

Fine-Grained Image Style Transfer with Visual Transformers

Jianbo Wang¹, Huan Yang², Jianlong Fu²,
Toshihiko Yamasaki¹, and Baining Guo²

¹ The Univerisity of Tokyo {jianbowang815,yamasaki@cvm.t.u-tokyo.ac.jp}

² Microsoft Research {huayan,jianf,bainguo@microsoft.com}

Abstract. With the development of the convolutional neural network, image style transfer has drawn increasing attention. However, most existing approaches adopt a global feature transformation to transfer style patterns into content images (e.g., AdaIN and WCT). Such a design usually destroys the spatial information of the input images and fails to transfer fine-grained style patterns into style transfer results. To solve this problem, we propose a novel **STyle TRansformer (STTR)** network which breaks both content and style images into visual tokens to achieve a fine-grained style transformation. Specifically, two attention mechanisms are adopted in our STTR. We first propose to use self-attention to encode content and style tokens such that similar tokens can be grouped and learned together. We then adopt cross-attention between content and style tokens that encourages fine-grained style transformations. To compare STTR with existing approaches, we conduct user studies on Amazon Mechanical Turk (AMT), which are carried out with 50 human subjects with 1,000 votes in total. Extensive evaluations demonstrate the effectiveness and efficiency of the proposed STTR in generating visually pleasing style transfer results³.

Keywords: Style Transfer · Vision Transformer

1 Introduction

Image style transfer has been receiving increasing attention in the creation of artistic images. Given one content and one style reference image, the model will produce an output image that retains the core elements of the content image but appears to be “painted” in the style of the reference image. It has many industrial applications, for example, clothes design [1], photo and video editing [2,3,4,5], material changing [6], fashion style transfer [7,8], virtual reality [9], and so on.

In recent years, style transfer employs deep neural networks. Those methods could be divided into three categories: 1) optimization-based methods, 2) feed-forward approximation, and 3) zero-shot style transfer. Gatys et al. [10] propose to optimize the pixel values of a given content image by jointly minimizing

³ Code is available at <https://github.com/researchmm/STTR>.

the feature reconstruction loss and style loss. It could produce impressive results, but it requires optimizing the content image for many iterations for any content-style image pairs which are computationally expensive. To solve this problem, other researchers addressed this problem by building a feed-forward network [11,12,13,14,15,16,17], to explicitly learn the mapping from a photograph to a stylized image with a particular painting style. Thus it requires retraining a new model for any unseen new styles. Zero-shot style transfer is more effective since it could handle various styles and transfer images given even unseen styles. Huang et al. [18] propose an arbitrary style transfer method by matching the mean-variance statistics between content and style features, usually named adaptive instance normalization (AdaIN). AdaIN first normalizes the input content image, then scales and shifts it according to parameters calculated by different style images. A more recent work replaces the AdaIN layer with a pair of whitening and coloring transformations [19]. Many works follow the formulation of AdaIN and further improve it [20,21].

However, the common limitation of these methods is that simply adjusting the mean and variance of feature statistics makes it hard to synthesize complicated style patterns with rich details and local structures. As shown later in Fig. 2, Gatys et al.’s method [10], AdaIN [18], and WCT [19] always bring distorted style patterns in the transferred results, which makes it hard to recognize the original object from the style transferred image. In a more recent work [22], Deng et al. propose the StyTr² to learn semantic correlations between the content and style features by attention. Although StyTr² could produce visually pleasing results, there still exists structure distortion as shown later in Fig. 2. This is because StyTr² adopts a shallow feature extractor without pre-trained weights which limits its abilities to capture and transfer style patterns and makes foreground and background objects indistinguishable in the resulting image. Thus, how to learn a better representation to assemble the content and style patterns and transfer fine-grained style patterns while keeping content structure is still a challenging problem.

In this paper, we adopt a Transformer to tackle this problem. Recently, the breakthroughs of Transformer networks [23] in the natural language processing (NLP) domain has sparked great interest in the computer vision community to adapt these models for vision tasks. Inspired by those works which explicitly model long-range dependencies in language sequences and images, we develop a Transformer-based network to first break content and style images into visual tokens then learn the global context between them. As similar content tokens lean to match with the same style tokens, fine-grained style transformation can be obtained. Specifically, the proposed STTR mainly consists of four parts: two backbones to extract and downsample features from the inputs (i.e., content and style images), two self-attention modules to summarize the style or content features, a cross-attention module to match style patterns into content patches adaptively, and a CNN-based decoder to upsample the output of cross-attention module to reconstruct the final result. We implement this by Transformer since

its encoder consists of self-attention module while the decoder has cross-attention module to compute the correlations between content and style tokens.

