

Domain-Aware Universal Style Transfer

Kibeom Hong¹ Seogkyu Jeon¹ Huan Yang³ Jianlong Fu³ Hyeran Byun^{1,2*}

¹Department of Computer Science, Yonsei University

²Graduate school of AI, Yonsei University

³Microsoft Research

{cha2068, jone9312, hrbyun}@yonsei.ac.kr {huayan, jianf}@microsoft.com




Figure 1. Stylization results. Given input pairs (a fixed content and four references from photo and art domains), we compare our method (a) with both artistic transfer methods (b-f) and photo-realistic methods (g-h). While previous works fail to generate plausible results with references coming from the opposite domains, our model successfully conducts style transfer in both domains. (Best viewed in color.)

Abstract

Style transfer aims to reproduce content images with the styles from reference images. Existing universal style transfer methods successfully deliver arbitrary styles to original images either in an artistic or a photo-realistic way. However, the range of “arbitrary style” defined by existing works is bounded in the particular domain due to their structural limitation. Specifically, the degrees of content preservation and stylization are established according to a predefined target domain. As a result, both photo-realistic and artistic models have difficulty in performing the desired style transfer for the other domain. To overcome this limitation, we propose a unified architecture, **Domain-aware Style Transfer Networks (DSTN)** that transfer not only the style but also the property of domain (*i.e.*, domainness) from

a given reference image. To this end, we design a novel domainness indicator that captures the domainness value from the texture and structural features of reference images. Moreover, we introduce a unified framework with domain-aware skip connection to adaptively transfer the stroke and palette to the input contents guided by the domainness indicator. Our extensive experiments validate that our model produces better qualitative results and outperforms previous methods in terms of proxy metrics on both artistic and photo-realistic stylizations. All codes and pre-trained weights are available at [Kibeom-Hong/Domain-Aware-Style-Transfer](https://github.com/Kibeom-Hong/Domain-Aware-Style-Transfer).

1. Introduction

Recreating an image with the style of another image has been a long-standing research topic. As a seminal work, Gatys *et al.* [5] propose the neural style transfer with deep features extracted from pre-trained networks, *i.e.*, VGG-19

* Corresponding author

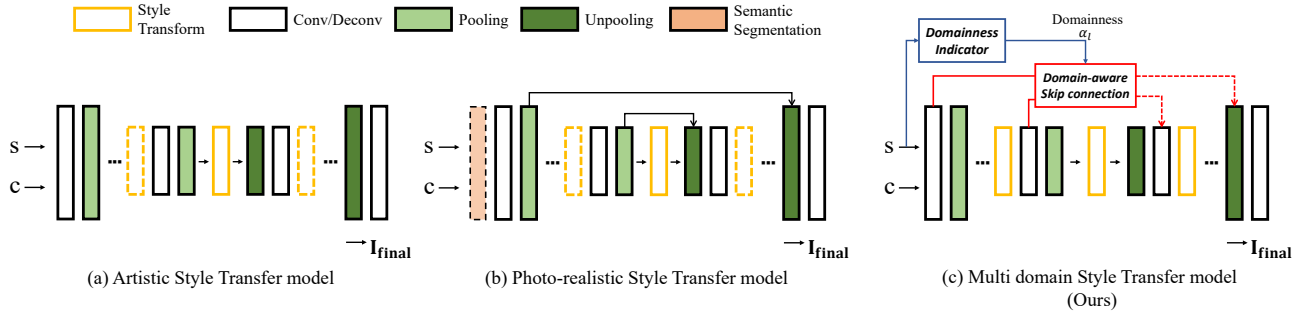


Figure 2. Conceptual differences between previous style transfer models and our model (DSTN). Artistic style transfer models (a) adopt encoder-decoder structures to transfer the global style pattern, but struggle to preserve the original contexts of input images. With the help of the skip connection, photo-realistic style transfer models (b) successfully preserve the structural information, but they lack the ability of transferring the delicate style. Different from the previous methods, our networks (c) deliver the structural information of content samples adaptively based on the domainness of a given style image. The dashed line of style transformation and semantic segmentation block indicate that they can be omitted. Note that given the content C and style S , each model outputs the final image (I_{final})

model [25]. After that, thanks to a series of previous studies, we can meet several stylized pictures of various painting styles of famous painters (e.g., Van Gogh).

Existing universal style transfer methods show the ability to deal with arbitrary reference images on either artistic or photo-realistic domain. On one hand, WCT [15] and AdaIN [10] transform the features of content images to match second-order statistics of reference features. Followed by these approaches, several artistic transfer studies [24, 33, 3, 28, 9] endeavor to transfer the global style pattern from reference images. Meanwhile, photo-realistic style transfer studies [19, 16, 31, 1] focus on preserving original structures while transferring target styles.

However, we observe the limitation that the meaning of “arbitrary” is restricted in a specific domain (i.e., either artistic or photo-realistic), and it comes from fundamental structural modifications for predefined target domains, as shown in Figure 2. In specific, artistic style transfer models have difficulty in maintaining clear details in the decoder because there is no clue directly coming from the content image. As a result, structural distortions of the content image occur when the artistic style transfer methods confront the photo-realistic reference images (Figure 1 (b)-(f)). On the other hand, photo-realistic style transfer models heavily constrain the transformation of input references with skip connections. Thus, they lack the ability to express delicate patterns (e.g., stroke pattern) of artistic references (Figure 1 (g)-(h)). In summary, existing arbitrary style transfer models generate undesired outputs when they receive samples of the other domain as references.

