

Capacity-Controlled Multi-View Stylization of 3D Gaussian Splatting

Zhihao Wen¹, Yixin Yang¹, Bojian Wu², Yang Zhou^{1*},
Dani Lischinski³, Daniel Cohen-Or¹, and Hui Huang¹

¹ Guangdong Provincial Key Laboratory of Visual Media and Multidimensional Intelligence, CSSE, Shenzhen University, China

² Tencent Games, China

³ The Hebrew University of Jerusalem, Israel
<https://vcc2310.github.io/SceneStyler/>

Abstract. While 3D Gaussian Splatting (3DGS) provides an efficient and explicit representation for novel view synthesis, enforcing stylistic coherence across viewpoints remains challenging. Existing 3D stylization methods typically apply 2D feature-matching losses independently per rendered view, which leads to unstable style allocation, many-to-one feature reuse, and limited cross-view consistency. We propose a capacity-controlled framework for multi-view stylization of 3DGS, grounded in optimal transport. Specifically, we reformulate local style matching as a semi-balanced optimal transport problem. By introducing explicit column-capacity constraints with tunable strength, our formulation mitigates many-to-one matching and enables controllable allocation of style features. This transport-based objective provides a principled mechanism for balancing feature coverage and stylistic diversity while maintaining stable correspondences across viewpoints. To further enhance cross-view coherence, we incorporate a novel cross-view matching guidance to constrain correspondences between scene content and style patterns. In addition, we introduce several geometric regularizations to enhance the vanilla 3DGS, thereby enabling optimized Gaussian primitives to represent finer-grained textures during stylization. Extensive experiments demonstrate that our approach significantly improves multi-view stylistic consistency and produces stable, expressive 3D stylizations while preserving the core semantic structure of the scene.

Keywords: 3D Gaussian Splatting · Style Transfer · Optimal Transport

1 Introduction

The advent of 3D Gaussian Splatting (3DGS) [15] has marked a significant breakthrough in novel view synthesis, enabling real-time photorealistic rendering with exceptional quality. The explicit and structured representation of 3DGS also enables intuitive appearance editing of scenes [24, 39]. This capability naturally

* Corresponding author

extends to 3D stylization: transferring the artistic style of a 2D reference image to a 3D scene while preserving its content structure. However, ensuring *detailed and consistent stylization across multiple viewpoints* remains a fundamental challenge, limiting the practical deployment of such technology in augmented reality, virtual production, and immersive digital art.

Existing 3DGS stylization methods typically apply 2D style losses independently to each rendered view [8, 11, 24, 25, 39]. While effective at transferring local texture details, this view-independent paradigm suffers from severe *multi-view inconsistency* for two main reasons: First, *within each individual view*, standard feature matching strategies such as Nearest Neighbor Feature Matching (NNFM) [40] suffer from the *many-to-one matching problem*: multiple content features are matched to the same dominant style feature, resulting in repetitive local textures and limited style diversity. Second, *across different views*, the same 3D scene content projected to different viewpoints may be matched to *different* style patterns due to viewpoint-dependent feature variations, causing blurred or inconsistent stylization when observed from novel angles.

Recent efforts to address multi-view consistency have explored using DINO features [22] or geometric constraints [24]. However, these approaches only enforce consistency *alongside* the matching process—either by augmenting image features with global descriptors or regularizing geometry—without addressing the *matching mechanism itself*. Consequently, they neither prevent many-to-one matching within individual views, nor explicitly constrain the same 3D content to match consistent style patterns across viewpoints.

In this work, we propose a principled framework that tackles both root causes of multi-view inconsistency. First, we introduce a **Capacity-Controlled Feature Transport (CCFT) loss** based on **semi-balanced optimal transport**. By formulating style matching as an optimal transport problem with explicit **column-capacity constraints**, our method *systematically guarantees* the suppression of many-to-one matching within each view, while providing *tunable control* over style pattern diversity via the capacity parameter. Second, to address cross-view misalignment, we propose a **Cross-View Matching Guidance** that regularizes feature correspondences across adjacent viewpoints, explicitly constraining the same 3D content to match consistent style patterns regardless of viewpoint variations. To support fine-grained texture representation, we additionally introduce a set of geometric constraints into vanilla 3DGS to yield smaller, more uniform Gaussian primitives, establishing a robust, finer geometric foundation for high-quality stylization.

In summary, our main contributions are:

- We introduce a novel Capacity-Controlled Feature Transport (CCFT) loss based on semi-balanced optimal transport. Our formulation suppresses many-to-one matching via explicit column-capacity constraints, with a tunable strength that controls the diversity of style patterns.
- We propose Cross-View Matching Guidance to enhance multi-view consistency, explicitly encouraging adjacent viewpoints to match the same 3D content to consistent style patterns.

