# Stylos: Multi-View 3D Stylization with Single-Forward Gaussian Splatting

Anonymous authors
Paper under double-blind review

# Abstract

We present Stylos, a single-forward 3D Gaussian framework for 3D style transfer that operates on unposed content, from a single image to a multiview collection, conditioned on a separate reference style image. Stylos synthesizes a stylized 3D Gaussian scene without per-scene optimization or precomputed poses, achieving geometry-aware, view-consistent stylization that generalizes to unseen categories, scenes, and styles. At its core, Stylos adopts a Transformer backbone with two pathways: geometry predictions retain self-attention to preserve geometric fidelity, while style is injected via global cross-attention to enforce visual consistency across views. With the addition of a voxel-based 3D style loss that aligns aggregated scene features to style statistics, Stylos enforces view-consistent stylization while preserving geometry. Experiments across multiple datasets demonstrate that Stylos delivers high-quality zero-shot stylization, highlighting the effectiveness of global style-content coupling, the proposed 3D style loss, and the scalability of our framework from single view to large-scale multi-view settings. Our codes will be fully open-sourced soon.

# 1 Introduction

Image guided 3D stylization aims to preserve scene geometry and cross-view consistency while transferring the reference style. With the rise of immersive content, augmented and virtual reality, demand for this capability is growing. Nevertheless, achieving reliable 3D stylization remains challenging and continues to attract significant research attention.

Early attempts rely on explicit 3D geometry such as meshes or point clouds, but their stylization quality is often limited by geometric reconstruction accuracy, resulting in noticeable artifacts in complex scenes (Kato et al., 2018; Cao et al., 2020). The introduction of implicit radiance fields, most notably NeRF (Mildenhall et al., 2020), enables higher-fidelity rendering, and subsequent works extended it to artistic stylization (Zhang et al., 2022; Huang et al., 2022; Liu et al., 2023), but typically require costly per-scene optimization, constraining their generalization to unseen scenes. More recently, 3D Gaussian Splatting (3DGS) has emerged as a promising explicit representation that combines high reconstruction quality with real-time rendering efficiency (Kerbl et al., 2023). Stylization approaches built on 3DGS achieve efficient multi-view consistency (Liu et al., 2024; Galerne et al., 2025), yet still struggle to generalize beyond scene-specific training.

In contrast, we introduce Stylos, a single-forward framework for 3D style transfer that eliminates the need for per-scene optimization and precomputed camera parameters, and generalizes effectively to novel categories, scenes, and styles. Stylos employs a shared Transformer backbone with two pathways: content and style images are projected into a shared feature space, where content retains self-attention for geometric reasoning, and style is injected via global cross-attention to condition aesthetics. Notably, the global cross-attention also benefits the preservation of geometric details. Geometry-related attributes, such as depth, camera intrinsics, and extrinsics, are derived from backbone features, whereas style conditioning guides the prediction of color coefficients. These outputs are estimated through prediction heads that serve as the interface between feature space and the final Gaussian representation. Prior works in 2D style transfer typically rely on distribution

alignment (Li et al., 2017), which operate on image statistics but lack explicit multi-view or 3D consistency. To address this, we explore alternative objectives and propose a voxel-based 3D style loss that aligns aggregated scene features with style statistics, providing stronger view-consistent stylization while preserving geometric fidelity.

We evaluate Stylos across different scenarios, including category-level transfer and cross-scene generalization. Our assessment spans challenging real-world benchmarks, where Stylos produces stylized renderings with high visual fidelity and consistent geometry, demonstrating robustness even in previously unseen environments. Our main contributions are threefold:

- We propose a shared-backbone design with two pathways: geometry predictions retain self-attention for geometric reasoning, while style is injected through global cross-attention to condition aesthetics.
- We introduce a voxel-based 3D style loss that aligns aggregated scene features with style statistics, enforcing cross-view coherence and geometry-aware stylization.
- We develop Stylos, a single-forward-pass pipeline for 3D style transfer from unposed inputs, scaling from a single to hundreds of views with a single style image, and achieving zero-shot generalization to unseen categories, scenes, and styles.

# 2 Related Work

## 2.1 Single-Forward 3D Reconstruction

Early CNN models (Ummenhofer et al., 2017; Zhou et al., 2018; Teed & Deng, 2020) demonstrate the feasibility of end-to-end 3D reconstruction. While the Transformer-based DUSt3R (Wang et al., 2024) has shown impressive performance, yet remains constrained by pairwise inputs and reliance on traditional optimization for pose estimation. VGGT (Wang et al., 2025a) advances this line by jointly predicting camera parameters, depth, point maps, and tracks from one to hundreds of views in a single forward pass, while AnySplat (Jiang et al., 2025) complements it with a feed-forward Gaussian splatting pipeline trained with VGGT priors to directly predict Gaussian parameters along with depth and camera parameters from uncalibrated images.

#### 2.2 3D Style Transfer and Losses

Explicit-representation approaches to 3D stylization (Kato et al., 2018; Cao et al., 2020; Guo et al., 2021; Huang et al., 2021; Höllein et al., 2022; Mu et al., 2022) have shown promising results, but their reliance on accurate geometry limits applicability in complex real-world scenes. Implicit neural representations like NeRF (Mildenhall et al., 2020) improve geometric consistency, inspiring a range of stylization methods based on stylized view supervision or auxiliary networks (Nguyen-Phuoc et al., 2022; Zhang et al., 2022; Chiang et al., 2022; Huang et al., 2022). StyleRF (Liu et al., 2023) further enhances fidelity and enables zero-shot transfer by applying style transformations in radiance-field feature space.

More recently, 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) has emerged as an efficient alternative to NeRF. StyleGaussian (Liu et al., 2024) extends 3DGS with efficient feature rendering to enable real-time stylization, while subsequent works refine 3DGS for specialized settings and efficiency improvements (Mei et al., 2024; Zhang et al., 2024; Yu et al., 2024; Galerne et al., 2025). In parallel, Styl3R (Wang et al., 2025b) disentangles geometry and style using separate encoders, but is restricted to 2–8 input views. In contrast, Stylos follows the shared-backbone paradigm of AdaIN (Huang & Belongie, 2017), leveraging VGGT for unified feature extraction, and scales from a single to hundreds of input views.

Underlying these architectures, style transfer objectives play a crucial role. Classical style transfer relies on Gram-matrix correlations (Gatys et al., 2016), while channel-wise statistics such as mean and variance (Li et al., 2017) offer efficient alternatives, forming the basis of AdaIN (Huang & Belongie, 2017). Extensions to video and multi-view stylization (Ruder et al., 2016; Gupta et al., 2017; Nguyen-Phuoc et al., 2022; Zhang et al., 2022) address temporal and cross-view consistency, and CLIP-based losses (Radford et al., 2021) introduce

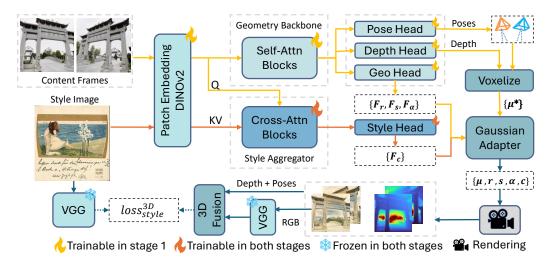


Figure 1: **Architecture overview**. Given multi-view content inputs and a style reference, Stylos enables instant 3D stylization without scene-specific training or post-optimization. A key component is the 3D style loss, matching voxelized 3D features with 2D style statistics.

semantic alignment with text or image prompts. However, most objectives remain inherently 2D and cannot guarantee consistency across complex 3D scenes. To this end, we extend statistics-based style losses into voxel space, introducing a 3D-aware objective that aligns feature distributions after multi-view fusion. This formulation complements perceptual and CLIP-based losses, enforcing geometry-aware and view-consistent stylization.

