

Multimodal Style Transfer via Graph Cuts

Yulun Zhang^{1,2}, Chen Fang^{2,3}, Yilin Wang², Zhaowen Wang², Zhe Lin², Yun Fu¹, Jimei Yang²
¹Northeastern University, ²Adobe Research, ³ByteDance AI Lab

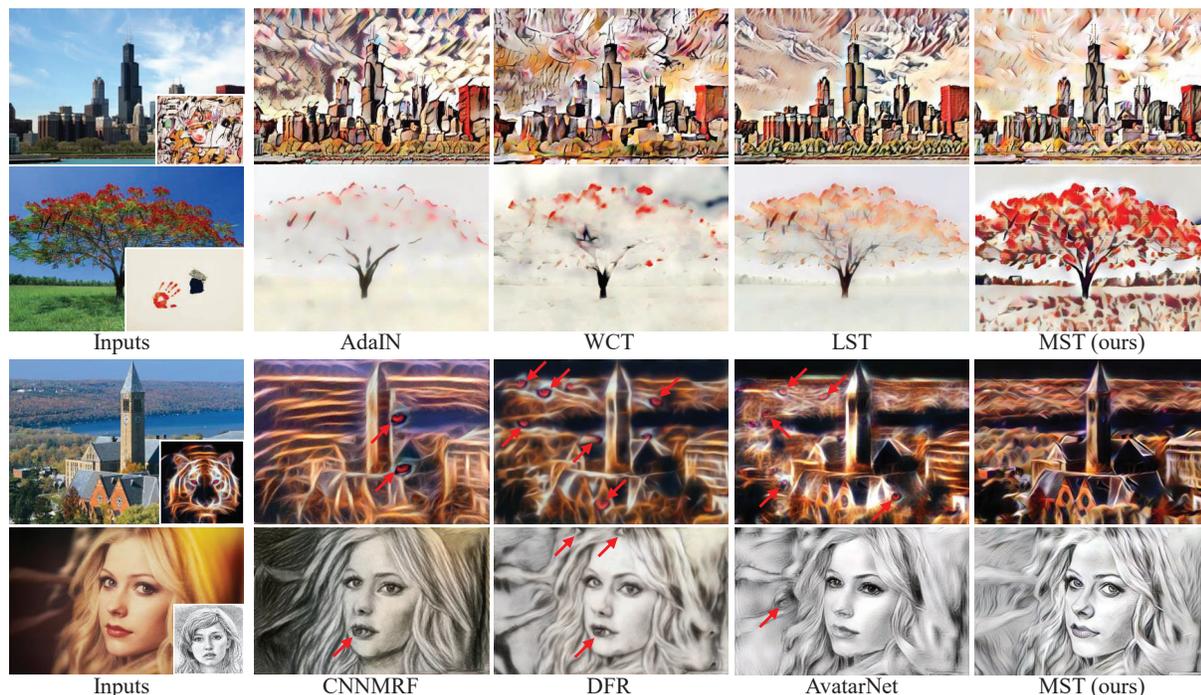


Figure 1: Gram matrix based style transfer methods (AdaIN [11], WCT [24], and LST [23]) may fail to distinguish style patterns (1st and 2nd rows). Patch-swap based methods (CNNMRF [19], DFR [10], and AvatarNet [36]) may copy some less desired style patterns (labeled with red arrows) to the results (3rd and 4th rows). Our MST alleviates all these limitations.

Abstract

An assumption widely used in recent neural style transfer methods is that image styles can be described by global statistics of deep features like Gram or covariance matrices. Alternative approaches have represented styles by decomposing them into local pixel or neural patches. Despite the recent progress, most existing methods treat the semantic patterns of style image uniformly, resulting in unpleasing results on complex styles. In this paper, we introduce a more flexible and general universal style transfer technique: multimodal style transfer (MST). MST explicitly considers the matching of semantic patterns in content and style images. Specifically, the style image features are clustered into sub-style components, which are matched with local content features under a graph cut formulation. A reconstruction network is trained to transfer each sub-style and render the final stylized result. We also generalize MST to improve some existing methods. Extensive experiments demonstrate the superior effectiveness, robustness, and flexibility of MST.

1. Introduction

Image style transfer is the process of rendering a content image with characteristics of a style image. Usually, it would take a long time for a diligent artist to create a stylized image with particular style. Recently, it draws a lot of interests [8, 14, 4, 39, 11, 24, 10, 36, 35, 3, 40, 23] since Gatys *et al.* [8] discovered that the correlations between convolutional features of deep networks can represent image styles, which would have been hard for traditional patch-based methods to deal with. These neural style transfer methods either use an iterative optimization scheme [8] or feed-forward networks [14, 4, 39, 11, 24, 36, 23] to synthesize the stylizations. Most of them are applicable for arbitrary style transfer with a pre-determined model. These universal style transfer methods [11, 24, 36, 23] inherently assume that the style can be represented by the global statistics of deep features such as gram matrix [8] and its approximates [11, 24]. Although these neural style transfer methods can preserve the content well and match the overall style

of the reference style images, they will also distort the local style patterns, resulting unpleasing visual artifacts.

