
Wild-GS: Real-Time Novel View Synthesis from Unconstrained Photo Collections

Jiacong Xu
Johns Hopkins University
Baltimore MD 21218, USA
jxu155@jhu.edu

Yiqun Mei
Johns Hopkins University
Baltimore MD 21218, USA
yme17@jhu.edu

Vishal M. Patel
Johns Hopkins University
Baltimore MD 21218, USA
vpatel136@jhu.edu

Abstract

Photographs captured in unstructured tourist environments frequently exhibit variable appearances and transient occlusions, challenging accurate scene reconstruction and inducing artifacts in novel view synthesis. Although prior approaches have integrated the Neural Radiance Field (NeRF) with additional learnable modules to handle the dynamic appearances and eliminate transient objects, their extensive training demands and slow rendering speeds limit practical deployments. Recently, 3D Gaussian Splatting (3DGS) has emerged as a promising alternative to NeRF, offering superior training and inference efficiency along with better rendering quality. This paper presents *Wild-GS*, an innovative adaptation of 3DGS optimized for unconstrained photo collections while preserving its efficiency benefits. *Wild-GS* determines the appearance of each 3D Gaussian by their inherent material attributes, global illumination and camera properties per image, and point-level local variance of reflectance. Unlike previous methods that model reference features in image space, *Wild-GS* explicitly aligns the pixel appearance features to the corresponding local Gaussians by sampling the triplane extracted from the reference image. This novel design effectively transfers the high-frequency detailed appearance of the reference view to 3D space and significantly expedites the training process. Furthermore, 2D visibility maps and depth regularization are leveraged to mitigate the transient effects and constrain the geometry, respectively. Extensive experiments demonstrate that *Wild-GS* achieves state-of-the-art rendering performance and the highest efficiency in both training and inference among all the existing techniques. The code can be accessed via <https://github.com/XuJiacong/Wild-GS>

1 Introduction

With the development of 3D scene representation technologies, novel view synthesis aiming to create photo-realistic images from arbitrary viewpoints is becoming increasingly popular in computer vision. Neural Radiance Field (NeRF) (Mildenhall et al., 2021), as a physically inspired approach, representing the scene by radiance field and density for volume rendering, has accomplished groundbreaking synthesis quality on a range of complex scenes. Many subsequent works extend the applications of NeRF (Haque et al., 2023; Huang et al., 2022; Poole et al., 2022; Yuan et al., 2022) and further improve its robustness (Verbin et al., 2022; Mildenhall et al., 2022; Ma et al., 2022). One central assumption of NeRF and other traditional novel view synthesis methods is that the geometry, material, and lighting conditions in the collected images should remain constant. However, the large number of tourist photos on the internet is usually captured at different times and weathers, or with various camera settings, containing variable appearance and transient occluders. Training NeRF with these *in-the-wild* image collections will result in ghosting and over-smoothing artifacts.

To tackle the aforementioned problem, NeRF-W (Martin-Brualla et al., 2021) introduces image-dependent appearance and transient embeddings to model the appearance variation and transient




Figure 1: Visual comparison between Wild-GS and other existing approaches (Chen et al., 2022b; Yang et al., 2023). Wild-GS presents superior computational efficiency (tested on single RTX3090), as well as better appearance and geometry reconstruction. Additionally, by modifying the appearance features defined by Wild-GS, one can freely adjust the visual appearance of the entire scene.

uncertainty, which successfully accomplishes high-fidelity and occluder-free rendering. Subsequent advanced variants (Chen et al., 2022b; Yang et al., 2023; Fridovich-Keil et al., 2023; Kassab et al., 2024) have further improved the appearance disentanglement and transfer capability across different views. These approaches decode the appearance variances for all the 3D points with the same latent global vector or directly redistribute the rendered 2D features by global statistics as style transfer (Yang et al., 2023), which cannot explicitly capture the positional-awareness local reflectance. Furthermore, the training cost for implicit representations is extremely high, and the volumetric ray-marching nature of existing methods prevents them from accomplishing real-time rendering.

Although various explicit and hybrid representations have been proposed in recent years to reduce the training cost and expedite the rendering speed of NeRF (Liu et al., 2020; Fridovich-Keil et al., 2022; Yu et al., 2021a; Müller et al., 2022), they usually come with the sacrifice of synthesis quality. Recently, 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) has revolutionized the realm of novel view synthesis by allowing high-quality and real-time rendering along with competitive training efficiency. Nevertheless, similar to NeRF, the original 3DGS struggles to handle the *in-the-wild* image collections, causing obvious ghosting artifacts and geometry errors. A large amount of Gaussians are placed in the highly occluded areas to model and shade the transient objects, which induces meaningless computations in both the training and rendering stages. Besides, without appearance encoding, 3DGS fails to distinguish the appearance variation between different views. In this paper, we introduce an adaptation of 3DGS, namely Wild-GS, which improves the robustness of 3DGS in dealing with unconstrained images without significant trade-offs on its efficiency merits.

Following the same paradigm as the previous *in-the-wild* approaches, the appearance of each view is decomposed into image-dependent and image-invariant components. Every 3D Gaussian in Wild-GS stores an intrinsic vector to represent the inherent material property around its dominant area, which is invariant to the external environment changes. In addition to the global appearance embedding utilized in prior works, which encodes the universal appearance for all the Gaussians, such as different global illuminations and camera ISP settings, we further align each Gaussian with its corresponding local appearance feature to describe the positional-awareness local variations of reflectance by combining triplane (Chan et al., 2022) and 3DGS representations. Following the nature of 3DGS, the local appearance modeling is implemented in an explicit manner and thus expedites the training process.

The contribution of this paper can be summarized as follows: *i)* We propose a hierarchical appearance decomposition strategy to handle the complicated appearance variances across different views; *ii)* We design an explicit local appearance modeling method to capture the high-frequency appearance details; *iii)* Our model accomplishes the best rendering quality and the highest efficiency in training and inference; *iv)* Our model presents high-quality appearance transfer from arbitrary images.