#### 3 Method

We propose **Stylos**, a transformer-based framework for stylized 3D scene reconstruction. Given a style reference image and one or more content views (also referred to as context views), Stylos predicts a set of stylized 3D Gaussian primitives together with camera parameters, enabling faithful reconstruction of the observed scene while transferring the desired style. We first formulate the problem in Sec. 3.1, then present the network architecture in Sec. 3.2, and finally detail our training strategy in Sec. 3.3 and the used style loss in Sec. 3.4.

## 3.1 Problem Formulation

We aim to disentangle geometry and style within a unified 3D scene representation. The input to Stylos consists of a set of N context views  $\{I_i\}_{i=1}^N$  of a scene, with  $N \geq 1$  and  $I_i \in \mathbb{R}^{H \times W \times 3}$ , together with a single style reference image  $S \in \mathbb{R}^{H \times W \times 3}$ . The context views provide geometric cues, while the style image specifies the desired aesthetic.

Formally, Stylos defines a conditional mapping,

$$f_{\theta}: (\{I_i\}_{i=1}^N, S) \mapsto (G, \{g_i\}_{i=1}^N),$$

where the scene representation  $G = \{(p_m, c_m)\}_{m=1}^M$  is parameterized by M anisotropic 3D Gaussians. Each Gaussian has geometry attributes  $p_m = (\mu_m, \alpha, r_m, s_m)$ , comprising 3D position  $\mu_m \in \mathbb{R}^3$ , opacity  $\alpha \in \mathbb{R}_+$ , orientation quaternion  $r_m \in \mathbb{R}^4$ , and anisotropic scale  $s_m \in \mathbb{R}^3$ , as well as a style-dependent color embedding  $c_m \in \mathbb{R}^{3 \times (k+1)^2}$  represented with spherical harmonics of degree k. In addition, Stylos jointly predicts camera parameters  $\{g_i \in \mathbb{R}^9\}_{i=1}^N$  for each input view, leveraging pose cues estimated by the VGGT backbone.

# 3.2 Network Architecture

**Architecture Overview.** At its core, Stylos employs a geometry backbone that alternates frame-wise and global attention to process multiple input views and infer geometry-related

parameters such as positions, scales, orientations, and opacities. To enable stylization, we introduce a dedicated conditioning branch, the Style Aggregator, which fuses content and style features through cross-attention and predicts style-aware color embeddings for each Gaussian. In this way, geometry remains derived solely from the backbone, while style representation is conditioned on the style reference. We next detail the geometric backbone in Sec. 3.2.1, the style aggregation module in Sec. 3.2.2, and the prediction heads in Sec. 3.2.3.

# 3.2.1 Geometric Backbone

163

164

165

166

167

168 169 170

171

172

173

174

175 176

177 178

179

180

181

182

183

184

185 186

187

188 189

195

196

197

199

200 201

202

203

204

205

206

207

208

209

210

211

212 213

214

215

The geometric backbone follows the alternating-attention design of VGGT (Wang et al., 2025a), which interleaves frame and global self-attention layers to capture both intra-frame structure and cross-view consistency. This component is kept unchanged to retain strong geometric reasoning, serving as the foundation on which we integrate style information.

#### 3.2.2 Style Aggregator

The Style Aggregator extends the geometric backbone by adapting its attention layers from self-attention to cross-attention. Both content views and the style image are embedded with a shared patch encoder, ensuring that their tokens lie in a common feature space. In this design, content tokens serve as queries while style tokens provide keys and values, so that style information conditions the token representations later used for predicting Gaussian colors. Geometry-related parameters, in contrast, are still predicted from the backbone pathway and remain unaffected by style conditioning.

Formally, let  $T \in \mathbb{R}^{B \times N \times K \times C}$  denote the content tokens from N views, each with K patches and C channels, and let  $S \in \mathbb{R}^{B \times K \times C}$  represent the style tokens. Let  $T_i$  be the token set of the i-th view, and let Attn(Q, KV) denote a standard cross-attention operation. Depending on how queries are formed, we distinguish three style-content coupling mechanisms,

(Frame-only): 
$$T_i^{\text{frame}} = \text{Attn}(T_i, S),$$
  $Q = T_i, K = V = S,$  (1)  
(Global-only):  $T^{\text{global}} = \text{Attn}(\{T_i\}_{i=1}^N, S),$   $Q = \{T_i\}_{i=1}^N, K = V = S,$  (2)  
(Hybrid):  $T^{\text{hybrid}} = [\{T_i^{\text{frame}}\}_{i=1}^N || T^{\text{global}}],$  concatenation along channel dim. (3)

(Global-only): 
$$T^{\text{global}} = \text{Attn}(\{T_i\}_{i=1}^N, S), \qquad Q = \{T_i\}_{i=1}^N, K = V = S,$$
 (2)

(Hybrid): 
$$T^{\text{hybrid}} = [\{T_i^{\text{frame}}\}_{i=1}^N || T^{\text{global}}],$$
 concatenation along channel dim. (3)

Frame-only injects style independently per view, global-only enforces a shared conditioning across all views, and hybrid concatenates the frame-level and global-level outputs along the channel dimension, yielding  $2 \times C$ -dimensional features. In our design, Stylos adopts the global-only configuration, implemented with CrossBlock layers where  $\{T_i\}_{i=1}^N$  serve as queries and S provides keys and values. The resulting features are passed to the prediction heads described in Sec. 3.2.3.

#### 3.2.3 Prediction Heads

To connect the geometric backbone and Style Aggregator with the Gaussian scene, we introduce prediction heads that translate feature tokens into explicit parameters. These heads serve as modular interfaces, keeping geometry and style distinct in feature space while ensuring their integration in the final Gaussian representation. We next describe the individual heads in detail.

Geometry Head. Geometric backbone features are passed through a DPT-style regression head that outputs the Gaussian geometry parameters  $p_m = (\mu_m, s_m, r_m, \alpha_m)$ , i.e., position, scale, orientation, and opacity (as defined in Sec. 3.1). By relying on this established design, structural predictions are derived from backbone features alone, without direct influence from style conditioning.

Style Head. The outputs of the Style Aggregator are subsequently processed by a color head to predict the spherical-harmonic coefficients  $c_m$  that define appearance. This pathway injects style information directly into Gaussian colors while leaving the geometry parameters  $p_m$  unaffected, enabling the two factors to be recombined seamlessly at the Gaussian level. Auxiliary Heads, Adapter, and Voxelization. We employ the existing VGGT camera head to estimate camera intrinsics and extrinsics, and a depth head to predict scene geometry cues, which are unprojected into 3D anchors for Gaussian placement (Kerbl et al., 2023; Ren et al., 2024). A Gaussian adapter then consolidates geometry and style outputs into a unified set of primitives  $\{(p_m, c_m)\}_{m=1}^M$  for differentiable rendering. Finally, to reduce redundancy and balance density, we follow the voxelization step introduced by AnySplat (Jiang et al., 2025), nearby points are clustered within a discretized 3D grid and fused using confidence-aware weighting. This operation depends only on the unprojected 3D points and features, and is independent of camera parameters.