Let’s start with some examples in Fig. 1. In the first row, where the style image consists of complex textures and strokes, these methods cannot tell them apart and neglect to match style patterns to content structures adaptively. This would introduce some less desired strokes in smooth content areas, e.g., the sky. In the second row, the style image has clear spatial patterns (e.g., large uniform background and blue/red hand). AdaIN, WCT, and LST failed to maintain the content structures and suffered from wash-out artifacts. This is mainly because the unified style background occupies a large proportion in the style image, resulting its domination in the global statistics of style features. These observations indicate that it may not be sufficient to represent style features as a unimodal distribution such as a Gram or covariance matrix. An ideal style representation should respect to the spatially-distributed style patterns.

Inherited from traditional patch-based methods, these neural patch-based algorithms could generate visually pleasing results when content and style images have similar structures. However, the greedy example matching usually employed by these methods will introduce less desired style patterns to the outputs. This is illustrated by the bottom two examples in Fig. 1, where some salient patterns in the style images, e.g., the eyes and lips, are improperly copied to the buildings and landscape. Moreover, the last row of Fig. 1 also illustrates shape distortion problem of these methods; e.g., the appearance for the girl has changed. This phenomenon apparently limits the choice of style images for these methods.

To address these issues, we propose the multimodal style transfer (MST), a more flexible and general style transfer method that seeks for a sweet spot between parametric (gram matrix based) and non-parametric (patch based) approaches. Specifically, instead of representing the style with a unimodal distribution, we propose a multimodal style representation with graph based style matching mechanism to adaptively match the style patterns to a content image.

Our main contributions are summarized as follows:

- We analyze the feature distributions of different style images (see Fig. 2) and propose a multimodal style representation that better models the style feature distribution. This multimodal representation consists of a mixture of clusters, each of which represents a particular style pattern. It also allows users to mix-and-match different styles to render diverse stylized results.
- We formulate style-content matching as an energy minimization problem with a graph and solve it via graph cuts. Style clusters are adapted to content features with respect to the content spatial configuration.



Figure 2: t-SNE [28] visualization for style features. The original high-dimension style features are extracted at layer Conv_4_1 in VGG-19 [37] and are reduced to 3 dimensions via t-SNE. We can see the feature distributions tend to fit as multimodal distributions rather than single-modal ones.

- We demonstrate the strength of MST by extensive comparison with several state-of-the-art style transfer methods. The robustness and flexibility of MST is shown with different sub-style numbers and multi-style mixtures. The general idea of MST can be extended to improve other existing stylization methods.

2. Related Works

Style Transfer. Originating from non-realistic rendering [18], image style transfer is closely related to texture synthesis [5, 7, 6]. Gatys *et al.* [8] were the first to formulate style transfer as the matching of multi-level deep features extracted from a pre-trained deep neural network, which has been widely used in various tasks [20, 21, 22]. Lots of improvements have been proposed based on the works of Gatys *et al.* [8]. Johnson *et al.* [14] trained feed-forward style-specific network and produced one stylization with one model. Sanakoyeu *et al.* [34] further proposed a style-aware content loss for high-resolution style transfer. Jing *et al.* [12] proposed a StrokePyramid module to enable controllable stroke with adaptive receptive fields. However, these methods are either time consuming or have to re-train new models for new styles.

The first arbitrary style transfer was proposed by Chen and Schmidt [4], who matched each content patch to the most similar style patch and swapped them. Luan *et al.* [27] proposed deep photo style transfer by adding a regularization term to the optimization function. Based on markov random field (MRF), Li and Wand [19] proposed CNNMRF to enforce local patterns in deep feature space. Ruder *et al.* [33] improved video stylization with temporal coherence. Although their visual stylizations for arbitrary style are appealing, the results are not stable [33].

Recently, Huang *et al.* [11] proposed real-time style transfer by matching the mean-variance statistics between content and style features. Li *et al.* [24] further introduced whitening and coloring (WCT) by matching the covariance matrices. Li *et al.* boosted style transfer with linear style transfer (LST) [23]. Gu *et al.* [10] proposed deep feature reshuffle (DFR), which connects both local and global style losses used in parametric and non-parametric methods. Sheng *et al.* [36] proposed AvatarNet to enable multi-scale transfer for arbitrary style. Shen *et al.* [35] built meta net-

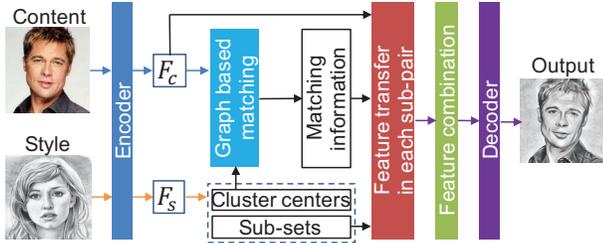


Figure 3: An overview of our MST algorithm.

works by taking style images as inputs and generating corresponding image transformation networks directly. Mechrez *et al.* [29] proposed contextual loss for image transformation. However, these methods fail to treat the style patterns distinctively and neglect to adaptively match style patterns with content semantic information. For more neural style transfer works, readers can refer to the survey [13].