To overcome this limitation, we focus on capturing the domain characteristics from a given reference image and adjust the degree of the stylization and the structural preservation adaptively. For this purpose, we propose *Domain-aware Style Transfer Networks* (DSTN) which are a unified architecture composed of an auto-encoder with domain-aware skip connections and the domainness indicator. First,

we introduce the *domain-aware skip connection* to balance between content preservation and texture stylization for the domain-aware universal style transfer. Unlike the conventional skip connection which conveys intact structural details, the proposed skip connection block adjusts the transmission clarity of the high-frequency component from the stylized feature maps according to the domain properties.

To obtain the domain property (i.e., domainness) from a given reference image, we design the *domainness indicator*. Our novel indicator analyzes the characteristics of the domain by utilizing both the texture and structural feature maps extracted from different levels of our encoder. In order to predict a continuous domain factor with the range of $[0, 1]$, we propose to augment the intermediate space between art and photo with mixed samples.

With the proposed domain-aware architecture, DSTNs deliver semantic and structural information enabling artistic and photo-realistic style transfer, respectively. Consequently, DSTNs generate impeccable stylized results for arbitrary style, regardless of the target domain (Figure 1 (a)).

In experiments, we train our decoder and the domainness indicator with Microsoft COCO [17] and WikiArt [23] datasets. Qualitatively, we show that DSTNs are capable of generating plausible stylized results on both domains. Besides, through quantitative proxy metrics and user study, we demonstrate that our method outperforms previous methods on both photo-realistic and artistic style transfer.

Our contributions are three-fold: 1) We propose the novel end-to-end unified architecture, domain-aware style transfer networks (DSTN), for multi-domain style transfer. 2) With the proposed domainness indicator and domain-aware skip connections, we capture the domain characteristics and adaptively balance between the content preservation and the texture transformation. 3) DSTNs achieve the admirable performance given references of both photo-realistic and artistic domains, and outperform previous methods in terms of preservation and stylization.

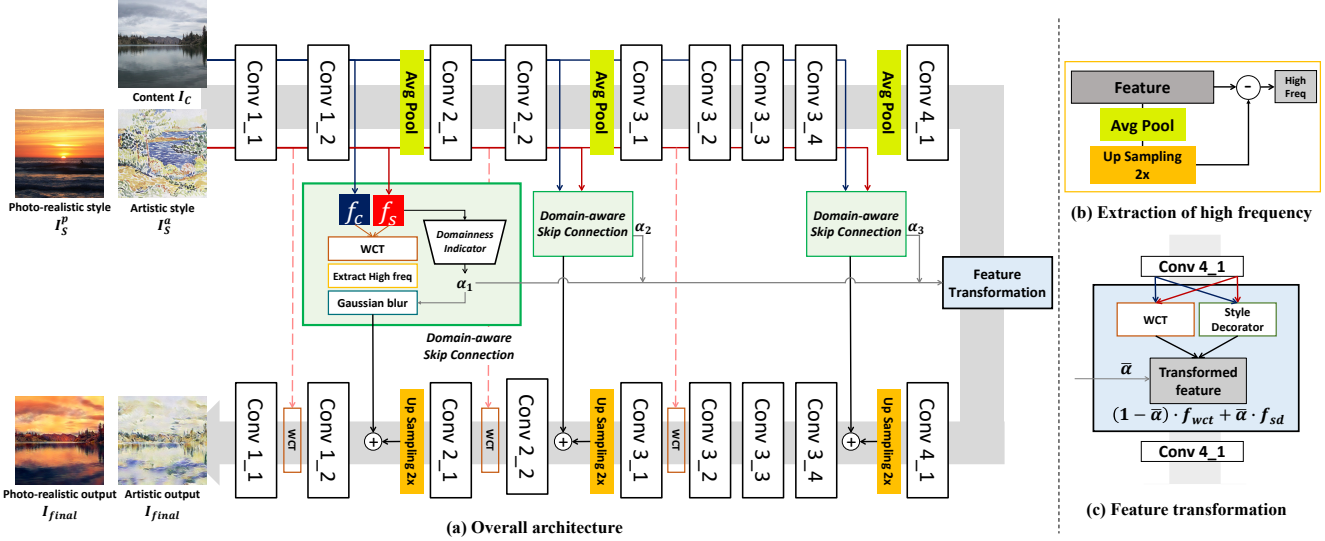


Figure 3. Overview of the domain-aware style transfer networks (DSTN). (b) shows the process of extracting the high-frequency components and (c) depicts the detail of feature transformation.

2. Related works

Since the pioneering work of Gatys *et al.* [5] opens up a new research area named *Neural Style Transfer*, many works have explored stylization methods with the power of representation ability of neural networks.

Artistic neural style transfer. Johnson *et al.* [11] and Ulyanov *et al.* [26] have directly trained feed-forward generative networks and achieve faster style transfer compared to image optimization methods. Still, these models need to be trained whenever they confront a new style. To alleviate this limitation, Li *et al.* [15] have proposed the whitening and coloring transformation (WCT) for arbitrary style transfer. Huang *et al.* [10] introduce the adaptive instance normalization (AdaIN) to simplify costly ZCA transformation by using the mean and standard deviation of the style features. Meanwhile, Chen *et al.* [2] and Sheng *et al.* [24] propose to swap the patches of content feature and normalized feature with the most correlated style features through the deconvolution methods, respectively. However, artistic transfer models tend to cause spatial distortion when generating photo-realistic styles.