## 3.3 Training Strategy

We adopt a two-stage training strategy for structure-aware stylization.

Stage 1: Geometry Pretraining. We initialize the geometric backbone of Stylos with VGGT weights (Wang et al., 2025a) and train the network end-to-end to learn geometry and photometric appearance. To avoid trivial identity mapping and improve robustness to color variations, one input view is randomly selected and color-jittered as the style reference. A frozen VGGT teacher provides pose and depth supervision. The objective combines reconstruction and distillation,

$$\mathcal{L}_{\mathrm{stage1}} = \mathcal{L}_{\mathrm{rec}} + \lambda_{\mathrm{distill}} \, \mathcal{L}_{\mathrm{distill}}.$$

Stage 2: Stylization Fine-tuning. We freeze all geometry-related modules and only update the Style Aggregator and the color head. Following ArtFlow (An et al., 2021), we use feature-level style and content losses in VGG space, matching channel-wise statistics for style and feature activations for content. We further extend these objectives with a 3D voxel-space style loss for cross-view consistency, and add a CLIP-based loss for semantic alignment. A total variation (TV) regularizer is also included to suppress high-frequency artifacts and stabilize optimization. The total loss is,

$$\mathcal{L}_{stage2} = \mathcal{L}_{rec} + \lambda_{style} \, \mathcal{L}_{stvle}^{3D} + \lambda_{cnt} \, \mathcal{L}_{content} + \lambda_{clip} \, \mathcal{L}_{clip} + \lambda_{tv} \, \mathcal{L}_{TV}.$$

#### 3.4 Style Losses

We denote by  $F_{b,s}^l \in \mathbb{R}^{C_l \times H_l \times W_l}$  the VGG feature map of the rendered (stylized) image for the b-th scene and s-th view at layer l, and by  $S_b^l \in \mathbb{R}^{C_l \times H_l \times W_l}$  the feature map of the style reference image. Building on the BN-statistics formulation (Li et al., 2017), which interprets channel-wise mean and variance as style descriptors, we progressively extend statistics matching beyond the standard 2D setting. This progression starts from independent image-level matching, moves to multi-view feature aggregation for enforcing cross-view consistency, and culminates in a voxel-space loss that directly constrains the fused 3D representation.

Image-Level Style Loss. The simplest baseline aligns each rendered frame independently with the style reference,

$$\mathcal{L}_{\text{sty}}^{\text{img}} = \frac{1}{BS} \sum_{b=1}^{B} \sum_{s=1}^{S} \sum_{l} \alpha_{l} \Big( \|\mu(F_{b,s}^{l}) - \mu(S_{b}^{l})\|_{2}^{2} + \|\sigma(F_{b,s}^{l}) - \sigma(S_{b}^{l})\|_{2}^{2} \Big). \tag{4}$$

This encourages per-frame stylization but does not enforce multi-view consistency.

Scene-Level Style Loss. To encourage consistent stylization across multiple views, we concatenate the per-view features  $\{F_{b,s}^l\}_{s=1}^S$  along the spatial dimension to form  $\widetilde{F}_b^l$ , and compute statistics on this aggregated map,

$$\mathcal{L}_{\text{sty}}^{\text{scn}} = \frac{1}{B} \sum_{b=1}^{B} \sum_{l} \alpha_{l} \Big( \|\mu(\widetilde{F}_{b}^{l}) - \mu(S_{b}^{l})\|_{2}^{2} + \|\sigma(\widetilde{F}_{b}^{l}) - \sigma(S_{b}^{l})\|_{2}^{2} \Big).$$
 (5)

This step enforces cross-view consistency but still operates in 2D feature space.

**3D Style Loss.** Finally, we fuse multi-view features into a voxel grid using differentiable unprojection, where features from different views are accumulated into spatial bins of a

discretized 3D volume. Let  $G_b^l$  denote the voxelized features for the b-th scene at layer l. We then compute style statistics directly in voxel space,

$$\mathcal{L}_{\text{sty}}^{3D} = \frac{1}{B} \sum_{b=1}^{B} \sum_{l} \alpha_{l} \Big( \|\mu(G_{b}^{l}) - \mu(S_{b}^{l})\|_{2}^{2} + \|\sigma(G_{b}^{l}) - \sigma(S_{b}^{l})\|_{2}^{2} \Big).$$
 (6)

By operating on voxelized features, this loss explicitly encodes geometry and enforces style consistency across both views and the underlying 3D structure.

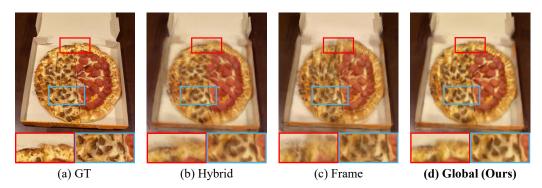


Figure 2: CO3D pizza scene comparing different style-content cross-attention strategies.

Table 1: Ablation on the style-content coupling strategies, comparing frame-level, global-level and hybrid cross-attention. The first frame of each scene is used as the style reference. Reconstruction quality is evaluated with PSNR↑, SSIM↑, and LPIPS↓ on the CO3D dataset.

Strategies		Skateboard			Pizza			Donut		
Global	Frame	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
<b>√</b>	✓	21.12	0.6858	0.2821	19.78	0.5939	0.3326	21.39	0.7198	0.3264
	$\checkmark$	20.93	0.6917	0.2912	19.72	0.5940	0.3405	21.40	0.7167	0.3340
$\checkmark$		21.68	0.7043	0.2684	20.57	0.6177	0.3110	22.09	0.7362	0.3125

#### 4 Experiment

**Datasets.** We evaluate cross-category generalization on the CO3D (Reizenstein et al., 2021) dataset by training on 17 categories and testing on 3 held-out ones, and cross-scene generalization by training on the full DL3DV-10K (Ling et al., 2024) and testing on Tanks and Temples (Knapitsch et al., 2017b). Style supervision is provided by WikiArt (WikiArt), with 50 images reserved for unseen-style evaluation.

Baselines. We compare with two representative open-source baselines, StyleRF (Liu et al., 2023) and StyleGS (Liu et al., 2024), using their released per-scene trained weights against our single DL3DV-trained model. Concurrent to Stylos, the other instant stylization solution Styl3R (Wang et al., 2025b) was proposed, but its official codes are not yet available at the time of our submission. We plan to provide comparisons once they are released.

Evaluation Metrics. Our evaluation covers three aspects. (1) To assess geometry reconstruction in the style-free setting, we report PSNR, SSIM, and LPIPS (Zhang et al., 2018) between content views and predictions, conditioned on the original first frame of each scene as the style input. (2) To measure stylization consistency, following prior work (Chiang et al., 2022), we compute LPIPS and RMSE in both short-range (consecutive frames t and t-1) and long-range (distant frames t and t-7) settings. (3) To further evaluate stylization quality, we report ArtScore (Chen et al., 2024), a recent metric specifically designed for reference-free evaluation of artness in generated images.

# 4.1 Ablation Study

In this section, we analyze the impact of each design choice on the CO3D dataset. To ensure fairness, all variants are trained with identical configurations in each experimental round.

