

## $\mathcal{Z}^*$ : Zero-shot Style Transfer via Attention Reweighting

Yingying Deng<sup>o,1</sup>, Xiangyu He<sup>o,1</sup>, Fan Tang<sup>✉,2</sup>, Weiming Dong<sup>1</sup><sup>1</sup> MAIS, Institute of Automation, Chinese Academy of Sciences<sup>2</sup> Institute of Computing Technology, Chinese Academy of Sciences

dyy15@outlook.com, xiangyu.he@nlpr.ia.ac.cn, tfan.108@gmail.com, weiming.dong@ia.ac.cn




Figure 1. Image style transfer results by the proposed  $\mathcal{Z}^*$ . **Top:** The stylized results by style/content references of different types. Our method can well balance the contents and styles in the results. **Bottom:** Comparisons with state-of-the-art methods, including diffusion-based models (VCT [9] and InST [51]), transformer-based model (StyTr<sup>2</sup> [13]), flow-based model (ArtFlow [2]), and CNN-based model (CAST [50]). Our method excels in generating stylized images with vivid style patterns and accurate content details.

### Abstract

Despite the remarkable progress in image style transfer, formulating style in the context of art is inherently subjective and challenging. In contrast to existing methods, this study shows that vanilla diffusion models can directly extract style information and seamlessly integrate the generative prior into the content image without retraining. Specifically, we adopt dual denoising paths to represent content/style references in latent space and then guide the content image denoising process with style latent codes. We further reveal that the cross-attention mechanism in latent diffusion models tends to blend the content and style images, resulting in stylized outputs that deviate from the original content image. To overcome this limitation, we introduce a cross-attention reweighting strategy. Through theoretical analysis and experiments, we demonstrate the effectiveness and superiority of the diffusion-based zero-shot style transfer via attention reweighting,  $\mathcal{Z}$ -STAR.

### 1. Introduction

The task of image style transfer has received significant attention in the research community, with numerous machine learning techniques utilized, such as convolutional neural networks (CNN) [11, 15, 18, 21, 32, 38], flow-based networks [2], visual transformers (ViT) [13, 46], and diffusion models [9, 51]. After completing the training process, the output stylized image is generated based on a content and style image, as shown in Figure 1. The generated image retains the content layout from the input content while adopting a similar style to that of the input style. In essence, the networks are trained by ensuring that the generated image and the content/style image exhibit content/style similarity.

While certain methods, such as [2, 13, 18], employ the Gram matrix [15] to measure global style similarity and achieve promising style transfer outcomes, the second-order statistics contained in the Gram matrix are limited in their ability to capture intricate style patterns and fail to transfer corresponding local features from the content image to the style image (e.g., hair and eyes in StyTr<sup>2</sup> [13],

<sup>o</sup>These authors contributed equally to this work.<sup>✉</sup>Corresponding author: Fan Tang.

ArtFlow [2], AdaIN [18] in Figure 1). To address this issue, CAST [50] proposes a contrastive loss that leverages the relationships between positive and negative examples to encourage the result to conform to the distribution of styles. However, CAST also faces the obstacle of generating stylized outcomes with vivid, fine-grained style details. In reality, *the contours and forms of a painting should be subject to the adaptable preferences of the artist’s painting techniques, rather than being rigidly determined by the content and style images*. In light of this, we rethink the role of training in style transfer and reveal that the generative model used to describe the distribution of images has already learned the art of transfer.

With the advent of diffusion models, text-controlled image editing and translation have gained unprecedented attention. When given an image as input, the diffusion model can generate an artistic image incorporating a style-related prompt. However, the textual prompt is often too coarse to effectively express the desired style details. While prior approaches such as InST [51] and VCT [9] endeavor to employ an image-controlled diffusion model for image style transfer and translation, they necessitate training a style embedding for each input style, leading to challenges in distilling precise style representations and resulting in deviations from the input style while failing to preserve content (see InST and VCT results in Figure 1). In contrast to [9, 51], where control information is encoded as text embedding, we propose that the vanilla diffusion model is capable of extracting style information directly from the desired style image and fusing it into the content image *without requiring re-training or tuning*.

In this paper, we leverage the prior knowledge from latent diffusion [36] and propose a zero-shot (*i.e.*, training-free) style transfer method via attention reweighting, namely  $\mathcal{Z}$ -STAR ( $\mathcal{Z}^*$ ), to address the issues above. To obtain generative image priors, we employ dual diffusing paths to invert the style and content images. The features obtained from diffusion models naturally represent content and style information and could be fused by the attention mechanism. However, without a training process, it is challenging to strike a balance between content and image influence. That is to say, the naive cross-attention operations are not optimal for directly integrating content and style latent status in the denoising process. Content structure may be compromised due to inaccurate cross-attention values (see Sec. 4.2 for more details). Therefore, we propose a multi-cross attention reweighted strategy that manipulates content and style information from images and seamlessly fuses them in the diffusion latent space. By leveraging a tailored attention mechanism, the diffusion model can naturally address the constraints of content and style without necessitating additional supervision. Experimental results demonstrate that our method generates satisfactory results

with well-preserved content and vivid styles adapted to content structures. In summary, our main contributions are as follows:

- A zero-shot image style transfer method leveraging the generative prior knowledge to conduct image stylization without retraining/tuning.
- A reweighted attention mechanism to disentangle and fuse content/style information in the diffusion latent space.
- Various experiments demonstrate that our method can generate outstanding style transfer results, naturally fusing and balancing content and style from two input images.

