

# Ref-NPR: Reference-Based Non-Photorealistic Radiance Fields for Controllable Scene Stylization

Yuechen Zhang<sup>1,2</sup>   Zexin He<sup>1</sup>   Jinbo Xing<sup>1</sup>   Xufeng Yao<sup>1</sup>   Jiaya Jia<sup>1,2</sup>

<sup>1</sup>The Chinese University of Hong Kong   <sup>2</sup>SmartMore

{yczhang21, zxhe22, jbxing, xfyao, leojia}@cse.cuhk.edu.hk




Figure 1. Given a pair of one reference view from a radiance field and its stylization, Ref-NPR propagates style more faithfully to novel views with semantic correspondence compared with state-of-the-art scene stylization method ARF [45].

## Abstract

Current 3D scene stylization methods transfer textures and colors as styles using arbitrary style references, lacking meaningful semantic correspondences. We introduce Reference-Based Non-Photorealistic Radiance Fields (Ref-NPR) to address this limitation. This controllable method stylizes a 3D scene using radiance fields with a single stylized 2D view as a reference. We propose a ray registration process based on the stylized reference view to obtain pseudo-ray supervision in novel views. Then we exploit semantic correspondences in content images to fill occluded regions with perceptually similar styles, resulting in non-photorealistic and continuous novel view sequences. Our experimental results demonstrate that Ref-NPR outperforms existing scene and video stylization methods regarding visual quality and semantic correspondence. The code and data are publicly available on the project page at <https://ref-npr.github.io>.

## 1. Introduction

In the past decade, there has been a rising demand for stylizing and editing 3D scenes and objects in various fields, including augmented reality, game scene design, and digital

artwork. Traditionally, professionals achieve these tasks by creating 2D reference images and converting them into stylized 3D textures. However, establishing direct cross-modal correspondence is challenging and often requires significant time and effort to obtain stylized texture results similar to the 2D reference schematics.

A critical challenge in the 3D stylization problem is to ensure that stylized results are perceptually similar to the given style reference. Benefiting from radiance fields [2, 10, 28, 29, 37, 46], recent novel-view stylization methods [5, 8, 15, 16, 30, 45] greatly facilitated style transfer from an arbitrary 2D style reference to 3D implicit representations. However, these methods do not provide explicit control over the generated results, making it challenging to specify the regions where certain styles should be applied and ensure the visual quality of the results. On the other hand, reference-based video stylization methods allow for the controllable generation of stylized novel views with better semantic correspondence between content and style reference, as demonstrated in works like [17, 38]. However, these methods may diverge from the desired style when stylizing a frame sequence with unseen content, even with the assistance of stylized keyframes.

To address the aforementioned limitations, we propose

a new paradigm for stylizing 3D scenes using a single stylized reference view. Our approach, called Reference-Based Non-Photorealistic Radiance Fields (Ref-NPR), is a controllable scene stylization method that takes advantage of volume rendering to maintain cross-view consistency and establish semantic correspondence in transferring style across the entire scene.

Ref-NPR utilizes stylized views from radiance fields as references instead of arbitrary style images to achieve both flexible controllability and multi-view consistency. A reference-based ray registration process is designed to project the 2D style reference into 3D space by utilizing the depth rendering of the radiance field. This process provides pseudo-ray supervision to maintain geometric and perceptual consistency between stylized novel views and the stylized reference view. To obtain semantic style correspondence in occluded regions, Ref-NPR performs template-based feature matching, which uses high-level semantic features as implicit style supervision. The correspondence in the content domain is utilized to select style features in the given style reference, which are then used to transfer style globally, especially in occluded regions. By doing so, Ref-NPR generates the entire stylized 3D scene from a single stylized reference view.

Ref-NPR produces visually appealing stylized views that maintain both geometric and semantic consistency with the given style reference, as presented in Fig. 1. The generated stylized views are perceptually consistent with the reference while also exhibiting high visual quality across various datasets. We have demonstrated that Ref-NPR, when using the same stylized view as reference, outperforms state-of-the-art scene stylization methods [30, 45] both qualitatively and quantitatively.

In summary, our paper makes three contributions. Firstly, we introduce a new paradigm for stylizing 3D scenes that allows for greater controllability through the use of a stylized reference view. Secondly, we propose a novel approach called Ref-NPR, consisting of a reference-based ray registration process and a template-based feature matching scheme to achieve geometrically and perceptually consistent stylizations. Finally, our experiments demonstrate that Ref-NPR outperforms state-of-the-art scene stylization methods such as ARF and SNeRF both qualitatively and quantitatively. More comprehensive results and a demo video can be found in the supplementary material and on our project page.

## 2. Related Works

### 2.1. Stylization in 2D

**Arbitrary style transfer** is a well-studied problem in Non-Photorealistic Rendering (NPR) [12, 23]. Gatys et al. [11] first introduced the idea of representing image style as high-

level features extracted from pre-trained deep neural networks. Since then, various parametric image style transfer methods [4, 18, 20, 22, 24, 25] have been developed to generate high-quality stylized images efficiently. Video stylization methods that use arbitrary style input [3, 7, 32, 41, 42] mainly focus on maintaining temporal coherence to achieve continuous stylized frames. Nonetheless, these stylization methods lack interpretability and controllability, even when using stylized keyframes as reference.

**Example-based stylization** methods use multiple style references to stylize images while ensuring semantic correspondences between them. Methods such as [14, 26] use multi-level semantic feature matching to establish dense correspondences between the content image and the style reference with similar semantics. To facilitate this feature matching, some methods use explicit alignments such as warping or content-aligned stylizing [17, 33, 34, 38]. These methods offer controllability by allowing editing of the style reference, but they are not suitable for stylizing novel views as they cannot correctly stylize unseen regions using 2D reference-based methods.

