
IBGS: Image-Based Gaussian Splatting

Hoang Chuong Nguyen¹ Wei Mao Jose M. Alvarez² Miaomiao Liu¹

¹Australian National University ²NVIDIA

hoangchuong.nguyen@anu.edu.au miaomiao.liu@anu.edu.au

wei.mao.research@gmail.com josea@nvidia.com

Abstract

3D Gaussian Splatting (3DGS) has recently emerged as a fast, high-quality method for novel view synthesis (NVS). However, its use of low-degree spherical harmonics limits its ability to capture spatially varying color and view-dependent effects such as specular highlights. Existing works augment Gaussians with either a global texture map, which struggles with complex scenes, or per-Gaussian texture maps, which introduces high storage overhead. We propose Image-Based Gaussian Splatting, an efficient alternative that leverages high-resolution source images for fine details and view-specific color modeling. Specifically, we model each pixel color as a combination of a base color from standard 3DGS rendering and a learned residual inferred from neighboring training images. This promotes accurate surface alignment and enables rendering images of high-frequency details and accurate view-dependent effects. Experiments on standard NVS benchmarks show that our method significantly outperforms prior Gaussian Splatting approaches in rendering quality, without increasing the storage footprint. Our project page is available at <https://hoangchuongnguyen.github.io/ibgs>.

1 Introduction

Recently, Neural Radiance Fields (NeRF) [23] and 3D Gaussian Splatting (3DGS) [16] have emerged as advanced techniques for Novel View Synthesis (NVS) thanks to their high-quality image rendering. 3DGS, in particular, outperforms NeRF in terms of rendering speed and faster optimization. However, since each Gaussian primitive in 3DGS can only represent a single color at a given camera viewpoint, 3DGS struggles to recover high-frequency details of the scene appearance without a large number of Gaussians [4]. Furthermore, due to the smooth characteristic of the color representation i.e., spherical harmonic (SH) functions, it is hard for 3DGS to capture complex view-dependent effects such as reflections and specular highlights [4].

To address this issue, recent works have attempted to model Gaussian’s spatially varying colors by either mapping each Gaussian-ray intersection to a global texture map [34], or a per-Gaussian texture map [4, 28, 33]. While global texture maps perform well for single-object scenes [34], they struggle with complex multi-object scenes due to the difficulty of learning a global mapping. Per-Gaussian texture map [4, 28, 33] can handle real-world scenes with multiple objects, but it incurs a significant storage overhead because the number of parameters per Gaussian grows quadratically as the texture-map resolution increases. This storage overhead constrains the resolution of per-Gaussian’s texture map, leading to inferior modeling of high-frequency details in the rendered images. Additionally, such texture map still cannot handle view-dependent effects.

In this work, we propose a drastically different approach to render high-frequency details and handle view-dependent effects while avoid significantly increasing the storage memory. Specifically, inspired by image-based rendering techniques [8, 22], we introduce an Image-Based Gaussian Splatting (IBGS) approach that utilizes the high-frequency details and view-dependent effects captured in training images. During rendering, the color of a pixel consists of two components i.e., the base color

from the SH functions following the standard rasterization process of 3DGS, and the residual color learnt from the corresponding pixel intensities of neighboring training images. The base color is used to handle most surface appearance while the color residual augments the base color with fine-grained details and view-dependent effects which are preserved in the training images. To model the color residuals, we propose a novel color residual prediction module. Specifically, for each ray/pixel, we first project the intersection points between the ray and Gaussians onto neighboring training images to obtain pixel colors, which are then aggregated to get warped colors. Then, the warped colors together with the base colors are processed by a lightweight neural network to predict the color residual for each pixel. We further introduce an image synthesis loss that leverages those warped colors, enforcing both geometric accuracy and image quality. This leads to more precise Gaussian parameters with high opacity centered around the true surface, allowing us to prune more Gaussians of low opacity while maintaining the rendering quality.

Furthermore, leveraging neighboring views enables our method to address inconsistent exposure across training views caused by the auto-exposure behavior of modern cameras. Different from prior work [6, 17] that jointly optimize an affine transformation matrix for each training view, we assume that images taken at nearby locations share similar global lighting conditions, and thus correct the camera exposure at novel views by mimicking the exposure setting of the closest source view. Unlike existing works [6, 17] that only correct the exposure at training views, our strategy can generalize to images rendered at any viewpoint.

In summary, our contributions are: i) We propose an image-based Gaussian Splatting pipeline that captures both high-frequency details and view-dependent effects that are challenging for prior methods to address. ii) We introduce a color residual module that leverages the training images to obtain better rendering quality with less number of Gaussians. iii) We introduce an exposure correction strategy, helping to improve the brightness of images rendered at any viewpoint by mimicking the exposure settings of their nearest neighbouring view. Our method sets a new state-of-the-art performance on three benchmark datasets: Tanks and Temples, Deep Blending, and Mip-NeRF360.

