatenated, producing four corresponding frequency sub-bands.

Method

To achieve high-fidelity 3D reconstruction from unconstrained image collections, we introduce a novel approach named Micro-macro Wavelet-based Gaussian Splatting that decouples appearance attributes into global, refined, and intrinsic features, achieving a more structured and explicit scene representation. Our method consists of the following key components: *i)* Global appearance features are extracted directly from a 2D reference image using CNN-based networks. *ii)* We propose a Micro-Macro Wavelet-based sampling technique to enhance the accuracy and diversity of refined appearance features by capturing multi-resolution information. *iii)* A Hierarchical Residual Fusion Network (HRFN) effectively integrates learnable intrinsic feature embeddings with global and refined features. *iv)* Visibility maps are employed to mitigate the influence of transient objects during the optimization process. This approach enables high-quality 3D reconstruction with greater attention to refined detail and overall appearance fidelity.

Structured and Explicit Appearance Decoupling

In unconstrained photo collections, global appearance variations arise from factors like lighting conditions during capture and post-processing techniques such as gamma correction, exposure adjustment, and tone mapping. Additionally, scene points may exhibit varying color details and the changing appearance due to directional lighting, such as distinct highlights and shadows. Intrinsic properties, such as material characteristics of the scene, remain constant. To effectively model these variations, we explicitly decouple the appearance into three distinct components: **Global Appearance Feature** ($f_g \in \mathbb{R}^{n_g}$): Captures information about the overall scene. **Refined Appearance Feature** ($f_r \in \mathbb{R}^{n_r}$): Contains details beyond global features, specific to regional positions within the scene, such as high-frequency details and local highlights and shadows. **Intrinsic Feature** ($f_v \in \mathbb{R}^{n_v}$): Represents the inherent properties of the scene points.

For a point v located at \mathbf{x}_i in 3D space, its dynamic appearance is characterized by these three components. The global and refined appearance features (f_g and f_r) are extracted from a reference image, while the intrinsic feature f_v is learned during the training process. This structured and explicit appearance decoupling effectively balances global




Figure 2: Sampling comparison between GS-W and the proposed MW-GS. Our approach integrates both narrow and broad conical frustums with wavelet-based sampling, allowing for a more comprehensive capture of features and resulting in enhanced accuracy.

scene information with local contextual details, while preserving intrinsic scene properties.

In our implementation, we utilize a voxel-based organization of Gaussians as described in Scaffold-GS (Lu et al. 2024). Each anchor point v at the center of a voxel is associated with a scaling factor $l_v \in \mathbb{R}^3$ and k learnable offsets $O_v \in \mathbb{R}^{k \times 3}$, which collectively manage the k Gaussians within the voxel. To model global appearance changes consistently across the scene, we assign a uniform global appearance feature f_g , extracted from a reference image, to all anchors. This extraction is performed by passing the feature map obtained from the UNet encoder through a global average pooling layer, followed by a trainable MLP MLP^G to obtain f_g .

Micro-macro Wavelet-based Sampling

To enhance the accuracy of scene representation, we introduce a novel approach called Micro-Macro Wavelet-Based Sampling (MWS). This method refines the appearance features of each 3D Gaussian by capturing more detailed and diverse information. This approach aligns more closely with real-world scene variations by incorporating both fine-grained and broad-scale features. The MWS strategy consists of two main components:

Micro-macro Projection (MP). In traditional MipMap techniques (Williams 1983), jitter sampling (Cook 1986) is used to introduce random perturbations, enhancing the representation of texture details. Building on this concept, we propose an adaptive jitter projection method to achieve micro-projections. Rather than projecting each 3D point directly onto a fixed position on the 2D feature map, we project each point within a narrow conical frustum. This approach allows each Gaussian along a ray to capture similar yet distinct features, effectively reflecting the unique properties of each point in 3D space.