### 2.2. Stylization in 3D

**Reference-based 3D stylization** methods have achieved promising results without the use of radiance fields, as seen in previous works such as Texture Map [1], Texture Field [31], StyLit [9, 36], and StyleProp [13]. However, these methods have their limitations. For instance, Texture Field is trained on a restricted ShapeNet dataset, limiting the stylization of 3D objects within the trained categories. StyLit, on the other hand, treats each view of a 3D object as multi-channel 2D guidance and applies texture mapping, resulting in flickering artifacts due to a lack of geometric prior. Although StyleProp utilizes multi-view correspondence maps to obtain stylized novel views, it can only work on viewing directions around the reference view.

**Stylizing radiance fields** has recently emerged as a popular topic in computer vision research, driven by the growing popularity of radiance fields [5, 8, 28]. Huang et al. [15] pioneered the application of scene stylization on implicit 3D representations, followed by StylizedNeRF [16], CLIP-NeRF [39], NeRF-Art [40], INS [8], SNeRF [30], and ARF [45], which focus on various degrees, like text-driven stylization in CLIP feature space, stylization using unified representation, memory efficiency, and stylization quality. However, these methods lack explicit correspondence modeling, leading to uncontrollable stylized novel views that differ perceptually from the style reference. To address this issue, we propose Ref-NPR. This reference-guided controllable scene stylization method generates stylized novel views with geometric and semantic consistency with a given stylized view reference.

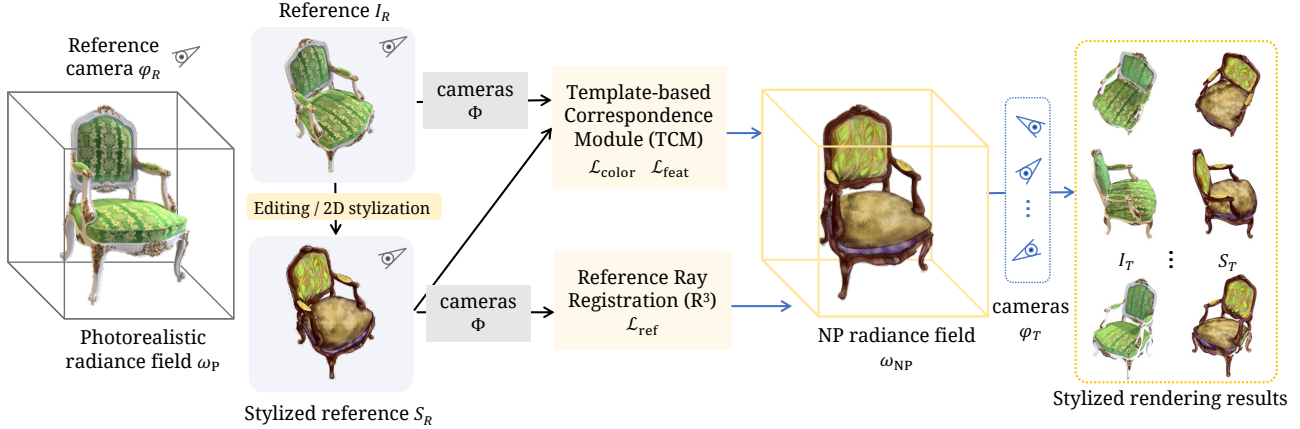


Figure 2. The workflow of **Ref-NPR**. Given a pre-trained photorealistic radiance field  $\omega_P$ , we can provide a stylized reference view  $S_R$  to obtain reference-based supervisions  $\mathcal{L}_{\text{ref}}$ ,  $\mathcal{L}_{\text{feat}}$ , and  $\mathcal{L}_{\text{color}}$ . Those loss constraints are used to optimize a non-photorealistic (NP) radiance field  $\omega_{\text{NP}}$ . During inference, with such NP radiance field, stylized results  $S_T$  could be rendered from an arbitrary set of camera poses  $\varphi_T$ , corresponding to the original views  $I_T$  rendered by  $\omega_P$ .

### 3. Ref-NPR

Fig. 2 outlines the process of Ref-NPR using a single reference view as an example. Ref-NPR aims to stylize a pre-trained photorealistic radiance field with the help of one or a few pairs of reference views and their corresponding stylizations. The first step involves rendering a reference view  $I_R$  from a specific reference camera  $\varphi_R$  using a pre-trained photorealistic radiance field  $\omega_P$ . Next, the reference view  $I_R$  is stylized by either manually editing or by utilizing structure-preserving 2D-stylization algorithms such as Gatys [11] or AdaIN [20] to obtain a stylized reference view  $S_R$  based on  $I_R$ .

To enable explicit supervision from the stylized reference view  $S_R$  to novel views, we propose a Reference Ray Registration ( $R^3$ ) process introduced in section 3.2. This process produces a set of reference-dependent pseudo-rays  $\Gamma_R$ , which are correlated rays produced between the reference camera  $\varphi_R$  and the set of training camera poses  $\Phi$ . Additionally, we use a Template-based Correspondence Module (TCM) in section 3.3 to obtain implicit style supervision in occluded regions of the reference view. Together,  $R^3$  and TCM provide explicit and implicit supervision to optimize a new non-photorealistic (NP) radiance field  $\omega_{\text{NP}}$  in section 3.4, allowing us to access stylized rendering results of arbitrary target views.

#### 3.1. Preliminary: Radiance Field Rendering

Volume rendering [19, 28] uses camera rays to sample a 3D radiance field  $\omega$  to render images. A camera ray  $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$  is defined by an origin  $\mathbf{o} \in \mathbb{R}^3$  and a direction  $\mathbf{d} \in \mathbb{R}^3$  pointing towards the center of a pixel in the image. The radiance field is sampled at  $N$  points along the ray, denoted by  $\{\mathbf{r}(t_i) | i = 1 \dots N, t_i < t_{i+1}\}$ . At each sample point, the

radiance field returns a density value  $\sigma_i = \sigma(\mathbf{r}(t_i))$  and a view-dependent color  $\mathbf{c}_i = c(\mathbf{r}(t_i), \mathbf{d})$ . The accumulated pixel color  $\hat{C}(\mathbf{r})$  of the ray  $\mathbf{r}$  is estimated in the discrete context, as formulated in [28]:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i(t_{i+1} - t_i))) \mathbf{c}_i, \quad (1)$$

$$\text{where } T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j(t_{j+1} - t_j)\right), \quad (2)$$

which estimates the accumulated transmittance along the ray from  $t_1$  to  $t_i$ . During training, the radiance field is optimized by directly minimizing the discrepancy between the predicted pixel color  $\hat{C}(\mathbf{r})$  and the ground truth pixel color  $C(\mathbf{r})$  for each ray, which is denoted by

$$\mathcal{L}_\omega = \sum_{\mathbf{r}} \|\hat{C}(\mathbf{r}) - C(\mathbf{r})\|_2^2. \quad (3)$$

