

# PhotoWCT<sup>2</sup>: Compact Autoencoder for Photorealistic Style Transfer Resulting from Blockwise Training and Skip Connections of High-Frequency Residuals

Tai-Yin Chiu  
The University of Texas at Austin

Danna Gurari  
University of Colorado Boulder

## Abstract

Photorealistic style transfer is an image editing task with the goal to modify an image to match the style of another image while ensuring the result looks like a real photograph. A limitation of existing models is that they have many parameters, which in turn prevents their use for larger image resolutions and leads to slower run-times. We introduce two mechanisms that enable our design of a more compact model that we call PhotoWCT<sup>2</sup>, which preserves state-of-art stylization strength and photorealism. First, we introduce blockwise training to perform coarse-to-fine feature transformations that enable state-of-art stylization strength in a single autoencoder in place of the inefficient cascade of four autoencoders used in PhotoWCT. Second, we introduce skip connections of high-frequency residuals in order to preserve image quality when applying the sequential coarse-to-fine feature transformations. Our PhotoWCT<sup>2</sup> model requires fewer parameters (e.g., 30.3% fewer) while supporting higher resolution images (e.g., 4K) and achieving faster stylization than existing models.

## 1. Introduction

Photorealistic style transfer is the task of rendering an image in the style of another image such that the result appears like a real photograph to end users (Figure 1a). A limitation of existing methods is that they are *parameter-heavy*, which results in a number of practical limitations. First, they cannot support images of 4K resolution (i.e., 8.3 megapixels) and above, also referred to as ultra high definition (UHD) media. Yet, advancements in technology have led UHD to become standard in commercial products, as exemplified by the increasing number of self-made UHD images and videos shared on online image stocks [4] and YouTube as well as the trend for more movies and TV series on streaming platforms (e.g., Netflix [3], Amazon Prime Video [2]) to support UHD resolution. Other practical concerns include the ability to run methods on memory-constrained or power-constrained devices and to support fast stylization. We aim to introduce a more compact model

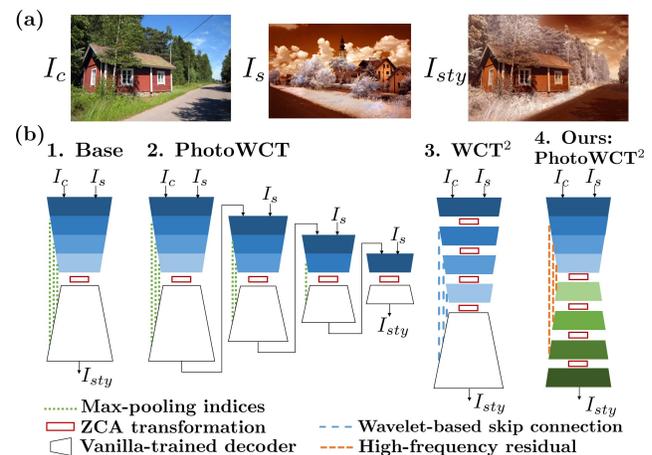


Figure 1: (a) A stylization example using our PhotoWCT<sup>2</sup>. (b) Comparison of methods for photorealistic style transfer: (b1) The base framework is an autoencoder that transforms the input content image,  $I_c$ , and style image,  $I_s$ , into a stylized image,  $I_{sty}$ , with a single feature transformation. (b2,b3,b4) The other models perform multi-scale feature transformation with (b2) PhotoWCT [24] using four cascaded autoencoders to achieve coarse-to-fine feature transformation, (b3) WCT<sup>2</sup> [41] performing fine-to-coarse feature transformation in a single autoencoder using wavelet-based skip connections, and (b4) our PhotoWCT<sup>2</sup> realizing coarse-to-fine feature transformation in a single autoencoder with skip connections of high-frequency residuals.

to address these practical limitations.

Our work, like state-of-art methods [24, 41], builds upon the predominant framework for style transfer methods: an autoencoder [12, 19, 23, 22, 34]. As exemplified in Figure 1b:1, this framework takes as input both a content image and style image ( $I_c$ ,  $I_s$ ), encodes each into a feature representation using a pre-trained network, applies a transformation to alter the content feature with respect to the style feature, and then finally decodes the resulting feature into the stylized image ( $I_{sty}$ ). Our work centers on redesigning two parameter-heavy mechanisms employed by state-of-art

methods into lighter weight representations.

Our first aim is to redesign the mechanism employed to strongly reflect the new style in the rendered image (i.e., to achieve strong stylization strength). This mechanism entails using multiple transformations of features of different scales. The state-of-art method, PhotoWCT [24], employs four cascaded autoencoders to transform the content image by recursively modifying its coarse feature to its fine feature with respect to the corresponding coarse to fine style features (illustrated in Figure 1b:2). Intuitively, this strengthens the transferred style by first modifying the big picture of the content image with respect to the style image and then gradually fine-tuning its fine-grained details. However, PhotoWCT’s use of multiple autoencoders makes it computationally expensive. In contrast, WCT<sup>2</sup> [23] uses a single autoencoder to progressively transform the content image. However, it modifies the content image from its fine feature to its coarse feature with respect to the corresponding fine to coarse style features before decoding the transformed feature into the stylized image (illustrated in Figure 1b:3). This fine-to-coarse feature transformation is shown experimentally [23] to result in weaker stylization strength than the coarse-to-fine feature transformation performed by PhotoWCT. Intuitively, this worse performance may be because initial fine-tuned details might get overshadowed by later big-picture modifications. We introduce a redesign that simultaneously embeds the strengths of PhotoWCT and WCT<sup>2</sup> while overcoming their limitations. We achieve this by introducing a novel technique, called blockwise training, that makes it possible to convert the PhotoWCT cascade into a *single compact autoencoder* that performs *coarse-to-fine feature stylization*.

Our second aim is to redesign the mechanism used to recover content information that gets lost by the autoencoder when rendering the stylized image. Existing methods employ skip connections from the autoencoder’s encoder to its decoder for this purpose. For instance, PhotoWCT [24] skip-connects from the encoder’s max-pooling layer the indices of computed maximum values (i.e., max-pooling indices) to the paired decoder’s max-unpooling layer (illustrated in Figure 1b:2). However, theoretically, max-pooling is lossy [39, 40] and so these max-pooling indices are not guaranteed to be sufficient for good image reconstruction, which in turn results in content distortion in stylized results. In contrast, WCT<sup>2</sup> [41] introduces skip connections based on wavelets (illustrated in Figure 1b:3) that are guaranteed with signal processing theories [39, 40] to yield better image reconstruction performance, and are shown experimentally to do so in Section 4.4. As will be discussed in Section 3, a key reason behind its advantage is that wavelet-based skip connections helps recover high-frequency information that can easily get lost in the encode-decode process. A limitation of wavelet-based skip connections though

is that they require many parameters. We introduce an architecture that we call *skip connections of high-frequency residuals* that makes it possible to achieve the advantage of WCT<sup>2</sup>’s wavelet-based skip connections for better image reconstruction with considerably fewer parameters.