Two backbones first extract compact feature representation from given content and style images and downsample them to a smaller spatial size. This could help to reduce the computation cost in the Transformer. Secondly, the self-attention modules encourage the style (or content) features with similar structures and semantic meanings to group together. Different from previous work, feeding content and style into the self-attention model before matching them could help to separate them into different distinctive semantic vectors (see Sec. 4.3 for more information). The cross-attention module in the decoder further incorporates style into content harmoniously without hurting the content structure. After that, a CNN-based decoder upsamples the output of the decoder to obtain the stylized results. The details of network architecture could be found in Sec. 3.

Our main contributions are as follows.

- We propose a fine-grained Transformer-based image style transfer method, capable of computing very local mapping between the content and style tokens.
- To better understand the effectiveness of STTR to model global context between style and content features, we provide a detailed ablation study to explore the contribution of the Transformer architecture and losses.
- We evaluated the performance of the proposed STTR with several state-of-the-art style transfer methods. Experimental results show that the STTR has a great capability of producing both perceptual quality and details in transferring style into a photograph. We also provide frame-wise video style transfer results which further demonstrate the effectiveness of our approach.

2 Related work

2.1 Image Style Transfer

Image style transfer is a technique that aims to apply the style from one image to another content image. Typically, neural style transfer techniques could be divided into three aspects: 1) optimization-based methods, 2) feed-forward approximation, and 3) zero-shot style transfer.

Optimization-based methods. Gatys et al. [10] are the first to formulate style transfer as the matching of multi-level deep features extracted from a pre-trained deep neural network. The authors achieve style transfer by jointly minimizing the feature loss [24] and the style loss formulated as the difference of Gram matrices. Optimizing an image is computationally expensive and contains no learned representations for the artistic style, which is inefficient for real-time applications.

Feed-forward approximation. Fast feed-forward approaches [11,12,13] address this problem by building a feed-forward neural network, i.e., style transfer

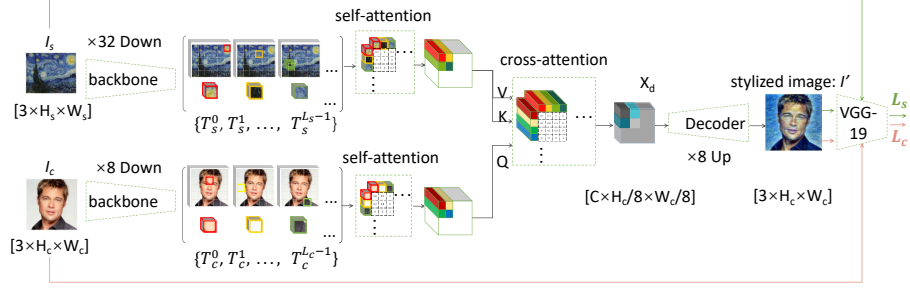


Fig. 1. An overview of our proposed Style Transformer (STTR). Our framework consists of four parts: two backbones which are used to extract features from the style and content image and break them into tokens, two self-attention modules to encourage the style (or content) tokens with similar structures and semantic meanings to group together, a cross-attention module observes the outputs and adaptively match the style code to the content features, and the CNN-based decoder upsamples the output of the Transformer decoder to reconstruct the final result. We use a fixed-weight VGG-19 network to compute the content and style loss as described in Sec. 3.2.

network, to explicitly learn the transformation from a photograph to a particular painting style. They minimize the same feature reconstruction loss and style reconstruction loss as [25]. However, they have to train individual networks for every style, which is very inefficient.

Zero-shot style transfer. Chen et al. [26] propose to match each content patch to the most similar style patch and swapped them. For fast arbitrary style transfer, Huang et al. [18] employ adaptive instance normalization (AdaIN) to normalize activation channels for stylization. Specifically, it follows an encoder-AdaIN-decoder formulation. The encoder uses the first few layers of a fixed VGG-19 network to encode the content and style images. The AdaIN layer is used to perform style transfer in the feature space. Then a decoder is trained to invert the AdaIN output to the image spaces. Unlike AdaIN, Whitening and Coloring Transform (WCT) [19] directly uses different layers of VGG network as the encoders, transfers multi-level style patterns by recursively applying whitening and coloring transformation, and trains the decoders to invert the feature into an image. Many works follow the formulation of AdaIN and further improve it [20,21,27,28]. In SANET [29] and AdaAttN [30], the authors propose to learn semantic correlations between the content and style features by a learnable non-local module. However, those feature transformation methods always fail to maintain content structures since they simply transfer features across all spatial locations for each channel. Hong et al. [31] propose a unified architecture for both artistic and photo-realistic stylizations. In StyTr², the authors propose a transformer-based style transfer model and design the tokenizer with a single convolutional layer (an unfold operation followed by a convolutional filter). Due

