

# DehazeGS: Seeing Through Fog with 3D Gaussian Splatting

Jinze Yu<sup>1</sup>, Yiqun Wang<sup>1\*</sup>, Aiheng Jiang<sup>1</sup>, Zhengda Lu<sup>2</sup>, Jianwei Guo<sup>3,4</sup>,  
Yong Li<sup>1</sup>, Hongxing Qin<sup>1</sup>, Xiaopeng Zhang<sup>4</sup>

<sup>1</sup>College of Computer Science, Chongqing University

<sup>2</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>3</sup>School of Artificial Intelligence, Beijing Normal University

<sup>4</sup>MAIS, Institute of Automation, Chinese Academy of Sciences

## Abstract

Current novel view synthesis methods are typically designed for high-quality and clean input images. However, in foggy scenes, scattering and attenuation can significantly degrade the quality of rendering. Although NeRF-based dehazing approaches have been developed, their reliance on deep fully connected neural networks and per-ray sampling strategies leads to high computational costs. Furthermore, NeRF’s implicit representation limits its ability to recover fine-grained details from hazy scenes. To overcome these limitations, we propose DehazeGS, the first physics-driven 3D Gaussian Splatting (3DGS) framework for dehazing. We adopt an explicit Gaussian representation to model fog formation via a physically consistent forward rendering process, enabling reconstruction and rendering of fog-free scenes using only multi-view foggy images as input. Specifically, based on the atmospheric scattering model, we simulate the formation of fog by establishing the transmission function directly on Gaussian primitives via depth-to-transmission mapping. During training, we jointly learn the atmospheric light and scattering coefficients while optimizing the Gaussian representation of foggy scenes. At inference time, we remove the effects of scattering and attenuation in Gaussian distributions and directly render the scene to obtain dehazed views. Experiments on both real-world and synthetic foggy datasets demonstrate that DehazeGS achieves state-of-the-art performance.

## Introduction

In recent years, Neural Radiance Fields (NeRF) (Mildenhall et al. 2021) have leveraged deep fully connected neural networks to implicitly represent 3D scenes and achieve impressive rendering quality through differentiable volumetric rendering techniques (Levoy 1990; Max 1995). Subsequent works have focused on enhancing NeRF’s performance, particularly in terms of training speed (Chen et al. 2022; Müller et al. 2022) and rendering efficiency (Garbin et al. 2021; Yu et al. 2021). A recent work, 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023), abandons NeRF’s implicit rendering approach and instead represents scenes explicitly by converting point clouds into 3D Gaussians (ellipsoids). Thanks to its tile-based multi-threaded parallel rendering mechanism, 3DGS achieves high-quality real-time rendering.

\*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

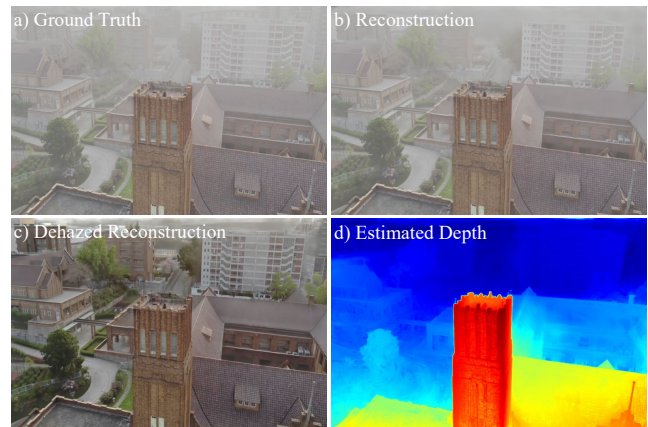


Figure 1: DehazeGS can generate accurate rendering results for scenes with participating media (b). By learning disentangled representations for the participating media and the underlying scene, it simultaneously recovers the clear scene (c) and obtains accurate depth estimation (d).

However, when capturing scenes containing scattering media, the light received by the detector is disrupted by the medium. The detected light mainly consists of two components: first, light reflected from the surfaces of objects that has been attenuated by particles of medium, and second, ambient light scattered by the particles of medium. This interference often leads to scenes suffering from low contrast and limited visibility. For both 3DGS and NeRF, most methods are designed for clear media, which is disadvantageous for applications such as autonomous driving and robotic operations under foggy scene conditions. Although some dehazing reconstruction methods based on NeRF (Ramazzina et al. 2023; Li et al. 2023) have been proposed, these approaches either suffer from extremely slow training and rendering speeds or are limited to specific indoor scenes. In contrast, traditional dehazing algorithms in low-level computer vision (Zhang and Patel 2018; Yang et al. 2022; Zhang, Zhou, and Li 2024) typically require large-scale paired or unpaired datasets for training. Moreover, these methods remain confined to a 2D image plane (Dong et al. 2020; Zheng et al. 2021; Qin et al. 2020), lacking consideration of 3D

spatial information. Therefore, they fail to meet the requirements of multi-view consistency across viewpoints necessary for 3D reconstruction. Feeding such inconsistent images into reconstruction models significantly degrades reconstruction quality (Detailed proofs can be found in the supplementary materials or Fig. 3). Thus, in the field of novel view synthesis, there is an urgent need for an end-to-end dehazing reconstruction model capable of effectively achieving physically accurate modeling and separating scattering effects in participating media.