- We develop an enhanced scene reconstruction procedure that produces smaller, more uniform Gaussians to represent fine-grained surface textures, providing a robust foundation for high-quality stylization.
- Extensive experiments demonstrate superior visual quality and multi-view consistency of our approach compared to state-of-the-art methods.

2 Related Work

Neural Style Transfer. Creating stylized/artistic content is a hot topic in vision and graphics. Since the seminal work of Gatys et al. [9], significant progress has been made in neural style transfer for images. Existing work mainly focuses on how to effectively model image content and represent style. Beyond the classical statistic-based Gram loss [9], many successful style losses have been proposed, including histogram loss [32], Wasserstein loss [10,18], nearest neighbor matching loss [20,42], GAN loss [43], token-based loss on DINO-ViT [34], and attention features-based losses on diffusion networks [44]. Despite these advances, directly applying such 2D methods to 3D stylization remains challenging, particularly in maintaining multi-view consistency and preserving scene geometry.

3D Style Transfer. The advent of Neural Radiance Field (NeRF) [27] has made neural style transfer feasible for 3D scenes. Leveraging this representation [3, 27, 29, 38, 41], NeRF-based 3D stylization has been explored through both iterative [4, 13, 21, 30, 40] and feed-forward methods [5, 16, 23, 36]. Among these works, ARF [40] achieves high-quality and view-consistent stylization by introducing a Nearest-Neighbor Feature Matching (NNFM) loss. However, NeRF-based methods are generally limited by their lengthy training and rendering costs.

Recently, 3DGS [15] has rapidly emerged as a cornerstone in 3D vision thanks to its outstanding rendering efficiency and quality, paving the way for 3DGS-based style transfer. Some works follow the feature statistics paradigm: Style-Gaussian [24] builds upon AdaIN [12], incorporating feature embedding and a 3D decoder to achieve feed-forward stylization, yet struggles to capture fine-grained patterns, resulting in limited style quality. SGSST [8] introduces a multiscale style loss based on the Gram matrix [9], but the global statistics it uses often lead to content-style mismatch. Other works adapt the nearest-neighbor feature matching to 3DGS: StylizedGS [39] augment the NNFM loss [40] with additional guidance to enhance user controllability; G-Style [19] incorporates low-frequency features and a densification strategy to improve the expression of style patterns. Nonetheless, these methods tend to produce monotonous styles due to the inherent “many-to-one” nature of NNFM. ABC-GS [25] mitigates this shortcoming via feature alignment, but also over-smooths high-frequency features. CLIPGaussian [11] takes a different approach by leveraging CLIP model [31] for multi-modal style guidance. While demonstrating strong generalization, the CLIP-based method falls short in style fidelity compared to VGG feature-based methods on image style references.

It is worth noting that, all the aforementioned 3DGS stylization methods ignored style coherence across viewpoints. MultiStyleGS [22] augmented image




Fig. 1: Method Overview. Our framework contains two stages: enhanced reconstruction and capacity-controlled stylization. In the reconstruction stage, we optimize the 3D Gaussians using estimated depth as geometric constraints to achieve more accurate scene reconstruction. Meanwhile, we regularize the scale and shape of the primitives to better represent stylized textures later. Then, in the stylization stage, we introduce a Capacity-Controlled Feature Transport (CCFT) loss and cross-view matching guidance to obtain coherent stylized views. With a content loss that preserves scene structure, we update the color parameters of 3D Gaussians, yielding the final stylized 3D scene.

features with DINO feature, aiming to improve the uniqueness of rendered features. StylizedGS [39] proposed a depth preservation loss to enhance geometric consistency. Nonetheless, there is still no work that directly improves multi-view consistency within the feature-matching mechanism.

In this work, we focus on addressing the primary obstacles affecting coherent 3DGS stylization. Our method improves reconstruction quality for better stylization, solves the many-to-one matching issue, and addresses multi-view consistency mechanistically, finally producing superior stylized results.

3 Method

Given a set of multi-view images, our goal is to reconstruct the scene with 3DGS [15] and transfer the scene’s appearance into the style of a user-provided reference image. As depicted in Fig. 1, we first train the 3DGS with additional geometric regularizers to get an enhanced scene representation; then, we stylize the Gaussian primitives view-by-view using a novel optimal-transport-based style loss, which addresses the many-to-one matching problem and multi-view consistency, yielding high-quality, coherent scene stylization.

