

GT²-GS: Geometry-aware Texture Transfer for Gaussian Splatting

Wenjie Liu¹, Zhongliang Liu², Junwei Shu¹, Changbo Wang³, Yang Li^{1*}

¹School of Computer Science and Technology, East China Normal University, Shanghai, China

²School of Software Engineering, East China Normal University, Shanghai, China

³School of Data Science and Engineering, East China Normal University, Shanghai, China
{51265901068, 10235101440, 51265901091}@stu.ecnu.edu.cn, {cbwang, yli}@cs.ecnu.edu.cn

Abstract

Transferring 2D textures onto complex 3D scenes plays a vital role in enhancing the efficiency and controllability of 3D multimedia content creation. However, existing 3D style transfer methods primarily focus on transferring abstract artistic styles to 3D scenes. These methods often overlook the geometric information of the scene, which makes it challenging to achieve high-quality 3D texture transfer results. In this paper, we present GT²-GS, a geometry-aware texture transfer framework for gaussian splatting. First, we propose a geometry-aware texture transfer loss that enables view-consistent texture transfer by leveraging prior view-dependent feature information and texture features augmented with additional geometric parameters. Moreover, an adaptive fine-grained control module is proposed to address the degradation of scene information caused by low-granularity texture features. Finally, a geometry preservation branch is introduced. This branch refines the geometric parameters using additionally bound Gaussian color priors, thereby decoupling the optimization objectives of appearance and geometry. Extensive experiments demonstrate the effectiveness and controllability of our method. Through geometric awareness, our approach achieves texture transfer results that better align with human visual perception.

Homepage — <https://vpx-ecnu.github.io/GT2-GS-website>

Introduction

3D style transfer (Zhang et al. 2022, 2024) aims to transfer the stylistic elements of a reference image onto a 3D scene while preserving the scene’s original structural and semantic information. With the rapid development of fields such as virtual reality, robotics, film, and gaming, the demand for high-quality 3D content has increased significantly. 3D stylization techniques offer a promising solution by accelerating the creation of 3D content, particularly in complex 3D scene environments.

Recently, the emergence of Neural Radiance Fields (NeRF) (Mildenhall et al. 2020) and 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) has significantly advanced the field of 3D stylization. NeRF-based stylization methods (Zhang et al. 2022, 2023; Nguyen-Phuoc, Liu, and

*Corresponding authors.




Figure 1: Previous state-of-the-art methods struggle to accurately transfer complex textures to 3D scene. In contrast, our GT²-GS framework incorporates geometric information into the optimization process, enabling geometry-aware texture transfer.

Xiao 2022), leveraging the advantages of implicit representations, allow for the decoupled optimization of appearance and geometry. In contrast, 3DGS-based approaches (Galerne et al. 2025; Liu et al. 2025, 2024) offer benefits such as explicit editability and real-time rendering. 3D stylization methods can be broadly categorized into feed-forward and optimization-based approaches, depending on whether they support zero-shot style transfer. While feed-forward methods (Liu et al. 2023; Huang et al. 2022; Liu et al. 2024) enable zero-shot stylization, they often exhibit lower rendering quality. In contrast, optimization-based methods (Zhang et al. 2022; Liu et al. 2025; Zhang et al. 2024; Galerne et al. 2025) typically produce higher-fidelity stylization results. Despite these differences, both types of methods commonly define optimization objectives within the VGG feature space during training.

Texture is a crucial component of various graphics pipelines and plays a vital role in producing high-quality 3D assets. Compared to abstract artistic styles, texture elements are more controllable for users. However, as illustrated in Fig. 1, existing methods struggle to accurately transfer the texture of the reference image onto the 3D scene. We analyze the limitations of current approaches from the perspective of the scene optimization process. First, during scene optimization, multi-view objectives should preserve correct geometric relationships. For example, using content images captured from the same scene as ground truth. However, ex-

isting style transfer methods define their objectives independently across views, without considering the rich geometric structure within the scene or the geometric consistency across viewpoints. Moreover, we observe a fine-grained mismatch between the VGG feature space and the pixel space. Commonly used style transfer losses fail to account for this mismatch, which leads to regions with high pixel-level information density being easily disrupted by coarse-grained feature representations. Addressing these issues is crucial for achieving high-quality texture transfer.

In this paper, we propose a novel framework GT^2 -GS for achieving geometry-aware texture transfer. Our method explicitly accounts for geometric information and fine-grained discrepancies, enabling texture transfer results that better align with human visual perception. We first propose a geometry-aware texture transfer loss, which is built upon texture features augmented with additional geometric parameters and incorporates cross-view geometric information to ensure consistent and controllable texture transfer. In addition, an adaptive fine-grained control module is proposed. It adaptively adjusts the strength of texture learning based on the information density of different pixel regions, thereby mitigating the fine-grained discrepancy between texture features and pixel-level representations. To address the coupling between Gaussian geometry and color parameters, we further introduce a geometry preservation branch. We bind additional color parameters to the Gaussians and optimize them using the content image as ground truth, in order to accurately refine the scene geometry. Extensive experiments demonstrate that our proposed framework achieves high-quality texture transfer results.

