

# A Closed-form Solution to Photorealistic Image Stylization

Yijun Li<sup>1</sup>, Ming-Yu Liu<sup>2</sup>, Xueting Li<sup>1</sup>, Ming-Hsuan Yang<sup>1,2</sup>, Jan Kautz<sup>2</sup>

<sup>1</sup>University of California, Merced <sup>2</sup>NVIDIA  
{yli62,xli75,mhyang}@ucmerced.edu {mingyul,jkautz}@nvidia.com


**Abstract.** Photorealistic image stylization concerns transferring style of a reference photo to a content photo with the constraint that the stylized photo should remain photorealistic. While several photorealistic image stylization methods exist, they tend to generate spatially inconsistent stylizations with noticeable artifacts. In this paper, we propose a method to address these issues. The proposed method consists of a stylization step and a smoothing step. While the stylization step transfers the style of the reference photo to the content photo, the smoothing step ensures spatially consistent stylizations. Each of the steps has a closed-form solution and can be computed efficiently. We conduct extensive experimental validations. The results show that the proposed method generates photorealistic stylization outputs that are more preferred by human subjects as compared to those by the competing methods while running much faster. Source code and additional results are available at <https://github.com/NVIDIA/FastPhotoStyle>.

**Keywords:** Image stylization, photorealism, closed-form solution.

## 1 Introduction

Photorealistic image stylization aims at changing style of a photo to that of a reference photo. For a faithful stylization, content of the photo should remain the same. Furthermore, the output photo should look like a real photo as it were captured by a camera. Figure 1 shows two photorealistic image stylization examples. In one example, we transfer a summery photo to a snowy one, while in the other, we transfer a day-time photo to a night-time photo.

Classical photorealistic stylization methods are mostly based on color/tone matching [1,2,3,4] and are often limited to specific scenarios (e.g., seasons [5] and headshot portraits [6]). Recently, Gatys et al. [7,8] show that the correlations between deep features encode the visual style of an image and propose an optimization-based method, the neural style transfer algorithm, for image stylization. While the method shows impressive performance for *artistic* stylization (converting images to paintings), it often introduces structural artifacts and distortions when applied to photorealistic image stylization as shown in Figure 1(c). In a follow-up work, Luan et al. [9] propose adding a regularization term to the optimization objective function of the neural style transfer algorithm



(a) Style (b) Content (c) Gatys et al. [8] (d) Luan et al. [9] (e) Ours

Fig. 1: Given a style photo (a) and a content photo (b), photorealistic image stylization aims at transferring style of the style photo to the content photo as shown in (c), (d) and (e). Comparing with existing methods [8,9], the output photos computed by our method are stylized more consistently and with fewer artifacts. Moreover, our method runs an order of magnitude faster.

for avoiding distortions in the stylization output. However, this often results in inconsistent stylizations in semantically uniform regions as shown in Figure 1(d). To address the issues, we propose a photorealistic image stylization method.

Our method consists of a stylization step and a smoothing step. Both have a closed-form solution<sup>1</sup> and can be computed efficiently. The stylization step is based on the whitening and coloring transform (WCT) [10], which stylizes images via feature projections. The WCT was designed for *artistic* stylization. Similar to the neural style transfer algorithm, it suffers from structural artifacts when applied to photorealistic image stylization. Our WCT-based stylization step resolves the issue by utilizing a novel network design for feature transform. The WCT-based stylization step alone may generate spatially inconsistent stylizations. We resolve this issue by the proposed smoothing step, which is based on a manifold ranking algorithm. We conduct extensive experimental validation with comparison to the state-of-the-art methods. User study results show that our method generates outputs with better stylization effects and fewer artifacts.

## 2 Related Work

Existing stylization methods can be classified into two categories: global and local. Global methods [1,2,11] achieve stylization through matching the means and variances of pixel colors [1] or their histograms [2]. Local methods [12,6,13,5,14] stylize images through finding dense correspondences between the content and style photos based on either low-level or high-level features. These approaches are slow in practice. Also, they are often developed for specific scenarios (e.g., day-time or season change).

<sup>1</sup> A closed-form solution means that the solution can be obtained in a fixed finite number of operations, including convolutions, max-pooling, whitening, etc.