2 Related Work

2.1 3D Scene Representation

Neural Radiance Field. Synthesizing arbitrary views of a scene from multi-view images is a long-standing research topic in computer vision and graphics. Photo-realistic rendering of novel view

requires accurate reconstruction of 3D geometry and appearance. Early approaches represent the 3D scenes by explicit mesh (Wachter et al., 2014; DEBEC, 1996; Liu et al., 2019) or voxel (Kutulakos & Seitz, 2000; Seitz & Dyer, 1999; Szeliski & Golland, 1998), leading to geometry and appearance inconsistency in complex scenarios (Gao et al., 2022b). Neural Radiance Field (NeRF) (Mildenhall et al., 2021) utilizes an interpolation approach between different views to parse the scene information by neural networks implicitly and provides revolutionary impact. Subsequent attempts (Barron et al., 2021, 2022; Yu et al., 2021b; Xu et al., 2022; Barron et al., 2023) further improve the modeling capability and rendering quality of the original NeRF. To mitigate the resource consumption of large MLP training and inference, advanced radiance field representations such as voxel grid (Liu et al., 2020; Sun et al., 2022a; Fridovich-Keil et al., 2022; Sun et al., 2022b; Hedman et al., 2021), octree (Yu et al., 2021a; Wang et al., 2022; Bai et al., 2023), planes (Chan et al., 2022; Cao & Johnson, 2023; Chen et al., 2022a; Fridovich-Keil et al., 2023) and hash grid (Müller et al., 2022), were investigated. However, the volumetric ray-marching nature of NeRF-based methods involves costly computation due to dense queries along each ray and thus restricts their rendering speed.

3D Gaussian Splatting. Recently, 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) is emerging as a promising alternative to NeRF by presenting impressive efficiency and higher rendering quality (Chen & Wang, 2024). Avoiding unnecessary computation in empty space, 3DGS represents the scene by millions of controllable 3D Gaussians and accomplishes real-time novel view rendering by directly projecting Gaussians within FOV to the 2D screen (Zwicker et al., 2001). Inspired by its efficiency advantages, 3DGS is being employed to replace NeRF in various vision and graphics applications, such as autonomous driving (Zhou et al., 2023b; Yan et al., 2024; Zhou et al., 2024), text-to-3D generation (Chen et al., 2023b; Chung et al., 2023; Liu et al., 2023; Ling et al., 2023), mesh extraction and physical simulation (Guédon & Lepetit, 2023; Xie et al., 2023), controllable 3D scene editing (Chen et al., 2023a; Fang et al., 2023; Zhou et al., 2023a), sparse view reconstruction (Szymanowicz et al., 2023; Charatan et al., 2023; Li et al., 2024), and human avatar (Hu et al., 2023; Shao et al., 2024; Li et al., 2023). This paper embraces the high efficiency of 3DGS and equips it with highly explicit appearance control to achieve fast training and rendering from unconstrained image collections. Notably, improvements in different aspects of the robustness of 3DGS are still under-explored (Zhao et al., 2024; Darmon et al., 2024; Meng et al., 2024).

2.2 Novel View Synthesis *in-the-wild*

Traditional novel view synthesis methodologies assume that the geometry, material, and lighting are static in the world, but the *in-the-wild* images collected from internet severely violate this assumption. To resolve this issue, NRW (Meshry et al., 2019) applies a rerendering network to merge the semantic mask and latent appearance embedding. Differently, NeRF-W (Martin-Brualla et al., 2021) leverages the implicit radiance field to represent the 3D scene and attaches transient and appearance embeddings for each training image to handle the environmental variations. Instead of optimizing the appearance for each inference image, Ha-NeRF (Chen et al., 2022b) encodes the images into latent appearance features by a CNN and employs a 2D visibility map conditioning on a learnable transient vector to emphasize static objects. Rendering the color of a pixel by a single ray may lose the global information across multiple pixels. CR-NeRF (Yang et al., 2023) utilizes grid sampling to render cross-ray features and transforms them by the global statistics derived from the reference view. More recently, K -planes (Fridovich-Keil et al., 2023), an extension of triplane (Chan et al., 2022), is implemented to factorize the variable radiance field into 2D planes and an appearance vector. After fitting the K -planes to the scene, RefinedFields (Kassab et al., 2024) additionally introduces a scene refining stage to refine the plane features by Stable Diffusion (Rombach et al., 2022). The aforementioned methods relying on implicit representations require costly training, and their rendering speed is also limited by the ray-marching nature of the radiance field.

3 Preliminaries

3.1 3D Gaussian Splatting

A collection of trainable 3D Gaussians is leveraged by 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) to represent the scene explicitly. The position and shape of each 3D Gaussian are controlled by its center $\mu \in \mathbb{R}^3$ and a 3D covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$ in world coordinates, respectively. Σ is decomposed into scaling s and rotation r to preserve positive semi-definite property. These Gaussians

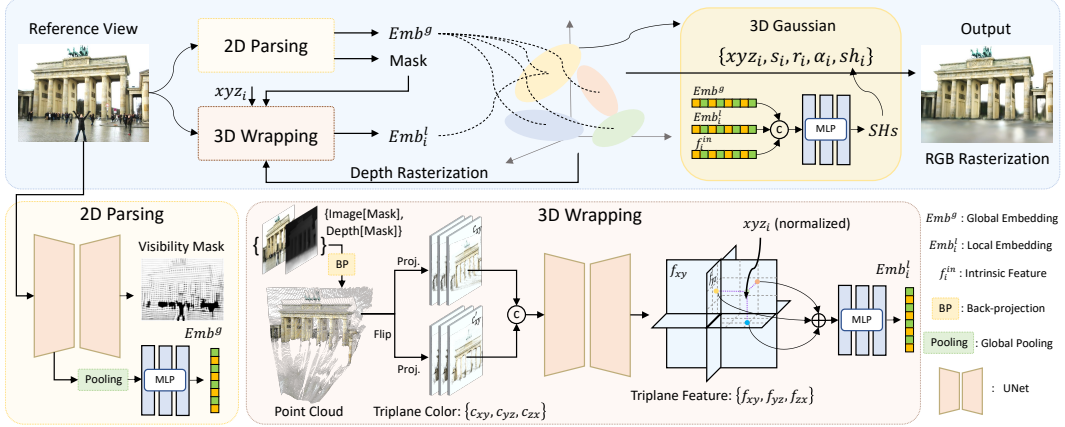


Figure 2: Overview of the architecture of our proposed Wild-GS. The reference view is first processed by a 2D Parsing module to extract the visibility mask and global appearance embedding. Given the mask and rendered depth from 3DGS, we back-project the 2D reference image without transient objects to the space and construct the static 3D point cloud. Then, these 3D points are re-projected to three predefined orthogonal planes using their normalized coordinates for generation of triplane features. Each 3D Gaussian queries its local appearance embedding by providing the spatial coordinate to the 3D Wrapping module. With the global and local embeddings and the stored intrinsic feature, we can predict the SH coefficients sh of every 3D Gaussian for RGB rasterization.