Algorithm 1: The training process of the proposed STTR

-
- Input:** Content Image I_c , Style Image I_s
Output: Stylized image I'
- 1 $X_c \leftarrow$ extracted feature maps from I_c by content backbone.
 - 2 $X_s \leftarrow$ extracted feature maps from I_s by style backbone.
 - 3 Linear flat X_c into L_c visual tokens: $\{T_c^0, T_c^1, \dots, T_c^{L_c-1}\}$. Linear flat X_s into L_s visual tokens: $\{T_s^0, T_s^1, \dots, T_s^{L_s-1}\}$.
 - 4 Compute style codes $\{C_s^0, C_s^1, \dots, C_s^{L_s-1}\}$ with $\{T_s^0, T_s^1, \dots, T_s^{L_s-1}\}$ by Transformer encoder.
 - 5 Match $\{T_c^0, T_c^1, \dots, T_c^{L_c-1}\}$ with semantic similar style codes $\{C_s^0, C_s^1, \dots, C_s^{L_s-1}\}$ by Transformer decoder, generate decoded visual tokens: $\{T_d^0, T_d^1, \dots, T_d^{L_c-1}\}$.
 - 6 Reshape $\{T_d^0, T_d^1, \dots, T_d^{L_c-1}\}$ with the same size of X_c , obtain X_d .
 - 7 Reconstruct X_d by CNN-based decoder, obtain the final output I' , it has the same size with I_c .
 - 8 Compute the content and style loss defined in Sec. 3.2 by a fixed weight VGG-19 network.
 - 9 Updates the weights by minimizing the loss value.
-

to the shallow feature extractor for the tokenizer, the learned correspondences between the content and style tokens are random.

2.2 Visual Transformer

Recently, there is impressive progress in the field of NLP, driven by the innovative architecture called Transformer [23]. In response, several attempts have been made to incorporate self-attention constructions into vision models. Those methods could be divided into two aspects: high-level vision and low-level vision.

High-level vision Transformer. Transformer can be applied to many high-level tasks like image classification. Dosovitskiy et al. propose a Vision Transformer (ViT) [32], dividing an image into patches and feeding these patches (i.e., tokens) into a standard Transformer. Recently, Transformers have also shown strong performance on other high-level tasks, including detection [33,34,35,36,37], segmentation [38,39,40], and pose estimation [41,42,43,44].

Low-level vision Transformer. Low-level vision tasks often take images as outputs (e.g., high-resolution or denoised images), which is more challenging than high-level vision tasks such as classification, segmentation, and detection, whose outputs are labels or boxes. iGPT [45] trains GPT v2 model [46] on flattened image sequences (1D pixel arrays) and shows that it can generate plausible image outputs without any external supervision. Zeng et al. [47] propose a Transformer-based Generation. Nowadays, researchers achieve good performance on Transformer networks for super-resolution and frame interpolation [48,49,50,51].

Image style transfer is also a sub-task of low-level vision tasks. Specifically, there exists a large appearance gap between input and output domains. This makes it difficult to produce a visually pleasing target output.

3 Approach

We show the pipeline of our proposed STTR in Fig. 1. Our method takes a content-style image pair as input and produces stylized results as output. In this section, we present the details of the network architecture and loss functions accordingly. The whole training process is illustrated in Alg. 1.

3.1 Network Architecture

As shown in Fig. 1, our model contains four main components. We will describe the network architecture as below: two CNN-based backbones to extract compact feature representations, two self-attention modules to encode the style or content features, a cross-attention module to match style patterns into content patches adaptively, and a CNN decoder to transform the combined output features from Transformer to the final output I' .

Tokenizer In our model, one image is divided into a set of visual tokens. Thus, we have to first convert the input image into a set of visual tokens. We assume that each of them represents a semantic concept in the image. We then feed these tokens to a Transformer. Let us denote the input feature map by $X \in \mathbb{R}^{H \times W \times C}$ (height H , width W , and channels C) and visual tokens by $T \in \mathbb{R}^{L \times C}$ where L indicates the number of tokens.

Filter-based tokenizer utilizes a deep CNN-based feature extractor to gradually downsample input images to obtain feature maps. After that, each position on the feature map represents a region in the input space (i.e., the receptive field). If the output size of feature maps is $H/32 \times W/32 \times C_2$ (height H , width W , and channels C_2), we then obtain $(H/32 \times W/32)$ patches (i.e., tokens). For each token, its dimension is C_2 . As CNN densely slides the filter over the input image, it could produce much more smooth features. We further provide a detailed illustration of different tokenizers in the supplementary.

Backbone To achieve fine-grained style transformation, inspired by the original Transformer [23], one image could be divided into a set of visual tokens as one sentence consisting of a few words. Thus, we have to first convert the input image into a set of visual tokens. We assume that each of them represents a semantic concept in the image. We then feed these tokens to a Transformer. Let us denote the input feature map by $X \in \mathbb{R}^{H \times W \times C}$ (height H , width W , and channels C) and visual tokens by $T \in \mathbb{R}^{L \times C}$ where L indicates the number of tokens.