3.1 Preliminary of Gaussian Splatting

3DGS [15] models a scene explicitly as a collection of anisotropic 3D Gaussian primitives. Each primitive is defined by a set of parameters: a center position μ_i ,

a covariance matrix Σ_i , an opacity α_i , and view-dependent color $\mathbf{c}_i$ represented using spherical harmonics coefficients. During rendering, each 3D Gaussian is projected into screen space through an affine transformation. The color of a pixel is then computed by blending the contributions of all Gaussians that overlap the pixel, ordered by their depth, using the alpha blending formula:

$$P_{pixel} = \sum_i \mathbf{c}_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j) \quad (1)$$

To achieve high performance, 3DGS employs a tile-based rasterization coupled with a CUDA-accelerated differentiable renderer, enabling rapid image synthesis and gradient computation while balancing high-quality rendering.

3.2 Enhanced Scene Reconstruction

While vanilla 3DGS can faithfully synthesize novel views of the original scene, its highly anisotropic nature often leads to irregular distributions and deformations of individual Gaussians, which can severely impact stylization quality. StylizedGS [39] and ABC-GS [25] identify a similar issue, and both propose periodically filtering out Gaussians with low opacity or large scale before the stylization stage. However, they didn't choose to holistically correct the scene geometry. In contrast, we introduce a series of regularization terms targeting geometric parameters during the reconstruction stage. Coupled with the density control mechanism, we perform an enhanced reconstruction to improve overall scene geometry and regularize the shapes of individual 3D Gaussians, facilitating stable, high-quality stylization of the next step.

First, we employ a depth loss to enhance the spatial accuracy of 3DGS, encouraging each Gaussian to concentrate on the valid objects, mitigating the floaters within the scene. Specifically, we compute the L1 distance between the rendered depth and the estimated depth:

$$\mathcal{L}_{depth} = \|D_{render} - D_{pred}\|_1, \text{ where } D_{render} = \sum_i d_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j) \quad (2)$$

Here d_i is the depth of the i -th Gaussian primitive, and the predicted depth D_{pred} is acquired from the training images using a pretrained depth estimator [37].

Beyond the depth, the shape and scale of Gaussian primitives directly affect the quality of stylization. To efficiently represent high-frequency textures in real views, vanilla 3DGS typically optimizes Gaussians into highly irregular shapes. However, extremely elongated primitives can introduce geometric artifacts in stylized views. We, instead, introduce a shape regularization term:

$$\mathcal{L}_{shape} = \frac{1}{N_{gs}} \sum_{i=1}^{N_{gs}} [S_{max}^i - \lambda_s S_{min}^i]_+, \text{ where } [x]_+ = \max\{x, 0\} \quad (3)$$




Fig. 2: Capacity-Controlled Feature Transport. The capacity of each style feature is represented by the size of its node, with larger nodes indicating greater capacity. (a) When there’s no capacity constraint, multiple rendering features are mapped to the same style feature, and matching degrades to nearest-neighbor search, resulting in the many-to-one issue. (b) With capacity control, the matching of each style feature tends to be equalized, and more style features are utilized.

where N_{gs} is the total number of Gaussian primitives, S_{max}^i and S_{min}^i denote the longest and the shortest axis of the i -th Gaussian primitive, respectively. The hyperparameter λ_s represents the expected ratio between the longest and the shortest axis. With this regularization, the shape uniformity of Gaussian primitives can be fairly enforced.

Furthermore, the original 3DGS tends to produce large-scale Gaussian primitives to efficiently represent regions with fewer textures in real views. However, large-scale primitives typically occupy large spaces and have low density, making it difficult to express detailed style features. Therefore, we also introduce a scale regularization term:

$$\mathcal{L}_{size} = \frac{1}{N_{gs}} \sum_{i=1}^{N_{gs}} \left([S_{max}^i - \beta_1]_+ + [\beta_2 - S_{min}^i]_+ \right) \quad (4)$$

where β_1 and β_2 are the hyperparameters that represent the upper and lower bounds of the longest axis, respectively. This regularization can normalize large-scale Gaussian primitives to smaller ones within a bounded range, thereby capturing stylistic details.

The total loss function used for our enhanced scene reconstruction is:

$$\mathcal{L}_{rec} = \mathcal{L}_{photo} + \lambda_{depth} \mathcal{L}_{depth} + \lambda_{shape} \mathcal{L}_{shape} + \lambda_{size} \mathcal{L}_{size} \quad (5)$$

where $\mathcal{L}_{photo}$ is the original training loss of 3DGS, *i.e.*, the linear combination of L1 loss and SSIM loss between the rendered image and the real image.