are directly projected to the screen for high-speed rendering, called Splatting (Zwicker et al., 2001). Given the world-to-camera transformation matrix \mathbf{W} and Jacobian of the affine approximation of perspective transformation \mathbf{J} , the projected 2D covariance matrix can be computed by:

$$\Sigma' = \mathbf{J}\mathbf{W}\Sigma\mathbf{W}^\top\mathbf{J}^\top. \quad (1)$$

Each Gaussian stores a learnable opacity α_i and a set of spherical harmonic (SH) coefficients for view-dependent color c_i computation. After sorting the Gaussians by depth, alpha compositing is utilized to compute the final color for each pixel, which can be written as:

$$C = \sum_{i=1}^n c_i \alpha'_i \prod_{j=1}^{i-1} (1 - \alpha'_j), \quad (2)$$

where α'_i is the multiplication of α_i and splatted 2D Gaussian. Heuristic point densification and pruning are employed in the training process for efficient 3D scene representation.

3.2 Triplane Representation

The hybrid explicit-implicit triplane representation was first introduced by EG3D (Chan et al., 2022) to enable 2D GAN (Goodfellow et al., 2020) to possess the capability of 3D generation. Many subsequent works (Gao et al., 2022a; Wang et al., 2023; Shue et al., 2023) further demonstrate the superiority of triplane on text-to-3D or 2D-to-3D generation tasks. A triplane consists of three individual 2D feature maps representing three orthogonal planes P_{XY} , P_{YZ} , and P_{ZX} in 3D space. The feature of a 3D point p can be queried by projecting the point onto three planes and summing up the interpolated individual plane features $f_p = f_{xy} + f_{yz} + f_{zx}$. In the original implementation, an MLP is attached to decode the feature into density and color for volume rendering.

4 Wild-GS

The objective of Wild-GS is to improve the robustness of 3DGS in handling *in-the-wild* photo collections without losing much on its efficiency benefits. Specifically, Wild-GS applies a heuristic decomposition on the appearances of space points to accomplish hierarchical and explicit appearance control for each 3D Gaussian. Figure 2 illustrates the pipeline of Wild-GS, which mainly consists of three components: *i*) 2D Parsing Module extracts the high-level 2D appearance information and predicts the mask for static objects; *ii*) 3D Wrapping Module constructs the positional-awareness local

appearance embedding for each 3D Gaussian; *iii*) A fusion network is shared for all the Gaussians merges and decodes the appearance features for adaptive SHs prediction.

4.1 Hierarchical Appearance Modeling

In this section, we propose a hierarchical appearance modeling framework for 3DGS that adaptively generates specific appearance embedding for individual 3D Gaussian. In this framework, the appearance of each Gaussian for a given reference view is determined by three components: (a) Global appearance embedding Emb^g capturing the illumination level or tone mapping of the entire scene; (b) Local appearance embedding Emb_i^l describing the positional-aware local reflectance for i -th 3D Gaussian; (c) Intrinsic feature f_i^{in} storing the inherent attributes of the material in the dominant area for each Gaussian. Before rasterization, a shared fusion network M_θ^F is leveraged to decode the view-dependent color sh from these three appearance components:

$$sh_i = M_\theta^F(Emb^g \oplus Emb_i^l \oplus f_i^{in}), \quad (3)$$

where \oplus refers to concatenate operation. This heuristic appearance decomposition allows efficient and effective appearance control from the entire scene to local areas, incorporating commonalities and specificities of different Gaussians. Without modifying the rasterization process, the superior inference efficiency of the original 3DGS is preserved after caching the sh for every Gaussian.

4.1.1 Global Appearance Encoding

Tourist photos collected on the internet are usually captured in different weathers and times or with various cameras, resulting in obvious appearance variations in the photo collections. Most of these variations are shared by different areas of the scene and determined by common environmental factors, e.g., the lighting level at the time of shooting. Furthermore, different post-processing processes of the photograph devices, such as gamma correction, exposure adjustment, and tone mapping, lead to various visual effects in the captured scene. These external and internal factors globally influence the appearance and are invariant to the positions of the space points. Therefore, we apply the global appearance embedding Emb^g extracted from the given reference image I_R to all the 3D Gaussians to model the low-frequency appearance changes among the entire scene. This encoding process is implemented by directing the feature maps F_{I_R} obtained from the UNet encoder in the 2D Parsing module through a global average pooling layer, followed by a trainable MLP M_θ^G :

$$Emb^g = M_\theta^G(AvgPooling(F_{I_R})). \quad (4)$$

4.1.2 Local Appearance Control

The physical interaction between 3D scenes and their environments presents long-standing challenges in computer graphics. For instance, varying directions of light can create distinct specular highlights and shadows in specific areas. The global appearance embedding is inadequate to model the detailed and high-frequency local appearance changes, especially for 3DGS with explicit and discrete representation. Thus, we design a local appearance control strategy to explicitly align the appearance clues from the reference image to the corresponding 3D Gaussians by combining the triplane and 3DGS representations. Specifically, each 3D Gaussian can query its local appearance embedding Emb_i^l from the generated triplane feature maps using its learned center position $\mu_i = xyz_i$.