Fig. 2: Our photorealistic image stylization method consists of two closed-form steps:  $\mathcal{F}_1$  and  $\mathcal{F}_2$ . While  $\mathcal{F}_1$  maps  $I_C$  to an intermediate image carrying the style of  $I_S$ ,  $\mathcal{F}_2$  removes noticeable artifacts, which produces a photorealistic output.

Gatys et al. [7,8] propose the neural style transfer algorithm for *artistic* stylization. The major step in the algorithm is to solve an optimization problem of matching the Gram matrices of deep features extracted from the content and style photos. A number of methods have been developed [15,16,17,18,19,20,21,10] to further improve its stylization performance and speed. However, these methods do not aim for preserving photorealism (see Figure 1(c)). Post-processing techniques [22,23] have been proposed to refine these results by matching the gradients between the input and output photos.

Photorealistic image stylization is related to the image-to-image translation problem [24,25,26,27,28,29,30,31] where the goal is to learn to translate an image from one domain to another. However, photorealistic image stylization does not require a training dataset of content and style images for learning the translation function. Photorealistic image stylization can be considered as a special kind of image-to-image translation. Not only can it be used to translate a photo to a different domain (e.g., from day to night-time) but also transfer style (e.g., extent of darkness) of a specific reference image to the content image.

Closest to our work is the method of Luan et al. [9]. It improves photorealism of stylization outputs computed by the neural style transfer algorithm [7,8] by incorporating a new loss term to the optimization objective, which has the effect of better preserving local structures in the content photo. However, it often generates inconsistent stylization with noticeable artifacts (Figure 1(d)). Moreover, the method is computationally expensive. Our proposed algorithm aims at efficient and effective photorealistic image stylization. We demonstrate that it performs favorably against Luan et al. [9] in terms of both quality and speed.

### 3 Photorealistic Image Stylization

Our photorealistic image stylization algorithm consists of two steps as illustrated in Figure 2. The first step is a stylization transform  $\mathcal{F}_1$  called PhotoWCT. Given a style photo  $I_S$ ,  $\mathcal{F}_1$  transfer the style of  $I_S$  to the content photo  $I_C$  while minimizing structural artifacts in the output image. Although  $\mathcal{F}_1$  can faithfully stylize  $I_C$ , it often generates inconsistent stylizations in semantically similar regions. Therefore, we use a photorealistic smoothing function  $\mathcal{F}_2$ , to eliminate these

artifacts. Our whole algorithm can be written as a two-step mapping function:

$$\mathcal{F}_2\left(\mathcal{F}_1(I_C, I_S), I_C\right), \quad (1)$$

In the following, we discuss the stylization and smoothing steps in details.

### 3.1 Stylization

The PhotoWCT is based on the WCT [10]. It utilizes a novel network design for achieving photorealistic image stylization. We briefly review the WCT below.

**WCT.** The WCT [10] formulates stylization as an image reconstruction problem with feature projections. To utilize WCT, an auto-encoder for general image reconstruction is first trained. Specifically, it uses the VGG-19 model [32] as the encoder  $\mathcal{E}$  (weights are kept fixed) and trains a decoder  $\mathcal{D}$  for reconstructing the input image. The decoder is symmetrical to the encoder and uses upsampling layers (pink blocks in Figure 3(a)) to enlarge the spatial resolutions of the feature maps. Once the auto-encoder is trained, a pair of projection functions are inserted at the network bottleneck to perform stylization through the whitening ( $P_C$ ) and coloring ( $P_S$ ) transforms. The key idea behind the WCT is to directly match feature correlations of the content image to those of the style image via the two projections. Specifically, given a pair of content image  $I_C$  and style image  $I_S$ , the WCT first extracts their vectorised VGG features  $H_C = \mathcal{E}(I_C)$  and  $H_S = \mathcal{E}(I_S)$ , and then transform the content feature  $H_C$  via

$$H_{CS} = P_S P_C H_C, \quad (2)$$

where  $P_C = E_C \Lambda_C^{-\frac{1}{2}} E_C^\top$ , and  $P_S = E_S \Lambda_S^{\frac{1}{2}} E_S^\top$ . Here  $\Lambda_C$  and  $\Lambda_S$  are the diagonal matrices with the eigenvalues of the covariance matrix  $H_C H_C^\top$  and  $H_S H_S^\top$  respectively. The matrices  $E_C$  and  $E_S$  are the corresponding orthonormal matrices of the eigenvectors, respectively. After the transformation, the correlations of transformed features match those of the style features, i.e.,  $H_{CS} H_{CS}^\top = H_S H_S^\top$ . Finally, the stylized image is obtained by directly feeding the transformed feature map into the decoder:  $Y = \mathcal{D}(H_{CS})$ . For better stylization performance, Li et al. [10] use a multi-level stylization strategy, which performs the WCT on the VGG features at different layers.

