

Towards Highly Realistic Artistic Style Transfer via Stable Diffusion with Step-aware and Layer-aware Prompt

Zhanjie Zhang¹, Quanwei Zhang¹, Huaizhong Lin^{1†}, Wei Xing^{1†}, Juncheng Mo¹, Shuaicheng Huang², Jinheng Xie³, Guangyuan Li¹, Junsheng Luan¹, Lei Zhao^{1†}, Dalong Zhang¹ and Lixia Chen^{4†}

¹ Zhejiang University

² Anhui University

³ National University of Singapore

⁴ Zhejiang Gongshang University

{cszsj, cszqw, linhz, wxing, csmjc, csly, l.junsheng121, cszh}@zju.edu.cn,
Z02114187@stu.ahu.edu.cn, xiejinheng2020@email.szu.edu.cn, zdlxing@126.com, clixia@163.com

Abstract

Artistic style transfer aims to transfer the learned artistic style onto an arbitrary content image, generating artistic stylized images. Existing generative adversarial network-based methods fail to generate highly realistic stylized images and always introduce obvious artifacts and disharmonious patterns. Recently, large-scale pre-trained diffusion models opened up a new way for generating highly realistic artistic stylized images. However, diffusion model-based methods generally fail to preserve the content structure of input content images well, introducing some undesired content structure and style patterns. To address the above problems, we propose a novel pre-trained diffusion-based artistic style transfer method, called LSAST, which can generate highly realistic artistic stylized images while preserving the content structure of input content images well, without bringing obvious artifacts and disharmonious style patterns. Specifically, we introduce a Step-aware and Layer-aware Prompt Space, a set of learnable prompts, which can learn the style information from the collection of artworks and dynamically adjusts the input images' content structure and style pattern. To train our prompt space, we propose a novel inversion method, called Step-aware and Layer-aware Prompt Inversion, which allows the prompt space to learn the style information of the artworks collection. In addition, we inject a pre-trained conditional branch of ControlNet into our LSAST, which further improved our framework's ability to maintain content structure. Extensive experiments demonstrate that our proposed method can generate more highly realistic artistic stylized images than the state-of-the-art artistic style transfer methods. Code is available at <https://github.com/Jamie-Cheung/LSAST>.

[†] Corresponding authors.

1 Introduction

Artistic style transfer has recently attracted widespread attention in academia and industry since the seminal work of CycleGAN [Zhu *et al.*, 2017]. Existing artistic style transfer methods can be divided into generative adversarial network-based approaches (GAN-based approaches) and large-scale pre-trained diffusion model-based approaches (Diffusion-based approaches).

More specifically, GAN-based methods [Zhu *et al.*, 2017; Sanakoyeu *et al.*, 2018; Kim *et al.*, 2019; Park *et al.*, 2020; Chen *et al.*, 2023; Zhang *et al.*, 2023d; Li *et al.*, 2023c] generally utilize generative adversarial network and a training set of aligned/unaligned image pairs to learn the mapping between an input image and an output image. For example, Zhu *et al.* [Zhu *et al.*, 2017] used two mirror generative adversarial network to learn and improve the mapping between the input image and output image, synthesizing artistic stylized images. Sanakoyeu *et al.* [Sanakoyeu *et al.*, 2018] proposed a style-aware content loss, to improve the quality of artistic stylized images by capturing how style patterns affect content structure. However, GAN-based methods are limited by the instability of adversarial training and the scarcity of training data, failing to generate highly realistic artistic stylized images and introducing the obvious artifacts and disharmonious patterns on the stylized images (Please see in Fig. 1 (e-h)).

Large-scale pre-trained diffusion model-based approaches [Nichol *et al.*, 2021; Wu, 2022; Ho *et al.*, 2020; Hu *et al.*, 2023; Li *et al.*, 2024] use massive parameters to learn and store the information from the large-scale training data, possessing the ability to generate highly realistic images. This opens up a new way for generating highly realistic artistic stylized images. For example, Zhang *et al.* [Zhang *et al.*, 2024b] proposed a global prompt space, a learnable parameter matrix, to learn and store the style information from the collection of artworks and condition pre-trained large-scale diffusion model to generate artistic stylized image. Zhang *et al.* [Zhang *et al.*, 2023b] introduced a step-aware prompt space, a set of learnable parameter matrixes, for the whole diffusion process to generate desired stylized image. Although these approaches could generate