Triplane Color Creation. Unlike previous works (Wang et al., 2023; Zou et al., 2023) that learn triplane features in generative ways, Wild-GS leverages the color information from the reference view to infer high-dimensional triplane maps. To capture the 3D local appearance, we first back-project the reference image into the space using the rendered depth \hat{D}_{I_R} and camera parameters ω_{I_R} . Here, the visibility mask M_{I_R} is required to exclude transient objects. This step can be implemented as:

$$\hat{D}_{I_R} = \sum_{i=1}^n d_i \alpha'_i \prod_{j=1}^{i-1} (1 - \alpha'_j), \quad (5)$$

$$\{C_{I_R}, P_{I_R}\} = BP(I_R[M_{I_R} > Th], \hat{D}_{I_R}[M_{I_R} > Th], \omega_{I_R}), \quad (6)$$

where Th defines the threshold distinguishing the transient and static objects in the visibility mask and $\{C_{I_R}, P_{I_R}\}$ refers to the generated point cloud with positions P_{I_R} and colors C_{I_R} . Then, the 3D points are normalized for re-projection onto the three orthogonal planes defined by the triplane:

$$\{c_{xy}, c_{yz}, c_{zx}\} = Proj(C_{I_R}, \tilde{P}_{I_R}), \quad (7)$$




Figure 3: (a) The point cloud from the reference image is projected along three axes and their reverses to generate the triplane color; (b) Illustration of the distribution of the 3D Gaussians on the original triplane and cropped one. Axis-aligned bounding box (AABB) is utilized to accomplish 3D cropping.

where \tilde{P}_{I_R} represents the normalized point positions. Viewing an object from one side can result in the loss of information from the opposite side. While three orthographic views capture most of the scene’s geometric information, we’ve empirically observed that a simple network often struggles to learn the complex multi-view correlations. Therefore, we concatenate the projections along each axis and their corresponding reverses, denoted by $\{c'_{xy}, c'_{yz}, c'_{zx}\}$, to form the triplane color (Figure 3-(a)).

The triplane color is further processed by a UNet U_{θ}^{3D} to extract the triplane feature maps $F_{I_R}^T$, which then are utilized to interpolate the three plane features $\{f_{xy}^i, f_{yz}^i, f_{zx}^i\}$ for each 3D Gaussian.

$$F_{I_R}^T = U_{\theta}^{3D}(\{c_{xy}, c_{yz}, c_{zx}\} \oplus \{c'_{xy}, c'_{yz}, c'_{zx}\}). \quad (8)$$

The positional-awareness local appearance embedding can be obtained by feeding the summation of three plane features to an MLP M_{θ}^L :

$$Emb_i^l = M_{\theta}^L(f_{xy}^i + f_{yz}^i + f_{zx}^i). \quad (9)$$

Efficient Triplane Sampling. The resolution of the triplane feature maps must be sufficiently high to accurately describe the detailed variations in local appearance. However, sampling from high-resolution maps is computationally expensive in terms of time and space. As depicted in Figure 3-(b), the Gaussian points projected onto the triplane maps tend to concentrate within a small area. To reduce the sampling cost and improve the utilization of the triplane maps, we apply an axis-aligned bounding box (AABB) to confine the 3D space containing most of the 3D Gaussians. As a compliment, we set $Emb_i^l = v$, where v is a learnable vector, for Gaussians outside the predefined AABB.

4.1.3 Learnable Intrinsic Feature

Beyond image-based appearance modeling strategies, Wild-GS maintains a learnable intrinsic appearance feature for each 3D Gaussian. This feature characterizes the inherent material properties within its dominant area, which remain consistent despite environmental changes. This approach is inspired by EAGLES (Girish et al., 2023), which compresses the attributes of the 3D Gaussian into a low-dimensional latent vector. By separating internal and external appearances, this heuristic decomposition effectively enhances the rendering quality of Wild-GS in our experiments.

4.2 Depth Regularization

Depth information of the training images has been widely employed in the sparse view reconstruction methods (Deng et al., 2022; Li et al., 2024; Zhu et al., 2023). Different from NeRF, 3DGS with unstructured representation is sensitive to geometric regularization. In Wild-GS, the rendered depth is leveraged to back-project the reference view, so it determines the precision of the generated point cloud. Thus, we also incorporate the depth regularization strategy to constrain the geometry of the scene. Specifically, Depth Anything (Yang et al., 2024) is employed here to estimate the monocular depth $D_{I_R}^{Est}$ for each reference view. Besides, we modify the Pearson correlation loss proposed by FSGS (Zhu et al., 2023) by masking out the transient objects and applying it in Wild-GS:

$$\mathcal{L}^D = \frac{Cov(\hat{D}_{I_R}[M_{I_R} > Th], D_{I_R}^{Est}[M_{I_R} > Th])}{\sqrt{Var(\hat{D}_{I_R}[M_{I_R} > Th]) \cdot Var(D_{I_R}^{Est}[M_{I_R} > Th])}}. \quad (10)$$

Method	Brandenburg Gate			Sacre Coeur			Trevi Fountain			Efficiency	
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	Training	Inference
3DGS	19.63	0.8817	0.1378	17.95	0.8455	0.1633	17.23	0.6963	0.2815	0.12	220
NeRF-W	24.17	0.8905	0.1670	19.20	0.8076	0.1915	18.97	0.6984	0.2652	60.3	0.05
Ha-NeRF	24.04	0.8873	0.1391	20.02	0.8012	0.1710	20.18	0.6908	0.2225	71.6	0.05
CR-NeRF	26.53	0.9003	0.1060	22.07	0.8233	0.1520	21.48	0.7117	0.2069	101	0.02
Wild-GS [†]	27.81	0.9180	0.1085	23.85	0.8567	0.1575	22.77	0.7629	0.2229	0.23	225
Wild-GS	29.65	0.9333	0.0951	24.99	0.8776	0.1270	24.45	0.8081	0.1622	0.52	227
w/o Crop	29.15	0.9324	0.0968	24.93	0.8761	0.1323	23.61	0.7915	0.1856	0.85	223
w/o {c'}	28.52	0.9286	0.1019	24.37	0.8623	0.1379	24.14	0.8036	0.1657	0.47	227
w/o Global	29.00	0.9295	0.1006	24.82	0.8761	0.1309	24.18	0.8045	0.1702	0.53	224
w/o Mask	28.85	0.9292	0.0981	24.71	0.8833	0.1216	23.91	0.8063	0.1590	0.53	205
w/o Depth	28.36	0.9261	0.1064	23.36	0.8573	0.1672	23.20	0.7878	0.1827	0.48	218

