This CVPR 2020 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# Two-Stage Peer-Regularized Feature Recombination for Arbitrary Image Style Transfer

Jan Svoboda<sup>1,2</sup>, Asha Anoosheh<sup>1</sup>, Christian Osendorfer<sup>1</sup>, Jonathan Masci<sup>1</sup> <sup>1</sup>NNAISENSE, Switzerland <sup>2</sup>Universita della Svizzera italiana, Switzerland

# Abstract

This paper introduces a neural style transfer model to generate a stylized image conditioning on a set of examples describing the desired style. The proposed solution produces high-quality images even in the zero-shot setting and allows for more freedom in changes to the content geometry. This is made possible by introducing a novel Two-Stage Peer-Regularization Layer that recombines style and content in latent space by means of a custom graph convolutional layer. Contrary to the vast majority of existing solutions, our model does not depend on any pre-trained networks for computing perceptual losses and can be trained fully end-to-end thanks to a new set of cyclic losses that operate directly in latent space and not on the RGB images.

An extensive ablation study confirms the usefulness of the proposed losses and of the Two-Stage Peer-Regularization Layer, with qualitative results that are competitive with respect to the current state of the art using a single model for all presented styles. This opens the door to more abstract and artistic neural image generation scenarios, along with simpler deployment of the model.

# 1. Introduction

Neural style transfer (NST), introduced by the seminal work of Gatys [8], is an area of research that focuses on models that transform the visual appearance of an input image (or video) to match the style of a desired target image. For example, a user may want to convert a given photo to appear as if Van Gogh had painted the scene.

NST has seen a tremendous growth within the deep learning community and spans a wide spectrum of applications e.g. converting time-of-day [43, 12], mapping among artwork and photos [1, 43, 12], transferring facial expressions [17], transforming animal species [43, 12], etc.

Despite their popularity and often-quality results, current NST approaches are not free from limitations. Firstly, the original formulation of Gatys *et al.* requires a new op-



Figure 1. Our model is able to take a content image and convert it into arbitrary target style in a forward manner. In this figure, a photo in the middle is converted into styles from 4 different painters, (left side, from top to bottom) Picasso, Kirchner, Roerich and Monet.

timization process for each transfer performed, making it impractical for many real-world scenarios. In addition, the method relies heavily on pre-trained networks, usually borrowed from classification tasks, that are known to be suboptimal and have recently been shown to be biased toward texture rather than structure [9]. To overcome this first limitation, deep neural networks have been proposed to approximate the lengthy optimization procedure in a single feed forward step thereby making the models amenable for realtime processing. Of notable mention in this regard are the works of Johnson *et al.* [15] and Ulyanov *et al.* [35].

Secondly, when a neural network is used to overcome the computational burden of [8], training of a model for every desired style image is required due to the limited-capacity of conventional models in encoding multiple styles into the weights of the network. This greatly narrows down the applicability of the method for use cases where the concept of style cannot be defined a-priori and needs to be inferred from examples. With respect to this second limitation, recent works attempted to separate style and content in feature space (latent space) to allow generalization to a style characterized by an additional input image, or set of images. The most widespread work in this family is AdaIN [11], a specific case of FiLM [30]. The current state of the art allows to control the amount of stylization applied, interpolating between different styles, and using masks to convert different regions of image into different styles [11, 32].

Beyond the study of new network architectures for improving NST, research has resulted in better suited loss functions to train the models. The perceptual loss [8, 15] with a pre-trained VGG19 classifier [33] is very commonly used in this task as it, supposedly, captures high-level features of the image. However, this assumption has been recently challenged in [9]. Cycle-GAN [43] proposed a new cycle consistent loss that does not require one-to-one correspondence between the input and the target images, thus lifting the heavy burden of data annotation.

The problem of image style transfer is challenging, because the style of an image is expressed by both *local* properties (e.g. typical shapes of objects, etc.) and *global* properties (e.g. textures, etc.). Of the many approaches for modelling the content and style of an image that have been proposed in the past, encoding of the information in a lower dimensional latent space has shown very promising results. We therefore advocate to model this hierarchy in latent space by local aggregation of pixel-wise features and by the use of metric learning to separate different styles. To the best of our knowledge, this has not been addressed by previous approaches explicitly.

