

FantasyStyle: Controllable Stylized Distillation for 3D Gaussian Splatting

Yitong Yang¹, Yinglin Wang^{1*}, Changshuo Wang³, Huajie Wang^{4,5}, Shuting He^{1,2*}

¹School of Computing and Artificial Intelligence, Shanghai University of Finance and Economics, Shanghai, China

²MoE Key Laboratory of Interdisciplinary Research of Computation and Economics, Shanghai, China

³Department of Computer Science University College London, London, United Kingdom

⁴Shandong University of Finance and Economics, Shandong, China

⁵Jinan Yunwei Software Technology Co., Ltd, Shandong, China

yangyitong@stu.sufe.edu.cn, wang.yinglin@shufe.edu.cn, wangchangshuo1@gmail.com, wanghuajie@sdufe.edu.cn, shuting.he@sufe.edu.cn

Abstract

The success of 3DGS in generative and editing applications has sparked growing interest in 3DGS-based style transfer. However, current methods still face two major challenges: (1) multi-view inconsistency often leads to style conflicts, resulting in appearance smoothing and distortion; and (2) heavy reliance on VGG features, which struggle to disentangle style and content from style images, often causing content leakage and excessive stylization. To tackle these issues, we introduce **FantasyStyle**, a 3DGS-based style transfer framework, and the first to rely entirely on diffusion model distillation. It comprises two key components: (1) **Multi-View Frequency Consistency**. We enhance cross-view consistency by applying a 3D filter to multi-view noisy latent, selectively reducing low-frequency components to mitigate stylized prior conflicts. (2) **Controllable Stylized Distillation**. To suppress content leakage from style images, we introduce negative guidance to exclude undesired content. In addition, we identify the limitations of Score Distillation Sampling and Delta Denoising Score in 3D style transfer and remove the reconstruction term accordingly. Building on these insights, we propose a controllable stylized distillation that leverages negative guidance to more effectively optimize the 3D Gaussians. Extensive experiments demonstrate that our method consistently outperforms state-of-the-art approaches, achieving higher stylization quality and visual realism across various scenes and styles.

1 Introduction

With the widespread adoption of 3D content (He et al. 2025; Ding et al. 2025; Wang et al. 2025a,b) in virtual reality (VR), augmented reality (AR), and related fields, user expectations for visual quality have shifted from traditional photorealism toward more artistic and personalized styles. In this context, style transfer has emerged as a key technique for altering the visual appearance of 3D content, attracting increasing attention from the research community. Previous works (Zhang et al. 2022; Jung et al. 2024; Zhang, Fernandez-Labrador, and Schroers 2024; Zhang et al. 2023a) have explored 3D style transfer based on Neural Radiance Fields (NeRF) (Mildenhall et al. 2021), typi-

*Corresponding author.

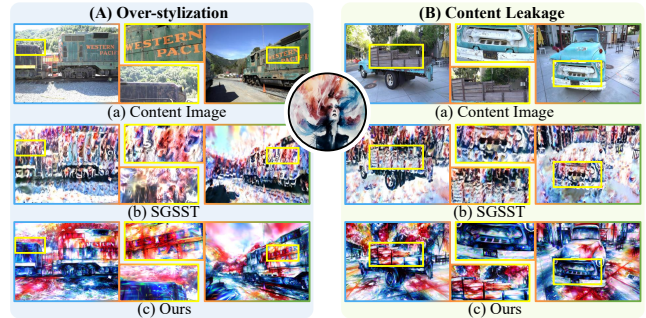


Figure 1: Limitations of the previous works. VGG-based methods often cause over-stylization and content leakage from the style image. In contrast, our approach ensures faithful style transfer and content preservation.

cally by introducing VGG-based stylization losses to ensure both style consistency and structural coherence in rendered views. However, NeRF suffers from slow rendering and long training times, limiting its practical applications. In contrast, 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023), as an emerging 3D representation, has become a research focus for 3D style transfer due to its faster rendering and superior visual quality.

Most recent studies on style transfer for 3DGS are based on VGG features, drawing inspiration from the early work on feature statistic matching (Gatys, Ecker, and Bethge 2016). These methods (Zhang et al. 2024; Saroha et al. 2024a; Kovács, Hermosilla, and Raidou 2024) typically extract VGG features and compute Gram matrices to achieve texture style transfer. Meanwhile, some studies (Zuo et al. 2024; Gu et al. 2024) attempt to leverage diffusion models to generate stylized 2D images, which are then further optimized back into the 3D scene to enable style transfer. Although these methods have shown promising results, they still face two critical challenges. (1) **Multi-view inconsistency**. Using 2D diffusion priors to guide 3D stylization often leads to inconsistent appearances across views, especially between adjacent ones. These inconsistencies can lead to style conflicts, hinder optimization, and distort geometry and appearance, as shown in Fig. 4 (b). (2) **Content leakage**

from style images and over-stylization. As shown in Fig. 1 (A) (B), VGG-based methods struggle to separate content and style, often causing unintended content transfer from the style image. Their emphasis on low-level textures often results in over-stylization, which obscures structural details.

To address the aforementioned challenges, we propose **FantasyStyle**, a 3DGS-based style transfer framework. Drawing inspiration from FreeU (Si et al. 2024) and FreeInit (Wu et al. 2024), which emphasize the critical role of frequency components in image and video generation, we perform a frequency-domain analysis of 2D multi-view stylized priors. We observe that low-frequency components primarily reflect view-dependent local details and tend to exhibit poor consistency across different viewpoints. In contrast, high-frequency components more stably capture texture features and demonstrate better consistency across views. Therefore, appropriately attenuating the low-frequency components can help alleviate inter-view discrepancies and improve overall multi-view consistency. Based on this observation, we introduce **Multi-View Frequency Consistency (MVFC)**, which applies a 3D frequency-domain filter to the multi-view latent after DDIM noise injection. MVFC preserves all high-frequency components and selectively suppresses low-frequency components, effectively reducing stylized prior conflicts and improving both consistency and stylization quality.