2 Related Works

Image-based rendering aims to generate novel views by “borrowing” pixels from a set of source images. The target pixel is a weighted blending of corresponding pixels obtained from those images. In early works [8, 11, 22], such blending weights are computed based on ray distance [22] or scene geometry [8]. Other works either tried to improve the scene geometry [5, 13] or use optical flow for better correspondence [1, 3, 9]. More recently, with the advance of neural rendering techniques [23], researchers have explored integrating it with image-based rendering [29, 31]. In particular, to render a target pixel/ray, Suhail et al. [29] first finds the corresponding epipolar lines in source views and sample points along such lines to obtain color features. These features are further fed into two feature aggregation modules subsequently for the final color. IBRNet [31] follows the volume rendering process as in NeRF [23]. The color and densities of the sampled points on target ray are computed by a transformer [30] with the features from source views as input. Despite their impressive results, their rendering is time-consuming due to the use of large feature aggregation networks (i.e., the transformers) and uniform sampling along the ray. By contrast, to the best of our knowledge, we propose the first image-based Gaussian splatting method that not only obtains fine-grained details from source images but also maintain fast rendering. Thanks to 3DGS, we only require projecting intersection points of ray with Gaussians, which are fairly sparse, to source views for aggregating image features. Moreover, instead of directly learning the final color from a large network, we propose to learn a residual to the base color, which only requires a lightweight network i.e., a nine-layer convolutional network with 3×3 with kernels.

Gaussian Splatting. 3DGS [16] renders images at novel-views by performing alpha blending of the Gaussians color splatted onto the image plane. Although each splatted Gaussian can have a large extent, its color is shared across all pixels, making it challenging for 3DGS to reconstruct fine-grained textures without using many Gaussians. To address this problem, prior works [4, 28, 33, 34] attempt to model Gaussian spatially varying color by learning a mapping from each Gaussian-ray intersection to a texture map. However, while learning a global texture map is challenging in scenes having multiple objects, learning per-Gaussian texture maps [4, 28, 33] leads to higher storage requirement since it requires to store texture maps of all Gaussians. Additionally, these methods struggle in recovering complex view-dependent color as they utilize low-degree SH functions which have limited capacity




Figure 1: Our pipeline. The color of each pixel (c^{final}) consists of two components: a base color c (in pink boxes) which follows the standard 3DGS rendering process and a color residual Δc predicted from the warped color of different source views c_m^{warp} . While the figure shows an example of using only two source views, in practice, our method can process an arbitrary number of source views.

to handle complex view-dependent colors [4]. Unlike these methods, we leverage information from training views to render photorealistic images by predicting high residuals for pixels whose color cannot be fully recovered via Gaussian rasterization (i.e., the base color), especially in regions with fine-grained details or view-dependent colors.

Apart from color modeling, existing works also target improving other aspects of Gaussian Splatting. In particular, [7, 12, 15] pursue more accurate 3D reconstruction by enforcing flat Gaussians via a hard [7, 15] or soft constraint [12]. On the other hand, [10, 25, 26] propose compression and quantization methods to reduce the memory requirements for storing optimized Gaussians while preserving rendering quality. [35] improves the adaptive density control strategy of 3DGS by exploiting the per-pixel gradient directions, whereas [18] views the Gaussian densification process as state transition of Markov Chain Monte Carlo samples. Our method is orthogonal to these works and can be integrated with them to further boost the performance or reduce memory usage.

3 Method

3.1 Preliminary: 3D Gaussian Splatting

3DGS [16] represents the scene as a set of 3D Gaussian primitives. Each 3D Gaussian $\mathcal{G}$ is parameterized by a 3D position $\mu \in \mathbb{R}^3$ and a covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$, where the covariance is decomposed into a rotation matrix $\mathbf{R} \in \text{SO}(3)$ and a diagonal scale matrix $\mathbf{S} \in \mathbb{R}^{3 \times 3}$ such that $\Sigma = \mathbf{R}\mathbf{S}\mathbf{S}^T\mathbf{R}^T$. To render an image at a viewpoint, each 3D Gaussian $\mathcal{G}$ is splatted on the image plane to obtain a 2D Gaussian $\mathcal{G}^{2D}$. The color of a pixel $\mathbf{p}$ can be then computed using the volume rendering equation,

$$c(\mathbf{p}) = \sum_{i=1}^K w_i c_i = \sum_{i=1}^K w_i \Psi_l(\mathbf{h}_i, \mathbf{v}_i), \quad \text{where } w_i = \alpha_i T_i, \quad \text{and } T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (1)$$

with $\mathbf{v}_i$ denoting the vector from the camera center to the center of the i^{th} Gaussian, $\Psi_l(\mathbf{h}, \mathbf{v})$ mapping the SH coefficients $\mathbf{h}$ of the Gaussian to a color c , conditioned on the direction $\mathbf{v}$. The scalar $l \in \mathbb{R}$ indicates the SH degree, which determines the expressivity of $\Psi_l(\mathbf{h}, \mathbf{v})$. The weight $\alpha_i = o_i \mathcal{G}_i^{2D}(\mathbf{p})$ is defined as the product of the Gaussian opacity o_i and 2D Gaussian value evaluated at pixel $\mathbf{p}$.

During training, the attributes of each Gaussian, including (1) its colors represented using spherical harmonic coefficients $\mathbf{h} \in \mathbb{R}^{3(l+1)^2}$, (2) the opacity $o \in \mathbb{R}$, (3) the center $\mu \in \mathbb{R}^3$, (4) the rotation matrix parameterized by a quaternion $\mathbf{q} \in \mathbb{R}^4$, (5) scaling factors $\mathbf{s} \in \mathbb{R}^3$ are optimized using the color rendering loss.