Our main innovations are as follows:

- We propose a geometry-aware texture transfer framework for general 3DGS scenes. The proposed framework enables controllable and high-quality texture transfer results.
- Our Geometry-aware Texture Transfer Loss and Adaptive Fine-Grained Control Module respectively account for geometric information and fine-grained discrepancies between texture features and pixels, enabling high-quality texture transfer. Moreover, the Geometry Preservation Branch provides a novel approach to preserving geometry during appearance editing.
- Extensive experiments demonstrate that, compared to SOTA methods, our proposed work achieves texture transfer results more consistent with human visual perception. Our proposed method maintains real-time rendering and multi-view consistency. The source code will be released.

Related Work

Texture Transfer

Early texture transfer algorithms (Ashikhmin 2003; Lee et al. 2010; Hertzmann et al. 2023; Efros and Freeman 2023) were based on non-parametric texture synthesis. These methods preserve the structural information of the target image to achieve texture transfer effects. Gatys et al. (Gatys,

Ecker, and Bethge 2016) pioneered the use of convolutional neural networks (CNNs) (Simonyan and Zisserman 2014) for image style transfer, achieving impressive results. Subsequently, a variety of deep learning-based style transfer methods (Chen and Schmidt 2016; Li and Wand 2016; Huang and Belongie 2017; Huo et al. 2021) have been proposed. These methods are capable of transferring abstract artistic elements such as color distribution, brushstroke characteristics, and overall composition, rather than focusing solely on texture transfer. Leveraging the powerful feature extraction capabilities of networks such as CNNs and Transformers (Vaswani et al. 2017), some deep learning-based texture transfer methods (Wang et al. 2022; Chen, Yin, and Fidler 2022; Pu et al. 2024) can transfer local textures from the reference image to the target image based on high-level semantic correspondence. In this work, we explore how to directly transfer texture features into 3D representations.

3D Style Transfer

3D style transfer aims to transfer the style from 2D images to the appearance of 3D scenes, while maintaining the scene’s content and multi-view consistency. Early research in 3D scene style transfer focused on explicit representations such as point clouds (Cao et al. 2020; Huang et al. 2021), meshes (Höllerin, Johnson, and Nießner 2022; Yin et al. 2021), and voxels (Guo et al. 2021; Klehm et al. 2014). Zhang et al. (Zhang et al. 2022) explored the potential of NeRF in 3D stylization tasks. By optimizing the scene representation with the proposed 3D style transfer loss, it achieved pleasing visual effects. Subsequently, a variety of NeRF-based 3D style transfer methods (Huang et al. 2022; Pang, Hua, and Yeung 2023; Liu et al. 2023; Zhang, Fernandez-Labrador, and Schroers 2024; Jung et al. 2024) have emerged. Recently, the emergence of 3DGS (Kerbl et al. 2023) has brought new possibilities for scene stylization. It features fast training speed, high rendering quality, and efficient performance. Some works (Liu et al. 2024; Zhang et al. 2024; Mei, Xu, and Patel 2024; Liu et al. 2025) have already explored the potential of stylization under this representation. Currently, 3DGS-based scene stylization methods can be categorized into feed-forward approaches (Liu et al. 2024) and optimization-based approaches (Zhang et al. 2024; Mei, Xu, and Patel 2024; Liu et al. 2025). Feed-forward methods enable zero-shot style transfer. For instance, StyleGaussian (Liu et al. 2024) embeds VGG features into Gaussians and leverages pre-trained scene representations combined with AdaIN (Huang and Belongie 2017) to achieve real-time stylization. On the other hand, optimization-based methods attain higher-quality style transfer results. Some works (Zhang et al. 2024; Liu et al. 2025) further exploit the explicit characteristic of Gaussian representations to enable region-controllable scene stylization. However, existing methods do not incorporate geometric information during the stylization process, making them unsuitable for texture images. Meanwhile, these methods overlook the fine-grained discrepancies between features and pixels. By incorporating geometric information and adaptive fine-grained control, our method is able to achieve high-quality texture transfer results.