Photo-realistic neural style transfer. Since artistic neural style transfer methods have difficulty in preserving original structures of content images, many attempts for photo-realistic style transfer have been made. Luan *et al.* [19] propose the deep photo style transfer with photorealism regularization term based on the Matting Laplacian [13]. Li *et al.* [16] replace the upsampling of the decoder with an unpooling layer and pass the index key of max pooling. In addition, they adopt additional post-processing steps, *e.g.*, smoothing and filtering. Yoo *et al.* [31] introduce the wavelet transform to preserve structural details. Yet, these

models have difficulty in expressing textures of art paintings since they heavily constrain the level of stylization.

Multi-domain neural style transfer. A few previous studies are capable of conducting style transfer on both artistic and photo-realistic domains. Li *et al.* [14] propose the linear style transfer (LST) with the learnable linear transformation matrix for universal style transfer. They additionally adopt the spatial propagation network [18] for photo-realistic style transfer. Chiu *et al.* [4] introduce an iterative transformation for stylized features with analytical gradient descent. However, these two models require additional information during inference, *i.e.*, the domain of a given reference, and they should adopt the auxiliary module or determine the number of iterative steps. In contrast, our DSTNs capture the domain-related characteristics from given reference inputs and achieve the multi-domain style transfer without any extra guidance.

3. Method

In this section, we describe the proposed domain-aware style transfer networks (DSTN) in detail. DSTNs consist of an auto-encoder with *domain-aware skip connections* and the *domaininess indicator*. Notably, our goal is transferring the style of an arbitrary reference image from artistic (I_S^a) or photo-realistic (I_S^p) domain to a content image (I_C).

The overview of our method is illustrated in Figure 3 (a). We first describe the proposed auto-encoder with domain-aware skip connection for multi-domain style transfer. Thereafter, we detail the proposed domaininess indicator which captures the property of domain from a given reference image.

3.1. Domain-aware Skip Connection

As discussed in earlier sections, the skip-connection is a fundamental difference that separates photo-realistic methods from artistic ones. It is advantageous in terms of structural preservation, but it constrains delicate artistic expressions. Based on this observation, we propose a domain-aware skip connection to conduct domain-aware universal style transfer on both photo-realistic and artistic domains. The domain-aware skip connection transforms the content feature f_c with the given reference feature f_s . Thereafter, we extract the high-frequency components of a stylized feature as shown in Figure 3 (b). We exploit this high-frequency components as a key of reconstruction, since it contains the structural information of I_c . Consequently, the decoder reconstructs an image with structural information coming from skip connections and texture information from the feature transformation block.

With this architecture design, we can adjust the level of structural preservation according to the domain properties of reference images. For artistic references, we deliberately blur the high-frequency components, so that our decoder reconstructs the image relying on deep texture features rather than structural details. We contort the high-frequency information with the Gaussian kernel of $\sigma = 16$ with the kernel size of $\lfloor \alpha_l \times 8 \rfloor + 1$. The α_l indicates the domain property obtained from the proposed domainness indicator which is detailed in the following section. As the blur increases in proportion to kernel size, a reference image with high α_l results in artistic stylization. In the opposite case of low α_l , the decoder utilizes the clear high-frequency components, thus resulting in photo-realistic results. Through ablation studies in section 4.3, we verify the effectiveness of each component in the proposed domain-aware skip connection.

3.2. Domainness Indicator

To capture the domain property (*i.e.*, domainness), we introduce the domainness indicator which exploits structural features as well as textural features. As shown in Figure 4, our indicator takes feature maps f_l from three different layers Φ_l (*i.e.*, Conv1_2, Conv2_2 and Conv3_4) of our encoder. Afterwards, we obtain the texture information by calculating a gram matrix [5] ($\mathcal{G}(f_l)$) of feature maps f_l . In addition, to encode structural information, we apply the channel-wise pooling on f_l through the 1×1 Conv layer. Lastly, we concatenate the texture and the structural information, which are forwarded through the weight-shared convolutional layers h to obtain the domainness α_l of each level l . These steps are formulated as:

$$\begin{aligned} f_l &= \Phi_l(I), l \in \{1, 2, 3\}, \\ F_l^{DI}(I) &= h([\psi_l(\mathcal{G}(f_l)) \odot f'_l]), \\ \alpha_l &= \sigma(F_l^{DI}(I)), \end{aligned} \quad (1)$$

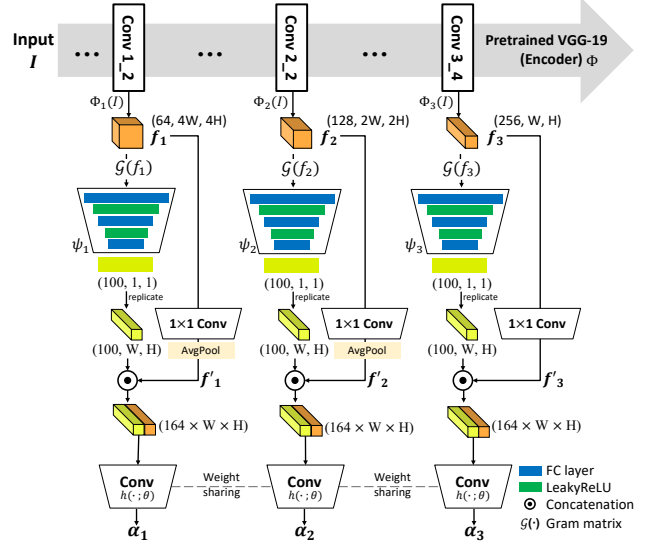


Figure 4. Overview of the proposed domainness indicator. It is designed to analyze the domainness based on multi-scale feature representations from VGG layers. We share the weights of the last convolutional layers for multi-level features.

where \odot denotes the channel-wise concatenation operator, ψ is fully-connected layers, f'_l is the channel-wise pooled feature maps of each level, and $\sigma(\cdot)$ is the sigmoid operation. We adopt the spatial average pooling for f_1 and f_2 to ensure the spatial sizes of features from different levels.

