StyleMamba : State Space Model for Efficient Text-Driven Image Style Transfer

Zijia Wang ^{a,c,*} and Zhi-Song Liu^b

^aImperial College London, the United Kingdom ^bLappeenranta-Lahti University of Technology, Lahti, Finland ^cDell Research UK, the United Kingdom

Abstract. We present StyleMamba , an efficient image style transfer framework that translates text prompts into corresponding visual styles while preserving the content integrity of the original images. Existing text-guided stylization requires hundreds of training iterations and takes a lot of computing resources. To speed up the process, we propose a conditional State Space Model for Efficient Text-driven Image Style Transfer, dubbed StyleMamba , that sequentially aligns the image features to the target text prompts. To enhance the local and global style consistency between text and image, we propose masked and second-order directional losses to optimize the stylization direction to significantly reduce the training iterations by $5\times$ and the inference time by $3\times$. Extensive experiments and qualitative evaluation confirm the robust and superior stylization performance of our methods compared to the existing baselines. Full code of this paper can be found in https://github.com/OliverDOU776/StyleMamba.

1 Introduction

In recent years, the intersection of computer vision and natural language processing has led to significant advancements in multimodal perception and understanding. One particularly intriguing area of research within this domain is text-driven or text-guided image style transfer. This emerging field explores the synthesis of images guided by textual descriptions, enabling the transformation of visual content to match desired artistic styles, scenes, or aesthetics as described in natural language. This fusion of vision and language offers promising avenues for creative expression and content generation.

Unlike image-driven style transfer, one notable advantage of textdriven style transfer is its enhanced flexibility and interpretability. Texts provide a more abstract and semantically rich representation of desired styles or attributes than explicit reference images. This abstraction enables users to express complex artistic concepts and emotions that are challenging to convey through images alone. Existing text-driven approaches, e.g., Clipstyler [19], DiffusionCLIP [18] and TxST [25], show tremendous promising results in visual diversity. The key idea is to align the text and image embeddings [31, 41], such that the stylized image can be faithful to the direction of the text. However, they need to consider the nuances of style-specific text description and visual style consistency. More importantly, they require hundreds of training iterations and considerable GPU resources, which is inefficient for real applications.

Thus, we introduce StyleMamba, a framework to incorporate a conditional State Space Model [9] into the AutoEncoder systems which can use text prompts to supervise the style fusion process. It can significantly improve the training converge with $10 \sim 20 \times$ speedup. Furthermore, we explore the Masked and Second-order directional loss to better align text and style. Figure 1 showcases the superiority of StyleMamba in performing style transfer, illustrating its capability to rapidly assimilate and apply complex styles to various content images with remarkable fidelity. Through a series of comparisons, StyleMamba demonstrates enhanced performance in not only fewer epochs (Figure 1 a) but also finer detail preservation (Figure 1 b) and adherence to diverse artistic prompts(Figure 1 c), ranging from mimicking the thick textures of Picasso's oil paintings to capturing the essence of atmospheric and lighting conditions. In Figure 1 d, these examples highlight StyleMamba 's better balance between content preservation and style transfer ability, enabling more generalized applications. Our contributions can be summarized below:

- We propose a simple framework, StyleMamba , incorporating the conditional Mamba into AutoEncoder to achieve fast text-driven style transfer (Section 3.1).
- Our proposed framework achieves better stylization via novel Masked directional loss and Second-order relational loss (Section 3.2), which can speed up overall stylization and better grasp global and local style consistency without compromising the contents.
- We empirically show the efficiency of the proposed StyleMamba over the existing state-of-the-art techniques in both quantitative and qualitative measurements. Owing to the simplistic nature of StyleMamba, we further experiment with video style transfer and multiple style transfer to show the versatility of our framework.

2 Related Work

Image-driven style transfer. Style transfer aims to fuse desirable style to the content images such that the resultant image can preserve the contents while showing the desired style, i.e., color and textures. Early approaches such as Gatys et al.'s work [7] on neural style transfer laid the groundwork by formulating style transfer as an optimization problem, leveraging deep convolutional neural networks to separate content and style representations within images. Building upon this foundation, subsequent research [22, 15, 12, 28, 24, 43] has explored various techniques to improve the efficiency, stability, and

^{*} Email: zijia.wang18@imperial.ac.uk



Figure 1. Comparative results showcasing the efficacy of the StyleMamba framework. a) highlights the rapid convergence and stylization capabilities with fewer of epochs. b) and c) demonstrate the detailed stylization fidelity and the transferability of various styles, including painting and lighting styles, as well as some complicated styles like environmental style. Finally, d) shows the great content preservation ability of StyleMamba with competitive style transfer performance compared with other style transfer models. All results reflect the superiority of StyleMamba .

fidelity of style transfer algorithms. For instance, Li et al. [15] introduced adaptive instance normalization (AdaIN), which dynamically adjusts the statistics of intermediate feature maps to better match the style of a reference image. Other approaches have incorporated perceptual loss functions inspired by human visual perception to preserve semantic content during style transfer better. Additionally, recent works [2, 11, 14, 27] have explored conditional and controllable style transfer, allowing users to specify desired style attributes or manipulate style transfer outcomes with good visual quality.