Table 1: Quantitative experimental results of existing methods (NeRF-W (Martin-Brualla et al., 2021), Ha-NeRF (Chen et al., 2022b), and CR-NeRF (Yang et al., 2023)) and Wild-GS on Phototourism dataset. Wild-GS[†] indicates that the model is trained by 15k iterations. The efficiencies of different methods are quantified by their training times (hours) and inference speeds (frame per second). Crop, Global, Mask, and Depth are abbreviations for triplane cropping, global appearance encoding, transient mask prediction, and depth regularization, respectively. {c'} refers to {c'_{xy}, c'_{yz}, c'_{zx}}.

4.3 Handling Transient Objects

The view-inconsistent transient objects widely exist in unconstrained image collections, requiring the model to put more useless and even malicious efforts into representing their appearance, especially for 3DGS, where the geometry and the appearance are highly entangled. Similar to CR-NeRF (Yang et al., 2023), we also leverage a visibility mask to indicate the easier exemplars, the static objects, by feeding the reference view to the UNet U_{θ}^{2D} in the 2D Parsing module and adaptively predicting the visibility mask M_{I_R} . The training of U_{θ}^{2D} is implemented in an unsupervised manner by forcing the rendering loss to focus only on the static objects. Additional mask regularization is utilized to prevent meaningful pixels from being masked. The detailed implementation can be written as:

$$\mathcal{L}^I = \lambda^I |I_R \odot M_{I_R} - \hat{I}_R \odot M_{I_R}| + (1 - \lambda^I) \cdot SSIM(I_R \odot M_{I_R}, \hat{I}_R \odot M_{I_R}). \quad (11)$$

$$\mathcal{L}^M = (1 - M_{I_R})^2. \quad (12)$$

For our explicit appearance control strategy, the accuracy of the visibility mask is crucial since the model may project the appearance of the transient objects to the static ones or, conversely, partially mask the appearance of the static objects. Thus, we provide a mask threshold to control the trade-off.

4.4 Training Objective

Incorporating all the aforementioned techniques, we can build the training objective for Wild-GS:

$$\mathcal{L}_{total} = \mathcal{L}^I + \lambda^M \mathcal{L}^M + \lambda^D \mathcal{L}^D. \quad (13)$$

Notably, in the initial training stage, the prediction of M_{I_R} is not sufficiently accurate. Therefore, to expedite the training speed and remedy the transient effect, the depth regularization and explicit appearance control strategies, whose functionalities are highly dependent on the mask, are not used.

5 Experimental Results

Implementation, Datasets, and Evaluation. We develop our method based on the original implementation of 3DGS (Kerbl et al., 2023). All the networks in Wild-GS are optimized by Adam optimizer (Kingma & Ba, 2014). The hyper-parameter λ^M is reduced linearly to effectively remove the transient objects and stabilize the training process. Following previous works (Chen et al., 2022b; Yang et al., 2023), we evaluate different methods on three *in-the-wild* datasets: "Brandenburg Gate", "Sacre Coeur", and "Trevi Fountain" extracted from the Phototourism dataset and downsample the images by 2 times ($R/2$). All the training times and inference speeds are tested on a single RTX3090 for fair comparison. In addition to comparing existing approaches, we also provide ablation studies for different model components to validate their effectiveness.

5.1 Comparison Experiments

Quantitative Comparison. Original 3DGS is not equipped with appearance modeling modules, resulting in lower rendering performance on *in-the-wild* datasets. As shown in Table 1, our adaptation



Figure 4: Visual comparison of rendering quality between different approaches. Red and blue crops mainly emphasize appearance and geometry differences, respectively.

Wild-GS successfully improves the capability and robustness of 3DGS and accomplishes superior rendering quality and efficiency compared with existing approaches. Specifically, Wild-GS presents around 3 PSNR increase along with $200\times$ shorter training time and $10000\times$ faster rendering speed versus the previous state-of-the-art model CR-NeRF. Furthermore, Wild-GS has achieved the highest PSNR and SSIM among all the methods with only 14 minutes of training (15k iterations).

Qualitative Comparison. The advanced representation of 3DGS has the potential to recover the high-frequency appearance and geometry details of the scene. Built upon 3DGS, our method Wild-GS also shows more accurate reconstructions of local appearance and geometry along with better global appearance modeling capability compared with other NeRF-based methods (Figure 4).

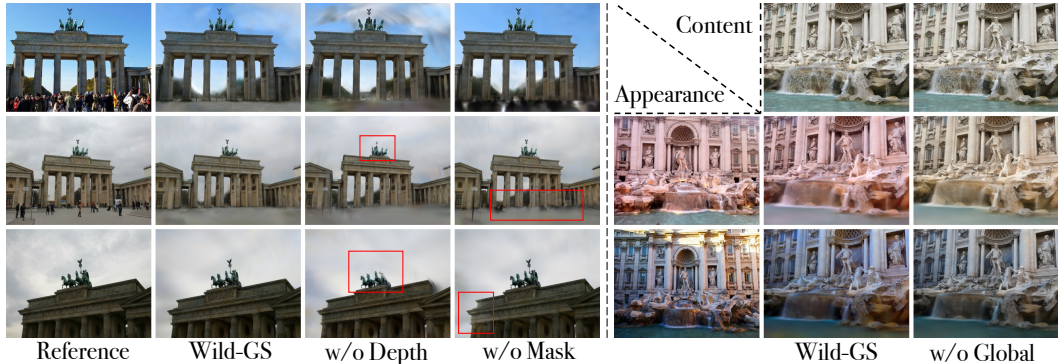


Figure 5: Rendering results of ablation study on Wild-GS when removing depth regularization, transient mask (left), and global appearance encoding (right). Red rectangles indicate the areas where geometry is missing or color inconsistency happens. Notations follow Table 1

5.2 Ablation Study

To explicitly model the complicated appearance variances of unconstrained photos and follow the nature of 3DGS, we leverage triplane to generate the position-awareness local appearance feature for each 3D Gaussian. Additional re-projections along opposite directions $\{c'_{xy}, c'_{yz}, c'_{zx}\}$ are combined with original $\{c_{xy}, c_{yz}, c_{zx}\}$ to constitute the color triplane. Table 1 indicates that this operation effectively improves the rendering quality with slight trade-offs in training efficiency. Besides, by



Figure 6: Appearance and style transfer to novel views using reference images inside and outside the training dataset. Two arbitrary style images are borrowed from Ha-NeRF (Chen et al., 2022b).

cropping the triplane and confining the sampling, Wild-GS significantly reduces the training time by around 39% while preserving and even improving the rendering quality.