To address the above challenges, we propose DehazeGS, the first differentiable physics-based 3DGS dehazing model. We have pioneered a novel representation of 3DGS for foggy scenes. Our method is primarily divided into three components. 1) Inspired by the transmission formula  $t(x) = e^{-\beta d(x)}$  of the Atmospheric Scattering Model (ASM) (McCartney 1976; Narasimhan and Nayar 2002), we propose to establish a mapping relationship between Gaussian depth  $d_G(x)$  and Gaussian transmission  $t_G(x)$ . By querying the depth of Gaussians as input to a convolutional neural network, we obtain the Gaussian transmission. The scattering coefficients are stored as the weights of the convolutional neural network. 2) After obtaining the Gaussian transmission, we model the foggy Gaussians based on the Atmospheric Scattering Model (ASM), in conjunction with the estimated atmospheric light coefficients. Once the foggy Gaussians are rasterized, the estimated fog map can be rendered. Using the real foggy map as supervision, we jointly optimize the parameters of the Gaussians while training the neural network and fully render the foggy scenes. 3) The accuracy of depth is crucial for estimating the transmission. Therefore, while optimizing the Gaussian representation, we adopt pseudo-depth maps as priors to impose depth regularization. To emphasize the restoration of distant details, we incorporate a depth-weighted reconstruction loss. In addition, note that for the Gaussian transmission obtained in the first stage, we directly perform alpha blending to render the transmission map. We propose to use the transmission map generated by the dark channel prior (He, Sun, and Tang 2010) and the bright channel prior (Zhang, Wang, and Wang 2021) as supervision to guide the optimization of the rendered transmission map.

Our main contributions are as follows:

- We propose the first framework for learning clear 3D Gaussian splatting solely from multi-view foggy images. This framework is capable of learning disentangled representations of participating media and clear scenes, achieving fast, high-quality dehazing reconstruction.
- We establish a mapping between Gaussian depth and its transmission, and based on the Atmospheric Scattering Model (ASM), we propose a representation method for Gaussian primitives in foggy scenes.
- We propose a novel Gaussian-based prior loss committee, leveraging these physical priors to guide the fine-tuning of our model. The combination of these prior methods further enhances the dehazing rendering performance of our model.
- Experiments on synthetic and real foggy datasets demon-

strate that our method surpasses existing approaches in both rendering quality and speed, achieving state-of-the-art results.

## Related Works

### Dehazing Based on Image Processing

In the field of image processing, one of the most classic dehazing algorithms is the Dark Channel Prior (DCP) algorithm (He, Sun, and Tang 2010), which is based on statistical observations from a large number of outdoor haze-free scenes. Subsequently, many other related algorithms have also been developed, such as the Bright Channel Prior (Zhang, Wang, and Wang 2021) and the Color Attenuation Prior (Zhu, Mai, and Shao 2015). With deep learning advancement, neural networks are increasingly applied to image processing, categorized into physics-based approaches and end-to-end methods. In physics-based algorithms, they primarily estimate hazy images' atmospheric light and transmission maps. For example, DCPDN (Zhang and Patel 2018) proposed a densely connected pyramid network capable of jointly estimating atmospheric light and transmission maps. PSD (Chen et al. 2021) introduced a dehazing network guided by physical priors. For end-to-end dehazing methods, for example, DehazeNet (Cai et al. 2016) utilizes a CNN to learn the mapping relationship between the original hazy images and the corresponding medium transmission maps. AOD-Net (Li et al. 2017) is a lightweight end-to-end dehazing network that can directly generate haze-free images without estimating transmission maps or atmospheric light independently. MSBDN (Dong et al. 2020) proposed a multi-scale enhanced dehazing network with dense feature fusion. DehazeFormer (Song et al. 2023) is designed as an end-to-end single-image dehazing method based on the Swin Transformer architecture.

### Neural Radiance Fields for Dehazing

NeRF (Mildenhall et al. 2021) employs a deep fully connected neural network for scene representation, which takes 5D sampled point coordinates as input and outputs volumetric density and view-dependent radiance (i.e., color). Volumetric rendering is then used to integrate the color and density along each viewing ray to generate pixel colors. ScatterNeRF (Ramazzina et al. 2023) builds upon NeRF by introducing an additional MLP to learn the opacity of the scattering medium and its associated color. A scattering term is added to the volumetric rendering equation, blending clear and blurred sampled points to render foggy images. Its primary limitation lies in the need for high-density sampling along each ray in space, which requires processing through a deep fully connected network to compute the volumetric density and color of sampled points. DehazeNeRF (Chen et al. 2023) extends the volumetric rendering equation via simulating atmospheric scattering physics, integrating regularization strategies for dehazing and 3D shape reconstruction. However, it is mainly designed for indoor scenes. Dehazing-NeRF's (Li et al. 2023) scattering coefficient/atmospheric light estimation network needs pre-training, lacking joint optimization of scattering medium and clear con-

tent. Additionally, it is confined to synthetic datasets with single objects (e.g., Lego, chair).