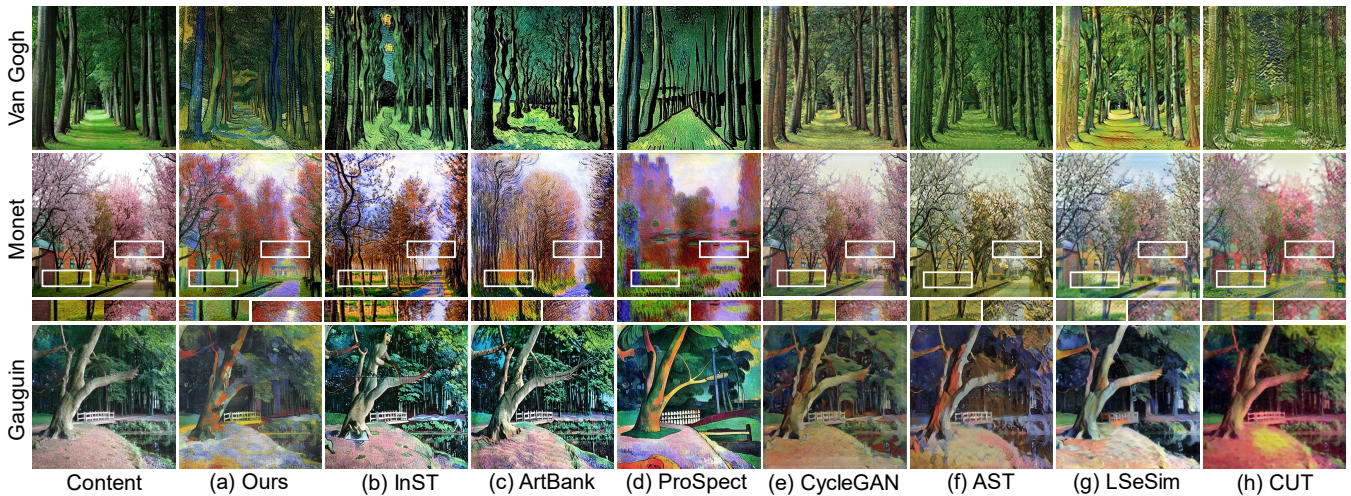


Figure 1: Stylization examples with three different styles (i.e., Van Gogh, Monet, Gauguin). Compared to existing state-of-the-art large-scale pre-trained diffusion model-based methods (b-d) and generative adversarial network-based methods (e-h), our proposed method (a) generates highly realistic artistic stylized images and preserves the content structure of input content images well.

highly realistic stylized images, they failed to preserve the content structure of input content image well, bringing some undesired content structure and style patterns (Please see in Fig. 1 (b-d)).

Motivated by the above observation, we propose a novel framework, called LSAST, which can dig out prior knowledge from the large-scale pre-trained diffusion model to generate highly realistic artistic stylized images while preserving the content structure of the input content image well. Specifically, we propose a Step-aware and Layer-aware Prompt Space P^* (i.e., a set of learnable prompts) to control the whole diffusion process from both step and layer dimensions, dynamically adjusting the input images' content structure and style pattern. In detail, to make P^* layer-aware, we divide U-Net into three layers and design different prompts for different layers to control content structure and style information at different scales. To make P^* step-aware, we divide the entire de-noising phase into 10 stages, each containing 100 steps, and each stage corresponds to a prompt, which dynamically adjusts the content structure and style information based on the de-noising steps. To train P^* , we propose a novel inversion method, called Step-aware and Layer-aware Prompt Inversion, which allows the prompt P^* to learn the style information of the artworks collection. Finally, we inject a pre-trained conditional branch of ControlNet [Zhang *et al.*, 2023a] into our LSAST, which further improved our framework's ability to maintain content structure.

To summarize, the main contribution of this paper is as follows:

- We propose a novel pre-trained diffusion-based artistic style transfer framework, which can generate highly realistic stylized images and preserve the content structure of input content images well without introducing obvious artifacts and disharmonious style patterns.
- We design a novel Step-aware and Layer-aware Prompt Space and conduct a Step-aware and Layer-aware

Prompt Inversion to train prompt space, which can learn the style information from the collection of artworks and dynamically adjusts the input images' content structure and style pattern.

- Extensive quantitative and qualitative experiments demonstrate that our proposed LSAST outperforms the state-of-the-art GAN-based and Diffusion-based methods.

2 Related Work

Generative Adversarial Network-based Methods. Generative adversarial network-based methods refer to the use of discriminator networks [Zhu *et al.*, 2017] to train a well-designed forward network that can bridge a mapping between the input image and output image. As the seminal work of artistic style transfer, Zhu *et al.* [Zhu *et al.*, 2017] expanded adversarial loss and proposed Cycle-Consistent loss to improve the quality of artistic stylized images. The cycle-consistent loss inspired a lot of researchers to explore a more effective way to enhance further the quality of artistic stylized images, including [Wang *et al.*, 2022; Park *et al.*, 2020; Zheng *et al.*, 2021; Fu *et al.*, 2019; Zuo *et al.*, 2023; Zhang *et al.*, 2023d; Zhang *et al.*, 2021]. For example, Park *et al.* [Park *et al.*, 2020] proposed a method to preserve the content structure of input content image by maximizing the mutual information between input image and output image via contrastive learning [Zhang *et al.*, 2024a]. Sanakoyeu *et al.* [Sanakoyeu *et al.*, 2018] proposed a style-aware content loss to improve the stylization of images by capturing how style affects content. Zheng *et al.* [Zheng *et al.*, 2021] exploited the spatial patterns of self-similarity to capture spatial relationships within an image rather than domain appearance to preserve the content structure. While generative adversarial network-based methods effectively transfer the learned style onto an arbitrary content image, they fail to generate highly realistic artistic stylized images, introducing obvious

artifacts and disharmonious patterns.

Large-scale Pre-trained Diffusion Model-based Approaches. Large-scale pre-trained diffusion model-based approaches use massive parameters to learn information from large-scale data. To dig out the prior knowledge from the pre-trained diffusion model, Zhang *et al.* [Zhang *et al.*, 2023c] proposed to learn a word embedding from a single style image and then guide the whole diffusion process to transfer learned style information onto a content image. Kim *et al.* [Kim *et al.*, 2022] performed text-driven image manipulation to preserve the content structure of input content image and learn style information from style image. Zhang *et al.* [Zhang *et al.*, 2024b] proposed a global prompt condition (i.e., a learnable parameter matrix) to learn and store the style information and condition pre-trained large-scale diffusion model to generate artistic stylized image. Zhang *et al.* [Zhang *et al.*, 2023b] introduced a step-aware prompt condition for the whole diffusion process to generate the stylized image. Xie *et al.* [Xie *et al.*, 2023] proposed to control object synthesis in the given spatial conditions. While the large-scale pre-trained model-based approaches can generate highly realistic stylized images, they always fail to preserve the content structure of the input content image well, bringing some undesired content structure and style patterns.

3 Proposed Method

Let I_s and I_c be the style image and content image respectively; our goal is to train a Step-aware and Layer-aware Prompt Space, a set of learnable prompts, to learn the style information from the style images and dig out the abundant prior knowledge from large-scale pre-trained diffusion model to transfer the learned style onto the content image I_c , synthesizing highly realistic stylized images I_{cs} . The pipeline of our proposed LSAST is shown in Fig. 2, which consists of two stages: Learning step-aware and layer-aware prompt space from the collection of artworks (training) and generating highly realistic artistic stylized images (inference).