Furthermore, we adopt a domain adaptation method [6] to utilize the intermediate space between photo-realistic and artistic. We define the intermediate samples as follows:

$$I^{mix} = Mix(I^p, I^a, \beta), \quad (2)$$

where I^p , I^a and I^{mix} represent photo-realistic, artistic and mixed samples respectively. Mix is the Mixup [8] method and β is the strength of interpolation from the beta distribution as in [6].

To make the proposed indicator learn domainness, we adopt the binary cross entropy which is formulated as:

$$\begin{aligned} \mathcal{L}_{bce} &= \frac{1}{|L|} \sum_{l \in L} (\mathbb{E}_{I^a} [\log(F_l^{DI}(I^a))] \\ &\quad + \mathbb{E}_{I^p} [\log(1 - F_l^{DI}(I^p))]), \end{aligned} \quad (3)$$

where L depicts aforementioned three layers.

For mixed samples of the intermediate domain, we train the domainness indicator to embed their features between two domains according to β :

$$\begin{aligned} \mathcal{L}_{dlow} &= \frac{1}{|L|} \sum_{l \in L} ((1-\beta) \cdot dist(F_l^{DI}(I^p), F_l^{DI}(I^{mix})) \\ &\quad + \beta \cdot dist(F_l^{DI}(I^a), F_l^{DI}(I^{mix}))). \end{aligned} \quad (4)$$

The distance between features ($dist$) is calculated with $L1$ distance, as follows:

$$dist(f_A, f_B) = \|f_A - f_B\|_1 \quad (5)$$

Lastly, we obtain our total loss function for domainness indicator as:

$$\mathcal{L}_{DI} = \lambda_{bce}\mathcal{L}_{bce} + \lambda_{dlow}\mathcal{L}_{dlow}, \quad (6)$$

where λ_{bce} and λ_{dlow} are weighting factors of losses.

3.3. Overall Architecture and Training

We adopt a pre-trained VGG-19 (up to `conv4_1`) as the encoder and substitute its max pooling layers with average pooling layers. The decoder mirrors the encoder and all pooling layers are replaced with up-sampling layers. Then, we set up our domain-aware skip connection between `Conv1_2`, `Conv2_2`, `Conv3_4` layers of both the encoder and decoder. Inspired by Avatar-Net [24], we add short-cut connections links at `Conv1_1`, `Conv2_1`, `Conv3_1` layers for better stylization.

Following previous studies [15, 10, 24, 30, 21, 27, 29, 32], we apply transformation on the final feature of the encoder with the mean of $\alpha_l(\bar{\alpha})$, as shown in Figure 3 (c). In the transformation block, we control the weight of features between the universal stylized feature [15] (f_{wct}) which is good at preserving the global context, and the style decorated feature [24] (f_{sd}) which represents the local style pattern well.

The decoder is trained to perform reconstruction, *i.e.*, to invert the feature maps to an RGB image with the perceptual distance [11] and the contextual similarity [20] as follows:

$$\mathcal{L}_{rec} = \lambda_p \sum_{k=1}^4 \|\phi_k(I_r) - \phi_k(I_o)\|_2^2 + \lambda_{cx}CX(I_r, I_o), \quad (7)$$

where ϕ_k denotes the activation maps after `ReLU_k_1` layers of VGG-19, CX is the contextual similarity [20] between input images (I_o) and reconstructed images (I_r). λ_p and λ_{cx} indicate the hyper-parameters for weighing the perceptual distance and the contextual similarity respectively.

We replace the conventional $L2$ distance with the contextual similarity for the better quality of stylized outputs (especially to prevent the blur issue that occurs in artistic stylization). We investigate the effect of this modification on the loss function in Section 4.3. To further elevate the stylization performance, we exploit the multi-scale discriminator with the adversarial loss (\mathcal{L}_{adv}) [7]. The details of the multi-scale discriminator and the adversarial loss are described in the supplementary material.

Finally, we train our networks with an end-to-end manner, and the overall loss function is calculated as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{DI} + \mathcal{L}_{rec} + \lambda_{tv}\mathcal{L}_{tv} + \lambda_{adv}\mathcal{L}_{adv}, \quad (8)$$

where \mathcal{L}_{tv} represents the total variation loss for enforcing pixel-wise smoothness.

4. Experiments

Implementation details. We train DSTNs on MS-COCO [17] and WikiART [23] datasets, each containing roughly 80,000 images of real photos and artistic images respectively. We use the Adam [12] optimizer to train the domainness indicator as well as the decoder with a batch size of 6 and the learning rate is initially set to $1e-4$. Throughout the experiments, we set 256×256 as the default image resolution. Weighting factors of loss functions are set as: $\lambda_p = 0.1$, $\lambda_{cx} = 1$, $\lambda_{tv} = 1$, $\lambda_{dlow} = 1$, $\lambda_{adv} = 0.1$ and $\lambda_{bce} = 5$. Hyperparameters of \mathcal{L}_{dlow} are set to the same as [6]. Our code is implemented with PyTorch [22] and our model is trained on a single GTX 2080Ti. The code and the trained model will be publicly available online.