## Method

### Preliminary

**Atmospheric Scattering Model.** The Atmospheric Scattering Model (McCartney 1976; Narasimhan and Nayar 2002) (ASM) is a physical model that describes the changes in light as it propagates through the atmosphere due to scattering and absorption. Based on the principles of light propagation, it reveals how particles in the atmosphere affect the quality of captured images. Its mathematical formulation is expressed as:

$$I(x) = J(x)t(x) + A(1 - t(x)). \quad (1)$$

Here,  $I$  represents the observed hazy image,  $J$  denotes the underlying clear image, and  $A$  is the global atmospheric light, which represents the background light intensity received by the imaging devices. It is generated by the scattering of ambient light within participating media and primarily affects the overall brightness of the image.  $t(x)$  represents the transmission map, where  $x$  denotes the position of a pixel. The transmission map primarily serves to characterize the attenuation of light as it travels through scattering media (fog) in the atmosphere, directly influencing the accuracy of the dehazing results.  $t(x)$  can be further expressed as  $t(x) = e^{-\beta d(x)}$ , where  $\beta$  represents the scattering coefficient and  $d(x)$  denotes the scene depth at pixel  $x$ .

**3D Gaussian Splatting.** The 3D Gaussian approach does not rely on neural radiance fields. Instead, it represents the scene as a series of 3D Gaussian distributions (Yu et al. 2024). Based on the initialized sparse point cloud, a set of 3D Gaussians, defined as  $G$ , is parameterized by its 3D coordinates  $\mathbf{x} \in \mathbb{R}^3$ , 3D covariance  $\Sigma \in \mathbb{R}^{3 \times 3}$ , opacity  $\alpha \in \mathbb{R}$  and color  $c \in \mathbb{R}^3$ .  $c$  is represented by spherical harmonics for view-dependent appearance. The distribution of each Gaussian is defined as:

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

In order to ensure the positive semi-definite property of the covariance matrix during the optimization, it is further expressed as:

$$\Sigma = \mathbf{R}\mathbf{S}\mathbf{S}^T \mathbf{R}^T, \quad (3)$$

where  $\mathbf{R}$  and  $\mathbf{S}$  denote the rotation and scaling matrix.

3D Gaussians are projected into the 2D image space from a given camera pose  $\mathbf{P}_c = \{\mathbf{R}_c \in \mathbb{R}^{3 \times 3}, \mathbf{t}_c \in \mathbb{R}^3\}$ . Given the viewing transformation  $\mathbf{W}$  and 3D covariance matrix  $\Sigma$ , the projected 2D covariance matrix  $\Sigma'$  is computed using, as described in (Zwicker et al. 2001):

$$\Sigma' = \mathbf{J}\mathbf{W}\Sigma\mathbf{W}^T \mathbf{J}^T, \quad (4)$$

where  $\mathbf{J}$  is the Jacobian of the affine approximation of the projective transformation.

Subsequently, the transformed Gaussians are sorted based on their depth, and the sorted Gaussians are rasterized to render pixel values using the volume rendering equation:

$$C = \sum_{i \in N} c_i \alpha'_i \prod_{j=1}^{i-1} (1 - \alpha'_j), \quad (5)$$

where  $c_i$  is the learned color and the  $\alpha'_i$  is the multiplication result of the learned opacity  $\alpha_i$  and the 2D projected Gaussian.

### DehazeGS

**Modeling Gaussian distributions in foggy scenes.** Fog is a natural phenomenon caused by the scattering effect of aerosol particles in the atmosphere. When photographing in foggy conditions, the light received by the sensor primarily originates from two sources: the light reflected from the object’s surface, which undergoes scattering and attenuation by the particles, along with the environmental light, which is also scattered by the particles. As light passes through the scattering medium, it experiences varying degrees of scattering and attenuation, with the attenuation increasing as the scene depth increases (i.e., the transmission rate decreases). Since the scene represented by 3DGS is modeled using 3D Gaussians, we propose directly defining the transmission function  $t(x)$  on each Gaussian distribution. The physical meaning of this function is the proportion of light intensity that remains unattenuated after passing through the particles. Furthermore, the transmission function is associated with the depth of the Gaussian. For any Gaussian  $\mathcal{G}(\mathbf{x})$ , its depth  $d_G(\mathbf{x})$  (representing the depth  $d$  of the Gaussian ellipsoid at coordinate  $x$ ) is queried, and its corresponding transmission  $t_G(\mathbf{x})$  (representing the transmittance  $t$  of the Gaussian ellipsoid at coordinate  $x$ ) is obtained through a single-layer one-dimensional convolutional neural network. The mapping relationship can be expressed as  $F_\beta : d_G(\mathbf{x}) \rightarrow t_G(\mathbf{x})$ , where  $\beta$  represents the learnable weight, that is, the scattering coefficient. The detailed structure of the network is shown in the upper-left part of Fig. 2.

Since the  $t_G(\mathbf{x})$  ranges between 0 and 1, with smaller values indicating more severe light attenuation, we apply a ReLU activation function to the exponential term to ensure the non-negativity of  $\beta d_G(\mathbf{x})$  (the product of the scattering coefficient  $\beta$  and Gaussian depth), thereby restricting the output range of the function to between 0 and 1. However, in practical experiments, we found that when ensuring the non-negativity of  $\beta d_G(\mathbf{x})$ , feeding the result of the entire negative exponential function into a sigmoid activation function can improve the training stability of the model. Although adding a sigmoid is not strictly necessary.

