

Painterly Image Harmonization in Dual Domains

Junyan Cao, Yan Hong, Li Niu*

MoE Key Lab of Artificial Intelligence, Shanghai Jiao Tong University
 {joy.c1, ustcnewly}@sjtu.edu.cn, yanHong.sjtu@gmail.com

Abstract

Image harmonization aims to produce visually harmonious composite images by adjusting the foreground appearance to be compatible with the background. When the composite image has photographic foreground and painterly background, the task is called painterly image harmonization. There are only few works on this task, which are either time-consuming or weak in generating well-harmonized results. In this work, we propose a novel painterly harmonization network consisting of a dual-domain generator and a dual-domain discriminator, which harmonizes the composite image in both spatial domain and frequency domain. The dual-domain generator performs harmonization by using AdaIN modules in the spatial domain and our proposed ResFFT modules in the frequency domain. The dual-domain discriminator attempts to distinguish the inharmonious patches based on the spatial feature and frequency feature of each patch, which can enhance the ability of generator in an adversarial manner. Extensive experiments on the benchmark dataset show the effectiveness of our method. Our code and model are available at <https://github.com/bcmi/PHDNet-Painterly-Image-Harmonization>.

1 Introduction

Image composition refers to cutting the foreground from one image and pasting it on another background image, producing a composite image. However, the foreground and background may have inconsistent color and illumination statistic, making the whole composite image inharmonious and unrealistic. Image harmonization (Lalonde and Efros 2007; Tsai et al. 2017; Cong et al. 2020) aims to adjust the foreground appearance to make it compatible with the background. In recent years, image harmonization has attracted growing research interest (Cong et al. 2021; Guo et al. 2021a; Hang et al. 2022). Besides combining foreground and background from photos, users may insert an object into a painting for creative painterly editing. This task is called painterly image harmonization, which has only received limited attention (Luan et al. 2018; Zhang, Wen, and Shi 2020; Peng, Wang, and Wang 2019). In particular, when a composite image is composed of photographic foreground object and painterly background image, painterly image harmonization aims to adjust the foreground style in the composite



Figure 1: Example of painterly image harmonization. From left to right are foreground object, background painting image, composite image, and harmonized image.

image to produce a harmonious image. Figure 1 shows an example of painterly image harmonization.

Existing painterly image harmonization methods can be divided into optimization-based methods (Luan et al. 2018; Zhang, Wen, and Shi 2020) and feed-forward methods (Peng, Wang, and Wang 2019). Optimization-based methods directly optimize the composite image to minimize the designed objective function. Specifically, the optimization-based methods (Luan et al. 2018; Zhang, Wen, and Shi 2020) employ a set of losses (e.g., content loss (Gatys, Ecker, and Bethge 2016), style loss (Huang and Belongie 2017), smoothness loss (Mahendran and Vedaldi 2015)) as the objective function. Then for each input composite image, they iteratively update its pixels and output the harmonized result, which does not rely on any training data. However, the optimization-based methods (Peng, Wang, and Wang 2019) are very time-consuming, which could not achieve real-time harmonization. Feed-forward methods pass the composite image through the network to produce a harmonized image. In particular, they train the network on the training set with a set of losses. However, the foregrounds are often not sufficiently stylized or not naturally blended into the background. Considering the demand of real-time application, we follow the research line of feed-forward method, which is also dominant in artistic style transfer (Huang and Belongie 2017; Park and Lee 2019; Liu et al. 2021).

In this work, we perform painterly image harmonization in two domains: spatial domain and frequency domain. Unlike previous works which only consider spatial domain (Luan et al. 2018; Zhang, Wen, and Shi 2020; Peng, Wang, and Wang 2019), we additionally explore frequency domain due to the following two concerns. Firstly, the convolution

*Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

operations in spatial domain have local receptive field, and lack the ability to capture long-range dependency (Wang et al. 2018). Meanwhile, the operations in frequency domain, *e.g.*, Fast Fourier Transform (*FFT*), have image-wise receptive field and thus could extract the global style of the whole image. Secondly, painterly image harmonization needs to transfer the style (*e.g.*, color, stroke, pattern, texture) of background image to the composite foreground. The background paintings often have periodic textures and patterns which appear regularly, which could be well captured in the frequency domain.

Motivated by the advantage of frequency domain, we propose a novel dual-domain network named **PHDNet** to accomplish Painterly image **H**armonization in **D**ual domains. Our PHDNet consists of a dual-domain generator and a dual-domain discriminator. Specifically, our generator is built upon UNet (Ronneberger, Fischer, and Brox 2015). We harmonize multi-scale encoder feature maps in the spatial domain and frequency domain sequentially in the skip connections. We first apply Adaptive Instance Normalization (AdaIN) (Huang and Belongie 2017) to align the statistics (*i.e.*, mean and variance) of composite foreground feature map with background feature map in the spatial domain. Then, we convert the normalized feature map to frequency feature map and apply our proposed ResFFT module to harmonize the frequency feature map in the frequency domain. For our dual-domain discriminator, we divide the composite image into different patches including foreground patches and background patches. We extract the spatial domain feature and frequency domain feature for each patch. Based on the dual-domain patch features, the discriminator strives to distinguish the foreground patches from background patches, while the generator attempts to fool the discriminator. The dual-domain discriminator can promote the harmonization ability of dual-domain generator in an adversarial manner, so that the foreground in the harmonized image is inseparable from the background and appears to exist in the original painting. We conduct extensive experiments to verify the effectiveness of our proposed dual-domain network. Our contributions are summarized as follows,

- To the best of our knowledge, we are the first to introduce frequency domain knowledge into painterly harmonization task.
- We accomplish painterly image harmonization in dual domains, and design a dual-domain network PHDNet. Our PHDNet contains a dual-domain generator with a novel ResFFT module to harmonize the composite image in both spatial and frequency domain, and a novel dual-domain discriminator to distinguish the inharmonious region in both spatial and frequency domain.
- Comprehensive experiments demonstrate that our PHDNet could produce more harmonious results with consistent style and intact content than previous methods.