Fig. 2a presents a comparison between our method and GS-W. The GS-W method starts by directly projecting onto the *projection feature map*, resulting in identical local appearance features for 3D points along the same ray. While GS-W attempts to mitigate this by adaptively sampling across multiple feature maps with varying positions, it lacks explicit control over targeting specific local regions. Consequently, the informative features within these regions are not fully utilized. In contrast, our micro projection method

employs a narrow conical frustum to project 3D points onto the 2D feature map, with the cross-sectional radius parameterized by \hat{r} . To further refine this process, we assign k_s learnable coordinate offsets $\{nc_i\}_{k_s}$ to each 3D point, enabling adaptive sampling within the frustum. The features obtained from these k_s samples are then averaged to produce the refined feature f_r^n . This approach ensures that feature sampling is diverse but still consistent, capturing rich fine-grained details while maintaining the coherence of rendered textures. By balancing variability and accuracy, our method significantly enhances representation performance.

Additionally, refined appearance features extend beyond fine details to encompass broader, long-range characteristics such as highlights. To capture these broader aspects, we also implement sampling through a broader conical frustum, as depicted on the right side of Fig. 2b. Inspired by the principle that the size of point projection is inversely proportional to its distance from the camera, we parameterize the projection radius of the broad conical frustum as $\hat{R} = \hat{R}_{max} / \|\mathbf{x}_i - \mathbf{x}_c\|_2$, where \mathbf{x}_c represents the camera center. To avoid excessively small radius, we impose a minimum threshold \hat{R}_{min} . By assigning k_s learnable coordinate scaling factors $\{bc_i\}_{k_s}$ to each 3D point, we facilitate adaptive sampling within the broad conical frustum using the implementation $bc_i \times \hat{p}_i$, where \hat{p}_i represents the projection center of the frustum. This design encourages a broader sampling range and enhances the representation of long-range features. The features derived from these k_s samples are averaged to generate the broad appearance feature f_r^b .

Wavelet-based Sampling (WS). In unconstrained image collections, camera parameters can vary significantly, making it challenging to handle large-scale variations across different viewpoints with fixed-resolution sampling alone. To address this, we introduce a Wavelet-based Sampling technique that enhances the capture of high-frequency and multi-scale information. Specifically, we apply the Discrete Wavelet Transform (DWT) to the feature map \mathbf{F}^{MAP} generated by the UNet, producing a series of feature maps. The DWT decomposes the feature map into four frequency bands while simultaneously reducing its resolution, effectively preserving the spatial information and enabling multi-scale sampling that captures diverse frequency information.

The process begins by segmenting the feature map \mathbf{F}^{MAP} into $2M + 2$ smaller feature maps $\{\mathbf{F}^1, \dots, \mathbf{F}^{2M+2}\}$, where each $\mathbf{F}^i \in \mathbb{R}^{\frac{n_r}{2^{M+2}} \times H^F \times W^F}$. Here, M represents the maximum number of downsampling operations or the highest level of Discrete Wavelet Transform (DWT), which serves as a critical hyper-parameter. The dimensions H^F and W^F correspond to the height and width of the smaller feature maps, respectively.

During the m -th downsampling stage, an m -level DWT is applied to the $(2m + 1)$ -th and $(2m + 2)$ -th feature maps, resulting in 4^m sub-feature maps as per Eq. (4). These sub-feature maps are then subjected to bilinear interpolation sampling through narrow and broad conical frustums, employing the Micro-macro Projection technique, which yields the feature sets $\{f_{r,m,j}^n\}_{4^m}$ and $\{f_{r,m,j}^b\}_{4^m}$, respectively.

Next, the refined features $f_{r,m}^n$ and $f_{r,m}^b$ for each down-

Method	Brandenburg Gate			Sacre Coeur			Trevi Fountain		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
NeRF-W	24.17	0.890	0.167	19.20	0.807	0.191	18.97	0.698	0.265
Ha-NeRF	24.04	0.887	0.139	20.02	0.801	0.171	20.18	0.690	0.223
CR-NeRF	26.53	0.900	0.106	22.07	0.823	0.152	21.48	0.712	0.207
WildGaussians	27.77	0.927	0.133	22.56	0.859	0.177	23.63	0.766	0.228
GS-W	<u>27.96</u>	<u>0.931</u>	<u>0.086</u>	<u>23.24</u>	<u>0.863</u>	<u>0.130</u>	<u>22.91</u>	<u>0.801</u>	<u>0.156</u>
Ours	29.37	0.942	0.052	24.64	0.897	0.073	24.07	0.821	0.120