In the training stage, we utilize a learnable parameter matrix to learn and store the style information from the collection of artworks. Then, we expand it into a set of learnable parameter matrices (Step-aware and Layer-aware Prompts Space). These learnable parameter matrixes dig out the prior knowledge of the pre-trained diffusion model from step and layer dimensions.

In the inference stage, the Step-aware and Layer-aware Prompts Space (i.e., a set of parameters trained in the training phase) are used to guide the pre-trained diffusion model to transfer the learned style onto the content images, generating highly realistic artistic stylized images. Besides, we use the conditional branch of ControlNet [Zhang *et al.*, 2023a] to extract the content structure of input content image as content prompt and guide pre-trained stable diffusion to preserve the content structure.

3.1 Step-aware and Layer-aware Prompt Space

Our goal is to dig out the abundant prior knowledge from the large-scale pre-trained Stable Diffusion [Rombach *et al.*, 2022], which can generate highly realistic images. To dig

out the prior knowledge from it, Gal *et al.* [Gal *et al.*, 2022] proposed to optimize the token prompt embedding space of CLIP [Radford *et al.*, 2021], outputting a suitable global prompt for the diffusion model. However, the denoising process within Stable Diffusion generally takes 1,000 steps. In the denoising process’s early, middle, and late stages [Agarwal *et al.*, 2023], the text prompt controls the coarse content structure, detail content structure, and detail style patterns, respectively. Only using a global prompt embedding has difficulty in controlling content structure and style patterns simultaneously. Zhang *et al.* [Zhang *et al.*, 2023b] divided the 1,000 denoising steps into ten stages on average, and each stage corresponds to a unique prompt. However, the typical U-Net model, within stable diffusion, can be divided into three layers [Agarwal *et al.*, 2023]: coarse, moderate, and fine, corresponding to coarse, detailed, and style patterns. Thus, only considering dividing the steps into ten stages is still not enough. To this end, we introduce a Step-aware and Layer-aware Prompt Space, which divides the 1,000 denoising steps into ten stages and the U-Net into three layers. Thus, to condition the pre-trained diffusion model from step and layer dimensions, the number of the learnable parameter matrix should be 30. To this end, we expand a learnable parameter matrix P (i.e., $P \in R^{1 \times 1 \times 768}$) into Step-aware and Layer-aware Prompt Space P^* (i.e., $P^* \in R^{30 \times 1 \times 768}$) with attention [Li *et al.*, 2023a; Cui *et al.*, 2022; Li *et al.*, 2023b; Li *et al.*, 2023c]. Specifically, we formulate Q (query), K (key) and V (value) as:

$$\begin{aligned} Q &= f(\text{Norm}(P)), \\ K &= g(\text{Norm}(P)), \\ V &= h(P), \end{aligned} \quad (1)$$

where f , g , and h are 1×1 learnable convolution layers to expand the dimension of P from $1 \times 1 \times 768$ to $30 \times 1 \times 768$. Norm here denotes channel-wise mean-variance normalization. Then, we calculate attention map A as:

$$A = \text{Softmax}(Q^\top \otimes K), \quad (2)$$

where \otimes denotes the matrix multiplication.

$$\hat{P} = (V \otimes A^\top), \quad (3)$$

where $\hat{P} = [p_1, p_2, \dots, p_n]$ (i.e., $n = 30$). Then, these textual conditions are grouped into ten groups. We refer to grouped textual conditions as Step-aware and Layer-aware Prompt Space P^* .

$$P^* = [[p_1^1, p_1^2, p_1^3], [p_2^1, p_2^2, p_2^3], \dots, [p_n^1, p_n^2, p_n^3]] (n = 10), \quad (4)$$

where P_i^j denotes the prompt corresponding to i -th stage of denoising process and j -th layer of U-Net. We define the above process as $P^* = \text{Self-atten}(P)$.

3.2 Step-aware and Layer-aware Prompt Inversion

To dig out knowledge from a large-scale pre-trained diffusion model, a naive method is to unfrozen the whole parameter to learn and store the style information from the collection of artworks. Specifically, we first add noise $\epsilon \sim \mathcal{N}(0, I)$