Depth regularization is utilized in the training process to constrain the geometry and stabilize our explicit appearance modeling strategy. Wild-GS cannot align appearance features to corresponding Gaussians without accurate depth information, causing performance degradation and missing geometry in the final reconstruction (Figure 5). Without the transient mask prediction, ghosting effects and color inconsistencies are observed in the areas highly occluded by transient objects.

In terms of the metrics, leveraging global appearance modeling cannot obtain obvious improvements in rendering quality. The underlying reason is that the re-projected color triplane already contained all the appearance information for the reference viewpoint. However, as shown in Figure 5, Wild-GS (w/o Global) fails to capture the global appearance statistics and struggles to transfer the global color tone to another view. Thus, both embeddings are critical to the robustness of Wild-GS.

5.3 Appearance Transfer

In addition to the reference-based view synthesis task, Ha-NeRF and CR-NeRF extended the application of *in-the-wild* methods to appearance transfer and even style transfer of 3D scenes, which further validates their appearance modeling capabilities. Figure 6 contains the qualitative comparison of Wild-GS and CR-NeRF on this new task. For most appearance (style) images, Wild-GS can successfully capture the overall color tone and transfer it to novel views. Compared with CR-NeRF, our method accomplishes more accurate and robust appearance modeling and presents more color-consistent renderings. Furthermore, by linearly combining the appearance features extracted from two different reference views, one can freely tune the appearance of the 3D scene (Figure 1).

6 Conclusion

In this paper, we introduce *Wild-GS*, which adapts 3DGS to handle unconstrained photo collections without significant trade-offs on its efficiency benefits. Specifically, Wild-GS hierarchically decomposes the appearance of a given reference view into image-based global and local appearance embeddings and image-invariant intrinsic appearance features for each Gaussian. Following the nature of 3DGS, we leverage triplane representation to accomplish explicit local appearance modeling and allow Gaussians to sample their triplane features according to their specific positions. Triplane generation and sampling modifications are proposed to improve the rendering quality and training efficiency. Depth regularization and transient object handling are employed for better geometry and color consistency. Extensive experiments demonstrate that *Wild-GS* achieves state-of-the-art rendering performance and the highest efficiency on training and inference among all the existing *in-the-wild* techniques. Besides, applications for appearance transfer and tuning are provided.

7 Acknowledgement

This research is based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA R&D Contract No. 140D0423C0076. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

References

- Haotian Bai, Yiqi Lin, Yize Chen, and Lin Wang. Dynamic plenotree for adaptive sampling refinement in explicit nerf. In *ICCV*, pp. 8785–8795, 2023.
- Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, pp. 5855–5864, 2021.
- Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, pp. 5470–5479, 2022.
- Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 19697–19705, 2023.
- Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *CVPR*, pp. 130–141, 2023.
- Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *CVPR*, pp. 16123–16133, 2022.
- David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. *arXiv preprint arXiv:2312.12337*, 2023.
- Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *ECCV*, pp. 333–350. Springer, 2022a.
- Guikun Chen and Wenguan Wang. A survey on 3d gaussian splatting. *arXiv preprint arXiv:2401.03890*, 2024.
- Xingyu Chen, Qi Zhang, Xiaoyu Li, Yue Chen, Ying Feng, Xuan Wang, and Jue Wang. Hallucinated neural radiance fields in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12943–12952, 2022b.
- Yiwen Chen, Zilong Chen, Chi Zhang, Feng Wang, Xiaofeng Yang, Yikai Wang, Zhongang Cai, Lei Yang, Huaping Liu, and Guosheng Lin. Gaussianeditor: Swift and controllable 3d editing with gaussian splatting. *arXiv preprint arXiv:2311.14521*, 2023a.
- Zilong Chen, Feng Wang, and Huaping Liu. Text-to-3d using gaussian splatting. *arXiv preprint arXiv:2309.16585*, 2023b.
- Jaeyoung Chung, Suyoung Lee, Hyeongjin Nam, Jaerin Lee, and Kyoung Mu Lee. Luciddreamer: Domain-free generation of 3d gaussian splatting scenes. *arXiv preprint arXiv:2311.13384*, 2023.
- Hiba Dahmani, Moussab Bennehar, Nathan Piasco, Luis Roldao, and Dzmitry Tsishkou. Swag: Splatting in the wild images with appearance-conditioned gaussians. *arXiv preprint arXiv:2403.10427*, 2024.
- François Darmon, Lorenzo Porzi, Samuel Rota-Bulò, and Peter Kotschieder. Robust gaussian splatting. *arXiv preprint arXiv:2404.04211*, 2024.