Style-Content Cross Attention. As reported in Table 1, restricting cross-attention to the frame level yields only moderate performance, with a PSNR of 20.9 dB on skateboard and 19.7 dB on pizza. The hybrid variant, which combines frame-level and global attention, provides observable improvement on donut with 21.4 dB, but brings limited gains on the other categories and overall underperforms the global-only setup. By contrast, applying cross-attention at the global level consistently achieves the best results, reaching a PSNR of 21.7 dB on skateboard, 20.6 dB on pizza, and 22.1 dB on donut. Qualitative comparisons in Fig. 2 highlight these differences on the pizza example. Frame-only attention produces a poorly defined crust with noticeable artifacts, the hybrid variant recovers part of the structure but leaves the box edges blurred and unstable, while global attention delivers the most faithful result, with clear toppings and a well-preserved crust boundary.

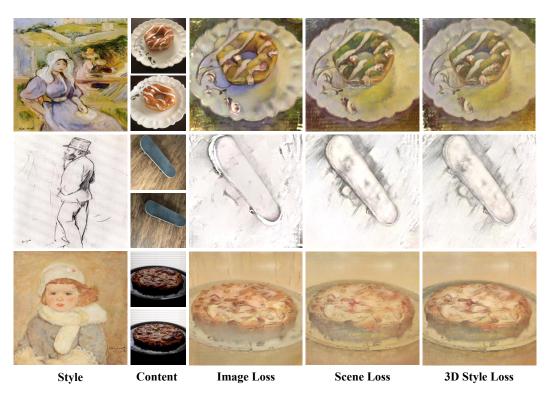


Figure 3: Comparison of style losses on unseen donut, skateboard, and pizza scenes from the CO3D dataset. Both scene and 3D style losses yield cleaner stylized textures compared to image-level matching, while the 3D loss further conveys a stronger sense of 3D geometry.

Table 2: Consistency comparison of different style loss designs on CO3D. The 3D loss achieves the strongest short-range and long-range consistency.

Style Loss	Short	range	Long-range		
2010 2000	LPIPS↓	$RMSE\downarrow$	$\overline{\text{LPIPS}}$	$\overline{\mathrm{RMSE}\!\!\downarrow}$	
Image loss (Eq. 4)	0.046	0.035	0.147	0.135	
Scene loss (Eq. 5)	0.048	0.037	0.147	0.133	
3D loss (Eq. 6)	0.046	0.033	0.140	0.128	

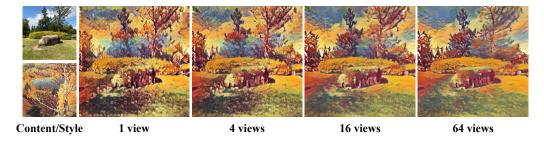


Figure 4: Effect of varying # views / batch on the Lighthouse scene from Tanks and Temples.

Table 3: Consistency comparison across methods. We report short-range (top) and long-range (bottom) LPIPS↓ and RMSE↓, on the four scenes from Tanks and Temples.

Method	Train		Truck		M60		Garden		
Westing	LPIPS↓	$\overline{\text{RMSE}}\downarrow$	LPIPS↓	$\overline{\text{RMSE}}\downarrow$	LPIPS↓	RMSE↓	LPIPS↓	RMSE↓	
Short-range consistency									
StyleRF 2023	0.075	0.065	0.060	0.062	0.039	0.042	_	_	
StyleGS 2024	0.033	0.038	0.031	0.034	0.038	0.037	0.069	0.061	
Stylos (ours)	0.030	0.026	0.028	0.021	0.035	0.024	0.047	0.044	
Long-range consistency									
StyleRF 2023	0.116	0.110	0.176	0.154	0.096	0.092	_	_	
StyleGS 2024	0.067	0.072	0.086	0.077	0.091	0.091	0.177	0.141	
Stylos (ours)	0.051	0.056	0.074	0.069	0.083	0.082	0.139	0.134	

Image vs. Scene vs. 3D Style Losses. As shown in Table 2, the image-level loss achieves short-range LPIPS of 0.046 and RMSE of 0.035, but long-range consistency rises to 0.147 and 0.135. The scene-level loss slightly improves long-range RMSE to 0.133 but degrades short-range performance to 0.048 and 0.037, indicating that simple 2D aggregation reduces some long-range inconsistencies but compromises short-range accuracy. In contrast, the 3D loss applies statistics after voxel-space fusion, matching the best short-range LPIPS of 0.046, further lowering short-range RMSE to 0.033, and improving long-range scores to 0.140 and 0.128. These results indicate that enforcing style consistency in a geometry-aligned space yields sharper boundaries and more coherent textures across views. As further illustrated in Fig. 3, image-level supervision often fails to faithfully transfer style. For example, the donut lacks stylization and the wooden surface is not well synthesized, whereas scene- and 3D-level losses successfully apply the target style. Among them, the 3D style loss produces sharper boundaries and a stronger sense of 3D geometry, as observed in the skateboard and pizza examples, leading to the most coherent multi-view stylization.

From a Single View to Dozens of Views. While Stylos can process hundreds of content inputs, we qualitatively observe a gradual decrease in visual quality once the number of views per batch (denoted as "# views / batch") exceeds 32. As shown in Fig. 4, very small batches (1 view) fail to synthesize parts of the scene such as the bench, while overly large batches (64 views) introduce edge artifacts on the tall building, potentially due to unstable Gaussian representations and the gap from training settings (no more than 24 views).

#### 4.2 Comparison with State-of-the-Art

For StyleRF (Liu et al., 2023) and StyleGS (Liu et al., 2024), we directly evaluate with their released per-scene weights. In contrast, our Stylos is trained once on the full DL3DV dataset (Ling et al., 2024) and tested in a zero-shot manner, without ever seeing the images of these scenes. **Quantitative Evaluation.** According to Table 3, Stylos achieves consistently lower LPIPS and RMSE than StyleGS across both short- and long-range view pairs. This

Table 4: ArtScore↑ comparison across methods on four Tanks and Temples scenes.

Method	Train	Truck	M60	Garden
StyleRF 2023	7.71	4.12	8.48	_
StyleGS 2024	0.78	5.76	8.63	9.38
Stylos (ours)	9.50	9.70	9.37	9.34

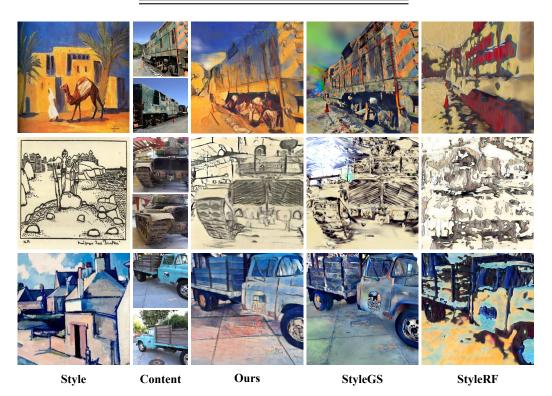


Figure 5: Visual comparison between Stylos and recent per-scene 3D stylization baselines.

demonstrates that Stylos provides stronger cross-view consistency. Furthermore, Table 4 reports that Stylos achieves the highest ArtScore across nearly all four Tanks and Temples scenes (Knapitsch et al. (2017a)), outperforming both StyleRF and StyleGS.