## 2. Related Work

**Image style transfer.** Since Gatys *et al.* [15] discovered that perceptual features can effectively represent content and Gram matrices may express style in CNNs, various frameworks for style transfer trained by content and style loss have been proposed [1, 11, 12, 18, 20, 23, 27, 34, 38, 42]. CNN-based methods [11, 12, 23, 38] achieved success by exploring the fusion of content and style representation. Some works [13, 39, 43, 44, 46, 49] utilize the long-range feature representation ability of Transformer [47] and enhance stylization effects. However, the Gram matrix measures second-order statistics of the entire image, which may not be sufficient for style representation. Recent works [7, 48, 50] use contrastive loss to replace style loss based on the Gram matrix, which is effective in processing fine-detail style patterns. Despite the continuous progress of existing methods, precise style representation remains challenging, and inaccurate style expression may lead to unsatisfactory stylized results. In light of this challenge, we aim to develop a zero-shot style transfer method that does not rely on explicit style constraints.

**Diffusion for image generation.** Diffusion models have demonstrated impressive results in text-to-image generation [29, 33, 36, 37] and image editing [3, 5, 6, 10, 16, 22, 28, 30, 40, 45, 52]. However, certain methods like Imagic [22] require fine-tuning the entire diffusion model for each instruction, which can be time-consuming and memory-intensive. To address this, Prompt-to-prompt [16] introduces cross-attention maps during the diffusion process by replacing or reweighting the attention map between text prompts and edited images. Additionally, NTI [28] proposes null-text optimization based on Prompt-to-prompt to enable real image editing. In order to reduce reliance on text prompts, StyleDiffusion [24] incorporates a mapping network to invert the input image to a context embedding, which is then utilized as a key in the cross-attention layers. However, manipulating the cross-attention solely between text and image can be challenging for achieving precise control. To address this, Plug-and-Play [40] and MasaCtrl [6] focus on spatial features using self-attention in

the U-Net of the latent diffusion model. While these methods can accomplish text-guided style transfer by inputting a text prompt like “a pencil drawing”, simple words may not be sufficient to describe fine-detail style patterns. To address this limitation, InST [51] and VCT [9] employ an inversion-based image style transfer/translation scheme that can train a style image into a style embedding to guide the generated results. In this paper, we demonstrate that style images alone (*i.e.*, without pseudo-text guidance) are adequate for latent diffusion models to achieve image-guided style transfer, without requiring additional training.

### 3. Preliminary

**Attention Mechanism** [4] was introduced as a powerful tool in neural network architectures for aggregating information and later adopted by Vaswani et al. [41] as a fundamental building block for machine translation:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V. \quad (1)$$

Attention-based vision transformers [14, 26] have demonstrated remarkable empirical results on mainstream benchmarks, solidifying attention mechanisms as a crucial component in modern deep neural networks. Moreover, by incorporating additional information, along with the utilization of *Key* and *Value* vectors, cross-attention has proved effective in latent diffusion models for applying conditions to the denoising process.

**Diffusion Model**, as described in the literature, belongs to a class of generative models that employ Gaussian noise to generate desired data samples. This is accomplished through an iterative process of noise removal, where a forward process is defined to add noise to an initial data sample  $x_0$ , resulting in a noisy sample  $x_t$  at time-step  $t$ , according to a predetermined noise-adding schedule  $\alpha_t$ :

$$x_t = \sqrt{\alpha_t} \cdot x_0 + \sqrt{1 - \alpha_t} \cdot z, \quad z \sim \mathcal{N}(0, \mathbf{I}). \quad (2)$$

Additionally, a corresponding reverse process is also defined:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(\mathbf{x}_t, t) \right) + \sigma_t z. \quad (3)$$

The backward process aims to gradually denoise  $x_T \sim \mathcal{N}(0, \mathbf{I})$ , where a cleaner image  $x_{t-1}$  is obtained at each step. This is accomplished by a neural network  $\epsilon_\theta(x_t, t)$  that predicts the added noise  $z$ .

Utilizing a U-Net integrated with an attention mechanism as  $\epsilon_\theta(x_t, t)$  is a common approach. This configuration allows for self-attention to capture long-range interactions among image features, while cross-attention receives a guiding signal from the given text prompt. The attention mechanisms are formulated as:

$$f_t^l = \text{Attention}(Q_t^l, K_t^l, V_t^l). \quad (4)$$

Even as *Key* and *Value* at  $l$ -th layer may vary from image spatial or text features, they still adhere to the standard format.

## 4. Method

Our research is based on the observation that the attention module within the stable diffusion can effectively align features  $K$  and  $V$  with the query  $Q$ . Previous studies, such as [6, 40], have leveraged the self-attention layer to extract information from key and value features, which represent spatial attributes in the DDIM inversion process, for image editing applications. However, in style transfer tasks, the simultaneous preservation of both style and content is crucial. Consequently, we need to address two important questions to practically apply this observation:

- How can appropriate style features  $K$  and  $V$  be obtained for the stable diffusion model without necessitating re-training? It is crucial to emphasize that our objective is to directly extract style information from images rather than depending on a proxy prompt like text embedding, which may lack the necessary level of detail.
- Mere utilization of style attention [31] results in poor content preservation under diffusion models without re-training. We aim to recreate the content image in a way that resembles the style image in terms of its contours, forms, and overall visual appearance.

To address these challenges, our approach adopts a two-fold strategy. Firstly, dual-path networks are designed to generate suitable  $K$  and  $V$  features during the reverse process. Secondly, we incorporate attention reweighting techniques to better align content features with style features.

### 4.1. Dual-path Networks

In our proposed approach, we address a key limitation in the conventional stable diffusion model, where the text embedding remains unchanged throughout the reverse process, from timestamp  $t \in [0, T]$ , despite the desirable need for the style feature to adapt to the denoising stylized image. This adaptation is crucial as the initial denoising process involves the reconstruction of the image’s shape and color, followed by the refinement of details such as contours and brushstrokes towards the end.