Text-image multimodality. Benefiting from the study of LLMs and VLMs, it becomes popular to use contrastive learning [3, 4] to align general texts and images. CLIP [31] is one of the pioneering works that can use natural language supervision for image representation. The key idea is to utilize pre-trained language and image encoders to project texts and images onto the compressed domain and maximize cosine similarity between paired text-image data. To further improve the generalization and text-image alignment, SigLIP [41] proposes to replace conservative learning with pairwise sigmoid loss. It can scale up the batch size and end up with memory-efficient and superior text-image alignment. Both CLIP and SigLIP offer zero-shot capacities that can be used for many text-image multimodal tasks. For example, [32] proposes a text-conditional image generation that can convert texts to desirable images. The Stable Diffusion Model [33] trains a CLIP-based autoencoder that can take text embeddings as conditions to further improve image quality. A similar concept can also be applied to video generation [26, 23], video retrieval [39], 3D generation [36, 20], 3D editing [16] and image captioning [8].

Text-driven style transfer. Different image-driven style transfer, text-driven style transfer uses text prompts to guide the stylization process. Clipstyler [19] is the first work to utilize CLIP to project both stylized images and texts to the latent space. It can use pairwise cosine similarity to maximize their distance. It further inspires several works on improving quality or speeding up the training iterations. For example, DiffusionCLIP [18] combines the state-of-the-art pre-trained diffusion model [33] and CLIP to learn stylization in the latent space. Gatha [17] proposes to modify the global directional loss between texts and images as a relational loss, which is calculated based on a style tensor that can question whether the generated image is truly aligned with the target texts. One of the key disadvantages of the aforementioned methods is the online training process. It usually takes up to 150 iterations (approximately 10 minutes on one standard GPU) to obtain one good stylized image. Recent works

either speed up the whole training times [35] or perform offline contrastive training [6, 25, 38, 40].

3 Approach

The overall flow of the proposed StyleMamba is shown in Figure 2. The input content image **X** is converted to latent representations $F_e(\mathbf{X})$ using a pretrained Variational AutoEncoder (VAE) $F_e(\cdot)$ from stable diffusion [33], while the style text t is converted into embeddings $F_T(t)$ using pretrained text encoder $F_T(\cdot)$ in SigLIP [41]. They are then fused in the proposed style fusion module to get a new feature map **M**. Guided by the text-to-image style loss, the decoder $F_d(\cdot)$ decodes the **M** into a stylized image **Y**. The training process is guided by the content loss ($\mathcal{L}_{content}$) based on VGG [34] or lpips [42], the second order loss (\mathcal{L}_{so}), the masked directional loss (\mathcal{L}_{md}) and global directional loss (\mathcal{L}_{dir}).

3.1 Overall framework

The whole architecture consists of three components: i) Auto Encoder, which is used to encode the content image X and reconstruct the stylized Image Y; ii) Style Fusion Module, which works as a fuser to selectively attend the $F_T(t)$ and $F_e(X)$ to derive M; iii) SigLIP Module, which is used to encode the style features contained in $F_T(t)$ and to act as a 'North Start' to guide the whole process.

Auto Encoder. The proposed masked directional loss, coupled with second-order loss, needs the powerful content-preserving ability of the Auto Encode, while the better reconstruction ability is where the pretrained VAE from the Stable Diffusion Model [33] excels, as it is adept at encoding the content image into a rich, compressed latent space while ensuring that the essential features are retained for accurate reconstruction. Moreover, the absence of shortcuts between the encoder and decoder in this architecture means that the stylization process relies entirely on the latent representation, fostering a more profound style integration without residual information from the input image. This approach aligns with the Information Bottleneck principle, striking an optimal balance between data compression and preserving relevant content features for the task. Therefore, the pretrained VAE in Stable Diffusion Model [33] is chosen here to be $F_e(\cdot)$.

Style Fusion Module. Inspired by DiT [29], we propose a conditional State Space Model for style fusion. Concretely, the Style Fusion Module in the StyleMamba framework effectively combines the



Style image Style image Style image

epoch 10

epoch 20

Figure 2. Workflow overview of StyleMamba framework. The process begins with a content image and a style prompt (e.g., "Paul Gauguin style"). An encoder converts the content image into a latent representation, which undergoes style fusion with features derived from the style prompt. This fusion is facilitated by the Style Fusion Module, incorporating masked and second-order directional losses to guide the text-to-image stylization. The result is a stylized image **Y**, which closely adheres to the style prompt while preserving content integrity.