2 Related Work

2.1 Image Harmonization

Image harmonization aims to harmonize a composite image with both foreground and background from photos. Early

traditional image harmonization methods (Song et al. 2020; Xue et al. 2012; Sunkavalli et al. 2010; Lalonde and Efros 2007) focused on manipulating the low-level statistics (*e.g.*, color, gradient, histogram) of foreground to match those of background. Then, unsupervised deep learning methods (Zhu et al. 2015) adopted adversarial learning to enforce the harmonized images to be indistinguishable from real images. More recently, abundant supervised deep learning methods (Tsai et al. 2017; Sofiuk, Popenova, and Konushin 2021; Cong et al. 2022) leveraged paired training data to train harmonization models. To name a few, Cun and Pun (2020) and Hao, Iizuka, and Fukui (2020) designed various attention mechanisms. Cong et al. (2020) and Cong et al. (2021) formulated image harmonization as domain translation task by treating foreground and background as two domains. Guo et al. (2021b) and Guo et al. (2021a) decomposed an image to reflectance map and illumination map, followed by adjusting the foreground illumination map. Ling et al. (2021) and Hang et al. (2022) introduced AdaIN (Huang and Belongie 2017) in style transfer to image harmonization. The above supervised image harmonization methods require ground-truth images as supervision, which are not applicable to our task.

2.2 Painterly Image Harmonization

When inserting an object into a painting, painterly image harmonization aims to transfer the background style to the foreground while retaining the foreground content, making the composite image as natural as possible. As far as we are concerned, there are only few works on painterly image harmonization. Luan et al. (2018) proposed to migrate relevant statistics of neural responses to the inserted object, ensuring both spatial and inter-scale statistical consistency. Zhang, Wen, and Shi (2020) developed a novel Poisson gradient loss jointly optimized with content and style loss. Peng, Wang, and Wang (2019) employed AdaIN to manipulate the foreground feature map, together with global and local discriminators for adversarial learning. All these methods only considered spatial domain, while we perform harmonization in both spatial domain and frequency domain.

2.3 Artistic Style Transfer

The goal of artistic style transfer is stylizing a content image according to the provided style image. The existing style transfer methods can also be divided into optimization-based methods and feed-forward methods. The optimization-based methods (Gatys, Ecker, and Bethge 2016; Li et al. 2017b; Kolkin, Salavon, and Shakhnarovich 2019; Du 2020) proposed to optimize over the content image to match its style with style image. The feed-forward methods combined the content of content image and the style of style image to produce a stylized image, during which the style-relevant statistics (*e.g.*, mean, variance) between foreground features and background features are matched in the network. According to global matching and local matching (matching corresponding regions), the feed-forward methods can be further divided into global transformation methods (Huang and Belongie 2017; Li et al. 2017a, 2018) and local transformation methods (Park and Lee 2019; Liu et al. 2021; Huo et al.

2021; Deng et al. 2022). Different from the above methods which stylize the entire content image, painterly image harmonization needs to consider the location of inserted object and harmonize it accordingly, achieving the goal that the object appears to be present in the original painting.

2.4 Frequency Domain Learning

Frequency domain information has been exploited in deep learning based methods for myriads of computer vision tasks, due to its enticing properties (*e.g.*, large receptive field, high and low frequency separation). For instance, a few works (Xu et al. 2020; Roy et al. 2021; Shen et al. 2021) converted the input image or output mask of network to frequency domain. Similarly, Suvorov et al. (2022) and Mardani et al. (2020) converted the intermediate features in the network to frequency domain, and processed the frequency features to achieve the goals of different tasks. By decomposing an image to low-frequency part and high-frequency part, some recent works (Bansal, Sheikh, and Ramanan 2018; Yang and Soatto 2020; Yu et al. 2021; Cai et al. 2021) proposed to manipulate the structural information and detailed information separately. In this work, we make the first attempt to introduce frequency domain into painterly image harmonization.

3 Our Method

The architecture of our PHDNet is shown in Figure 2. A composite image I_c is obtained by pasting foreground object I_c^f on a complete background painting I_b , and we use a foreground mask M to indicate the foreground region. Our goal is to train a model that transfers the style from I_b to I_c^f while keeping the content of I_c^f , generating a harmonized image \tilde{I}_c .

Our PHDNet consists of a dual-domain generator and a dual-domain discriminator, under the adversarial learning framework (Goodfellow et al. 2014). As demonstrated in Figure 2, the dual-domain generator G takes in I_c and I_b , and adjusts the style of I_c^f in both spatial domain and frequency domain. We also employ a dual-domain discriminator D , which predicts an inharmonious mask to distinguish the foreground patches from the background patches. The discriminator D is used to strengthen the generator G in an adversarial manner. Next, we will detail our dual-domain generator in Section 3.1 and dual-domain discriminator in Section 3.2.

3.1 Dual-Domain Generator

We employ the encoder-decoder architecture in (Huang and Belongie 2017) as our backbone, in which the encoder is pretrained VGG-19 network (Simonyan and Zisserman 2015) and the decoder is a symmetric structure of encoder. Note that we only use the first few layers (up to $ReLU-4_1$) of VGG-19 as our encoder, and freeze them to extract multi-scale encoder features. Following (Ronneberger, Fischer, and Brox 2015), we add skip connections on $ReLU-1_1$, $ReLU-2_1$, and $ReLU-3_1$ layers of the encoder. By feeding I_c and I_b into the encoder respectively, we could get the feature map F_{gc}^l and F_{gb}^l extracted by the l -th layer

($l \in \{1, 2, 3, 4\}$) of encoder. The four encoder layers contain $ReLU-1_1$, $ReLU-2_1$, and $ReLU-3_1$, and $ReLU-4_1$ (bottleneck). Then for the l -th layer, we feed F_{gc}^l and F_{gb}^l jointly with a downsampled mask \bar{M}^l into the AdaIN module followed by our ResFFT module, aiming to transfer the style from I_b to I_c^f in both spatial domain and frequency domain. Detailed architectures of these two modules will be introduced later. The harmonized encoder features are taken as the input of decoder or concatenated with decoder features via skip connection. At the end of decoder, we insert a blending layer similar to (Sofiiuk, Popenova, and Konushin 2021), which takes the concatenation of the decoder feature map and mask M as input, producing a soft mask \tilde{M} for the final blending.