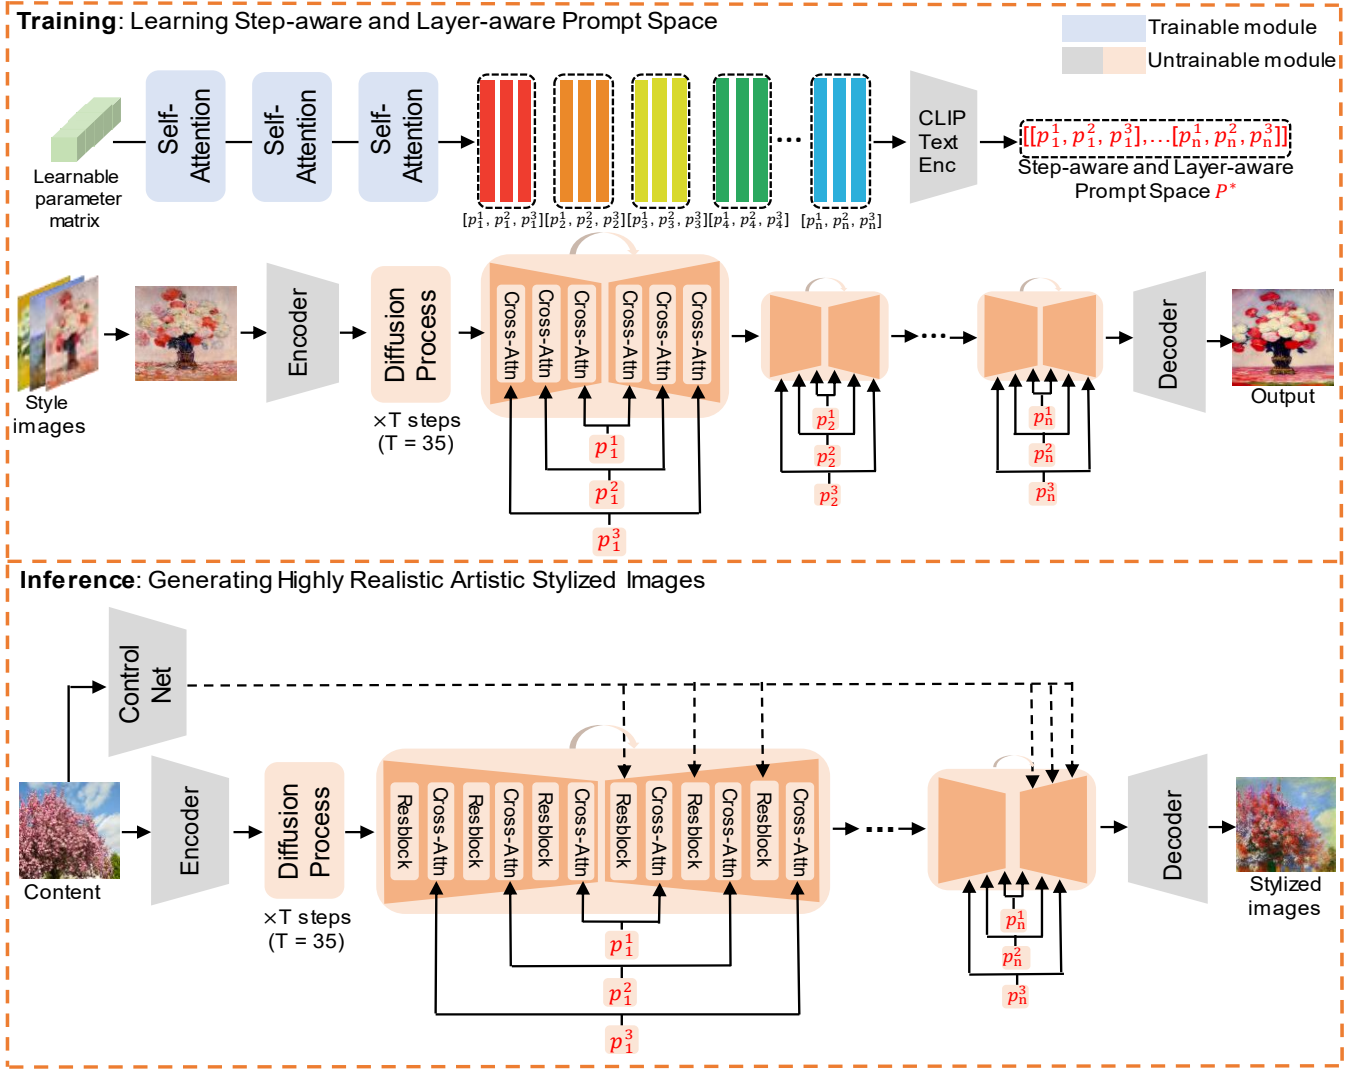


Figure 2: The overview of our proposed framework which consists of a training stage and an inference stage. In the training stage, the Step-aware and Layer-aware Prompts Spaces will learn and store the style information from the collection of artworks. In the inference stage, we utilize a pre-trained conditional branch of ControlNet, as content prompt, into our LSAST, which further improved our framework’s ability to maintain content structure.

onto style images, obtaining the noisy style image $z_{s,t}$. Then, feeding the $z_{s,t}$ into Stable Diffusion and using the following loss:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{z,t} \left[\|\epsilon - \epsilon_{\theta}(z_{s,t}, t)\|_2^2 \right], \quad (5)$$

Once the $\mathcal{L}_{\text{diff}}$ converges, the fine-tuned Stable Diffusion will possess the knowledge to generate highly realistic artistic stylized images. However, such a naive way costs extra computation burden. To this end, we propose Step-aware and Layer-aware Prompt Inversion to learn and store the knowledge from the collection of artworks and dig out the abundant prior knowledge from the large-scale pre-trained diffusion model. In detail, we feed the $z_{s,t}$ into pre-trained Stable Diffusion (version 1.5), and train the Step-aware and Layer-aware Prompt Space P^* using the following loss:

$$\mathcal{L}_p = \mathbb{E}_{z,p,t} \left[\|\epsilon - \epsilon_{\theta}(z_{s,t}, P^*, t)\|_2^2 \right], \quad (6)$$

Once the \mathcal{L}_p converges, the P^* will store the style information from the collection of artworks and can be used to condition pre-trained stable diffusion to generate the highly realistic artistic stylized image.

3.3 Content Prompt

ControlNet [Zhang *et al.*, 2023a] is widely used to condition large-scale pre-trained diffusion model to preserve the content structure of input content image. We utilize an additional branch (Canny edge) of the pre-trained ControlNet as an additional content prompt and inject content prompt into pre-trained stable diffusion to further improve the ability to maintain the content structure of the input image

Metrics	Datasets	Ours	InST	ArtBank	ProSpect	CycleGAN	AST	LSeSim	CUT
FID↓	Van Gogh	94.17	113.26	121.62	118.22	112.04	96.10	105.21	95.61
	Morisot	176.70	221.37	225.47	188.11	204.74	199.78	230.37	201.41
	Ukiyoe	92.03	104.73	111.50	93.15	140.78	148.17	130.77	115.44
	Monet	130.25	159.39	166.70	133.73	170.50	123.17	162.25	150.08
	Cezanne	130.28	150.24	143.90	139.51	151.14	133.47	141.33	141.34
	Gauguin	123.40	145.65	155.27	125.03	160.00	164.73	172.18	165.69
	Peploe	158.44	188.37	182.51	186.45	167.77	222.30	192.88	164.39
Time/sec ↓	-	4.1325	4.0485	3.7547	3.925	0.0312	0.0312	0.0365	0.0312
Preference↑	-	0.634*	0.581 /0.419	0.532 /0.468	0.545 /0.455	0.692 /0.308	0.674 /0.326	0.689 /0.311	0.726 /0.274
Deception↑	-	0.752	0.716	0.614	0.605	0.582	0.579	0.547	0.508

Table 1: Quantitative comparisons. The lower the FID score, the better the image quality. * denotes the average user preference.