In parallel, to eliminate potential content leakage from the style image, we incorporate negative guidance during the denoising phase, ensuring that the derived 2D stylized prior remains free of content information. While methods (Chung et al. 2023) such as Score Distillation Sampling (SDS) and Delta Denoising Score (DDS) have been employed for 3D scene optimization, they are not directly applicable in our context, as they tend to produce overly smooth images that lose important brushstroke details present in style images. To overcome this limitation, we remove the reconstruction term from these objectives and incorporate negative guidance to design a **Controllable Stylized Distillation (CSD)**, which leverages 2D stylized priors to facilitate faithful and structurally coherent style transfer in 3D scenes.

Unlike existing approaches that primarily rely on VGG-based features, to the best of our knowledge, our method is the first 3DGS-based style transfer framework that relies solely on diffusion model distillation. It enables flexible extension of 2D style transfer techniques to 3D scenes, bridging the gap between 2D and 3D stylization. We evaluate our method through both qualitative and quantitative comparisons against baseline approaches, demonstrating superior performance in style transfer quality and content preservation. In summary, our key contributions are as follows:

- We propose a Multi-View Frequency Consistency that enhances cross-view consistency by preserving high-frequency components and selectively suppressing low-frequency components.
- We propose a Controllable Stylized Distillation that prevents content leakage from 2D stylized priors and optimizes 3D scenes.
- Extensive experiments demonstrate that our method out-

performs existing state-of-the-art approaches in both qualitative and quantitative evaluations.

- To the best of our knowledge, FantasyStyle is the first 3DGS-based style transfer method that relies entirely on diffusion model distillation. It enables the flexible extension of 2D style transfer techniques to 3D domains.

2 Related Work

2D Style Transfer Based on Diffusion Models. Style transfer aims to separate style and content representations from a style image and apply the style to another image. With the rise of diffusion models, 2D style transfer has seen notable advances. Recent studies leverage cross-attention (Zhang et al. 2023b) and self-attention mechanisms (Hertz et al. 2024; Chung, Hyun, and Heo 2024) to inject style features into the generation process. In addition, frequency-domain techniques (Xu et al. 2025; Gao et al. 2024) have been incorporated into the diffusion process, effectively enhancing the model’s ability to control stylistic attributes. More recently, methods (Frenkel et al. 2024; Ouyang, Li, and Hou 2025) based on Low-Rank Adaptation (LoRA) (Hu et al. 2022) have been proposed to achieve a more effective separation of content and style features, thereby improving the flexibility and diversity of style transfer. We propose a Controllable Stylized Distillation approach that utilizes 2D stylized priors to optimize 3D scenes, effectively bridging the gap between 2D and 3D style transfer.

3DGS Style Transfer Based on VGG. 3D style transfer is a conditional generative task that applies the visual style of a 2D image to a 3D representation while preserving the original geometric structure. The VGG network, one of the earliest architectures used in style transfer, remains a widely adopted feature extractor in both 2D and 3D tasks. Saroha et al. (Saroha et al. 2024b) employ VGG and AdaIN to transfer styles to specific views. StyleSplat (Jain et al. 2024) introduces a nearest neighbor feature matching loss between VGG features of rendered and reference images. SGSST (Galerie et al. 2024) proposes a Simultaneously Optimized Scales loss to enable ultra-high-resolution 3D stylization. StyleGaussian (Liu et al. 2024) embeds 2D VGG scene features into transformed 3D Gaussian features, which are fused with reference features and decoded into stylized RGB outputs for more coherent, view-consistent stylization. However, these methods often suffer from content leakage from the style image. To address this, we propose a negative guidance mechanism to suppress the transfer of irrelevant content information from the style image.

3DGS Style Transfer Based on Diffusion Model. With the powerful generative capacity of text-to-image (T2I) diffusion models, recent 3D style transfer methods increasingly incorporate pre-trained diffusion priors to guide stylization. InstantStyleGaussian (Yu et al. 2024) accelerates the process by generating target-style images via diffusion and integrating them into the training set to iteratively optimize the 3D Gaussian Splatting scene. ArtNVG (Gu et al. 2024) enhances content-style disentanglement using CSGO (Xing et al. 2024) and Tile-ControlNet (Zhang, Rao, and Agrawala 2023), and introduces an attention-based view