We adopt ResNet-50 [52] as our backbone since it gradually downsamples the features from the input into a smaller spatial size. The original ResNet-50 [52]

has four layers and downsamples features by 4, 8, 16, and 32 times. The output channel of each layer is 256, 512, 1024, and 2048, respectively. For extracting style features, we choose the output from the 4th layer while for content features we choose the 2nd layer. We would like to utilize the ability of the shallow network (the 2nd layer) to extract and retain detailed information of the spatial structure. For style features, higher-level semantic features are extracted by a deeper network (the 4th layer).

In summary, starting from the input content image $I_c \in \mathbb{R}^{H_c \times W_c \times 3}$ and style image $I_s \in \mathbb{R}^{H_s \times W_s \times 3}$, two backbones will extract features independently with the size of $H_c/8 \times W_c/8 \times 512$ and $H_s/32 \times W_s/32 \times 2048$ for content and style, respectively.

The extracted style and content features have a different number of channels as we described above. Thus we first use a 1×1 convolution to make them have the same channel number d . Specifically, we set $d = 256$. Next, we flatten the 3D features in spatial dimensions, resulting in 2D features, $(H_c/8 \cdot W_c/8) \times 256$ for content features while $(H_s/32 \cdot W_s/32) \times 256$ for style features.

Attention Layer Our proposed STTR consists of an encoder module and a decoder module with several encoder or decoder layers of the same architecture. The encoder layer is composed of a self-attention layer and a feed-forward neural network while the decoder layer is composed of a self-attention layer, a feed-forward neural network and a cross-attention layer and the other feed-forward neural network.

In each attention layer, following the original Transformer [23], the input features X are first projected onto these weight matrices to get $Q = XW^Q$, $K = XW^K$ and $V = XW^V$, where the Q , K , and V denote the triplet of query, key, and value through one 1×1 convolution on each input features from an image (i.e. either content or style image). Then the output is given by:

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

where the Q , K , and V denote query, key and value as described above. All of the three inputs to attention layer maintain the same dimension d .

The detailed network architecture of STTR's transformer could be found in the supplementary.

Transformer Encoder The encoder is used to encode style features, it consists of six encoder layers. Each encoder layer has a standard architecture and consists of a self-attention module and a feed-forward network (FFN). Here, the number of the head is set to be eight. To supplement the image features with position information, fixed positional encodings are added to the flattened features before the features are fed into each attention layer.

Transformer Decoder The decoder is used to reason relationships between content and style tokens while being able to model the global context among

them. It consists of two multi-head attention layers and a feed-forward network (FFN). The first multi-head attention layer mainly learns the self-attention within content features while the second one learns the cross-attention between content and style features.

As shown in Fig. 1, the difference from that of the original Transformer is the input query of the decoder. We adopt content visual tokens as the query. Additionally, we also introduce a fixed positional encoding with the same size of style features.

CNN Decoder The CNN decoder is used to predict the final stylized image. It upsamples the output of the Transformer (with the size of $H_c/8 \times W_c/8 \times 512$) into the original size the same with the input content image. To achieve this, the CNN decoder should upsample it three times. Each upsample module consists of three parts: one standard residual block defined in [52], one bilinear interpolation layer upsample by a factor of 2, and one 3×3 convolution.

Let $RBCk$ denotes a **R**esidual block-**B**ilinear interpolation-**C**onvolution layer where k indicates the channel number of the output features. The CNN decoder architecture consists of $RBC256 - RBC128 - RBC64 - RBC3$.

3.2 Objective Function

Following existing works [13,53,14,18], we use the pre-trained VGG-19 [54] to compute the loss function to train the STTR. The loss function \mathcal{L} is defined as a combination of two loss terms, content loss \mathcal{L}_c and the style loss \mathcal{L}_s :

$$\mathcal{L} = \mathcal{L}_c + \lambda \mathcal{L}_s, \quad (2)$$

where λ denote the weight of style loss. We empirically set $\lambda=10$.

The content loss \mathcal{L}_c is defined by the squared Euclidean distance between the feature representations of the content image I_c and the stylized image I' . Thus we define the content loss \mathcal{L}_c as follows:

$$\mathcal{L}_c = \|f_4(I_c) - f_4(I')\|_2, \quad (3)$$

where we interpret $f_i(\cdot)$ as feature map extracted by a pre-trained VGG-19 network [54] at the layer i . Here, we adopt *relu1_1*, *relu2_1*, *relu3_1*, and *relu4_1* as the layers for extracting features.