To address this limitation, we introduce a novel dual-path scheme that simultaneously generates the denoised style image and the stylized content image at the same timestamp  $T$ . This is achieved through the following equations:

$$I_s = \mathcal{G}_\theta(\epsilon_{I_s}, \{f_s\}, T), \quad I_c = \mathcal{G}_\theta(\epsilon_{I_c}, \{f_c\}, T). \quad (5)$$

This ensures that the features in both networks are naturally aligned in the time dimension. Specifically, given a content image  $I_c$  and a style image  $I_s$ , our objective is to obtain a stylized result  $\hat{I}_c$  that retains the content of  $I_c$  while

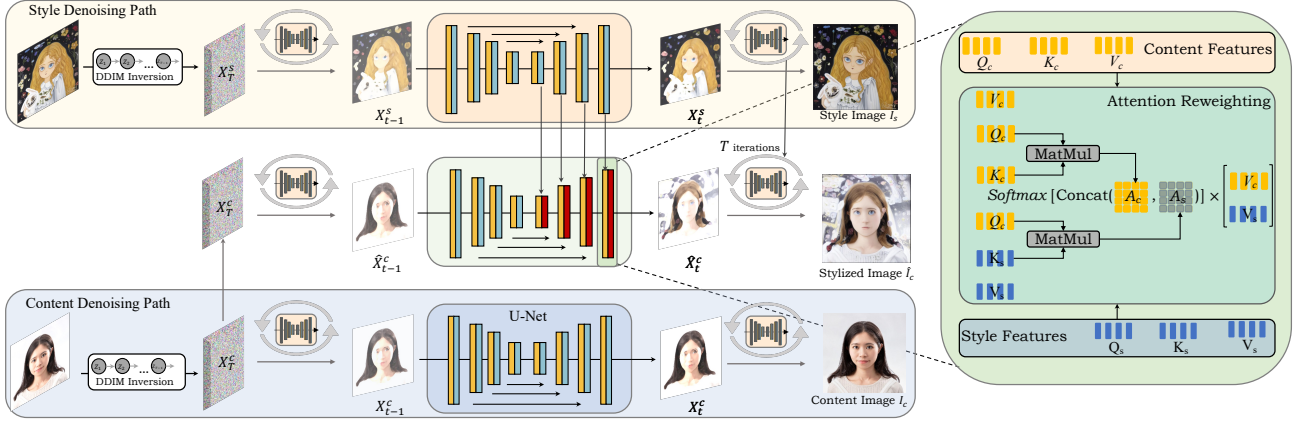


Figure 2. Overall pipeline of our style transfer framework. The stylization process operates in the latent space. We perform DDIM inversion separately for the content and style images. During the denoising process, our Cross-attention Reweighting is employed to integrate style patterns into the content structure. By iteratively performing 50 denoising steps, we are able to achieve the final stylized output.

incorporating stylistic patterns from  $I_s$ . This is achieved through the following equation:

$$\hat{I}_c = \mathcal{G}_\theta(\epsilon_{I_c}, \{f_s, f_c\}, T), \quad (6)$$

where  $\mathcal{G}_\theta(\cdot, \cdot, T)$  represents the application of denoising for  $T$  steps in the diffusion model, using fixed pre-trained weights  $\theta$ . The term  $\epsilon_{I_s}$  denotes the noisy  $x_T$  generated in the forward process by progressively adding Gaussian noise to  $I_c$  or  $I_s$ , as described in Eq. (2). The notation  $\{f_s, f_c\}$  refers to the spatial U-Net features in the diffusion models from the style and content images, respectively, which are utilized in cross-attention.

As shown in Figure 2, we utilize DDIM inversion to invert the style image and content image, obtaining the diffusion trajectories  $x_{[0:T]}^c$  and  $x_{[0:T]}^s$ . Subsequently, we introduce a novel cross-attention arrangement to disentangle and fuse content and style information (i.e.,  $f_c$  and  $f_s$ , denoted as Query, Key and Value), within the diffusion latent space using U-Net at timestamp  $t$ . Through  $T$  denoising steps, we convert the stylized latent features  $\hat{f}_c$ , generated by the reweighted attention, into the style transfer result  $\hat{I}_c$ .

## 4.2. Attention Reweighting

As demonstrated in Eq. (6), our attention mechanism incorporates two types of attention calculation between  $f_s$  and  $f_c$ : style-cross attention for merging content and style features, and content self-attention for preserving structure. Since we use the standard self-attention, we mainly discuss the proposed style-cross attention in this section.

**Naive Setting** It is intuitive to represent content information, such as image structure, using the Query, and represent style information, such as color, texture, and object shape, using the Key and Value features. The style-cross attention

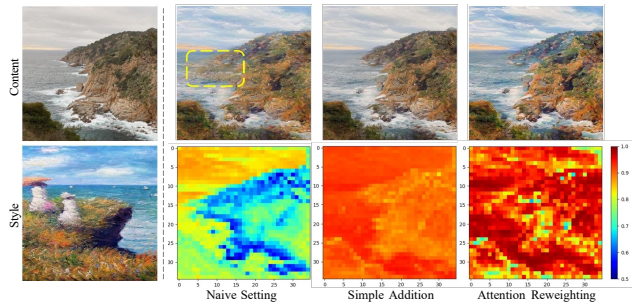



Figure 3. The outcome of the “Naive Setting” exhibits a bias towards emphasizing the style patterns while neglecting the original content structures. The distorted area corresponds to the low similarity portion between self-attention and cross-attention results, suggesting that the  $V_s$  fails to adequately reconstruct the target region, resulting in a loss of content. “Simple Addition” preserves an excessive number of content features, whereas the re-arranged attention achieves a more favorable trade-off.

then uses the content features to query the information from style images that best suit the input patch. Formally, the inputs of style-cross attention are features from the content latent space  $c$  and style latent space  $s$  where

$$\hat{f}_c = \text{Attn}(Q_c, K_s, V_s) = \text{Softmax}\left(\frac{Q_c K_s^T}{\sqrt{d}}\right) V_s. \quad (7)$$

Despite the simplicity, we observe that the naive fusion setting in Eq. (7) tends to prioritize the style patterns at the expense of the original content structures. Figure 3 displays the heatmap representing the cosine similarity between the results of cross-attention  $\text{Attn}(Q_c, K_s, V_s)$  and self-attention  $\text{Attn}(Q_c, K_c, V_c)$ . It is observed that regions with low similarity scores correspond to pixels that have experienced a loss of content information.