To summarize our key contributions, we introduce a new photorealistic style transfer model, which we call PhotoWCT<sup>2</sup> (illustrated in Figure 1b:4), alongside two new mechanisms used to create this model. The first mechanism is *blockwise training* for redesigning the coarse-to-fine feature transformations in PhotoWCT’s cascade of autoencoders into a single decoder. The second mechanism is *skip connections of high-frequency residuals* that serves as a lightweight representation of wavelet-based skip connections and enables the success of our blockwise training. Experiments show our model preserves state-of-art stylization strength and photorealism while achieving a 30.3% and 15.6% parameter reduction compared to PhotoWCT and WCT<sup>2</sup> respectively. Moreover, experiments show it can support higher resolution images (i.e., UHD) and achieve faster stylization than existing methods [24, 41, 6]. Ablation studies demonstrate that existing mechanisms for training and skip connections are insufficient to produce our compact model and so underscore the critical need of our two new mechanisms, blockwise training and skip connections of high-frequency residuals.

## 2. Related works

**Photorealistic style transfer.** In 2017, the seminal neural network-based method for photorealistic style transfer was introduced [27]. To address it is relatively slow due to its need for many iterations of forward passing and backpropagation, new methods [24, 41, 6, 37] incurred speed gains by using one forward pass. Among them, PhotoWCT [24] achieves the strongest stylization strength at the expense of a parameter-heavy architecture of four autoencoders. WCT<sup>2</sup> [41] and PhotoNAS [6], in contrast, offer single autoencoder architectures, with WCT<sup>2</sup> being superior due to its use of fewer layers and wavelet-based skip connections (which support theory-backed image reconstruction). We introduce a model that achieves comparable stylization strength to the state-of-art PhotoWCT while realizing further advantages over existing models [24, 41, 6]<sup>1</sup>, by requiring fewer parameters, supporting stylization of larger images, and providing faster stylization.

**Greedy layerwise training.** Like traditional greedy layerwise training for autoencoders [9, 21, 30, 8], our blockwise training entails splitting an autoencoder into a sequence of sub-model pairs and then training the pairs se-

<sup>1</sup>We cannot compare to [37] because the code has not been released. With that said, we expect poorer performance from it because it downsizes images to support high resolution (4K) images and so discards information.

quentially. However, the traditional approach pairs an encoder layer with a decoder layer while our approach pairs an encoder block with a decoder block. In addition, they are used for different purposes: while the traditional approach centers on *learning an encoder* that represents a specific dataset [11, 15, 31, 25, 33] and so uses the decoder as a disposable accessory needed to achieve this aim, our approach instead *fixes a pretrained encoder* during training (e.g., VGG [35]) in order to learn a decoder that can reproduce the coarse-to-fine features. Our experiments demonstrate that blockwise training is non-trivial with existing neural network architecture components, failing to produce an effective compact autoencoder (Sections 4.3 and 4.4). We introduce skip connections of high-frequency residuals and demonstrate that it overcomes this limitation, enabling the effective use of blockwise training to develop a compact autoencoder for photorealistic style transfer.

**Skip connections.** A challenge is how to employ skip connections [32, 18] within autoencoders for photorealistic style transfer. In particular, when an autoencoder consists of an encoder that is a fixed pre-trained model and a decoder that learns its inverse function, directly connecting an output from a layer  $l_e$  in the encoder to some layer  $l_d$  of the decoder results in a short circuit phenomenon [6]. This means the connection is so informative that it overshadows the middle layers between  $l_e$  and  $l_d$ , such that the middle layers will not affect the pixel values in the output of the decoder after training. Numerous variants of skip connections address this issue, including indices of maximal values between max-pooling/unpooling layers for PhotoWCT [24], instance-normalized skip-connected features for PhotoNAS [6], and a theoretically motivated wavelet-based approach for WCT<sup>2</sup> [41]. We simplify the wavelet-based architecture into a computationally light variant we call skip connections of high-frequency residuals.

As will be detailed in Section 3.2.3, the computation of a high-frequency residual is similar to that of the first difference image in a Laplacian pyramid [10]. However, while a Laplacian pyramid is built upon an image for multiple levels, a high-frequency residual is computed from a feature map and does not form a pyramid. Moreover, most previous works [14, 17, 20, 38, 16, 7] that integrate a Laplacian pyramid into neural networks heuristically leverage the concept that a Laplacian pyramid preserves high-frequency details from the input image to generate images of better quality. We extend prior work by providing a theoretical explanation why our approach can realize an autoencoder for coarse-to-fine feature transformation for photorealistic style transfer.

### 3. Method

We now introduce our new model PhotoWCT<sup>2</sup> and our two mechanisms that enable its creation: blockwise training

and skip connections of high-frequency residuals.

#### 3.1. Background

To begin, we describe the parameter-heavy mechanisms used in state-of-art photorealistic style transfer methods that we aim to redesign into compact representations.

**PhotoWCT’s coarse-to-fine feature transformations.**

As summarized in the Introduction and illustrated in Figure 1b:2, PhotoWCT [24] consists of a cascade of four autoencoders AEC<sub>N</sub>’s ( $N = 1, 2, 3, 4$ ), where each includes an encoder  $enc_N$  and decoder  $dec_N$ .  $enc_N$  is a pretrained network, specifically VGGNet, from the input layer to the  $reluN_1$  layer.  $dec_N$  is structurally symmetric to  $enc_N$ . To realize the coarse-to-fine feature transformation, the cascade of four autoencoders is in the order from  $N = 4$  to  $N = 1$ . Specifically, content and style images are first encoded by  $enc_4$  into the  $relu4_1$  features. The  $relu4_1$  content feature is then transformed with reference to the  $relu4_1$  style feature using a ZCA feature transformation [23, 13]. The transformed feature is then decoded by  $dec_4$  to become an image  $I_4$ . The three steps of encoding, transformation, and decoding repeat in the next three rounds of  $N = 3, 2, 1$ , with  $I_{N+1}$  as the content image, until the stylized image  $I_1$  is decoded by  $dec_1$ . Finally, image smoothing (using guided filtering) is applied as a post-processing step to  $I_1$  to remove undesired artifacts in the final stylized image.<sup>2</sup>