Content

image X

epoch 0

Stylized image Y

textual style features with the content image features by incorporating elements of Adaptive Layer Norm (AdaLN) and Mamba process [9]. Mamba's selective process accelerates the style transfer process by focusing only on relevant features, which speeds up training and inference. AdaLN's ability is to adaptively adjust style characteristics from textual descriptions using learnable shifting and scaling values. The AdaLN approach allows for the adjustment of the normalized content features to resonate with the stylistic elements described in the text, achieving adaptive style manipulation. Specifically, the Style Fusion Module incorporates AdaLN by adjusting feature distributions based on the style features derived from the text, ensuring the content image's integrity while infusing the desired style. The normalization of content feature $F_e(\mathbf{X})$ is depicted mathematically as:

Style Mamba

Lsecond_order

Style prompt (t):

Paul Gauauin style

Content (X)

Lconten

$$\mathbf{M} = (LN(F_e(\mathbf{X}) + \alpha_1 \cdot SSM(LN(F_e(\mathbf{X}))) \cdot \mu_1 + \sigma_1)) + \alpha_2 + \sigma_2$$
⁽¹⁾

In Equation 1, LN denotes Layer Normalization, and SSM represents the State Space Model. In our case, we use Mamba [9], α and σ are scaling and shifting parameters, respectively, which are learned from the style features. The Mamba, with its selective state space model, allows for content-dependent transformation, selectively attending to relevant information and dynamically adjusting to the content based on the current context. This allows for efficient and adaptive style transfer while maintaining content consistency.

SigLIP Module. For the SigLIP Module model section, the introduction of the pretrained SigLIP [41] is central to enhancing the style fusion process, which is critically dependent on the rich semantic encoding capabilities of the pretrained CLIP models. SigLIP achieves better zero-shot classification accuracy on ImageNet [5] than CLIP [31]. SigLIP can better generalize to unseen image categories based on natural language descriptions. To demonstrate its superior performance over CLIP, we conduct ablations in the Experiment section.

3.2 Proposed Text-driven style losses

Global directional loss. In the proposed net framework, the global directional loss \mathcal{L}_{dir} is the key guidance in aligning the transformation of the image content with the specified text-driven style. Like the direction loss in [19], it operates by computing the cosine similarity between the direction vectors of text and image features, ensuring that the stylization process is coherent with the textual description. Formally, this loss is defined as:

$$\mathcal{L}_{dir} = 1 - \frac{\mathbf{T}_{dir} \cdot \mathbf{I}_{dir}}{\|\mathbf{T}_{dir}\| \|\mathbf{I}_{dir}\|}, \text{ where}$$

$$\mathbf{T}_{dir} = F_T(t) - F_T(t_{src}), \quad \mathbf{I}_{dir} = F_I(\mathbf{Y}) - F_I(\mathbf{X})$$
(2)

Fused Feature (F

where F_T and F_I are the text encoder and image encoder of the SigLIP, respectively, then t and src are style text and source text (here we define it as "a plain photo"), \mathbf{T}_{dir} and \mathbf{I}_{dir} are the direction vectors for the text and image modalities, respectively.

Masked directional Loss. As shown in [10], randomly masking patches can speed up the reconstruction process and serve as a self-supervised training paradigm for image reconstruction tasks. In StyleMamba , the mask can work efficiently to enforce the convergence regarding style loss. Furthermore, we consider that the style in the images is locally consistent and independent from the contents. Hence, we mark partial contents to enforce the style similarity close to the complete image results. Therefore, we extend the directional loss with a masking step to incorporate robustness against partial visibility of style features. Given the content image **X** and style text *t*, the stylized image **Y** is generated and subsequently masked randomly by 50% using 16×16 patches to obtain the masked image *Z*. The features of *Z*, $F_I(Z)$, and the content features $F_I(\mathbf{X})$ are then computed to calculate the masked directional loss. The proposed masked directional loss.

$$\mathcal{L}_{\rm md} = 1 - \frac{\mathbf{T}_{\rm dir} \cdot \mathbf{I}_{\rm dir}^{\rm mask}}{\|\mathbf{T}_{\rm dir}\| \|\mathbf{I}_{\rm dir}\|} , \text{ where } \mathbf{I}_{\rm dir}^{\rm mask} = F_I(\mathbf{Z}) - F_I(\mathbf{X}) \quad (3)$$

It leverages the reconstruction challenge as a self-supervised signal, enhancing the model's ability to maintain style fidelity even when significant portions of the image are masked, drawing insights from the effectiveness of high masking ratios in self-supervised learning settings.

Second-order directional loss. However, to speed up the convergence speed, the second-order loss \mathcal{L}_{so} introduces a novel approach by considering the progression of the stylized image across successive epochs. As shown in Figure 3, the second-order directional loss (L_{so}) is designed to speed up the optimization of the alignment between the generated image and the textual style description. It achieves this by considering the changes in stylization direction across successive epochs, ensuring a coherent and swift transition towards the desired visual style. This is quantified by the squared norm

Mamba

Laver Norm



Figure 3. Illustration of the proposed second-order directional loss. It shows how L_{so} allows for rapid adjustments in the direction of stylization. Notably, it facilitates refined stylistic shifts, ensuring a swift and coherent transition towards the desired visual style.