(a)  $\vec{q}_c K_s$  (b)  $\text{Softmax}(\vec{q}_c K_s)$  (c) Ours

Figure 4. Visualization of the distribution of  $\vec{q}_c K_s$  for two content feature points (represented by blue and red bars) before and after the Softmax in (a) and (b). In (c), we display the distribution of  $\vec{q}_c K_s$  normalized by Eq.(11). It is observed that Softmax operation tends to overly amplify smaller values of  $\vec{q}_c K_s$  (e.g., blue bars being shifted further to the right compared to the red bars in (b), even though the original  $\vec{q}_c K_s$  values are mostly less than 0). In contrast, the negatively correlated values are limited to small magnitudes before and after normalization in (c).

**Simple Addition** To tackle the aforementioned issue, we propose a simple solution by enhancing the content information in  $\hat{f}_c$  through the reintroduction of content self-attention. The equation is formulated as follows:

$$\hat{f}_c = \lambda \cdot \text{Attn}(Q_c, K_s, V_s) + (1 - \lambda) \cdot \text{Attn}(Q_c, K_c, V_c), \quad (8)$$

where  $\lambda \in [0, 1]$ . However, we found the selection of  $\lambda$  can be delicate. For instance, certain content pixels exhibit a weak correlation with style, as indicated by small values of  $\vec{q}_c K_s^T$  (represented by blue bars in Figure 4(a), where  $\vec{q}_c K_s^T$  is less than 0). It is desirable to assign smaller attention weights to these pixels in order to minimize their negative impact. Conversely, in other scenarios where  $\vec{q}_c K_s^T$  yields meaningful values (represented by red bars in Figure 4(a)), we aim to assign larger attention weights to these pixels. Unfortunately, due to the inherent nature of the Softmax function, which disregards the absolute magnitudes and solely amplifies the differences between  $\vec{q}_c K_s^T$  values, we observe counter-intuitive results as depicted in Figure 4(b), where smaller  $\vec{q}_c K_s^T$  values result in larger attention weights after Softmax normalization. In such cases, we need to introduce an additional variable, denoted as  $\lambda$ , to compensate for this deficiency. However, it is worth noting that a pre-defined  $\lambda$  value cannot cater to the requirements of every content/style image pair.

**Cross-attention Reweighting** Though a handcraft  $\lambda$  outside of  $\text{Attn}(\cdot, \cdot, \cdot)$  can not adapt to input images, we find that it can be achieved with a  $\lambda$  inside  $\text{Softmax}(\cdot)$ . By expressing Eq. (8) in matrix form through an equivalent reformulation, we obtain:

$$\hat{f}_c = \left[ \lambda \cdot \sigma\left(\frac{Q_c K_s^T}{\sqrt{d}}\right), (1 - \lambda) \cdot \sigma\left(\frac{Q_c K_c^T}{\sqrt{d}}\right) \right] * \begin{bmatrix} V_s \\ V_c \end{bmatrix} \quad (9)$$

$$= A * V'^T. \quad (10)$$

Here,  $\sigma(\cdot)$  represents the Softmax function, and each row in  $A \in \mathbb{R}^{N \times 2N}$ , denoted as  $\vec{a} \in \mathbb{R}^{2N}$ , is normalized, i.e.,  $\vec{a} \cdot \vec{1}^T = 1$ . This normalization inspires us to re-construct the matrix  $A$  in the form of applying Softmax to rows, i.e.,

$$A' = \sigma\left(\left[\lambda \cdot \frac{Q_c K_s^T}{\sqrt{d}}, \frac{Q_c K_c^T}{\sqrt{d}}\right]\right) \quad (11)$$

$$\hat{f}_c = A' * V'^T = \sigma\left(\left[\lambda \cdot \frac{Q_c K_s^T}{\sqrt{d}}, \frac{Q_c K_c^T}{\sqrt{d}}\right]\right) * \begin{bmatrix} V_s \\ V_c \end{bmatrix}. \quad (12)$$

Specifically, Eq. (11) integrates  $\lambda$  within Softmax, which inherently normalizes the output to the  $[0, 1]$  range, eliminating the requirement for  $1 - \lambda$ . In contrast to the previous attention formulation presented in Eq. (9), the newly proposed reweighted attention matrix  $A' \in \mathbb{R}^{N \times 2N}$  takes into account both the intra-content feature differences and the inter-content and style feature differences simultaneously during the application of the  $\text{Softmax}(\cdot)$  function for output normalization. The reweighted attention matrix effectively enhances significant values of both  $\vec{q}_c K_s^T$  and  $\vec{q}_c K_c^T$  at each pixel, while automatically suppressing the small values of  $\vec{q}_c K_s^T$  when the content pixel corresponding to  $\vec{q}_c$  is irrelevant to all style pixels.