In the presence of a well structured latent space where style and content are fully separated, transfer can be easily performed by exchanging the style information in latent space between the input and the conditioning style images, without the need to store the transformation in the decoder weights. Such an approach is independent with respect to feature normalization and further avoids the need for rather problematic pre-trained models.

However, the content and style of an image are not completely separable. The content of an image exhibits changes in geometry depending on what style it is painted with. Recently, Kotovenko *et al.* [20] has proposed a content transformer block in an adversarial setting, where the model is trained in two stages. First a style transfer network is optimized. Then it is fixed and content transformer block is optimized instead, learning to account for changes in geometry relevant to a given style. The style exchange therefore becomes two-stage. Modeling such dependence has been shown to improve the visual results dramatically.

This paper addresses the NST setting where style is externally defined by a set of input images to allow transfer from arbitrary domains and to tackle the challenging zeroshot style-transfer scenario by introducing a novel feature regularization layer capable of recombining global and local style content from the input style image. It is achieved by borrowing ideas from geometric deep learning (GDL) [2] and modelling pixel-wise graph of peers on the feature maps in the latent space. To the best of our knowledge, this is the first work that successfully leverages the power of GDL in the style transfer scenario. We successfully demonstrate this in a series of zero-shot style transfer experiments, whose generated result would not be possible if the style was not actually inferred from the respective input images.

This work addresses the aforementioned limitations of NST models by making the following contributions:

- A state-of-the-art approach for NST using a custom graph convolutional layer that recombines style and content in latent space;
- Novel combination of existing losses that allows endto-end training without the need for any pre-trained model (e.g. VGG) to compute perceptual loss;
- Constructing a globally- and locally-combined latent space for content and style information and imposing structure on it by means of metric learning.

## 2. Background

The key component of any NST system is the modeling and extraction of the "style" from an image (though the term is partially subjective). As style is often related to texture, a natural way to model it is to use visual texture modeling methods [39]. Such methods can either exploit texture image statistics (e.g. Gram matrix) [8] or model textures using Markov Random Fields (MRFs) [6]. The following paragraphs provide an overview of the style transfer literature introduced by [14].

**Image-Optimization-Based Online Neural Methods.** The method from Gatys *et al.* [8] may be the most representative of this category. While experimenting with representations from intermediate layers of the VGG-19 network, the authors observed that a deep convolutional network is able to extract image content from an arbitrary photograph, as well as some appearance information from works of art. The content is represented by a low-level layer of VGG-19, whereas the style is expressed by a combination of activations from several higher layers, whose statistics are described using the network features Gram matrix. Li [25] later pointed out that the Gram matrix representation can be generalized using a formulation based on Maximum Mean Discrepacy (MMD). Using MMD with a quadratic polynomial kernel gives results very similar to the Gram matrix-based approach, while being computationally more efficient. Other non-parametric approaches based on MRFs operating on patches were introduced by Li and Wand [21].

**Model-Optimization-Based Offline Neural Methods.** These techniques can generally be divided into several subgroups [14]. *One-Style-Per-Model* methods need to train a separate model for each new target style [15, 35, 22], rendering them rather impractical for dynamic and frequent use. A notable member of this family is the work by Ulyanov *et al.* [36] introducing Instance Normalization (IN), better suited for style-transfer applications than Batch Normalization (BN).

*Multiple-Styles-Per-Model* methods attempt to assign a small number of parameters to each style. Dumoulin [5] proposed an extension of IN called Conditional Instance Normalization (CIN), StyleBank [3] learns filtering kernels for different styles, and other works instead feed the style and content as two separate inputs [23, 40] similarly to our approach.