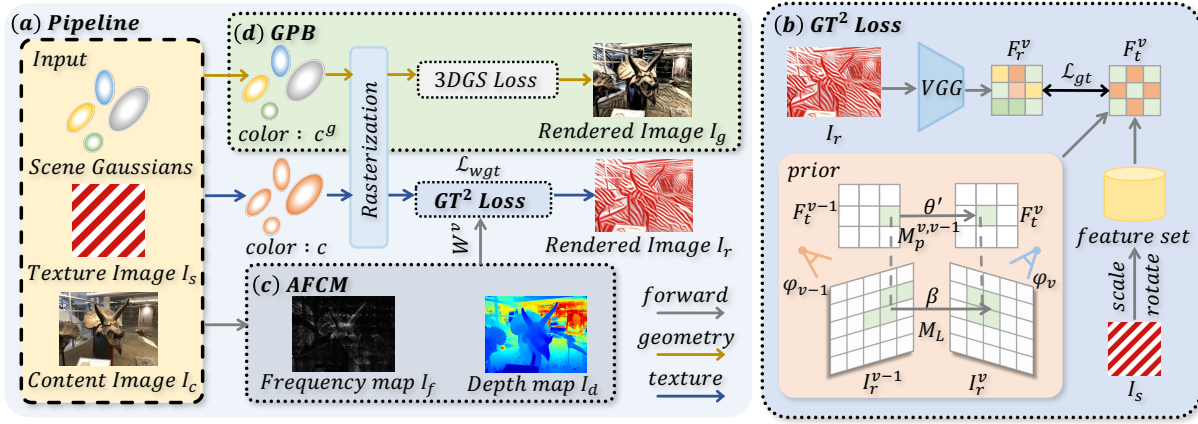


Figure 2: Overview of GT^2 -GS. The overall pipeline is illustrated in (a). The input to our proposed framework includes the scene Gaussians, content images, and a texture image. The texture image is transformed into a texture feature set, which is used to construct the target feature map in the GT^2 Loss, as shown in (b). We extract pixel-wise information from the input and transform it into an adaptive weight matrix to control texture learning, as illustrated in (c). After embedding the parameter c^g , the Gaussians are optimized using two branches. The additional geometry preservation branch is shown in (d). As a result, the texture features are integrated into the Gaussian representation.

Preliminaries

3D Gaussian Splatting

3D Gaussian Splatting (Kerbl et al. 2023) is a novel explicit 3D representation method. It uses a set of parameterized Gaussians $G = \{g_i\}$ to represent 3D scenes. The parameters of each Gaussian g_i include a mean vector μ_i representing its center position, a covariance matrix Σ_i describing its shape, an opacity parameter σ_i and a color parameter c_i , which is represented as spherical harmonic coefficients. Among these, the covariance matrix Σ_i is decomposed into rotation parameters r_i and scaling parameters s_i to effectively maintain its positive semi-definite property during the optimization process. The optimization parameters of Gaussian g_i are actually represented as $g_i = \{\mu_i, r_i, s_i, \sigma_i, c_i\}$. The 3D Gaussians are efficiently rendered through a fast differentiable rasterization (Lassner and Zollhofer 2021). Specifically, the Gaussians are grouped into tiles and sorted by depth. The color C of a pixel is computed through α -blending:

$$C = \sum_{i \in N} T_i \alpha_i c_i, T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (1)$$

where T_i is the transmittance and α_i is the alpha-compositing weight for the i -th Gaussian. The depth value D can be obtained using a method similar to rendering color,

$$D = \sum_{i \in N} T_i \alpha_i d_i, \quad (2)$$

where d_i is the depth for i -th Gaussian. During optimization, 3DGS adopts an adaptive densification strategy to control the distribution of Gaussians. Specifically, this strategy performs cloning or splitting of Gaussians based on their positional gradients and sizes.

Style Transfer Loss

By computing the loss on the feature maps of rendered views, the learnable parameters in the scene representation can be optimized. Existing optimization-based 3D style transfer methods (Zhang et al. 2022; Pang, Hua, and Yeung 2023; Zhang, Fernandez-Labrador, and Schroers 2024; Zhang et al. 2024) typically adopt the nearest neighbor feature matching (NNFM) loss (Zhang et al. 2022) as their style loss function. It matches the rendered features with the nearest neighbor features in the style feature set and minimizes the cosine distance between them. Specifically, a random viewpoint is selected to obtain the rendered image I_r . The same feature extractor (e.g., VGG (Simonyan and Zisserman 2014)) is used to extract the feature maps F_r and F_s from I_r and the style image I_s , respectively. Let $F_r(i, j)$ denote the feature vector at the pixel location (i, j) of the rendered feature map F_r , the NNFM loss can be expressed as

$$L_{style} = \frac{1}{N} \sum_{i,j} \min_{i',j'} dist(F_r(i, j), F_s(i', j')), \quad (3)$$

where N is the number of pixels in F_r , $dist(a, b)$ is the cosine distance between two feature vectors a and b :

$$dist(a, b) = 1 - \frac{a \cdot b}{\|a\| \|b\|}. \quad (4)$$

Method

GT^2 -GS Framework Overview

In this section, we provide a detailed overview of the proposed GT^2 -GS framework. Geometry-aware Texture Transfer Loss (GT^2 Loss) is proposed to enable controllable and high-quality texture transfer. It is computed using texture features bound with geometric parameters and the prior information from the previous viewpoints. Furthermore, we




Figure 3: Differences in Perspective. The same textured region (highlighted in red) on the 3D object exhibits different texture orientations under varying viewpoints.

propose the Adaptive Fine-grained Control Module (AFCM) to regulate the texture learning intensity based on the scene information density in the pixel space. This prevents excessive information from being corrupted by low-granularity features. Finally, for the input Gaussian, an additional color parameter c^g is embedded and initialized using the current color parameter c . The Geometry Preservation Branch (GPB) is introduced based on c^g , with the aim of decoupling the optimization processes of the geometry and appearance of the scene. The pipeline is shown in Fig. 2.

Geometry-aware Texture Transfer Loss

Existing methods primarily focus on artistic style transfer and are mostly based on NFM loss. As shown in Eq. 3, the objective considers only the relationship between the rendered feature map and the style feature map. The optimization objectives across different viewpoints are constructed independently, without accounting for the underlying geometric relationships between views. However, texture is inherently tied to the geometry of the scene. To achieve accurate texture transfer, we incorporate geometric information into the loss function computation during the optimization process.