Table 1: Quantitative results on three datasets. **Bold** and underlined values correspond to the best and the second-best value, respectively. Our method outperforms the previous methods across all datasets on PSNR, SSIM, and LPIPS.

sampling level are computed by applying learnable weight parameters to the sampled features:

$$f_{r,m}^n = \sum_{j=1}^{4^m} \omega_{m,j}^n \cdot f_{r,m,j}^n, \quad f_{r,m}^b = \sum_{j=1}^{4^m} \omega_{m,j}^b \cdot f_{r,m,j}^b, \quad (5)$$

where $\omega_{m,j}^n$ and $\omega_{m,j}^b$ are the learnable weights for the $(2m+1)$ -th and $(2m+2)$ -th feature maps, respectively. Finally, the refined appearance features corresponding to each anchor are obtained by concatenating the features from all scales:

$$f_r = f_{r,0}^n \oplus f_{r,0}^b \oplus \dots \oplus f_{r,M}^n \oplus f_{r,M}^b. \quad (6)$$

The combination of Micro-macro Projection and Wavelet-based Sampling supplements scene representation with multi-scale and high-frequency features. It accurately captures refined appearance variations, providing a comprehensive understanding of scene structures at multiple scales.

Hierarchical Residual Fusion Network

After obtaining the global appearance, refined appearance, and intrinsic features, it is essential to effectively combine these with the position and view direction to generate the final k Gaussian colors corresponding to each anchor. Given that these features exist in different high-dimensional spaces, simple concatenation would not achieve an effective fusion. To address this, we propose a Hierarchical Residual Fusion Network (HRFN), which consists of four Multi-Layer Perceptrons (MLPs), denoted as $\mathcal{M}^H = \{\mathcal{M}_1^H, \mathcal{M}_2^H, \mathcal{M}_3^H, \mathcal{M}_4^H\}$.

The inputs to the network include the position of the anchor center \mathbf{x}_i , global appearance feature f_g , refined appearance feature f_r , intrinsic features f_v , and the direction vector $\vec{\mathbf{d}}_{ic} = \frac{\mathbf{x}_i - \mathbf{x}_c}{\|\mathbf{x}_i - \mathbf{x}_c\|_2}$. The network processes these inputs to infer the output colors $\{\hat{c}_k\}$ for the k Gaussians as follows:

$$\begin{aligned} Emb &= \mathcal{M}_1^H(\gamma(\mathbf{x}_i) \oplus f_v \oplus f_r \oplus f_g) \oplus \omega_r f_r, \\ \{\hat{c}_k\} &= \mathcal{M}_4^H \left(\mathcal{M}_3^H \left(\mathcal{M}_2^H (Emb \oplus \omega_v f_v) \right) \oplus \vec{\mathbf{d}}_{ic} \right), \end{aligned} \quad (7)$$

where \oplus denotes concatenation, $\gamma(\cdot)$ is the positional encoding function, and ω_r and ω_v are learnable adaptive weights. This hierarchical structure enables the network to effectively fuse features from different levels, ensuring a comprehensive integration of global, refined, and intrinsic information.

This design improves the network’s ability to accurately predict the final Gaussian colors by capturing the complex interactions among the various input features.

Handling Transient Objects and Optimization

Transient objects with view inconsistencies are widespread in unconstrained photo collections, posing significant challenges for 3DGS, which aim to fit each view accurately. To address this issue, we reuse the UNet encoder that extracts appearance features, and adding a separate decoder to predict the visibility map vm . Once the vm is obtained, it is applied to the encoded intermediate features of the appearance feature extraction network. Additionally, a visibility map regularization is applied to prevent meaningful pixels from being masked out. The overall loss function between the rendered image I_r and the reference image I_{gt} , incorporating the visibility map, can be expressed as follows:

$$\begin{aligned} \mathcal{L} &= \lambda_{SSIM} \mathcal{L}_{SSIM}(vm \odot I_r, vm \odot I_{gt}) \\ &+ \lambda_1 \mathcal{L}_1(vm \odot I_r, vm \odot I_{gt}) + \lambda_{vm} \mathcal{L}_2(vm, 1), \end{aligned} \quad (8)$$

where \odot denotes the Hadamard product.