4.1. Qualitative results

We compare our networks with several state-of-the-art models qualitatively in terms of artistic and photo-realistic style transfer. Figure 5 shows the stylized results of photo-realistic (row 1,2) and artistic (row 3,4,5) references respectively. AvatarNet [24], AAMS [30], and Collaborative Distillation [27] show good artistic stylized outputs, but exhibit some distortions on photographic references (*e.g.*, the detail of leaves or skyline of mountains). On the other hand, PhotoWCT [16] and WCT² [31] successfully generate photo-realistic outputs, but they fail to express specific painting styles (*e.g.*, brush pattern or detailed texture) of artistic references. In other words, previous methods stylize content images according to the predefined target domain without considering the domain characteristics of the input references.

On the other hand, our DSTNs show plausible results on both domains (Figure 5 (a)). These samples demonstrate that the domainness indicator extracts the domain property well from given reference samples, and the decoder inverts features to images adaptively according to the given domainness. Concretely, DSTNs conduct the artistic style transfer while preserving semantic information via the domain-aware skip connection rather than costly attention modules [30, 21]. Moreover, structural details delivered from the proposed skip connection enables both intact content representation and stylization in photo-realistic style transfer. We note that our networks do not require any pre-processing or post-processing step.

Furthermore, we compare our method with previous multi-domain neural style transfer models [14, 4] as shown

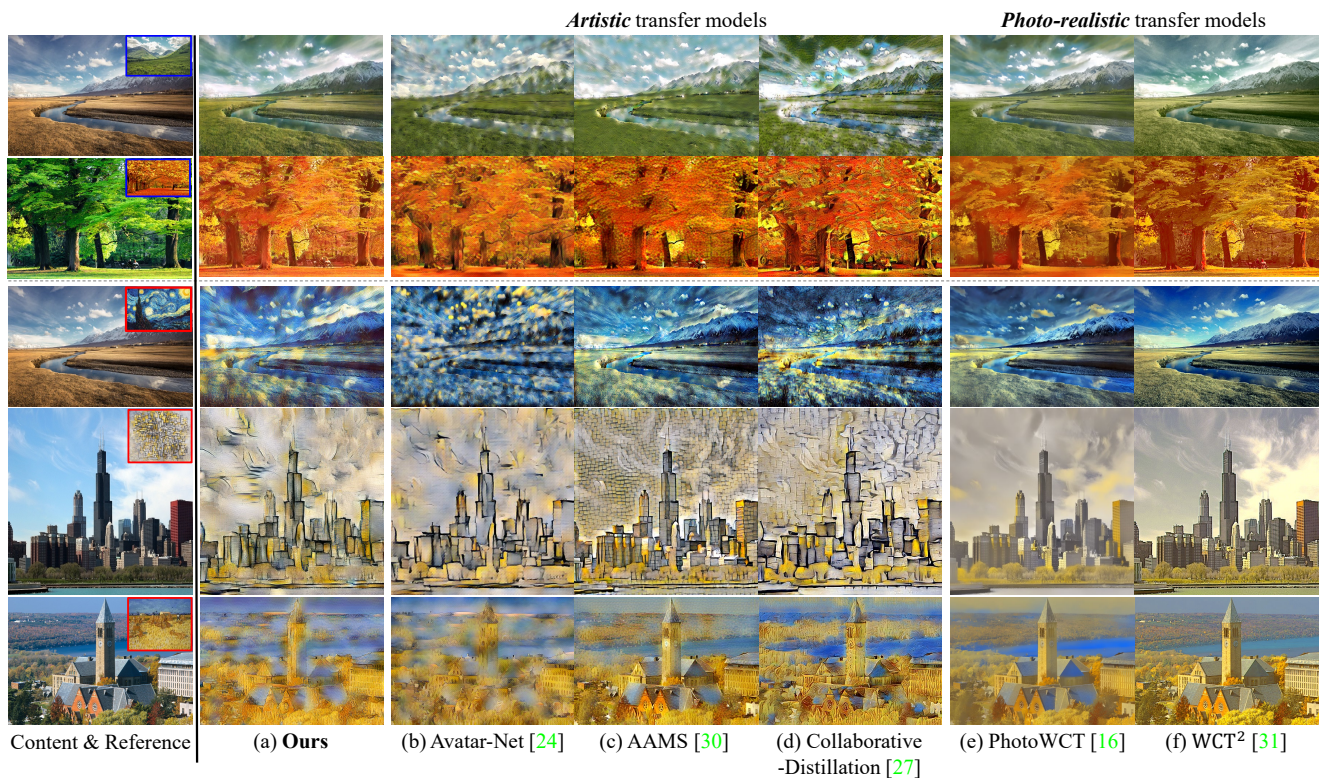


Figure 5. Qualitative comparisons with state-of-the-art models. The blue box indicates photo-realistic reference images and the red box indicates artistic ones. We depict the stylized results from existing artistic style transfer models (b-d) and photo-realistic ones (e-f). Previous models produce unsatisfactory results when they receive images from different domains, not the target of each model. The results of (a) demonstrate that our DSTNs produce both photo-realistic and artistic well regardless of the domain of given images.