AdaIN Module Firstly, we apply AdaIN (Huang and Belongie 2017) in the spatial domain. As stated above, the input of AdaIN module contains three parts: the foreground mask, the encoder feature maps of both composite image and background image.

Inspired by (Huang and Belongie 2017; Ling et al. 2021), for the l -th layer of VGG-19 encoder, we pass F_{gc}^l and F_{gb}^l jointly with \bar{M}^l through the AdaIN module in Figure 2, aiming to align the channel-wise mean and standard deviation of the foreground region of F_{gc}^l to those of the whole region of F_{gb}^l . The process could be expressed as

$$F_{gs}^l = \left(\sigma(F_{gb}^l) \frac{F_{gc}^l - \mu(F_{gc}^l \circ \bar{M}^l)}{\sigma(F_{gc}^l \circ \bar{M}^l)} + \mu(F_{gb}^l) \right) \circ \bar{M}^l + F_{gc}^l \circ (1 - \bar{M}^l), \quad (1)$$

where \circ means element-wise multiplication, $\mu(\cdot)$ and $\sigma(\cdot)$ denote the formulas to compute the mean and standard deviation of the feature map within the masked region (see (Huang and Belongie 2017; Ling et al. 2021) for details).

ResFFT Module Then we feed the normalized feature map F_{gs}^l into our ResFFT module for harmonization in the frequency domain. Following (Suvorov et al. 2022), we apply *Real FFT* to feature map F_{gs}^l with size $h^l \times w^l \times c_g^l$ and drop the redundant negative frequency terms due to the symmetric property, leading to the frequency feature map. The obtained frequency feature map is in the complex form with two parts, *i.e.*, real and imaginary part, both of which have the size $h^l \times \frac{w^l}{2} \times c_g^l$. We concatenate two parts channel-wisely and obtain the frequency feature map F_{gf}^l with size $h^l \times \frac{w^l}{2} \times 2c_g^l$.

Then we pass frequency feature map F_{gf}^l through the residual block (He et al. 2016). In the residual block, we learn the residual and add it to the input frequency feature map. Intuitively, we hope that the learned residual could harmonize the frequency feature map, *e.g.*, restoring the missing or corrupted texture and pattern within the foreground region in the frequency domain. Through the residual block, we get the harmonized frequency feature map \hat{F}_{gf}^l . Finally we convert \hat{F}_{gf}^l back to the spatial domain. In detail, we first

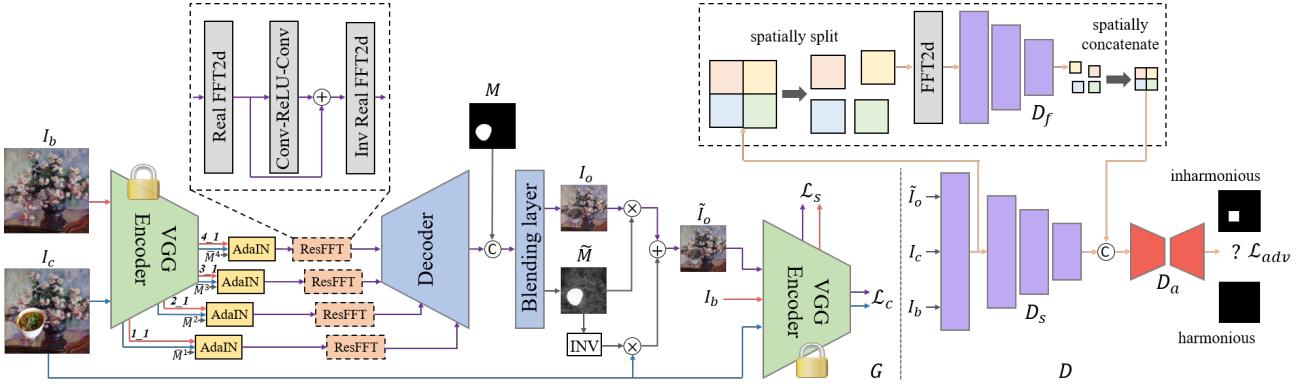


Figure 2: The architecture of our PHDNet, which consists of a dual-domain generator G and a dual-domain discriminator D . Pretrained VGG encoder is freezed. ‘INV’ means ‘inverse’.

convert \hat{F}_{gf}^l to complex form and then apply *inverse Real FFT* to get the spatial feature map \hat{F}_{gs}^l with size $h^l \times w^l \times c_g^l$.

After AdaIN module and ResFFT module, we obtain the harmonized feature map \hat{F}_{gs}^l , which is delivered to the decoder to generate the coarse output I_o . Then we blend I_o with the composite image I_c using the soft mask \tilde{M} , producing the harmonized image \tilde{I}_o :

$$\tilde{I}_o = I_o \circ \tilde{M} + I_c \circ (1 - \tilde{M}), \quad (2)$$

where \tilde{M} is produced by the blending layer as mentioned above.