Experiment

Datasets and Metrics

Following previous works (Chen et al. 2022b; Zhang et al. 2024), we evaluate different methods on three datasets: *Brandenburg Gate*, *Sacre Coeur*, and *Trevi Fountain*, with all images downsampled by a factor of 2 during both training and evaluation. We use PSNR, SSIM (Wang et al. 2004), and LPIPS (Zhang et al. 2018) as metrics to assess the performance of our method. We compare it against NeRF-W (Martin-Brualla et al. 2021), Ha-NeRF (Chen et al. 2022b), CR-NeRF (Yang et al. 2023), WildGaussians (Kulhanek et al. 2024) and GS-W (Zhang et al. 2024).

Comparison Experiments

Quantitative comparison. The quantitative results presented in Tab. 1 highlight the efficacy of our MW-GS method. NeRF-W introduces global appearance and transient embeddings, providing a basic level of scene reconstruction. Ha-NeRF and CR-NeRF build on this foundation, achieving improvements, but they still fall short in adequately capturing the local contextual information of various scene points. Similarly, WildGaussians exhibit limitations in

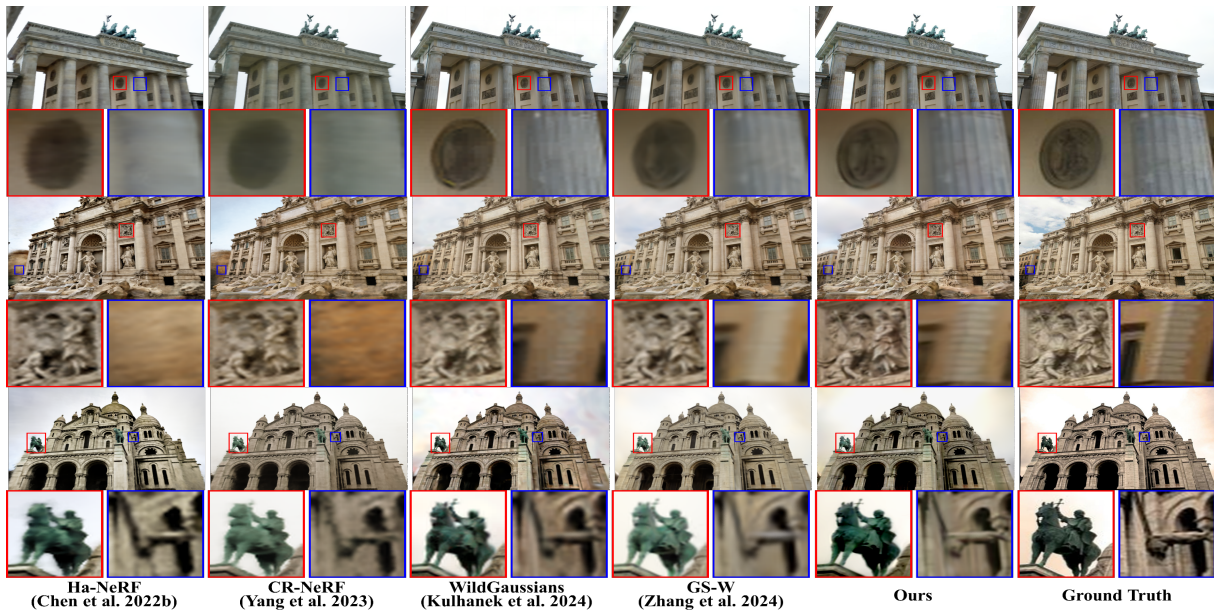


Figure 3: Qualitative comparison on three datasets. Red and blue crosses emphasize that MW-GS can recover finer details.