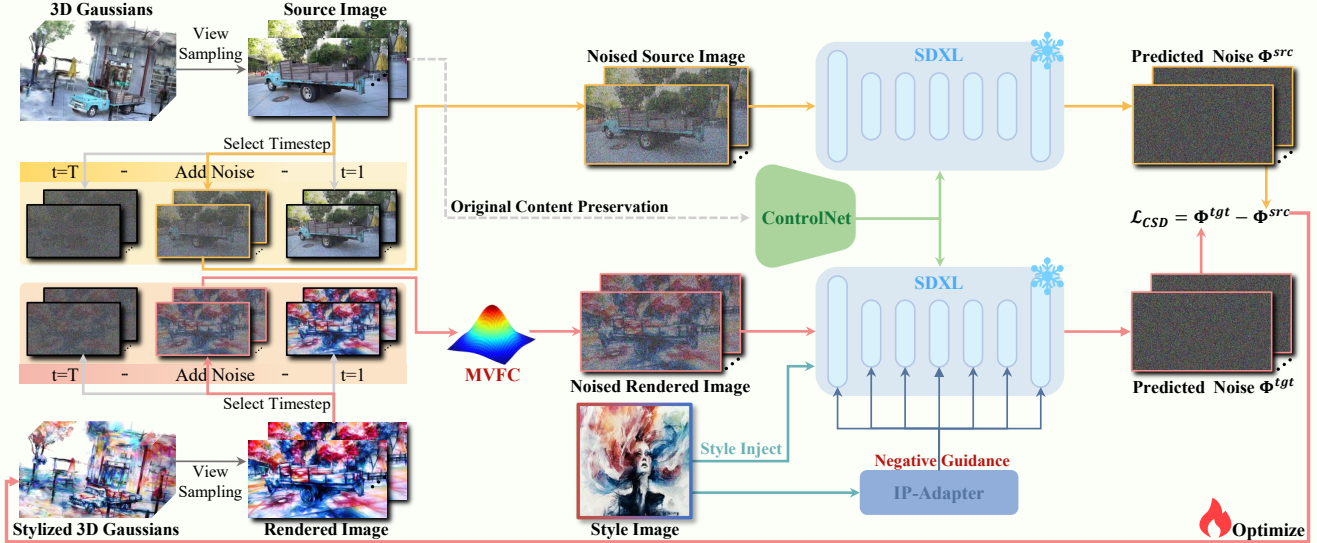


Figure 2: Overview of our proposed method. Inspired by DDS (Hertz, Aberman, and Cohen-Or 2023), FantasyStyle consists of two distinct pathways: Source Image and Rendered Image, each highlighted in different colors. We propose Controllable Stylized Distillation (CSD) to optimize the 3D scene. In Rendered Image pathway, we introduce a multi-view frequency consistency (MVFC) and inject style image features to obtain a multi-view consistent 2D stylized prior. Additionally, we incorporate negative guidance to suppress potential content leakage from the style image.

alignment module for consistent local texture and color. StyleMe3D (Zhuang et al. 2025) proposes dynamic style score distillation in the latent space of Stable Diffusion for improved semantic alignment. These methods, when guided by 2D images for 3D scene stylization, do not explicitly enforce multi-view consistency. To address this, we propose a Multi-View Frequency Consistency to improve consistency across views.

3 Preliminary

Score Distillation Sampling (SDS) (Poole et al. 2022) is widely used in 3D generation, leveraging pretrained diffusion models with rich 2D priors to distill knowledge into 3DGS models. Ignoring the UNet Jacobian, it is formulated as:

$$\nabla_{\theta} \mathcal{L}_{\text{SDS}} = \mathbb{E}_{t, \epsilon} \left[\omega(t) (\epsilon_{\phi}(z_t | \mathcal{P}) - \epsilon) \frac{\partial z_t}{\partial \theta} \right], \quad (1)$$

where $\omega(t)$ denotes a time-dependent weighting function, $\epsilon \sim \mathcal{N}(0, \mathbf{I})$, $z_t = \sqrt{\alpha_t} \mathbf{x} + \sqrt{1 - \alpha_t} \epsilon$ represents the noised input at timestep t , and $\epsilon_{\phi}(z_t | y)$ is the predicted denoising score conditioned on the text prompt \mathcal{P} .

Delta Denoising Score (DDS) (Hertz, Aberman, and Cohen-Or 2023) extends SDS for editing tasks by estimating the score difference between target and source distributions to guide updates. This mitigates the noisy gradients of SDS that often lead to blurred or unclear images. It is defined as:

$$\nabla_{\theta} \mathcal{L}_{\text{DDS}} = \mathbb{E}_{t, \epsilon} \left[\omega(t) (\epsilon_{\phi}(z_t^{\text{tgt}}, \mathcal{P}^{\text{tgt}}, t) - \epsilon_{\phi}(z_t^{\text{src}}, \mathcal{P}^{\text{src}}, t)) \frac{\partial z_t^{\text{tgt}}}{\partial \theta} \right], \quad (2)$$

where z_t^{tgt} and z_t^{src} denote the latent representations at timestep t , sharing the same noise ϵ_t .

Task Definition. The 3DGS reconstruction process can be formally defined as follows:

$$\min_{\Theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{L1+SSIM}(\mathcal{R}(C_i; \Theta), V_i^{\text{gt}}), \quad (3)$$

where C_i denotes the i -th camera, and $\Theta = \{(u_j, \Sigma_j, \alpha_j, c_{j,0}, (c_{j,k,l})_{k,l})\}_{j=1}^M$ represents the set of 3D Gaussian scene parameters. $c_{i,0}$ corresponds to the spherical harmonics (SH) basis functions, while $c_{j,k,l}$ are the SH coefficients that control the view-dependent color appearance. $\mathcal{R}(\cdot)$ denotes the rasterization-based rendering process that produces the corresponding view, and V_i^{gt} refers to the ground truth image for camera C_i .

The objective of style transfer is to modify the appearance of a scene without altering its underlying geometry. Accordingly, we reformulate Eq. 3 by optimizing only the color-related terms while keeping other parameters fixed. This can be formally expressed as:

$$\min_{\Theta_c} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(\mathcal{R}(C_i; \Theta_c); P, I_{\text{style}}), \quad (4)$$

where P denotes the prompt describing the original scene, and I_{style} represents the reference style image.

4 Method

We present the proposed FantasyStyle framework in detail (see Fig. 2). To address the issue of multi-view inconsistency, we apply a Multi-View Frequency Consistency (Sec. 4.1) on the DDIM-noised multi-view latent,

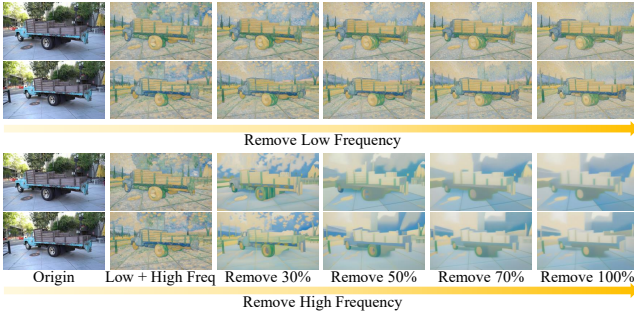


Figure 3: The role of different frequency components. For intuitive visualization, we present multi-view 2D stylization results. We observe that selectively removing low-frequency components slightly reduces local detail while significantly improving multi-view consistency, whereas removing high-frequency components severely degrades texture features, resulting in blurred appearances.

which enhances view consistency and mitigates conflicts arising from inconsistent stylized priors. Furthermore, we introduce a Controllable Stylized Distillation (Sec. 4.2) that suppresses content information from the style image and leverages 2D stylized priors to optimize the 3D scene, enabling effective 3D style transfer.