of the difference between the image features at consecutive epochs, relative to the text direction vector, and is weighted by a dynamic term, (α_{shift}), to adjust the influence of the style shift on the loss function. Concretely, given a content image **X**, style text *t*, and stylized images **Y**(*i*) and **Y**(*i* + 1) obtained at epochs *i* and *i* + 1, the second-order directional loss \mathcal{L}_{so} is formulated as follows:

$$\mathcal{L}_{so} = 1 - \frac{\mathbf{T}_{dir} \cdot \mathbf{I}_{dir}^{so}}{\|\mathbf{T}_{dir}\| \|\mathbf{I}_{dir}\|} \cdot \lambda_{shift} , \text{ where } \mathbf{I}_{dir}^{so} = F_I(\mathbf{Y}(i+1)) - F_I(\mathbf{Y}(i))$$
(4)

which measures the change in stylization from epoch *i* to i + 1 relative to the change in text features. Specifically, λ_{shift} is a dynamic weighting term that could be formulated as:

$$\lambda_{\text{shift}} = \alpha \cdot \left(1 - e^{-\beta \|\mathbf{F}_I(\mathbf{Y}(\mathbf{i})) - \mathbf{F}_I(\mathbf{X})\|} \right) \tag{5}$$

where α and β are hyperparameters that control the sensitivity of the loss to the magnitude of the style shift. Concretely, α serves as an upper limit or a scaling factor for the weight. It can be used to adjust the maximum influence the style shift weight can have on the secondorder loss. β is a scaling hyperparameter that controls the sensitivity of the weight to changes in the style feature difference. A larger β would make the weight more responsive to smaller differences in style features. The exponential term ensures that the weight increases as the difference between the target and source style features grows, but it is bound to prevent the loss from escalating excessively. When the style feature difference is large (meaning the current style is far from the target), this term is close to 0, making $\alpha_{shift} = \alpha$, which in turn makes the weight larger, allowing the second-order loss to have a smaller impact on the training update, and vice versa. Furthermore, to enhance the training stability and ensure fast updates, a threshold is applied to the direction loss \mathcal{L}_{dir} . The second-order loss \mathcal{L}_{so} is employed only if \mathcal{L}_{dir} falls below a predefined threshold θ , allowing for the refinement of style transfer in later stages of training. Including the aforementioned masked directional loss, the overall proposed text-driven style loss is defined as follows.

$$\mathcal{L}_{\text{style}} = \mathcal{L}_{\text{dir}} + \mathcal{L}_{\text{md}} + \mathbb{I}(\mathcal{L}_{\text{dir}} < \theta) \cdot \mathcal{L}_{\text{so}}$$
(6)

where $\mathbb{I}(\cdot)$ is the indicator function, which is 1 if the condition is true, and 0 otherwise.

Content losses. The VGG content loss [28] (L_{vgg}) is used to supervise content similarity. It is defined as the sum of squared errors between the feature maps of the content and the stylized image, extracted from various layers of the VGG network. Meanwhile, we also

use the LPIPS [42] loss L_{lpips} to align the stylized results more closely with human perceptual judgments.

Total loss. We incorporate the content loss and the style losses as $L_{total} = \omega_1 L_{style} + \omega_2 L_{lpips} + \omega_3 L_{vgg}$, where $\omega_{1,2,3}$ are the weights that balance the contribution of each loss component to the total loss. This formulation ensures that the stylization aligns with the textual description while preserving the integrity of the content image.

4 Experiments

• **Datasets.** To evaluate our StyleMamba framework, we utilized two distinct datasets: COCO[21] and WikiArt[1]. It is important to note that these datasets were not employed during the training phase of our model. Instead, they are used in the testing and inference stages, providing a diverse range of images and artistic styles to assess the effectiveness of our text-driven image style transfer approach. We use the COCO dataset as content images while applying various textual styles to evaluate the model's performance in capturing and rendering complex artistic styles.

• **Parameter setting.** For our experiments, we adapted the pretrained VAE from Stable Diffusion [33], opting to keep the encoder fixed while only training the decoder. To ensure stable training, we introduced a two-stage learning rate strategy for the content loss weight. Initially, we set the weight at 9000, which is then reduced to 150 after 5 epochs. This approach was taken to prevent early overfitting to content features and to gradually refine the content alignment as the training progressed. Additionally, our training procedure included a masking component, where we randomly masked 50% of the patches. With these settings, we trained our models using the Adam optimizer with a starting learning rate of 5×10^{-4} , halving it at epoch 10, and terminating the training at 20 epochs.

• Metrics and evaluation. To assessour textimage model's effectiveness in content consistency and style alignment, we employed three metrics: **1. CLIP score**: This is quantified by computing the cosine similarity between the features extracted by the CLIP (SigLIP) model from the text and the image. A **higher** CLIP similarity score indicates better style alignment. **2. SSIM**: It is used to measure the similarity between the content of the original image and the stylized image. A **higher** SSIM index signifies a better preservation of content. **3. VGG Loss**: It captures the content differences between content and stylized images by examining the feature responses at various layers of the VGG network. A **lower** VGG content loss implies better performance.