3.3 Capacity-Controlled Scene Stylization

As discussed previously, while the NNFM loss [40] produces more visually appealing results in multi-view stylization, it still suffers from overuse of a limited

subset of style features, i.e., the many-to-one issue, and it completely ignores multi-view consistency. To address these issues, we first reformulate the nearest neighbor matching problem as a local feature assignment task solved via optimal transport, and then introduce a novel cross-view guidance to further enhance matching coherence across different viewpoints.

Capacity-Controlled Feature Transport. To prevent many-to-one matching, we formulate style transfer as a capacity-constrained optimal transport problem, where each style feature has a tunable capacity limit restricting the number of matched rendered features. This gives rise to our Capacity-Controlled Feature Transport (CCFT) loss (Fig. 2), which inherently prevents degenerate matching while enabling flexible control over style diversity via the capacity parameter.

Formally, let $X = \{x_i\}_{i=1}^N$ denote the set of rendered image features and $Y = \{y_j\}_{j=1}^M$ denote the set of style image features. For each rendered feature, we search for K nearest neighbor style features using cosine similarity, with indices denoted as $\text{KNN}(i)$, which forms a bipartite sparse support:

$$E = \{(i, j) \mid j \in \text{KNN}(i)\} \quad (6)$$

We then solve an entropy-regularized semi-balanced optimal transport problem on the bipartite sparse support:

$$\begin{aligned} T = \arg \min_{T \geq 0} & \left[\langle T, C \rangle_E - \varepsilon H(T)_E + \tau \text{KL}(T^\top \vec{1} \| b) \right] \\ \text{s.t. } & T \vec{1} = a, \quad a_i = \frac{1}{N} \end{aligned} \quad (7)$$

Here, T is the transport matrix to be optimized, and C is the cost matrix representing the pairwise distances between rendered and style features. $H(T)$ is the entropy regularization term, with parameter ε controls its strength. a is the row capacity constraint, set to a uniform distribution. b is the column capacity constraint, defined as a uniform distribution over the reachable column set:

$$b_j = \begin{cases} \frac{1}{|J|}, & j \in J \\ 0, & j \notin J \end{cases} \quad \text{where } J = \{j : \exists i, (i, j) \in E\} \quad (8)$$

The column capacity constraint is softly enforced using KL divergence, with parameter τ controlling the degree of uniform allocation among style features. After iteratively solving for the optimal transport matrix T using the Sinkhorn-Knopp algorithm [7], our CCFT loss is finally computed as the transport cost under the optimal plan:

$$\mathcal{L}_{\text{CCFT}} = \langle T, C \rangle_E = \sum_{(i,j) \in E} T_{ij} C_{ij} \quad (9)$$

Cross-view Matching Guidance. Ensuring multi-view consistency is crucial for high-quality stylization. Although the CCFT loss enables smooth, controllable feature assignment at each iteration, content changes across viewpoints can




Fig. 3: Cross-view Matching Guidance. We introduce a guidance map to improve matching coherence across viewpoints by reaggregating rendered features according to the transport matrix. Here, for example, at viewpoint $v - 1$, suppose rendered features $1^{(v-1)}$ & $2^{(v-1)}$ are assigned to style feature C in the optimal transport. Then, the guidance feature $g^{(v)}$ for C is aggregated from $1^{(v-1)}$ & $2^{(v-1)}$ according to the transport plan. Finally, at viewpoint v , the guided distance between $1^{(v)}$ and C is defined as the weighted sum of the distance between $1^{(v)}$ & C and the distance between $1^{(v)}$ & $g^{(v)}$.

lead to unstable bipartite sparse support when searching for nearest-neighbor solely based on the similarity between rendered and style features, thereby compromising multi-view consistency. Inspired by Guided Correspondence Distance [42], we incorporate a cross-view matching guidance in the distance computation of nearest neighbor searching:

$$\begin{aligned} \text{Distance}(x_i^{(v)}, y_j) &= \text{Sim}(x_i^{(v)}, y_j) + \lambda_{\text{guide}} \text{Sim}(x_i^{(v)}, g_j^{(v)}) \\ \text{where } g_j^{(v)} &= \sum_i T_{ij}^{(v-1)} x_i^{(v-1)} \end{aligned} \quad (10)$$

Here $\text{Sim}(\cdot, \cdot)$ represents cosine distance, and the superscript (v) denotes the iteration at viewpoint v . The modified/guided distance metric jointly considers the distance between rendered and style features, and the distance of rendered features to a guidance map computed from the previous iteration. As illustrated in Fig. 3, the guidance map for viewpoint v is the weighted projection of the rendered features from viewpoint $v - 1$ (i.e., the viewpoint of previous iteration) onto the style features according to the optimal transport matrix. The element of the guidance map indicates which rendered features are matched to the corresponding style feature at viewpoint $v - 1$. This guidance encourages similar rendered features across different viewpoints to match the same corresponding style feature, thereby stabilizing sparse bipartite support throughout optimization and improving multi-view consistency.