4.1 Multi-View Frequency Consistency

3D Frequency Analysis. Inspired by FreeU (Si et al. 2024) and FreeInit (Wu et al. 2024), which emphasize the critical role of frequency components in image and video generation, we perform a frequency-domain analysis of the multi-view stylized priors to better understand their behavior. To visualize more intuitively, we use IP-Adapter (Ye et al. 2023) to generate 2D multi-view stylized images. Moreover, since our task only involves fine-tuning color parameters to achieve texture-level style consistency without the need to preserve geometric structure, we introduce ControlNet to help guide the structural information in the images. Given N rendered views from a 3D scene, we first apply DDIM (Ho, Jain, and Abbeel 2020) to add noise and obtain z_t :

$$z_t^N = \sqrt{\bar{\alpha}_t} z_0^N + \sqrt{1 - \bar{\alpha}_t} \epsilon^N, \quad (5)$$

where $\bar{\alpha}_t$ is a constant, and $\epsilon^N \sim \mathcal{N}(0, \mathbf{I})$ represents Gaussian white noise. We investigate the impact of different frequency components on style consistency across 3D views by decomposing z_t into its low-frequency and high-frequency components, as formulated below:

$$\begin{aligned} F_{z_t}^L &= FFT_{3D}(z_t^N) \odot (1 - \mathcal{H}), \\ F_{z_t}^H &= FFT_{3D}(z_t^N) \odot \mathcal{H}, \\ z_t^{\prime N} &= IFFT_{3D}(\alpha * F_{z_t}^L + F_{z_t}^H), \quad \text{where } \alpha \in [0, 1] \\ z_t^{\prime\prime N} &= IFFT_{3D}(F_{z_t}^L + \alpha * F_{z_t}^H), \quad \text{where } \alpha \in [0, 1] \end{aligned} \quad (6)$$

where FFT_{3D} denotes the Fast Fourier Transform applied to both multi-view spatial and batch dimensions, $IFFT_{3D}$ is its inverse, and \mathcal{H} is a High Pass Filter (e.g., Gaussian).




Figure 4: Visual results of MVFC on multi-view 2D stylization. Without MVFC, inconsistent stylization conflicts arise across views.

\odot denotes element-wise multiplication. We control the parameter α to selectively remove specific frequency components, and the corresponding results are shown in Fig. 3. We observe that removing low-frequency components does not significantly affect stylization performance. Moderately reducing low-frequency components slightly diminishes local details but significantly improves multi-view consistency. However, when low-frequency components are overly suppressed, the image loses most local details and appears noticeably brighter. In contrast, as high-frequency components are progressively removed, the images lose fine-grained texture details and become overly smooth.

Multi-View Frequency Control. Based on the above observations, we propose a Multi-View Frequency Consistency (MVFC) mechanism to address view inconsistency in 3D style transfer. Specifically, we retain all high-frequency components to preserve texture details, while selectively preserving the low-frequency components to balance local details and multi-view consistency. Additionally, we introduce a shared low-frequency component of random Gaussian noise across all views to explicitly enforce consistency. The mathematical formulation is as follows:

$$\begin{aligned} F_{\epsilon'}^L &= FFT_{3D}(\epsilon') \odot (1 - \mathcal{H}), \quad \epsilon' \sim \mathcal{N}(0, \mathbf{I}) \\ \hat{z}_t^N &= IFFT_{3D}(\gamma * F_{z_t}^L + (1 - \gamma) * F_{\epsilon'}^L + F_{z_t}^H), \end{aligned} \quad (7)$$

where γ is a balancing factor. We use 2D stylized images to visually demonstrate the issue of view inconsistency, as shown in Fig. 4. Without MVFC, significant style conflicts occur across different views. In contrast, incorporating MVFC substantially improves multi-view style consistency, enabling us to obtain coherent 2D stylized priors that can guide the stylization of 3D scenes through subsequent Controllable Stylized Distillation.

4.2 Controllable Stylized Distillation

Style Injection. As shown in Fig. 2, we build upon Stable Diffusion XL (SDXL) and incorporate IP-Adapter for style injection, aiming to obtain 2D stylized priors. Although the 3D scene’s geometry remains fixed during optimization, generating 2D stylized priors often causes geometric information loss, leading to structural inconsistencies across views. To address this, we incorporate ControlNet (Zhang, Rao, and Agrawala 2023) to explicitly guide structural features during stylization, improving the geometric consistency of the 2D stylized priors.

Negative Guidance. The 2D stylized priors obtained through the above method, similar to previous approaches based on VGG features, often suffer from inevitable content leakage from the style image, leading to suboptimal performance in 3D style transfer. Inspired by recent advances in 2D image generation, we introduce negative prompts into the Classifier-Free Guidance (CFG) mechanism to suppress content interference from the style image. The formulation of the classifier-free guidance is as follows:

$$\tilde{\epsilon}_\phi(z_t, t, \mathcal{P}, \emptyset) = \epsilon_\phi(z_t, t, \emptyset) + \beta (\epsilon_\phi(z_t, t, \mathcal{P}) - \epsilon_\phi(z_t, t, \emptyset)), \quad (8)$$

where \mathcal{P} denotes the text prompt, \emptyset represents the null-text embedding, and β is the CFG scale. To suppress specific content information, \emptyset is replaced with the negative prompt \mathcal{P}_{neg} , thereby enabling negative guidance. Negative prompts can be generated using large multimodal language models. For convenience, we leverage the IPAdapter-Instruct (IP) (Rowles et al. 2024) to extract both style and content features from the style image, which are then used to construct the negative prompt. Accordingly, Eq. 8 is reformulated as follows:

$$\hat{\epsilon}_\phi(z_t, t, \mathcal{P}, \text{IP}(I_r)^s, \text{IP}(I_r)^c) = \epsilon_\phi(z_t, t, \text{IP}(I_r)^c) + \beta (\epsilon_\phi(z_t, t, [\mathcal{P}, \text{IP}(I_r)^s]) - \epsilon_\phi(z_t, t, \text{IP}(I_r)^c)), \quad (9)$$

where $\text{IP}(I_r)^s$ and $\text{IP}(I_r)^c$ denote the content and style features extracted from the reference style image, respectively, and $[\cdot, \cdot]$ represents the concatenation operation.