Qualitative Evaluation. Fig. 5 shows that Stylos successfully transfers the target style while preserving scene geometry and delivering consistent textures. In contrast, StyleGS introduces noticeable geometric distortions and texture discontinues, while StyleRF tends to over-impose style patterns, negatively impacting the object structure. Overall, Stylos achieves the best trade-off between structural fidelity and faithful stylization.

# 5 Conclusion

In this work, we propose Stylos, a single-forward-pass framework for multi-view 3D stylization. By keeping geometry predictions on the self-attentive backbone path and conditioning color via global cross-attention, together with a voxel-space 3D style loss that aggregates multi-view features, Stylos achieves geometry-aware and view-consistent stylization. Our ablations show that the global-only cross-attention between style and content better preserves geometric details than alternative coupling strategies. Extensive experiments across category-level and large-scale scene datasets demonstrate the generalization ability of our approach, achieving aesthetically pleasing stylization with strong cross-view consistency. In the future, we aim to scale Stylos to support higher-resolution inputs while improving efficiency, paving the way for practical 3D content creation.

# 6 Reproducibility Statement

We have taken several steps to ensure the reproducibility of our work. The main paper describes all model architectures in Sec. 3. We release detailed pseudo-code for all loss functions in Appendix A, and will release the full implementation upon acceptance. However, to support reproducibility, we provide detailed pseudo codes of the proposed style loss functions (image-level, scene-level, and 3D-level) in Appendix A.1. All datasets used are publicly available, as described in Sec. 4. These resources should allow researchers to reproduce and extend our findings. In addition, we provide extensive qualitative results in Appendix A.3, and analyze efficiency trends with respect to the number of views per batch in Appendix A.2. Efficiency experiments are averaged over 100 iterations per setting.

# References

- Jie An, Siyu Huang, Yibing Song, Dejing Dou, Wei Liu, and Jiebo Luo. Artflow: Unbiased image style transfer via reversible neural flows. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 862–871, 2021.
- Xu Cao, Weimin Wang, Katashi Nagao, and Ryosuke Nakamura. Psnet: 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*, pp. 3337–3345, 2020.
- Junyu Chen, Jie An, Hanjia Lyu, Christopher Kanan, and Jiebo Luo. Learning to evaluate the artness of ai-generated images. *IEEE Transactions on Multimedia*, 26:10731–10740, 2024.
- Pei-Ze Chiang, Meng-Shiun Tsai, Hung-Yu Tseng, Wei-Sheng Lai, and Wei-Chen Chiu. Stylizing 3d scene via implicit representation and hypernetwork. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 1475–1484, 2022.
- Bruno Galerne, Jianling Wang, Lara Raad, and Jean-Michel Morel. Sgsst: Scaling gaussian splatting style transfer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 26535–26544, 2025.
- Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2414–2423, 2016.
- Jie Guo, Mengtian Li, Zijing Zong, Yuntao Liu, Jingwu He, Yanwen Guo, and Ling-Qi Yan. Volumetric appearance stylization with stylizing kernel prediction network. ACM Trans. Graph., 40(4):162-1, 2021.
- Agrim Gupta, Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Characterizing and improving stability in neural style transfer. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4067–4076, 2017.
- Lukas Höllein, Justin Johnson, and Matthias Nießner. Stylemesh: Style transfer for indoor 3d scene reconstructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6198–6208, June 2022.
- Hsin-Ping Huang, Hung-Yu Tseng, Saurabh Saini, Maneesh Singh, and Ming-Hsuan Yang. Learning to stylize novel views. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 13869–13878, 2021.
- Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pp. 1501–1510, 2017.
- Yi-Hua Huang, Yue He, Yu-Jie Yuan, Yu-Kun Lai, and Lin Gao. 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*, pp. 18342–18352, 2022.

543

544

545

546

548

549 550

551

552

553

554

555

556

558

559

560

561

562

563

565

566

567

568

569

570

571 572

573

574

575

576

577

578 579

580

581

582

583

584

585 586

588

589

- Lihan Jiang, Yucheng Mao, Linning Xu, Tao Lu, Kerui Ren, Yichen Jin, Xudong Xu, Mulin Yu, Jiangmiao Pang, Feng Zhao, et al. Anysplat: Feed-forward 3d gaussian splatting from 542 unconstrained views. arXiv preprint arXiv:2505.23716, 2025.
  - Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3907– 3916, 2018.
    - Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph., 42(4):139-1, 2023.
    - Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. ACM Transactions on Graphics, 36(4), 2017a.
  - Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. ACM Transactions on Graphics (ToG), 36(4):1-13, 2017b.
    - Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. Demystifying neural style transfer. arXiv preprint arXiv:1701.01036, 2017.
  - Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. Dl3dv-10k: A large-scale scene dataset for deep learningbased 3d vision. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 22160–22169, 2024.
  - Kunhao Liu, Fangneng Zhan, Yiwen Chen, Jiahui Zhang, Yingchen Yu, Abdulmotaleb El Saddik, Shijian Lu, and Eric P Xing. Stylerf: Zero-shot 3d style transfer of neural radiance fields. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8338–8348, 2023.
  - Kunhao Liu, Fangneng Zhan, Muyu Xu, Christian Theobalt, Ling Shao, and Shijian Lu. Stylegaussian: Instant 3d style transfer with gaussian splatting. In SIGGRAPH Asia 2024 Technical Communications, pp. 1–4. ACM, 2024.
  - Yiqun Mei, Jiacong Xu, and Vishal Patel. Regs: Reference-based controllable scene stylization with gaussian splatting. Advances in Neural Information Processing Systems, 37: 4035–4061, 2024.
  - Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In ECCV, 2020.
  - Fangzhou Mu, Jian Wang, Yicheng Wu, and Yin Li. 3d photo stylization: Learning to generate stylized novel views from a single image. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 16273–16282, 2022.
  - Thu Nguyen-Phuoc, Feng Liu, and Lei Xiao. Snerf: stylized neural implicit representations for 3d scenes. ACM Transactions on Graphics, 41:1-11, 07 2022. doi: 10.1145/3528223. 3530107.
  - Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference* on machine learning, pp. 8748–8763. PmLR, 2021.
- Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-592 life 3d category reconstruction. In Proceedings of the IEEE/CVF international conference on computer vision, pp. 10901-10911, 2021.

- Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Linning Xu, Zhangkai Ni, and Bo Dai. Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. arXiv preprint arXiv:2403.17898, 2024.
  - Manuel Ruder, Alexey Dosovitskiy, and Thomas Brox. Artistic style transfer for videos. In German conference on pattern recognition, pp. 26–36. Springer, 2016.
  - Zachary Teed and Jia Deng. Deepv2d: Video to depth with differentiable structure from motion. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=HJeO7RNKPr.
  - Benjamin Ummenhofer, Huizhong Zhou, Jonas Uhrig, Nikolaus Mayer, Eddy Ilg, Alexey Dosovitskiy, and Thomas Brox. Demon: Depth and motion network for learning monocular stereo. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5038–5047, 2017.
  - Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 5294–5306, 2025a.
  - Peng Wang, Xiang Liu, and Peidong Liu. Styl3r: Instant 3d stylized reconstruction for arbitrary scenes and styles, 2025b. URL https://arxiv.org/abs/2505.21060.
  - Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20697–20709, 2024.
  - WikiArt. Wikiart: Visual art encyclopedia. https://www.wikiart.org/. Accessed: 2025-09.
  - Xin-Yi Yu, Jun-Xin Yu, Li-Bo Zhou, Yan Wei, and Lin-Lin Ou. Instantstylegaussian: Efficient art style transfer with 3d gaussian splatting. arXiv preprint arXiv:2408.04249, 2024.
  - Dingxi Zhang, Yu-Jie Yuan, Zhuoxun Chen, Fang-Lue Zhang, Zhenliang He, Shiguang Shan, and Lin Gao. Stylizedgs: Controllable stylization for 3d gaussian splatting. arXiv preprint arXiv:2404.05220, 2024.
  - Kai Zhang, Nick Kolkin, Sai Bi, Fujun Luan, Zexiang Xu, Eli Shechtman, and Noah Snavely. Arf: Artistic radiance fields. In *European Conference on Computer Vision*, pp. 717–733. Springer, 2022.
  - Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.
  - Huizhong Zhou, Benjamin Ummenhofer, and Thomas Brox. Deeptam: Deep tracking and mapping. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 822–838, 2018.

# A Appendix

648

649 650

651

652

653

654 655

656

675 676

677

678

679

680

681

682

683

684

685

686

687

688

689

690 691

692

693

694 695 696

697 698

699

700

701

#### A.1 Style Loss Pseudo Codes

We provide pseudo codes to the three variants for style losses introduced in the main paper, image-level, scene-level, and 3d style losses. This is for reproducibility purpose, and please follow the implementation in our codes which will be released in the near future.

## Algorithm 1 Image-level AdaIN Style Loss

```
Require: Content images I_c \in \mathbb{R}^{B \times S \times C \times H \times W}, Rendered images I_r, Style images I_s
657
             Ensure: Image-level style loss \mathcal{L}_{img}
658
               1: for i = 1 to B do
659
                       Extract style features F_s \leftarrow \text{VGG}(I_s[i])
660
               3:
                       for j = 1 to S do
661
                           Extract content features F_c \leftarrow \text{VGG}(I_c[i,j])
               4:
662
                           Extract rendered features F_r \leftarrow \text{VGG}(I_r[i,j])
               5:
663
               6:
                           Content loss:
664
                                                               \mathcal{L}_c \leftarrow \|F_r^{(4)} - F_c^{(4)}\|_2^2 + \|F_r^{(5)} - F_c^{(5)}\|_2^2
665
                           Style loss initialization: \mathcal{L}_s \leftarrow 0
               7:
666
                           for k = 1 to 5 do
               8:
667
                              (\mu_r, \sigma_r) \leftarrow \text{MeanStd}(F_r^{(k)})
               9:
668
                              (\mu_s, \sigma_s) \leftarrow \text{MeanStd}(F_s^{(k)})
\mathcal{L}_s \leftarrow \mathcal{L}_s + \|\mu_r - \mu_s\|_2^2 + \|\sigma_r - \sigma_s\|_2^2
             10:
669
             11:
670
             12:
671
                       end for
             13:
672
             14: end for
673
             15: return \mathcal{L}_{img} = \lambda_c \mathcal{L}_c + \lambda_s \mathcal{L}_s + \lambda_{mse} ||I_r - I_c||^2
674
```

# Algorithm 2 Scene-level AdaIN Style Loss

```
Require: Content images I_c \in \mathbb{R}^{B \times S \times C \times H \times W}, Rendered images I_r, Style images I_s
Ensure: Scene-level style loss \mathcal{L}_{\text{scene}}
 1: for i = 1 to B do
          Extract style features F_s \leftarrow \text{VGG}(I_s[i])
 2:
          Extract batched content features F_c \leftarrow \text{VGG}(I_c[i,:])
 3:
          Extract batched rendered features F_r \leftarrow \text{VGG}(\tilde{I}_r[i,:])
 4:
 5:
          Content loss:
                                                  \mathcal{L}_c \leftarrow \|F_r^{(4)} - F_c^{(4)}\|_2^2 + \|F_r^{(5)} - F_c^{(5)}\|_2^2
          Style loss initialization: \mathcal{L}_s \leftarrow 0
 6:
          for k = 1 to 5 do
Reshape F_r^{(k)} into (B, S, C, H, W)
 7:
 8:
              (\mu_r, \sigma_r) \leftarrow \text{MeanStd}(F_r^{(k)})
 9:
              (\mu_s, \sigma_s) \leftarrow \text{MeanStd}(F_s^{(k)})
\mathcal{L}_s \leftarrow \mathcal{L}_s + \|\mu_r - \mu_s\|_2^2 + \|\sigma_r - \sigma_s\|_2^2
10:
          end for
12:
13: end for
14: return \mathcal{L}_{\text{scene}} = \lambda_c \mathcal{L}_c + \lambda_s \mathcal{L}_s + \lambda_{\text{mse}} ||I_r - I_c||^2
```

#### A.2 Efficiency Scaling

We further analyze the efficiency of Stylos by varying the number of views processed per batch. The results in Fig. 6 indicate a clear scaling pattern. For very small batches, such as 1–2 views, the average runtime per batch is low, around 0.3 s, but the measurements fluctuate heavily due to overheads and GPU scheduling. This makes single-view or extremely small-batch inference unstable, even though the raw latency is minimal.

# Algorithm 3 3D Voxel-based AdaIN Style Loss

702

703

704

705

706

707

708

709

710

711

712

713

714

715

716

717 718

719 720 721

722

723

724

725

726

727

728

729

730 731

733

734

735 736

738

739

740

741

742

743 744

745 746

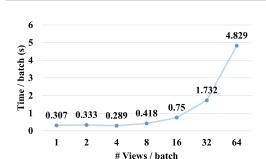
747

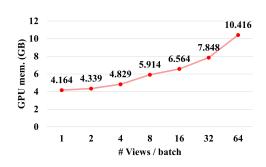
748

749

750

```
Require: Rendered features F_r, Style features F_s, Content features F_c, 3D points P \in
      \mathbb{R}^{B \times S \times H \times W \times 3}, Confidence C, Valid mask M
Ensure: 3D style loss \mathcal{L}_{3d}
 1: for i = 1 to B do
         for k = 1 to 5 do
 2:
             Resize (C, M, P) to match the resolution of F_r^{(k)}
 3:
             for s = 1 to S do
 4:
                 \mathcal{V} \leftarrow \text{Voxelize}(F_r^{(k)}[s], P[s], C[s], M[s])
 5:
                 (\mu_v, \sigma_v) \leftarrow \text{MeanStd}(\mathcal{V})
 6:
                 (\mu_s, \sigma_s) \leftarrow \text{MeanStd}(F_s^{(k)})
 7:
                 \mathcal{L}_s \leftarrow \mathcal{L}_s + \|\mu_v - \mu_s\|_2^2 + \|\sigma_v - \sigma_s\|_2^2
 8:
             end for
 9:
10:
         end for
         Content loss:
11:
                                              \mathcal{L}_c \leftarrow \|F_r^{(4)} - F_c^{(4)}\|_2^2 + \|F_r^{(5)} - F_c^{(5)}\|_2^2
12: end for
13: return \mathcal{L}_{3d} = \lambda_c \mathcal{L}_c + \lambda_s \mathcal{L}_s + \lambda_{\text{mse}} ||I_r - I_c||^2
```





- (a) Time per batch vs. # views per batch
- (b) GPU memory usage vs. # views per batch

Figure 6: Scaling of inference time and GPU memory with the number of views per batch. Time is averaged over 100 iterations, and memory is reported as the peak allocation, without rendering cost.