To match multi-scale style statistics between the background image and the foreground of harmonized image, we employ the style loss in (Huang and Belongie 2017), which could be expressed as

$$\begin{aligned} \mathcal{L}_s = \sum_{l=1}^L \left\| \mu \left(\phi^l(\tilde{I}_o) \circ \bar{M}^l \right) - \mu \left(\phi^l(I_b) \right) \right\|^2 + \\ \sum_{l=1}^L \left\| \sigma \left(\phi^l(\tilde{I}_o) \circ \bar{M}^l \right) - \sigma \left(\phi^l(I_b) \right) \right\|^2, \end{aligned} \quad (3)$$

where each $\phi^l, l \in \{1, 2, 3, 4\}$ denotes the l -th *ReLU-l*-1 layer in VGG-19 encoder.

Besides, we utilize a content loss (Gatys, Ecker, and Bethge 2016) to ensure that the content of \tilde{I}_o is close to that of I_c :

$$\mathcal{L}_c = \left\| \phi^4(\tilde{I}_o) - \phi^4(I_c) \right\|^2. \quad (4)$$

3.2 Dual-Domain Discriminator

To improve the quality of harmonized image \tilde{I}_o , we resort to adversarial learning and design a dual-domain discriminator D . Given an input image uniformly split into $n \times n$ patches, D contains an encoder with spatial (*resp.*, frequency) branch D_s (*resp.*, D_f) to extract the spatial (*resp.*, frequency) feature for each patch, followed by a light-weighted auto-encoder D_a to identify the inharmonious patch. Detailed architectures could be found in the Supplementary.

As shown in Figure 2, given an input image I , we pass it through the spatial branch D_s to get the bottleneck feature map F_{ds} with size $n \times n \times c_{ds}$, in which each pixel-wise feature vector in F_{ds} is deemed as the spatial feature vector for one patch.

Then we choose one intermediate feature map F_{dm} in D_s and derive the frequency feature for each patch. Supposing that F_{dm} has size $m \times m \times c_{dm}$, we uniformly divide F_{dm} into $n \times n$ non-overlapped patches with patch size being $(\frac{m}{n}) \times (\frac{m}{n}) \times c_{dm}$. We denote the (i, j) -th patch in F_{dm} as $F_{dm}^{i,j}$, in which $i, j \in [1, n]$. Similar to the ResFFT module in Section 3.1, we apply *FFT* to each patch separately and convert it to frequency domain. In particular, for $F_{dm}^{i,j}$, we obtain the converted frequency feature map $F_{df}^{i,j}$ containing the real part and the imaginary part, both of which have the size $(\frac{m}{n}) \times (\frac{m}{n}) \times c_{dm}$. Note that we use both positive and negative frequency terms here for regular feature map size. Then we concatenate the real and imaginary parts of $F_{df}^{i,j}$ channel-wisely and feed it into D_f to get a c_{df} -dim frequency feature vector $\hat{f}_{df}^{i,j}$. Each frequency feature vector $\hat{f}_{df}^{i,j}$ encodes the frequency domain information of the (i, j) -th patch.

We spatially combine $\hat{f}_{df}^{i,j}$ according to the spatial position (i, j) , yielding a frequency feature map \hat{F}_{df} with size $n \times n \times c_{df}$. We concatenate \hat{F}_{df} with F_{ds} to form a feature map with size $n \times n \times (c_{df} + c_{ds})$, in which each pixel-wise feature vector contains both spatial domain information and frequency domain information of one patch. Then, the concatenated feature map is delivered to D_a to predict a $n \times n$ inharmonious region mask, in which 0 indicates harmonious patches and 1 indicates inharmonious patches.

By taking the harmonized result \tilde{I}_o , the composite image I_c , and the background image I_b as the input of D separately, we could get a $n \times n$ inharmonious region mask for each input. The loss function to update D could be written as

$$\mathcal{L}_{adv}^D = \|D(\tilde{I}_o) - \bar{M}\|^2 + \|D(I_c) - \bar{M}\|^2 + \|D(I_b)\|^2, \quad (5)$$

where \bar{M} means the downsampled mask with size $n \times n$. For I_c and \tilde{I}_o , we expect that D_a could distinguish the foreground (inharmonious) patches from the background (har-

monious) patches, so the predicted inharmonious region mask aligns with \tilde{M} . For I_b , since there is no inharmonious patch, the predicted inharmonious region mask is supposed to be an all-zero mask.

Under the adversarial learning framework (Goodfellow et al. 2014), we update the dual-domain generator G and the dual-domain discriminator D alternately. When updating G , we hope that the harmonized output \tilde{I}_o could confuse D , so that D is unable to distinguish the inharmonious patches. Therefore, the adversarial loss to update G is given as $\mathcal{L}_{adv}^G = \|D(\tilde{I}_o)\|^2$.

So far, the total loss for training G is summarized as

$$\mathcal{L}_G = \mathcal{L}_s + \lambda_c \mathcal{L}_c + \lambda_{adv} \mathcal{L}_{G_{adv}}, \quad (6)$$

where the trade-off parameters λ_c and λ_{adv} are set to 2 and 10 respectively in our experiments.

4 Experiments

We conduct experiments on COCO (Lin et al. 2014) and WikiArt (Tan et al. 2019). Refer to the Supplementary for more implementation details.

4.1 Baselines

We divide baselines into two groups: painterly image harmonization methods (Luan et al. 2018; Zhang, Wen, and Shi 2020; Peng, Wang, and Wang 2019) and artistic style transfer methods (Huang and Belongie 2017; Liu et al. 2021).

The first group of baselines process the foreground region of composite image. We compare with Deep Image Blending (Zhang, Wen, and Shi 2020) (“DIB” for short), Deep Painterly Harmonization (Luan et al. 2018) (“DPH” for short) and E2STN (Peng, Wang, and Wang 2019). We also include traditional image blending method Poisson Image Editing (Pérez, Gangnet, and Blake 2003) (“Poisson” for short) for comparison.