The WCT performs well for artistic image stylization. However it generates structural artifacts (e.g., distortions on object boundaries) for photorealistic image stylization (Figure 4(c)). The proposed PhotoWCT is designed to suppress these structural artifacts.

**PhotoWCT.** Our PhotoWCT design is motivated by the observation that the max-pooling operation in the WCT reduces spatial information in feature maps. Simply upsampling feature maps in the decoder fails to recover detailed structures of the input image. That is, we need to pass the lost spatial information to the decoder to facilitate reconstructing these fine details. Inspired by the success of




Fig. 3: The PhotoWCT and WCT share the same encoder architecture and projection steps. In the PhotoWCT, we replace the upsampling layers (pink) with unpooling layers (green). Note that the unpooling layer is used together with the pooling mask (yellow) which records *where* carries the *maximum* over each max pooling region in the corresponding pooling layer [33].

the unpooling layer [33,34,35] in preserving spatial information, the PhotoWCT replaces the upsampling layers in the WCT with unpooling layers. The PhotoWCT function is formulated as

$$Y = \mathcal{F}_1(I_C, I_S) = \overline{\mathcal{D}}(P_S P_C H_C), \quad (3)$$

where  $\overline{\mathcal{D}}$  is the decoder, which contains unpooling layers and is trained for image reconstruction. Figure 3 illustrates the network architecture difference between the WCT and the proposed PhotoWCT.

Figure 4(c) and (d) compare the stylization results of the WCT and PhotoWCT. As highlighted in close-ups, the straight lines along the building boundary in the content image becomes zigzagged in the WCT stylization result but remains straight in the PhotoWCT result. The PhotoWCT-stylized image has much fewer structural artifacts. We also perform a user study in the experiment section to quantitatively verify that the PhotoWCT generally leads to better stylization effects than the WCT.

### 3.2 Photorealistic Smoothing

The PhotoWCT-stylized result (Figure 4(d)) still looks less like a photo since semantically similar regions are often stylized inconsistently. As shown in Figure 4, when applying the PhotoWCT to stylize the day-time photo using the night-time photo, the stylized sky region would be more photorealistic if it were uniformly dark blue instead of partly dark and partly light blue. It is based on this observation, we employ the pixel affinities in the content photo to smooth the PhotoWCT-stylized result.

We aim to achieve two goals in the smoothing step. First, pixels with similar content in a local neighborhood should be stylized similarly. Second, the output should not deviate significantly from the PhotoWCT result in order to maintain the global stylization effects. We first represent all pixels as nodes in a graph and define an affinity matrix  $W = \{w_{ij}\} \in \mathbb{R}^{N \times N}$  ( $N$  is the number of pixels) to



Fig. 4: The stylization output generated by the PhotoWCT better preserves local structures in the content images, which is important for the image smoothing step as shown in (e) and (f).

describe pixel similarities. We define a smoothness term and a fitting term that model these two goals in the following optimization problem:

$$\operatorname{argmin}_r \frac{1}{2} \left( \sum_{i,j=1}^N w_{ij} \left\| \frac{r_i}{\sqrt{d_{ii}}} - \frac{r_j}{\sqrt{d_{jj}}} \right\|^2 + \lambda \sum_{i=1}^N \|r_i - y_i\|^2 \right), \quad (4)$$

where  $y_i$  is the pixel color in the PhotoWCT-stylized result  $Y$  and  $r_i$  is the pixel color in the desired smoothed output  $R$ . The variable  $d_{ii} = \sum_j w_{ij}$  is the diagonal element in the degree matrix  $D$  of  $W$ , i.e.,  $D = \operatorname{diag}\{d_{11}, d_{22}, \dots, d_{NN}\}$ . In (4),  $\lambda$  controls the balance of the two terms.