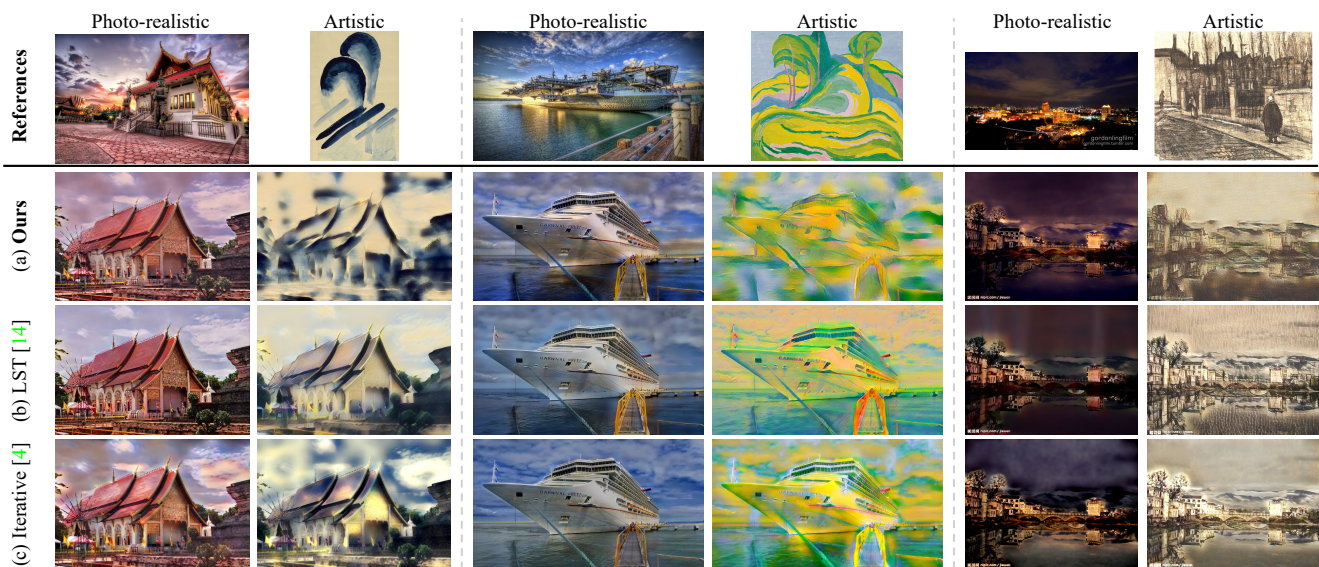


Figure 6. Comparisons with previous state-of-the-arts for multi-domain style transfer.

in Figure 6. We observe that DSTNs express not only the texture of the reference image but also the domain characteristics well. (More qualitative comparisons are included in our supplementary materials.)

4.2. Quantitative results

Statistics. To measure both photorealism and the degree of style representation, we employ two proxy metrics for

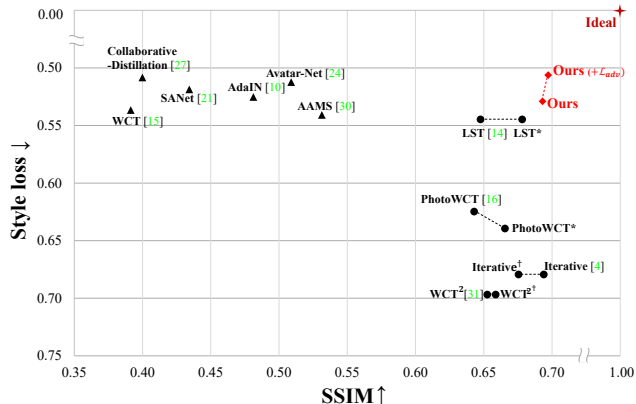


Figure 7. Quantitative results. SSIM index (higher is better) versus Style loss (lower is better). Ideal case for multi-domain style transfer is the top-right corner (red star). Dashed lines depict the gap before and after the post-processing steps or additional information, e.g., smoothing step (*) or segmentation maps (†).

structural preservation and stylization. Figure 7 shows the plot of SSIM (X-axis) versus style loss (Y-axis) as quantitative results. We calculate the structural similarity (SSIM) index between edge responses of content images and **photo-realistic** stylized outputs following WCT² [31]. In addition, we compute the style loss [5] between **artistic** stylized outputs and reference samples over five crop patches with 128×128 resolution. We use the 60 pairs of content and photo-realistic references provided by DPST [19] for quantitative evaluation. Also, we sample 60 artistic references randomly from the test split of WikiART dataset.

As shown in Figure 7, artistic style transfer models (left top) achieve decent style loss scores but they fail to preserve original structures with low SSIM values. On the other hand, photo-realistic transfer models (bottom right) conserve the original structure information well, but have trouble in expressing the style of a reference image. Though LST [14] achieves the good performance on both sides, we note that they require an additional spatial propagation network (SPN) [18] as well as domain information of references for transfer. In contrast, our DSTNs conduct the style transfer without any pre-processing (e.g., determining domains of given reference images or semantic segmentation) or post-processing (e.g., smoothing and filtering), and outperform other methods including the multi-domain models [14, 4] through a single training. Remarkably, our DSTNs accomplish the state-of-the-art performance on both domains with or even without the adversarial loss (\mathcal{L}_{adv}).