First, we associate geometric parameters with each texture feature. Considering the perspective geometry within the scene, edge geometric structures, and transformations across viewpoints, we apply scaling and rotation operations to the texture images. Texture features are then extracted from the transformed images and aggregated into a feature set. Specifically, we obtain the original scene depth map using Eq. 2. The depth values are then sorted and discretized into K groups based on predefined depth intervals. The scaling factor for each group is computed as the ratio between the lowest group depth value Z_1 and the group’s depth value Z_k . Next, each scaled image is rotated by an angle θ to capture multi-directional information present in rendered views. Each feature in the texture feature set is denoted as $\{f_{k,\theta}\}$. The scaling parameter k and rotation angle θ are retained for subsequent computations.

Based on the constructed set of texture features, we propose the GT² Loss to enable geometry-consistent texture transfer. We construct a per-pixel matched target feature map F_t for the rendered feature map F_r . Texture transfer is achieved by minimizing the cosine similarity between the corresponding feature vectors of F_r and F_t . To ensure that the feature maps F_t constructed from multiple viewpoints maintain correct geometric relationships, we incorporate the

construction result of the previous viewpoint’s feature map F_t^{v-1} when building the target feature map F_t^v for the current viewpoint. The computation process of the homography matrix $M_p^{v,v-1}$ between two viewpoints is formulated as

$$M_p^{v,v-1} = K_{v-1}[R_{v-1}|T_{v-1}][R_v|T_v]^{-1}K_v^{-1}, \quad (5)$$

where K_v and $[R_v|T_v]$ represent the intrinsic matrix and world-to-camera extrinsic parameters of the v -th viewpoint, respectively. Through the homography matrix, the current viewpoint can sample the prior feature $f_{k',\theta'}$ in the screen coordinate system of the prior viewpoint. Considering occlusion relationships, we filter the projected points using the depth map I_d^{v-1} of the prior viewpoint.

However, as shown in Fig. 3, the orientation of the same texture region varies across different viewing angles. Each texture feature vector inherently contains scale and orientation. Directly using this feature vector as a prior fails to account for the impact of viewpoint changes in 3D space. To address this, we utilize upsampling to obtain the pixel set $\{p_v\}$ corresponding to each feature map location in the pixel coordinate system. Through the transformation relationship $M_p^{v-1,v}$, we can determine its corresponding pixel set $\{p_{v-1}\}$ in the previous viewpoint. We use the least squares method to compute the linear transformation matrix $M_L \in R^{2 \times 2}$ for texture variation between viewpoints. The obtained transformation matrix M_L can be decomposed using SVD to extract the rotation angle β . The construction method of the texture feature vector at position (i, j) in the target feature map F_t is formulated as

$$F_t(i, j) = \arg \min_{f_{k,\theta}} \text{dist}(F_r(i, j), f_{k,\theta}) + \lambda_p |\theta' + \beta - \theta|, \quad (6)$$

where λ_p is the prior texture orientation control coefficient. Subsequently, texture transfer can be achieved by minimizing the cosine similarity between the feature vectors at corresponding positions in the rendered feature map F_r^v and the target feature map F_t^v for the current viewpoint v . The geometry-aware texture transfer loss function is formulated as follows:

$$L_{gt} = \frac{1}{N} \sum_{i,j} \text{dist}(F_r^v(i, j), F_t^v(i, j)). \quad (7)$$

Adaptive Fine-grained Control Module

Most existing optimization-based style transfer methods rely on VGG features. However, feature maps extracted through multiple layers of convolutional neural networks exhibit significantly lower granularity compared to the original image pixels. When optimizing the scene using only texture transfer losses, this can lead to the following issues: First, due to perspective projection, regions with greater depth tend to concentrate more scene information. Learning coarse-grained texture features may overwrite these important details. Second, scenes often contain geometrically fine-grained structures such as stairs and railings, which can be degraded or lost under coarse texture representations. To address these challenges, we propose an Adaptive Fine-Grained Control Module.

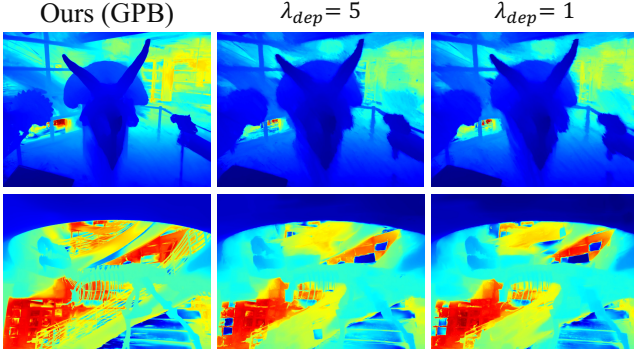


Figure 4: Comparison of Geometry Preservation Results. It can be observed that as the number of Gaussians increases, depth regularization struggles to preserve the geometric information of the scene. λ_{dep} is the weighting coefficient for the depth regularization term.