Our formulation is motivated by the graph-based ranking algorithms [36,37]. In the ranking algorithms,  $Y$  is a binary input where each element indicates if a specific item is a query ( $y_i = 1$  if  $y_i$  is a query and  $y_i = 0$  otherwise). The optimal solution  $R$  is the ranking values of all the items based on their pairwise affinities. In our method, we set  $Y$  as the PhotoWCT-stylized result. The optimal solution  $R$  is the smoothed version of  $Y$  based on the pairwise pixel affinities,




Fig. 5: Smoothing with different affinities. To refine the PhotoWCT result in (c), it is hard to find an optimal  $\sigma$  for the Gaussian Affinity that performs globally well as shown in (e)-(f). In contrast, using the Matting Affinity can simultaneously smooth different regions well as shown in (d).

which encourages consistent stylization within semantically similar regions. The above optimization problem is a simple quadratic problem with a closed-form solution, which is given by

$$R^* = (1 - \alpha)(I - \alpha S)^{-1}Y, \quad (5)$$

where  $I$  is the identity matrix,  $\alpha = \frac{1}{1+\lambda}$  and  $S$  is the normalized Laplacian matrix computed from  $I_C$ , i.e.,  $S = D^{-\frac{1}{2}}WD^{-\frac{1}{2}} \in \mathbb{R}^{N \times N}$ . As the constructed graph is often sparsely connected (i.e., most elements in  $W$  are zero), the inverse operation in (5) can be computed efficiently. With the closed-form solution, the smoothing step can be written as a function mapping given by:

$$R^* = \mathcal{F}_2(Y, I_C) = (1 - \alpha)(I - \alpha S)^{-1}Y. \quad (6)$$

**Affinity.** The affinity matrix  $W$  is computed using the content photo based on an 8-connected image graph assumption. While several choices of affinity metrics exist, a popular one is to define the affinity (denoted as GaussianAff) as  $w_{ij} = e^{-\|I_i - I_j\|^2 / \sigma^2}$  where  $I_i, I_j$  are the RGB values of adjacent pixels  $i, j$  and  $\sigma$  is a global scaling hyper-parameter [38]. However, it is difficult to determine the  $\sigma$  value in practice. It often results in either over-smoothing the entire photo (Figure 5(e)) or stylizing the photo inconsistently (Figure 5(f)). To avoid selecting one global scaling hyper-parameter, we resort to the matting affinity [39,40] (denoted as MattingAff) where the affinity between two pixels is based on means and variances of pixels in a local window. Figure 5(d) shows that the matting affinity is able to simultaneously smooth different regions well.

**WCT plus Smoothing.** We note that the smoothing step can also remove structural artifacts in the WCT as shown in Figure 4(e). However, it leads to unsatisfactory stylization. The main reason is that the content photo and the WCT result are severely misaligned due to spatial distortions. For example, a stylized pixel of the building in the WCT result may correspond to a pixel of the sky in the content photo. Consequently this causes wrong queries in  $Y$  for the smoothing step. This shows why we need to use the PhotoWCT to remove distortions first. Figure 4(f) shows that the combination of PhotoWCT and smoothing leads to better photorealism while still maintaining faithful stylization.

## 4 Experiments

In the section, we will first discuss the implementation details. We will then present visual and user study evaluation results. Finally, we will analyze various design choices and run-time of the proposed algorithm.

**Implementation details.** We use the layers from *conv1\_1* to *conv4\_1* in the VGG-19 network [32] for the encoder  $\mathcal{E}$ . The encoder weights are given by ImageNet-pretrained weights. The decoder  $\overline{\mathcal{D}}$  is the inverse of the encoder. We train the decoder by minimizing the sum of the  $L_2$  reconstruction loss and perceptual loss [17] using the Microsoft COCO dataset [41]. We adopt the multi-level stylization strategy proposed in the WCT [10] where we apply the PhotoWCT to VGG features in different layers.

Similar to the state-of-the-art methods [42,9], our algorithm can leverage semantic label maps for obtaining better stylization results when they are available. When performing PhotoWCT stylization, for each semantic label, we compute a pair of projection matrices  $P_C$  and  $P_S$  using the features from the image regions with the same label in the content and style photos, respectively. The pair is then used to stylize these image regions. With a semantic label map, content and style matching can be performed more accurately. We note that the proposed algorithm does not need precise semantic label maps for obtaining good stylization results. Finally, we also use the efficient filtering step described in Luan et al. [9] for post-processing.