Graph Cuts based Matching. Many problems that arose in early vision can be naturally expressed in terms of energy minimization. For example, a large number of computer vision problems attempt to assign labels to pixels based on noisy measurements. Graph cuts is a powerful method to solve such discrete optimization problems. Greig *et al.* [9] were firstly successful solving graph cuts by using powerful min-cut/max-flow algorithms from combinatorial optimization. Roy and Cox [32] were the first to use these techniques for multi-camera stereo computation. Later, a growing number of researches in computer vision use graph-based energy minimization for a wide range of applications, which includes stereo [16], texture synthesis [17], image segmentation [38], object recognition [1], and others. In this paper, we formulate the matching between content and style features as an energy minimization problem. We approximate its global minimum via efficient graph cuts algorithms. To the best of our knowledge, we are the first to formulate style matching as energy minimization problem and solve it via graph cuts.

3. Proposed Method

We show the pipeline of our proposed MST in Fig. 3.

3.1. Multimodal Style Representation

In previous CNN-based image style transfer works, there are two main ways to represent style. One is to use the features from the whole image and assume that they are in the same distribution (e.g., AdaIN [11] and WCT [24]). The other one treats the style patterns as individual style patches (e.g., Deep Feature Reshuffle [10]). Equal treatments to different style patterns lack flexibility in the real cases, where there has several distributions among style features. Let's see the t-SNE [28] visualization for style features in Fig. 2, where the style features are clustered to multiple groups. Therefore, if a cluster dominates the feature space, e.g., second example of Fig. 1, the Gram matrix based methods [24, 23, 11] will fail to capture the the overall style pat-

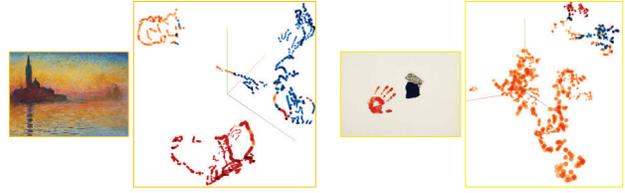


Figure 4: t-SNE [28] visualization for style features with cluster labels. For each style-visualization pair, we set $K = 3$ and label style features with corresponding cluster labels.

terns. On the other hand, patch-based methods which treat each sub-patch distinctly would suffer from copying multiple same style patterns to the results directly. For example, in Fig. 1, the eyes in the style images are copied multiple times, causing displeasing stylization results.

Based on the observations and analyses above, we argue that neither a global statics of deep features nor local neural patches could be a suitable way to represent the complex real-world cases. As a result, we propose multimodal style representation, a more efficient and flexible way to represent different style patterns.

For a given style image I_s , we can extract its deep features $F_s \in \mathbb{R}^{C \times H_s \times W_s}$ via a pre-trained encoder $E_{\theta_{enc}}(\cdot)$, like VGG-19 [37]. H_s and W_s are the height and width of the style feature. To achieve multimodal representation in high-dimension feature space, we target to segment the style patterns into multiple subsets. Technically, we simply apply K-means to cluster all the style feature points into K clusters without considering spatial style information

$$F_s = F_s^{l_1} \cup F_s^{l_2} \cup \dots \cup F_s^{l_k} \cup \dots \cup F_s^{l_K}, \quad (1)$$

where $F_s^{l_k} \in \mathbb{R}^{C \times N_k}$ is the k -th cluster with N_k features and we assign this cluster a label l_k . In the clustered space, features in the same cluster have similar visual properties and are likely drawn from the same distribution (resembling Gaussian Mixture Model [31]). This process helps us obtain a multimodal representation of style.

We visualize multimodal style representation in Fig. 4. For each style image, we extract its VGG feature (at layer Conv_4.1 in VGG-19) and cluster it into $K = 3$ clusters. Then, we conduct t-SNE [28] visualization with the cluster labels. As shown in Fig. 4, clustering results match our assumption of multimodal style representation well. Nearby feature points tend to be in the same cluster. These observation not only shows the multimodal style distribution, but also demonstrates that clustering is a proper way to model such a multimodal distribution.

3.2. Graph Based Style Matching

Like style feature extraction, we extract deep content features $F_c \in \mathbb{R}^{C \times H_c \times W_c}$ from a content image I_c . H_c and W_c are the height and width of the content feature. Distance measurement is the first step before matching. To reach a good distance metric, we should consider the scale differ-

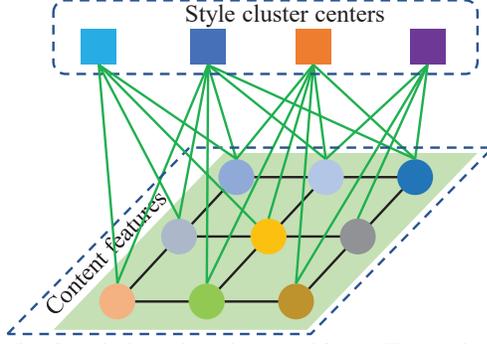


Figure 5: Graph based style matching. Example of the graph containing content features and style cluster centers. We match content features with style cluster in pixel level.

ence between the content and style features. Computation complexity should also be taken into consideration, since all the content features will be used to match. Based on above analysis, we calculate the cosine distance between content feature $F_{c,p} \in \mathbb{R}^{C \times 1}$ and style cluster center $F_{s,l_k} \in \mathbb{R}^{C \times 1}$ as follows

$$D(F_{c,p}, F_{s,l_k}) = 1 - \frac{F_{c,p}^T F_{s,l_k}}{\|F_{c,p}\| \|F_{s,l_k}\|}, \quad (2)$$

where $(\cdot)^T$ is transpose operation and $\|\cdot\|$ is magnitude of the feature vector.