3.2 Modeling Spatially Varying and View-Dependent Color

In the color modeling of 3DGS [16], at a given viewpoint, although a Gaussian $\mathcal{G}_i$ can cover multiple pixels in the image, it represents only a single color, as the SH coefficients $\mathbf{h}$ and view direction $\mathbf{v}$ are shared across all pixels. This limits the model's capacity in modeling color of regions with high-frequency details. Moreover, due to the low-degree SH function ($l \leq 3$) utilized by 3DGS [16] to model view-dependent color, this method struggles to capture complex view-dependent effects, such as reflections or specular highlights. A naive solution is to increase the SH degree l , which quadratically increases the number of SH coefficients, thereby leading to high storage requirement.

In this work, we present a solution to model high-frequency details and view-dependent color of the image without increasing the storage footprint. Specifically, we propose IBGS, an image-based Gaussian splatting method that models (1) the base color from SH functions, (2) a color residual term capturing view-specific and high-frequency information from neighboring source images. We formulate it as follows:

$$\mathbf{c}^{\text{final}}(\mathbf{p}) = \underbrace{\sum_{i=1}^N w_i \Psi_l(\mathbf{h}_i, \mathbf{v}_i)}_{\text{Base color } \mathbf{c}(\mathbf{p})} + \underbrace{\mathcal{F}(\mathbf{c}(\mathbf{p}), \mathbf{d}(\mathbf{p}), \{\Delta \mathbf{c}_m(\mathbf{p})\}_{m=1}^M, \{\Delta \mathbf{d}_m(\mathbf{p})\}_{m=1}^M)}_{\text{Color residual } \Delta \mathbf{c}(\mathbf{p})}, \quad (2)$$

where $\mathbf{d}(\mathbf{p}) \in \mathbb{R}^3$ is the direction of the camera ray passing through pixel $\mathbf{p}$. $\Delta \mathbf{c}_m \in \mathbb{R}^3$ and $\Delta \mathbf{d}_m \in \mathbb{R}^4$ denotes appearance and camera features extracted from the m^{th} nearby source view, respectively. $\mathcal{F}(\cdot)$ is a lightweight network that takes the extracted features as input to predict a residual term $\Delta \mathbf{c}(\mathbf{p})$ supplementing the base color, $\mathbf{c}(\mathbf{p})$, produced by Gaussian rasterization.

Given the extracted multi-view features $\Delta \mathbf{c}_m$ that capture high-frequency details and color variation across different viewpoints, we use this features to predict the pixel color in the current view. By utilizing the multi-view color observations, our model learns how lighting effect changes across viewpoint, thus being able to produce accurate view-dependent color for the target view.

In the following sections, we first describe the two main components of our method: feature extraction from source views (Sec. 3.3) and color residual prediction (Sec. 3.4). Then, we present an exposure correction approach (Section. 3.5) to correct the exposure of the rendered images caused by inconsistent exposure camera setting. Fig. 1 depicts the overall pipeline of our method.

3.3 Feature Extraction from Source Views

To predict color residual, $\Delta \mathbf{c}(\mathbf{p})$, we first extract features from multiple source images. For a target pixel $\mathbf{p}$, we obtain the color information $\Delta \mathbf{c}_m(\mathbf{p})$ from each source image by first finding the intersection between (1) the camera ray originating at the camera center $\mathbf{o}$ with direction $\mathbf{d}(\mathbf{p})$ and (2) the plane parameterized by the Gaussian center μ_i and its normal vector $\mathbf{n}_i$,

$$\mathbf{x}_i(\mathbf{p}) = \mathbf{o} + \frac{\mathbf{n}_i^T (\mu_i - \mathbf{o})}{\mathbf{n}_i^T \mathbf{d}(\mathbf{p})} \mathbf{d}(\mathbf{p}). \quad (3)$$

Here we incorporate a normal vector $\mathbf{n}_i \in \mathbb{R}^3$ as an additional learnable attribute of each Gaussian $\mathcal{G}_i$. Given the intersection point $\mathbf{x}_i(\mathbf{p})$, we project it onto the image plane of nearby source views from which the color information is extracted. For the m^{th} source view, this is achieved as,

$$\mathbf{c}_{i,m}^{\text{warp}}(\mathbf{p}) = \mathcal{B}(\pi_m(\mathbf{x}_i(\mathbf{p})), \mathbf{C}_m^{\text{real}}) \quad (4)$$

with $\pi_m(\mathbf{x})$ denoting a function that projects the intersection point to the image plane of the source view. The function $\mathcal{B}(\cdot)$ takes an image coordinate as input to produce bilinear-interpolated color obtained from the input source image $\mathbf{C}_m^{\text{real}} \in \mathbb{R}^{H \times W \times 3}$.

Based on Eq. 4, the warped color $\mathbf{c}_{i,m}^{\text{warp}}(\mathbf{p})$ is only accurate if the intersection point $\mathbf{x}_i(\mathbf{p})$ is close to the actual surface. This implies that it is not necessary to project all the Gaussian-ray intersections to the source views, as floating Gaussians that are far from true surface introduce noise into the extracted appearance features. Following 2DGS [15], we consider the actual surface lies near the median intersections such that the accumulated transmittance T_i (see Eq. 1) is close to 0.5, and thus only project the K median intersection points $\{\mathbf{x}_{k,m}\}_{k=1}^K$ to the source views. As a result, we obtain a set of K warped colors $\{\mathbf{c}_{k,m}^{\text{warp}}\}_{k=1}^K$ for each source view m .