**Superiority of Cross-attention Reweighting** The Cross-attention Reweighting can be considered as a more versatile formulation. Its properties can be demonstrated as follows: (i) In cases where the correlation between the style and content images is weak, i.e., when each element  $\vec{q}_c \vec{k}_s^T$  in  $Q_c K_s^T$  approaches  $-\infty$ , the modified attention  $\hat{f}_c = A' * V'^T$  reduces to the standard self-attention of content images, denoted as  $\text{Attention}(Q_c, K_c, V_c)$ . (ii) When the correlation between the style and content images is strong, if the maximum value of  $\vec{q}_c \vec{k}_s^T$  is approximately equal to the maximum value of  $\vec{q}_c \vec{k}_c^T$ , and the Softmax operation generates an approximate one-hot probability distribution, then  $\hat{f}_c = A' * V'^T$  is equivalent to Eq. (8). (iii) Last but not least, Eq. (8) can be rewritten using  $A'$  as follows:

$$\hat{f}_c = \frac{1}{2} \cdot \text{Attn}(Q_c, K_s, V_s) + \frac{1}{2} \cdot \text{Attn}(Q_c, K_c, V_c), \quad (13)$$

$$= \sigma\left(\left[\frac{Q_c K_s^T}{\sqrt{d}} + C, \frac{Q_c K_c^T}{\sqrt{d}}\right]\right) * \begin{bmatrix} V_s \\ V_c \end{bmatrix}, \quad (14)$$

where

$$C = \ln \frac{\sum_j \exp([Q_c K_c^T]_{\cdot, j})}{\sum_j \exp([Q_c K_s^T]_{\cdot, j})}. \quad (15)$$

In Eq. (14), the simple addition of self-attention to cross-attention introduces an additional term  $C$ . This variable serves to magnify all elements within  $Q_c K_s^T$ , including small values that represent weak correlations between style and content features, which are deemed inconsequential and




Figure 5. The illustrations of regional control. The control map is represented by the binary mask located at the bottom right of the content reference. Black pixels mark the areas designated for style transfer.

should be disregarded. Consequently, the incorporation of  $C$  may introduce an increased level of noise into the Softmax function, thereby resulting in sub-optimal outcomes.

**Conditional Control** The simplicity of Eq. (11) allows for easy extension of the attention reweighting technique to more complex downstream applications. To illustrate this, we introduce an additional mapping function  $\phi(\cdot)$  on  $\frac{Q_c K_c^T}{\sqrt{d}}$ , which provides enhanced control over a specific region  $\Omega$  for image style transfer. The modified equation is given by:

$$A' = \sigma\left(\left[\phi\left(\frac{Q_c K_c^T}{\sqrt{d}}\right), \frac{Q_c K_c^T}{\sqrt{d}}\right]\right). \quad (16)$$

Here,  $\phi(x_{i,j})$  is defined as:

$$\phi(x_{i,j}) = \begin{cases} -\infty & \{i, j\} \notin \Omega \\ x_{i,j} & \text{otherwise} \end{cases}. \quad (17)$$

It is important to note that directly setting the values to  $-\infty$  may lead to discontinuous style switching, resulting in an artificial sharp boundary. To achieve a more natural gradient effect, we utilize a linear gradient for  $\phi(x_{i,j})$ , transitioning from  $-\infty$  to  $x_{i,j}$ . We visually demonstrate the effect of conditional control in Figure 5, where  $\Omega$  region corresponds to the white pixels in the binary mask.

Furthermore, the extension of style transfer from a one-to-one content-style image pair to a one-to-many setting can be easily achieved using Eq. (11). In the case where we aim to transfer the style of  $N$  style images to a single content image, the equation is modified as follows:

$$A' = \sigma\left(\left[\frac{Q_c K_{s1}^T}{\sqrt{d}}, \dots, \frac{Q_c K_{sN}^T}{\sqrt{d}}, \frac{Q_c K_c^T}{\sqrt{d}}\right]\right). \quad (18)$$

Owing to the limited memory footprint, we consider this aspect as a potential area for future research and leave it for further investigation.

## 5. Experiments

### 5.1. Implementation Details

Our research builds upon the concept of Stable Diffusion [35] and utilizes the v1.5 checkpoint. In our experimental setup, the text prompts are configured as null character strings. The denoising process consists of a total of 30 steps. We introduce our cross-attention reweighting module between layers 10-15 during the 5<sup>th</sup>-30<sup>th</sup> denoising steps.

### 5.2. Evaluation

We conduct a comparative analysis of our proposed method against state-of-the-art style transfer approaches, including ArtFlow [2], AdaAttN [25], IEST [8], StyTr<sup>2</sup> [13], CAST [50], QuanArt [17], InST [51] and VCT [9].

**Qualitative evaluation.** The qualitative comparisons presented in Figure 6 provide a visual assessment of the outcomes achieved by different style transfer methods. ArtFlow [2] exhibits limitations regarding style representation and the smoothness of stylization results. AdaAttN [25] and IEST [8] exhibit inconsistencies in the generated output style compared to the input style reference. The results of QuanArt [17] suffer from a weakening of the style performance due to the prominence of the content appearance. CAST [50] and StyTr<sup>2</sup> [13] still fall short in terms of faithfully reproducing the artistic qualities, as “not real compared with artworks” due to unfaithful style loss constraints.

Diffusion-based approaches, InST [51] and VCT [9], encounter challenges in converging towards the optimal style embedding, leading to failure in generating a similar style and content as the input (as evident from the content deviation in Figure 1 and style deviation in Figure 6). VCT [9] demonstrates better style performance, however, still suffers from the issue of content bias, as evidenced by the complete alteration of the girl’s facial identity while not precisely matching the desired input style. Additionally, it is worth noting that InST [51] and VCT [9] require a training time of approximately 20 minutes per style image, whereas our proposed method does not require any training/tuning.

In contrast, our proposed approach does not rely on a style loss to enforce conformity of the original image to a different distribution. Unlike CNN-based methods, our model does not include a fixed pre-trained encoder, thereby reducing the loss of content and style information. As demonstrated in Figure 6, our method achieves captivating stylized results by effectively transferring style patterns, such as painting strokes and lines, onto input content images. These style patterns are skillfully adapted to the content semantics, as exemplified in the 1<sup>st</sup>, 3<sup>rd</sup>, and 7<sup>th</sup> rows of Figure 6.