Arbitrary-Style-Per-Model methods either treat the style information in a non-parametric, i.e. as in StyleSwap [4], or parametric manner using summary statistics, such as in Adaptive Instance Normalization (AdaIN) [11]. AdaIN, instead of learning global normalization parameters during training, uses first moment statistics of the style image features as normalization parameters. Later, Li *et al.* [24] introduced a variant of AdaIN using Whitening and Coloring Transformations (WTC). Going towards zero-shot style transfer, ZM-Net [38] proposed a transformation network with dynamic instance normalization to tackle the zero-shot transfer problem. More recently, Avatar-Net [32] proposed the use of a "style decorator" to re-create content features by semantically aligning input style features with those derived from the style image.

**Other methods.** Cycle-GAN [43] introduced a cycleconsistency loss on the reconstructed images that delivers very appealing results without the need for aligned input and target pairs. However, it still requires one model per style. The approach was extended in Combo-GAN [1], which lifted this limitation and allowed for a practical multistyle transfer; however, also this method requires a decoderper-style. Sanakoyeu *et al.* [31] observed that the applying the cycle consistency loss in image space might be overrestricting the stylization process. They also show how to use a set of images, rather than a single one, to better express the style of an artwork. In order to provide more accurate style transfer with respect to an image content, Kotovenko *et al.* [20] designed the so called content transformer block, an additional sub-network that is supposed to finetune a trained style transfer model to a particular content. The strict consistency loss was later relaxed by MU-NIT [12], a multi-modal extension of UNIT [26], which imposes it in latent space instead, providing more freedom to the image reconstruction process. Later, the same authors proposed FUNIT [27], a few-shot extension of MUNIT.

Towards exploring the disentanglement of different styles in the latent space, Kotovenko *et al.* [19] proposed a fixpoint triplet loss in order to perform metric learning in the style latent space, showing how to separate two different styles within a single model.

# 3. Method

The core idea of our work is a region-based mechanism that exchanges the style between input and target style images, similarly to StyleSwap [4], while preserving the semantic content. To successfully achieve this, style and content information must be well separated, disentangled. We advocate the use of metric learning to directly enforce separation among different styles, which has been experimentally shown to greatly reduce the amount of style dependent information retained in the decoder. Furthermore, in order to account for geometric changes in content that are bound to a certain style, we model the style transfer as a two-stage process, first performing style transfer, and then in the second step modifying the content geometry accordingly. This is done using the Two-stage Peer-regularized Feature Recombination (TPFR) module presented below.

#### **3.1.** Architecture and losses

The proposed system architecture is shown in Figure 2. To prevent the main decoder from encoding the stylization in its weights, auxiliary decoder [7] is used during training to optimize the parameters of the encoder and decoder independently. The yellow module in Figure 2 is trained as an autoencoder (AE) [29, 41, 28] to reconstruct the input. The green module, instead, is trained as a GAN[10] to generate the stylized version of the input using the encoder from the yellow modules is interleaved together with the discriminator. Additionally, following the analysis from Martineau *et al.* [16], the Relativistic Average GAN (Ra-GAN) is used as our adversarial loss formulation, which has shown to be more stable and to produce more natural-looking images than traditionally used GAN losses.



Figure 2. The proposed architecture with two decoder modules. The yellow decoder is the auxiliary decoder, while the main decoder is depicted in green. Dashed lines indicate lack of gradient propagation.

Let us now describe the four main building blocks of our approach in detail<sup>1</sup>. We denote  $x_i, x_t, x_f$  an input image, a target and fake image, respectively. Our model consists of an encoder  $E(\cdot)$  generating the latent representations, an auxiliary decoder  $\tilde{D}(\cdot)$  taking a single latent code as input, a main decoder  $D(\cdot)$  taking one latent code as input, and TPFR module  $T(\cdot, \cdot)$  receiving two latent codes. Generated latent codes are denoted  $z_i = E(x_i), z_t = E(x_t)$ . We further denote  $(\cdot)_C, (\cdot)_S$  the content and style part of the latent code, respectively.