The second group of baselines stylize the whole photographic (content) image which provides the foreground object. To adapt artistic style transfer methods to our task, we first stylize the entire content image according to the background (style) image. Then we cut the stylized foreground object and paste it on the background image. We compare with typical or recent style transfer methods: WCT (Li et al. 2017a), AdaIN (Huang and Belongie 2017), SANet (Park and Lee 2019), AdaAttN (Liu et al. 2021), and StyTr2 (Deng et al. 2022).

4.2 Qualitative Analysis

To compare with the first group of baselines, the results of different methods are illustrated in Figure 3. Although Poisson (Pérez, Gangnet, and Blake 2003) could smoothen the boundary, the styles of foreground and background are still dramatically different. E2STN (Peng, Wang, and Wang 2019) is also struggling to transfer the style (*e.g.*, row 2, 4). DIB (Zhang, Wen, and Shi 2020) could transfer the style to some extent, but it severely distorts the content information of foreground object (*e.g.*, row 2, 5). DPH (Luan et al. 2018) achieves competitive results among the baselines. Compared with DPH (Luan et al. 2018), our PHDNet can preserve the

content structure better (row 1) and transfer the style better (row 2). Our PHDNet also has stronger ability to transfer the texture and pattern from background image. For example, our PHDNet can transfer the colorful strips to the umbrella (row 3) and quadrangle color blocks with sharp edges to the truck (row 4). Our PHDNet can also restore the vertical strips in the foreground region (row 5).

To compare with the second group of baselines, the results of different methods are illustrated in Figure 4. Since style transfer methods stylize the entire content image with limited attention paid to the foreground object, the foreground object may not be sufficiently stylized (row 1, 4), which makes the foreground very obtrusive and easily separated from the background. In contrast, our PHDNet focuses on the foreground stylization. By taking the locality into account, in our harmonized results, the foreground object has more consistent style with its neighboring regions and thus appears to be more naturally blended into the background. Moreover, our PHDNet has stronger style transfer ability. For example, in row 1, the background has several green spots, so the foreground cat also has green spots. In row 2, 4, 5, the foreground objects of other methods are very smooth while our foreground objects own the fine-grained texture transferred from background images.

The advantages of our PHDNet come from two aspects. Firstly, we perform harmonization in both spatial domain and frequency domain. As claimed in Section 1, the frequency feature can capture the global style and periodic texture/pattern, so our PHDNet is able to reconstruct the missing or corrupted textures and patterns in the foreground. Secondly, the discriminator helps the generator in an adversarial manner, so that the foreground in the harmonized image is more compatible with the background.

4.3 User Study

As there is no ground-truth harmonized image, we cannot use evaluation metrics (*e.g.*, MSE, PSNR) to evaluate the model performance quantitatively. Therefore, we conduct user study to compare different methods. We randomly select 100 content images from COCO and 100 style images from WikiArt to generate 100 composite images for user study. We compare the harmonized results generated by SANet (Park and Lee 2019), AdaAttN (Liu et al. 2021), StyTr2 (Deng et al. 2022), DPH (Luan et al. 2018), E2STN (Peng, Wang, and Wang 2019), and our PHDNet.

Specifically, for each composite image, we can obtain 6 harmonized outputs generated by 6 above-mentioned methods. Then we select 2 images from these 6 images to construct image pairs. Based on 100 composite images, we could construct 1,500 image pairs. Then we invite 20 users to see one image pair each time and pick out the more harmonious one. Finally, we collect 30,000 pairwise results and employ the Bradley-Terry (B-T) model (Bradley and Terry 1952; Lai et al. 2016) to obtain the overall ranking of all methods. The results are reported in Table 1 in the left sub-table, in which we can observe that our PHDNet achieves the highest B-T score and outperforms other baselines.