Finally, based on equation  $t(x) = e^{-\beta d(x)}$  from Section 3.1, we can obtain the transmission function formula based on Gaussians as follows:

$$t_G(\mathbf{x}) = \sigma(e^{-\max(0, \beta d_G(\mathbf{x}))}), \quad (6)$$

the depth of the Gaussian ellipsoid  $d_G(\mathbf{x})$  is normalized to the range  $[0, 1]$ , and its gradient is detached to prevent the gradients from flowing through it. This approach is designed to prevent the gradient from affecting the latent representation of clear Gaussians, while allowing for better optimization of the network. For atmospheric light estimation, it is generally assumed that the atmospheric light in foggy scenes is global. The learnable global atmospheric light parameter  $A_G$  applies to all Gaussians. Consequently, the Gaussian distribution in foggy scenes is represented as follows,

$$\mathcal{G}_f(\mathbf{x}) = \mathcal{G}_c(\mathbf{x})t_G(\mathbf{x}) + A_G(1 - t_G(\mathbf{x})), \quad (7)$$

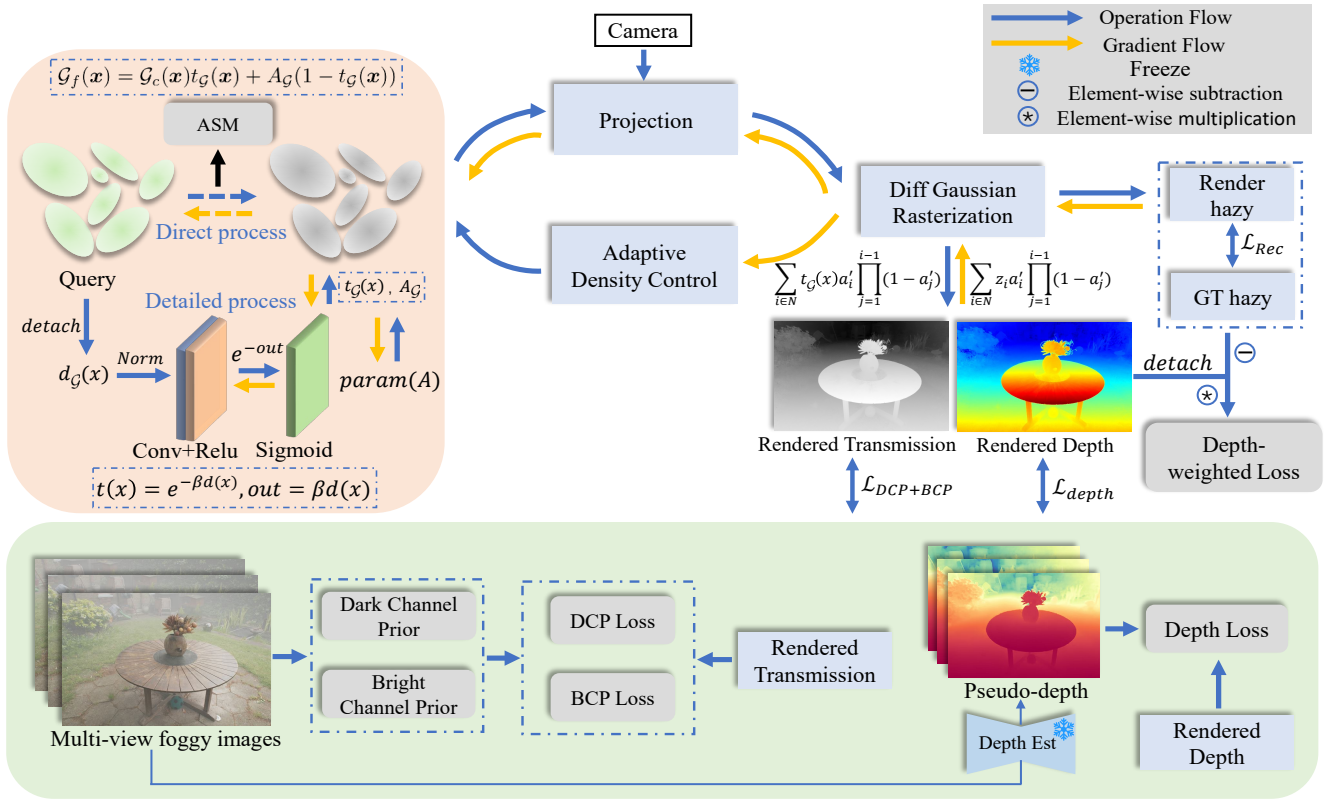


Figure 2: DehazeGS overview. We first obtain the Gaussian distributions in foggy scenes and perform alpha blending on the transmission of each Gaussian distribution to render the transmission map, which is guided and optimized using DCP and BCP priors. We utilize pseudo-depth maps as prior information for depth estimation when optimizing each input image.

where  $\mathcal{G}_f(\mathbf{x})$  and  $\mathcal{G}_c(\mathbf{x})$  represent the fogged 3D Gaussians and the latent clear 3D Gaussians, respectively.