**Wavelet-based skip connections.** The architecture of wavelet-based skip connections, which were introduced as part of WCT<sup>2</sup> [41], is shown in Figure 2a. Note that WCT<sup>2</sup> is an autoencoder which emulates AEC<sub>4</sub> in PhotoWCT while replacing its max-pooling/unpooling layers with wavelet pooling/unpooling layers for better image reconstruction. As exemplified in Figure 2a, its wavelet pooling layer, when given a feature  $\mathbf{F}$ , produces a low-frequency component  $\mathbf{F}_{LL}$  and three high-frequency components  $\mathbf{F}_{LH}$ ,  $\mathbf{F}_{HL}$ , and  $\mathbf{F}_{HH}$ . Structurally,  $\mathbf{F}_{LL}$  propagates through the middle layers ( $enc_{part}-dec_{part} \triangleq \text{AEC}_{part}$ ) of the network. The skip connections  $\mathbf{F}_{LH}$ ,  $\mathbf{F}_{HL}$ ,  $\mathbf{F}_{HH}$  are then aggregated with  $\mathbf{F}_{LL}$  at the corresponding wavelet unpooling layer in the decoder. The intuition is that most information of  $\mathbf{F}$  gets carried by  $\mathbf{F}_{LL}$ , and the high-frequency information  $\mathbf{F}_{LH}$ ,  $\mathbf{F}_{HL}$ , and  $\mathbf{F}_{HH}$  can be supplemented to improve the model’s reconstruction.

Mathematically, a wavelet pooling/unpooling performs four depthwise convolutions/deconvolutions with stride 2 using the following Haar wavelet kernels:

$$\mathcal{K}_{LL} = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \mathcal{K}_{LH} = \frac{1}{2} \begin{bmatrix} -1 & 1 \\ -1 & 1 \end{bmatrix}, \mathcal{K}_{HH} = \frac{1}{2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \quad (1)$$

<sup>2</sup>The original code for this step has a bug. We describe this issue and our fix in the Supplementary Materials.

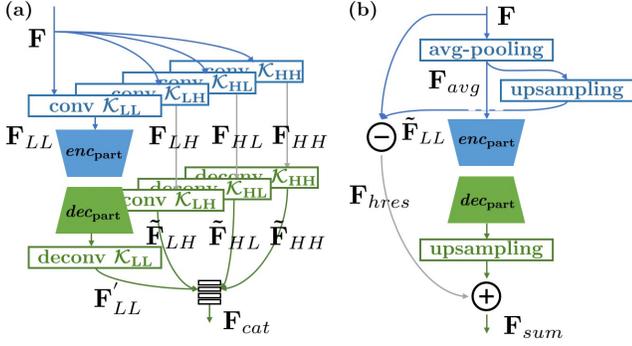


Figure 2: Shown are the middle layers of an autoencoder between a pooling and an unpooling layer for (a) WCT<sup>2</sup>[41] and (b) our method. (a) The wavelet-based skip connection uses the wavelet pooling (blue rectangles) and wavelet unpooling (green rectangles) to improve the image reconstruction ability. (b) Our skip connection of the high-frequency residual simplifies this wavelet-based skip connection into a more compact representation.

and  $\mathcal{K}_{HL} = \mathcal{K}_{LH}^T$ . Let  $(\mathcal{K}_{ij} * \mathbf{f})_{\downarrow 2}$  and  $(\mathcal{K}_{ij} * \mathbf{f})_{\uparrow 2}$  denote the 2-strided convolution and deconvolution of  $\mathcal{K}_{ij}$  ( $i, j \in \{L, H\}$ ) and a feature  $\mathbf{f}$ , respectively.

The wavelet unpooling output  $\mathbf{F}_{cat}$  in Figure 2a is the concatenation of four components  $\mathbf{F}'_{LL}$ ,  $\tilde{\mathbf{F}}_{LH}$ ,  $\tilde{\mathbf{F}}_{HL}$ , and  $\tilde{\mathbf{F}}_{HH}$ , mathematically described as:

$$\mathbf{F}'_{LL} = (\mathcal{K}_{LL} * \text{AEC}_{part}((\mathcal{K}_{LL} * \mathbf{F})_{\downarrow 2}))_{\uparrow 2}, \quad (2)$$

$$\tilde{\mathbf{F}}_{ij} = (\mathcal{K}_{ij} * (\mathcal{K}_{ij} * \mathbf{F})_{\downarrow 2})_{\uparrow 2}, \quad i, j \in \{L, H\}. \quad (3)$$

While this type of skip connection prevents the loss of high-frequency information of the input image and so leads to better image reconstruction [41], it is computationally expensive. For comparison, it requires four times as many parameters as the max-pooling indices skip connection used by PhotoWCT, since the wavelet unpooling output  $\mathbf{F}_{cat}$  has four times the channel length of PhotoWCT's max-unpooling output.

## 3.2. Our approach: PhotoWCT<sup>2</sup>

### 3.2.1 Model architecture

We design our model as an autoencoder  $\text{AEC}_{bt}$ . An overview of its architecture is shown in Figure 3a.

For the encoder, we rely on PhotoWCT's encoder  $enc_4$  (described in Section 3.1). As exemplified in Figure 3a,  $enc_4$  is split into the series of blocks  $enc_4blk_1$ ,  $enc_4blk_2$ ,  $enc_4blk_3$ , and  $enc_4blk_4$ . The output layer of the block  $enc_4blk_N$  is the  $relu_{N-1}$  layer in VGGNet.

We design the decoder, which we call  $dec_{bt}$ , to be structurally symmetric to  $enc_4$ . As exemplified in Figure 3a,  $dec_{bt}$  is split into the series of blocks  $dec_{bt}blk_4$ ,  $dec_{bt}blk_3$ ,

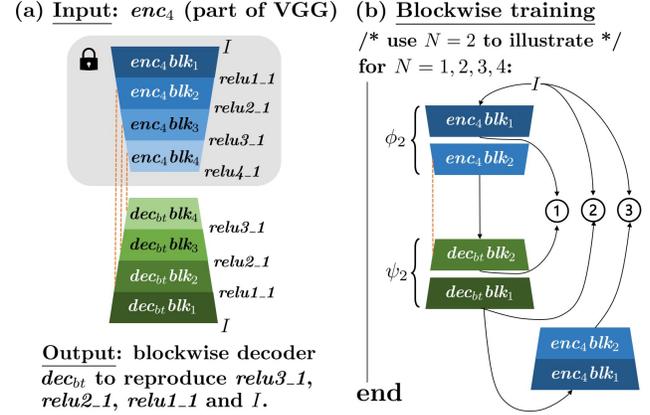


Figure 3: (a) Overview of our PhotoWCT<sup>2</sup> model architecture and (b) illustration of our blockwise training strategy needed to effectively support coarse-to-fine feature transformation in a single autoencoder. The three shown circles indicate the losses blockwise training minimizes, with 1 representing function inversion loss, 2 representing image reconstruction loss, and 3 representing perceptual loss.