**Visual comparison.** We compare the proposed algorithm to two categories of stylization algorithms: photorealistic and artistic. The evaluated photorealistic stylization algorithms include Reinhard et al. [1], Pitié et al. [2], and Luan et al. [9]. Both Reinhard et al. [1] and Pitié et al. [2] represent classical techniques that are based on color statistics matching, while Luan et al. [9] is based on neural style transfer [8]. On the other hand, the set of evaluated artistic stylization algorithms include Gatys et al.[8], Huang et al.[21], and the WCT [10]. They all utilize deep networks.




Fig. 6: Visual comparisons with photorealistic stylization methods. In addition to color transfer, our method also synthesizes patterns in the style photos (e.g., the dark cloud in the top example, the snow at the bottom example).




Fig. 7: Visual comparison with artistic stylization algorithms. Note the structural distortions on object boundaries (e.g., building) and detailed edges (e.g., sea, cloud) generated by the competing stylization methods.

Figure 6 shows visual results of the evaluated photorealistic stylization algorithms. Overall, the images generated by the proposed algorithm exhibit better stylization effects. While both Reinhard et al. [1] and Pitié et al. [2] change colors of the content photos, they fail to transfer the style. We argue that photorealistic stylization cannot be purely achieved via color transfer. It requires adding new patterns that represent the style photo to the content photo. For example, in the third example of Figure 6 (bottom), our algorithm not only changes the color of ground regions to white but also synthesizes the snow patterns as they appear in the style photo. The method of Luan et al. [9] achieves good stylization effects at first glance. However, a closer look reveals that the generated photos contain noticeable artifacts, e.g., the irregular brightness on buildings and trees. Several semantically similar regions are stylized inconsistently.

Figure 7 shows the visual comparison between the proposed algorithm and artistic stylization algorithms. Although the other evaluated algorithms are

Table 1: User preference: proposed vs. Luan et al. and proposed vs. Pitié et al.

	Luan et al. [9] / proposed	Pitié et al. [2] / proposed
Better stylization	36.9% / <b>63.1%</b>	44.8% / <b>55.2%</b>
Fewer artifacts	26.5% / <b>73.5%</b>	48.8% / <b>51.2%</b>

Table 2: User preference: proposed versus *artistic* stylization algorithms.

	Gatys et al. [8]	Huang et al. [21]	Li et al. [10]	proposed
Better stylization	19.2%	8.4%	16.0%	<b>56.4%</b>
Fewer artifacts	21.6%	6.0%	6.8%	<b>65.6%</b>

able to transfer the style well, they render noticeable structural artifacts and inconsistent stylizations across the images. In contrast, our method produces more photorealistic results.

**User studies.** We resort to user studies for performance evaluation since photorealistic image stylization is a highly subjective task. Our benchmark dataset consists of a set of 25 content–style pairs provided by Luan et al. [9]. We use the Amazon Mechanical Turk (AMT) platform for evaluation. In each question, we show the AMT workers a content–style pair and the stylized results from the evaluated algorithms displayed in random order. The AMT workers (lifetime Human Intelligent Task approval rate greater than 98%) are asked to select a stylized result based on the instructions. Each question is answered by 10 different workers. Hence, the performance score for each study is computed based on 250 questions. We compute the average number of times the images from an algorithm is selected, which is used as the preference score of the algorithm.

We conduct two user studies. In one study, we ask the AMT workers to select which stylized photo better carries the target style. In the other study, we ask the workers to select which stylized photo looks more like a real photo (containing fewer artifacts). Through the studies, we would like to answer which algorithm better stylizes content images and which renders better photorealistic outputs.

In Table 1, we compare the proposed algorithm to Luan et al. [9], which is the current state-of-the-art. The results show that 63.1% of the users prefer the stylization results generated by our algorithm and 73.5% regard our output photos as more photorealistic. We also compare our algorithm to the classical algorithm of Pitié et al. [2]. From Table 1, our results are as photorealistic as those computed by the classical algorithm (which simply performs color matching), and 55.2% of the users consider our stylization results better.