Similar to [18], the style loss is defined as follows:

$$\begin{aligned} \mathcal{L}_s = \sum_{i=1}^4 (&\|\mu(f_i(I_s)) - \mu(f_i(I'))\|_2 \\ &+ \|\sigma(f_i(I_s)) - \sigma(f_i(I'))\|_2), \end{aligned} \quad (4)$$

where $\mu(\cdot)$, $\sigma(\cdot)$ are the mean and standard deviation, computed across batch size and spatial dimensions independently for each feature channel.




Fig. 2. Visual comparison. All of content and style images are collected from websites and copyright-free. They are never observed during training.

4 Experiments and Discussions

4.1 Training Details

We train our network by using images from MSCOCO [55] and WikiArt [56] as content and style data, respectively. Each dataset contains about 80,000 images. In each training batch, we randomly choose one pair of content and style images with the size of 512×512 as inputs. We implement our model with PyTorch and apply the Adam optimizer [57] with a learning rate of 10^{-5} .

4.2 Comparison Results

Qualitative Evaluations There are many quantitative evaluation methods for image quality assessment [58,59]. However, as style transfer is a very subjective task and does not have a reference image, to evaluate our method, we first evaluate the quality of the results produced by our model. We visually compare our results to the state-of-the-art style transfer methods: the optimization-based method proposed by Gatys et al. [10], three feature transformation-based methods WCT [19], AdaIN [18], AAST [21], patch-based method Avatar-Net [60], attention based methods (SANET [29] and StyTr [22]), ArtFlow [27], and MCCNet [28]. Fig. 2 shows the results.




Fig. 3. User preference result of seven style transfer algorithms.

As shown in the figure, the proposed STTR performs favorably against the state-of-the-art methods. Both the content structure and style patterns appear well since we match the style tokens onto content tokens in a fine-grained manner. Although the optimization-based method [10] allows universal-style transfer, it is computationally expensive due to its iterative optimization process (see analysis about the speed in Sec. 4.2). Also, the results highly depend on many hyper-parameters, including the number of iterations, the trade-off loss weight, and the initial image. Thus, the results are not robust (e.g., see the 1st, 3rd, 6th, and 7th rows in Fig. 2). The AdaIN [18] method presents efficient solution for arbitrary style transfer, but it generate sub-optimal results (e.g., see the 1st and 7th rows in Fig. 2) as it globally stylized the whole content image. The WCT [19] and Avatar-Net usually perform well, however, in some cases, they would bring strong distortions which makes it hard to recognize the original object from the style transferred image (e.g., see the 1st, 2nd, and 7th rows in Fig. 2). AAST [21] transfers color and texture in two independent paths but sometimes cannot produce well-colored results (see the 2nd and 4th rows in Fig. 2). In contrast, our method learns semantic relationships between content and style tokens and performs appealing results for arbitrary style, especially working very well on preserving content structure. SANET [29], ArtFlow [27], and MCCNet [28] also hard to generate appealing results especially when the content and style features could not be matched very well (e.g., see the 1st and 3rd rows in Fig. 2).

StyTr² [22] also adopt a transformer to build a style transfer model and it could produce visually pleasing results. However, it still cannot preserve the fine structures. The main difference is that our method learns correspondence on the semantic-level while StyTr² focuses on the region-level. This is because StyTr² adopts a shallow feature extractor to build the tokenizer. Thus the learned correspondences between the content and style tokens are random. For example, on the 6th row, StyTr² generates fires in the sky (marked in the red box) while in

our method, the flame area is smoother and appears on the edge of the building. It looks like StyTr² tends to copy-paste the texture in the style image onto the content image. The same thing happens on the 7th row. Our model understands the sail and the body of the boat are different objects so the result is clean. However, StyTr² produces results with large yellow areas across the sea and the sail (marked in the white box). It looks like StyTr² tends to copy-paste the texture in the style image onto the content image. Taking advantage of knowledge, our method could learn semantic-level correspondences between the content and style images and reduce distortion.

Perceptual Study As the evaluation of style transfer is highly subjective, we conduct a user study to further evaluate the seven methods shown in Fig. 2. We use 10 content images and 20 style images. For each method, we use the released codes and default parameters to generate 200 results.

We hire 50 volunteers on Amazon Mechanical Turk (AMT) for our user study. Twenty of 200 content-style pairs are randomly selected for each user. For each style-content pair, we display the stylized results of seven methods on a web page in random order. Each user is asked to vote for the one that he/she likes the most. Finally, we collect 1000 votes from 50 users and calculate the percentage of votes that each method received.

The results are shown in Fig. 3, where our STTR obtains 19.6% of the total votes. It is much higher than that of AAST [21] whose stylization results are usually thought to be high-quality. Attention-based methods, StyTr² [22] and SANET [29] get sub-optimal results. This user study result is consistent with the visual comparisons in Fig. 2 and further demonstrates the superior performance of our proposed STTR.

Table 1. Execution time comparison (in seconds).