Specifically, we first obtain the depth map I_d and frequency density map I_f for each viewpoint from the original Gaussian scene and content image, respectively. Both are rescaled to match the spatial dimensions of the feature maps. Furthermore, we introduce a geometric distortion map Φ , defined as the discrepancy in geometric information between the learned texture features and rendered features. The geometric distortion map is computed by measuring the angular difference between the features obtained using Eq. 6 and those derived without any prior information. The adaptive fine-grained control map is formulated as

$$W^v = \lambda_d(1 - I_d^v) + \lambda_f(1 - I_f^v) + \lambda_\Phi(1 - \Phi), \quad (8)$$

where I_d^v , I_f^v and Φ are all normalized, and λ_d , λ_f and λ_Φ are weight coefficients. In most cases, we aim to alter the appearance of foreground objects while simultaneously satisfying requirements for both shallow depth and high frequency. We therefore employ an additive formulation to balance the relationship between depth and frequency. Furthermore, we regularize the learning process to favor textures with low distortion, thereby better preserving geometric fidelity during texture adaptation.

The derived adaptive weight matrix W^v is subsequently applied to Eq. 7 as

$$L_{wgt} = \frac{1}{N} \sum_{i,j} W^v(i,j) \text{dist}(F_r^v(i,j), F_t^v(i,j)), \quad (9)$$

where L_{wgt} is weighted GT² Loss. The total loss during the texture transfer phase is expressed as

$$L_{tot} = \lambda_{wgt}L_{wgt} + \lambda_cL_{content} + \lambda_{tv}L_{tv}, \quad (10)$$

where L_{tv} is the total variation loss, λ_{wgt} , λ_c , λ_{tv} are the coefficients of the corresponding loss functions.

Geometry Preservation Branch

3DGS is an explicit representation, where the scene’s ability to learn textures is closely related to the distribution of Gaussians. In low-texture regions, the density of Gaussians

is typically lower. While these regions achieve good rendering quality in the original rendering process, they struggle to learn new texture appearances. The Gaussian densification strategy increases the number of Gaussians in low-texture regions during the texture migration phase, enabling the learning of complex textures. However, Gaussians are densified solely based on gradients. Since texture transfer lacks ground truth during optimization, the densification strategy may introduce erroneous Gaussians floating in space. As shown in Fig. 4, simply adding depth regularization does not solve this problem. Unlike NeRF, explicit 3D Gaussians jointly encode both geometric and color parameters. To address this issue, we propose a specialized branch for optimizing the geometric parameters of Gaussians.

Our key insight is to introduce an additional optimization objective focused on geometry preservation during training, in order to balance appearance optimization with geometric integrity. Specifically, we associate each Gaussian with an additional color parameter c^g , initialized using the original color values from the scene. During optimization, we render an image I_g using these color parameters and optimize the Gaussian parameters by treating the content image I_c as ground truth. The 3D Gaussian Splatting reconstruction loss function is formulated as

$$\mathcal{L}_{rec} = (1 - \lambda)\mathcal{L}_1 + \lambda\mathcal{L}_{D-SSIM}, \quad (11)$$

where \mathcal{L}_1 and \mathcal{L}_{D-SSIM} are calculated between the rendered image I_g and the content image I_c . Through an optimization process with ground truth, the Gaussians in the scene are moved to their correct geometric positions.

Experiment

Datasets. For the scene datasets, we utilize the LLFF dataset (Mildenhall et al. 2019) and the Tanks and Temples (T&T) dataset (Knapitsch et al. 2017), which are collected from the real world. Additionally, we use images from the ARF (Zhang et al. 2022) style dataset and the DTD dataset (Cimpoi et al. 2014) as reference image datasets.

Baseline. We compare our method with the state-of-the-art 3D stylization methods, including SGSST (Galerie et al. 2025), ABC-GS (Liu et al. 2025), StyleGaussian (Liu et al. 2024), ARF (Zhang et al. 2022), Ref-NPR (Zhang et al. 2023), and SNeRF (Nguyen-Phuoc, Liu, and Xiao 2022). Specifically, SGSST, ABC-GS, and StyleGaussian are based on 3DGS, and ARF, Ref-NPR, and SNeRF are based on NeRF. StyleGaussian is a feed-forward-based method, while others are optimization-based methods.

Implementation Details

We perform view-consistent color transfer (Zhang et al. 2022) on the rendered images before and after texture transfer and use them as content images to optimize the Gaussian parameters. For the VGG (Simonyan and Zisserman 2014) feature extractor, we employ the conv3 block of VGG-16. When associating color parameters with texture features, we set the default number of depth groups K to 4, and the rotation angle θ is sampled across the full 360 degrees. The hyperparameter for AFCM is denoted as $\{\lambda_d, \lambda_f, \lambda_\Phi\} =$