Table 2 compares our algorithm with the artistic stylization algorithms for user preference scores. We find our algorithm achieves a score of 56.4% and 65.6% for the stylization effect and photorealism, which are significantly better than the other algorithms. The artistic stylization algorithms do not perform well since they are not designed for the photorealistic stylization task.




Fig. 8: Visualization of effects of using different  $\lambda$  values in the photorealistic smoothing step. We show the edge maps of different stylization results (inset) at bottom and compare them with the edge map of the content in terms of the ODS and OIS metric (rightmost).




Fig. 9: Comparison between using our photorealistic smoothing step and other refinement methods (b)-(d).

**WCT versus PhotoWCT.** We compare the proposed algorithm with a variant where the PhotoWCT step is replaced by the WCT [10]. Again, we conduct two user studies on stylization effects and photorealism as described earlier. The result shows that the proposed algorithm is favored over its variant for better stylization 83.6% of the times and favored for better photorealism 83.2% of the times.

**Sensitivity analysis on  $\lambda$ .** In the photorealistic smoothing step, the  $\lambda$  balances between the smoothness term and fitting term in (4). A smaller  $\lambda$  renders smoother results, while a larger  $\lambda$  renders results that are more faithful to the queries (the PhotoWCT result). Figure 8 shows results of using different  $\lambda$  values. In general, decreasing  $\lambda$  helps remove artifacts and hence improves photorealism. However, if  $\lambda$  is too small, the output image tends to be over-smoothed. In order to find the optimal  $\lambda$ , we perform a grid search. We use the similarity between the boundary maps extracted from stylized and original content photos as the criteria since object boundaries should remain the same despite the stylization [44]. We employ the HED method [43] for boundary detection and use two standard boundary

detection metrics: ODS and OIS. A higher ODS or OIS score means a stylized photo better preserves the content in the original photo. The average scores over the benchmark dataset are shown on the rightmost of Figure 8. Based on the results, we use  $\lambda = 10^{-4}$  in all the experiments.

**Alternative smoothing techniques.** In Figure 9, we compare our photorealistic smoothing step with two alternative approaches. In the first approach, we use the PhotoWCT-stylized photo as the initial solution for solving the second optimization problem in the method of Luan et al. [9]. The result is shown in Figure 9(b). This approach leads to noticeable artifacts as the road color is distorted. In the second approach, we use the method of Mechrez et al. [23], which refines stylized results by matching the gradients in the output photo to those in the content photo. As shown in Figure 9(c), we find this approach performs well for removing structural distortions on boundaries but does not remove visual artifacts. In contrast, our method (Figure 9(d)) generates more photorealistic results with an efficient closed-form solution.

**Run-time.** In Table 3, we compare the run-time of the proposed algorithm to that of the state-of-the-art [9]. We note that while our algorithm has a closed-form solution, Luan et al. [9] rely on non-convex optimization. To stylize a photo, Luan et al. [9] solve two non-convex optimization problems sequentially where the solution to the first optimization problem is used as an initial solution to solve the second optimization problem. We report the total run-time required for obtaining the final stylization results. We resize the content images in the benchmark dataset to different sizes and report the average run-time for each image size. The experiment is conducted on a PC with an NVIDIA Titan X Pascal GPU. To stylize images of  $1024 \times 512$  resolution, our algorithm takes 13.16 seconds, which is 49 times faster than 650.45 seconds achieved by Luan et al. [9].

In Table 3, we also report the run-time of each step in our algorithm. We find the smoothing step takes most of the computation time, since it involves inverting the sparse matrix  $W$  in (5) using the LU decomposition. By employing efficient LU-decomposition algorithms developed for large sparse matrices, the complexity can be roughly determined by the number of non-zero entries in the matrices only. In our case, since each pixel is only connected to its neighbors (e.g.,  $3 \times 3$  window), the number of non-zero values in  $W$  grows linearly with the image size.

For further speed-up, we can approximate the smoothing step using guided image filtering [45], which can smooth the PhotoWCT output based on the content photo. We will refer to this version of our algorithm **approx**. Although approximating the smoothing step with guided image filtering results in slightly degraded performance as comparing to the original algorithm, it leads to a large speed gain as shown in Table 3. To stylize images of  $1024 \times 512$  resolution, **approx** only takes 0.64 seconds, which is 1,016 times faster than 650.45 seconds achieved by Luan et al. [9]. To quantify the performance degradation due to the approximation, we conduct additional user studies comparing the proposed algorithm and its approximation. We use the same evaluation protocol as described above. The results are shown in Table 4. In general, the stylization results rendered by

Table 3: Run-time comparison. We compute the average run time (in seconds) of the evaluated algorithms across various image resolutions.