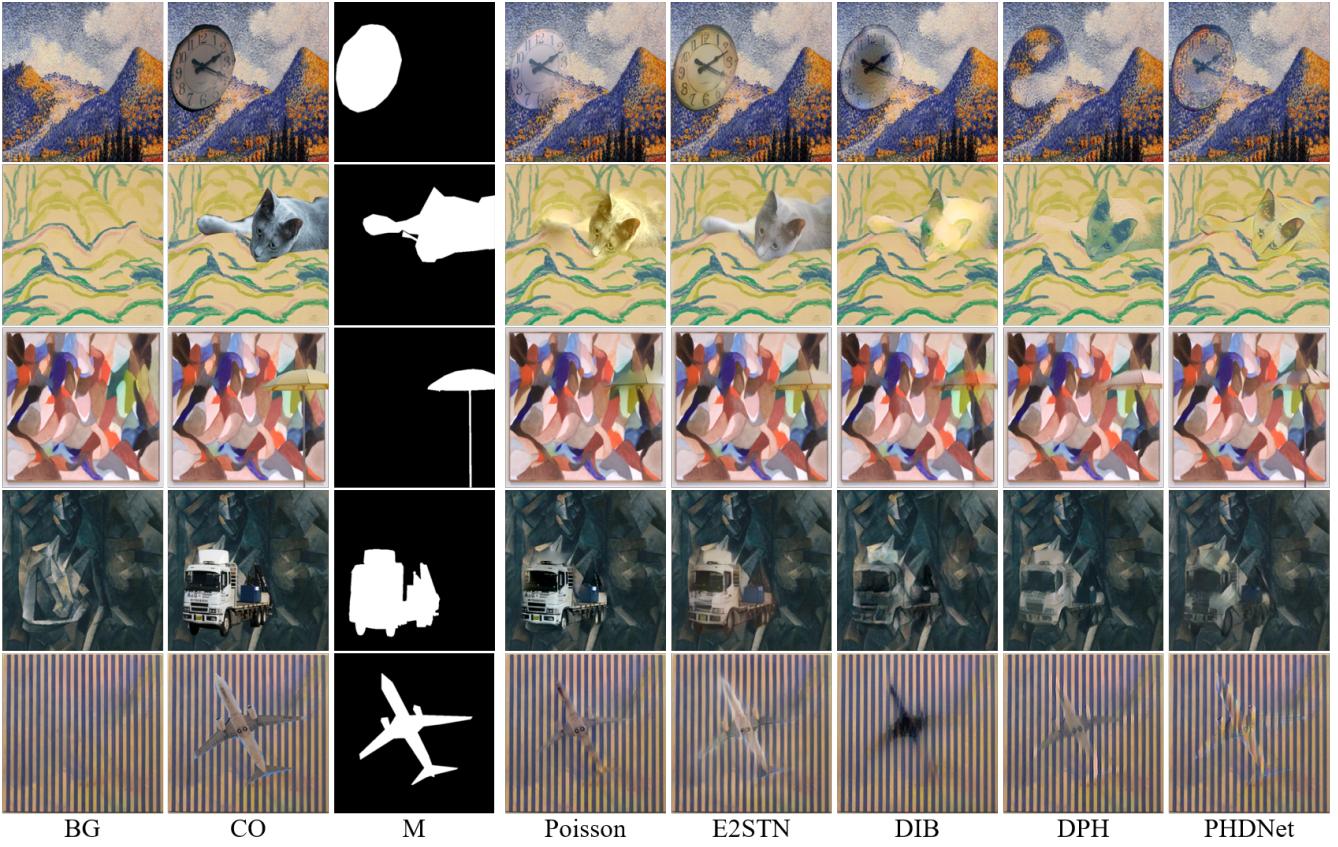


Figure 3: Example results of painterly image harmonization baselines and our PHDNet. “BG” (*resp.*, “CO”) means “background” (*resp.*, “composite”).

Method	Type	B-T	#	B-T	
				G w/ f.	D w/ f.
DPH	OP	0.555	V1	-	-1.729
E2STN	FF	-1.811	V2	✓	-
SANet	FF	-0.168	V3	✓	-0.626
AdaAttN	FF	0.029	V4	✓	0.179
StyTr2	FF	0.343	V5	✓	0.827
PHDNet	FF	1.052			1.349

Table 1: B-T scores. Left sub-table: B-T scores of different baselines and our PHDNet. In “Type” column, “OP” means optimization-based method, while “FF” means feed-forward method. Right sub-table: B-T scores of different network structures, in which G (*resp.*, D) means generator (*resp.*, discriminator), w/ f. means “with frequency-related module”, “-” means without discriminator.

4.4 Ablation Studies

We ablate each frequency-related module in our PHDNet, *i.e.*, the ResFFT module in generator G and the frequency branch D_f in discriminator D . We construct different ablated versions according to whether using ResFFT module, whether using discriminator, and whether using frequency branch in the discriminator, leading to in total 5 versions. As summarized in the right sub-table in Table 1, we first

remove the ResFFT module in the generator and remove the whole discriminator, which is referred to as “V1”. Then, we add ResFFT module in the generator, leading to “V2”. Based on “V2”, we add the discriminator without frequency branch, leading to “V3”. Next, we further add frequency branch to the discriminator, arriving at our full version “V5”. Additionally, based on “V5”, we remove the ResFFT module in the generator and get “V4”. Following the way in Section 4.3, we conduct user study and employ the B-T model (Bradley and Terry 1952; Lai et al. 2016) to obtain the overall ranking of all versions.

From the right sub-table in Table 1, we can see that the performances without using discriminator (“V1”, “V2”) are very poor. Based on “V2” and “V3”, we can see that even using the simplified discriminator without frequency branch can significantly improve the performance, which demonstrates that it is useful to push the foreground to be indistinguishable from the background. The comparison between “V1” (*resp.*, “V4”) and “V2” (*resp.*, “V5”) verifies the effectiveness of the ResFFT module in the generator. The comparison between “V3” and “V5” verifies the effectiveness of the frequency branch in the discriminator. Together with two frequency-related modules, our full version “V5” achieves the highest score, which proves that the frequency branch in the discriminator can help the ResFFT module learn to