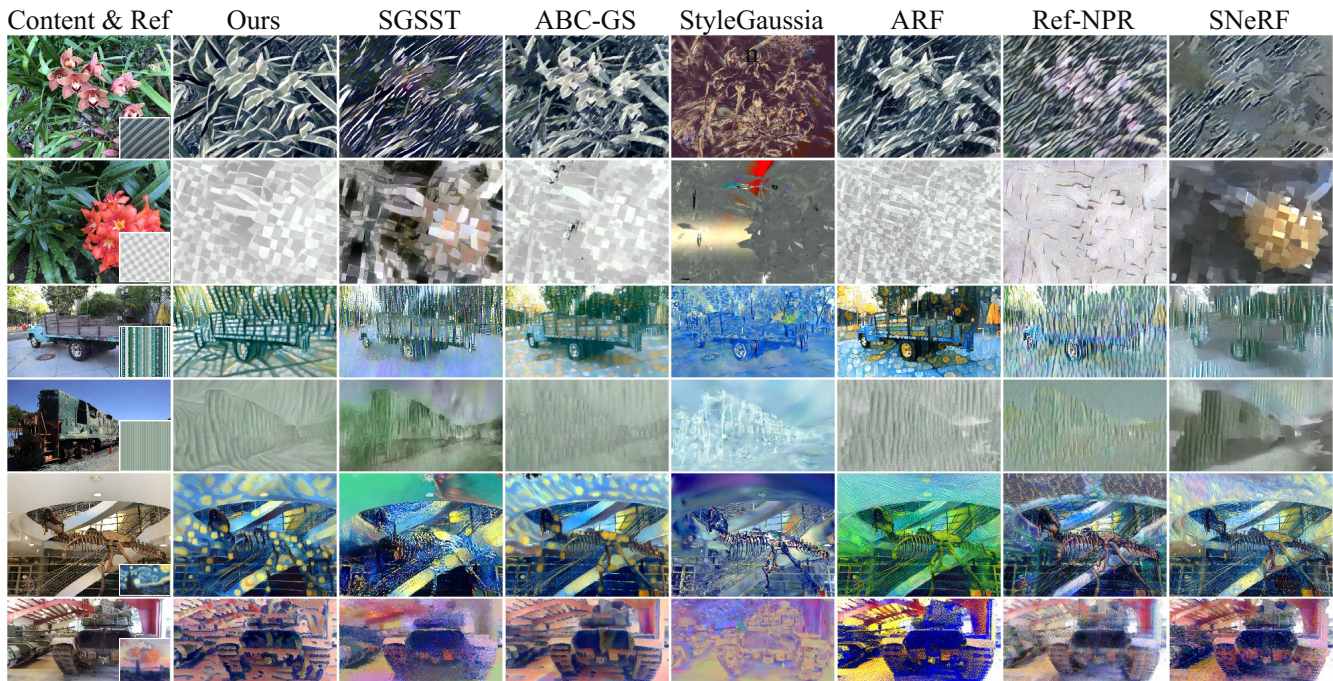


Figure 5: Qualitative Comparison of Texture Transfer and Style Transfer. The first four rows in the figure show the texture transfer results, while the last two rows present the style transfer results.

$\{0.8, 0.8, 0.25\}$. For texture transfer optimization, we set $\{\lambda_{wgt}, \lambda_c, \lambda_{tv}\} = \{2, 0.005, 0.02\}$. All our experiments are conducted on a single NVIDIA RTX 4090 GPU.

Qualitative Evaluation

To comprehensively evaluate the effectiveness of our method, the qualitative experiments are divided into two parts: texture transfer and style transfer. The results are shown in Fig. 5.

Texture Transfer. Fig. 5 shows the texture transfer results of our method compared with other methods. Visually, our results exhibit higher fidelity to the reference texture and better alignment with human visual perception. For example, in the geometrically complex orchids scene (row 1 in Fig. 5), our method produces transfer results that appear as if the texture is wrapped naturally around the surface of the orchids. StyleGaussian, as a zero-shot style transfer method, struggles to handle such complex texture transfers. The results of other optimization-based methods exhibit texture discontinuities and appear noticeably blurry, significantly deteriorating overall visual quality. This is primarily due to the inherent coupling between texture and geometry, which these methods fail to consider during optimization. In contrast, our geometry-aware approach ensures consistent texture transfer across multiple views and coherent 3D results.

Style Transfer. To validate the generalizability of our method, we selected artistic style images as reference inputs and conducted a qualitative comparison of style transfer with other methods. As shown in Fig. 5, our approach is more effective in transferring the texture elements of the

style image to the 3D scene. SGSST and ABC-GS disable the Gaussian densification strategy during style transfer optimization, making it challenging to capture rich and detailed style features. For instance, in the trex scene (second-to-last row of Fig. 5), the appearance of the white wall region fails to be effectively optimized. In contrast, our method benefits from the joint effect of the AFCM and GPB. This design enables the model to preserve scene geometry while effectively learning style textures from the reference image. NeRF-based methods tend to capture more abstract stylistic elements, such as brush strokes.

Quantitative Evaluation

In the 3D scene appearance editing task, maintaining multi-view consistency and preserving scene content information are crucial. To this end, we conducted extensive quantitative experiments from both aspects. We randomly selected 100 scene-reference image pairs to quantitatively evaluate our method.

Multi-view Consistency. We evaluate multi-view consistency (Liu et al. 2024) using both short-term and long-term consistency metrics. The results are shown in Tab. 1. It can be observed that both our method and ABC-GS achieve high-quality multi-view consistency. Notably, ABC-GS disables the Gaussian densification strategy during optimization. In contrast, our method maintains multi-view consistency even after applying the densification strategy, demonstrating the effectiveness of the proposed GPB.