As the batch size increases, runtime grows roughly linearly, reaching 0.75 s at 16 views and 4.83 s at 64 views. Importantly, once the batch size exceeds 4–8 views, the runtime becomes much more stable, showing only minor variation across repeated runs.

Memory usage shows a similarly consistent trend: GPU memory rises steadily from 4.16 GB at 1 view to 10.42 GB at 64 views. The growth is close to linear with respect to batch size, which makes resource requirements easy to estimate when scaling to larger workloads.

#### A.3 Additional Visual Results

Figure 7-16 presents additional qualitative results of Stylos on diverse style—content pairs. Across different styles and object categories, Stylos produces visually pleasing stylizations that respect the input content geometry while transferring the target artistic appearance. These examples further demonstrate the versatility and robustness of our approach beyond the main results shown in the paper.



Figure 7: Additional visual results of Stylos on diverse style—content pairs on Tanks and Temples Dataset (Knapitsch et al. (2017a)). Stylos consistently produces visually pleasing stylizations that preserve content geometry while transferring the target style.

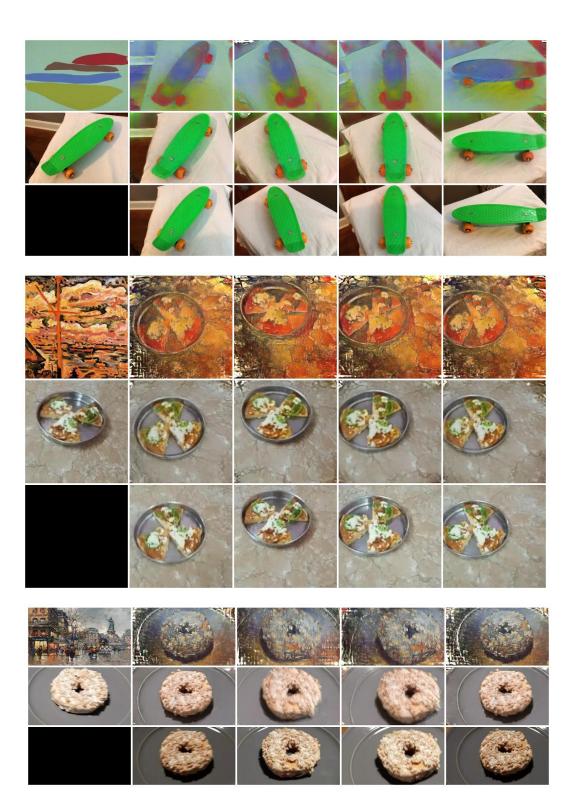


Figure 8: Additional visual results of Stylos on the CO3D dataset. The top left is the style image while the rest of the top row are the rendered images. The mid row are the rendered images conditioned on the mid left image. The bottom row is the content inputs.

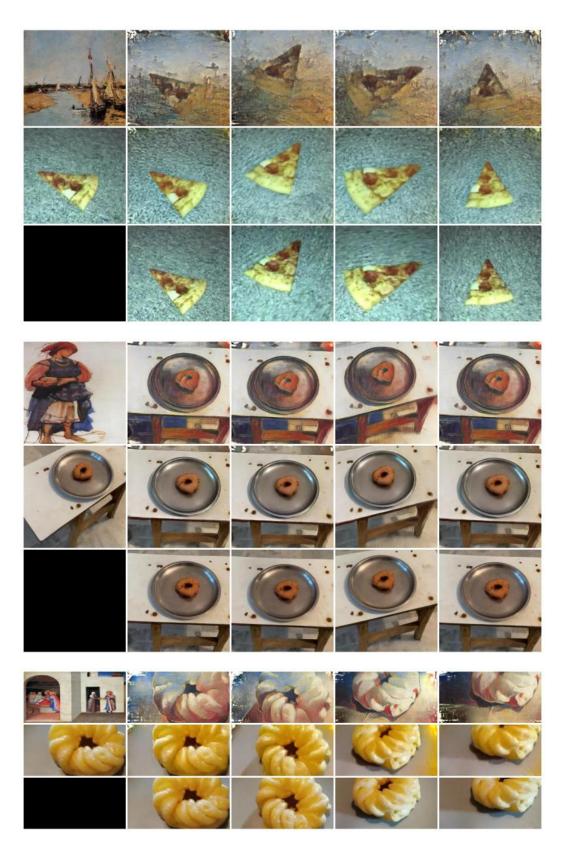


Figure 9: Additional visual results of Stylos on diverse style–content pairs.

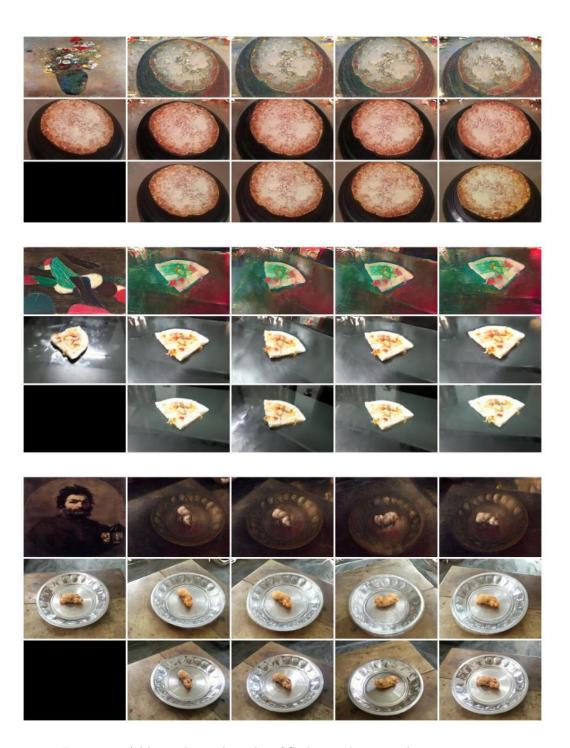


Figure 10: Additional visual results of Stylos on diverse style–content pairs.

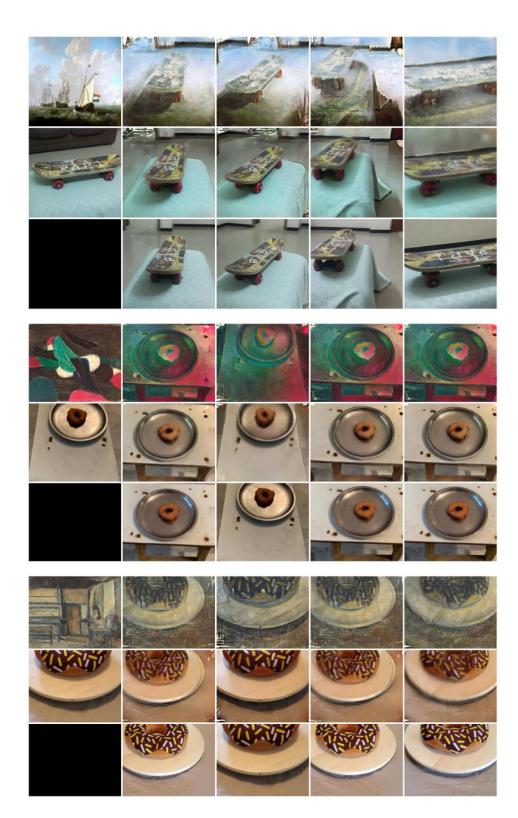


Figure 11: Additional visual results of Stylos on diverse style—content pairs.