Then we target to find a labeling f that assigns each content feature $F_{c,p}$ with a style cluster center label $f_p \in \{l_1, l_2, \dots, l_K\}$. We formulate the disagreement between f and content features as follows

$$E_{data}(f) = \sum_{p=1}^{HcWc} D(F_{c,p}, F_{s,f_p}), \quad (3)$$

where we name E_{data} as data energy. Minimizing E_{data} encourages f to be consistent with content features.

However, the spatial content information here is not considered, failing to preserve discontinuity and producing some unpleasing structures in the stylized results. Instead, we hope pixels in the same content local region have same labels. Namely, we want f to be piecewise smooth and discontinuity preserving. So, we further introduce another smooth term $E_{smooth}(f)$ as follows

$$E_{smooth}(f) = \sum_{\{p,q\} \in \Omega} V_{p,q}(f_p, f_q), \quad (4)$$

where Ω is the position set of direct interacting pairs of content features. $V_{p,q}$ denotes the distinct penalty for each position pair of features $\{p, q\}$. This has been investigated to be important in various computer vision applications [2]. Also, various forms of energy functions have been investigated before. Here, we take the discontinuity preserving function given by the Potts model

$$V_{p,q}(f_p, f_q) = \lambda \cdot T(f_p \neq f_q), \quad (5)$$

where $T(\cdot)$ is 1 if its argument is true, and otherwise 0. λ is a smooth constant. This model encourages the labeling f to



Figure 6: Visualization of style matching. Here, we cluster style features into $K = 2$ subsets for better understanding.

pursue several regions, where content features in the same region have same style cluster labels.

By taking Eqs. (3) and (4) into consideration, we naturally formulate the style matching problem as a minimization of the following energy function:

$$E(f) = E_{data}(f) + E_{smooth}(f). \quad (6)$$

The whole energy $E(f)$ measures not only the disagreement between f and content features, but also the extent to which f is not piecewise smooth. However, the global minimization of such an energy function is NP-hard even in the simplest discontinuity-preserving case [2].

To solve the energy minimization problem in Eq. (6), we propose to build a graph by regarding content features as p -vertices and style cluster centers as l -vertices (shown in Fig. 5). Then the energy minimization is equal to min-cut/max-flow problem, which can be efficiently solved via graph cuts [2]. After finding a local minimum, the whole content features can be re-organized as follows

$$F_c = F_c^{l_1} \cup F_c^{l_2} \cup \dots \cup F_c^{l_k} \cup \dots \cup F_c^{l_K}, \quad (7)$$

where $F_c^{l_k}$ denotes the sub-set whose content features are matched with the same style label l_k .

We show visualization details about graph based style matching in Fig. 6. We extract style and content features from Conv_4_1 layer in VGG-19. Due to several downsampling modules in VGG-19, the spatial resolution of the features is much smaller than that of the inputs. We label the spacial style feature pixels with their corresponding cluster labels and obtain the style cluster maps. According to the style cluster maps in Fig. 6, we find that style feature clustering grasps semantic information from style images.

After style matching in pixel level, we get the content-style matching map, which also reflects the semantic information, matching the content structures adaptively. Such an adaptive matching alleviates the wash-out artifacts, when the style is very simple or has large area of unified background. Then, we are able to conduct feature transform in each content-style pair group.

3.3. Multimodal Style Transfer

For each content-style pair group $\{F_c^{l_k}, F_s^{l_k}\}$, we first center them by subtracting their mean vectors $\mu(F_c^{l_k})$ and $\mu(F_s^{l_k})$ respectively. We conduct feature whitening and coloring as used in WCT [24].

$$F_{cs}^{l_k} = C_s W_c F_c^{l_k} + \mu(F_s^{l_k}), \quad (8)$$

where $W_c = E_c^{l_k} (D_c^{l_k})^{-\frac{1}{2}} (E_c^{l_k})^T$ is a whitening matrix and $C_s = E_s^{l_k} (D_s^{l_k})^{\frac{1}{2}} (E_s^{l_k})^T$ is a coloring matrix. $E_c^{l_k}$ and $D_c^{l_k}$ are diagonal matrix of eigenvalues and the orthogonal matrix of eigenvectors of the covariance matrix $F_c^{l_k} (F_c^{l_k})^T$. For style covariance matrix $F_s^{l_k} (F_s^{l_k})^T$, the corresponding matrices are $E_s^{l_k}$ and $D_s^{l_k}$. The reasons why we choose WCT to transfer features is its robustness and efficiency [24, 23]. More details about whitening and coloring are introduced in [24].