Content Preservation. For 3D texture transfer, it is essential to ensure that the original scene content remains

Methods	SSIM(\uparrow)	CLIP-score(\uparrow)	ST-LPIPS(\downarrow)	ST-RMSE(\downarrow)	LT-LPIPS(\downarrow)	LT-RMSE(\downarrow)
Ours	<u>0.51</u>	0.47	<u>0.054</u>	<u>0.048</u>	<u>0.087</u>	<u>0.077</u>
SGSST	0.45	0.44	0.075	0.072	0.119	0.108
ABC-GS	0.56	<u>0.46</u>	0.049	0.041	0.080	0.068
StyleGaussian	0.41	0.40	0.058	0.052	0.097	0.082
ARF	0.37	0.45	0.109	0.072	0.152	0.108
Ref-NPR	0.35	0.42	0.092	0.069	0.137	0.102
SNeRF	0.48	0.36	0.075	0.057	0.127	0.090

Table 1: Quantitative Experiment on Multi-view Consistency and Content Preservation.

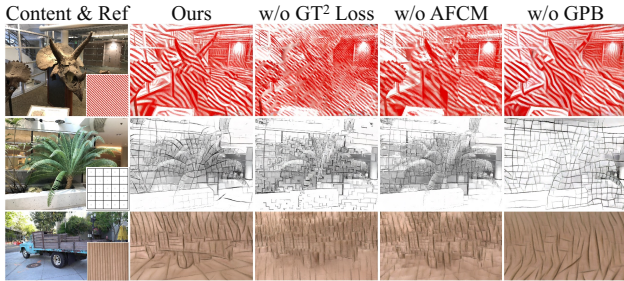


Figure 6: Ablation study for GT^2 Loss, AFCM, and GPB.

recognizable while editing the scene’s appearance. We use SSIM (Wang et al. 2004) and CLIP-score (Radford et al. 2021) to assess the preservation of content information. Specifically, SSIM evaluates the structural and informational similarity between two images at the pixel level. In contrast, CLIP-score measures the semantic similarity by computing the cosine similarity between the CLIP embeddings of the two images. As shown in Tab. 1, our proposed method achieves significantly higher scores on both evaluation metrics. In particular, our method outperforms previous approaches in terms of the CLIP-score. It indicates that incorporating geometric information enables accurate texture transfer to the scene appearance while preserving the semantic content of the scene.

Ablation Study

Impact of GT^2 Loss. GT^2 Loss is introduced to incorporate geometric information into the optimization objective of texture transfer. As shown in Fig. 6, removing the GT^2 Loss results in noticeable texture discontinuities and blurring in the texture transfer outputs. Tab. 2 further demonstrates the significance of GT^2 Loss from a quantitative perspective.

Impact of AFCM. The AFCM is introduced to address the granularity mismatch between texture features and pixel-level details. It encourages texture learning in foreground and low-frequency regions, while preserving scene content in background and high-frequency areas. As illustrated in Fig. 6, in the fern scene (second row), low-texture foreground regions fail to capture style patterns without the presence of AFCM. In the truck scene (third row), which exhibits significant depth variation as a 360° environment, removing AFCM leads to noticeable degradation of geometric and ap-

	Ours	w/o GT^2 Loss	w/o AFCM	w/o GPB
SSIM(\uparrow)	<u>0.41</u>	0.38	0.45	0.31
CLIP-score(\uparrow)	0.39	0.36	<u>0.38</u>	0.37

Table 2: Quantitative Ablation Study on Content Protection. The results are obtained from 25 randomly selected experiments conducted on LLFF scenes.

pearance fidelity.

Impact of GPB. The 3D texture transfer tasks inherently lack ground truth supervision, which can introduce incorrect geometry during the scene appearance editing process. To address this, we propose GPB, which leverages the original content images to optimize the geometric parameters of Gaussians and correct inaccurate geometry. As illustrated in Fig. 6, removing GPB results in noticeable artifacts across the scene. Tab. 2 highlights the importance of GPB in maintaining content fidelity. Through GPB, our texture transfer framework achieves a balance between learning texture features and preserving scene geometry.

Conclusion

In this paper, we introduced GT^2 -GS, a novel Geometry-aware Texture Transfer framework for Gaussian Splatting that achieves high-quality texture transfer results. Unlike the previous 3D stylization methods, our approach explicitly considered the intrinsic relationship between geometry and texture. We first propose the GT^2 Loss, which leverages features augmented with geometric parameters and utilizes cross-view priors to guide scene optimization, enabling geometrically consistent texture transfer. AFCM addresses the granularity mismatch between features and pixels by adaptively controlling the strength of texture learning. Additionally, GPB introduces a geometry optimization objective grounded in ground-truth appearance, effectively preserving scene geometry during complex stylization. Extensive quantitative and qualitative experiments demonstrated the effectiveness of our proposed method. Moreover, our framework was capable of generating high-quality stylization results, showcasing its generalizability.

As a limitation, because optimizing rendered VGG features requires minimizing texture cosine distance while preserving content loss, which results in a texture interpolated between scene and texture geometry.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (62572194, 62472178, 62376244) and the Fundamental Research Funds for the Central Universities. This work was supported by the Shanghai Frontiers Science Center of Molecule Intelligent Syntheses. Meanwhile, we thank the financial support from Shanghai Chinafortune Co., Ltd.