**User study.** To quantitatively evaluate the impact of different stylization methods, we conducted a user study to gather

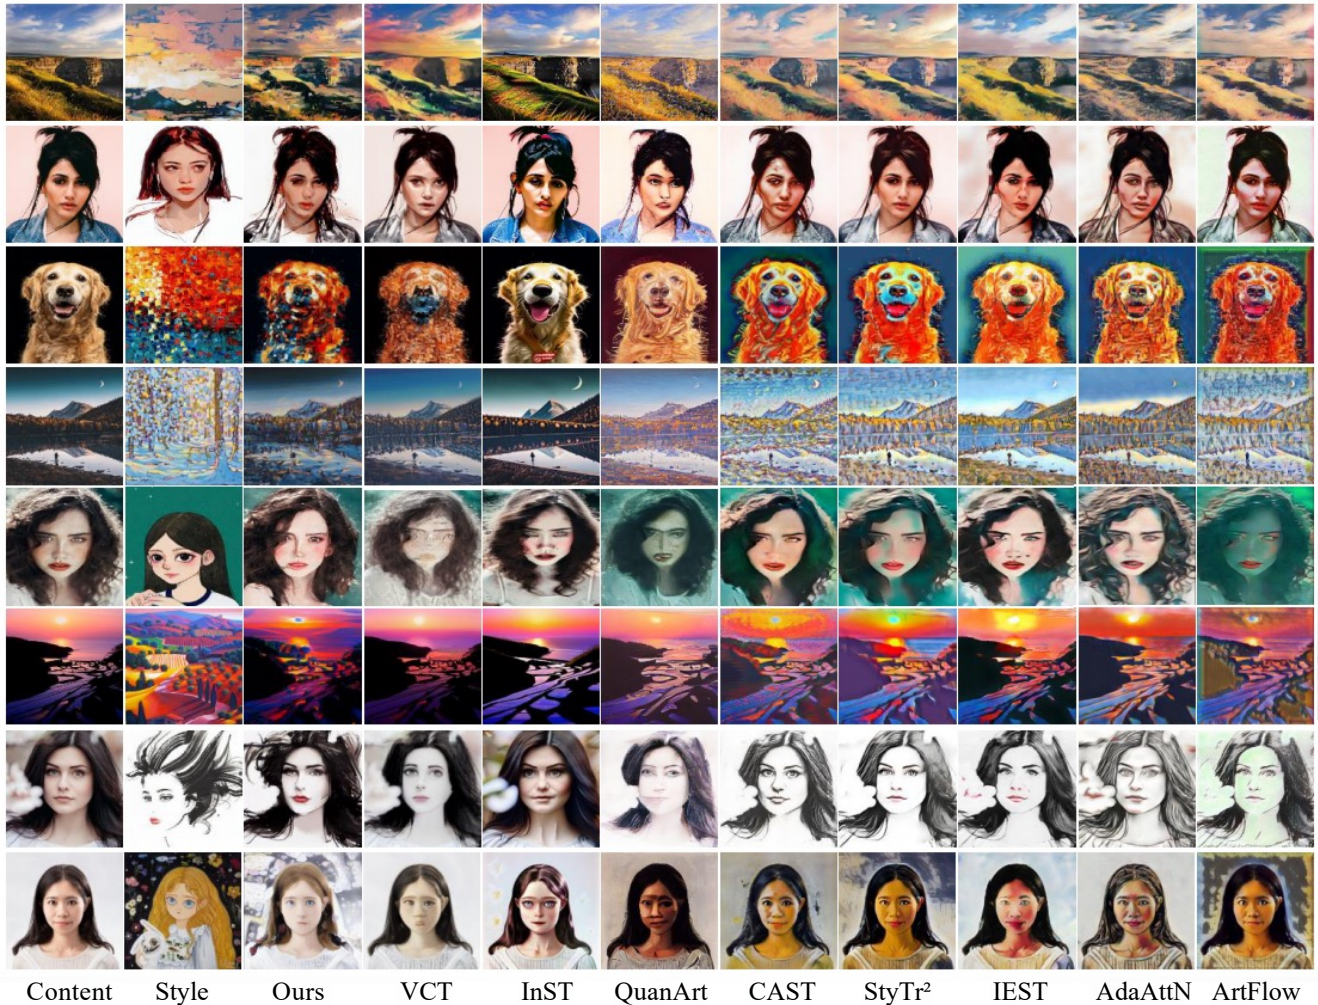


Figure 6. Compared with other style transfer methods. The content images are presented in first column, the style images are presented in second column, and the stylized results generated by different methods are presented in the rest.

Table 1. User study results. Each number represents the percentage of votes that the corresponding method is preferred to ours, using the criteria of overall quality, preservation of content and style, respectively.

	VCT [9]	StyleDiff [19]	InST [51]	QuanArt [17]	CAST [50]	StyTr <sup>2</sup> [13]	IEST [8]	AdaAttN [25]	ArtFlow [2]
content	20.6%	16.4%	18.6%	35.5%	23.6%	45.0%	45.9%	32.3%	37.7%
style	30.7%	25.9%	32.3%	28.6%	56.4%	51.4%	40.9%	40.5%	23.6%
overall	29.5%	16.4%	15.9%	35.9%	29.5%	35.5%	36.4%	35.0%	22.7%

public preferences. A total of 55 participants were randomly selected for the study. We provided them with 10 content images and 10 style images and generated 100 stylized results using both our method and comparison methods. Each participant was shown 32 groups of questions, where they were presented with a random content/style image and its corresponding stylized result from both our method and a random comparison method. The participants were then asked to answer three questions: 1) which stylization result better preserves the content, 2) which stylization result

exhibits better style patterns and 3) which stylization result has a better overall effect. In total, we collected 1,760 votes. The vote counts are summarized in Table 1. We can observe that  $Z^*$  beats all contrast methods on average in terms of content preservation. Only CAST and StyTr<sup>2</sup> outperform  $Z^*$  regarding style representation, while the difference is insignificant. In terms of the overall effect,  $Z^*$  clearly outperforms the comparative methods; showing that  $Z^*$  achieves a balance between style and content, leading to satisfactory results.




Figure 7. Injecting attention module in different denoising steps.




Figure 8. Injecting attention module in different U-Net layers.