This approach of directly establishing the transmission function and atmospheric light on the Gaussians effectively avoids the high sampling and rendering costs associated with NeRF-based methods (Chen et al. 2023; Ramazzina et al. 2023). Moreover, this approach effectively leverages the reconstruction capability in 3DGS, enhancing the structural consistency of the dehazed scene across multiple viewpoints.

After rendering the fog image, the reconstruction loss between the rendered image  $\hat{I}$  and the real fog image  $I$  is calculated. Consistent with the original 3D Gaussian reconstruction loss, The loss function is  $\mathcal{L}_1$  combined with a  $D - SSIM$  term,  $\lambda = 0.2$ :

$$\mathcal{L}_{rec} = (1 - \lambda)\mathcal{L}_1 + \lambda\mathcal{L}_{D-SSIM}. \quad (8)$$

**Gaussian Transmission Optimization.** We directly perform alpha blending on the transmission values  $t_G(\mathbf{x})$  of each Gaussian to obtain the rendered transmittance map  $\hat{\mathcal{T}}(\mathbf{P})$ :

$$\hat{\mathcal{T}}(\mathbf{P}) = \sum_{i \in N} t_G(\mathbf{x}) \alpha'_i \prod_{j=1}^{i-1} (1 - \alpha'_j), \quad (9)$$

where  $\mathbf{P}$  represents pose. We propose to use the Dark Channel Prior (DCP) (He, Sun, and Tang 2010) and further con-

sider Bright Channel Prior (BCP) (Sun and Guo 2016) algorithms to guide the optimization of our estimated transmission map. The DCP is one of the most well-known dehazing algorithms, leveraging the statistical properties of clear images to estimate the transmission map and atmospheric light. we follow the prior of traditional methods (Chen et al. 2021; Golts, Freedman, and Elad 2019) and reformulate the prior as a 3DGS-based energy function:

$$\mathcal{L}_{DCP} = \mathcal{T}_D^T(\mathbf{P})L\mathcal{T}_D(\mathbf{P}) + \lambda\tilde{\mathcal{T}}(\mathbf{P})^T\tilde{\mathcal{T}}(\mathbf{P}), \quad (10)$$

$$\tilde{\mathcal{T}}(\mathbf{P}) = (\mathcal{T}_D(\mathbf{P}) - \hat{\mathcal{T}}(\mathbf{P})).$$

Here,  $\mathcal{T}_D(\mathbf{P})$  represents the transmission map under the corresponding pose  $\mathbf{P}$ , estimated by the DCP algorithm.  $L$  and  $\lambda$  are the Laplacian-like matrix and hyper-parameter, respectively.

To mitigate the darkening effect caused by DCP, we further introduce BCP, a brightness enhancement algorithm based on the bright channel prior, which helps to enhance the overall brightness of the rendered dehazed images. The loss function is formulated as follows,

$$\mathcal{L}_{BCP} = \left\| \mathcal{T}_B(\mathbf{P}) - \hat{\mathcal{T}}(\mathbf{P}) \right\|_1 \quad (11)$$

where  $\mathcal{T}_B(\mathbf{P})$  is estimated from the BCP.

**Depth Supervision Loss.** In real-world scenarios, as the depth increases, the influence of scattering and attenuation

on light becomes more significant, resulting in reduced visibility of distant details. The accuracy of depth estimation is a key factor in obtaining a more precise transmission map. In 3DGS, depth maps can be rendered by performing alpha blending on the depths of Gaussians.

$$\hat{D} = \sum_{i \in N} z_i \alpha'_i \prod_{j=1}^{i-1} (1 - \alpha'_j). \quad (12)$$

We leverage DepthAnything V2 (Yang et al. 2024) to predict the depth maps from foggy images, generating pseudo-depth maps to guide the optimization of the 3DGS-rendered depth. Since the pseudo-depth maps are estimated from the foggy inputs and inherently contain certain errors, the alignment between  $D_{pseudo}$  and  $\hat{D}$  is relative rather than absolute. To address this, we apply weight decay during training to gradually reduce the weight of this loss. Specifically, we adopt a continuous learning rate decay schedule. The initial weight is set to 1 at step 0 and gradually reduced to 0.01 at the specified maximum step, with logarithmic interpolation applied between these steps.

$$\mathcal{L}_d = \left\| D_{pseudo} - \hat{D} \right\|_1 \quad (13)$$

Furthermore, to enhance the recovery of distant details, we incorporate a depth-weighted reconstruction loss  $\mathcal{L}_{d_{rec}}$ :

$$\mathcal{L}_{d_{rec}} = \left\| \hat{D}_{detach} \cdot (\hat{I} - I) \right\|_1 \quad (14)$$

Finally, our total loss is formulated as follows:

$$\mathcal{L} = \mathcal{L}_{rec} + \lambda_D \mathcal{L}_{DCP} + \lambda_B \mathcal{L}_{BCP} + \lambda_d \mathcal{L}_d + \lambda_{d'} \mathcal{L}_{d_{rec}}, \quad (15)$$

where  $\lambda_D$ ,  $\lambda_B$ ,  $\lambda_d$  and  $\lambda_{d'}$  are trade-off weights.