4.1 Comparison with the state of the art

	CLIP score ↑	SSIM ↑	VGG loss \downarrow	Aesthetic score↓
Clipstyler[19]	0.180	0.988	2.106	5.379
DiffusionCLIP [18]	0.327	0.979	1.272	5.742
Diffstyler [13]	0.351	0.982	1.182	4.660
TxST [25]	0.409	0.990	0.975	4.552
ZeCon [40]	0.417	0.971	1.218	5.327
Our final model	0.492	0.996	0.547	4.172

 Table 1.
 Comparison with Sate-of-the-art methods.
 StyleMamba is

 consistently better than other text-guided style transfer models regarding the CLIP score, SSIM, VGG loss and Aesthetic score.
 StyleMamba is

To show the superiority of our proposed StyleMamba, we compare it with several state-of-the-art approaches, including Diffusion-CLIP [18], Diffstyler [13], TxST [25] and ZeCon [40]. We first quantitively compare the performance of these models, the experiment is



Figure 4. Qualitive comparison with SOTA algorithms. We show three cases of text-guided style transfer. For reference, we use the text to retrieve a reference image for style comparison and the input to InstantStyle.

done on 10 random content images and 10 text prompts of various descriptions (100 results each time), and we repeat this process around 20 times for final experiment results. Among them, TxST claims to be real-time stylization without online optimization, and ZeCon claims to be a faster version of Clipstyler [19]. The CLIP score of StyleMamba is 0.492, the highest in the group, indicating proficiency in generating semantically coherent images in response to textual descriptions. StyleMamba also shows higher values on SSIM and VGG content loss. These metrics underscore the model's ability to retain the structural integrity of the content images. Lastly, measured by the Aesthetic differences metric [30], shows the lowest score among the compared models, suggesting that StyleMamba is capable of producing images that are aesthetically appealing to humans. Overall, our proposed StyleMamba effectively produces images that are visually aligned with text prompts, structurally similar to reference images and superior in aesthetic quality. A qualitative comparison is also conducted based on three cases as shown in Figure 4. The retrieved image is derived based on the prompt text, which is used as a reference for the style transfer performance and the input to InstantStyle [37]. Apparently, StyleMamba can achieve consistently better results in all cases regarding content preservation and style transfer ability. Concretely, StyleMamba outperforms Clipstyler, Diffusion-CLIP, Diffstyler and ZeCon regarding both style transfer and content preservation capabilities. Compared with InstantStyle ([37]), Style-Mamba shows competitive style transfer ability without taking in a reference style image, while the content preservation ability is better than InstantStyle.

Model	Training time	Inference Time	Model Param.
Clipstyler[19]	403 sec	4 sec	28M
DiffusionCLIP [18]	293 sec	56 sec	214 M
Diffstyler [13]	-	42 sec	207M
TxST [25]	-	27 sec	51M
ZeCon [40]	-	36 sec	67M
Ours	23 sec	2 sec	17M

Table 2. Computational complexity comparison of different models. We record the training time, inference time, and model parameters.

In Table 2, we show the computation complexity among different approaches. We record both computation time and model parameters to demonstrate the efficiency of our proposed StyleMamba . Note that Diffstyler, TxST and Zecon require long offline training hours to achieve real-time stylization. The approximate training time could be hundreds of hours. We can see that ours achieves $10 \sim 20 \times$ speedup compared to Clipstyler and DiffusionCLIP on training time. Due to the proposed fast Mamba based style fusion module, Style-Mamba also achieves a faster inference process, approximately $2 \times$ faster than Clipstyler.

4.2 Ablation study

Style Fusion module. To test the effectiveness of the style fusion module, we compare two settings: ours with the Mamba module, and the cross-attention module used in [25]. Table 3 shows the experiment results. The combination of SigLIP and Mamba achieves the best results regarding the CLIP similarity and the VGG content loss. We can also observe that using our proposed style fusion module achieves better stylization and runs faster than the cross-attention module, approximately $3 \times$ speedup. This is due to the efficiency of Mamba's sequential process over the cross-attention's global correlation. It can also be observed from the model complexity in Table 2.

Modules	CLIP score ↑	SSIM ↑	Running time↓
Clipstyler	0.180	0.988	4 sec
CLIP+cross-attention	0.387	0.979	7 sec
SigLIP+cross-attention	0.391	0.987	7 sec
CLIP+Mamba	0.422	0.990	2 sec
SigLIP+Mamba	0.452	0.991	2 sec

Table 3. Ablation study on Style fusion module. We compare different

 Style Fusion modules used in the StyleMamba . The combination of SigLIP

 and Mamba achieves the best results in style transfer accuracy and speed.

Text-to-image models. In Section 3, we propose to replace CLIP [31] with SigLIP [41] because SigLIP generally performs better in text-to-image alignment. Hence, it has the potential for a better understanding of style representation in the images. To demonstrate its efficiency, we randomly select 10 content images and 10 text prompts to test different models during the training and testing phases on the performance of text-image similarity, content preservation, and style transfer in the generated images. The average results are displayed in Table 4.



Figure 5. Speed comparison between StyleMamba and Clipstyler. It depicts the style loss and corresponding intermediate results over time, illustrating the convergence of each method. The result of StyleMamba at epoch 15 shows a more detailed stylization result than Clipstyler's 100 epoch result.

Training & evaluation	CLIP score ↑	SSIM \uparrow	VGG loss ↓
CLIP & CLIP	0.276	0.989	3.920
CLIP & SigLIP	0.066	0.989	3.920
SigLIP & CLIP	0.325	0.990	3.899
SigLIP & SigLIP	0.127	0.990	3.899

 Table 4. Ablation study on text-to-image models. To measure the alignment between texts and images, we use CLIP and SigLIP for model training and evaluation. The CLIP score is calculated by using corresponding evaluation models.