Method	Venue	Time (512 px)
Gatys et al. [10]	CVPR16	106.63
AdaIN [18]	CVPR17	0.13
WCT [19]	CVPR17	1.63
Avatar-Net [60]	CVPR18	0.49
SANET [29]	CVPR19	0.19
AAST [21]	ACMMM20	0.49
ArtFlow [27]	CVPR21	0.40
MCCNet [28]	AAAI21	0.15
StyTr ² [22]	CVPR22	0.57
STTR (ours)	0.14	0.28

Efficiency Table 1 shows the run time performance of the proposed method and other methods at 512 pixels image scales. All the methods are tested on a PC




Fig. 4. Visual comparison for video style transfer.

with an Intel(R) Xeon(R) Gold 6226R 2.90GHz CPU and a Tesla V100 GPU. The optimization-based method [10] is impractically computationally expensive because of its iterative optimization process. WCT [19] is also relatively slower as it conducts multi-level coarse-to-fine stylization.

In contrast, our proposed STTR runs at 4 frames per second (fps) for 512×512 images. Therefore, our method could feasibly process style transfer in real time. Our model is much faster than Gatys et al. [10] and achieves better qualitative performance as shown in Fig. 2.

Memory usage Compared with two current SOTAs, our method has fewer parameters than StyTr² (ours: 45.643M v.s. StyTr²: 48.339M) and takes less memory than AdaAttn (ours: 18391M v.s. AdaAttn: 25976M) on a 512×512 image. However, such memory costs could be further reduced by using fewer channel numbers.

Results for Video Style Transfer Here we also provide results of video style transfer. As shown in Fig. 4, our model can perform video stylization on video sequence frame-by-frame. On the contrary, AdaIN [18] produces results with high motion blur. To show the stable results produced by our model, we further compute the difference between neighboring frames to show the smoothness between frames. It is clearly shown that the difference generated by our method is much more close to that from the input frames. It is because our method could well preserve the content structure and reduce the motion-blurred video frames.

4.3 Ablations

To verify the effect of each component in the STTR, we conduct an ablation study by visually comparing the generated results without each of them. All of the compared models are trained by the same 400,000 content-style pairs (i.e., trained for 5 epochs). More ablation experiments could be found in the supplementary.



Fig. 5. Encoder attention maps for different sampling points on style image. The style image (e.g., Van Gogh’s painting) are shown in the center. We further show four sampled points and their according attention maps with different colors.



Fig. 6. Decoder attention maps for different sampling points on the output. The stylized image is shown in the center. We further show three sampled points and their according attention maps with different colors. On the upper left is the input style image.

Learned Attention Maps We can observe model with deeper encoder and decoder has stronger capability to preserve semantic similarity, so that similar style patterns (e.g., fire) can be transferred to similar content regions.

The encoder is used for encoding the style pattern while the decoder learns a matching between the content and style features.

In Fig. 5, we provide visualized attention maps to show that self-attention could encode style features and similar visual tokens can be grouped together. For the top right marked point (which locates the center of the sun), all of the regions related to other suns have a higher value (in yellow color) in their corresponding attention map. In Fig. 6, attention maps show the learned fine-grained relationships between content and style features. The lower left point in the content image has strong relationships (e.g., see its attention map) with the



Fig. 7. Trading off between the content and style images.

black tree in the style image. So the pixels around that point in the output are very dark (see the stylized results located around that point).

The Loss Weight λ The degree of stylization can be modified by the hyperparameter λ as described in Sec. 3.2. As shown in Fig. 7, changing λ in the training stage could control the magnitude of stylization.

5 Conclusion

Image style transfer mainly aims to transfer style patterns into the content images. Taking advantage of the breakthrough of vision Transformers, in this work, we propose a style Transformer for solving image style transfer tasks. Since STTR breaks content and style features into tokens and computes very local relationships between them, fine-grained style transformation could be obtained. Experiments demonstrate that the proposed model can produce visually pleasing results and preserve content structure very well without bringing a heavy timing cost. The proposed method also has the potential to be applied to video style transfer frameworks. There are still some limitations that call for continuing efforts. In the future, we would like to extend STTR to handle multiple style mixtures and work on light-weight Transformer architecture for style transfer.