After feature transformation, we may also want to blend transferred features with content features as did in previous works (e.g., AdaIN [11] and WCT [24]). Most previous works have to blend the whole transferred features with a unified content-style trade-off, which treats different content parts equally and is not flexible to the real-world cases. Instead, our multimodal style representation and matching make it possible to adaptively blend features. Namely, for each content-style pair group, we blend them via

$$F_{cs}^{l_k} = \alpha_k F_{cs}^{l_k} + (1 - \alpha_k) F_c^{l_k}, \quad (9)$$

where $\alpha_k \in [0, 1]$ is a content-style trade-off for specific labeled content features. After blending all the features, we obtain the whole transferred features

$$F_{cs} = F_{cs}^{l_1} \cup F_{cs}^{l_2} \cup \dots \cup F_{cs}^{l_k} \cup \dots \cup F_{cs}^{l_K}. \quad (10)$$

F_{cs} is then fed into the decoder $D_{\theta_{dec}}(\cdot)$ to reconstruct the final output I_{cs} .

3.4. Implementation Details

Now, we specify the implementation details about our proposed MST. Similar to some previous works (e.g., AdaIN, WCT, DFR), we incorporate the pre-trained VGG-19 (up to Conv_4.1) [37] as the encoder $E_{\theta_{enc}}(\cdot)$. We obtain decoder $D_{\theta_{dec}}(\cdot)$ by mirroring the encoder, whose pooling layers are replaced by nearest up-scaling layers.

To train the decoder, we use the pre-trained VGG-19 [37] to compute perceptual loss $l_{total} = l_c + \gamma l_s$, which combines content loss l_c and style loss l_s . We simply set the weighting constant as $\gamma = 10^{-2}$. Inspired by the loss designations in [14, 25, 11], we formulate content loss l_c as

$$l_c = \|\phi_{4.1}(I_s) - \phi_{4.1}(I_{cs})\|_2, \quad (11)$$

where $\phi_{4.1}(\cdot)$ extracts features at layer Conv_4.1 in VGG-19. We then formulate the style loss l_s as

$$l_s = \sum_{i=1}^4 (\|\mu(\phi_{i.1}(I_c)) - \mu(\phi_{i.1}(I_{cs}))\|_2) + \sum_{i=1}^4 (\|\sigma(\phi_{i.1}(I_c)) - \sigma(\phi_{i.1}(I_{cs}))\|_2), \quad (12)$$

where $\phi_{i.1}(\cdot)$ extracts features at layer Conv_ i .1 in VGG-19. We use $\mu(\cdot)$ and $\sigma(\cdot)$ to compute the mean and standard deviation of the content and style features.

We train our network by using images from MSCOCO [26] and WikiArt [30] as content and style data re-



Figure 7: Distance measurement investigation.

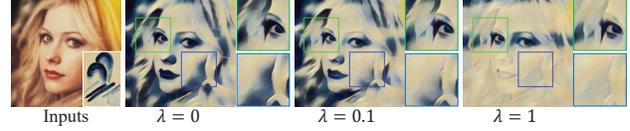


Figure 8: Discontinuity preservation investigation.

spectively. Each dataset contains about 80,000 images. In each training batch, we randomly crop one pair of content and style images with the size of 256×256 as input. We implement our model with TensorFlow and apply Adam optimizer [15] with learning rate of 10^{-4} .

4. Discussions

To better position MST among the whole body of style transfer works, we further discuss and clarify the relationship between MST and some representative works.

Differences to CNNMRF. CNNMRF [19] extracts a pool of neural patches from style images, with which patch matching is used to match content. MST clusters style features into multiple sub-sets and matches style cluster centers with content feature points via graph cuts. CNNMRF uses smoothness prior for reconstruction, while MST uses it for style matching only. CNNMRF minimizes energy function to synthesize the results. MST generates stylization results with a decoder.

Differences to MT-Net. Both color and luminance are treated as a mixture of modalities in MT-Net [39]. MST obtains multimodal representation from style features via clustering. It should also be noted that MT-Net has to train new models for new style images. While, MST is designed for arbitrary style transfer with a single model.

Differences to WCT. In WCT [24], the decoder is trained by using only content data and loss. MST introduces additional style images for training. WCT uses multiple layers of VGG features and conducts multi-level coarse-to-fine stylization, which costs much more time and sometimes distorts structures. While, MST only transfers single-level content and style features. Consequently, even we set $K = 1$ in MST, we achieve more efficient stylizations.

5. Experiments

We conduct extensive experiments to validate the contributions of each component in our method, the effectiveness of our method, and the flexibility for user control.

5.1. Ablation Study

Distance Measurement. We first investigate the choice of the distance measurement, as it is critical for graph build-

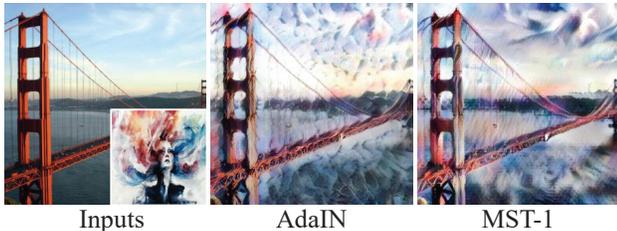


Figure 9: Feature transformation investigation.

ing. Here, we mainly investigate Euclidean distance and cosine distance (shown in Eq. (2)). As shown in Fig. 7, MST with Euclidean distance is affected by the huge background and may fail to transfer desired style patterns, leading to wash-out artifacts. This is mainly because there is no normalization for the deep features. As a result, the weight of style cluster center is proportional to its spatial proportion, weakening its semantic meaning. Instead, MST with cosine distance performs much better.