In Table 4, we show the training and evaluation combinations of using different text-to-image models. It shows that the experiments with the SigLIP model during evaluation (rows 3 and 4) outperform the corresponding experiments using the CLIP model (rows 1 and 2). Specifically, the CLIP similarity score increases from 0.276 to 0.325 and from 0.066 to 0.127, when comparing the CLIP testing model against the SigLIP testing model. This suggests that the SigLIP model is more effective at capturing and translating stylistic elements from text to image. Moreover, comparing rows 1 and 3, we can see that using SigLIP to train the model can significantly improve the CLIP score, which indicates that SigLIP is better at understanding the style representation and aligning the stylized image closer to the style prompts. This is also indicated by comparing rows 2 and 4. **Content Encoder** F_e . Furthermore, we explored the synergistic effects of combining different text-to-image models with VGG and stable diffusion autoencoders (SD VAE) [33]. As shown in Table 5, the results indicated that SD VAE outperforms VGG AE in conjunction with the same CLIP model, yielding a higher CLIP similarity score and lower VGG content loss. Specifically, we can see improvements with SD VAE (row 2 > row 1, and row 4 > row 3), suggesting that SD VAE is more effective in capturing and preserving style and contents. Further, the SigLIP model consistently delivered higher CLIP scores than CLIP models (1&2, 3&4), affirming the superiority of SigLIP in image-text alignment.

Loss functions \mathcal{L}_{style} . As shown in Table 6, to assess the impact of various losses on the stylization process, we conducted a series of experiments by systematically integrating different loss functions and evaluating their influence on stylization. For comparison, the base-line model only utilizes the global directional loss (\mathcal{L}_{dir}) and VGG

Modules	CLIP score ↑	SSIM ↑	VGG loss \downarrow
SigLIP+VGG AE	0.322	0.990	3.874
SigLIP+SD VAE	0.349	0.985	2.759
CLIP+VGG AE	0.126	0.789	3.852
CLIP+SD VAE	0.140	0.785	2.937

 Table 5. Ablation study on Encoder model. To better preserve the contents, we test using different pretrained autoencoder (VGG [34] and SD VAE [33]) for style transfer.

content loss (\mathcal{L}_{vgg}) in row 1. Subsequently, in row 2, we enhanced the loss function framework of Clipstyler by incorporating the proposed masked directional loss (\mathcal{L}_{md}), aimed at refining the model's ability to capture directional attributes within the content image. This addition sought to bolster the CLIP score, a metric indicative of style-content alignment. Further augmentation, in row 3, involved the amalgamation of LPIPS (\mathcal{L}_{lpips}) with the masked directional loss(\mathcal{L}_{md}), facilitating simultaneous improvements in both content preservation, as reflected by the content loss metric, and style alignment, as indicated by the improved CLIP score. Finally, in row 4, we integrated a second-order loss \mathcal{L}_{so} , implemented at an interval of every 5 epochs. It is pertinent to note that narrower intervals were avoided to prevent potential instability in the convergence pattern. The experimental outcomes, as presented in Table 6, were generated under a consistent training duration of only 50 epochs across all models.

Comparing rows 1 and 2, we can see that masked directional loss markedly enhances the style similarity. Rows 2 and 3 show that adding LPIPS loss can significantly improve content preservation, with approximately 2.2 reductions in VGG loss. The last row is the complete losses used in StyleMamba, we can see that adding secondorder directional loss can further improve the style similarity without compromising the content information.

To assess the training efficiency of the proposed StyleMamba compared to Clipstyler, we plot the progression of style loss over 100 epochs (Figure 5). The red line in the plot indicates that StyleMamba begins to exhibit signs of convergence by epoch 15 while Clipstyler achieves similar style loss at epoch 100. Approximately, we achieve $7 \times$ speedup. Qualitatively, we show the enlarged patches generated by the StyleMamba and Clipstyler, and we can see that ours can bet-

Loss terms	CLIP score ↑	SSIM ↑	VGG loss ↓
$\mathcal{L}_{dir} + \mathcal{L}_{vgg}$ (baseline)	0.329	0.986	0.624
baseline+ \mathcal{L}_{md}	0.422	0.889	2.779
baseline+ $\mathcal{L}_{md} + \mathcal{L}_{lpips}$	0.428	0.985	0.593
baseline+ $\mathcal{L}_{md} + \mathcal{L}_{lpips} + \mathcal{L}_{so}$	0.475	0.986	0.572

Table 6. Ablation study on loss functions. To compare different loss terms, we gradually add more loss terms to evaluate the performance. The final model combining \mathcal{L}_{md} , \mathcal{L}_{lpips} and \mathcal{L}_{so} achieves the best results in terms of the style similarity and content similarity.

ter transfer the "bee hives" patterns to the content image.



Figure 6. StyleMamba in the wild. Our better balance between content preservation and style transfer ability can enable many applications with high content preservation requirements, while with good style transfer ability.