Controllable Stylized Distillation. Directly applying SDS or DDS to optimize 3D scenes using 2D stylized priors results in overly smooth and blurry images, losing crucial brushstroke details essential for style transfer, as shown in Fig. 9. To better understand this limitation, we define $\delta_{z_t}^{\text{SDS}} := \epsilon_\phi(z_t, t, \mathcal{P}) - \epsilon$ in Eq. 1 and $\delta_{z_t}^{\text{DDS}} := \epsilon_\phi(z_t^{\text{tgt}}, \mathcal{P}^{\text{tgt}}, t) - \epsilon_\phi(z_t^{\text{src}}, \mathcal{P}^{\text{src}}, t)$ in Eq. 2, and extend it based on Eq. 8 as follows:

$$\begin{aligned} \delta_{z_t}^{\text{SDS}} &:= \underbrace{\epsilon_\phi(z_t, t, \emptyset) - \epsilon}_{\delta_{z_t}^{\text{recon}}} + \beta \underbrace{(\epsilon_\phi(z_t, t, \mathcal{P}) - \epsilon_\phi(z_t, t, \emptyset))}_{\delta_{z_t}^{\text{cfg}}}, \\ \delta_{z_t}^{\text{DDS}} &:= \underbrace{\epsilon_\phi(z_t^{\text{tgt}}, t, \emptyset) - \epsilon_\phi(z_t^{\text{src}}, t, \emptyset)}_{\delta_{z_t}^{\text{recon}}} + \beta (\delta_{z_t}^{\text{cfg}} - \delta_{z_t^{\text{src}}}^{\text{cfg}}). \end{aligned} \quad (10)$$

The two formulations above are highly similar and can be decoupled into a reconstruction term $\delta_{z_t}^{\text{recon}}$ and a classifier-free guidance term $\delta_{z_t}^{\text{cfg}}$. Since our task involves only fine-tuning color-related parameters without considering geometric structure or identity preservation, the reconstruction term becomes a limiting factor. It not only leads to overly smooth and blurry results but also significantly slows down the optimization process. To address this issue, we remove the reconstruction term and introduce negative guidance. Building upon Eq. 2, Eq. 9, and Eq. 10, we propose Controllable Stylized Distillation (CSD), formulated as follows:

$$\begin{aligned} \Phi^{\text{tgt}} &= \beta (\epsilon_\phi(z_t^{\text{tgt}}, t, [\mathcal{P}, \text{IP}(I_r)^s]) - \epsilon_\phi(z_t^{\text{tgt}}, t, \text{IP}(I_r)^c)), \\ \Phi^{\text{src}} &= \beta (\epsilon_\phi(z_t^{\text{src}}, t, \mathcal{P}) - \epsilon_\phi(z_t^{\text{src}}, t, \emptyset)), \end{aligned} \quad (11)$$

$$\nabla_\theta \mathcal{L}_{\text{CSD}} = \mathbb{E}_{t, \epsilon} [\omega(t) (\Phi^{\text{tgt}} - \Phi^{\text{src}}) \frac{\partial z_t^{\text{tgt}}}{\partial \theta}].$$

Algorithm 1: FantasyStyle

- 1: **Input:** Diffusion model ϵ_ϕ , iterations M , guidance scale β , origin text embedding \mathcal{P} , content and style embeddings of the style image $\text{IP}(I_r)^c$ and $\text{IP}(I_r)^s$.
 - 2: **Output:** 3D Stylized Scene.
 - 3: **for** $m = 1, 2, \dots, M$ **do**
 - 4: Sample: $z_0^N = \mathcal{R}(C_N; \Theta)$, $\epsilon^N \sim \mathcal{N}(0, \mathbf{I})$, random timestep t , view number N
 - 5: **for** $k = [\text{src}, \text{tgt}]$, $P_{pos} = [\mathcal{P}, [\mathcal{P}, \text{IP}(I_r)^s]]$, $P_{neg} = [\emptyset, \text{IP}(I_r)^c]$ **do**
 - 6: $z_t^N = \sqrt{\alpha_t} z_0^N + \sqrt{1 - \alpha_t} \epsilon^N$
 - 7: **if** $k = \text{tgt}$ **then**
 - 8: $\hat{z}_t^N = \text{MVFC}(z_t^N)$
 - 9: **end if**
 - 10: Predict $\epsilon_\phi(z_t^N, t, P_{pos})$ and $\epsilon_\phi(z_t^N, t, P_{neg})$
 - 11: $\Phi = \beta (\epsilon_\phi(z_t^N, t, P_{pos}) - \epsilon_\phi(z_t^N, t, P_{neg}))$
 - 12: **end for**
 - 13: $\nabla_\theta \mathcal{L}_{\text{CSD}} = \omega(t) (\Phi^{\text{tgt}} - \Phi^{\text{src}})$ Eq. (11)
 - 14: Update $z_0^{\text{tgt}^N}$ with $\nabla_\theta \mathcal{L}_{\text{CSD}}$
 - 15: Update 3D Gaussians
 - 16: **end for**
 - 17: **Return:** 3D Stylized Scene.
-

Unlike SDS and DDS, which randomly sample from the entire diffusion timestep range, we perform random sampling from a fixed set of discrete timesteps to better simulate the DDIM denoising process of the diffusion model. As illustrated in Fig. 2, a timestep t is randomly selected to perturb both the source and rendered images, after which Eq. 11 is computed for each perturbed pair to optimize the 3D Gaussian scene. FantasyStyle is detailed in Alg. 1.