Subsequently, we compute the weighted average color for each source view and measure its deviation from the Gaussian-rasterized color $\mathbf{c}(\mathbf{p})$ (computed via Eq. 1) as below,

$$\Delta \mathbf{c}_m(\mathbf{p}) = \mathbf{c}_m^{\text{warp}}(\mathbf{p}) - \mathbf{c}(\mathbf{p}), \quad \mathbf{c}_m^{\text{warp}}(\mathbf{p}) = \sum_{k=1}^K \frac{w_k \mathbf{c}_{k,m}^{\text{warp}}(\mathbf{p})}{\sum_{k=1}^K w_k}, \quad (5)$$

where w_k is the same blending weight of the Gaussian colors computed from Eq. 1. Intuitively, $\mathbf{c}_m^{\text{warp}}(\mathbf{p})$ is used to approximate the true pixel color by leveraging the information from neighboring views. Thus, it constrains the weight w_k of the Gaussians near the true surface to be larger than those of others. Apart from the appearance features, we also compute $\Delta \mathbf{d}_m$, the difference in the camera position and orientation between the target and each source view,

$$\Delta \mathbf{d}_m = \begin{bmatrix} \mathbf{o}_m - \mathbf{o} \\ \mathbf{d}_m(\mathbf{p})^T \mathbf{d}(\mathbf{p}) \end{bmatrix}, \quad \mathbf{d}_m(\mathbf{p}) = \frac{\mathbf{x}(\mathbf{p}) - \mathbf{o}_m}{\|\mathbf{x}(\mathbf{p}) - \mathbf{o}_m\|_2}, \quad \mathbf{x}(\mathbf{p}) = \sum_{k=1}^K \frac{w_k \mathbf{x}_k(\mathbf{p})}{\sum_{k=1}^K w_k}, \quad (6)$$

with $\mathbf{o}_m$ being the camera center of the m^{th} source view. We repeat this feature extraction process for the M nearby source views, yielding a set of color features $\{\Delta \mathbf{c}_m\}_{m=1}^M$ and camera features $\{\Delta \mathbf{d}_m\}_{m=1}^M$ which are used as input to the color residual prediction network.

3.4 Color Residual Prediction

We employ a lightweight network to predict the color residuals, consisting of two main components: a per-pixel feature extractor $\mathcal{E}(\cdot)$ and a CNN decoder $\mathcal{D}(\cdot)$. The extractor has a PointNet-style structure [27] to handle an arbitrary number of source views M . For each view m , it processes the color feature $\Delta \mathbf{c}_m(\mathbf{p})$ and camera features $\Delta \mathbf{d}_m(\mathbf{p})$ through two linear layers of 32 output dimension followed by ReLU activation function,

$$\mathbf{f}_m(\mathbf{p}) = \mathcal{E}(\Delta \mathbf{c}_m(\mathbf{p}), \Delta \mathbf{d}_m(\mathbf{p})). \quad (7)$$

We then apply max-pooling to the set of vectors $\{\mathbf{f}_m(\mathbf{p})\}_{m=1}^M$ to obtain the aggregated feature $\bar{\mathbf{f}}(\mathbf{p}) \in \mathbb{R}^{32}$. Assembling these features across all pixels yields a feature map $\mathbf{F} \in \mathbb{R}^{H \times W \times 32}$. Similarly, by stacking $\mathbf{c}(\mathbf{p})$ and $\mathbf{d}(\mathbf{p})$ over all pixels, we form the Gaussian-rasterized image $\mathbf{C} \in \mathbb{R}^{H \times W \times 3}$ and the ray-direction map $\mathbf{D} \in \mathbb{R}^{H \times W \times 3}$, respectively. This information are then fed into a nine-layer convolutional decoder (with kernel size of 3) to predict the color residual map $\Delta \mathbf{C} \in \mathbb{R}^{H \times W \times 3}$.

$$\Delta \mathbf{C} = \mathcal{D}(\mathbf{C}, \mathbf{D}, \mathbf{F}). \quad (8)$$

Finally, we obtain the final image by adding the predicted residuals to the Gaussian-rasterized image,

$$\mathbf{C}^{\text{final}} = \mathbf{C} + \Delta \mathbf{C} \quad (9)$$

3.5 Exposure Correction

Due to varying lighting conditions at different locations, cameras with an auto-exposure setting can capture images with inconsistent brightness, which introduces noise into the optimization of the Gaussians. To stabilize the training, prior works [6, 17] optimize an color affine transformation matrix for each training view. This approach, however, can not generalize to correct the exposure of images rendered at novel views. To address this issue, we assume that images taken at nearby locations share similar global lighting conditions and thus propose to correct the exposure of the Gaussian-rasterized image by mimicking the exposure setting of the closest source view. In particular, we first obtain an affine transformation matrix $\mathbf{A}^*$ representing the exposure at the target view by solving the following least-square problem,

$$\mathbf{A}^* = \arg \min_{\mathbf{A} \in \mathbb{R}^{3 \times 4}} \sum_{\mathbf{p} \in \chi} \left\| \mathbf{A} \begin{bmatrix} \mathbf{c}(\mathbf{p}) \\ 1 \end{bmatrix} - \mathbf{c}_1^{\text{warp}}(\mathbf{p}) \right\|_2^2 \quad (10)$$