4.3 Style transfer in the wild

To test the generalization of our proposed StyleMamba , we show six real applications in Figure 6, which showcases the versatile capabilities of the proposed network across a range of creative domains, demonstrating its robust generalization ability. Figure 6 a shows the multiple style transfer. Given the two text prompts, "cyberpunk" and "impressionism", we see a coherent stylized result that can reflect both style features. Figure 6 b highlights the network's application in product design. Knowing the mask of the contents, we can customize the appearance of the desktop by providing any text prompts. Figure 6 c shows that the proposed StyleMamba can understand the contents from sketches and generate desirable colors and textures to complete the painting, indicating its ability for painting assistance. This ability could speed up the process of art creation, providing artists with a powerful tool to visualize and experiment with different color palettes and textural effects. The network's prowess extends to user interface (UI) design, where in Figure 6 d, it reimagines the main page of ECAI 2024 with a watercolor painting style. This adaptation is not just cosmetic but conceptual, showcasing how a digital platform can embody the aesthetics of traditional art, potentially making digital experiences more organic and visually engaging. Then the cinematic style transformation in Figure 6 e highlights StyleMamba 's capacity to alter not just the mood but the narrative tone of a scene.

It adeptly applies a "dark tone" to a typical urban setting, demonstrating the potential to guide the viewer's emotional response and enrich tvisual media storytelling. Lastly, Figure 6 f explores the frontier of fashion design. Here, StyleMamba demonstrates a keen eye for the transposition of iconic art onto wearable designs. The network takes the quintessential patterns of a painting and coherently wraps them around a three-dimensional garment, bridging the gap between classical art and modern apparel. This instance exemplifies how Style-Mamba could be a transformative tool for fashion designers, enabling them to create innovative prints and textiles inspired by a vast array of visual art forms.

These six applications not only affirm the adeptness of Style-Mamba in understanding and executing complex visual prompts and illuminate its potential as a catalyst in creative industries, from enriching visual storytelling to revolutionizing product and fashion design.

4.4 Limitations and discussion

As shown in Figure 7, StyleMamba has limitations on understanding content-guided or less commonly used texts. For example, it can not understand "*Pixar styles*" for face editing (Figure 7 a) or "*watermelon*" (Figure 7 b) for shape transformation. Moreover, a controled style transfer can not be done with StyleMamba , while the CLIP doesn't have the segmentation capcity.

These limitations highlight areas for future research and development, such as improving the model's ability to handle diverse facial features and expanding its understanding of novel and abstract concepts for more accurate and varied style transfers.



Figure 7. Failure examples. StyleMamba has some limitations on face images and object manipulation.

5 Conclusion

This paper presented StyleMamba, an innovative framework for textdriven image style transfer that utilizes a conditional State Space Model integrated within an AutoEncoder architecture. Our framework demonstrates a significant advancement in the field by reducing the number of training iterations and also the inference time and training time required for each epoch, as well as enhancing the efficiency of the stylization process without sacrificing the quality of the stylized images. The results from our extensive experiments illustrate that StyleMamba outperforms existing baselines in terms of speed, stylization accuracy, and content preservation. The introduction of novel loss functions, masked and second-order directional losses, have been particularly effective in achieving high-quality stylization that aligns closely with textual descriptions. This capability allows for greater flexibility and creative expression in the applications of style transfer, making StyleMamba a valuable tool for both artistic endeavors and practical applications like UI design, clothes design, etc. In future work, we aim to explore the integration of more diverse linguistic inputs and expand the model's ability to handle a broader range of visual styles in a finer, controlled way.