Once the radiance field is optimized, depth estimation can be performed by mapping the ray  $\mathbf{r}$  to an exact 3D position in the scene. This is achieved by setting a threshold  $\sigma_z$  on the accumulated density  $T_i$  calculated in Equation 2. The first sample that exceeds this threshold is interpreted as the intersection point between the ray and the scene. The length of the ray  $\mathbf{r}$  is then defined as the distance between the camera origin  $\mathbf{o}$  and the intersection point, denoted by

$$l(\mathbf{r}) = \min\{t_i | \sum_{j=1}^{i-1} \sigma_j(t_{j+1} - t_j) \geq \sigma_z\}. \quad (4)$$

Then we write the intersection point corresponding to ray  $\mathbf{r}$  as  $\mathbf{x}(\mathbf{r}) = \mathbf{o} + l(\mathbf{r})\mathbf{d}$ , which is our desired mapping.

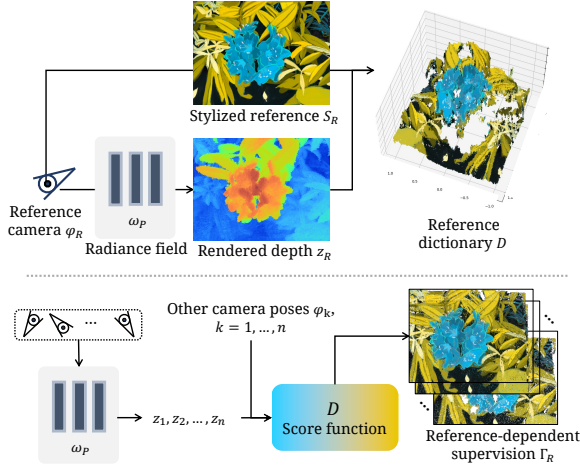


Figure 3. **Reference Ray Registration.** To create a reference dictionary  $D$  for a stylized reference view  $S_R$  with camera pose  $\varphi_R$ , we estimate pseudo-depths from  $\omega_P$ . Then, we apply a ray registration process to train camera poses and acquire a collection  $\Gamma_R$  consisting of pseudo-rays and their assigned colors.

### 3.2. Reference Ray Registration

Ref-NPR differs from existing scene stylization methods [16, 30, 45] by incorporating an additional objective of establishing semantic consistency between the stylized reference view and novel views. This objective is also pursued by some scene modeling methods [6, 43, 44], which use 3D information such as depth to enhance rendering quality. Notably, precise 3D positions and camera parameters enable pixel-wise correspondence between views with a closed-form solution. This is utilized in many methods, such as homographic warping. In Ref-NPR, the radiance field  $\omega_P$  is used to estimate pseudo-depth information based on Eq. (4).

$R^3$ , or Reference Ray Registration, utilizes pseudo-depth information to obtain reference-related novel view supervision as shown in Fig. 3. Using the depth rendering property, pixels in the stylized reference view  $S_R$  can be mapped to 3D space by estimating their respective ray lengths. A reference dictionary  $D$  is constructed, where the element at index  $(x, y, z)$  contains all the rays that terminate in the voxel with the same index. This is achieved using a quantization operator  $Q(\cdot)$ , which maps 3D positions to their corresponding voxels. The reference dictionary is formally defined as

$$D_{(x,y,z)} = \{\mathbf{r}_i \in \varphi_R \mid Q(\mathbf{x}(\mathbf{r}_i)) = (x, y, z)\}, \quad (5)$$

in which  $\varphi_R$  is the reference camera and  $\mathbf{x}(\mathbf{r}_i)$  is the intersection point of  $\mathbf{r}_i$  estimated from the pseudo-depth  $l(\mathbf{r}_i)$  according to Eq. (4). By splitting 3D space into discrete voxels, each entry in  $D$  may be mapped with multiple rays or none at all, depending on the estimated pseudo-depths. For each ray  $\mathbf{r}_i \in \varphi_R$ , we use  $\hat{C}_R(\cdot)$  to denote the stylized color according to the stylized reference view  $S_R$ .

Now that  $S_R$  is propagated to 3D space with such reference dictionary, we thus register each ray  $\mathbf{r}_j \in \Phi$  from the training views as a pseudo-ray  $\hat{\mathbf{r}}_j \in \varphi_R$  in a best-matching manner, through minimizing the Euclidean distance of two corresponding intersection points in 3D space. Further, to avoid the over-matching problem coming from the gap of ray directions, we deploy a constraint that the angle spanned between directions of matched rays should not exceed a certain threshold  $\theta$ , i.e.,  $\angle(\mathbf{d}_{\mathbf{r}_i}, \mathbf{d}_{\mathbf{r}_j}) < \theta$ . We formulate the ray registration process as

$$\hat{\mathbf{r}}_j = \arg \min_{\substack{\mathbf{r}_i \in D_{(x,y,z)}, \\ \angle(\mathbf{d}_{\mathbf{r}_i}, \mathbf{d}_{\mathbf{r}_j}) < \theta}} \|\mathbf{x}(\mathbf{r}_i) - \mathbf{x}(\mathbf{r}_j)\|_2, \quad (6)$$

$$\text{where } Q(\mathbf{x}(\mathbf{r}_j)) = (x, y, z), \mathbf{r}_j \in \Phi. \quad (7)$$

Ray registration finds a ray in the dictionary whose intersection point drops into the same voxel  $D_{(x,y,z)}$  as the ray of interest. Eventually, we construct reference-dependent pseudo-ray supervision as  $\Gamma_R$  by collecting each validated, registered ray  $\mathbf{r}_j$  and assign its color in accordance to the corresponding reference ray  $\hat{\mathbf{r}}_j$ . Such a collection of pseudo-rays and their stylized colors with the aforementioned  $\hat{C}_R(\cdot)$  is defined as

$$\Gamma_R = \{(\mathbf{r}_j, \hat{C}_R(\hat{\mathbf{r}}_j)) \mid \mathbf{r}_j \in \Phi \cup \varphi_R, \hat{\mathbf{r}}_j \neq \emptyset\}, \quad (8)$$

where  $\Phi$  is the collection of all accessible camera poses.