Discontinuity Preservation. In Fig. 8, we show the effectiveness of the smooth term $E_{smooth}(f)$ in Eq. (6). Specifically, we set λ as 0, 0.1, and 1 respectively. In real-world style transfer, people would like to smooth the facial area, as they do in the real photos. Here, we select one portrait to investigate how λ affects smoothness. When we set $\lambda = 0$, the energy function in Eq. (6) is minimized by only considering the data term. This would introduce some unpleasing artifacts in the facial area near edges and demonstrate the necessity of smooth term. However, large smooth term (e.g., $\lambda = 1$) would over-smooth the stylization results, decreasing the style diversity. A proper value of λ would not only keep better smoothness, but also preserve style diversity. We empirically set $\lambda = 0.1$ through the whole experiments.

Feature Transform. Here we set $K = 1$ in MST and compare with AdaIN [11] to show the effectiveness of WCT for feature transfer. As shown in Fig. 9, AdaIN produces some stroke artifacts in the smooth area. These artifacts may make the cloud unnatural. This is mainly because AdaIN uses the mean and variance of the whole content/style features. Instead, by using whitening and coloring in a more optimized way, our MST-1 achieves more natural stylization and cleaner smoothed area. As a result, we introduce whitening and coloring for feature transform.

5.2. Comparisons with Prior Arts

After investigating the effects of each component in our method, we turn to validate the effectiveness of our proposed MST. We compare with 7 state-of-the-art methods: method by Gatys *et al.* [8], CNNMRF [19], AdaIN [11], WCT [24], Deep feature reshuffle (DFR) [10], AvatarNet [36], and LST [23]. We obtain results using their official codes and default parameters, except for Gatys *et al.*¹.

Qualitative Comparisons. We show extensive compar-

isons² in Fig. 10. Gatys *et al.* [8] transfer style with iterative optimization, being likely to falling in local minimum (e.g., 1st and 3rd columns). AdaIN [11] often produces less desired artifacts in the smooth area and some halation around the edges (e.g., 1st, 5th, and 6th columns). CNNMRF [19] may suffer from distortion effects and not preserve content structure well. Due to the usage of higher-level deep feature (e.g., Conv_5_1), WCT [24] would generate distorted results, failing to preserve the main content structures (e.g., 1st and 2nd columns). DRF [10] reconstructs the results by using the style patches, which could also distort the content structure (e.g., 1st and 3rd columns). In some cases (e.g., 5th, 6th, and 7th columns), some tiny style patterns (e.g., the eyes in the flowers and tree) would be copied to the results, leading unpleasing stylizations. AvatarNet [36] would introduce some less desired style patterns in the smooth area (e.g., 1st column) and also copy some style patterns in the results (e.g., 6th and 7th columns). LST [23] could generate very good result in some cases (e.g., 6th column). However, it may suffer from wash-out artifacts (e.g., 3rd and 4th columns) and halation around the edges (e.g., 5th column). These compared methods mainly treat the style patterns as a whole, lacking distinctive ability to style patterns.

Instead, we treat style features as multimodal presentations in high-dimension space. We match each content feature to its most related style cluster and adaptively transfer features according to the content semantic information. These advantages help explain why MST generates clearer results (e.g., 1st, 3rd, 5th, and 7th columns), performs more semantic matching with style patterns (e.g., 2nd column), and alleviates wash-out artifacts (e.g., 4th column). Such superior results demonstrate the effectiveness of our MST.

Table 1: Percentage of the votes that each method received.

Method	Gatys	AdaIN	WCT	DFR	AvatarNet	MST
Perc./%	21.41	11.31	12.67	11.55	9.61	33.45

User Study. To further evaluate the 6 methods shown in Fig 10, we conduct a user study like [24]. We use 15 content images and 30 style images. For each method, we use the released codes and default parameters to generate 450 results. 20 content-style pairs are randomly selected for each user. For each style-content pair, we display the stylized results of 6 methods on a web-page in random order. Each user is asked to vote the one that he/she like the most. Finally, we collect 2,000 votes from 100 users and calculate the percentage of votes that each method received. The results are shown in Tab. 1, where our MST ($K=3$) obtains 33.45% of the total votes. It’s much higher than that of Gatys *et al.* [8], whose stylization results are usually thought to be high-quality. This user study result is consistent with the visual comparisons (in Fig. 10) and further demonstrate the superior performance of our MST.

¹We use code from <https://github.com/jcjohnson/neural-style>

²The fifth content is from <https://www.mordeo.org>



Figure 10: Visual comparison. MST ($K = 3$) and all compared methods use default parameters.

Table 2: Running time (s) comparisons.