5 Experiments

5.1 Experimental Setup

Datasets. Our evaluation utilizes several distinct scene datasets (IN2N (Haque et al. 2023), Tandt DB (Kerbl et al. 2023), and Mip-NeRF 360 (Barron et al. 2022)), with 8 different image styles applied to each scene.

Metrics. We use Art-FID (Wright and Ommer 2022) to jointly evaluate content and style preservation. LPIPS (Zhang et al. 2018) and RMSE measure short-term and long-term content consistency with respect to the input, while FID assesses both style fidelity to the style image and content preservation relative to the original input.

Comparison Methods. We compare our method with recent state-of-the-art 3DGS-based style transfer methods, including SGSST (Galerie et al. 2024) and StyleGaussian (Liu et al. 2024).

Implementation Detail. We deploy 3DGS, SDXL (Podell et al. 2023), ControlNet (Zhang, Rao, and Agrawala 2023), and IP-Adapter (Ye et al. 2023) in all experiments. DDIM is used for noise sampling, with the CFG scale β set to 7.5. For MVFC, the parameter γ is set to 0.9. All experiments are conducted on 2 NVIDIA L20 (48GB) GPUs.




Figure 5: Qualitative comparison of different methods. Our approach achieves superior style transfer quality compared to existing methods. Zoom in for better view.

Method	ArtFID ↓	FID _{style} ↓	FID _{content} ↓	Short-range consistency		Long-range consistency	
				LPIPS ↓	RMSE ↓	LPIPS ↓	RMSE ↓
StyleGaussian	45.31	398.17	331.53	<u>0.290</u>	<u>0.267</u>	<u>0.542</u>	0.425
SGSST	<u>44.70</u>	<u>370.03</u>	<u>314.09</u>	0.295	0.27	0.569	0.502
Ours	43.52	347.61	261.71	0.285	0.241	0.529	<u>0.462</u>

Table 1: Quantitative comparison of different methods in 3DGS style transfer. FID_{style} measures the style fidelity to style image, while FID_{content} evaluates content preservation to the original input. **Bold**: best; underline: second best.

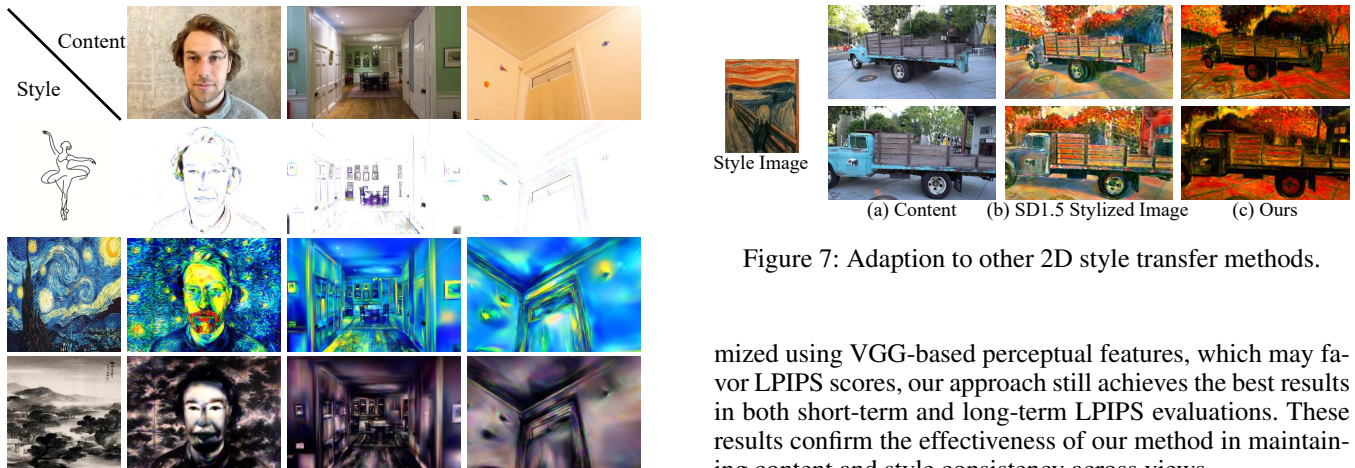


Figure 6: More qualitative results from our method.

5.2 Comparison Results

Quantitative Comparisons. To comprehensively evaluate the performance of our method, we conduct systematic comparisons with existing approaches across multiple metrics, as shown in Tab. 1. Our method consistently outperforms SGSST and StyleGaussian on key metrics such as ArtFID and FID, demonstrating superior style transfer quality and content preservation. Although prior methods are often opti-




Figure 7: Adaption to other 2D style transfer methods.

mized using VGG-based perceptual features, which may favor LPIPS scores, our approach still achieves the best results in both short-term and long-term LPIPS evaluations. These results confirm the effectiveness of our method in maintaining content and style consistency across views.

Qualitative Comparisons. Fig. 5 presents qualitative comparisons across four different scenes and various styles. Although StyleGaussian achieves LPIPS scores comparable to ours, its visual quality is significantly inferior. It fails to effectively transfer the target style’s color and brushstroke characteristics and often introduces artifacts that compromise the integrity of the original content. In contrast, SGSST demonstrates better style transfer in terms of color and stroke patterns; however, it tends to over-stylize the results. As seen in the first, second, seventh, and eighth columns of the third row, it introduces excessive stylistic elements, such as dense brushstrokes, which overwhelm and distort




Figure 8: Ablation visual results on MVFC.