References

1. Date, P., Ganesan, A., Oates, T.: Fashioning with networks: neural style transfer to design clothes. In: KDD ML4Fashion workshop. Volume 2. (2017) [1](#)
2. Chen, D., Liao, J., Yuan, L., Yu, N., Hua, G.: Coherent online video style transfer. In: ICCV. (2017) 1105–1114 [1](#)
3. Zhang, W., Cao, C., Chen, S., Liu, J., Tang, X.: Style transfer via image component analysis. TMM **15** (2013) 1594–1601 [1](#)
4. Liu, J., Yang, W., Sun, X., Zeng, W.: Photo Stylistic Brush: Robust style transfer via superpixel-based bipartite graph. TMM **20** (2017) 1724–1737 [1](#)
5. Virtusio, J.J., Ople, J.J.M., Tan, D.S., Tanveer, M., Kumar, N., Hua, K.L.: Neural Style Palette: A multimodal and interactive style transfer from a single style image. TMM (2021) [1](#)
6. Matsuo, S., Shimoda, W., Yanai, K.: Partial style transfer using weakly supervised semantic segmentation. In: ICME Workshops, IEEE (2017) 267–272 [1](#)
7. Kim, B.K., Kim, G., Lee, S.Y.: Style-Controlled synthesis of clothing segments for fashion image manipulation. TMM **22** (2019) 298–310 [1](#)
8. Liu, Y., Chen, W., Liu, L., Lew, M.S.: Swapgan: A multistage generative approach for person-to-person fashion style transfer. TMM **21** (2019) 2209–2222 [1](#)
9. Castillo, C., De, S., Han, X., Singh, B., Yadav, A.K., Goldstein, T.: Son of Zorn’s Lemma: Targeted style transfer using instance-aware semantic segmentation. In: ICASSP, IEEE (2017) 1348–1352 [1](#)
10. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: CVPR. (2016) 2414–2423 [1](#), [2](#), [3](#), [9](#), [10](#), [11](#), [12](#)
11. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV, Springer (2016) 694–711 [2](#), [3](#)
12. Li, C., Wand, M.: Precomputed real-time texture synthesis with markovian generative adversarial networks. In: ECCV, Springer (2016) 702–716 [2](#), [3](#)
13. Ulyanov, D., Lebedev, V., Vedaldi, A., Lempitsky, V.S.: Texture Networks: Feed-forward synthesis of textures and stylized images. In: ICML. (2016) 4 [2](#), [3](#), [8](#)
14. Dumoulin, V., Shlens, J., Kudlur, M.: A learned representation for artistic style. arXiv preprint arXiv:1610.07629 (2016) [2](#), [8](#)
15. Chen, D., Yuan, L., Liao, J., Yu, N., Hua, G.: Stylebank: An explicit representation for neural image style transfer. In: CVPR. (2017) 1897–1906 [2](#)
16. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Diversified texture synthesis with feed-forward networks. In: CVPR. (2017) 3920–3928 [2](#)
17. Zhang, H., Dana, K.: Multi-style generative network for real-time transfer. In: ECCV Workshops. (2018) 0–0 [2](#)
18. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV. (2017) 1501–1510 [2](#), [4](#), [8](#), [9](#), [10](#), [11](#), [12](#)
19. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Universal style transfer via feature transforms. arXiv preprint arXiv:1705.08086 (2017) [2](#), [4](#), [9](#), [10](#), [11](#), [12](#)
20. Kitov, V., Kozlovstev, K., Mishustina, M.: Depth-Aware Arbitrary style transfer using instance normalization. arXiv preprint arXiv:1906.01123 (2019) [2](#), [4](#)
21. Hu, Z., Jia, J., Liu, B., Bu, Y., Fu, J.: Aesthetic-Aware image style transfer. In: ACM MM. (2020) 3320–3329 [2](#), [4](#), [9](#), [10](#), [11](#)
22. Deng, Y., Tang, F., Pan, X., Dong, W., Ma, C., Xu, C.: Stytr²: Unbiased image style transfer with transformers. CVPR (2021) [2](#), [9](#), [10](#), [11](#)
23. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. arXiv preprint arXiv:1706.03762 (2017) [2](#), [5](#), [6](#), [7](#)