Figure 9. Results of using different attention arrangement.

### 5.3. Ablation Study

**Influence of the attention injection step and layers.** The impact of the attention injection step and layers on stylization outcomes is analyzed. As depicted in Figure 7, initiating the denoising step too early results in the loss of content structure information, whereas increasing the start denoising step enhances content preservation but sacrifices style patterns (as observed in the first three columns of results). On the other hand, increasing the overall denoising steps makes style patterns more prominent without compromising content structure (as evident in the comparison between the second and fourth columns of results). Consequently, the denoising process is initiated at the 5<sup>th</sup> step and concluded at the 30<sup>th</sup> step to achieve the most optimal stylized results.

In Figure 8, the stylized results obtained by injecting the attention module at different layers of the U-Net architecture are presented. Utilizing all layers in U-Net adversely affects content structure in the results. The low-resolution layers in U-Net effectively preserve content structure but transfer fewer style patterns (layers 5-10). Conversely, the high-resolution layers in U-Net can extract style features (layers 0-5 and layers 10-15), but encoder layers (layers 0-5) damage content structure. Hence, the attention module is injected into the high-resolution layers of the decoder.

**Influence of the cross-attention reweighting.** To evaluate the efficacy of the cross-attention reweighting, we provide a visual representation of the stylized outcome in Figure 9. We conduct a comparative analysis by considering two ablation scenarios: 1) the complete removal of content,



Figure 10. The effect of  $\lambda$  in Eq. (12).

as described in Eq. (7), and 2) the utilization of only the summation operation, denoted by Eq. (8). The coefficient  $\lambda$  in Eq. (8) is assigned a value of 0.5. As illustrated in Figure 9, eliminating the content component leads to compromised preservation of content structures within the stylized outputs. Similarly, incorporating solely the Query feature of the content image results in the loss of content details and insufficient style patterns. In contrast, the cross-attention reweighting demonstrates a superior ability to strike a balance between content and style in the stylized results.

**Influence of the style scaling factor.** We present a visualization of the impact of various values of  $\lambda$  in Eq. (12). The results, depicted in Figure 10, demonstrate the robustness of our method to different choices of  $\lambda$ . Specifically, our approach exhibits satisfactory performance when  $\lambda$  is set to values greater than or equal to 1.2. Based on these findings, we adopt a fixed value of  $\lambda = 1.2$  for all experiments conducted in this study.

## 6. Conclusion

This paper introduces a novel zero-shot style transfer approach that leverages ample prior knowledge within a pre-trained diffusion model. By incorporating a Key and Value features attention layer, we modify the self-attention mechanism in the diffusion model, enabling the use of Query features to retrieve style-related information from the Key and Value features. To enhance the preservation of content structures in stylized outputs, we propose a cross-attention reweighting technique that incorporates additional content information and achieves a more favorable balance between content preservation and style rendering. Extensive experimental evaluations demonstrate the superiority of our proposed method in terms of stylization results, outperforming existing state-of-the-art approaches.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under nos. 62102162, and U20B2070, in part by the Beijing Science and Technology Plan Project under no. Z231100005923033, and in part by Beijing Natural Science Foundation under no. L221013.



## References

- [1] Jie An, Haoyi Xiong, Jun Huan, and Jiebo Luo. Ultrafast photorealistic style transfer via neural architecture search. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 10443–10450, 2020. [2](#)
- [2] Jie An, Siyu Huang, Yibing Song, Dejing Dou, Wei Liu, and Jiebo Luo. ArtFlow: Unbiased image style transfer via reversible neural flows. In *IEEE/CVF Conferences on Computer Vision and Pattern Recognition (CVPR)*, pages 862–871, 2021. [1](#), [2](#), [6](#), [7](#)
- [3] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18208–18218, 2022. [2](#)
- [4] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [3](#)
- [5] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. [2](#)
- [6] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. *CoRR*, abs/2304.08465, 2023. [2](#), [3](#)
- [7] Haibo Chen, Zhizhong Wang, Huiming Zhang, Zhiwen Zuo, Ailin Li, Wei Xing, Dongming Lu, et al. Artistic style transfer with internal-external learning and contrastive learning. *Advances in Neural Information Processing Systems*, 34:26561–26573, 2021. [2](#)
- [8] Haibo Chen, Zhizhong Wang, Huiming Zhang, Zhiwen Zuo, Ailin Li, Wei Xing, Dongming Lu, et al. Artistic style transfer with internal-external learning and contrastive learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [6](#), [7](#)
- [9] Bin Cheng, Zuhao Liu, Yunbo Peng, and Yue Lin. General image-to-image translation with one-shot image guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22736–22746, 2023. [1](#), [2](#), [3](#), [6](#), [7](#)
- [10] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*, 2022. [2](#)
- [11] Yingying Deng, Fan Tang, Weiming Dong, Wen Sun, Feiyue Huang, and Changsheng Xu. Arbitrary style transfer via multi-adaptation network. In *ACM International Conference on Multimedia*, pages 2719–2727, 2020. [1](#), [2](#)
- [12] Yingying Deng, Fan Tang, Weiming Dong, Haibin Huang, Chongyang Ma, and Changsheng Xu. Arbitrary video style transfer via multi-channel correlation. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 1210–1217, 2021. [2](#)
- [13] Yingying Deng, Fan Tang, Weiming Dong, Chongyang Ma, Xingjia Pan, Lei Wang, and Changsheng Xu. Stytr2: Image style transfer with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11326–11336, 2022. [1](#), [2](#), [6](#), [7](#)
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021. [3](#)
- [15] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2414–2423, 2016. [1](#), [2](#)
- [16] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross-attention control. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. [2](#)
- [17] Siyu Huang, Jie An, Donglai Wei, Jiebo Luo, and Hanspeter Pfister. Quantart: Quantizing image style transfer towards high visual fidelity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2023. [6](#), [7](#)
- [18] Xun Huang and Belongie Serge. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1501–1510, 2017. [1](#), [2](#)
- [19] Jaeseok Jeong, Mingi Kwon, and Youngjung Uh. Training-free style transfer emerges from h-space in diffusion models. *arXiv preprint arXiv:2303.15403*, 2023. [7](#)
- [20] Yongcheng Jing, Yining Mao, Yiding Yang, Yibing Zhan, Mingli Song, Xinchao Wang, and Dacheng Tao. Learning graph neural networks for image style transfer. In *European Conference on Computer Vision*, 2022. [2](#)
- [21] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, pages 694–711. Springer, 2016. [1](#)
- [22] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Magic: Text-based real image editing with diffusion models. *CoRR*, abs/2210.09276, 2022. [2](#)
- [23] Chuan Li and Michael Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. In *European Conference on Computer Vision (ECCV)*, pages 702–716, 2016. [2](#)
- [24] Senmao Li, Joost van de Weijer, Taihang Hu, Fahad Shahbaz Khan, Qibin Hou, Yaxing Wang, and Jian Yang. Stylediffusion: Prompt-embedding inversion for text-based editing. *arXiv preprint arXiv:2303.15649*, 2023. [2](#)
- [25] Songhua Liu, Tianwei Lin, Dongliang He, Fu Li, Meiling Wang, Xin Li, Zhengxing Sun, Qian Li, and Errui Ding. Adaattn: Revisit attention mechanism in arbitrary neural style transfer. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 6649–6658, 2021. [6](#), [7](#)
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [3](#)