Content Loss. For content preservation of the 3D scene, we calculate the MSE distance between rendered features $X = \{x_i\}_{i=1}^N$ and content features $\hat{X} =$

$\{\hat{x}_i\}_{i=1}^N$ as the content loss, which has the following form:

$$\mathcal{L}_{content} = \frac{1}{N} \sum_{i=1}^N (x_i - \hat{x}_i)^2 \quad (11)$$

Finally, the total loss used for stylization is a weighted combination of the CCFT Loss and the content loss. To further preserve the scene’s structure, we fix the geometric attributes of 3DGS during the stylization, optimizing only its color attributes. Additionally, we apply color transfer following [40] before and after the stylization stage to maintain color consistency with the style image.

4 Experiments

Implementation Details. Our implementation is based on the original 3DGS and introduces three regularizations during reconstruction, where the parameters $\{\lambda_s, \beta_1, \beta_2, \lambda_{shape}, \lambda_{size}\}$ are set to $\{3.5, 0.1, 0.001, 10, 1\}$, and λ_{depth} decays exponentially from 1 to 0.01. For stylization, we use a pre-trained VGG-16 network [33] as the feature extractor, and we set parameters $\{K, \varepsilon, \tau, \lambda_{guide}, \lambda_{CCFT}, \lambda_{content}\}$ to $\{16, 0.05, 0.5, 1, 30, 0.005\}$. All experiments are performed on a workstation equipped with a single NVIDIA RTX 4090 GPU.

Datasets. We conduct comprehensive experiments on multiple datasets, including LLFF [26], Tanks and Temples (T&T) [17], and Mip-NeRF 360 [2], encompassing both forward-facing scenes and unbounded 360° environments. To validate robustness across diverse scenarios, we use stylistic images from various artistic categories from the WikiArt dataset [24] as style references.

Baselines. For comparative analysis, we select state-of-the-art 3DGS-based style transfer methods as baselines, including StylizedGS [39], ABC-GS [25], SGSST [8], CLIPGaussian [11], and StyleGaussian [24]. With the exception of StyleGaussian, which adopts a feed-forward approach, all comparative methods are optimization-based. For all baselines, we use their officially released code.

4.1 Experimental Results

Qualitative Results. Fig. 4 shows a gallery of visual comparisons between our method and the baselines. StylizedGS employs the NNFM Loss to stylize scenes, where the many-to-one issue leads to a large amount of repetitive textures; for example, numerous yellow stripe-like textures appear (column 1). ABC-GS uses FAST Loss to present more style patterns, but the global linear transformation tends to cause erroneous matches in regions with insufficient local semantics; for example, sunflower seed patterns are incorrectly matched to the triceratops’s mouth (column 3). SGSST achieves a highly global stylization, but when there are significant differences in local structure between content and style, the content structure is severely degraded; for instance, it is difficult to distinguish the shapes of petals and leaves (column 4). CLIPGaussian constructs the style loss using the CLIP model; however, the model exhibits semantic illusions between



Fig. 4: Qualitative comparison. Our method demonstrates a significant advantage in reproducing the texture details and brushstrokes in the style references while still preserving the semantic structure of the input scene.

images, leading to patterns and colors that are inconsistent with the style image. StyleGaussian utilizes AdaIN for instant stylization, but this alignment of low-order statistics makes it difficult to capture the specific style details.

Compared to the baselines, our stylization results demonstrate significant advantages in style matching and content preservation. For example, in the first and second scenes, our method accurately reproduces the painting’s brushstrokes; in the last scene, it reproduces texture details, particularly the white droplets in the waves, more evenly than other methods. Meanwhile, our results clearly preserve the semantic content of the input scene; see, e.g., the truck’s wheels (column 2) and the staircases’ handrails (columns 3 & 5). These qualitative results provide strong validation of the effectiveness of our stylization strategy.

Table 1: Quantitative comparison, where ArtFID evaluates style quality, Structure Loss measures content preservation, and the last two metrics indicate multi-view consistency. “Ours (3DGS)” denotes the ablated version of our method using the original 3DGS for scene reconstruction. Our full method achieved the best across all metrics.