Method	Short-range consistency	Long-range consistency
	LPIPS ↓	LPIPS ↓
w/o MVFC	0.253	0.587
Ours	0.250	0.574

Table 2: Quantitative results of the ablation study on MVFC.

the original structure. Although these methods only optimize color parameters, they still disrupt the geometry of the rendered images. Moreover, due to the inherent limitations of VGG-based features, existing methods often focus excessively on the appearance of the style image, rather than extracting abstract, transferable style representations adaptable to diverse input content. In comparison, our approach benefits from diffusion-based priors, enabling it to extract high-level, abstract style features and apply them consistently across varied scenes, while preserving the original structure of the image. We also present additional results, as shown in Fig. 6. Overall, our method achieves a better visual balance between stylization quality and content preservation, offering both faithful style transfer and structural integrity.

Other Extensions. As shown in Fig. 7, our method can be easily extended to incorporate various 2D style transfer models. Enhancements in 2D stylization directly translate to improved 3D visual quality.

Discussion. While VGG-based methods have fallen out of favor in 2D style transfer due to their limited ability to capture high-level style semantics and prevent content leakage, most 3DGS-based approaches still rely on them, inheriting these limitations. In contrast, recent 2D methods have shifted toward diffusion-based techniques for their flexibility and expressive power. However, 3DGS-based style transfer frameworks distilled purely from diffusion models remain unexplored. Our method introduces a novel pipeline that bridges this gap, enabling the extension of diffusion-based 2D style transfer techniques to 3D scenes.

5.3 Ablation Study

Ablation Study on Multi-View Frequency Consistency.

To thoroughly validate the effectiveness of MVFC, we conduct both quantitative and qualitative experiments. As shown in Tab. 2, our method achieves consistent improvements on LPIPS metrics under both short-term and long-term consistency evaluations. Notably, since MVFC is specifically designed to address multi-view inconsistency,



Figure 9: Ablation results on optimization strategy and CFG.

the gains are more pronounced in the long-term LPIPS results. Qualitative comparisons in Fig. 8 further support these findings. Without MVFC, conflicting 2D stylized priors often lead to blurred edge details and structural distortion in the reconstructed content. In contrast, MVFC effectively mitigates these conflicts, preserving sharp object boundaries and maintaining structural integrity across views.

Ablation Study on Optimization Strategy. Due to the presence of the reconstruction term, SDS and DDS often produce overly smooth and blurry outputs. In style transfer tasks, this is undesirable, as it leads to the loss of essential stylistic features, particularly brushstroke details. As illustrated in Fig. 9, while SDS and DDS successfully transfer the color characteristics of the style image, they fail to preserve brushstroke textures and significantly blur the structural details of the original content. Moreover, other optimization strategies require careful CFG scale tuning across styles or scenes, increasing complexity. In contrast, our method removes the reconstruction term, making the CFG scale less sensitive. As a result, adjusting the CFG scale does not affect stylistic expressiveness.

5.4 Limitations

Although SDXL-based optimization is time-consuming, it can be mitigated by using smaller models, lower resolutions, or adjusted color-related learning rates.

6 Conclusion

We propose FantasyStyle, a 3DGS-based style transfer framework that fully exploits diffusion model distillation. It features two key components: (1) Multi-View Frequency Consistency, which applies a 3D frequency filter to DDIM-noised multi-view latents to preserve high-frequency details and suppress low-frequency components for better view consistency; and (2) Controllable Stylized Distillation, which introduces negative guidance to prevent content leakage and uses 2D stylized priors to optimize the 3D scene. Experiments show that our method outperforms existing approaches in visual quality and content preservation.

Acknowledgments

This project is sponsored by Shanghai Pujiang Programme 24PJJD030 and Natural Science Foundation of Shanghai 25ZR1402138. Changshuo Wang is supported by the European Union’s Horizon 2024 Research and Innovation Programme for the Marie Skłodowska-Curie Actions under Grant No. 101211118.