- [27] Ming Lu, Hao Zhao, Anbang Yao, Yurong Chen, Feng Xu, and Li Zhang. A closed-form solution to universal style transfer. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5952–5961, 2019. 2
- [28] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. *CoRR*, abs/2211.09794, 2022. 2
- [29] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 2
- [30] Zhihong Pan, Riccardo Gherardi, Xiufeng Xie, and Stephen Huang. Effective real image editing with accelerated iterative diffusion inversion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15912–15921, 2023. 2
- [31] Dae Young Park and Kwang Hee Lee. Arbitrary style transfer with style-attentional networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5880–5888, 2019. 3
- [32] Dae Young Park and Kwang Hee Lee. Arbitrary style transfer with style-attentional networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5880–5888, 2019. 1
- [33] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *CoRR*, abs/2204.06125, 2022. 2
- [34] Eric Risser, Pierre Wilmot, and Connelly Barnes. Stable and controllable neural texture synthesis and style transfer using histogram losses. *arXiv preprint arXiv:1701.08893*, 2017. 2
- [35] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 10674–10685. IEEE, 2022. 6
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 2
- [37] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022. 2
- [38] Lu Sheng, Ziyi Lin, Jing Shao, and Xiaogang Wang. Avatar-net: Multi-scale zero-shot style transfer by feature decoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8242–8250. IEEE, 2018. 1, 2
- [39] Hao Tang, Songhua Liu, Tianwei Lin, Shaoli Huang, Fu Li, Dongliang He, and Xinchao Wang. Master: Meta style transformer for controllable zero-shot and few-shot artistic style transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18329–18338, 2023. 2
- [40] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. *CoRR*, abs/2211.12572, 2022. 2, 3
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017. 3
- [42] Huan Wang, Yijun Li, Yuehai Wang, Haoji Hu, and Ming-Hsuan Yang. Collaborative distillation for ultra-resolution universal style transfer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1860–1869, 2020. 2
- [43] Jianbo Wang, Huan Yang, Jianlong Fu, Toshihiko Yamasaki, and Baining Guo. Fine-grained image style transfer with visual transformers. In *Proceedings of the Asian Conference on Computer Vision*, pages 841–857, 2022. 2
- [44] Hua-Peng Wei, Ying-Ying Deng, Fan Tang, Xing-Jia Pan, and Wei-Ming Dong. A comparative study of cnn-and transformer-based visual style transfer. *Journal of Computer Science and Technology*, 37(3):601–614, 2022. 2
- [45] Chen Henry Wu and Fernando De la Torre. A latent space of stochastic diffusion models for zero-shot image editing and guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7378–7387, 2023. 2
- [46] Xiaolei Wu, Zhihao Hu, Lu Sheng, and Dong Xu. Style-former: Real-time arbitrary style transfer via parametric style composition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14618–14627, 2021. 1, 2
- [47] Yifan Xu, Huapeng Wei, Minxuan Lin, Yingying Deng, Kekai Sheng, Mengdan Zhang, Fan Tang, Weiming Dong, Feiyue Huang, and Changsheng Xu. Transformers in computational visual media: A survey. *Computational Visual Media*, 8(1): 33–62, 2022. 2
- [48] Serin Yang, Hyunmin Hwang, and Jong Chul Ye. Zero-shot contrastive loss for text-guided diffusion image style transfer. *arXiv preprint arXiv:2303.08622*, 2023. 2
- [49] Chiyu Zhang, Jun Yang, Lei Wang, and Zaiyan Dai. S2wat: Image style transfer via hierarchical vision transformer using strips window attention. *arXiv preprint arXiv:2210.12381*, 2022. 2
- [50] Yuxin Zhang, Fan Tang, Weiming Dong, Haibin Huang, Chongyang Ma, Tong-Yee Lee, and Changsheng Xu. Domain enhanced arbitrary image style transfer via contrastive learning. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–8, 2022. 1, 2, 6, 7
- [51] Yuxin Zhang, Nisha Huang, Fan Tang, Haibin Huang, Chongyang Ma, Weiming Dong, and Changsheng Xu. Inversion-based style transfer with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10146–10156, 2023. 1, 2, 3, 6, 7

- [52] Zhixing Zhang, Ligong Han, Arnab Ghosh, Dimitris N Metaxas, and Jian Ren. Sine: Single image editing with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6027–6037, 2023. [2](#)