## Experiments

### Experimental Setup

**Synthetic Foggy Dataset.** We chose four representative scenes from the Mip-NeRF (Barron et al. 2021) dataset, including two indoor scenes (bonsai and counter) and two outdoor scenes (garden and stump). For each scene, we employed the Depth Anything V2 model (Yang et al. 2024) to estimate ground-truth depths of original clear images. Subsequently, we stochastically sampled scattering coefficients and global atmospheric light for each scene. Finally, haze synthesis was performed on the images based on principles of the atmospheric scattering model. Additionally, we performed dehazing tests on a driving scene from the Waymo (Sun et al. 2020) dataset.

**Real Foggy Dataset.** We used three indoor foggy scenes provided by DehazeNeRF (Chen et al. 2023), captured with two professional fog machines and one mobile phone. The three scenes are bear, elephant, and lion. Each scene contains a certain number of clear views and foggy views. We combined the clear and foggy views from the same scene and input them into COLMAP to obtain the sparse point cloud and corresponding poses.

**Evaluation methods.** During the training process, we only take the multi-view foggy views of the scene as input to render clear views, and compute evaluation metrics (PSNR /

SSIM / LPIPS) using the rendered clear images and their corresponding real clear images.

**Baselines.** In addition to comparing with the original 3DGS (Kerbl et al. 2023), we also compared our method with ScatterNeRF (Ramazzina et al. 2023) and DehazeNeRF (Chen et al. 2023). We further evaluated the reconstruction performance of combining low-level computer vision dehazing models with 3DGS. Specifically, we first dehazed the multi-view images using the pretrained models of Dehazer (Guo et al. 2022) or DehazeFormer (Song et al. 2023) (as shown in the supplementary materials), which were then fed into 3DGS for reconstruction. Furthermore, since SeaSplat (Yang, Leonard, and Girdhar 2024) (a real-time underwater rendering method) also has some dehazing reconstruction capabilities, we included it in the comparison as well. Since DehazeNeRF has not released its source code, its qualitative and quantitative results are both sourced from its original paper—for the Lion and Bear scenes (where DehazeNeRF lacks quantitative data), we calculated such results for corresponding viewpoints for fair comparison. Note that the real dataset views in DehazeNeRF were specifically cropped (removing scene edges). For fair comparison, we similarly cropped edges of our rendered views.

**Training Details.** ScatterNeRF was trained on three 4090 GPUs with a batch size of 2048, while all other models were trained on a single 4090 GPU. Remaining parameters followed the default settings. For our model, training was performed on a single GPU. The number of iterations was set to 30K for synthetic datasets and 3K for real-world datasets.  $\lambda_D$ ,  $\lambda_B$ , and  $\lambda_{d'}$  were all set to 0.1. The learning rate for the convolutional network was set to  $1 \times 10^{-7}$ , with a weight decay factor of 0.1, decaying over the corresponding number of iterations for synthetic and real scenes.

### Quantitative and Qualitative Results

Table 1 and Table 2 present the quantitative comparison results on real hazy datasets and synthetic hazy datasets. Since DehazeNeRF does not provide quantitative metrics for the bear and lion scenes, we calculated them based on the rendered images presented in its paper. On real foggy datasets, our method outperforms existing methods in rendering details, with average improvements of 37.2% (LPIPS), 28.8% (SSIM), and 5.0% (PSNR) over the second-best method. For the synthetic foggy dataset, it achieves 24.3% (LPIPS), 1.9% (SSIM), and 4.7% (PSNR) gains over the same baseline.

The qualitative comparison results are shown in Fig. 3 and Fig. 4. SeaSplat’s dehazing results show noticeable color anomalies and missing edge regions. Compared to DehazeNeRF, our reconstruction results outperform in both dehazing quality and the details of near and far scenes. Meanwhile, ScatterNeRF fails to reconstruct clear scenes on both datasets. Additionally, our model is capable of achieving dehazing reconstruction in both indoor and outdoor complex scenarios. For additional rendering results, please refer to the supplementary materials.

### Ablation Study

In this section, we perform ablation studies on the various components of our framework to evaluate the role and con-




Figure 3: Qualitative comparison of novel view synthesis (input foggy images and render clear images under novel views unseen by the model) results on real datasets, the images of DehazeNeRF are taken from its original paper. We encourage readers to refer to the supplementary materials to view rendering results from more views.

Dataset Method/Metric	Bear			Elephant			Lion			Average			Train
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	
ScatterNeRF	10.34	0.351	0.728	11.01	0.431	0.727	9.92	0.320	0.763	10.42	0.367	0.739	> 16 hours
SeaSplat	12.13	0.586	<u>0.263</u>	11.09	0.539	0.314	12.16	0.530	0.285	11.79	0.552	0.287	~ 25 mins
3DGS	13.09	<u>0.597</u>	<u>0.276</u>	13.28	0.605	0.307	12.01	0.507	0.337	12.79	0.570	0.307	~ 5 mins
Dehamer+3DGS	13.88	0.393	0.315	13.36	0.444	0.368	13.50	0.431	0.328	13.58	0.423	0.337	~ 5 mins
DehazeNeRF	<u>14.92</u>	0.439	0.341	<u>17.87</u>	<u>0.730</u>	<u>0.150</u>	<b>18.01</b>	<u>0.558</u>	<u>0.259</u>	<u>16.93</u>	<u>0.576</u>	<u>0.250</u>	> 4 hours
Ours	<b>16.79</b>	<b>0.725</b>	<b>0.173</b>	<b>19.19</b>	<b>0.791</b>	<b>0.113</b>	<u>17.33</u>	<b>0.709</b>	<b>0.186</b>	<b>17.77</b>	<b>0.742</b>	<b>0.157</b>	~ 1.2 mins