### 3.3. Template-Based Semantic Correspondence

While  $R^3$  is effective in generating pseudo-ray supervision around the given reference camera pose  $\varphi_R$ , it struggles to register reference rays to occluded regions under such a camera, especially for 360° scenes in [21, 28]. To overcome this limitation, we leverage the common assumption in scene stylization that *the semantic correspondence in the entire scene should be consistent before and after stylization*. Based on this, we use the stylized reference view  $S_R$  based on source content  $I_R$  to establish a content-style mapping as a template. We then broadcast such reference style to novel views with semantically similar content. To achieve this, we introduce a Template-based Correspondence Module (TCM) that utilizes this content-style correspondence to construct a semantic correlation within the content domain.

As illustrated in Fig. 4, for each content domain view  $I$  rendered from  $\omega_P$  under some certain camera pose  $\varphi \in \Phi$ , we obtain its high-level semantic feature map  $F_I$  from a pre-trained semantic feature extractor (e.g., VGG16 [35]). Similarly, we denote the extracted feature maps for content reference  $I_R$  and style reference  $S_R$  by  $F_{I_R}$  and  $F_{S_R}$ , respectively, and use a superscript to index each element in the feature maps. We therefore construct our desired guidance feature  $F_G$  for further supervision, described by a search-and-replace process on 2D position  $(i, j)$  as

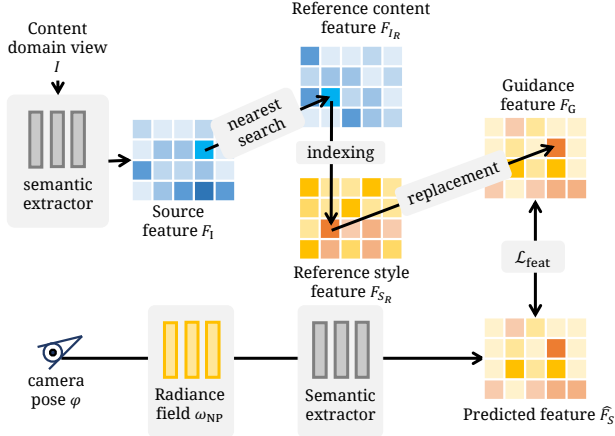


Figure 4. **Template-based Correspondence Module (TCM)**. To get the implicit style supervision of one view, the content domain view  $I$  is passed to a semantic extractor to obtain its semantic feature  $F_I$ . Then a search-replacement process is conducted to replace the reference feature according to the correspondence in the content domain. The resulting guidance feature  $F_G$  serves as implicit supervision to optimize the radiance field  $\omega_{NP}$ .

$$F_G^{(i,j)} = F_{S_R}^{(i^*,j^*)}, \quad (9)$$

$$\text{where } (i^*, j^*) = \arg \min_{i', j'} \text{dist}(F_I^{(i,j)}, F_{I_R}^{(i',j')}). \quad (10)$$

Here we use  $\text{dist}(\mathbf{a}, \mathbf{b})$  to denote the distance between two feature vectors  $\mathbf{a}$  and  $\mathbf{b}$ , which is proved to be effective [22, 45] by taking the form of cosine distance when evaluating semantic features.

Finally, for a stylized view  $S$  rendered from the NP radiance field  $\omega_{NP}$  under the camera pose  $\varphi$ , we conduct implicit feature-level supervision by forcing its semantic feature  $\hat{F}_S$  to imitate the aforesaid guidance feature  $F_G$ .

### 3.4. Ref-NPR Optimization

With the previously discussed collection  $\Gamma_R$  and guidance feature  $F_G$  as supervision, we optimize Ref-NPR and obtain the stylized scene representation  $\omega_{NP}$ .

In each training iteration, we sample a subset from  $\Gamma_R$  and denote the set of reference-dependent rays as  $N_s$ . For each sampled reference ray  $\mathbf{r}_k \in N_s$ ,  $\hat{C}_R(\hat{\mathbf{r}}_k)$  is the assigned color of the corresponding pseudo-ray  $\hat{\mathbf{r}}_k$ , as defined in Eq. (6), and we further denote  $\hat{C}_{NP}(\mathbf{r}_k)$  to be the stylized color rendered from  $\omega_{NP}$ . This explicit supervision is formulated as the reference loss

$$\mathcal{L}_{\text{ref}} = \frac{1}{|N_s|} \sum_{\mathbf{r}_k \in N_s} \|\hat{C}_{NP}(\mathbf{r}_k) - \hat{C}_R(\hat{\mathbf{r}}_k)\|_2^2. \quad (11)$$

As for implicit supervision, the discrepancy between  $F_G$  and  $\hat{F}_S$  should be minimized, as discussed in Sec. 3.3. To maintain the original content structure during such implicit

stylization, we also minimize the mean squared distance between content feature  $F_I$  and stylized feature  $\hat{F}_S$ , according to [11]. This implicit supervision is formulated as the feature loss

$$\mathcal{L}_{\text{feat}} = \frac{1}{N} \sum_{i,j} (\text{dist}(F_G, \hat{F}_S) + \lambda' \|F_I - \hat{F}_S\|_2^2), \quad (12)$$

where  $\lambda'$  is a balancing factor.