$dec_{bt}blk_2$ , and  $dec_{bt}blk_1$ . We design the  $dec_{bt}blk_N$ 's to be structurally symmetric to the  $enc_4blk_N$ 's, with the goal that it will learn the inverse function of  $enc_4blk_N$ , i.e., to convert  $relu_{N-1}$  features to  $relu_{(N-1)-1}$  features. As such, after training, the decoder taking the  $relu_{4-1}$  feature at the bottleneck should be able to sequentially reproduce the  $relu_{3-1}$ ,  $relu_{2-1}$ , and  $relu_{1-1}$  features and the input image.

In order to realize the coarse-to-fine feature transformations, we embed feature transformations at the bottleneck between the encoder and decoder as well as the outputs of  $dec_{bt}blk_4$ ,  $dec_{bt}blk_3$ , and  $dec_{bt}blk_2$  blocks (illustrated in Figure 1b:4). Recall that the purpose of each transformation is to alter each content feature with respect to each style feature at a different scale. Following PhotoWCT, we employ ZCA transformations. We will describe in Section 3.2.2 blockwise training, which is the critical ingredient to make this compact design possible.

We also integrate skip connections into the autoencoder in order to improve not only the image reconstruction, as shown for WCT<sup>2</sup> [41], but also the feature reconstruction in our blockwise training (as will be shown in Section 4.4). We insert our new skip connections of high-frequency residuals in the same positions as used for PhotoWCT's max-pooling indices skip connections. Consequently, in  $enc_4blk_N$  for  $N = 2, 3, 4$ , we replace the original max-pooling with an average-pooling layer from which a skip connection of the high-frequency residual links to the counterpart upsampling layer in  $dec_{bt}blk_N$ . Our new, compact skip connection design will be described in Section 3.2.3.

Finally, as done for PhotoWCT, image smoothing (via guided filtering) is applied as a post-processing step.

### 3.2.2 Blockwise training

We propose two methods: end-to-end training and blockwise training to realize the four function inversions for the decoder. End-to-end training enables the decoder to learn the four function inversions at once. While end-to-end training is good enough to invert the functions, we improve upon it by proposing blockwise training. Our blockwise training approach is illustrated in Figure 3b. As shown, learning of the four function inversions for the decoder is distributed into four steps such that, at each step, a decoder block  $dec_{bt}blk_N$  learns the inverse function of  $enc_{A}blk_N$ . It will be demonstrated in Section 4.3 that this training approach enables more faithful feature reconstruction and image reconstruction in the decoder than alternatives including end-to-end training.

Mathematically, the decoder blocks  $dec_{bt}blk_N$ 's are trained in the order from  $N = 1$  to  $N = 4$  by minimizing the loss  $\mathcal{L}_N$ :

$$\begin{aligned} \mathcal{L}_N(I) = & \mathbb{1}_{N \neq 1} \|\phi_{N-1}(I) - dec_{bt}blk_N(\phi_N(I))\|_2^2 \\ & + \|I - \psi_N(\phi_N(I))\|_2^2 \\ & + \lambda \|\phi_N(I) - \phi_N(\psi_N(\phi_N(I)))\|_2^2, \end{aligned} \quad (4)$$

where  $\phi_N$  and  $\psi_N$  are the functions of the series  $\{enc_{A}blk_1, \dots, enc_{A}blk_N\}$  and  $\{dec_{bt}blk_N, \dots, dec_{bt}blk_1\}$ , respectively,  $\mathbb{1}_{N \neq 1}$  is an indicator function equal to one when  $N \neq 1$  and zero when  $N = 1$ , and  $\lambda$  is set to one for  $N \neq 1$  and zero for  $N = 1$ . The three terms in Equation 4 are the function inversion, image reconstruction, and perceptual losses, respectively. When training a decoder block, the previously trained blocks and the encoder are fixed.<sup>3</sup>

### 3.2.3 Skip connections of high-frequency residuals

Our skip connection of a high-frequency residual is illustrated in Figure 2b. It helps achieve the aim of end-to-end and blockwise trainings of feature/image reconstruction, utilizing and simplifying wavelet-based skip connections into a less computationally expensive representation by replacing the Haar convolutions with average pooling, upsampling and subtraction and the Haar deconvolutions with upsampling and addition (as observed when comparing Figure 2a to Figure 2b). In doing so, it redesigns the concatenation of  $\tilde{\mathbf{F}}_{LL}$ ,  $\tilde{\mathbf{F}}_{LH}$ ,  $\tilde{\mathbf{F}}_{HL}$ , and  $\tilde{\mathbf{F}}_{HH}$  for wavelet-based skip connections into an addition for our approach, thereby enabling the channel length of our outcome  $\mathbf{F}_{sum}$  to become one fourth that of  $\mathbf{F}_{cat}$  from the wavelet-based skip connection.

<sup>3</sup>Due to space constraints, we show in the Supplementary Materials two advantages of training from  $N = 1$  to  $N = 4$  over the reversed order ( $N = 4$  to  $N = 1$ ): better image reconstruction and ease of redesigning a cascade of fewer autoencoders (e.g., 3) into a single autoencoder.