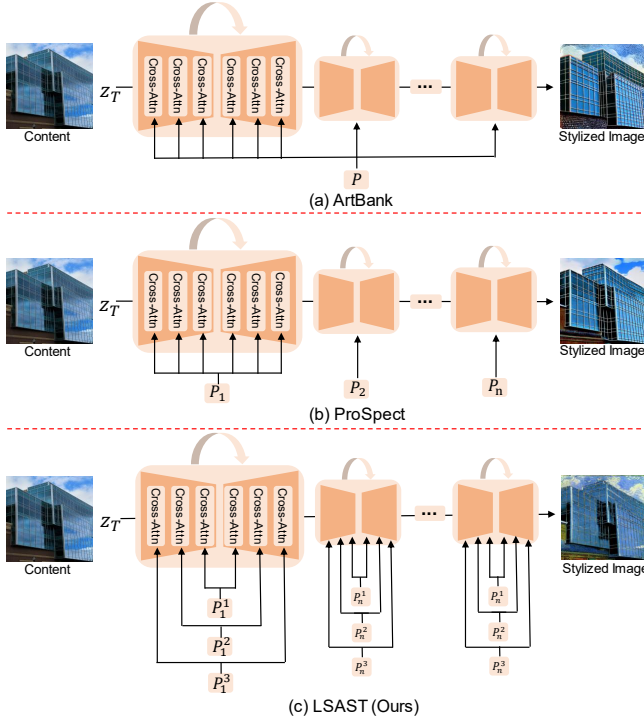


Figure 3: Differences between (a) ArtBank. (b) ProSpect. (c) LSAST (Ours). Unlike existing methods that use only a global/step-aware prompt to condition the whole diffusion process, LSAST utilizes a step-aware and layer-aware prompt to dynamically adjust the input images’ content structure and style pattern.

4 Experiments

4.1 Implementation Details

We adopt the large-scale pre-trained stable diffusion model (version 1.5) as our backbone due to its strong ability to generate highly realistic images. We implement our LSAST on two NVIDIA RTX 3090 GPUs via Pytorch [Paszke *et al.*, 2019]. The learning rate is initially set to 0.000005 and then decreased by a factor of 10 after 20,000 iterations (100,000 iterations in total). In training, we collect seven collections of artworks (e.g., Van Gogh, Cezanne, Peploe, Ukiyoe, Monet, Morisot, and Gauguin) from AST [Sanakoyeu *et al.*, 2018] as artistic style images to train our LSAST. Before feeding

artistic style images into LSAST, all the style images will be resized to 512×512 pixels. In inference, we randomly select some content images from DIV2K [Agustsson and Timofte, 2017] and resize them to 512×512 as the initial input content image.

4.2 Qualitative Comparison

We compare our method with the SOTA GAN-based and Diffusion-based artistic style transfer methods as shown in Fig. 4, including InST [Zhang *et al.*, 2023c], Art-Bank [Zhang *et al.*, 2024b], ProSpect [Zhang *et al.*, 2023b], AST [Sanakoyeu *et al.*, 2018], CycleGAN [Zhu *et al.*, 2017], LSeSim [Zheng *et al.*, 2021] and CUT [Park *et al.*, 2020]. As the representative of GAN-based approaches, CycleGAN generally introduces some obvious artifacts into the stylized images (e.g., 4th and 5th rows). AST sometimes shows blurry content structure and abrupt colors (e.g., 5th and 7th rows). LSeSim introduces some obvious artifacts and disharmonious patterns (e.g., 5th and 6th rows). CUT always introduces some unwanted content structure and style patterns (e.g., 5th and 7th rows). As the representative of Diffusion-based model, the stylized image of InST fails to preserve the content structure of input content image, even deviating from the original content image’s semantic (e.g., 6th rows). Art-Bank generally fails to capture some detailed style patterns (e.g., 6th and 7th rows). ProSpect sometimes has limitations in preserving the content structure of input content image (e.g., 4th rows). In comparison, our proposed LSAST not only can generate highly realistic artistic stylized images without introducing obvious artifacts and disharmonious patterns while preserving the content structure of the input content image well.

4.3 Quantitative Comparisons

Fréchet Inception Distance. Fréchet Inception Distance (FID) [Heusel *et al.*, 2017] utilizes Inception-V3 model [Szegedy *et al.*, 2016] to extract image features and calculate the distance between generated artistic stylized image’s distribution and style image’s distribution. The lower FID score means that the distribution of the generated artistic stylized images is closer to that of real artistic images. In other words, the generated artistic stylized image is prone to be regarded as a human-created artistic image. Take Van Gogh’s style, for example, we collect 10,000 content images and synthesize 10,000 stylized images for our LSAST and