However, as discussed in [22, 45], optimizing the cosine distance between feature vectors cannot effectively eliminate color mismatches. To address this issue, we transfer the average color in a patch by a coarse color-matching loss

$$\mathcal{L}_{\text{color}} = \frac{1}{N} \sum_{i,j} \|\bar{C}_{NP}^{(i,j)} - \bar{C}_R^{(i^*,j^*)}\|_2^2, \quad (13)$$

in which  $\bar{C}_{NP}^{(i,j)}$  is the  $\omega_{NP}$ -rendered average color of the patch at feature-level index  $(i, j)$ , and  $\bar{C}_R^{(i^*,j^*)}$  is the average color of reference patch matched by minimizing feature distance, as described in Eq. (10). Since semantic features are extracted at the image level, considering the memory limitation caused by back-propagation, we follow the gradient cache strategy in [45] and optimize  $\omega_{NP}$  patch-wisely.

Ultimately, the overall objective is  $\mathcal{L}_{NP} = \lambda_f \mathcal{L}_{\text{feat}} + \lambda_r \mathcal{L}_{\text{ref}} + \lambda_c \mathcal{L}_{\text{color}}$ , where  $\lambda_{(\cdot)}$  are the balancing factors. Once  $\omega_{NP}$  is optimized, we may consider it to be a normal radiance field and render stylized novel views with arbitrary camera poses.

## 4. Experiments

### 4.1. Implementation Details

Ref-NPR is based on the ARF codebase [45] and uses Plenoxels [10] as the radiance field for scene representation. We follow Plenoxels's training scheme to obtain the photorealistic radiance field  $\omega_P$ . As we do not expect a view-dependent color change in stylized scenes, following [16, 45], we discard view-dependent rendering and apply a view-independent fitting on training views for two epochs before optimizing  $\omega_{NP}$ . Then, content domain views  $I$  are rendered from  $\omega_P$  after this view-independent training. In addition, to keep the same geometrical structure of the scene, we do not optimize the density function  $\sigma(\mathbf{r}(t_i))$  in  $\omega_{NP}$  [10, 45].

In  $\mathbb{R}^3$ , we set reference dictionary  $D$  as a cube containing  $256^3$  voxels. To parallelize the registration, we store at most 8 rays at each entry  $D_{(x,y,z)}$ . The angle constraint of directions in Eq. (6) is empirically set to  $\cos(\theta) > 0.6$  for ray registration. For each training step, we set the number of pseudo-ray samples to be  $|N_s| = 10^6$ , with half of them from  $\varphi_R$  and the other half from rays registered in correlated views. In TCM, we use VGG16 [35] as the semantic feature extractor. We concatenate features that