User study. As stylization is quite subjective, we conduct user study to better evaluate our method in terms of both photo-realistic and artistic criteria. We use 24 pairs of content and reference images, and compare our DSTNs to representative models of each domain, i.e., artistic [24, 30, 27] and photo-realistic [16, 31]. Besides, as competitors, we in-

Methods	Preservation ^p	Stylization ^a	Domainness ^o
Avatar-Net [24]	0.76%	19.66%	16.54%
AAMS [30]	5.30%	18.59%	20.51%
Colla-Distil [27]	1.52%	26.07%	19.62%
DSTN (ours)	92.42%	35.68%	43.33%
Photo-WCT [16]	13.26%	13.37%	14.87%
WCT ² [31]	39.01%	6.01%	20.64%
DSTN (ours)	47.73%	80.62%	64.49%
LST [14]	43.94%	27.71%	29.74%
Iterative [4]	26.89%	9.88%	19.49%
DSTN (ours)	29.17%	62.41%	50.77%

Table 1. The results of user studies. We ask the participants three questions of preference: 1) preservation of content with photo-realistic references (Preservation^p), 2) expression of style pattern with artistic references (Stylization^a) and 3) looks most similar to the domain of both photo-realistic and artistic references (Domainness^o).

Methods	256×256	512×512	1024×1024
AdaIN [10]	0.227	0.224	0.220
WCT [15]	0.640	1.415	4.111
Avatar-Net [24]	0.678	0.853	1.758
Photo-WCT [16]	2.579	10.870	OOM
WCT ² [31]	0.414	0.466	0.773
LST [14]	0.151	0.213	0.540
Iterative [4]	0.104	0.274	1.049
DSTN (ours)	0.195	0.396	1.441

Table 2. Execution time analysis in seconds. We compare under different resolutions. OOM denotes the out-of-memory error.

clude two previous models [14, 4] that enable multi-domain style transfer. The stylized results are provided to the users in random order with content and style images. In user study, we ask the participants questions about the preservation of content, the expression of texture and domainness. Consequently, we collect 2340 responses from 65 subjects and the results are shown in Table 1. Overall, it can be noticed that most of the participants favor our DSTNs over all evaluated methods.

Execution time analysis. Table 2 shows the execution time comparison with other methods under different resolutions. Results are estimated with a single NVIDIA RTX 2080Ti 11GB. Since DSTNs can transfer styles of both photo and artistic domains, we conduct 60 transfers for each domain and report the average time of total 120 transfers. Regardless of the resolutions, our method achieves the comparable execution time.

4.3. Ablation studies

In this section, we conduct several ablation studies on the domainness, domain-aware skip connections and reconstruction losses.




Figure 8. Ablation study on the effect of domainness α . More results of ablation studies are included in our supplementary materials.

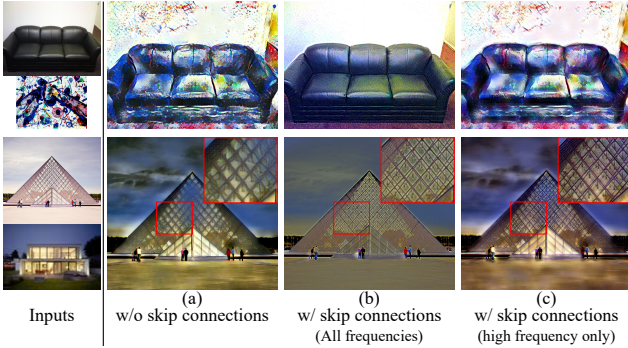


Figure 9. Ablation study of domain-aware skip connections. We compare each case with artistic and photo-realistic references. We also display the zoomed patch in the red box for better evaluation on photo-realism.

Effect of changing the domainness α . To analyze the relation between the domainness α and the stylized output, we intentionally set the α_l of every layer with the same value $\bar{\alpha}$ between the range of $[0, 1]$. As shown in Figure 8, outputs change smoothly between artistic and photo-realistic stylizations. Noticeably, DSTNs control the level of stylization and handle intermediate domains with a continuous domainness value.

Analysis on domain-aware skip connections. Our skip connections pass the high-frequency component of content images from the encoder to the decoder with the adaptive domainness value of reference images. In Figure 9, we qualitatively conduct ablation studies on domain-aware skip connections and high-frequency components. Without domain-aware skip connections, DSTNs fail to preserve structural details of the original image on photo-realistic style transfer, e.g., window frames in Figure 9 (a). Though establishing domain-aware skip connections on entire encoded features enables preserving the structure of original features, degrees of transformation on artistic references are insignificant since encoded features contain too abundant features for reconstruction as in Figure 9 (b). By passing and transforming high-frequency components only, we found a sweet-spot between the content preservation and the stylization as in Figure 9 (c).

Analysis of reconstruction loss. In DSTNs, the amount of structural clues coming from high-frequency components



(a) Content & Reference



(b) L2 distance



(c) Contextual similarity

Figure 10. Ablation study of the reconstruction loss, \mathcal{L}_{rec} .

is less for artistic images, thus their reconstruction results get blurry. Therefore, we employ the contextual similarity [20] instead of pixel-wise L_2 distance for reconstruction. As shown in Figure 10 (b), the stylized outputs from networks trained with L_2 distance are blurry and fail to express the detailed texture of the reference images. Contrarily, the proposed decoder trained with contextual similarity (CX) transfers the texture of oil painting on canvas effectively (Figure 10 (c)).