where χ is a set of pixels with valid coordinate when mapping to the source view, and $\mathbf{c}_1^{\text{warp}}(\mathbf{p})$ is color warped from the nearest source view. After that, we use $\mathbf{A}^*$ to correct the exposure of the rendered image as, $\mathbf{c}^{\text{expo}}(\mathbf{p}) = \mathbf{A}^* \begin{bmatrix} \mathbf{c}(\mathbf{p}) \\ 1 \end{bmatrix}$. Note that if exposure correction is applied, the color with corrected exposure $\mathbf{c}^{\text{expo}}(\mathbf{p})$ is used in place of the originally rendered color $\mathbf{c}(\mathbf{p})$ for computing appearance features (Eq. 5), residual prediction (Eq. 8) and obtaining the final image (Eq. 9).

3.6 Optimization

The overall loss function used to train our method is,

$$\mathcal{L} = \mathcal{L}_{\text{rgb}} + \lambda_1 \mathcal{L}_{\text{photo}} + \lambda_2 \mathcal{L}_{\text{normal}} \quad (11)$$

where λ_1, λ_2 are loss weights.

Color Rendering Loss $\mathcal{L}_{\text{rgb}}$. We compute the loss for both the final image and the Gaussian-rasterized image (i.e., the base image). The two terms are balanced by the weight γ as follows,

$$\mathcal{L}_{\text{rgb}} = \gamma \mathbb{L}(\mathbf{C}, \mathbf{C}^{\text{real}}) + (1 - \gamma) \mathbb{L}(\mathbf{C}^{\text{final}}, \mathbf{C}^{\text{real}}). \quad (12)$$

In particular, $\mathbb{L}$ is defined as

$$\mathbb{L}(\mathbf{C}, \mathbf{C}^{\text{gt}}) = \beta \mathcal{L}_1(\mathbf{C}, \mathbf{C}^{\text{gt}}) + (1 - \beta) \mathcal{L}_{\text{SSIM}}(\mathbf{C}, \mathbf{C}^{\text{gt}}). \quad (13)$$

where β is set to 0.8 and $\mathbf{C}, \mathbf{C}^{\text{gt}} \in \mathbb{R}^{H \times W \times 3}$ are the rendered and ground-truth image, respectively.

Multi-view Color Consistency Loss $\mathcal{L}_{\text{photo}}$. We also enforce photometric consistency across neighboring views to encourage accurate pixel matching,

$$\mathcal{L}_{\text{photo}} = \frac{1}{M} \sum_{m=1}^M \mathbb{L}(\mathbf{C}_m^{\text{warp}}, \mathbf{C}^{\text{real}}), \quad (14)$$

where $\mathbf{C}_m^{\text{warp}}$ is the warped image obtained by stacking all $\mathbf{c}_m^{\text{warp}}(\mathbf{p})$ computed in Eq. 5.

Normal Consistency Loss $\mathcal{L}_{\text{normal}}$. Following 2DGS [15], we apply the normal consistency loss to improve the overall geometry,

$$\mathcal{L}_{\text{normal}} = \frac{1}{|\Omega|} \sum_{\mathbf{p} \in \Omega} (1 - \mathbf{N}(\mathbf{p})^T \mathbf{N}_{\text{depth}}(\mathbf{p})) \quad (15)$$

with Ω being a set of all pixel coordinates and $\mathbf{N}$ denoting the rasterized normal map. $\mathbf{N}_{\text{depth}}$ is the normal map derived via finite difference of the point map $\mathbf{X}$ constructed from $\mathbf{x}(\mathbf{p})$ (see Eq. 6).

Visibility-based Source Views Selection. To find the nearby source views, we first compute the distance between the target and each source view, then use the closest S views as candidates to search for M visible source views (i.e., $M \leq S$). Specifically, for each pixel $\mathbf{p}$, we only use the s^{th} source view for feature extraction if it satisfies the following the condition,

$$\frac{|z(\mathbf{x}(\mathbf{p})) - z(\mathbf{x}_s^{\text{warp}}(\mathbf{p}))|}{z(\mathbf{x}(\mathbf{p})) + z(\mathbf{x}_s^{\text{warp}}(\mathbf{p}))} \leq \tau, \quad \text{with } \mathbf{x}_s^{\text{warp}}(\mathbf{p}) = \mathcal{B}(\pi_s(\mathbf{x}(\mathbf{p})), \tilde{\mathbf{X}}_s) \quad (16)$$

where τ is a depth error threshold, $z(\mathbf{x})$ denotes the depth value of the 3D point $\mathbf{x}$. The point map $\tilde{\mathbf{X}}_s$ can be obtained by transforming the point map $\mathbf{X}_s$ of source view to the coordinate system of the target view. Intuitively, this approach performs depth consistency check to exclude the source views in which the point $\mathbf{x}(\mathbf{p})$ is not visible.

4 Experiments

4.1 Experimental Setup

Dataset. Following 3DGS [16], we evaluate the NVS performance of our method using 2 scenes in the Tanks and Temples (TNT) [20] dataset, 2 scenes in the Deep Blending [14] dataset, and 9 scenes in the Mip-NeRF360 [2] dataset. We also show the results on 3 scenes in the Shiny dataset [32] which pose challenging view-dependent effects including specular highlight, reflection and disc diffraction. For all scenes, we use every 8th image for evaluation, and the rest for training.