728

- [1] K. nichol. painter by numbers, wikiart. https://www.kaggle.com/c/ [29] painter-by-numbers, 2016.
- J. An, S. Huang, Y. Song, D. Dou, W. Liu, and J. Luo. Artflow: Unbiased [2] image style transfer via reversible neural flows. In CVPR, 2021.
- [3] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. 2020.
- [4] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. E. Hinton. Big self-supervised models are strong semi-supervised learners. In NeurIPS, 2020.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In CVPR09, 2009.
- [6] T.-J. Fu, X. E. Wang, and W. Y. Wang. Language-driven artistic style transfer. In ECCV, 2022.
- [7] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In CVPR, 2016.
- [8] S. Ge, T. Park, J.-Y. Zhu, and J.-B. Huang. Expressive text-to-image generation with rich text. In IEEE International Conference on Computer Vision (ICCV), 2023.
- [9] A. Gu and T. Dao. Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752, 2023.
- [10] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick. Masked autoencoders are scalable vision learners. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 16000–16009, 2022.
- [11] K. Hong, S. Jeon, H. Yang, J. Fu, and H. Byun. Domain-aware universal style transfer. In ICCV, 2021.
- [12] N. Huang, F. Tang, W. Dong, and C. Xu. Draw your art dream: Diverse digital art synthesis with multimodal guided diffusion. MM '22, page 1085-1094, 2022. ISBN 9781450392037.
- [13] N. Huang, Y. Zhang, F. Tang, C. Ma, H. Huang, W. Dong, and C. Xu. Diffstyler: Controllable dual diffusion for text-driven image stylization. IEEE Transactions on Neural Networks and Learning Systems, pages 1-14, 2024, doi: 10.1109/TNNLS.2023.3342645.
- [14] S. Huang, J. An, D. Wei, J. Luo, and H. Pfister. Quantart: Quantizing image style transfer towards high visual fidelity. In 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 5947-5956, jun 2023. doi: 10.1109/CVPR52729.2023.00576.
- [15] X. Huang and S. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In ICCV, 2017.
- [16] J. Hyung, S. Hwang, D. Kim, H. Lee, and J. Choo. Local 3d editing via 3d distillation of clip knowledge. In 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 12674-12684, 2023. doi: 10.1109/CVPR52729.2023.01219.
- [17] S. Jandial, S. Deshmukh, A. Java, S. Shahid, and B. Krishnamurthy. Gatha: Relational loss for enhancing text-based style transfer. In 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 3546-3551, 2023.
- [18] G. Kim, T. Kwon, and J. C. Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 2416-2425, 2022.
- [19] G. Kwon and J. C. Ye. Clipstyler: Image style transfer with a single text condition. arXiv preprint arXiv:2112.00374, 2021.
- [20] C.-H. Lin and et al. Magic3d: High-resolution text-to-3d content creation. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2023
- [21] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In ECCV, 2014.
- [22] S. Liu, T. Lin, D. He, F. Li, M. Wang, X. Li, Z. Sun, Q. Li, and E. Ding. Adaattn: Revisit attention mechanism in arbitrary neural style transfer. In ICCV, 2021.
- [23] X. Liu, C. Gong, L. Wu, S. Zhang, H. Su, and Q. Liu. Fusedream: Training-free text-to-image generation with improved clip+gan space optimization, 2021.
- [24] Z.-S. Liu, V. Kalogeiton, and M.-P. Cani. Multiple style transfer via variational autoencoder. In ICIP, 2021.
- [25] Z.-S. Liu, L.-W. Wang, W.-C. Siu, and V. Kalogeiton. Name your style: text-guided artistic style transfer. In 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 3530-3534, 2023. doi: 10.1109/CVPRW59228.2023.00359.
- [26] H. Luo, L. Ji, M. Zhong, Y. Chen, W. Lei, N. Duan, and T. Li. CLIP4Clip: An empirical study of clip for end to end video clip retrieval. arXiv preprint arXiv:2104.08860, 2021.
- [27] W. Lötzsch, M. Reimann, M. Büßemeyer, A. Semmo, J. Döllner, and

M. Trapp. Wise: Whitebox image stylization by example-based learning. ECCV, 2022

- [28] D. Y. Park and K. H. Lee. Arbitrary style transfer with style-attentional networks. In CVPR, 2019.
- W. Peebles and S. Xie. Scalable diffusion models with transformers. In 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pages 4172-4182, 2023. doi: 10.1109/ICCV51070.2023.00387.
- [30] J. D. Pressman, K. Crowson, and S. C. Contributors. Simulacra aesthetic captions. Technical Report Version 1.0, Stability AI, 2022. https:// github.com/crowsonkb/simulacra-aesthetic-models.
- A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, [31] G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision, 2021.
- [32] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. ArXiv, abs/2204.06125, 2022. URL https://api.semanticscholar.org/CorpusID: 248097655
- [33] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. Highresolution image synthesis with latent diffusion models. In CVPR, pages 10674-10685, 2022. doi: 10.1109/CVPR52688.2022.01042.
- [34] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In ICLR, 2015.
- [35] A. P. Suresh, S. Jain, P. Noinongyao, A. Ganguly, U. Watchareeruetai, and A. Samacoits. Fastclipstyler: Optimisation-free text-based image style transfer using style representations. In 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pages 7301-7310, 2024. doi: 10.1109/WACV57701.2024.00715.
- J. Tang, J. Ren, H. Zhou, Z. Liu, and G. Zeng. Dreamgaussian: Gener-[36] ative gaussian splatting for efficient 3d content creation. arXiv preprint arXiv:2309.16653, 2023
- [37] H. Wang, Q. Wang, X. Bai, Z. Qin, and A. Chen. Instantstyle: Free lunch towards style-preserving in text-to-image generation. arXiv preprint arXiv:2404.02733, 2024.
- [38] Z. Wang, L. Zhao, and W. Xing. Stylediffusion: Controllable disentangled style transfer via diffusion models. In 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pages 7643-7655, 2023.
- [39] H. Xue, Y. Sun, and et al. Clip-vip: Adapting pre-trained image-text model to video-language alignment. In International Conference on Learning Representations, 2022.
- [40] S. Yang, H. Hwang, and J. C. Ye. Zero-shot contrastive loss for textguided diffusion image style transfer. In 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pages 22816–22825, 2023.
- [41] X. Zhai, B. Mustafa, A. Kolesnikov, and L. Beyer. Sigmoid loss for language image pre-training. In 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pages 11941-11952, 2023.
- [42] R. Zhang and et al. The unreasonable effectiveness of deep features as a perceptual metric. CVPR, 2018.
- [43] Y. Zhang, F. Tang, W. Dong, H. Huang, C. Ma, T.-Y. Lee, and C. Xu. Domain enhanced arbitrary image style transfer via contrastive learning. In ACM SIGGRAPH, 2022.