Method	ArtFID↓	Structure Loss↓	Short Met3R↓	Long Met3R↓
StylizedGS	23.068	0.0462	0.1201	0.2858
ABC-GS	25.026	0.0357	<u>0.1198</u>	<u>0.2822</u>
SGSST	25.143	0.0568	0.1384	0.3064
CLIPGaussian	27.293	0.0431	0.1353	0.2903
StyleGaussian	31.855	0.0425	0.1284	0.2982
Ours (3DGS)	24.133	<u>0.0320</u>	<u>0.1198</u>	0.2795
Ours	22.801	0.0318	0.1196	0.2795

Table 2: User study. We presented the 6-method results for each scene and asked users to select the top three for each evaluation aspect, with the best one receiving a score of 5, the second 3, the third 2, and 1 for the remaining results. We showed each user 6 scenes of different styles, and, in total, 25 participants were involved, yielding 150 rankings per evaluation aspect. Note that, for view consistency, we show rendered videos, while for style and content performance, we only display rendered views.

	StylizedGS	ABC-GS	SGSST	CLIPGaussian	StyleGaussian	Ours
Style	13.25%	13.08%	<u>25.04%</u>	8.55%	7.86%	32.22%
Content	8.29%	<u>19.04%</u>	10.26%	12.91%	11.54%	37.61%
Consistency	11.02%	<u>19.83%</u>	11.03%	11.03%	9.32%	37.61%

Quantitative Results. We conduct quantitative comparisons in three aspects: stylization quality, content preservation, and view consistency. Specifically, we use ArtFID [35] to evaluate the stylization quality as in previous work. To evaluate content preservation, considering the substantial domain gap between stylized and real images, we adopt a cross-domain structural loss [34] to measure the structural similarity between the stylized and real scene. For multi-view consistency, we utilize Met3R [1] to compute consistency error between two views at both short range (one-view gap) and long range (seven-view gap). We select five forward-facing scenes from LLFF [26] and five 360-degree scenes from T&T [17] and Mip-NeRF 360 [2], and stylize each scene using six different style images, producing 60 test cases per method. As evidenced in Table 1, our method outperforms baselines across all metrics.

Since style transfer is inherently subjective, we also conducted a user study to assess human preferences. Our user study contains three evaluation aspects: style quality, content preservation, and multi-view consistency. For style quality and content preservation, we present the user with a rendering view before stylization, along with the corresponding stylized views produced by the five baselines and our approach, in random order. For view consistency, we show users rendered videos. A total of 25 participants have been involved, with 450 result rankings




Fig. 5: Ablation study on loss functions.

collected. The final preference ratio is summarized in Table 2. Again, our method leads the user study by a large margin, validating its high perceptual quality.

4.2 Ablation Study

Loss function. We first compare our CCFT Loss to the NNFM Loss [40], FAST Loss [25] and GRAM Loss [9]. Specifically, we replace the style loss with each alternative loss and use the style image and scene from the fifth column in Fig. 4. As shown in Fig. 5, our CCFT loss successfully captures the sea-wave and water-splash patterns in the style image, faithfully preserving fine stylistic details. Although the FAST loss can also capture relevant style patterns, it produces less detailed results and causes artifacts. In contrast, the NNFM Loss only matches monotonous background patterns of the style image, while the GRAM Loss fails to transfer specific local style patterns and results more artifacts.

Capacity control. To demonstrate the capacity control capability of our CCFT Loss, Fig. 6 presents the stylization results with different capacity constraint strengths under the same style and scene. When the capacity constraint parameter τ is set to 0.01, the constraint is weak, and the transport plan tends to simply minimize the matching cost, assigning a large number of content features to a few lowest-cost style features, resulting in highly repetitive striped textures across large areas of the scene. When the parameter τ is increased to 0.1, the capacity constraint on lowest-cost style features is enhanced, which encourages content features to be distributed across a broader set of style features, leading to the emergence of swirl-like textures in the scene. When the parameter τ is





Fig. 6: Ablation study for capacity-controlled feature transport.

further increased to 1, content features are assigned to style features in a more balanced manner, thereby introducing more detailed circular textures.

Cross-view matching guidance. We evaluate the effectiveness of our cross-view guidance in improving multi-view consistency through additional quantitative experiments. Table 3 reports the qualitative results of the cross-view guidance under different parameter settings. The results show that as the strength of guidance increases, both short-range and long-range consistency exhibit improvements. This indicates that cross-view guidance can effectively mitigate inconsistent textures arising from viewpoint changes during multi-view training, aligning with our expectations. However, we also observe that as the parameter is gradually increased, the fidelity of the scene structure declines slightly. To explain this phenomenon, we further analyze the stylization results in Fig. 7. Essentially, this trade-off in content stems from the “2D-3D gap” between 2D views and 3D scenes: cross-view guidance tends to cover part of the content details in the scene with consistent 2D style patterns, making the projected textures of the 3D scene appear more consistent across viewpoints, which “damages” some original content boundaries. Therefore, in practice, we set the guidance weight to a reasonable value ($\lambda_{guide} = 1$) during stylization.