Our motivation for this addition-based approach is an approximation resulting from the observation that the low-frequency parts of an image are much better reconstructed by an autoencoder than the high-frequency edges, as exemplified in Figure 5b. Taking advantage of this observation, we assume a low-frequency feature  $\mathbf{f}$  can be approximately reconstructed by  $AEC_{part}$  (described in Section 3.1), i.e.,  $AEC_{part}(\mathbf{f}) \approx \mathbf{f}$ . This assumption implies the following approximation:

$$\begin{aligned} \mathbf{F}'_{LL} &= (\mathcal{K}_{LL} * AEC_{part}((\mathcal{K}_{LL} * \mathbf{F})_{\downarrow 2}))_{\uparrow 2} \\ &\approx (\mathcal{K}_{LL} * (\mathcal{K}_{LL} * \mathbf{F})_{\downarrow 2})_{\uparrow 2} = \tilde{\mathbf{F}}_{LL}. \end{aligned} \quad (5)$$

With the Haar wavelet expansion  $\mathbf{F} = \tilde{\mathbf{F}}_{LL} + \tilde{\mathbf{F}}_{LH} + \tilde{\mathbf{F}}_{HL} + \tilde{\mathbf{F}}_{HH}$  and the substitution  $\tilde{\mathbf{F}}_{LL} \approx \mathbf{F}'_{LL}$ , we arrive at the following approximation  $\mathbf{F}'_{LL} + \tilde{\mathbf{F}}_{LH} + \tilde{\mathbf{F}}_{HL} + \tilde{\mathbf{F}}_{HH} \approx \mathbf{F}$ . This implies that with addition as feature aggregation, a wavelet pooling input  $\mathbf{F}$  in the encoder is likely to be reconstructed at the wavelet unpooling layer in the decoder.

In describing the implementation of the *encoder* part in Figure 2b, let  $\mathbf{F}_{hres}$  be the sum of high-frequency components  $\tilde{\mathbf{F}}_{LH} + \tilde{\mathbf{F}}_{HL} + \tilde{\mathbf{F}}_{HH}$ . We call  $\mathbf{F}_{hres}$  the *high-frequency residual* of  $\mathbf{F}$  since it is the difference between  $\mathbf{F}$  and the low-frequency component  $\tilde{\mathbf{F}}_{LL}$ . By substituting  $\mathcal{K}_{LL}$  in Equation 3 with its matrix form,  $\tilde{\mathbf{F}}_{LL}$  can be rewritten as follows:

$$\begin{aligned} \tilde{\mathbf{F}}_{LL} &= \left( \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} * \left( \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} * \mathbf{F} \right)_{\downarrow 2} \right)_{\uparrow 2} \\ &= \left( \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} * \left( \frac{1}{4} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} * \mathbf{F} \right)_{\downarrow 2} \right)_{\uparrow 2} \\ &= \text{upsampling}(\text{avgpooling}(\mathbf{F})). \end{aligned} \quad (6)$$

Therefore, the computation of  $\mathbf{F}_{hres}$  simplifies to the following:  $\mathbf{F}$  is first average-pooled to become a low-frequency feature  $\mathbf{F}_{avg} = \text{avgpooling}(\mathbf{F})$ , which in turn is upsampled and subtracted from  $\mathbf{F}$ . Note that if we generalize to build a pyramid on a feature map but not an image and replace the 2-by-2 matrices of ones in Equation 6 by 5-by-5 Gaussian matrices,  $\mathbf{F}_{hres}$  becomes the first difference ‘image’ in the Laplacian pyramid [10] of  $\mathbf{F}$ . Also note that different from the Haar pyramid [5] of  $\mathbf{F}$ , where the first level saves  $\mathbf{F}_{ij}$  ( $i, j \in \{L, H\}$ ) of the half size of  $\mathbf{F}$ , our framework saves  $\tilde{\mathbf{F}}_{LL}$  and  $\mathbf{F}_{hres}$  of the same size as  $\mathbf{F}$ .

To reproduce the feature  $\mathbf{F}$  in the *decoder* by skip-connecting  $\mathbf{F}_{hres}$ ,  $\mathbf{F}_{avg}$  is forward-passed through  $AEC_{part}$ , upsampled, and added to  $\mathbf{F}_{hres}$  as illustrated in Figure 2b. The sum feature  $\mathbf{F}_{sum}$  reproduces  $\mathbf{F}$  under the assumption  $AEC_{part}(\mathbf{f}) \approx \mathbf{f}$  for a low-frequency feature  $\mathbf{f}$ :

$$\begin{aligned} \mathbf{F}_{sum} &= \text{upsampling}(AEC_{part}(\mathbf{F}_{avg})) + \mathbf{F}_{hres} \\ &\approx \text{upsampling}(\mathbf{F}_{avg}) + \mathbf{F}_{hres} \\ &= \tilde{\mathbf{F}}_{LL} + (\tilde{\mathbf{F}}_{LH} + \tilde{\mathbf{F}}_{HL} + \tilde{\mathbf{F}}_{HH}) = \mathbf{F}. \end{aligned} \quad (7)$$

## 4. Experiments

We now evaluate our style transfer model and two mechanisms for creating it: blockwise training (BT) and skip connections of high-frequency residuals. For comparison, we also include our model resulting from end-to-end (E2E) training. We refer to our two models as PhotoWCT<sup>2</sup> when distinction is not necessary, and PhotoWCT<sup>2</sup> (E2E) and PhotoWCT<sup>2</sup> (BT) otherwise. We conduct all experiments on an Nvidia 1080-Ti GPU with 11GB memory.

### 4.1. Model size and speed

First, we assess model size and two benefits that arise from a more compact model: model speed and support for higher resolution images.

**Baselines.** For comparison, we evaluate two top-performing models, PhotoWCT [24] and WCT<sup>2</sup> [41], as well as the more recent PhotoNAS model [6].

**Dataset.** We test all models on five resolutions: 1024 × 512, 1280 × 720 (HD), 1920 × 1080 (Full HD), 2560 × 1440 (Quad HD), and 3840 × 2160 (4K UHD). To efficiently collect images, we download a 4K video [1] from YouTube and sample a frame per second to collect 100 frames. We then downsample each frame to the other lower resolutions.

**Results.** Table 1(a,b) shows the number of parameters, the number of layers<sup>4</sup> on the mainstream path of each model (i.e., skip connections are excluded), and the models’ stylization speed for different resolutions.

As observed in Table 1a, our PhotoWCT<sup>2</sup> model is the most lightweight. It uses 82.5% fewer parameters than PhotoNAS (7.05M vs. 40.24M), 30.3% fewer parameters than WCT<sup>2</sup> (7.05M vs. 10.12M), and 15.6% fewer parameters than PhotoWCT (7.05M vs. 8.35M).

PhotoWCT<sup>2</sup> is also the only model that can handle all the tested resolutions. This implies that it is the only approach that can stylize UHD (i.e., 4K) images.