Table 1: Quantitative comparison for novel view synthesis on the real foggy dataset.

tribution of each part. Specifically, we investigate the effects of depth regularization ( $\mathcal{L}_d$ ), depth-weighted ( $\mathcal{L}_{d_{rec}}$ ), and relevant physical priors ( $\mathcal{L}_{DCP}$ ,  $\mathcal{L}_{BCP}$ ) on rendering results. The quantitative results are presented in Table 3. Note that Vanilla 3DGS (Kerbl et al. 2023) refers to the original and unmodified 3DGS, without the ASM model. (the results

of its novel view synthesis are shown as 3DGS in Fig. 3 and Fig. 4). Starting from  $\mathcal{L}_{rec}$ , which corresponds to our framework with only the reconstruction loss retained and all other loss terms removed, we observe that the performance progressively improves as additional losses are incorporated.



Figure 4: Qualitative comparison of dehazing results on the synthetic foggy dataset. Our method exhibits finer texture details and is closer to the ground truth (GT) compared to existing methods.

Dataset Method/Metric	Garden			Bonsai			Counter			Stump			Average			Train
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	
3DGS	12.49	0.620	<u>0.255</u>	12.43	0.609	0.161	13.06	0.621	0.193	11.10	0.516	0.298	12.27	0.592	0.227	~ 10 mins
ScatterNeRF	14.15	0.346	0.609	18.47	0.492	0.586	18.32	0.506	0.633	11.82	0.224	0.771	15.69	0.392	0.649	> 20 hours
SeaSplat	18.42	<u>0.687</u>	0.275	<u>20.40</u>	<u>0.705</u>	0.121	<u>19.16</u>	<b>0.736</b>	0.183	<u>17.78</u>	<u>0.631</u>	<u>0.293</u>	18.94	0.689	0.218	~ 30 mins
Ours	<b>18.72</b>	<b>0.742</b>	<b>0.183</b>	<b>22.19</b>	<b>0.726</b>	<b>0.104</b>	<b>19.31</b>	<u>0.704</u>	<b>0.153</b>	<b>19.10</b>	<b>0.637</b>	<b>0.219</b>	<b>19.83</b>	<b>0.702</b>	<b>0.165</b>	~ 25 mins

Table 2: Quantitative comparison of rendering results on the synthetic fog dataset.




Figure 5: Qualitative ablation results show progressive improvements from vanilla GS to the complete model, with enhanced details and dehazing in distant and nearby scenes.

## Conclusion

We propose the first framework for learning clear 3D Gaussian splatting solely from multi-view hazy images. Our approach leverages explicit Gaussian representation, based on the atmospheric scattering model, to interpret the

	PSNR	SSIM	LPIPS
Vanilla 3DGS	12.27	0.591	0.227
$\mathcal{L}_{rec}$ (Physical Scattering Model)	19.35	0.675	0.189
$\mathcal{L}_{rec} + \mathcal{L}_{DCP}$	19.66	0.687	0.173
$\mathcal{L}_{rec} + \mathcal{L}_{DCP} + \mathcal{L}_{BCP}$	19.74	<u>0.697</u>	0.175
$\mathcal{L}_{rec} + \mathcal{L}_{DCP} + \mathcal{L}_{BCP} + \mathcal{L}_d$	<u>19.77</u>	0.692	0.166
All losses	<b>19.83</b>	<b>0.702</b>	<b>0.164</b>

Table 3: Ablation study on the synthetic foggy dataset.

formation of hazy images through a physically accurate forward-rendering process. This enables joint optimization of 3D scene representation while learning the participating medium. Overall, our method outperforms existing methods in terms of rendering quality, training speed, and rendering efficiency. However, in certain scenarios, the dehazing effect on distant scenes still needs further improvement. We hope that our approach can inspire future research in this field.

## Acknowledgements

This research was supported by the National Natural Science Foundation of China (62202076), the National Key R&D Program of China under Grant 2022ZD0160804, National Natural Science Foundation of China (U21A20515, 52238003, U22B2034, 82174224), Shenzhen S&T programme (No. CJGJZD20240729141906008).