- PE DEBEC. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *Proc. SIGGRAPH'96*, 1996.
- Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12882–12891, 2022.
- Jiemin Fang, Junjie Wang, Xiaopeng Zhang, Lingxi Xie, and Qi Tian. Gaussianeditor: Editing 3d gaussians delicately with text instructions. *arXiv preprint arXiv:2311.16037*, 2023.
- Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, pp. 5501–5510, 2022.
- Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *CVPR*, pp. 12479–12488, 2023.
- Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. *Advances In Neural Information Processing Systems*, 35:31841–31854, 2022a.
- Kyle Gao, Yina Gao, Hongjie He, Dening Lu, Linlin Xu, and Jonathan Li. Nerf: Neural radiance field in 3d vision, a comprehensive review. *arXiv preprint arXiv:2210.00379*, 2022b.
- Sharath Girish, Kamal Gupta, and Abhinav Shrivastava. Eagles: Efficient accelerated 3d gaussians with lightweight encodings. *arXiv preprint arXiv:2312.04564*, 2023.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. *arXiv preprint arXiv:2311.12775*, 2023.
- Ayaan Haque, Matthew Tancik, Alexei A Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 19740–19750, 2023.
- Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *ICCV*, pp. 5875–5884, 2021.
- Liangxiao Hu, Hongwen Zhang, Yuxiang Zhang, Boyao Zhou, Boning Liu, Shengping Zhang, and Liqiang Nie. Gaussianavatar: Towards realistic human avatar modeling from a single video via animatable 3d gaussians. *arXiv preprint arXiv:2312.02134*, 2023.
- Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High dynamic range neural radiance fields. In *CVPR*, pp. 18398–18408, 2022.
- Karim Kassab, Antoine Schnepf, Jean-Yves Franceschi, Laurent Caraffa, Jeremie Mary, and Valérie Gouet-Brunet. Refinedfields: Radiance fields refinement for unconstrained scenes, 2024.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)*, 42(4):1–14, 2023.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *International journal of computer vision*, 38:199–218, 2000.
- Jiahe Li, Jiawei Zhang, Xiao Bai, Jin Zheng, Xin Ning, Jun Zhou, and Lin Gu. Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization. *arXiv preprint arXiv:2403.06912*, 2024.

- Zhe Li, Zerong Zheng, Lizhen Wang, and Yebin Liu. Animatable gaussians: Learning pose-dependent gaussian maps for high-fidelity human avatar modeling. *arXiv preprint arXiv:2311.16096*, 2023.
- Huan Ling, Seung Wook Kim, Antonio Torralba, Sanja Fidler, and Karsten Kreis. Align your gaussians: Text-to-4d with dynamic 3d gaussians and composed diffusion models. *arXiv preprint arXiv:2312.13763*, 2023.
- Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020.
- Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *ICCV*, pp. 7708–7717, 2019.
- Xian Liu, Xiaohang Zhan, Jiayang Tang, Ying Shan, Gang Zeng, Dahua Lin, Xihui Liu, and Ziwei Liu. Humangaussian: Text-driven 3d human generation with gaussian splatting. *arXiv preprint arXiv:2311.17061*, 2023.
- Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12861–12870, 2022.
- Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *CVPR*, pp. 7210–7219, 2021.
- Jiarui Meng, Haijie Li, Yanmin Wu, Qiankun Gao, Shuzhou Yang, Jian Zhang, and Siwei Ma. Mirror-3dgs: Incorporating mirror reflections into 3d gaussian splatting. *arXiv preprint arXiv:2404.01168*, 2024.
- Moustafa Meshry, Dan B Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural rerendering in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6878–6887, 2019.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P Srinivasan, and Jonathan T Barron. Nerf in the dark: High dynamic range view synthesis from noisy raw images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16190–16199, 2022.
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4): 1–15, 2022.
- Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations*, 2022.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- Steven M Seitz and Charles R Dyer. Photorealistic scene reconstruction by voxel coloring. *International Journal of Computer Vision*, 35:151–173, 1999.
- Zhijing Shao, Zhaolong Wang, Zhuang Li, Duotun Wang, Xiangru Lin, Yu Zhang, Mingming Fan, and Zeyu Wang. Splattingavatar: Realistic real-time human avatars with mesh-embedded gaussian splatting. *arXiv preprint arXiv:2403.05087*, 2024.
- J Ryan Shue, Eric Ryan Chan, Ryan Po, Zachary Anknor, Jiajun Wu, and Gordon Wetzstein. 3d neural field generation using triplane diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20875–20886, 2023.
- Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, pp. 5459–5469, 2022a.

- Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, pp. 5459–5469, 2022b.
- Richard Szeliski and Polina Golland. Stereo matching with transparency and matting. In *ICCV*, pp. 517–524. IEEE, 1998.
- Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Splatter image: Ultra-fast single-view 3d reconstruction. *arXiv preprint arXiv:2312.13150*, 2023.
- Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *CVPR*, pp. 5481–5490, 2022.
- Michael Waechter, Nils Moehrle, and Michael Goesele. Let there be color! large-scale texturing of 3d reconstructions. In *ECCV*, pp. 836–850. Springer, 2014.
- Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Fourier plenoctrees for dynamic radiance field rendering in real-time. In *CVPR*, pp. 13524–13534, 2022.
- Renke Wang, Guimin Que, Shuo Chen, Xiang Li, Jun Li, and Jian Yang. Creative birds: Self-supervised single-view 3d style transfer. In *ICCV*, pp. 8775–8784, 2023.
- Tianyi Xie, Zeshun Zong, Yuxin Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. Physgaussian: Physics-integrated 3d gaussians for generative dynamics. *arXiv preprint arXiv:2311.12198*, 2023.
- Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5438–5448, 2022.
- Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians for modeling dynamic urban scenes. *arXiv preprint arXiv:2401.01339*, 2024.
- Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. *arXiv preprint arXiv:2401.10891*, 2024.
- Yifan Yang, Shuhai Zhang, Zixiong Huang, Yubing Zhang, and Mingkui Tan. Cross-ray neural radiance fields for novel-view synthesis from unconstrained image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15901–15911, 2023.
- Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenoctrees for real-time rendering of neural radiance fields. In *ICCV*, pp. 5752–5761, 2021a.
- Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, pp. 4578–4587, 2021b.
- Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18353–18364, 2022.
- Dongbin Zhang, Chuming Wang, Weitao Wang, Peihao Li, Minghan Qin, and Haoqian Wang. Gaussian in the wild: 3d gaussian splatting for unconstrained image collections. *arXiv preprint arXiv:2403.15704*, 2024.
- Lingzhe Zhao, Peng Wang, and Peidong Liu. Bad-gaussians: Bundle adjusted deblur gaussian splatting. *arXiv preprint arXiv:2403.11831*, 2024.
- Hongyu Zhou, Jiahao Shao, Lu Xu, Dongfeng Bai, Weichao Qiu, Bingbing Liu, Yue Wang, Andreas Geiger, and Yiyi Liao. Hugs: Holistic urban 3d scene understanding via gaussian splatting. *arXiv preprint arXiv:2403.12722*, 2024.

- Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejie Xu, Pradyumna Chari, Suya You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. *arXiv preprint arXiv:2312.03203*, 2023a.
- Xiaoyu Zhou, Zhiwei Lin, Xiaojun Shan, Yongtao Wang, Deqing Sun, and Ming-Hsuan Yang. Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. *arXiv preprint arXiv:2312.07920*, 2023b.
- Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. Fsgs: Real-time few-shot view synthesis using gaussian splatting. *arXiv preprint arXiv:2312.00451*, 2023.
- Zi-Xin Zou, Zhipeng Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Yan-Pei Cao, and Song-Hai Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. *arXiv preprint arXiv:2312.09147*, 2023.
- Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa volume splatting. In *Proceedings Visualization, 2001. VIS'01.*, pp. 29–538. IEEE, 2001.

A Comparison with Concurrent *in-the-wild* 3DGS

There are two concurrent works, GS-W (Zhang et al., 2024) and SWAG (Dahmani et al., 2024), focusing on the adaptation of 3DGS to the *in-the-wild* setting at the time when we submit this paper. While GS-W leverages adaptive sampling on 2D feature maps to get dynamic appearance embedding for each Gaussian, their methods are still constrained in the 2D space without fully explicit appearance control. In terms of rendering performance, our method surpasses GS-W by around 1.5 PSNR and SWAG by 2 PSNR on Phototourism datasets. Besides, GS-W requires 2 hours for training on a single RTX3090 (stated in their paper), while our method only takes around half an hour (32 mins).

Method	Palace of Westminster			Pantheon Exterior			Buckingham Palace		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
3DGS-AE	21.0781	0.8189	0.2289	22.4432	0.8389	0.1681	23.1788	0.8489	0.2141
GS-W	24.7343	0.8615	0.1782	26.1668	0.8888	0.1022	25.6356	0.8634	0.1671
Wild-GS	25.8281	0.8677	0.1635	26.7969	0.8888	0.1018	26.7160	0.8799	0.1592
w/o Local	21.8409	0.8093	0.2164	21.7973	0.8094	0.1568	23.6144	0.8606	0.1860
w/o f^{in}	22.2028	0.8014	0.2648	21.1427	0.7751	0.2531	24.9937	0.8500	0.2426

Table 2: Quantitative experimental results on three extra datasets. 3DGS-AE replaces the appearance encoding of Wild-GS (ours) with a learnable embedding as NeRF-W and optimizes it in an autoencoder way. Local and f^{in} refer to the triplane local appearance embedding and the intrinsic feature.

Besides the three datasets used in the main paper, we also extracted three subsets: "Palace of Westminster", "Pantheon Exterior", and "Buckingham Palace" from Phototourism dataset and implemented more experiments to further demonstrate the strengths of Wild-GS. As shown in Figure 7, Wild-GS provides more accurate appearance modeling compared with GS-W Zhang et al. (2024).




Figure 7: Visual comparison of Wild-GS (ours) and GS-W (concurrent work) and ablation study.

The quantitative results in Table 2 show that Wild-GS archives around 1 PSNR increase compared with GS-W and the local appearance modeling and intrinsic feature inside each Gaussian are required for its superior performance.

B More Implementation Details

For semantically meaningful transient mask prediction, we utilize the ResNet-18 pre-trained by ImageNet as the encoder of the UNet in the 2D Parsing module. The ratio of triplane cropping is simply set to 0.5, and fine-tuning this parameter can achieve a better trade-off in efficiency and rendering quality. The dimensions for global & local appearance embeddings and intrinsic features are set to be 16 and 32, respectively. For a fair comparison with 3DGS, we only optimize the hyper-parameters in our attached framework and do not introduce any modifications to their original setting. λ^M is linearly reduced from 0.4 to 0.1 to stabilize the training process, while λ^D is kept constant 0.05 during the entire training process. Figure 8 demonstrates the transient object handling capability of Wild-GS.




Figure 8: Visualization of the transient masks (learned in an unsupervised way) predicted by Wild-GS.

C Detailed Appearance Control

As shown in Figure 9, our method Wild-Gs can capture the high-frequency local appearance details and accomplishes a more accurate local appearance modeling than CR-NeRF, which further demonstrates the effectiveness of our proposed explicit local appearance control strategy.



Figure 9: Comparison of Wild-GS and CR-NeRF on Local appearance modeling.

D Limitation

Similar to previous approaches, Wild-GS still cannot recover the detailed geometry and appearance of the ground (road or sidewalk), causing blur and useless computations in corresponding areas. Besides, since the transient masks are learned in an unsupervised manner, they tend to mask the areas with unusual appearance (hard to model but easy to mask out) and lead to color inconsistency in these areas. We suggest using more advanced segmentation networks to obtain accurate transient masks, which is beneficial to appearance modeling and geometry reconstruction. Even though Wild-GS has achieved high efficiency in training and rendering, it still requires at least double the training time of the original 3DGS to achieve comparable rendering performance. Therefore, further reducing training time without sacrificing rendering quality is a potential and meaningful work in the future.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We state our contribution in the final paragraph of the introduction part.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We provide the limitation discussion in the appendix.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This is a learning-based method, so no theoretical proof is required.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All the details for method and implementation are introduced, so one can easily re-implement it.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We will release our code upon the publication of this paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We describe the training details in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: No significant differences are observed in different runs.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Only one RTX3090 is needed for our experiments and we indicated that in paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We fully conform to the NeurIPS code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: This paper addresses very common problem and can be applied widely, we indicate this in the introduction part.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper focuses on novel view synthesis and has very little risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The model we build our method on and the dataset we conduct experiments on are indicated in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We introduce a new model here and provide all the implementation details. (we do not know if this is a new "assets")

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We use public dataset and therefore no human involved in the experiments.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This work is a traditional task with no potential risk.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.