Additionally, our PhotoWCT<sup>2</sup> can stylize images at the fastest speeds among all models. For instance, PhotoWCT<sup>2</sup> is 0.27 and 0.22 seconds faster than WCT<sup>2</sup> and PhotoWCT for FHD images, taking 0.53 seconds. PhotoWCT<sup>2</sup> also saves 0.42 seconds compared to PhotoWCT (0.82 vs. 1.24) for QHD image stylization.

### 4.2. Image quality and stylization strength

We next assess to what extent our compressed model can preserve the advantages of existing methods, specifically the ability to generate high quality images and render images with a strong stylization strength.

<sup>4</sup>The following types of layers are counted: convolution, deconvolution, max-pooling, average pooling, and upsampling layers.

**Baselines.** For comparison, we again evaluate PhotoNAS, PhotoWCT, and WCT<sup>2</sup>. As an upper bound, we also evaluate the quality of the original content images, which we refer to as “pristine” images. This enables examination of the extent to which stylization downgrades the original quality.

**Dataset.** We use the modified version of the DPST dataset [27].<sup>5</sup> We swap the roles of content and style images in each pair to generate more examples, resulting in 100 stylized images per model.

**Metrics.** Four no-reference *image quality assessment* algorithms are adopted for the quality evaluation: BRISQUE [28] (**0**, 100), NIQE [29] (**0**, ∞), NIMA-quality [36] (1, **10**), and NIMA-aesthetic (1, **10**), where the range of a metric value is shown in a parenthesis pair with the bold number as the best. We use regularized style loss from DPST [27] to evaluate *stylization strength* (description in the Supplementary Materials). The lower the mean style loss  $\bar{\mathcal{L}}_{s,m}$  of a method  $m$ , the stronger the method’s stylization strength.

**Results.** Table 1(c) shows the image quality assessment<sup>6</sup> and stylization strength results. Note that all reported scores are the mean of values across the 100 stylized images.

For image quality, we observe that our PhotoWCT<sup>2</sup> performs similarly to the other top-performing methods: PhotoWCT and WCT<sup>2</sup>. Our PhotoWCT<sup>2</sup> performs slightly better with respect to two of the metrics (NIQE and NIMA-a) and slightly worse for the other two metrics. This highlights that our PhotoWCT<sup>2</sup> can preserve the quality of the top-performing PhotoWCT and WCT<sup>2</sup> methods while using considerably fewer model parameters.

For stylization strength, we observe again that our PhotoWCT<sup>2</sup> performs comparably to the top-performing PhotoWCT; i.e., -0.70 for PhotoWCT versus -0.66 and -0.69 for our PhotoWCT<sup>2</sup> (E2E) and PhotoWCT<sup>2</sup> (BT). We attribute the stronger stylization strength of PhotoWCT<sup>2</sup> (BT) than that of PhotoWCT<sup>2</sup> (E2E) to the better feature/image reconstruction ability of PhotoWCT<sup>2</sup> (BT) (see Section 4.3). Moreover, our models considerably outperform the other two baselines. For example, PhotoWCT<sup>2</sup> (BT) results in scores that are 1.71 and 1.49 better than that of PhotoNAS and WCT<sup>2</sup> respectively. We illustrate qualitatively the stylization strength of our approach compared to the baselines in Figure 4a.

### 4.3. Ablation study on training strategies

Next we compare using blockwise training to train our AEC<sub>bt</sub> architecture (described in Section 3.2.1) with two other methods: end-to-end and vanilla training methods.

<sup>5</sup>The original dataset consists of 60 content-style image pairs including some toy examples, which are excluded here.

<sup>6</sup>We suspect the NIQE metric is inferior, since the scores for WCT<sup>2</sup>, PhotoWCT, and our model are better than those for pristine images.

Model	(a) Size		(b) Speed performance					(c) Image quality & Stylization strength				
	# par	# layer	1024×512	HD 1280×720	FHD 1920×1080	QHD 2560×1440	4K 3840×2160	BRIS (27.4)	NIQE (3.19)	NIMA-q (5.11)	NIMA-a (5.27)	$\bar{\mathcal{L}}_{s,m}$
PhNAS	40.24M	35	0.23	OOM	OOM	OOM	OOM	33.0	3.24	4.75	4.92	1.02
WCT <sup>2</sup>	10.12M	24	0.30	0.43	0.80	OOM	OOM	<b>30.8</b>	3.07	<b>4.91</b>	5.01	0.80
PhWCT	8.35M	48	0.21+0.03	0.32+0.06	0.61+0.14	1.01+0.23	OOM	31.8	<b>2.90</b>	4.88	5.06	<b>-0.70</b>
Ours (E2E)	<b>7.05M</b>	<b>24</b>	<b>0.18+0.03</b>	<b>0.24+0.06</b>	<b>0.39+0.14</b>	<b>0.59+0.23</b>	<b>1.22+0.54</b>	31.7	2.91	4.90	<b>5.10</b>	-0.66
Ours (BT)								31.6	<b>2.90</b>	4.90	<b>5.10</b>	-0.69

Table 1: Characteristics of our models (PhotoWCT<sup>2</sup> (E2E) and PhotoWCT<sup>2</sup> (BT)) and three baselines PhotoNAS (PhNAS), WCT<sup>2</sup>, and PhotoWCT (PhWCT) in terms of model size, speed performance, image quality, and stylization performance. Our models are the most lightweight, the fastest, able to handle the largest resolution (4K), while preserving the good image quality and strong stylization strength of existing state-of-art methods. BRIS: BRISQUE. NIMA-q/a: NIMA-quality/aesthetic. OOM: out-of-memory. (The four values in the parentheses for (c) are scores for pristine images.)



Figure 4: Stylized results from (a) different models and (b) our model trained with vanilla training. Our end-to-end trained model (PhotoWCT<sup>2</sup> (E2E)) and blockwisely trained model (PhotoWCT<sup>2</sup> (BT)) can produce visually pleasant results and strong stylization strength comparable to that of PhotoWCT. See more results in Supplementary Material.

Recall that end-to-end training enables the decoder to learn the four function inversions at once. *Vanilla training*, used for WCT<sup>2</sup> and PhotoNAS, uses  $AEC_{bt}$  as a vanilla auto-encoder that only realizes image reconstruction (i.e., without  $relu_{N-1}$  feature reproduction) in the decoder.