5. Conclusion

In this work, we proposed domain-aware style transfer networks (DSTN) for the domain-aware universal style transfer. We introduced the novel domainness indicator which captures the domain characteristics (*i.e.*, domainness) from the arbitrary references. Moreover, our domain-aware skip connection adjusts the clarity of transferred information by Gaussian blur in accordance with the domainness. As a result, DSTN generated admirable stylized outputs on both domains without any additional user-guidance. Qualitative and quantitative experiments verified that DSTNs achieve the state-of-the-art stylization performance in both photo-realistic and artistic domains without any pre-processing or post-processing.

6. Acknowledgements

This project was partly supported by the National Research Foundation of Korea grant funded by the Korea government (MSIT) (No. 2019R1A2C2003760) and the Institute for Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (No. 2020-0-01361: Artificial Intelligence Graduate School Program (YONSEI UNIVERSITY)). We sincerely appreciate all participants for the user study.

References

- [1] Jie An, Haoyi Xiong, Jun Huan, and Jiebo Luo. Ultrafast photorealistic style transfer via neural architecture search. *AAAI*, abs/1912.02398, 2020.
- [2] Tian Qi Chen and Mark Schmidt. Fast patch-based style transfer of arbitrary style. *arXiv preprint arXiv:1612.04337*, 2016.
- [3] Jiaxin Cheng, Ayush Jaiswal, Yue Wu, Pradeep Natarajan, and Prem Natarajan. Style-aware normalized loss for improving arbitrary style transfer. In *CVPR*, pages 134–143, June 2021.
- [4] Tai-Yin Chiu and Danna Gurari. Iterative feature transformation for fast and versatile universal style transfer. In *ECCV*, pages 169–184, 2020.
- [5] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, pages 2414–2423, 2016.
- [6] Rui Gong, Wen Li, Yuhua Chen, and Luc Van Gool. Dlow: Domain flow for adaptation and generalization. In *CVPR*, pages 2477–2486, 2019.
- [7] Ian J. Goodfellow, Jean Pouget-Abadie, M. Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.
- [8] Yann N. Dauphin Hongyi Zhang, Moustapha Cisse and David Lopez-Paz. mixup: Beyond empirical risk minimization. *ICLR*, 2018.
- [9] Zhiyuan Hu, Jia Jia, Bei Liu, Yaohua Bu, and Jianlong Fu. Aesthetic-aware image style transfer. *ACM MM*, 2020.
- [10] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, pages 1501–1510, 2017.
- [11] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711. Springer, 2016.
- [12] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- [13] Anat Levin, Dani Lischinski, and Yair Weiss. A closed-form solution to natural image matting. *IEEE transactions on pattern analysis and machine intelligence (PAMI)*, 30(2):228–242, 2007.
- [14] Xueting Li, Sifei Liu, Jan Kautz, and Ming-Hsuan Yang. Learning linear transformations for fast image and video style transfer. In *CVPR*, pages 3809–3817, 2019.
- [15] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *NeurIPS*, pages 386–396, 2017.
- [16] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *ECCV*, pages 453–468, 2018.
- [17] 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*, pages 740–755. Springer, 2014.
- [18] Sifei Liu, Shalini De Mello, Jinwei Gu, Guangyu Zhong, Ming-Hsuan Yang, and Jan Kautz. Learning affinity via spatial propagation networks. In *NeurIPS*, pages 1520–1530, 2017.
- [19] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *CVPR*, pages 4990–4998, 2017.
- [20] Roey Mechrez, Itamar Talmi, and Lihl Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *ECCV*, pages 768–783, 2018.
- [21] Dae Young Park and Kwang Hee Lee. Arbitrary style transfer with style-attentional networks. In *CVPR*, pages 5880–5888, 2019.
- [22] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NeurIPS Autodiff Workshop*, 2017.
- [23] Fred Phillips and Brandy Mackintosh. Wiki art gallery, inc.: A case for critical thinking. *Issues in Accounting Education*, 26(3):593–608, 2011.
- [24] Lu Sheng, Ziyi Lin, Jing Shao, and Xiaogang Wang. Avatar-net: Multi-scale zero-shot style transfer by feature decoration. In *CVPR*, pages 8242–8250, 2018.
- [25] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [26] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *CVPR*, pages 6924–6932, 2017.
- [27] Huan Wang, Yijun Li, Yuehai Wang, Haoji Hu, and Ming-Hsuan Yang. Collaborative distillation for ultra-resolution universal style transfer. In *CVPR*, pages 1860–1869, 2020.
- [28] Pei Wang, Yijun Li, and Nuno Vasconcelos. Rethinking and improving the robustness of image style transfer. In *CVPR*, pages 124–133, June 2021.
- [29] Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Bain-ing Guo. Learning texture transformer network for image super-resolution. In *CVPR*, pages 5790–5799, June 2020.
- [30] Yuan Yao, Jianqiang Ren, Xuansong Xie, Weidong Liu, Yong-Jin Liu, and Jun Wang. Attention-aware multi-stroke style transfer. In *CVPR*, pages 1467–1475, 2019.
- [31] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. In *ICCV*, pages 9036–9045, 2019.
- [32] Yanhong Zeng, Jianlong Fu, and Hongyang Chao. Learning joint spatial-temporal transformations for video inpainting. In *ECCV*, 2020.
- [33] Yulun Zhang, Chen Fang, Y. Wang, Zhaowen Wang, Zhe Lin, Yun Fu, and Jimei Yang. Multimodal style transfer via graph cuts. *ICCV*, pages 5942–5950, 2019.