Table 3: Ablation study for cross-view matching guidance.

λ_{guide}	ArtFID↓	Structure Loss↓	Short M _{Et} 3R↓	Long M _{Et} 3R↓
0	22.809	0.0312	0.1204	0.2807
1	22.801	<u>0.0318</u>	<u>0.1196</u>	<u>0.2795</u>
10	<u>22.805</u>	0.0323	0.1191	0.2786

**Fig. 7:** Ablation on different strengths of cross-view matching guidance.

Enhanced reconstruction. To evaluate the contribution of our enhanced reconstruction, we use the original 3DGS with our capacity-controlled stylization. As shown in Tab. 1, this ablated version achieves quite competitive results, especially in multi-view consistency, compared to the baselines. We further present a more detailed ablation study of each constraint in the supplementary.

5 Conclusions

In this paper, we present a capacity-controlled framework for multi-view stylization of 3D Gaussian Splatting. Addressing multi-view inconsistency in existing 3DGS stylization methods, we propose a Capacity-Controlled Feature Transport (CCFT) loss that prevents many-to-one matching while enabling tunable control over style diversity. Furthermore, we introduce cross-view matching guidance to constrain the same 3D content to match consistent style patterns. We also develop an enhanced scene reconstruction with geometric regularizations to establish a robust foundation for high-quality stylization.

While our method produces superior stylization results, several limitations remain. First, the cross-view matching guidance introduces a trade-off between multi-view consistency and content preservation. Second, our method requires per-scene optimization. Future work will explore feed-forward approaches to enable real-time stylization while preserving the controllability and consistency benefits of our framework.

Acknowledgments

This work was supported in part by ICFCRT (W2441020), Guangdong Basic and Applied Basic Research Foundation (2023B1515120026), SZU Teaching Reform Research Project (JG2026011), and Guangdong Provincial Key Laboratory of Visual Media and Multidimensional Intelligence.

References

1. Asim, M., Wewer, C., Wimmer, T., Schiele, B., Lenssen, J.E.: Met3r: measuring multi-view consistency in generated images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 6034–6044 (2025)
2. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5460–5469 (2022)
3. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: *European conference on computer vision (ECCV)*. pp. 333–350 (2022)
4. Chen, W., Zha, Z., Wang, S., Ali, L., Wen, B., Yuan, X., Zhou, J., Zhu, C.: Texture-consistent 3d scene style transfer via transformer-guided neural radiance fields. *IEEE Transactions on Image Processing (TIP)* **34**, 7193–7208 (2025)
5. Chiang, P.Z., Tsai, M.S., Tseng, H.Y., Lai, W.S., Chiu, W.C.: Stylizing 3d scene via implicit representation and hypernetwork. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 1475–1484 (2022)
6. Chizat, L., Peyré, G., Schmitzer, B., Vialard, F.X.: Scaling algorithms for unbalanced transport problems. *arXiv preprint arXiv:1607.05816* (2016)
7. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. In: *Advances in Neural Information Processing Systems (NIPS)*. pp. 2292–2300 (2013)
8. Galerne, B., Wang, J., Raad, L., Morel, J.M.: Sgsst: Scaling gaussian splatting style transfer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 26535–26544 (2025)
9. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*. pp. 2414–2423 (2016)
10. Heitz, E., Vanhoey, K., Chambon, T., Belcour, L.: A sliced wasserstein loss for neural texture synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9412–9420 (2021)
11. Howil, K., Borycki, P., Dziarmaga, T., Mazur, M., Spurek, P., et al.: Clipgaussian: Universal and multimodal style transfer based on gaussian splatting. *arXiv preprint arXiv:2505.22854* (2025)

12. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE international conference on computer vision (ICCV). pp. 1501–1510 (2017)
13. Huang, Y.H., He, Y., Yuan, Y.J., Lai, Y.K., Gao, L.: 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 (CVPR). pp. 18342–18352 (2022)
14. Kantorovich, L.V.: On the translocation of masses. Dokl. Akad. Nauk. USSR (NS) **37**, 199–201 (1942)
15. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics (TOG) **42**(4) (July 2023)
16. Kim, G., Youwang, K., Oh, T.H.: Fprf: Feed-forward photorealistic style transfer of large-scale 3d neural radiance fields. In: Proceedings of the aaai conference on artificial intelligence (AAAI). pp. 2750–2758 (2024)
17. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: benchmarking large-scale scene reconstruction. ACM Transactions on Graphics (TOG) **36**(4), 1–13 (2017)
18. Kolkin, N., Salavon, J., Shakhnarovich, G.: Style transfer by relaxed optimal transport and self-similarity. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10051–10060 (2019)
19. Kovács, Á.S., Hermosilla, P., Raidou, R.G.: G-style: Stylized gaussian splatting. In: Computer Graphics Forum. vol. 43, p. e15259. Wiley Online Library (2024)
20. Li, C., Wand, M.: Combining markov random fields and convolutional neural networks for image synthesis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2479–2486 (2016)
21. Li, W., Wu, T., Zhong, F., Oztireli, C.: Arf-plus: Controlling perceptual factors in artistic radiance fields for 3d scene stylization. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 2301–2310 (2025)
22. Lin, Y., Lei, J., Jia, K.: Multi-stylegs: stylizing gaussian splatting with multiple styles. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). pp. 5289–5297 (2025)
23. Liu, K., Zhan, F., Chen, Y., Zhang, J., Yu, Y., El Saddik, A., Lu, S., Xing, E.P.: Stylerf: Zero-shot 3d style transfer of neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8338–8348 (2023)
24. Liu, K., Zhan, F., Xu, M., Theobalt, C., Shao, L., Lu, S.: Stylegaussian: Instant 3d style transfer with gaussian splatting. In: SIGGRAPH Asia 2024 Technical Communications. pp. 1–4 (2024)
25. Liu, W., Liu, Z., Yang, X., Sha, M., Li, Y.: Abc-gs: Alignment-based controllable style transfer for 3d gaussian splatting. IEEE International Conference on Multi-media and Expo (ICME) (2025)
26. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: practical view synthesis with prescriptive sampling guidelines. ACM Transactions on Graphics (TOG) **38**(4), 1–14 (2019)
27. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European Conference on Computer Vision (ECCV). pp. 405–421 (2020)

28. Monge, G.: Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.* pp. 666–704 (1781)
29. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)* **41**(4), 1–15 (2022)
30. Nguyen-Phuoc, T., Liu, F., Xiao, L.: Snerf: stylized neural implicit representations for 3d scenes. *ACM Transactions on Graphics (TOG)* **41**(4), 1–11 (2022)
31. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning (ICML)*. pp. 8748–8763 (2021)
32. Risser, E., Wilmot, P., Barnes, C.: Stable and controllable neural texture synthesis and style transfer using histogram losses. *arXiv preprint arXiv:1701.08893* (2017)
33. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: *International Conference on Learning Representations (ICLR)* (2015)
34. Tumanyan, N., Bar-Tal, O., Bagon, S., Dekel, T.: Splicing vit features for semantic appearance transfer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10748–10757 (2022)
35. Wright, M., Ommer, B.: Artfid: Quantitative evaluation of neural style transfer. In: *DAGM German Conference on Pattern Recognition (GCPR)*. p. 560–576 (2022)
36. Xu, H., Chen, W., Xiao, F., Sun, B., Kang, W.: Styledyrf: Zero-shot 4d style transfer for dynamic neural radiance fields. *arXiv preprint arXiv:2403.08310* (2024)
37. Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H.: Depth anything v2. In: *Proceedings of the 38th International Conference on Neural Information Processing Systems (NeurIPS)*. pp. 21875–21911 (2024)
38. Yu, A., Fridovich-Keil, S., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*. pp. 5501–5510 (2022)
39. Zhang, D., Yuan, Y.J., Chen, Z., Zhang, F.L., He, Z., Shan, S., Gao, L.: Stylizedgs: Controllable stylization for 3d gaussian splatting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **47**(12), 11961–11973 (2025)
40. Zhang, K., Kolkin, N., Bi, S., Luan, F., Xu, Z., Shechtman, E., Snavely, N.: Arf: Artistic radiance fields. In: *European Conference on Computer Vision (ECCV)*. pp. 717–733 (2022)
41. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492* (2020)
42. Zhou, Y., Chen, K., Xiao, R., Huang, H.: Neural texture synthesis with guided correspondence. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 18095–18104 (2023)
43. Zhou, Y., Chen, Z., Huang, H.: Deformable one-shot face stylization via dino semantic guidance. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7787–7796 (2024)
44. Zhou, Y., Gao, X., Chen, Z., Huang, H.: Attention distillation: A unified approach to visual characteristics transfer. In: *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*. pp. 18270–18280 (2025)