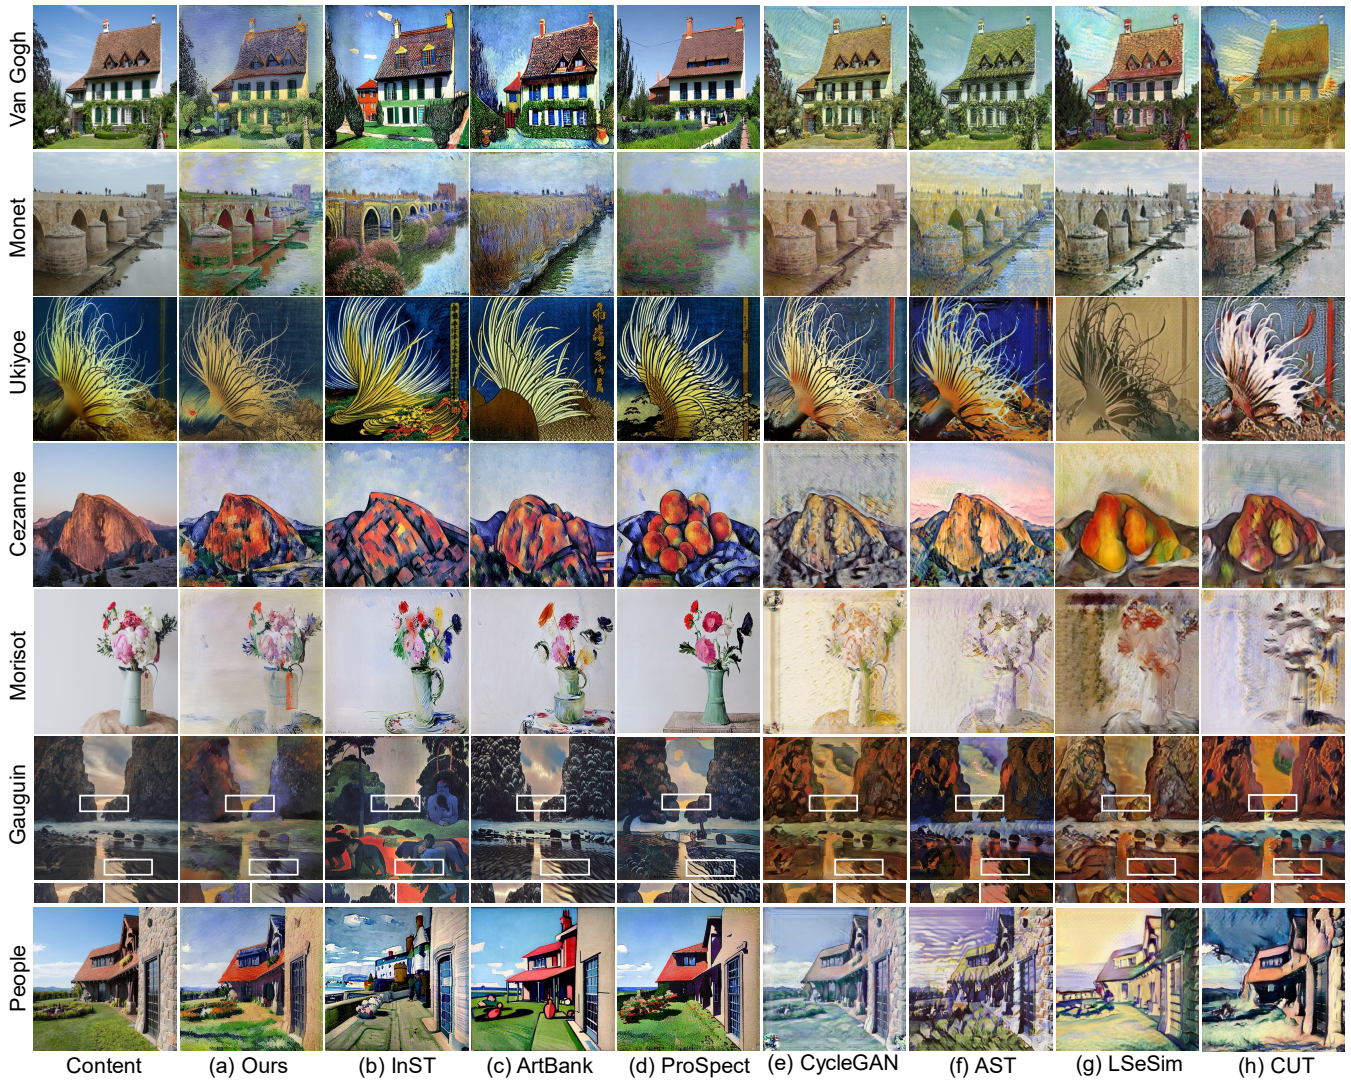


Figure 4: Qualitative comparisons with other state-of-the-art artistic style methods. The first column shows the input content image. The (b-d) column shows the stylized image from large-scale pre-trained diffusion model-based methods, and the (e-h) column presents the stylized images generated by generative adversarial network-based methods.

other SOTA artistic style transfer methods to calculate the FID score. As shown in Table. 1, we proposed LSAST to achieve a better FID score, which means our proposed method can generate highly realistic artistic stylized images closer to human-created artistic images.

Preference Score. The preference score effectively conducted an A/B Test user study to compare the stylization effects between our method and the other SOTA method, reflecting the popularity of our method compared to other state-of-the-art methods. We randomly choose 100 content images to synthesize 100 stylized images for each method. Then, the content image and corresponding two stylized image pairs generated by our method and a randomly selected SOTA are displayed. The subjects need to choose their preferred stylization effects. We collected 5000 votes from 50 subjects and showed the preference score in Table. 1 (“Preference”

row). The results indicate that our LSAST achieves the best popularity.

Deception Score. The deception score is an effective means for evaluating whether artistic stylized images are prone to be regarded as human-created from people’s perspective. A higher deception score means the artistic stylized image is more prone to be considered a human-created artistic image. To compute the deception score, we first randomly choose 20 content images. Then, we feed these content images into our proposed LSAST and other SOTA artistic style transfer methods, generating stylized images. We asked 50 subjects to distinguish whether the generated image was human-created. The result is shown in Tab. 1. A higher deception score means the stylized images are more likely to be considered as human-created artistic images. Besides, we report the deception score of the human-created artistic style image, cho-




Figure 5: Ablation studies of our LSAST. Please zoom-in for better comparison.