**Dataset.** We randomly sample 500 images from the MSCOCO [26] dataset.

**Metrics.** We evaluate with respect to two metrics. First, we assess the *image reconstruction* ability by computing a pixelwise  $L2$  loss between an input image  $I$  of shape  $H \times W \times C$  and a reconstructed image  $I_{rec}$ :  $\frac{\|I - I_{rec}\|_2^2}{HWC}$ . Second, we assess the *feature reconstruction* ability by computing a relative  $L2$  loss between a feature  $\mathbf{F}_N$  ( $N = 1, 2, 3$ ) from the  $relu_{N-1}$  layer in the encoder block  $enc_4blk_N$  and a reproduced feature  $\mathbf{F}_{N,r}$  from the decoder block  $dec_{bt}blk_{N+1}$ :  $\frac{\|\mathbf{F}_N - \mathbf{F}_{N,r}\|_2^2}{\|\mathbf{F}_N\|_2^2}$ . The reason for using the relative error instead of the absolute error is that an element value in  $\mathbf{F}_N$  or  $\mathbf{F}_{N,r}$  can be large, resulting in a misleading large absolute difference  $\|\mathbf{F}_N - \mathbf{F}_{N,r}\|_2^2$  even if the relative error is small and so suggests a good reconstruction. Mean values are reported for each training strategy.

(a) Different training strategies w/ high-freq residuals

Strategy	$relu_{3.1}$	$relu_{2.1}$	$relu_{1.1}$	image
Vanilla training	3.887	1.163	1.211	<b>0.0003</b>
End-to-end training	0.048	0.022	<b>0.008</b>	0.0008
Blockwise training	<b>0.035</b>	<b>0.021</b>	<b>0.008</b>	0.0006

(b) Different skip connection types w/ blockwise training

Type	$relu_{3.1}$	$relu_{2.1}$	$relu_{1.1}$	image
no skip connect.	0.144	0.183	0.162	0.0052
indices of max.[24]	0.090	0.092	0.065	0.0028
instance norm.[6]	0.045	0.048	0.037	0.0012
wavelet-based skip [41]	0.043	0.030	0.010	<b>0.0006</b>
high-freq residuals	<b>0.035</b>	<b>0.021</b>	<b>0.008</b>	<b>0.0006</b>

Table 2: Loss values for feature and image reconstruction in the decoder resulting from (a) different training strategies and (b) different skip connection types. Results show that blockwise training and high-frequency residuals together achieve the best reconstruction performance.

**Results.** Quantitative results are reported in Table 2a. We observe that our blockwise training results in the best fea-

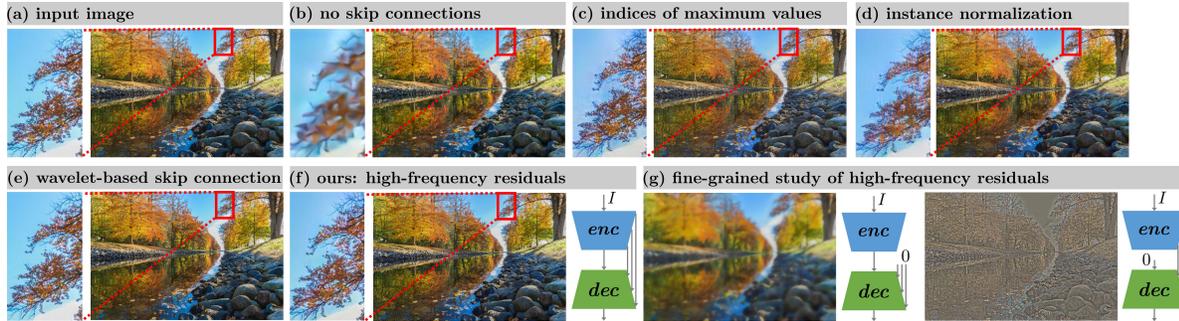


Figure 5: Image reconstruction ability of blockwisely trained autoencoders with different skip connection types. The wavelet-based skip connections in (e) and our high-frequency residuals in (f) achieve the best result among (b-f). The fine grained study in (g) shows its good reconstruction results from the supply of high-frequency information in the input image by the high-frequency residuals. Our high-frequency residuals are the lightweight version of wavelet-based skip connections.

ture reconstruction and second best image reconstruction. We attribute the better reconstruction ability of blockwise training than that of end-to-end training to the distribution of the load of four function inversions into four individual steps. Also, the high feature reconstruction losses of vanilla training indicates that it fails to reproduce any  $reluN_1$  features and so fails to realize the PhotoWCT algorithm. Qualitative results from vanilla training are shown in Figure 4b.

#### 4.4. Ablation study on skip connection types

Next we compare our skip connections based on high-frequency residuals to three alternative lightweight options (1-3) and the heavyweight option (4): (1) no skip connections, (2) skip connections of indices of maximum values used in PhotoWCT [24], (3) instance-normalized skip connections used in PhotoNAS [6], and (4) wavelet-based skip connections used in WCT<sup>2</sup> [41]. To do so, we train four variants of  $AEC_{bt}$  with our blockwise training, each with high-frequency residuals replaced by one of the above skip connection types. We employ the same experimental metrics and dataset as in Section 4.3.

**Results.** The average loss across all images for feature and image reconstruction is shown in Table 2b. We observe that our skip connections of high-frequency residuals consistently leads to considerable improvements over the three alternative lightweight options. Our skip connections even perform slightly better than the heavyweight wavelet-based skip connections. The autoencoder without skip connections has the worst performance, as indicated by the highest loss values. Compared to this no skip connection case, our high-frequency residuals improves by 75.7%, 88.5%, 95.0%, and 88.5% for  $relu3_1$ ,  $relu2_1$ ,  $relu1_1$ , and image reconstruction, respectively. When skip connections of indices of maximal values or instance-normalized skip connections are applied, the performance improves slightly.

In Figure 5, we exemplify the image reconstruction ability of each model with different skip connection types. We notice in Figure 5b that the reconstructed result without

skip connections captures the general idea, such as color, of the input image but fails to reconstruct the high-frequency edges. Although skip connections of indices of maximum values and instance-normalized skip connections improve the reconstruction of high-frequency components, failures still arise including uneven blue sky around leaves and artifacts within leaves as shown in the zoomed-in crops in Figures 5c and 5d. In contrast, the results from wavelet-based skip connections and high-frequency residuals in Figures 5(e,f) are reconstructed better and closer to the original image.

We exemplify the influence of high-frequency residuals on image reconstruction in Figure 5g. We show what a reconstructed image looks like when there are no high-frequency residuals and when there are only high-frequency residuals. In the former case, the skip connections to the decoder are replaced by zero tensors, while in the latter case, the connection between the encoder and decoder at the bottleneck is cut and instead the input to the decoder is a zero tensor. As shown in Figure 5g, the result from no high-frequency residuals is a low-frequency, blurry image, while the result from high-frequency residuals is only the high-frequency edges of the input image. This implies that high-frequency residuals reinforce image reconstruction by supplying high-frequency components from the input image.