Method	Brandenburg Gate			Sacre Coeur			Trevi Fountain		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o WS	28.88	0.939	0.055	24.35	0.895	0.075	23.38	0.810	0.127
only micro	28.49	0.939	0.059	23.83	0.888	0.081	23.47	0.809	0.129
w/o HRFN	28.69	0.936	0.057	24.45	0.896	0.074	23.68	0.813	0.125
w/o vm	28.66	0.934	0.057	24.83	<u>0.893</u>	0.074	<u>23.65</u>	<u>0.815</u>	<u>0.122</u>
Full model	29.37	0.942	0.052	<u>24.64</u>	0.897	0.073	24.07	0.821	0.120

Table 2: Ablation studies on three datasets. **Bold** and underlined values correspond to the best and the second-best value.

appearance modeling due to their dependence on global appearance embeddings. GS-W addresses this by employing adaptive sampling of local features, leading to enhanced focus on detailed information, as reflected in consistently improved metrics. Our MW-GS method advances beyond GS-W by seamlessly integrating long-range contextual information with detailed localized features within the Gaussian representations. Furthermore, it enriches multi-scale information and captures high-frequency details more effectively. This approach yields significant improvements, outperforming GS-W by 1.41 dB, 1.4 dB, and 1.16 dB in PSNR across the three evaluated scenes.

Qualitative comparison. The qualitative results in Fig. 3 clearly demonstrate the strengths of our approach. Existing methods often struggle to capture intricate scene details and complex textures accurately. In contrast, our method excels by utilizing micro-macro wavelet-based sampling to enhance feature extraction, effectively integrating frequency domain and multi-scale information. Furthermore, the hierarchical fusion of structured features allows for a more precise recovery of appearance details and clear structural representation. This results in superior visual fidelity and accuracy in the reconstructed scenes. For example, our method

successfully captures finer details and more accurate colors in the reliefs of the Trevi Fountain and the bronze statues at Sacre Coeur, outperforming current techniques.

Component Analysis

Ablation Study. The ablation study results achieved on three datasets are summarized in Tab. 2 and qualitative outcomes are illustrated in Fig. 5. Key findings include: 1) MP significantly enhances the diversity of refined appearance sampling. As demonstrated in Fig. 5, relying solely on micro projection, without full MP and WS, results in noticeable blurring of distant objects. 2) WS further refines attention to high-frequency and multi-scale information. Without WS, there is a marked loss of detail, as evidenced by the blurring of the Trevi Fountain sculptures and a corresponding 0.69 dB decrease in PSNR. 3) HRFN facilitates a comprehensive fusion of diverse information compared to simple concatenation. This results in a 0.68 dB increase in PSNR on the *Brandenburg Gate*. 4) The removal of the visibility map introduces significant artifacts in the rendered images, as seen in Fig. 5, underscoring its importance in maintaining visual integrity.

Analysis of Appearance Transfer. Our method demon-



(a) Appearance transfer using reference images inside and outside the dataset. (b) Appearance tuning by global and local features.

Figure 4: Qualitative results of appearance transfer.





Figure 5: Ablation studies by visualization.

strates strong support for appearance transfer in 3D scenes, showing its advanced appearance modeling capabilities. As illustrated in Fig. 4a, a qualitative comparison between GS-W and our approach reveals that our method not only transfers both foreground and background elements to novel views but also preserves the intricate details of the scene, rather than merely replicating the overall scene tone. This capability underscores the accuracy and robustness of our appearance modeling. Furthermore, our approach allows for the adjustment of dynamic appearance features during rendering by modulating their corresponding weights. As depicted in Fig. 4b, altering these weights can effectively adjust the overall exposure of the rendered image. Additionally, the variation in intrinsic features across different regions results in corresponding changes in local highlights, emphasizing the importance and effectiveness of disentangling and capturing intrinsic features.

Analysis of Sampling. We project the sampling positions onto corresponding camera images to generate an attention map, where areas with dense sampling points indicate high attention. As shown in Fig. 6a, our sampling approach effectively captures fine-grained local details and long-range information by utilizing narrow and broad conical frustum projections, respectively. We further visualize the features of interest by examining the refined narrow features f_r^n and



(a) Sampling attention. (b) Sampling features rendering result.

Figure 6: Visualization of sampling analysis.

broad features f_r^b at different resolutions, as depicted in Fig. 6b. The f_r^n features, focused on high-resolution details, effectively capture local texture intricacies, while the 0.5x resolution features, processed through DWT, attend to varying details across different frequency bands. In contrast, the f_r^b features primarily target low-texture regions such as water surfaces or highlights, which correspond to long-range features. The combination of both enables our refined appearance feature to effectively model dynamic appearances.