Figure 12: Additional visual results of Stylos on diverse style–content pairs.

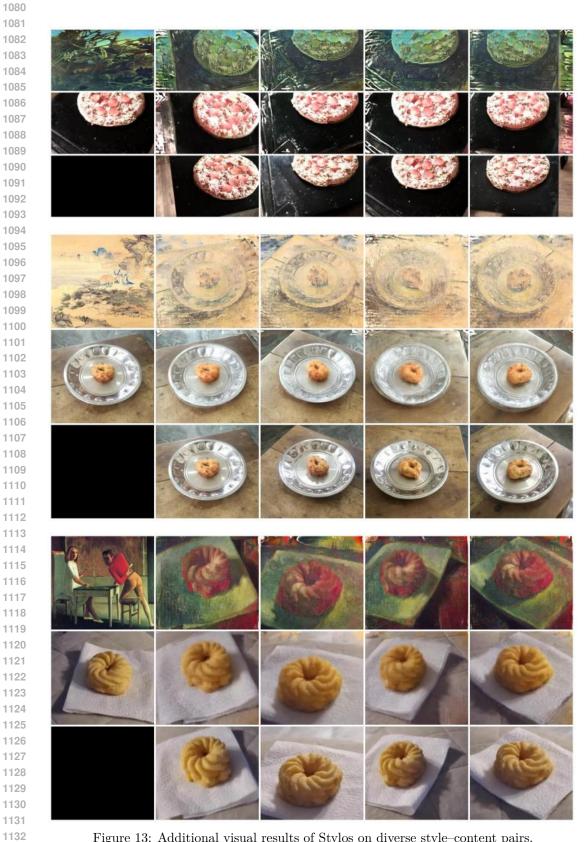


Figure 13: Additional visual results of Stylos on diverse style—content pairs.

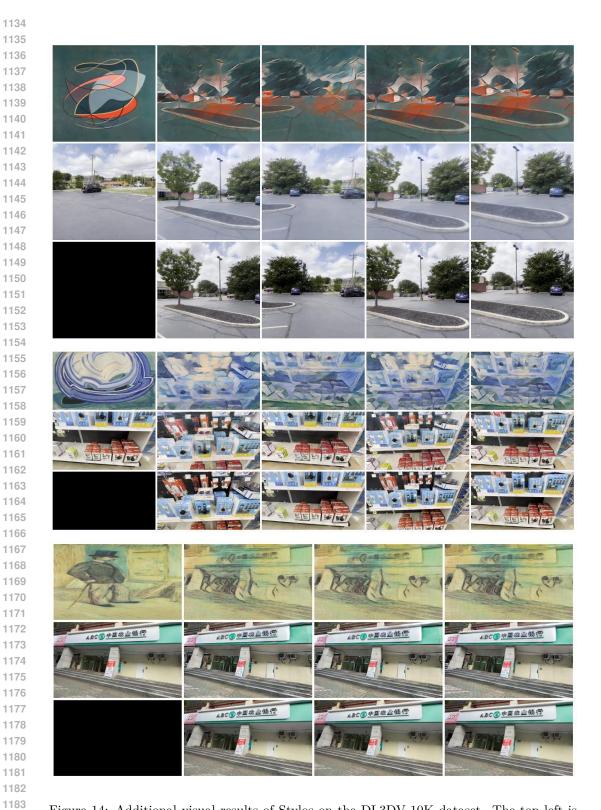
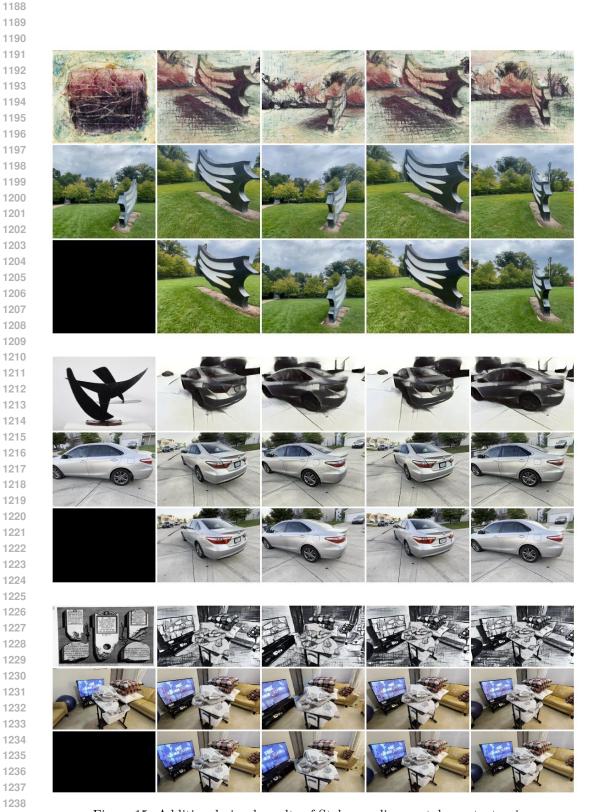


Figure 14: Additional visual results of Stylos on the DL3DV-10K dataset. The top left is the style image while the rest of the top row are the rendered images. The mid row are the rendered images conditioned on the mid left imamge. The bottom row is the content inputs.



 $Figure \ 15: \ Additional \ visual \ results \ of \ Stylos \ on \ diverse \ style-content \ pairs.$ 

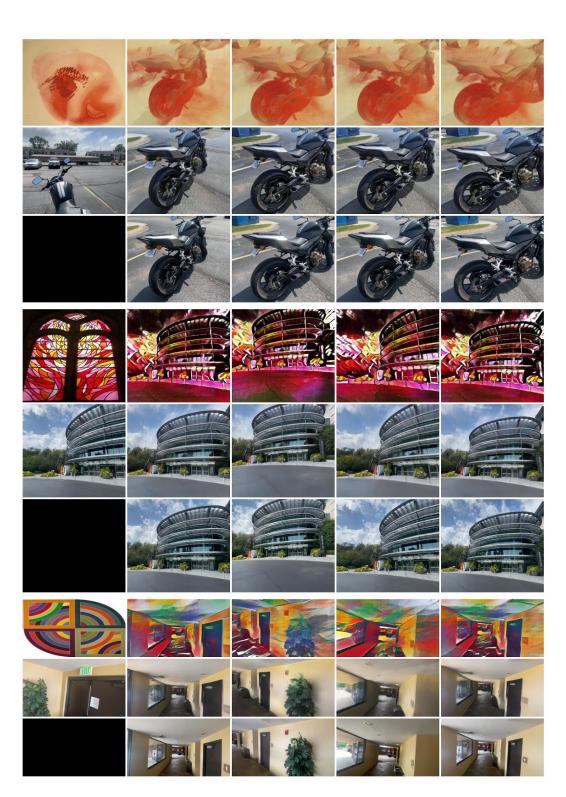


Figure 16: Additional visual results of Stylos on diverse style–content pairs.