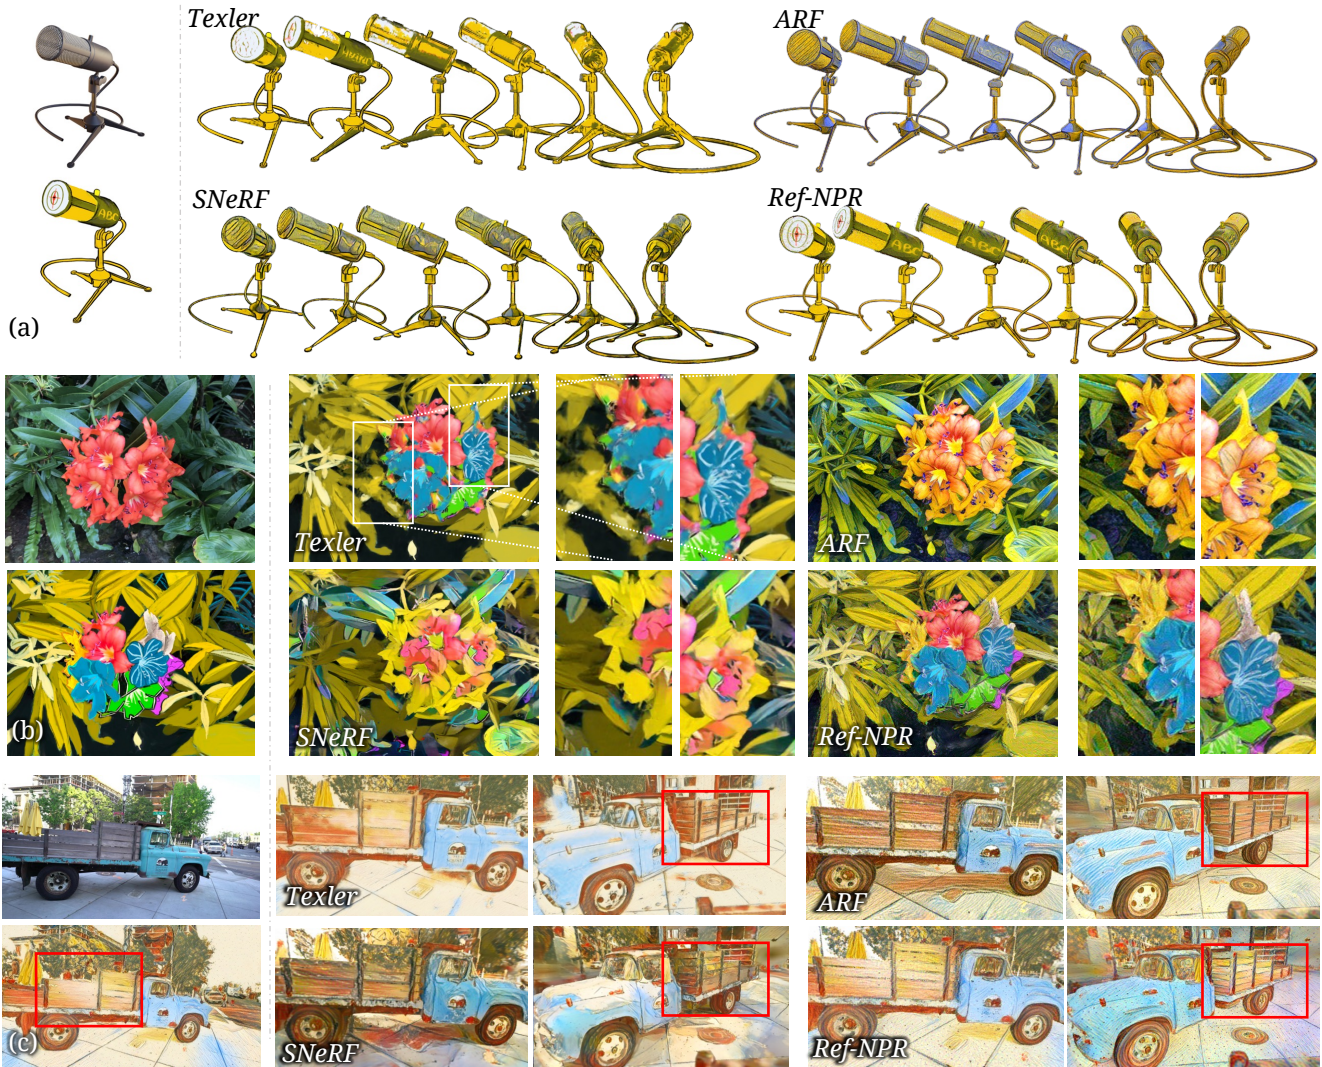


Figure 5. Qualitative comparisons in novel-view stylization. For each example, we provide the reference view (top) and its corresponding stylized reference view (bottom) on the left. We compare Ref-NPR with other methods in Synthetic (a) [28], LLFF (b) [27], zoomed-in on the flower and the occluded regions, and T&T (c) [21]. Semantic consistencies are highlighted.

have passed through the activation layers in stages 3 and 4 ( $relu\_3\_*$ ,  $relu\_4\_*$ ), and use them for  $\mathcal{L}_{feat}$ . Balancing factors among loss terms are set to  $\lambda' = 5 \times 10^{-3}$ ,  $\lambda_c = 5$ , and  $\lambda_r = \lambda_f = 1$ . We train each scene on one NVIDIA 3090 GPU for 10 epochs. Before training the final 3 epochs, for a smoother content update, we replace  $\omega_P$  with a frozen  $\omega_{NP}$  as the content view generator in TCM, and minimize  $\mathcal{L}_{ref}$  and  $\mathcal{L}_{feat}$  only, with  $\lambda_f = 0.2$  and  $\lambda' = 0$ .

## 4.2. Datasets

We evaluate the performance of Ref-NPR on three standard datasets.

**Synthetic** [28] is a well-defined synthetic dataset for 3D objects. Stylizing the views of 3D models with a foreground mask using a full-image style reference is a challenging yet

meaningful task, which is missed in many existing scene stylization methods [15, 16, 30, 45]. To avoid overflow of the stylized view on the mask boundary, we apply a 2D morphological erosion on the foreground mask in  $\mathbb{R}^3$ .

**LLFF** [27] is a real-world high-resolution dataset for novel view synthesis, and we choose a resolution that is  $4\times$  down-sampled, following [10, 45].

**Tanks and Temples (T&T)** [21] is a  $360^\circ$  scene dataset for novel view synthesis and scene reconstruction with 200 to 300 high-resolution training views for each scene. In  $\mathbb{R}^3$ , we register only the rays in Plenoxels’s foreground model [10] apart from the reference view to make the dictionary more compact. We follow the official training and testing splits for all datasets.




Figure 6. Ablations on the effectiveness of respective loss components.

### 4.3. Comparisons

We compare our Ref-NPR with two recent scene stylization methods: ARF [45] and a reimplemented SNeRF on Plenoxels [10, 30]. Besides, we include a reference-based video stylization method by Texler et al. [38] for a more comprehensive comparison.

**Qualitative comparison.** We present qualitative comparisons with state-of-the-art video and scene stylization methods in Fig. 5. The method of Texler et al. [38] produces novel views with proper low-level color distribution but suffers from limited correspondence in the scene, leading to flickering effects as shown in Fig. 5(a). ARF [45] and SNeRF [30] maintain geometric consistency in novel views but lack content-style correlation, whether the style reference is manually created (Fig. 5(a-b)) or generated by a 2D-neural stylization method [22] in Fig. 5(c). In contrast, Ref-NPR significantly improves both geometric and semantic-style consistency in each example. Additionally, Ref-NPR benefits from TCM to fill occluded regions with perceptually reasonable visual contents. More detailed comparisons can be found in the supplementary materials.

**Quantitative comparison.** In order to evaluate the quality of reference-based stylization, we adopt a metric used by CCPL [42] that measures the consistency between the style reference image  $S_R$  and the top 10 nearest test novel views using LPIPS [47]. Additionally, we report the long-range cross-view consistency ability as a stability metric, following the approach of SNeRF. We also evaluate the robustness of the methods by running them iteratively and measuring the PSNR between rendered results. Further details on the evaluation metrics are provided in the appendix. The results presented in Tab. 1 demonstrate that Ref-NPR achieves higher perceptual similarity and cross-view geometric consistency than state-of-the-art scene stylization methods. Moreover, we conduct a user study to evaluate the perceptual quality of the stylization sequences. We apply ten stylization sequences of four shuffled methods and

metric	Texler	SNeRF	ARF	Ref-NPR
Stability	✗	✓	✓	✓
Ref-LPIPS ↓	<b>0.335</b>	0.405	0.394	<b>0.339</b>
Robustness ↑	18.69	26.03	26.34	<b>28.11</b>

Table 1. Quantitative comparisons. We evaluate the stability, reference similarity, and robustness among stylization methods. Methods except Texler get the same stability as they adopt the same optimized density field.

metric	Texler	SNeRF	ARF	Ref-NPR
Avg. rank ↓	2.92	2.83	2.52	<b>1.73</b>
Ours preference	<b>79.3%</b>	<b>78.1%</b>	<b>70.6%</b>	-

Table 2. User study results. We report the average ranking and pairwise preference rate of Ref-NPR.

collect 33 responses for each sequence based on the visual quality. The user study results are presented in Tab. 2.