Method	Gatys	AdaIN	WCT	DFR	AvatarNet
Time (s)	116.46	0.09	0.92	54.32	0.33
Method	MST-1	MST-2	MST-3	MST-4	MST-5
Time (s)	0.20	1.10	1.40	1.97	2.27

Efficiency. We further compare the running time of our methods with previous ones [8, 11, 24, 10, 36]. Tab. 2 gives the average time of each method on 100 image pairs with size of 512×512 . All the methods are tested on a PC with an Intel i7-6850K 3.6 GHz CPU and a Titan Xp GPU. Our MST with different K performs relatively faster than meth-

ods by Gatys *et al.* [8] and DFR [10]. Even using SVD in CPU, MST-1 is faster than AvatarNet [36] and WCT [24]. It should be noted that WCT conducts multi-level stylization, which costs much more time than that of MST-1. MST- K ($K > 1$) becomes much slower with larger K . This is mainly because our cluster operation is executed in CPU and consumes much more time. On the other hand, although MST with larger K would consume more time, its stylized results would be very robust. So, in general, we don't have to choose very large K , of which we give more details about the effects later.

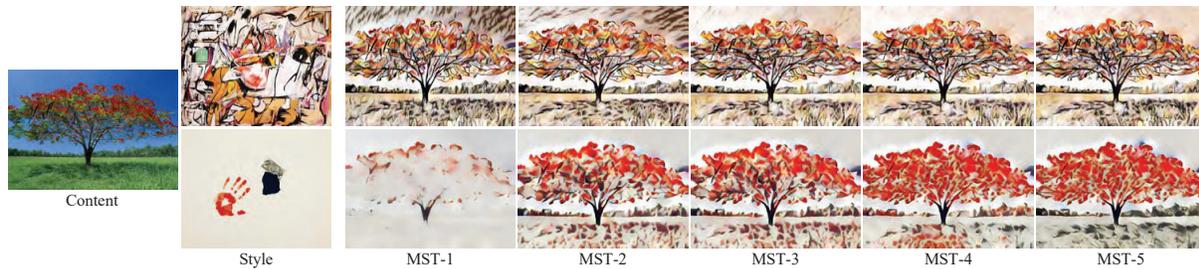


Figure 11: Style cluster number investigation. Same content image with complex and simple style images.



Figure 12: Multi-style transfer. MST treats patterns from different style images distinctively and transfers them adaptively.

5.3. Style Cluster Number

We investigate how style cluster number K affects the stylization in Fig. 11. When $K = 1$, our MST performs style transfer by taking the whole style features equally, resulting in either very complex (1st row) or simple (2nd row) stylizations. These results are not consistent with the content structures and lack flexibility, leading to unpleasant feelings to users. Instead, we can produce multiple results with different K . When we enlarge K with multimodal style representation, stylization results would either throw unnecessary style patterns (1st row) or introduce more matched style patterns (2nd row). The stylizations become more matched with the content structures. This is mainly because multimodal style representation allows distinctive and adaptive treatment for the style patterns. More important, MST reconstructs several stylization results with different K , providing multiple selections for the users.

5.4. Adaptive Multi-Style Transfer

Most previous style transfer methods enable style interpolation, which blends the content image with a set of weighted stylizations. However, we don't fix the weights for each style image, but adaptively interpolate the style patterns to the content. As shown in Fig. 12, the content image³ is stylized by two style images simultaneously. We use AdaIN [11] and WCT [24] for reference (because it's not strictly fair comparisons) by setting equal weight for each style image. In Fig. 12, AdaIN and WCT suffer from wash-out artifacts. While, our MST preserves the content structures well. MST transfers more portrait hair style to the cat body and more cloud style to the cat eyes and green leaves. Our adaptive multi-style transfer is also similar to spatial control in previous methods [11, 24]. But, they need additional manually designed mask as input, consuming more user efforts. Instead, MST automatically allows good matching between content and style features.

³It's from <https://wallpaperstream.com>



Figure 13: Generalization of MST to AdaIN [11].

5.5. Generalization of MST

We further investigate the generalization of our proposed MST to improve some existing style transfer methods. Here, we take the popular AdaIN [11] as an example. We apply style clustering and graph based style matching to AdaIN, which is then denoted as “AdaIN + MST- K ”. As shown in Fig. 13, AdaIN may distort some content structures (e.g., mouth) by switching the global mean and standard deviation between style and content features. When we cluster the style feature into K sub-sets and match them with content features via graph cuts, such a phenomenon can be obviously alleviated (see 3rd and 4th columns in Fig. 13). According to these observations and analyses, we can learn that our MST can be generalized and will benefit to some other existing style transfer methods.

6. Conclusion

We propose multimodal style representation to model the complex style distribution. We then formulate the style matching problem as an energy minimization one and solve it using our proposed graph based style matching. As a result, we propose multimodal style transfer to transform features in a multimodal way. We treat the style patterns distinctively, and also consider the semantic content structure and its matching with style patterns. We also investigate that MST can be generalized to some existing style transfer methods. We conduct extensive experiments to validate the effectiveness, robustness, and flexibility of MST.

Acknowledgements: This work was supported by Adobe Research funding.