Implementation Details. Similar to 3DGS [16], we train our method for 30,000 iterations. During the first 7,000 iterations, we set $\lambda_1 = \lambda_2 = 0$ and only activate the photometric and normal consistency loss thereafter with $\lambda_1 = 0.3$ and $\lambda_2 = 0.03$. The weight γ is initially set to 1, and then decreased to 0.5 during the last 20,000 iterations. Regarding hyper-parameters, we set SH degree $l = 2$, number of median intersection points $K = 4$, number of candidate source views $S = 4$, number of visible source views $M = 3$ and depth error threshold $\tau = 0.001$. We also prune Gaussians with opacity lower than 0.05. Following [6], we apply the exposure compensation from [6] and our proposed exposure correction method only to the TNT dataset. We use Adam optimizer [19] to train the residual prediction network. The initial learning rate is 0.001, which halves at iterations 18,000 and 25,000. All experiments are conducted using a single RTX 4090 GPU.

Table 1: Comparison between our method and previous works in three datasets. We measure the storage memory (Mem) in MB and the number of Gaussians (#Gauss) in millions. We report two released result from TexturedGaussian [4], with and without total memory usage.

Dataset Method Metric	Mip-NeRF 360					Tanks&Temples (TNT)					Deep blending				
	PSNR↑	SSIM↑	LPIPS↓	#Gauss	Mem	PSNR	SSIM	LPIPS	#Gauss	Mem	PSNR	SSIM	LPIPS	#Gauss	Mem
Mip-NeRF 360 [2]	<u>27.69</u>	0.792	0.237	×	8.6	22.22	0.759	0.257	×	8.6	29.40	0.901	0.245	×	8.6
Instant-NGP [24]	25.30	0.671	0.371	×	13	21.72	0.723	0.330	×	13	23.62	0.797	0.423	×	13
3DGS [16]	<u>27.69</u>	0.825	<u>0.203</u>	3.22	764	23.11	0.840	0.184	1.75	415	<u>29.53</u>	<u>0.904</u>	<u>0.242</u>	3.14	745
SuperGauss [33]	27.31	0.815	0.209	3.04	1021	23.72	0.847	0.179	1.50	502	28.83	0.901	0.250	2.27	762
TexturedGauss [4]	27.35	<u>0.827</u>	0.186	—	—	24.26	<u>0.854</u>	<u>0.168</u>	—	—	28.33	0.891	0.270	—	—
TexturedGauss* [4]	27.26	—	—	3.50	1047	<u>24.28</u>	—	—	1.30	691	28.52	—	—	1.00	668
Ours	28.33	0.837	0.186	1.59	291	24.84	0.869	0.148	0.75	143	30.12	0.912	0.237	1.11	197

Table 2: Comparison on three scenes in the Shiny dataset with challenging view-dependent colors.

Scene (effect) Method Metric	Guitars (specular highlight)					Lab (reflection)					CD (diffraction)				
	PSNR↑	SSIM↑	LPIPS↓	#Gauss	Mem	PSNR	SSIM	LPIPS	#Gauss	Mem	PSNR	SSIM	LPIPS	#Gauss	Mem
3DGS [16]	29.37	0.947	0.131	0.41	97	29.17	0.927	0.123	0.63	150	29.10	0.935	0.110	0.51	121
SuperGauss [33]	<u>30.43</u>	<u>0.952</u>	<u>0.121</u>	0.39	131	<u>29.38</u>	<u>0.932</u>	<u>0.107</u>	0.61	204	<u>29.49</u>	<u>0.944</u>	<u>0.091</u>	0.70	234
Ours	35.65	0.953	0.105	0.18	46	35.06	0.966	0.056	0.27	66	35.23	0.955	0.060	0.27	69

4.2 Results

In Tab. 1, we show the comparison of our method with prior methods in terms of NVS performance, number of Gaussians and storage usage. The results reveal that our method consistently achieve the best image quality across all datasets. For the PSNR metrics, our method gains an improvement of at least **0.64**, **0.56** and **0.59** dB in the Mip-NeRF 360, TNT and Deep blending datasets, respectively. Notably, on the Mip-NeRF 360 and TNT datasets, our method reduces the number of Gaussians and the storage by at least **62%** and **42%**, respectively, compared to existing Gaussian Splatting methods [4, 16, 33] and still outperforms them. For the Deep blending dataset, although we use slightly more Gaussians than TexturedGauss [4], our method consumes **70%** less storage. This is because TexturedGauss needs to store the texture maps of all Gaussians, which requires significantly more memory compared to storing the source images as in our method.

Tab. 2 presents the comparisons on three scenes in the Shiny [32] dataset with challenging view-dependent effects. For this dataset, we train 3DGS [16] and SuperGaussian [33] using their official implementations. Despite using fewer Gaussians, our method achieves significantly better NVS performance, with at least a **5.22** dB gain in PSNR. This demonstrates the superior capability of our method in modeling view-dependent color compared to previous works.