References

- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2022. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5470–5479.
- Chung, J.; Hyun, S.; and Heo, J.-P. 2024. Style injection in diffusion: A training-free approach for adapting large-scale diffusion models for style transfer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8795–8805.
- Chung, J.; Lee, S.; Nam, H.; Lee, J.; and Lee, K. M. 2023. Luciddreamer: Domain-free generation of 3d gaussian splatting scenes. *arXiv preprint arXiv:2311.13384*.
- Ding, H.; Liu, C.; He, S.; Ying, K.; Jiang, X.; Loy, C. C.; and Jiang, Y.-G. 2025. MeViS: A multi-modal dataset for referring motion expression video segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Frenkel, Y.; Vinker, Y.; Shamir, A.; and Cohen-Or, D. 2024. Implicit style-content separation using b-lora. In *European Conference on Computer Vision*, 181–198. Springer.
- Galerie, B.; Wang, J.; Raad, L.; and Morel, J.-M. 2024. SGSST: Scaling Gaussian Splatting StyleTransfer. *arXiv preprint arXiv:2412.03371*.
- Gao, X.; Xu, Z.; Zhao, J.; and Liu, J. 2024. Frequency-controlled diffusion model for versatile text-guided image-to-image translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1824–1832.
- Gatys, L. A.; Ecker, A. S.; and Bethge, M. 2016. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2414–2423.
- Gu, Z.; Li, M.; Chen, R.; Ji, Z.; Guo, S.; Zhang, Z.; Ye, G.; and Hu, Z. 2024. ArtNVG: Content-Style Separated Artistic Neighboring-View Gaussian Stylization. *arXiv preprint arXiv:2412.18783*.
- Haque, A.; Tancik, M.; Efros, A. A.; Holynski, A.; and Kanazawa, A. 2023. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *Proceedings of the IEEE/CVF international conference on computer vision*, 19740–19750.
- He, S.; Ji, P.; Yang, Y.; Wang, C.; Ji, J.; Wang, Y.; and Ding, H. 2025. A survey on 3d gaussian splatting applications: Segmentation, editing, and generation. *arXiv preprint arXiv:2508.09977*.
- Hertz, A.; Aberman, K.; and Cohen-Or, D. 2023. Delta denoising score. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2328–2337.
- Hertz, A.; Voynov, A.; Fruchter, S.; and Cohen-Or, D. 2024. Style aligned image generation via shared attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4775–4785.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Jain, S.; Kuthiala, A.; Sethi, P. S.; and Saxena, P. 2024. Stylesplat: 3d object style transfer with gaussian splatting. *arXiv preprint arXiv:2407.09473*.
- Jung, H.; Nam, S.; Sarafianos, N.; Yoo, S.; Sorkine-Hornung, A.; and Ranjan, R. 2024. Geometry transfer for stylizing radiance fields. In *proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8565–8575.
- 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.
- Kovács, Á. S.; Hermosilla, P.; and Raidou, R. G. 2024. G-Style: Stylized Gaussian Splatting. In *Computer Graphics Forum*, volume 43, e15259. Wiley Online Library.
- Liu, K.; Zhan, F.; Xu, M.; Theobalt, C.; Shao, L.; and Lu, S. 2024. Stylegaussian: Instant 3d style transfer with gaussian splatting. In *SIGGRAPH Asia*.
- 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.
- Ouyang, Z.; Li, Z.; and Hou, Q. 2025. K-lora: Unlocking training-free fusion of any subject and style loras. *arXiv preprint arXiv:2502.18461*.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Poole, B.; Jain, A.; Barron, J. T.; and Mildenhall, B. 2022. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*.
- Rowles, C.; Vainer, S.; De Nigris, D.; Elizarov, S.; Kutsy, K.; and Donné, S. 2024. Ipadapter-instruct: Resolving ambiguity in image-based conditioning using instruct prompts. In *European Conference on Computer Vision*, 54–70. Springer.
- Saroha, A.; Gladkova, M.; Curreli, C.; Muhle, D.; Yenamandra, T.; and Cremers, D. 2024a. Gaussian splatting in style. In *DAGM German Conference on Pattern Recognition*, 234–251. Springer.
- Saroha, A.; Gladkova, M.; Curreli, C.; Muhle, D.; Yenamandra, T.; and Cremers, D. 2024b. Gaussian splatting in style. *arXiv preprint arXiv:2403.08498*.
- Si, C.; Huang, Z.; Jiang, Y.; and Liu, Z. 2024. Freeu: Free lunch in diffusion u-net. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4733–4743.

- Wang, C.; He, S.; Fang, X.; Han, J.; Liu, Z.; Ning, X.; Li, W.; and Tiwari, P. 2025a. Point Clouds Meets Physics: Dynamic Acoustic Field Fitting Network for Point Cloud Understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 22182–22192.
- Wang, C.; He, S.; Fang, X.; Wu, M.; Lam, S.-K.; and Tiwari, P. 2025b. Taylor series-inspired local structure fitting network for few-shot point cloud semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 7527–7535.
- Wright, M.; and Ommer, B. 2022. Artfid: Quantitative evaluation of neural style transfer. In *DAGM German Conference on Pattern Recognition*, 560–576. Springer.
- Wu, T.; Si, C.; Jiang, Y.; Huang, Z.; and Liu, Z. 2024. Freeinit: Bridging initialization gap in video diffusion models. In *European Conference on Computer Vision*, 378–394. Springer.
- Xing, P.; Wang, H.; Sun, Y.; Wang, Q.; Bai, X.; Ai, H.; Huang, R.; and Li, Z. 2024. Csgo: Content-style composition in text-to-image generation. *arXiv preprint arXiv:2408.16766*.
- Xu, R.; Xi, W.; Wang, X.; Mao, Y.; and Cheng, Z. 2025. StyleSSP: Sampling StartPoint Enhancement for Training-free Diffusion-based Method for Style Transfer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 18260–18269.
- Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Yu, X.-Y.; Yu, J.-X.; Zhou, L.-B.; Wei, Y.; and Ou, L.-L. 2024. Instantstylegaussian: Efficient art style transfer with 3d gaussian splatting. *arXiv preprint arXiv:2408.04249*.
- Zhang, D.; Fernandez-Labrador, C.; and Schroers, C. 2024. Coarf: Controllable 3d artistic style transfer for radiance fields. In *2024 International Conference on 3D Vision (3DV)*, 612–622. IEEE.
- Zhang, D.; Yuan, Y.-J.; Chen, Z.; Zhang, F.-L.; He, Z.; Shan, S.; and Gao, L. 2024. Stylizedgs: Controllable stylization for 3d gaussian splatting. *arXiv preprint arXiv:2404.05220*.
- Zhang, K.; Kolkin, N.; Bi, S.; Luan, F.; Xu, Z.; Shechtman, E.; and Snavely, N. 2022. Arf: Artistic radiance fields. In *European Conference on Computer Vision*, 717–733. Springer.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhang, Y.; He, Z.; Xing, J.; Yao, X.; and Jia, J. 2023a. Ref-npr: Reference-based non-photorealistic radiance fields for controllable scene stylization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4242–4251.
- Zhang, Y.; Huang, N.; Tang, F.; Huang, H.; Ma, C.; Dong, W.; and Xu, C. 2023b. Inversion-based style transfer with diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10146–10156.
- Zhuang, C.; Hu, Y.; Zhang, X.; Cheng, W.; Bao, J.; Liu, S.; Yang, Y.; Zeng, X.; Yu, G.; and Li, M. 2025. Styleme3d: Stylization with disentangled priors by multiple encoders on 3d gaussians. *arXiv preprint arXiv:2504.15281*.
- Zuo, Y.; Xiao, J.; Chan, K.-C.; Dong, R.; Yang, C.; He, Z.; Xie, H.; and Lam, K.-M. 2024. Towards multi-view consistent style transfer with one-step diffusion via vision conditioning. In *European Conference on Computer Vision*, 296–313. Springer.