References

- [1] Yuri Boykov and Daniel P Huttenlocher. A new bayesian framework for object recognition. In *CVPR*, 1999. 3
- [2] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *TPAMI*, 2001. 4
- [3] Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, and Gang Hua. Stereoscopic neural style transfer. In *CVPR*, 2018. 1
- [4] Tian Qi Chen and Mark Schmidt. Fast patch-based style transfer of arbitrary style. In *NIPS*, 2016. 1, 2
- [5] Alexei A Efros and Thomas K Leung. Texture synthesis by non-parametric sampling. In *ICCV*, 1999. 2
- [6] Michael Elad and Peyman Milanfar. Style transfer via texture synthesis. *TIP*, 2017. 2
- [7] Leon Gatys, Alexander S Ecker, and Matthias Bethge. Texture synthesis using convolutional neural networks. In *NIPS*, 2015. 2
- [8] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016. 1, 2, 6, 7
- [9] Dorothy M Greig, Bruce T Porteous, and Allan H Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 271–279, 1989. 3
- [10] Shuyang Gu, Congliang Chen, Jing Liao, and Lu Yuan. Arbitrary style transfer with deep feature reshuffle. In *CVPR*, 2018. 1, 2, 3, 6, 7
- [11] Xun Huang and Serge J Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017. 1, 2, 3, 5, 6, 7, 8
- [12] Yongcheng Jing, Yang Liu, Yezhou Yang, Zunlei Feng, Yizhou Yu, Dacheng Tao, and Mingli Song. Stroke controllable fast style transfer with adaptive receptive fields. In *ECCV*, 2018. 2
- [13] Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, Yizhou Yu, and Mingli Song. Neural style transfer: A review. *arXiv preprint arXiv:1705.04058*, 2017. 3
- [14] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 1, 2, 5
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [16] Vladimir Kolmogorov and Ramin Zabih. Multi-camera scene reconstruction via graph cuts. In *ECCV*, 2002. 3
- [17] Vivek Kwatra, Arno Schödl, Irfan Essa, Greg Turk, and Aaron Bobick. Graphcut textures: image and video synthesis using graph cuts. *TOG*, 2003. 3
- [18] Jan Eric Kyprianidis, John Collomosse, Tinghuai Wang, and Tobias Isenber. State of the” art”: A taxonomy of artistic stylization techniques for images and video. *TVCG*, 2013. 2
- [19] Chuan Li and Michael Wand. Combining markov random fields and convolutional neural networks for image synthesis. In *CVPR*, 2016. 1, 2, 5, 6
- [20] Kai Li, Martin Renqiang Min, and Yun Fu. Rethinking zero-shot learning: A conditional visual classification perspective. In *ICCV*, 2019. 2
- [21] Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Guided attention inference network. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2019. 2
- [22] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *ICCV*, 2019. 2
- [23] Xueting Li, Sifei Liu, Jan Kautz, and Ming-Hsuan Yang. Learning linear transformations for fast arbitrary style transfer. In *CVPR*, 2019. 1, 2, 3, 5, 6
- [24] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *NIPS*, 2017. 1, 2, 3, 4, 5, 6, 7, 8
- [25] Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. Demystifying neural style transfer. *arXiv preprint arXiv:1701.01036*, 2017. 5
- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5
- [27] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *CVPR*, 2017. 2
- [28] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 2008. 2, 3
- [29] Roey Mechrez, Itamar Talmi, and Lihi Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *ECCV*, 2018. 3
- [30] K Nichol. Painter by numbers, wikiart. <https://www.kaggle.com/c/painter-by-numbers>, 2016. 5
- [31] Douglas Reynolds. Gaussian mixture models. *Encyclopedia of biometrics*, 2015. 3
- [32] Sébastien Roy and Ingemar J Cox. A maximum-flow formulation of the n-camera stereo correspondence problem. In *ICCV*, 1998. 3
- [33] Manuel Ruder, Alexey Dosovitskiy, and Thomas Brox. Artistic style transfer for videos. In *German Conference on Pattern Recognition*, 2016. 2
- [34] Artsiom Sanakoyeu, Dmytro Kotovenko, Sabine Lang, and Björn Ommer. A style-aware content loss for real-time hd style transfer. In *ECCV*, 2018. 2
- [35] Falong Shen, Shuicheng Yan, and Gang Zeng. Neural style transfer via meta networks. In *CVPR*, 2018. 1, 2
- [36] Lu Sheng, Ziyi Lin, Jing Shao, and Xiaogang Wang. Avatar-net: Multi-scale zero-shot style transfer by feature decoration. In *CVPR*, 2018. 1, 2, 6, 7
- [37] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2, 3, 5
- [38] Olga Veksler. Image segmentation by nested cuts. In *CVPR*, 2000. 3
- [39] Xin Wang, Geoffrey Oxholm, Da Zhang, and Yuan-Fang Wang. Multimodal transfer: A hierarchical deep convolutional neural network for fast artistic style transfer. In *CVPR*, 2017. 1, 5
- [40] Yexun Zhang, Ya Zhang, and Wenbin Cai. Separating style and content for generalized style transfer. In *CVPR*, 2018. 1