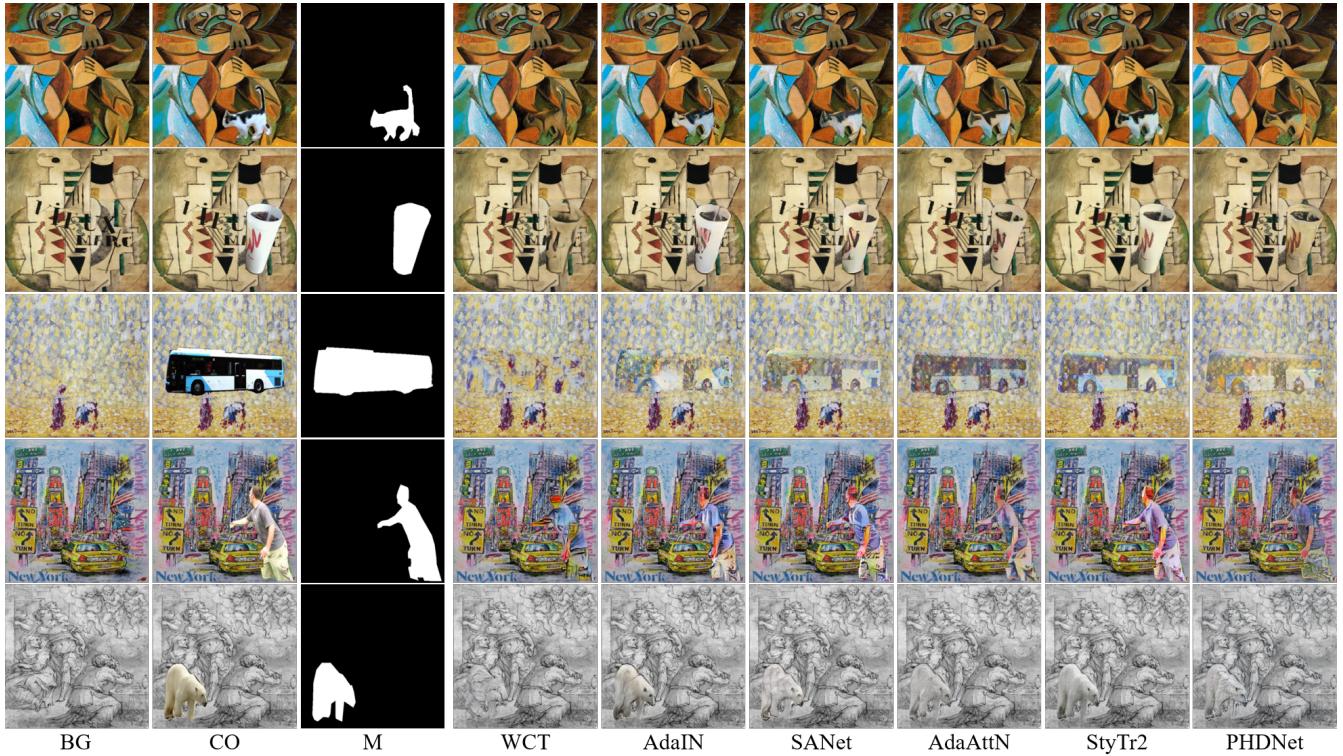


Figure 4: Example results of artistic style transfer baselines and our PHDNet.

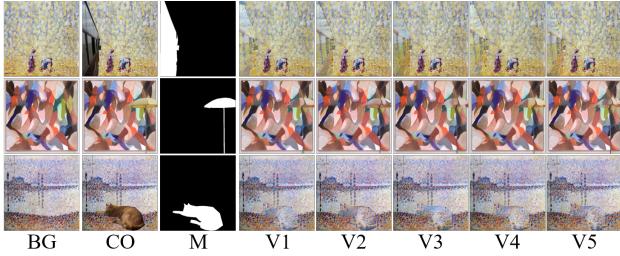


Figure 5: Example results of each ablated version.

harmonize the frequency feature map.

In addition, we show the harmonized results of different versions in Figure 5. One observation is that the generated results without using discriminator (“V1”, “V2”) are prone to have artifacts and the discriminator can help enhance the quality of generated images. Another observation is that the frequency-related modules (ResFFT module in the generator and the frequency branch in the discriminator) can collaborate with each other to better transfer the textures/patterns from background image to composite foreground, resulting in more harmonious images.

4.5 Hyper-Parameter Analyses

We investigate the impact of the hyper-parameter in PHDNet, *i.e.*, the patch number n^2 in our dual-domain discriminator (see Section 3.2). We provide the visualization results

when varying n . Details are left to the Supplementary.

4.6 Visualization of Frequency Maps

In order to demonstrate the effectiveness of frequency domain learning in our PHDNet intuitively, we visualize the different frequency maps of our frequency-related modules. The comparison results show that our PHDNet can well transfer the textures from the background style image to the foreground of the composite image, and generate the harmonized image. Details are also left to the Supplementary.

4.7 Limitations

Although our PHDNet can generally produce visually appealing and harmonious results, it may also generate under-stylized results when handling certain types of background styles. More discussions and detailed results can be found in the Supplementary.

5 Conclusion

In this work, we have introduced frequency domain learning into painterly image harmonization task. We have proposed a novel dual-domain network PHDNet, which contains a dual-domain generator and a dual-domain discriminator. Extensive experiments have demonstrated that our PHDNet has very strong style transfer ability and the stylized foreground is compatible with the background.

Acknowledgments

The work was supported by the National Science Foundation of China (62076162), and the Shanghai Municipal Science and Technology Major/Key Project, China (2021SHZDZX0102, 20511100300).

References

Bansal, A.; Sheikh, Y.; and Ramanan, D. 2018. Pixelnn: Example-based image synthesis. *ICLR*.

Bradley, R. A.; and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*.

Cai, M.; Zhang, H.; Huang, H.; Geng, Q.; Li, Y.; and Huang, G. 2021. Frequency domain image translation: More photo-realistic, better identity-preserving. In *ICCV*.

Cong, W.; Niu, L.; Zhang, J.; Liang, J.; and Zhang, L. 2021. BargainNet: Background-Guided Domain Translation for Image Harmonization. In *ICME*.

Cong, W.; Tao, X.; Niu, L.; Liang, J.; Gao, X.; Sun, Q.; and Zhang, L. 2022. High-Resolution Image Harmonization via Collaborative Dual Transformations. In *CVPR*.

Cong, W.; Zhang, J.; Niu, L.; Liu, L.; Ling, Z.; Li, W.; and Zhang, L. 2020. Dovenet: Deep image harmonization via domain verification. In *CVPR*.

Cun, X.; and Pun, C. 2020. Improving the Harmony of the Composite Image by Spatial-Separated Attention Module. *IEEE Transactions on Image Processing*, 29: 4759–4771.

Deng, Y.; Tang, F.; Dong, W.; Ma, C.; Pan, X.; Wang, L.; and Xu, C. 2022. StyTr2: Image Style Transfer with Transformers. In *CVPR*.

Du, L. 2020. How much deep learning does neural style transfer really need? an ablation study. In *WACV*.

Gatys, L. A.; Ecker, A. S.; and Bethge, M. 2016. Image style transfer using convolutional neural networks. In *CVPR*.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *NeurIPS*, 27.

Guo, Z.; Guo, D.; Zheng, H.; Gu, Z.; Zheng, B.; and Dong, J. 2021a. Image harmonization with transformer. In *ICCV*.

Guo, Z.; Zheng, H.; Jiang, Y.; Gu, Z.; and Zheng, B. 2021b. Intrinsic Image Harmonization. In *CVPR*.