Metric	Datasets	Full Model	w/o content	w/o step&layer	w/o step-aware	w/o layer-aware
FID↓	Van Gogh	94.17	95.05	96.33	94.46	94.58
	Morisot	176.70	177.51	178.29	177.54	177.61
	Ukiyoe	92.03	93.85	94.97	93.26	93.17
	Monet	130.25	131.78	133.34	130.64	130.72
	Cezanne	130.28	131.07	132.84	131.16	131.78
	Gauguin	143.40	144.02	145.90	144.27	144.51
	Peploe	158.44	159.15	160.32	159.49	159.71

Table 2: The quantitative ablation studies of LSAST. The lower the FID score, the better the image quality.

sen randomly from WikiArt [Phillips and Mackintosh, 2011] is 0.863. From Tab. 1, the deception of our method is closer to 0.863, which indicates that our LSAST achieves a better deception score.

Time Information. The “Time/sec” row of Tab. 1 shows the inference time [Sun *et al.*, 2024; Sun *et al.*, 2023; Yang *et al.*, 2022] comparison on images with a scale of 512×512 pixels. The GAN-based method is faster than diffusion-based methods because diffusion-based methods commonly possess massive parameters. Our approach is also slower than other diffusion-based approaches due to the additional use of pre-trained content prompts.

4.4 Abalation Studies

To verify the effectiveness of our proposed Step-aware and Layer-aware Prompt Inversion, we split it into Step-aware Prompt Inversion and Layer-aware Prompt Inversion.

With and without Step-aware Prompt Inversion. As introduced in Section. 3.2, the Step-aware Prompt Inversion divides the 1,000 steps of denoising into three stages from step dimension: coarse content structure stage, detailed content structure stage, and detailed style patterns stage. Here, we train an LSAST model that does not involve step-aware prompt inversion. The experimental results are shown in Fig. 5. Compared with the full LSAST model, we can see that the LSAST model without a step-aware prompt introduces some obvious artifacts and disharmonious style patterns into the stylized image. The step-aware prompt inversion is vital to generating highly realistic stylized images.

With and without Layer-aware Prompt Inversion. The Layer-aware Prompt Inversion divides the architecture of U-Net into three layers from layer dimension: coarse content structure layer, detailed content structure layers, and detail style patterns layer. We show stylized images without layer-aware prompt inversion, as shown in Fig. 5. The stylized image shows some obvious artifacts and disharmonious style patterns.

With and without Step-aware and Layer-aware Prompt Inversion. We train an LSAST model without step-aware and layer-aware prompt inversion. As shown in Fig. 5, the stylized shown shows some unwanted content structure and style patterns.

Without content prompt. If removing the content prompt in inference, the stylized image fails to preserve the content structure of the input content image.

5 Conclusion

In this paper, we propose a novel framework, which can learn the style information from the collection of artworks and dynamically adjusts the input images’ content structure and style pattern, generating highly realistic stylized images (i.e., harmonious patterns and detailed texture) and preserving the content structure of the input content image well. Extensive quantitative and qualitative experiments verify that our proposed LSAST outperforms the SOTA GAN-based and Diffusion-based methods.

Acknowledgments

This work was supported in part by National Social Science Foundation Major Project “Research on Virtual Restoration of Tang and Song Painting Colors and Construction of Traditional Color Resource Library from a Digital Perspective” (19ZDA046), Zhejiang Province Program (2022C01222, 2023C03199, 2023C03201), the National Program of China (62172365, 2021YFF0900604, 19ZDA197), Ningbo Science and Technology Plan Project (2022Z167, 2023Z137), and MOE Frontier Science Center for Brain Science & Brain-Machine Integration (Zhejiang University).

Contribution Statement

Zhanjie Zhang and Quanwei Zhang contributed equally to this research.

References

- [Agarwal *et al.*, 2023] Aishwarya Agarwal, Srikrishna Karanam, Tripti Shukla, and Balaji Vasani Srinivasan. An image is worth multiple words: Multi-attribute inversion for constrained text-to-image synthesis. *arXiv preprint arXiv:2311.11919*, 2023.
- [Agustsson and Timofte, 2017] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 126–135, 2017.
- [Chen *et al.*, 2023] Jiafu Chen, Boyan Ji, Zhanjie Zhang, Tianyi Chu, Zhiwen Zuo, Lei Zhao, Wei Xing, and Dongming Lu. Testnerf: text-driven 3d style transfer via cross-modal learning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 5788–5796, 2023.
- [Cui *et al.*, 2022] Xuelian Cui, Zhanjie Zhang, Tao Zhang, Zhuoqun Yang, and Jie Yang. Attention graph: Learning effective visual features for large-scale image classification. *Journal of Algorithms & Computational Technology*, 16:17483026211065375, 2022.
- [Fu *et al.*, 2019] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, Kun Zhang, and Dacheng Tao. Geometry-consistent generative adversarial networks for one-sided unsupervised domain mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2427–2436, 2019.
- [Gal *et al.*, 2022] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022.
- [Heusel *et al.*, 2017] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637, 2017.
- [Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [Hu *et al.*, 2023] Teng Hu, Jiangning Zhang, Liang Liu, Ran Yi, Siqu Kou, Haokun Zhu, Xu Chen, Yabiao Wang, Chengjie Wang, and Lizhuang Ma. Phasic content fusing diffusion model with directional distribution consistency for few-shot model adaption. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2406–2415, 2023.
- [Kim *et al.*, 2019] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwanghee Lee. U-gat-it: unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. *arXiv preprint arXiv:1907.10830*, 2019.
- [Kim *et al.*, 2022] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2426–2435, 2022.
- [Li *et al.*, 2023a] Guangyuan Li, Wei Xing, Lei Zhao, Zehua Lan, Jiakai Sun, Zhanjie Zhang, Quanwei Zhang, Huaizhong Lin, and Zhijie Lin. Self-reference image super-resolution via pre-trained diffusion large model and window adjustable transformer. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 7981–7992, 2023.
- [Li *et al.*, 2023b] Guangyuan Li, Wei Xing, Lei Zhao, Zehua Lan, Zhanjie Zhang, Jiakai Sun, Haolin Yin, Huaizhong Lin, and Zhijie Lin. Dudoinet: Dual-domain implicit network for multi-modality mr image arbitrary-scale super-resolution. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 7335–7344, 2023.
- [Li *et al.*, 2023c] Guangyuan Li, Lei Zhao, Jiakai Sun, Zehua Lan, Zhanjie Zhang, Jiafu Chen, Zhijie Lin, Huaizhong Lin, and Wei Xing. Rethinking multi-contrast mri super-resolution: Rectangle-window cross-attention transformer and arbitrary-scale upsampling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21230–21240, 2023.
- [Li *et al.*, 2024] Guangyuan Li, Chen Rao, Juncheng Mo, Zhanjie Zhang, Wei Xing, and Lei Zhao. Rethinking diffusion model for multi-contrast mri super-resolution. *arXiv preprint arXiv:2404.04785*, 2024.
- [Nichol *et al.*, 2021] 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.
- [Park *et al.*, 2020] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16*, pages 319–345. Springer, 2020.