References

- Ashikhmin, N. 2003. Fast texture transfer. *IEEE computer Graphics and Applications*, 23(4): 38–43.
- Cao, X.; Wang, W.; Nagao, K.; and Nakamura, R. 2020. Psnnet: A style transfer network for point cloud stylization on geometry and color. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer vision*, 3337–3345.
- Chen, T. Q.; and Schmidt, M. 2016. Fast patch-based style transfer of arbitrary style. *arXiv preprint arXiv:1612.04337*.
- Chen, Z.; Yin, K.; and Fidler, S. 2022. Auv-net: Learning aligned uv maps for texture transfer and synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1465–1474.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3606–3613.
- Efros, A. A.; and Freeman, W. T. 2023. Image quilting for texture synthesis and transfer. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 571–576.
- Galerie, B.; Wang, J.; Raad, L.; and Morel, J.-M. 2025. SGSST: Scaling Gaussian Splatting Style Transfer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 26535–26544.
- 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.
- Guo, J.; Li, M.; Zong, Z.; Liu, Y.; He, J.; Guo, Y.; and Yan, L.-Q. 2021. Volumetric appearance stylization with stylizing kernel prediction network. *ACM Trans. Graph.*, 40(4): 162–1.
- Hertzmann, A.; Jacobs, C. E.; Oliver, N.; Curless, B.; and Salesin, D. H. 2023. Image analogies. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 557–570.
- Höllein, L.; Johnson, J.; and Nießner, M. 2022. Stylemesh: Style transfer for indoor 3d scene reconstructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6198–6208.
- Huang, H.-P.; Tseng, H.-Y.; Saini, S.; Singh, M.; and Yang, M.-H. 2021. Learning to stylize novel views. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13869–13878.
- Huang, X.; and Belongie, S. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, 1501–1510.
- Huang, Y.-H.; He, Y.; Yuan, Y.-J.; Lai, Y.-K.; and Gao, L. 2022. Stylizednerf: consistent 3d scene stylization as stylized nerf via 2d-3d mutual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18342–18352.
- Huo, J.; Jin, S.; Li, W.; Wu, J.; Lai, Y.-K.; Shi, Y.; and Gao, Y. 2021. Manifold alignment for semantically aligned style transfer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14861–14869.
- 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 Transactions on Graphics*, 42(4).
- Klehm, O.; Ihrke, I.; Seidel, H.-P.; and Eisemann, E. 2014. Property and lighting manipulations for static volume stylization using a painting metaphor. *IEEE Transactions on Visualization and Computer Graphics*, 20(7): 983–995.
- Knapitsch, A.; Park, J.; Zhou, Q.-Y.; and Koltun, V. 2017. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4): 1–13.
- Lassner, C.; and Zollhofer, M. 2021. Pulsar: Efficient sphere-based neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1440–1449.
- Lee, H.; Seo, S.; Ryoo, S.; and Yoon, K. 2010. Directional texture transfer. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*, 43–48.
- Li, C.; and Wand, M. 2016. Combining markov random fields and convolutional neural networks for image synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2479–2486.
- Liu, K.; Zhan, F.; Chen, Y.; Zhang, J.; Yu, Y.; El Saddik, A.; Lu, S.; and Xing, E. P. 2023. Stylerf: Zero-shot 3d style transfer of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8338–8348.
- 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 2024 Technical Communications*, 1–4.
- Liu, W.; Liu, Z.; Yang, X.; Sha, M.; and Li, Y. 2025. ABC-GS: Alignment-Based Controllable Style Transfer for 3D Gaussian Splatting. *arXiv preprint arXiv:2503.22218*.
- Mei, Y.; Xu, J.; and Patel, V. 2024. ReGS: Reference-based Controllable Scene Stylization with Gaussian Splatting. *Advances in Neural Information Processing Systems*, 37: 4035–4061.
- Mildenhall, B.; Srinivasan, P. P.; Ortiz-Cayon, R.; Kalantari, N. K.; Ramamoorthi, R.; Ng, R.; and Kar, A. 2019. Local light field fusion: Practical view synthesis with prescriptive

sampling guidelines. *ACM Transactions on Graphics (ToG)*, 38(4): 1–14.

Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.

Nguyen-Phuoc, T.; Liu, F.; and Xiao, L. 2022. Snerf: stylized neural implicit representations for 3d scenes. *arXiv preprint arXiv:2207.02363*.

Pang, H.-W.; Hua, B.-S.; and Yeung, S.-K. 2023. Locally stylized neural radiance fields. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 307–316. IEEE Computer Society.

Pu, G.; Xu, S.; Cao, X.; and Lian, Z. 2024. Dynamic Texture Transfer using PatchMatch and Transformers. *arXiv preprint arXiv:2402.00606*.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.

Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.

Wang, Z.; Zhao, L.; Chen, H.; Li, A.; Zuo, Z.; Xing, W.; and Lu, D. 2022. Texture reformer: Towards fast and universal interactive texture transfer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2624–2632.

Yin, K.; Gao, J.; Shugrina, M.; Khamis, S.; and Fidler, S. 2021. 3dstylenet: Creating 3d shapes with geometric and texture style variations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12456–12465.

Zhang, D.; Chen, Z.; Yuan, Y.-J.; Zhang, F.-L.; He, Z.; Shan, S.; and Gao, L. 2024. StylizedGS: Controllable Stylization for 3D Gaussian Splatting. *arXiv preprint arXiv:2404.05220*.

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, 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, Y.; He, Z.; Xing, J.; Yao, X.; and Jia, J. 2023. 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.