Qualitative Results. Fig. 2 shows the qualitative comparison between our method, 3DGS [16] and SuperGauss [33]. The results in the first two scenes reveal that 3DGS and SuperGauss struggle in reconstructing high-frequency details, while our method delivers more photorealistic results. We also show the two color components of our method, including a base image and a predicted residual map. While the base image alone exhibits the same detail deficiencies as 3DGS, adding our predicted color residuals helps to restore realistic textures. For scenes with complex view-dependent color, our method also demonstrates more compelling visual results, as shown in the last two scenes of Fig. 2. In these cases, 3DGS [16] and SuperGaussian [4] fail to capture the specular highlights and reflection effects, while our method successfully recovers the complex view-dependent colors in the final rendered images by predicting high residuals for these challenging regions. Interestingly, in the zoomed-in region of the Guitars scene, our method can decompose the color into a diffuse and specular component modeled via the base color and predicted residuals, respectively. Comparisons with more baseline methods and additional qualitative results can be found in the supplementary material.

4.3 Ablation Studies

Tab. 3 presents the ablation study results. **Base color only:** Discarding the predicted color residual leads to a significant drop in the image quality. This highlights the importance of the color residual module. **W/o color consistency loss $\mathcal{L}_{\text{photo}}$:** Training without the loss $\mathcal{L}_{\text{photo}}$ results in less accurate projections onto the source views, thereby reducing the quality of the rendered images. **Use source colors $\mathbf{c}_m$ as network’s input:** Here, we use the full colors $\mathbf{c}_m^{\text{warp}}$ obtained from the source images




Figure 2: Qualitative results. Our method can render images with both high-frequency details (first two scenes) and view-dependent effect (last two scenes). However this cannot be achieved by 3DGS [16] and SuperGaussian [33]

Table 3: Ablation study on the exposure correction, color consistency loss and the network’s input.

Method	Tanks&Temples (TNT)			MipNeRF-360		
	PSNR↑	SSIM↑	LPIPS↓	PSNR	SSIM	LPIPS
Full	24.84	0.869	0.148	28.33	0.837	0.186
Base color only	23.06	0.836	0.202	27.08	0.814	0.227
W/o color consistency loss $\mathcal{L}_{\text{photo}}$	<u>24.70</u>	0.866	0.152	<u>28.31</u>	0.833	0.192
Use source color $\mathbf{c}_m^{\text{warp}}$ as network’s input	24.61	<u>0.867</u>	<u>0.150</u>	28.21	0.837	<u>0.187</u>
W/o exposure correction	24.28	0.866	0.152	–	–	–

Table 4: Ablation study on opacity threshold.

Dataset Method (threshold)	Mip-NeRF 360					Tanks&Temples					Deep blending				
	PSNR↑	SSIM↑	LPIPS↓	#Gauss	Mem	PSNR	SSIM	LPIPS	#Gauss	Mem	PSNR	SSIM	LPIPS	#Gauss	Mem
3DGS (0.005)	27.69	0.825	0.203	3.22	764	23.11	0.840	0.184	1.75	415	29.53	0.904	0.242	3.14	745
Ours (0.005)	28.42	0.836	0.183	2.61	456	24.76	0.869	0.146	1.23	220	29.91	0.907	0.234	2.41	405
3DGS (0.05)	27.51	0.818	0.221	1.46	346	23.52	0.837	0.202	0.74	175	29.16	0.902	0.256	1.03	243
Ours (0.05)	28.33	0.837	0.186	1.59	291	24.84	0.869	0.148	0.75	143	30.12	0.912	0.237	1.11	197

as input to the residual prediction network, instead of their difference $\Delta \mathbf{c}_m$ from the base color. As a result, this approach performs consistently worse than our full model. **W/o exposure correction:** Removing the exposure correction also results in performance drop in the TNT dataset which exhibits inconsistent camera exposure across viewpoints [2]. Fig. 3 illustrates that our method can improve the exposure of the rendered image in both underexposure and overexposure cases, leading to higher image quality. More ablation studies can be found in the supplementary materials.

Additionally, we compare the sensitivity of our method and 3DGS [17] to the total number of Gaussians. For this experiment, we train our method and 3DGS with different opacity thresholds (0.005 and 0.05) used for pruning the Gaussians. Tab. 4 reveals that 3DGS requires a large number of low-opacity Gaussians to achieve good rendering quality, as its performance degrades when a higher opacity threshold is used. In contrast, with the same large threshold, our method can reduce the number of Gaussians while retaining most of the image quality.

5 Conclusion

In this paper, we present IBGS, an image-based Gaussian Splatting pipeline that is capable of rendering photorealistic images with both high-frequency details and view-dependent effects. Our key contribution is the color residual module, which leverages fine-grained textures and view-dependent information in nearby source images to predict a residual term added to the base Gaussian-rasterized color. Moreover, we introduce the exposure correction module to improve the brightness of the rendered image by mimicking the exposure of the closest source view. Extensive experimental results show that our method consistently outperforms previous works across different datasets.



Figure 3: Comparison of rendered images with and without exposure correction. Our method can correct the exposure in both underexposure (top) and overexposure (bottom) cases.

Limitations. Our method may struggle in a sparse-view setting, in which obtaining dense pixel correspondences used for residual prediction is challenging. Additionally, due to the additional computations in the rendering process, our method achieves lower rendering speed and requires higher runtime memory compared to 3DGS [16]. We discuss this in more detail in the supplementary material.

Broader Impacts. Our method has no immediate societal impacts. However, its downstream applications, such as 3D reconstruction [15] or controllable human modeling [21], can potentially be abused for malicious purposes such as unauthorized reconstructions, identity fraud.