Image resolution	Luan et al. [9]	proposed	PhotoWCT	smoothing	approx
256×128	79.61	0.96	0.40	0.56	0.41
512×256	186.52	2.95	0.42	2.53	0.47
768×384	380.82	7.05	0.53	6.52	0.55
1024×512	650.45	13.16	0.56	12.60	0.64

Table 4: User preference score comparison: comparing **approx** (the fast approximation of the proposed algorithm) to the proposed algorithm as well as other photorealistic stylization algorithms.

	proposed/approx	Luan et al. [9]/approx	Pitié et al. [2]/approx
Better stylization	<b>59.6%</b> / 40.4	36.4 / <b>63.6%</b>	46.0 / <b>54.0%</b>
Fewer artifacts	<b>52.8%</b> / 47.2	20.8 / <b>79.2%</b>	46.8 / <b>53.2%</b>



Fig. 10: Failure case. Both the proposed and other photorealistic stylization algorithms fail to transfer the flower patterns to the pot.

**approx** are less preferred by the users as compared to those generated by the full algorithm. However, the results from **approx** are still preferred over other methods in terms of both stylization effects and photorealism.

**Failure case.** Figure 10 shows a failure case where the proposed method fails to transfer the flower patterns in the style photo to the content photo. Similar limitations also apply to the other photorealistic stylization methods [2,1,9]. Since the proposed method uses the pixel affinity of the content in the photorealistic smoothing step, it favors a stylization output with smooth color transition on the pot surface as in the input photo.

## 5 Conclusions

We presented a novel fast photorealistic image stylization method. It consists of a stylization step and a photorealistic smoothing step. Both steps have efficient closed-form solutions. Experimental results show that our algorithm generates stylization outputs that are much more preferred by human subject as compared to those by the state-of-the-art, while running much faster.

## References

1. Reinhard, E., Ashikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. *IEEE Computer graphics and applications* **21**(5) (2001) 34–41 [1](#), [2](#), [8](#), [9](#), [10](#), [14](#)
2. Pitié, F., Kokaram, A.C., Dahyot, R.: N-dimensional probability density function transfer and its application to color transfer. In: *ICCV*. (2005) [1](#), [2](#), [8](#), [9](#), [10](#), [11](#), [14](#)
3. Sunkavalli, K., Johnson, M.K., Matusik, W., Pfister, H.: Multi-scale image harmonization. *ACM Transactions on Graphics* **29**(4) (2010) 125 [1](#)
4. Bae, S., Paris, S., Durand, F.: Two-scale tone management for photographic look. *ACM Transactions on Graphics* **25**(3) (2006) 637–645 [1](#)
5. Laffont, P.Y., Ren, Z., Tao, X., Qian, C., Hays, J.: Transient attributes for high-level understanding and editing of outdoor scenes. *ACM Transactions on Graphics* **33**(4) (2014) 149 [1](#), [2](#)
6. Shih, Y., Paris, S., Barnes, C., Freeman, W.T., Durand, F.: Style transfer for headshot portraits. In: *SIGGRAPH*. (2014) [1](#), [2](#)
7. Gatys, L.A., Ecker, A.S., Bethge, M.: Texture synthesis using convolutional neural networks. In: *NIPS*. (2015) [1](#), [3](#)
8. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: *CVPR*. (2016) [1](#), [2](#), [3](#), [8](#), [10](#), [11](#)
9. Luan, F., Paris, S., Shechtman, E., Bala, K.: Deep photo style transfer. In: *CVPR*. (2017) [1](#), [2](#), [3](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#), [14](#)
10. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Universal style transfer via feature transforms. In: *NIPS*. (2017) [2](#), [3](#), [4](#), [6](#), [8](#), [10](#), [11](#), [12](#)
11. Freedman, D., Kisilev, P.: Object-to-object color transfer: Optimal flows and smsp transformations. In: *CVPR*. (2010) [2](#)
12. Shih, Y., Paris, S., Durand, F., Freeman, W.T.: Data-driven hallucination of different times of day from a single outdoor photo. In: *SIGGRAPH*. (2013) [2](#)
13. Wu, F., Dong, W., Kong, Y., Mei, X., Paul, J.C., Zhang, X.: Content-based colour transfer. In: *Computer Graphics Forum*. Volume 32. (2013) 190–203 [2](#)
14. Tsai, Y.H., Shen, X., Lin, Z., Sunkavalli, K., Yang, M.H.: Sky is not the limit: Semantic-aware sky replacement. *ACM Transactions on Graphics* **35**(4) (2016) 149 [2](#)
15. Li, C., Wand, M.: Combining markov random fields and convolutional neural networks for image synthesis. In: *CVPR*. (2016) [3](#)
16. Ulyanov, D., Lebedev, V., Vedaldi, A., Lempitsky, V.: Texture networks: Feed-forward synthesis of textures and stylized images. In: *ICML*. (2016) [3](#)
17. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *ECCV*. (2016) [3](#), [8](#)
18. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Diversified texture synthesis with feed-forward networks. In: *CVPR*. (2017) [3](#)
19. Dumoulin, V., Shlens, J., Kudlur, M.: A learned representation for artistic style. In: *ICLR*. (2017) [3](#)
20. Ghiasi, G., Lee, H., Kudlur, M., Dumoulin, V., Shlens, J.: Exploring the structure of a real-time, arbitrary neural artistic stylization network. In: *BMVC*. (2017) [3](#)
21. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: *ICCV*. (2017) [3](#), [8](#), [10](#), [11](#)
22. Li, S., Xu, X., Nie, L., Chua, T.S.: Laplacian-steered neural style transfer. In: *ACM MM*. (2017) [3](#)
23. Mechrez, R., Shechtman, E., Zelnik-Manor, L.: Photorealistic style transfer with screened poisson equation. In: *BMVC*. (2017) [3](#), [12](#), [13](#)

24. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR. (2017) [3](#)
25. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: CVPR. (2018) [3](#)
26. Liu, M.Y., Tuzel, O.: Coupled generative adversarial networks. In: NIPS. (2016) [3](#)
27. Taigman, Y., Polyak, A., Wolf, L.: Unsupervised cross-domain image generation. In: ICLR. (2017) [3](#)
28. Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: CVPR. (2017) [3](#)
29. Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: NIPS. (2017) [3](#)
30. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV. (2017) [3](#)
31. Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: ECCV. (2018) [3](#)
32. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR. (2015) [4](#), [8](#)
33. Zhao, J., Mathieu, M., Goroshin, R., LeCun, Y.: Stacked what-where auto-encoders. In: ICLR Workshop. (2016) [5](#)
34. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: ECCV. (2014) [5](#)
35. Noh, H., Hong, S., Han, B.: Learning deconvolution network for semantic segmentation. In: ICCV. (2015) [5](#)
36. Zhou, D., Weston, J., Gretton, A., Bousquet, O., Schölkopf, B.: Ranking on data manifolds. In: NIPS. (2004) [6](#)
37. Yang, C., Zhang, L., Lu, H., Ruan, X., Yang, M.H.: Saliency detection via graph-based manifold ranking. In: CVPR. (2013) [6](#)
38. Shi, J., Malik, J.: Normalized cuts and image segmentation. PAMI **22**(8) (2000) 888–905 [7](#)
39. Levin, A., Lischinski, D., Weiss, Y.: A closed-form solution to natural image matting. PAMI **30**(2) (2008) 228–242 [7](#)
40. Zelnik-Manor, L., Perona, P.: Self-tuning spectral clustering. In: NIPS. (2005) [7](#)
41. 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. (2014) [8](#)
42. Gatys, L.A., Ecker, A.S., Bethge, M., Hertzmann, A., Shechtman, E.: Controlling perceptual factors in neural style transfer. In: CVPR. (2017) [8](#)
43. Xie, S., Tu, Z.: Holistically-nested edge detection. In: ICCV. (2015) [12](#)
44. Cutzu, F., Hammoud, R., Leykin, A.: Estimating the photorealism of images: Distinguishing paintings from photographs. In: CVPR. (2003) [12](#)
45. He, K., Sun, J., Tang, X.: Guided image filtering. PAMI **35**(6) (2013) 1397–1409 [13](#)