## 5. Conclusion

We designed a compact photorealistic style transfer model to embed lightweight representations of PhotoWCT’s coarse-to-fine feature transformations and WCT<sup>2</sup>’s wavelet-based skip connections. Two novel mechanisms, blockwise training and skip connections of high frequency residuals made this compact representation possible. Experiments demonstrate that our PhotoWCT<sup>2</sup> preserves the strong stylization strength of PhotoWCT and good image/feature reconstruction ability of WCT<sup>2</sup> while supporting stylization of higher resolution images and faster stylization speed without loss to image quality.

## References

- [1] See the most beautiful villages in switzerland: lakes, mountains, and green rolling hills. <https://www.youtube.com/watch?v=Ww4Wc34s-fA>. Accessed: 2020-11-15.
- [2] 4K UHD on Amazon Prime Video. <https://reurl.cc/N6gNoQ>. Accessed: 2020-11-15.
- [3] 4K UHD on Netflix. <https://www.netflix.com/search?q=4K>. Accessed: 2020-11-15.
- [4] 4K wallpaper on Pexel. <https://www.pexels.com/search/4k%20wallpaper/>. Accessed: 2020-11-15.
- [5] Edward H Adelson, Eero Simoncelli, and Rajesh Hingorani. Orthogonal pyramid transforms for image coding. In *Visual Communications and image processing II*, volume 845, pages 50–58. International Society for Optics and Photonics, 1987.
- [6] Jie An, Haoyi Xiong, Jun Huan, and Jiebo Luo. Ultrafast photorealistic style transfer via neural architecture search. In *AAAI*, pages 10443–10450, 2020.
- [7] Saeed Anwar and Nick Barnes. Densely residual laplacian super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [8] Eugene Belilovsky, Michael Eickenberg, and Edouard Oyallon. Greedy layerwise learning can scale to imagenet. In *International conference on machine learning*, pages 583–593. PMLR, 2019.
- [9] Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. Greedy layer-wise training of deep networks. In *Advances in neural information processing systems*, pages 153–160, 2007.
- [10] Peter J Burt and Edward H Adelson. The laplacian pyramid as a compact image code. In *Readings in computer vision*, pages 671–679. Elsevier, 1987.
- [11] Min Chen, Xiaobo Shi, Yin Zhang, Di Wu, and Mohsen Guizani. Deep features learning for medical image analysis with convolutional autoencoder neural network. *IEEE Transactions on Big Data*, 2017.
- [12] Tian Qi Chen and Mark Schmidt. Fast patch-based style transfer of arbitrary style. *arXiv preprint arXiv:1612.04337*, 2016.
- [13] Tai-Yin Chiu. Understanding generalized whitening and coloring transform for universal style transfer. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4452–4460, 2019.
- [14] Emily Denton, Soumith Chintala, Arthur Szlam, and Rob Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. *arXiv preprint arXiv:1506.05751*, 2015.
- [15] Fahimeh Farahnakian and Jukka Heikkonen. A deep auto-encoder based approach for intrusion detection system. In *2018 20th International Conference on Advanced Communication Technology (ICACT)*, pages 178–183. IEEE, 2018.
- [16] Xueyang Fu, Borong Liang, Yue Huang, Xinghao Ding, and John Paisley. Lightweight pyramid networks for image de-raining. *IEEE transactions on neural networks and learning systems*, 31(6):1794–1807, 2019.
- [17] Golnaz Ghiasi and Charless C Fowlkes. Laplacian pyramid reconstruction and refinement for semantic segmentation. In *European conference on computer vision*, pages 519–534. Springer, 2016.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [19] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017.
- [20] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 624–632, 2017.
- [21] Hugo Larochelle, Yoshua Bengio, Jérôme Louradour, and Pascal Lamblin. Exploring strategies for training deep neural networks. *Journal of machine learning research*, 10(1), 2009.
- [22] Xueting Li, Sifei Liu, Jan Kautz, and Ming-Hsuan Yang. Learning linear transformations for fast image and video style transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3809–3817, 2019.
- [23] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *Advances in neural information processing systems*, pages 386–396, 2017.
- [24] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 453–468, 2018.
- [25] Peng Liang, Wenzhong Shi, and Xiaokang Zhang. Remote sensing image classification based on stacked denoising autoencoder. *Remote Sensing*, 10(1):16, 2018.
- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [27] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4990–4998, 2017.
- [28] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 21(12):4695–4708, 2012.
- [29] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012.
- [30] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *ICML*, 2011.
- [31] Chathurara Sri Nadith Pathirage, Jun Li, Ling Li, Hong Hao, Wanquan Liu, and Pinghe Ni. Structural damage identification based on autoencoder neural networks and deep learning. *Engineering structures*, 172:13–28, 2018.

- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [33] Alaa Sagheer and Mostafa Kotb. Unsupervised pre-training of a deep lstm-based stacked autoencoder for multivariate time series forecasting problems. *Scientific reports*, 9(1):1–16, 2019.
- [34] Lu Sheng, Ziyi Lin, Jing Shao, and Xiaogang Wang. Avatar-net: Multi-scale zero-shot style transfer by feature decoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8242–8250, 2018.
- [35] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [36] Hossein Talebi and Peyman Milanfar. Nima: Neural image assessment. *IEEE Transactions on Image Processing*, 27(8):3998–4011, 2018.
- [37] Xide Xia, Meng Zhang, Tianfan Xue, Zheng Sun, Hui Fang, Brian Kulis, and Jiawen Chen. Joint bilateral learning for real-time universal photorealistic style transfer. *arXiv preprint arXiv:2004.10955*, 2020.
- [38] Kai Xu, Zhikang Zhang, and Fengbo Ren. Lapran: A scalable laplacian pyramid reconstructive adversarial network for flexible compressive sensing reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 485–500, 2018.
- [39] Jong Chul Ye, Yoseob Han, and Eunju Cha. Deep convolutional framelets: A general deep learning framework for inverse problems. *SIAM Journal on Imaging Sciences*, 11(2):991–1048, 2018.
- [40] Rujie Yin, Tingran Gao, Yue M Lu, and Ingrid Daubechies. A tale of two bases: Local-nonlocal regularization on image patches with convolution framelets. *SIAM Journal on Imaging Sciences*, 10(2):711–750, 2017.
- [41] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9036–9045, 2019.