## References

- Barron, J. T.; Mildenhall, B.; Tancik, M.; Hedman, P.; Martin-Brualla, R.; and Srinivasan, P. P. 2021. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5855–5864.
- Cai, B.; Xu, X.; Jia, K.; Qing, C.; and Tao, D. 2016. Dehazenet: An end-to-end system for single image haze removal. *IEEE transactions on image processing*, 25(11): 5187–5198.
- Chen, A.; Xu, Z.; Geiger, A.; Yu, J.; and Su, H. 2022. Tensorf: Tensorial radiance fields. In *European conference on computer vision*, 333–350. Springer.
- Chen, W.-T.; Yifan, W.; Kuo, S.-Y.; and Wetzstein, G. 2023. Dehazenerf: Multiple image haze removal and 3d shape reconstruction using neural radiance fields. *arXiv preprint arXiv:2303.11364*.
- Chen, Z.; Wang, Y.; Yang, Y.; and Liu, D. 2021. PSD: Principled synthetic-to-real dehazing guided by physical priors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7180–7189.
- Dong, H.; Pan, J.; Xiang, L.; Hu, Z.; Zhang, X.; Wang, F.; and Yang, M.-H. 2020. Multi-scale boosted dehazing network with dense feature fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2157–2167.
- Garbin, S. J.; Kowalski, M.; Johnson, M.; Shotton, J.; and Valentin, J. 2021. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14346–14355.
- Golts, A.; Freedman, D.; and Elad, M. 2019. Unsupervised single image dehazing using dark channel prior loss. *IEEE transactions on Image Processing*, 29: 2692–2701.
- Guo, C.-L.; Yan, Q.; Anwar, S.; Cong, R.; Ren, W.; and Li, C. 2022. Image dehazing transformer with transmission-aware 3d position embedding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5812–5820.
- He, K.; Sun, J.; and Tang, X. 2010. Single image haze removal using dark channel prior. *IEEE transactions on pattern analysis and machine intelligence*, 33(12): 2341–2353.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.*, 42(4): 139–1.
- Levoy, M. 1990. Efficient ray tracing of volume data. *ACM Transactions on Graphics (ToG)*, 9(3): 245–261.
- Li, B.; Peng, X.; Wang, Z.; Xu, J.; and Feng, D. 2017. Aod-net: All-in-one dehazing network. In *Proceedings of the IEEE international conference on computer vision*, 4770–4778.
- Li, T.; Li, L.; Wang, W.; and Feng, Z. 2023. Dehazing-NeRF: neural radiance fields from hazy images. *arXiv preprint arXiv:2304.11448*.
- Max, N. 1995. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2): 99–108.
- McCartney, E. 1976. Optics of the atmosphere: scattering by molecules and particles.
- 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.
- Narasimhan, S. G.; and Nayar, S. K. 2002. Vision and the atmosphere. *International journal of computer vision*, 48: 233–254.
- Qin, X.; Wang, Z.; Bai, Y.; Xie, X.; and Jia, H. 2020. FFA-Net: Feature fusion attention network for single image dehazing. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 11908–11915.
- Ramazzina, A.; Bijelic, M.; Walz, S.; Sanvito, A.; Scheuble, D.; and Heide, F. 2023. Scatternerf: Seeing through fog with physically-based inverse neural rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17957–17968.
- Song, Y.; He, Z.; Qian, H.; and Du, X. 2023. Vision transformers for single image dehazing. *IEEE Transactions on Image Processing*, 32: 1927–1941.
- Sun, P.; Kretschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; et al. 2020. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2446–2454.
- Sun, S.; and Guo, X. 2016. Image enhancement using bright channel prior. In *2016 International Conference on Industrial Informatics-Computing Technology, Intelligent Technology, Industrial Information Integration (ICIICII)*, 83–86. IEEE.
- Yang, D.; Leonard, J. J.; and Girdhar, Y. 2024. Seasplat: Representing underwater scenes with 3d gaussian splatting and a physically grounded image formation model. *arXiv preprint arXiv:2409.17345*.
- Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024. Depth Anything V2. *arXiv preprint arXiv:2406.09414*.
- Yang, Y.; Wang, C.; Liu, R.; Zhang, L.; Guo, X.; and Tao, D. 2022. Self-augmented unpaired image dehazing via density and depth decomposition. In *Proceedings of*

*the IEEE/CVF conference on computer vision and pattern recognition*, 2037–2046.

Yu, A.; Li, R.; Tancik, M.; Li, H.; Ng, R.; and Kanazawa, A. 2021. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5752–5761.

Yu, J.; Peng, X.; Lu, Z.; Kneip, L.; and Wang, Y. 2024. SpikeGS: Learning 3D Gaussian Fields from Continuous Spike Stream. In *Proceedings of the Asian Conference on Computer Vision*, 4280–4298.

Zhang, H.; and Patel, V. M. 2018. Densely connected pyramid dehazing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3194–3203.

Zhang, L.; Wang, S.; and Wang, X. 2021. Single image dehazing based on bright channel prior model and saliency analysis strategy. *IET Image Processing*, 15(5): 1023–1031.

Zhang, Y.; Zhou, S.; and Li, H. 2024. Depth Information Assisted Collaborative Mutual Promotion Network for Single Image Dehazing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2846–2855.

Zheng, Z.; Ren, W.; Cao, X.; Hu, X.; Wang, T.; Song, F.; and Jia, X. 2021. Ultra-high-definition image dehazing via multi-guided bilateral learning. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16180–16189. IEEE.

Zhu, Q.; Mai, J.; and Shao, L. 2015. A fast single image haze removal algorithm using color attenuation prior. *IEEE transactions on image processing*, 24(11): 3522–3533.

Zwicker, M.; Pfister, H.; Van Baar, J.; and Gross, M. 2001. EWA volume splatting. In *Proceedings Visualization, 2001. VIS'01.*, 29–538. IEEE.