### 4.4. Ablation Studies

We perform several ablation studies to investigate the effectiveness of the different components of Ref-NPR. First, we conduct a module-wise ablation to evaluate the contributions of the supervision components in  $R^3$  and TCM. Our results in Fig. 6 (a-d) demonstrate that the removal of any of these components leads to a degradation in the quality of the stylized images. For instance, removing the pseudo-ray supervision  $\mathcal{L}_{ref}$  (a) or supervising on the reference view only (b) results in a loss of detailed textures. While discarding the feature supervision  $\mathcal{L}_{feat}$  (c) causes a photo-realistic colorization in occluded regions. Moreover, neglecting the color supervision  $\mathcal{L}_{color}$  (d) can cause color mismatches in occluded regions. The combination of  $R^3$  and TCM in the full Ref-NPR model compensates for these shortcomings and leverages the strengths of each module.

We also validate the effectiveness of implicit supervision in TCM. As depicted in Fig. 7, we observed that without




Figure 7. Ablation study on TCM. Benefiting from multi-view correspondence, TCM can effectively fill meaningful styles in occluded regions.

TCM, the occluded regions often failed to obtain desired correspondence due to the lack of supervision in the semantic feature space. TCM effectively matches features within the same content domain. This demonstrates the crucial role of TCM in achieving accurate and consistent stylization.

## 5. Discussions

**Adapt to general style transfer.** We have demonstrated the ability of Ref-NPR to use arbitrary style images as reference for style transfer. In particular, as shown in Fig. 8, we generate three reference views of the same scene using different 2D stylization methods [11, 20, 22] and feed them into Ref-NPR to render three sets of stylized novel views, each preserving the characteristics of the corresponding style reference. This extension allows Ref-NPR to be more flexible in working with various style reference images and provides a more controllable solution than other scene stylization methods [15, 16, 30, 45].

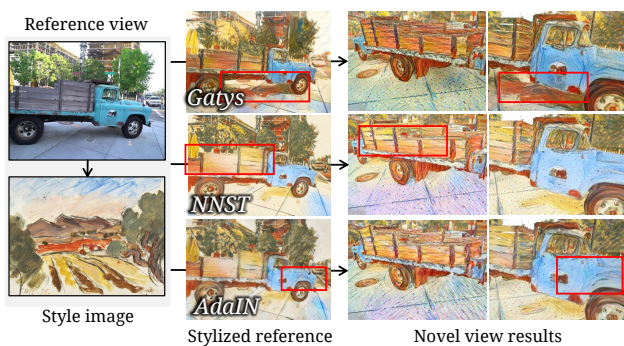


Figure 8. The pipeline of Ref-NPR naturally extends to arbitrary style reference. Images are cropped for a better presentation. Method-related style textures are highlighted.

**Multi-reference.** To accommodate the stylization of large-scale scenes, it is essential to extend Ref-NPR to handle multiple style references. This can be achieved by reg-

istering rays using all stylized reference views in  $R^3$  and expanding the capacity of styles and content features in TCM. An example of multi-reference input in the Playground scene [21] is shown in Fig. 9. Using two additional stylized views, Ref-NPR achieves better feature matching and richer style content.

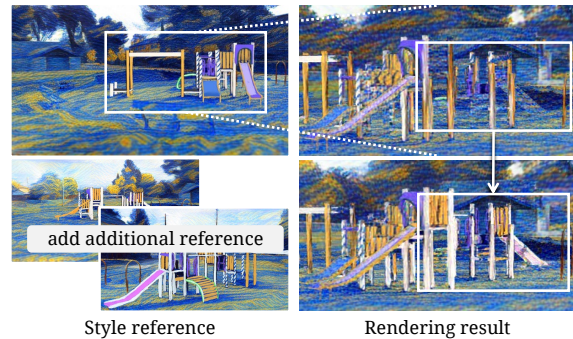


Figure 9. Multi-reference results of a 360° scene.

**Limitations.** Ref-NPR is a versatile method. However, it may not perform well when no meaningful semantic correspondence is found in the reference view, as depicted in Fig. 10. Additionally, feature matching may also fail in stylizing objects with intricate geometric structures.




Figure 10. Failure cases of Ref-NPR. Feature matching may fail when (a) stylizing intricate geometric structures, or (b) stylizing a large scene with a single reference.

## 6. Conclusion

This paper introduces Ref-NPR, a new framework for controllable non-photorealistic 3D scene stylization based on radiance fields. Ref-NPR can generate high-quality semantic correspondence and geometrically consistent stylizations for novel views by utilizing a stylized reference view. The proposed framework can potentially enhance the efficiency of human creativity in professional visual content creation.

**Acknowledgement.** This work is partially supported by ITF Partnership Research Programme (No.PRP/65/20FX) and Shenzhen Science and Technology Program KQTD20210811090149095. We are grateful to Shaozuo Yu and Mingxuan Zuo for their meaningful discussions.