- [Paszke *et al.*, 2019] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [Phillips and Mackintosh, 2011] Fred Phillips and Brandy Mackintosh. Wiki art gallery, inc.: A case for critical thinking. *Issues in Accounting Education*, 26(3):593–608, 2011.
- [Radford *et al.*, 2021] 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*, pages 8748–8763. PMLR, 2021.
- [Rombach *et al.*, 2022] 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.
- [Sanakoyeu *et al.*, 2018] Artsiom Sanakoyeu, Dmytro Kotovenko, Sabine Lang, and Bjorn Ommer. A style-aware content loss for real-time hd style transfer. In *proceedings of the European conference on computer vision (ECCV)*, pages 698–714, 2018.
- [Sun *et al.*, 2023] Jiakai Sun, Zhanjie Zhang, Jiafu Chen, Guangyuan Li, Boyan Ji, Lei Zhao, and Wei Xing. Vgos: voxel grid optimization for view synthesis from sparse inputs. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 1414–1422, 2023.
- [Sun *et al.*, 2024] Jiakai Sun, Han Jiao, Guangyuan Li, Zhanjie Zhang, Lei Zhao, and Wei Xing. 3dgsstream: On-the-fly training of 3d gaussians for efficient streaming of photo-realistic free-viewpoint videos. *arXiv preprint arXiv:2403.01444*, 2024.
- [Szegedy *et al.*, 2016] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [Wang *et al.*, 2022] Zhizhong Wang, Zhanjie Zhang, Lei Zhao, Zhiwen Zuo, Ailin Li, Wei Xing, and Dongming Lu. Aesust: Towards aesthetic-enhanced universal style transfer. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1095–1106, 2022.
- [Wu, 2022] Xianchao Wu. Creative painting with latent diffusion models. *arXiv preprint arXiv:2209.14697*, 2022.
- [Xie *et al.*, 2023] Jinheng Xie, Yuexiang Li, Yawen Huang, Haozhe Liu, Wentian Zhang, Yefeng Zheng, and Mike Zheng Shou. Boxdiff: Text-to-image synthesis with training-free box-constrained diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7452–7461, 2023.
- [Yang *et al.*, 2022] Fan Yang, Haibo Chen, Zhanjie Zhang, Lei Zhao, and Huaizhong Lin. Gating pattern pyramid for diversified image style transfer. *Journal of Electronic Imaging*, 31(6):063007, 2022.
- [Zhang *et al.*, 2021] Tao Zhang, Zhanjie Zhang, Wenjing Jia, Xiangjian He, and Jie Yang. Generating cartoon images from face photos with cycle-consistent adversarial networks. *Computers, Materials and Continua*, 2021.
- [Zhang *et al.*, 2023a] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023.
- [Zhang *et al.*, 2023b] Yuxin Zhang, Weiming Dong, Fan Tang, Nisha Huang, Haibin Huang, Chongyang Ma, Tong-Yee Lee, Oliver Deussen, and Changsheng Xu. Prospect: Prompt spectrum for attribute-aware personalization of diffusion models. *ACM Transactions on Graphics (TOG)*, 42(6):1–14, 2023.
- [Zhang *et al.*, 2023c] Yuxin Zhang, Nisha Huang, Fan Tang, Haibin Huang, Chongyang Ma, Weiming Dong, and Changsheng Xu. Inversion-based creativity transfer with diffusion models. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [Zhang *et al.*, 2023d] Zhanjie Zhang, Jiakai Sun, Jiafu Chen, Lei Zhao, Boyan Ji, Zehua Lan, Guangyuan Li, Wei Xing, and Duanqing Xu. Caster: Cartoon style transfer via dynamic cartoon style casting. *Neurocomputing*, 556:126654, 2023.
- [Zhang *et al.*, 2024a] Zhanjie Zhang, Jiakai Sun, Guangyuan Li, Lei Zhao, Quanwei Zhang, Zehua Lan, Haolin Yin, Wei Xing, Huaizhong Lin, and Zhiwen Zuo. Rethink arbitrary style transfer with transformer and contrastive learning. *Computer Vision and Image Understanding*, page 103951, 2024.
- [Zhang *et al.*, 2024b] Zhanjie Zhang, Quanwei Zhang, Guangyuan Li, Wei Xing, Lei Zhao, Jiakai Sun, Zehua Lan, Junsheng Luan, Yiling Huang, and Huaizhong Lin. Artbank: Artistic style transfer with pre-trained diffusion model and implicit style prompt bank. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.
- [Zheng *et al.*, 2021] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. The spatially-correlative loss for various image translation tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16407–16417, 2021.
- [Zhu *et al.*, 2017] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
- [Zuo *et al.*, 2023] Zhiwen Zuo, Lei Zhao, Ailin Li, Zhizhong Wang, Zhanjie Zhang, Jiafu Chen, Wei Xing, and Dongming Lu. Generative image inpainting with segmentation confusion adversarial training and contrastive learning. *arXiv preprint arXiv:2303.13133*, 2023.