Hang, Y.; Xia, B.; Yang, W.; and Liao, Q. 2022. SCS-Co: Self-Consistent Style Contrastive Learning for Image Harmonization. In *CVPR*.

Hao, G.; Iizuka, S.; and Fukui, K. 2020. Image Harmonization with Attention-based Deep Feature Modulation. In *BMVC*.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.

Huang, X.; and Belongie, S. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*.

Huo, J.; Jin, S.; Li, W.; Wu, J.; Lai, Y.-K.; Shi, Y.; and Gao, Y. 2021. Manifold alignment for semantically aligned style transfer. In *ICCV*.

Kolkin, N.; Salavon, J.; and Shakhnarovich, G. 2019. Style transfer by relaxed optimal transport and self-similarity. In *CVPR*.

Lai, W.; Huang, J.; Hu, Z.; Ahuja, N.; and Yang, M. 2016. A Comparative Study for Single Image Blind Deblurring. In *CVPR*.

Lalonde, J.; and Efros, A. A. 2007. Using Color Compatibility for Assessing Image Realism. In *ICCV*.

Li, X.; Liu, S.; Kautz, J.; and Yang, M.-H. 2018. Learning linear transformations for fast arbitrary style transfer. *arXiv preprint arXiv:1808.04537*.

Li, Y.; Fang, C.; Yang, J.; Wang, Z.; Lu, X.; and Yang, M.-H. 2017a. Universal style transfer via feature transforms. *NeurIPS*.

Li, Y.; Wang, N.; Liu, J.; and Hou, X. 2017b. Demystifying neural style transfer. In *IJCAI*.

Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. 740–755. Springer.

Ling, J.; Xue, H.; Song, L.; Xie, R.; and Gu, X. 2021. Region-aware adaptive instance normalization for image harmonization. In *CVPR*.

Liu, S.; Lin, T.; He, D.; Li, F.; Wang, M.; Li, X.; Sun, Z.; Li, Q.; and Ding, E. 2021. Adaattn: Revisit attention mechanism in arbitrary neural style transfer. In *ICCV*.

Luan, F.; Paris, S.; Shechtman, E.; and Bala, K. 2018. Deep painterly harmonization. In *Computer graphics forum*, 95–106. Wiley Online Library.

Mahendran, A.; and Vedaldi, A. 2015. Understanding deep image representations by inverting them. In *CVPR*.

Mardani, M.; Liu, G.; Dundar, A.; Liu, S.; Tao, A.; and Catanzaro, B. 2020. Neural ffts for universal texture image synthesis. *NeurIPS*.

Park, D. Y.; and Lee, K. H. 2019. Arbitrary style transfer with style-attentional networks. In *CVPR*.

Peng, H.-J.; Wang, C.-M.; and Wang, Y.-C. F. 2019. Element-Embedded Style Transfer Networks for Style Harmonization. In *BMVC*.

Pérez, P.; Gangnet, M.; and Blake, A. 2003. Poisson image editing. *ACM Transactions on Graphics*, 22(3): 313–318.

Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.

Roy, H.; Chaudhury, S.; Yamasaki, T.; and Hashimoto, T. 2021. Image inpainting using frequency-domain priors. *Journal of Electronic Imaging*, 30(2): 023016.

Shen, X.; Yang, J.; Wei, C.; Deng, B.; Huang, J.; Hua, X.-S.; Cheng, X.; and Liang, K. 2021. Dct-mask: Discrete cosine transform mask representation for instance segmentation. In *CVPR*.

Simonyan, K.; and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. *ICLR*.

Sofiuk, K.; Popenova, P.; and Konushin, A. 2021. Foreground-aware Semantic Representations for Image Harmonization. In *WACV*.

Song, S.; Zhong, F.; Qin, X.; and Tu, C. 2020. Illumination Harmonization with Gray Mean Scale. In *Computer Graphics International Conference*.

Sunkavalli, K.; Johnson, M. K.; Matusik, W.; and Pfister, H. 2010. Multi-scale image harmonization. *ACM Transactions on Graphics*, 29(4): 125:1–125:10.

Suvorov, R.; Logacheva, E.; Mashikhin, A.; Remizova, A.; Ashukha, A.; Silvestrov, A.; Kong, N.; Goka, H.; Park, K.; and Lempitsky, V. 2022. Resolution-robust large mask inpainting with fourier convolutions. In *WACV*.

Tan, W. R.; Chan, C. S.; Aguirre, H.; and Tanaka, K. 2019. Improved ArtGAN for Conditional Synthesis of Natural Image and Artwork. *IEEE Transactions on Image Processing*, 28(1): 394–409.

Tsai, Y.; Shen, X.; Lin, Z.; Sunkavalli, K.; Lu, X.; and Yang, M. 2017. Deep Image Harmonization. In *CVPR*.

Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *CVPR*.

Xu, K.; Qin, M.; Sun, F.; Wang, Y.; Chen, Y.-K.; and Ren, F. 2020. Learning in the frequency domain. In *CVPR*.

Xue, S.; Agarwala, A.; Dorsey, J.; and Rushmeier, H. E. 2012. Understanding and improving the realism of image composites. *ACM Transactions on Graphics*, 31(4): 84:1–84:10.

Yang, Y.; and Soatto, S. 2020. Fda: Fourier domain adaptation for semantic segmentation. In *CVPR*.

Yu, Y.; Zhan, F.; Lu, S.; Pan, J.; Ma, F.; Xie, X.; and Miao, C. 2021. Wavefill: A wavelet-based generation network for image inpainting. In *ICCV*.

Zhang, L.; Wen, T.; and Shi, J. 2020. Deep image blending. In *WACV*.

Zhu, J.; Krähenbühl, P.; Shechtman, E.; and Efros, A. A. 2015. Learning a Discriminative Model for the Perception of Realism in Composite Images. In *ICCV*.