## References

- [1] Sai Bi, Nima Khademi Kalantari, and Ravi Ramamoorthi. Patch-based optimization for image-based texture mapping. *ACM Trans. Graph.*, 36(4), 2017. 2
- [2] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *ECCV, 2022*. 1
- [3] Dongdong Chen, Jing Liao, Lu Yuan, Nenghai Yu, and Gang Hua. Coherent online video style transfer. In *ICCV*, pages 1105–1114, 2017. 2
- [4] Tian Qi Chen and Mark W. Schmidt. Fast patch-based style transfer of arbitrary style. *ArXiv*, abs/1612.04337, 2016. 2
- [5] 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 *CVPR*, pages 1475–1484, 2022. 1, 2
- [6] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised NeRF: Fewer views and faster training for free. In *CVPR*, June 2022. 4
- [7] Yingying Deng, Fan Tang, Weiming Dong, Haibin Huang, Chongyang Ma, and Changsheng Xu. Arbitrary video style transfer via multi-channel correlation. In *AAAI*, volume 35, pages 1210–1217, 2021. 2
- [8] Zhiwen Fan, Yifan Jiang, Peihao Wang, Xinyu Gong, Dejia Xu, and Zhangyang Wang. Unified implicit neural stylization. In *ECCV*, pages 636–654. Springer, 2022. 1, 2
- [9] Jakub Fišer, Ondřej Jamriška, Michal Lukáč, Eli Shechtman, Paul Asente, Jingwan Lu, and Daniel Šykora. Stylit: illumination-guided example-based stylization of 3d renderings. *ACM Trans. Graph.*, 35(4):1–11, 2016. 2
- [10] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 1, 5, 6, 7
- [11] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, pages 2414–2423, 2016. 2, 3, 5, 8
- [12] Bruce Gooch and Amy Gooch. *Non-photorealistic rendering*. AK Peters/CRC Press, 2001. 2
- [13] Filip Hauptfleisch, Ondrej Texler, Aneta Texler, Jaroslav Krivánek, and Daniel Šykora. Styleprop: Real-time example-based stylization of 3d models. In *Computer Graphics Forum*, volume 39, pages 575–586. Wiley Online Library, 2020. 2
- [14] Mingming He, Jing Liao, Dongdong Chen, Lu Yuan, and Pedro V Sander. Progressive color transfer with dense semantic correspondences. *ACM Trans. Graph.*, 38(2):1–18, 2019. 2
- [15] Hsin-Ping Huang, Hung-Yu Tseng, Saurabh Saini, Maneesh Singh, and Ming-Hsuan Yang. Learning to stylize novel views. In *ICCV*, 2021. 1, 2, 6, 8
- [16] 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 *CVPR*, 2022. 1, 2, 4, 5, 6, 8
- [17] Ondřej Jamriška, Šárka Sochorová, Ondřej Texler, Michal Lukáč, Jakub Fišer, Jingwan Lu, Eli Shechtman, and Daniel Šykora. Stylizing video by example. *ACM Trans. Graph.*, 38(4):1–11, 2019. 1, 2
- [18] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711. Springer, 2016. 2
- [19] James T Kajiya and Brian P Von Herzen. Ray tracing volume densities. *ACM Trans. Graph.*, 18(3):165–174, 1984. 3
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019. 2, 3, 8
- [21] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Trans. Graph.*, 36(4), 2017. 4, 6, 8
- [22] Nicholas Kolkin, Michal Kucera, Sylvain Paris, Daniel Šykora, Eli Shechtman, and Greg Shakhnarovich. Neural neighbor style transfer. *arXiv e-prints*, pages arXiv–2203, 2022. 2, 5, 7, 8
- [23] Jan Eric Kyprianidis, John Collomosse, Tinghui Wang, and Tobias Isenber. State of the” art”: A taxonomy of artistic stylization techniques for images and video. *TVCG*, 19(5):866–885, 2012. 2
- [24] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. *NeurIPS*, 30, 2017. 2
- [25] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *ECCV*, pages 453–468, 2018. 2
- [26] Jing Liao, Yuan Yao, Lu Yuan, Gang Hua, and Sing Bing Kang. Visual attribute transfer through deep image analogy. *ACM Trans. Graph.*, 36(4):120, 2017. 2
- [27] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Trans. Graph.*, 2019. 6
- [28] 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. 1, 2, 3, 4, 6
- [29] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. 1
- [30] Thu Nguyen-Phuoc, Feng Liu, and Lei Xiao. Snerf: stylized neural implicit representations for 3d scenes. *arXiv preprint arXiv:2207.02363*, 2022. 1, 2, 4, 6, 7, 8
- [31] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4531–4540, 2019. 2
- [32] Manuel Ruder, Alexey Dosovitskiy, and Thomas Brox. Artistic style transfer for videos and spherical images. *IJCV*, 126(11):1199–1219, 2018. 2
- [33] Ahmed Selim, Mohamed Elgharib, and Linda Doyle. Painting style transfer for head portraits using convolutional neural networks. *ACM Trans. Graph.*, 35(4):1–18, 2016. 2

- [34] YiChang Shih, Sylvain Paris, Connelly Barnes, William T Freeman, and Frédo Durand. Style transfer for headshot portraits. *ACM Trans. Graph.*, 2014. [2](#)
- [35] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015. [4](#), [5](#)
- [36] Daniel Šykora, Ondřej Jamriška, Ondrej Texler, Jakub Fišer, Michal Lukáč, Jingwan Lu, and Eli Shechtman. Styleblit: Fast example-based stylization with local guidance. In *Computer Graphics Forum*, volume 38, pages 83–91. Wiley Online Library, 2019. [2](#)
- [37] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P. Srinivasan, Jonathan T. Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. In *CVPR*, pages 8248–8258, June 2022. [1](#)
- [38] Ondřej Texler, David Futschik, Michal Kučera, Ondřej Jamriška, Šárka Sochorová, Menglei Chai, Sergey Tulyakov, and Daniel Šykora. Interactive video stylization using few-shot patch-based training. *ACM Trans. Graph.*, 39(4):73, 2020. [1](#), [2](#), [7](#)
- [39] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *CVPR*, pages 3835–3844, 2022. [2](#)
- [40] Can Wang, Ruixiang Jiang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Nerf-art: Text-driven neural radiance fields stylization. *arXiv preprint arXiv:2212.08070*, 2022. [2](#)
- [41] Wenjing Wang, Shuai Yang, Jizheng Xu, and Jiaying Liu. Consistent video style transfer via relaxation and regularization. *TIP*, 29:9125–9139, 2020. [2](#)
- [42] Zijie Wu, Zhen Zhu, Junping Du, and Xiang Bai. Ccpl: Contrastive coherence preserving loss for versatile style transfer. In *ECCV*, 2022. [2](#), [7](#)
- [43] Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Humphrey Shi, and Zhangyang Wang. Sinnerf: Training neural radiance fields on complex scenes from a single image. In *ECCV*, 2022. [4](#)
- [44] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *CVPR*, pages 5438–5448, 2022. [4](#)
- [45] Kai Zhang, Nick Kolkin, Sai Bi, Fujun Luan, Zexiang Xu, Eli Shechtman, and Noah Snavely. Arf: Artistic radiance fields. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, pages 717–733. Springer, 2022. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [46] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020. [1](#)
- [47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [7](#)