24. Mahendran, A., Vedaldi, A.: Understanding deep image representations by inverting them. In: CVPR. (2015) 5188–5196 [3](#)
25. Gatys, L.A., Ecker, A.S., Bethge, M.: A neural algorithm of artistic style. arXiv preprint arXiv:1508.06576 (2015) [4](#)
26. Chen, T.Q., Schmidt, M.: Fast patch-based style transfer of arbitrary style. arXiv preprint arXiv:1612.04337 (2016) [4](#)
27. An, J., Huang, S., Song, Y., Dou, D., Liu, W., Luo, J.: Artflow: Unbiased image style transfer via reversible neural flows. In: CVPR. (2021) 862–871 [4](#), [9](#), [10](#), [11](#)
28. Deng, Y., Tang, F., Dong, W., Huang, H., Ma, C., Xu, C.: Arbitrary video style transfer via multi-channel correlation. In: AAAI. (2021) 1210–1217 [4](#), [9](#), [10](#), [11](#)
29. Park, D.Y., Lee, K.H.: Arbitrary style transfer with style-attentional networks. In: CVPR. (2019) 5880–5888 [4](#), [9](#), [10](#), [11](#)
30. Liu, S., Lin, T., He, D., Li, F., Wang, M., Li, X., Sun, Z., Li, Q., Ding, E.: AdaAttN: Revisit attention mechanism in arbitrary neural style transfer. In: ICCV. (2021) 6649–6658 [4](#)
31. Hong, K., Jeon, S., Yang, H., Fu, J., Byun, H.: Domain-Aware Universal Style Transfer. In: ICCV. (2021) [4](#)
32. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An Image is Worth 16x16 Words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020) [5](#)
33. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-End Object Detection with Transformers. In: ECCV, Springer (2020) 213–229 [5](#)
34. Beal, J., Kim, E., Tzeng, E., Park, D.H., Zhai, A., Kislyuk, D.: Toward transformer-based object detection. arXiv preprint arXiv:2012.09958 (2020) [5](#)
35. Pan, X., Xia, Z., Song, S., Li, L.E., Huang, G.: 3d object detection with point-former. arXiv preprint arXiv:2012.11409 (2020) [5](#)
36. Yuan, Z., Song, X., Bai, L., Zhou, W., Wang, Z., Ouyang, W.: Temporal-Channel Transformer for 3D Lidar-Based Video Object Detection in Autonomous Driving. arXiv preprint arXiv:2011.13628 (2020) [5](#)
37. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: Deformable Transformers for End-to-End Object Detection. arXiv preprint arXiv:2010.04159 (2020) [5](#)
38. Wang, H., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: MaX-DeepLab: End-to-End Panoptic Segmentation with Mask Transformers. arXiv preprint arXiv:2012.00759 (2020) [5](#)
39. Wang, Y., Xu, Z., Wang, X., Shen, C., Cheng, B., Shen, H., Xia, H.: End-to-End Video Instance Segmentation with Transformers. arXiv preprint arXiv:2011.14503 (2020) [5](#)
40. Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., et al.: Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. arXiv preprint arXiv:2012.15840 (2020) [5](#)
41. Huang, L., Tan, J., Liu, J., Yuan, J.: Hand-Transformer: Non-autoregressive structured modeling for 3d hand pose estimation. In: ECCV, Springer (2020) 17–33 [5](#)
42. Huang, L., Tan, J., Meng, J., Liu, J., Yuan, J.: HOT-Net: Non-Autoregressive Transformer for 3D Hand-Object Pose Estimation. In: ACM MM. (2020) 3136–3145 [5](#)
43. Lin, K., Wang, L., Liu, Z.: End-to-End Human Pose and Mesh Reconstruction with Transformer. arXiv preprint arXiv:2012.09760 (2020) [5](#)

44. Yang, S., Quan, Z., Nie, M., Yang, W.: TransPose: Towards Explainable Human Pose Estimation by Transformer. arXiv preprint arXiv:2012.14214 (2020) 5
45. Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Luan, D., Sutskever, I.: Generative pretraining from pixels. In: ICML, PMLR (2020) 1691–1703 5
46. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners. OpenAI blog 1 (2019) 9 5
47. Zeng, Y., Yang, H., Chao, H., Wang, J., Fu, J.: Improving visual quality of image synthesis by a token-based generator with transformers. In: NeurIPS. (2021) 5
48. Yang, F., Yang, H., Fu, J., Lu, H., Guo, B.: Learning texture transformer network for image super-resolution. In: CVPR. (2020) 5791–5800 5
49. Liu, C., Yang, H., Fu, J., Qian, X.: Learning trajectory-aware transformer for video super-resolution. In: CVPR. (2022) 5
50. Qiu, Z., Yang, H., Fu, J., Fu, D.: Learning spatiotemporal frequency-transformer for compressed video super-resolution. arXiv preprint arXiv:2208.03012 (2022) 5
51. Liu, C., Yang, H., Fu, J., Qian, X.: TTVFI: Learning Trajectory-Aware Transformer for Video Frame Interpolation. arXiv preprint arXiv:2207.09048 (2022) 5
52. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. (2016) 770–778 6, 8
53. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Improved Texture Networks: Maximizing Quality and Diversity in Feed-forward Stylization and Texture Synthesis. In: CVPR. (2017) 6924–6932 8
54. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014) 8
55. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common Objects in Context. In: ECCV, Springer (2014) 740–755 9
56. Nichol, K.: Painter by numbers, wikiart (2016) 9
57. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv preprint arXiv:1412.6980 (2014) 9
58. Zheng, H., Yang, H., Fu, J., Zha, Z.J., Luo, J.: Learning conditional knowledge distillation for degraded-reference image quality assessment. In: ICCV. (2021) 9
59. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. TIP 13 (2004) 600–612 9
60. Sheng, L., Lin, Z., Shao, J., Wang, X.: Avatar-Net: Multi-scale Zero-shot Style Transfer by Feature Decoration. In: CVPR. (2018) 8242–8250 9, 11