Acknowledgement. This research was supported in part by the Australia Research Council ARC Discovery Grant (DP200102274).

References

- [1] Floating textures. In *Computer graphics forum*, pages 409–418. Wiley Online Library, 2008.
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5470–5479, 2022.
- [3] Dan Casas, Christian Richardt, John Collomosse, Christian Theobalt, and Adrian Hilton. 4d model flow: Precomputed appearance alignment for real-time 4d video interpolation. In *Computer Graphics Forum*, pages 173–182. Wiley Online Library, 2015.
- [4] Brian Chao, Hung-Yu Tseng, Lorenzo Porzi, Chen Gao, Tuotuo Li, Qinbo Li, Ayush Saraf, Jia-Bin Huang, Johannes Kopf, Gordon Wetzstein, et al. Textured gaussians for enhanced 3d scene appearance modeling. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.
- [5] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM transactions on graphics (TOG)*, 32(3):1–12, 2013.
- [6] Danpeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin Liu, Hujun Bao, and Guofeng Zhang. Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 2024.
- [7] Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. High-quality surface reconstruction using gaussian surfels. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024.
- [8] Paul E Debevec, Camillo J Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 465–474. 2023.
- [9] Ruofei Du, Ming Chuang, Wayne Chang, Hugues Hoppe, and Amitabh Varshney. Montage4d: interactive seamless fusion of multiview video textures. In *I3D*, pages 5–1, 2018.
- [10] Sharath Girish, Kamal Gupta, and Abhinav Shrivastava. Eagles: Efficient accelerated 3d gaussians with lightweight encodings. In *European Conference on Computer Vision*, pages 54–71. Springer, 2024.
- [11] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 453–464. 2023.
- [12] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5354–5363, 2024.
- [13] Peter Hedman, Tobias Ritschel, George Drettakis, and Gabriel Brostow. Scalable inside-out image-based rendering. *ACM Transactions on Graphics (TOG)*, 35(6):1–11, 2016.
- [14] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (ToG)*, 37(6):1–15, 2018.
- [15] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pages 1–11, 2024.
- [16] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

- [17] Bernhard Kerbl, Andreas Meuleman, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, and George Drettakis. A hierarchical 3d gaussian representation for real-time rendering of very large datasets. *ACM Transactions on Graphics (TOG)*, 43(4):1–15, 2024.
- [18] Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Yang-Che Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3d gaussian splatting as markov chain monte carlo. *Advances in Neural Information Processing Systems*, 37:80965–80986, 2024.
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [20] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017.
- [21] Muhammed Kocabas, Jen-Hao Rick Chang, James Gabriel, Oncel Tuzel, and Anurag Ranjan. Hugs: Human gaussian splats. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 505–515, 2024.
- [22] Marc Levoy and Pat Hanrahan. Light field rendering. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 441–452. 2023.
- [23] 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.
- [24] 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.
- [25] KL Navaneet, Kossar Pourahmadi Meibodi, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. Compgs: Smaller and faster gaussian splatting with vector quantization. In *European Conference on Computer Vision*, pages 330–349. Springer, 2024.
- [26] Simon Niedermayr, Josef Stumpffegger, and Rüdiger Westermann. Compressed 3d gaussian splatting for accelerated novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10349–10358, 2024.
- [27] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [28] Victor Rong, Jingxiang Chen, Sherwin Bahmani, Kiriakos N Kutulakos, and David B Lindell. Gstex: Per-primitive texturing of 2d gaussian splatting for decoupled appearance and geometry modeling. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3508–3518. IEEE, 2025.
- [29] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Light field neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8269–8279, 2022.
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [31] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2021.
- [32] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8534–8543, 2021.
- [33] Rui Xu, Wenyue Chen, Jiepeng Wang, Yuan Liu, Peng Wang, Lin Gao, Shiqing Xin, Taku Komura, Xin Li, and Wenping Wang. Supergaussians: Enhancing gaussian splatting using primitives with spatially varying colors. *arXiv preprint arXiv:2411.18966*, 2024.
- [34] Tian-Xing Xu, Wenbo Hu, Yu-Kun Lai, Ying Shan, and Song-Hai Zhang. Texture-gs: Disentangling the geometry and texture for 3d gaussian splatting editing. In *European Conference on Computer Vision*, pages 37–53. Springer, 2024.

- [35] Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. Absgs: Recovering fine details in 3d gaussian splatting. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 1053–1061, 2024.

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: Our claims made in the abstract and introduction accurately reflect the paper's contributions and scope. We further demonstrate it in our methodology and experiment sections.

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: Yes, we discuss the limitation of our method in the main paper.

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: The paper does not include theoretical results.

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: We provide all the information needed to reproduce the main experimental results in the main paper, supplementary materials

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: Our code will be released upon acceptance.

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, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide experimental details in the main paper and supplementary material.

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: This paper does not report error bars.

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: We provide experimental details, including the computer resources needed to reproduce the experimental results in the main paper and supplementary material.

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: This work conforms with 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: We provide discussion on the potential societal impacts of our method in the main paper.

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: The paper poses no such risks.

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: We have cited the original owners of all assets used in this paper properly.

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: [NA]

Justification: That the paper does not release 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: The paper does not involve crowdsourcing nor research with human subjects.

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: The paper does not involve crowdsourcing nor research with human subjects.

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.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.