Conclusion

In this paper, we present MW-GS, a novel approach for handling varying appearances in unconstrained photo collections for novel view synthesis. By decomposing appearance features into global, refined, and intrinsic components and employing Micro-macro Wavelet-based Sampling, our method captures precise and diverse appearance features while preserving multi-scale and frequency domain details. Extensive experiments show that MW-GS outperforms existing methods, setting a new standard in rendering quality and adaptability to dynamic scenes.

Acknowledgments

This work is partly supported by the National Key R&D Program of China (2022ZD0161902), the National Natural Science Foundation of China (No. 62202031, 62176012), the Fundamental Research Funds for the Central Universities, the Research Program of State Key Laboratory of Critical Software Environment, and the Fundamental Research Funds for the Central Universities.

References

- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2022. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5470–5479.
- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2023. Zip-nerf: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19697–19705.
- Chen, A.; Xu, Z.; Geiger, A.; Yu, J.; and Su, H. 2022a. Tensorf: Tensorial radiance fields. In *European conference on computer vision*, 333–350. Springer.
- Chen, S.; Yao, T.; Chen, Y.; Ding, S.; Li, J.; and Ji, R. 2021. Local relation learning for face forgery detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 1081–1088.
- Chen, X.; Zhang, Q.; Li, X.; Chen, Y.; Feng, Y.; Wang, X.; and Wang, J. 2022b. Hallucinated neural radiance fields in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12943–12952.
- Chen, Y.; Chen, Z.; Zhang, C.; Wang, F.; Yang, X.; Wang, Y.; Cai, Z.; Yang, L.; Liu, H.; and Lin, G. 2024a. Gaussianeditor: Swift and controllable 3d editing with gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21476–21485.
- Chen, Z.; Funkhouser, T.; Hedman, P.; and Tagliasacchi, A. 2023. Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16569–16578.
- Chen, Z.; Wang, F.; Wang, Y.; and Liu, H. 2024b. Text-to-3d using gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21401–21412.
- Chung, J.; Lee, S.; Nam, H.; Lee, J.; and Lee, K. M. 2023. Luciddreamer: Domain-free generation of 3d gaussian splatting scenes. *arXiv preprint arXiv:2311.13384*.
- Cook, R. L. 1986. Stochastic sampling in computer graphics. *ACM Transactions on Graphics (TOG)*, 5(1): 51–72.
- 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*.
- Daubechies, I. 1992. *Ten lectures on wavelets*. Houthalen-Helchteren: SIAM.
- Duan, Y.; Liu, F.; Jiao, L.; Zhao, P.; and Zhang, L. 2017. SAR image segmentation based on convolutional-wavelet neural network and Markov random field. *Pattern Recognition*, 64: 255–267.
- Groueix, T.; Fisher, M.; Kim, V. G.; Russell, B. C.; and Aubry, M. 2018. A papier-mâché approach to learning 3d surface generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 216–224.
- Hu, W.; Wang, Y.; Ma, L.; Yang, B.; Gao, L.; Liu, X.; and Ma, Y. 2023. Tri-miprf: Tri-mip representation for efficient anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19774–19783.
- Huang, B.; Yu, Z.; Chen, A.; Geiger, A.; and Gao, S. 2024. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, 1–11.
- Kanazawa, A.; Tulsiani, S.; Efros, A. A.; and Malik, J. 2018. Learning category-specific mesh reconstruction from image collections. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 371–386.
- Kato, H.; Ushiku, Y.; and Harada, T. 2018. Neural 3d mesh renderer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3907–3916.
- Kerbl, B.; Kopanas, G.; Leimkuehler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics (TOG)*, 42(4): 1–14.
- Kulhanek, J.; Peng, S.; Kukulova, Z.; Pollefeys, M.; and Sattler, T. 2024. WildGaussians: 3D Gaussian Splatting in the Wild. *arXiv preprint arXiv:2407.08447*.
- Lee, J. C.; Rho, D.; Sun, X.; Ko, J. H.; and Park, E. 2024. Compact 3d gaussian representation for radiance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21719–21728.
- Li, J.; Xie, H.; Yu, L.; and Zhang, Y. 2022. Wavelet-enhanced Weakly Supervised Local Feature Learning for Face Forgery Detection. In *Proceedings of the 30th ACM International Conference on Multimedia*, 1299–1308.
- 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20654–20664.
- Martin-Brualla, R.; Radwan, N.; Sajjadi, M. S.; Barron, J. T.; Dosovitskiy, A.; and Duckworth, D. 2021. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7210–7219.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Müller, T.; Evans, A.; Schied, C.; and Keller, A. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4): 1–15.

- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30.
- Qin, M.; Li, W.; Zhou, J.; Wang, H.; and Pfister, H. 2024. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20051–20060.
- Reiser, C.; Szeliski, R.; Verbin, D.; Srinivasan, P.; Mildenhall, B.; Geiger, A.; Barron, J.; and Hedman, P. 2023. Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *ACM Transactions on Graphics (TOG)*, 42(4): 1–12.
- Schwarz, K.; Sauer, A.; Niemeyer, M.; Liao, Y.; and Geiger, A. 2022. Voxgraf: Fast 3d-aware image synthesis with sparse voxel grids. *Advances in Neural Information Processing Systems*, 35: 33999–34011.
- Shi, J.-C.; Wang, M.; Duan, H.-B.; and Guan, S.-H. 2024. Language embedded 3d gaussians for open-vocabulary scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5333–5343.
- Shi, S.; Guo, C.; Jiang, L.; Wang, Z.; Shi, J.; Wang, X.; and Li, H. 2020. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10529–10538.
- Snavely, N.; Seitz, S. M.; and Szeliski, R. 2006. Photo tourism: exploring photo collections in 3D. In *ACM siggraph 2006 papers*, 835–846.
- Strang, G.; and Nguyen, T., eds. 1996. *Wavelets and filter banks*. Wellesley: SIAM.
- Tancik, M.; Casser, V.; Yan, X.; Pradhan, S.; Mildenhall, B.; Srinivasan, P. P.; Barron, J. T.; and Kretzschmar, H. 2022. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8248–8258.
- Wang, J.; Fang, J.; Zhang, X.; Xie, L.; and Tian, Q. 2024. Gaussianeditor: Editing 3d gaussians delicately with text instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20902–20911.
- Wang, Y.; Wang, J.; and Qi, Y. 2024. WE-GS: An In-the-wild Efficient 3D Gaussian Representation for Unconstrained Photo Collections. *arXiv preprint arXiv:2406.02407*.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wen, C.; Zhang, Y.; Li, Z.; and Fu, Y. 2019. Pixel2mesh++: Multi-view 3d mesh generation via deformation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1042–1051.
- Williams, L. 1983. Pyramidal parametrics. In *Proceedings of the 10th annual conference on Computer graphics and interactive techniques*, 1–11.
- Wu, Z.; Song, S.; Khosla, A.; Yu, F.; Zhang, L.; Tang, X.; and Xiao, J. 2015. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1912–1920.
- Xie, T.; Zong, Z.; Qiu, Y.; Li, X.; Feng, Y.; Yang, Y.; and Jiang, C. 2024. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4389–4398.
- 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, Y.; Lin, H.; Zhou, C.; Wang, W.; Sun, H.; Zhan, K.; Lang, X.; Zhou, X.; and Peng, S. 2024. Street gaussians for modeling dynamic urban scenes. *arXiv preprint arXiv:2401.01339*.
- Yang, Y.; Zhang, S.; Huang, Z.; Zhang, Y.; and Tan, M. 2023. Cross-ray neural radiance fields for novel-view synthesis from unconstrained image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15901–15911.
- Yu, Z.; Chen, A.; Huang, B.; Sattler, T.; and Geiger, A. 2024. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19447–19456.
- 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*.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21336–21345.
- Zhou, S.; Chang, H.; Jiang, S.; Fan, Z.; Zhu, Z.; Xu, D.; Chari, P.; You, S.; Wang, Z.; and Kadambi, A. 2024b. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21676–21685.
- 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21634–21643.
- Zwicker, M.; Pfister, H.; Van Baar, J.; and Gross, M. 2001. EWA volume splatting. In *Proceedings Visualization, 2001. VIS'01.*, 29–38. IEEE.