The distance f between two latent representations is defined as the smooth L1 norm [13] in order to stabilize training and to avoid exploding gradients:

$$f[d] = \frac{1}{W \times H} \sum_{i=1}^{W \times H} \sum_{j=1}^{N} d_{i,j},$$

$$d_{i,j} = \begin{cases} 0.5 \, d_{i,j}^2 & \text{if } d_{i,j} < 1 \\ |d_{i,j}| - 0.5 & \text{otherwise} \end{cases},$$
(1)

where  $d = z_1 - z_2$ ,  $z_1$  and  $z_2$  are two different feature embeddings with N channels and  $W \times H$  is the spatial dimension of each channel.

**Encoder.** The encoder used to produce latent representation of all input images is composed of several stridedconvolutional layers for downsampling followed by multiple ResNet blocks. The latent code z is composed by two parts: the content part,  $(z)_C$ , which holds information about the image content (e.g. objects, position, scale, etc.), and the style part,  $(z)_S$ , which encodes the style that the content is presented in (e.g. level of detail, shapes, etc.). The style component of the latent code  $(z)_S$  is further split equally into  $(z)_S = [(z)_S^{loc}, (z)_S^{glob}]$ . Here,  $(z)_S^{loc}$  encodes local style information per pixel of each feature map, while  $(z)_{S}^{glob}$  undergoes further downsampling via a small subnetwork to generate a single value per feature map.

**Auxiliary decoder.** The Auxiliary decoder reconstructs an image from its latent representation and is used only during training to train the encoder module. It is composed of several ResNet blocks followed by fractionally-strided convolutional layers to reconstruct the original image. The loss is inspired by [31, 19] and is composed of the following parts. A content feature cycle loss that pulls together latent codes representing the same content:

$$L_{z_{cont}} = f[E(D(T(z_i, z_t)))_C - (z_i)_C] + f[E(D(T(z_i, z_i)))_C - (z_i)_C]$$
(2)

A metric learning loss, enforcing clustering of the style part of the latent representations:

$$\begin{split} L_{z_{style}}^{pos} &= f[(z_{i_1})_S - (z_{i_2})_S] + f[(z_{t_1})_S - (z_{t_2})_S] \\ L_{z_{style}}^{neg} &= f[(z_{i_1})_S - (z_{t_1})_S] + f[(z_{i_2})_S - (z_{t_2})_S] \\ L_{z_{style}} &= L_{z_{style}}^{pos} + max(0.0, \mu - L_{z_{style}}^{neg}). \end{split}$$
(3)

A classical reconstruction loss used in autoencoders, which forces the model to learn perfect reconstructions of its inputs:

$$\tilde{L}_{idt} = f[\tilde{D}(E(x_i)) - x_i] + f[\tilde{D}(E(x_t)) - x_t].$$
 (4)

And a latent cycle loss enforcing the latent codes of the inputs to be the same as latent codes of the reconstructed images:

$$\tilde{L}_{z_{cycle}} = f[E(\tilde{D}(z_i)) - (z_i)] + f[E(\tilde{D}(z_t)) - (z_t)],$$
 (5)

which has experimentally shown to stabilize the training. The total auxiliary decoder loss  $L_{\tilde{D}}$ , is then defined as:

$$L_{\tilde{D}} = L_{z_{cont}} + L_{z_{style}} + \tilde{L}_{z_{cycle}} + \lambda \tilde{L}_{idt}, \qquad (6)$$

where  $\lambda$  is a hyperparameter weighting importance of the reconstruction loss  $\tilde{L}_{idt}$  with respect to the rest.

<sup>&</sup>lt;sup>1</sup>The full architecture details are found in the supplementary material.



Figure 3. One step of peer regularization takes as input latent representations of content and style images. The content part of the latent representation is used to induce a graph structure on the style latent space, which is then used to recombine the style part of the content image's latent representation. This results in a new style latent code. The Two-Stage Peer Regularization Layer performs the peer regularization operation twice. In the second step, the roles of content and style information are swapped.

**Main decoder.** This network replicates the architecture of the auxiliary decoder, and uses the output of the Two-stage Peer-regularized Feature Recombination module (see Section 3.2). During training of the main decoder the encoder is kept fixed, and the decoder is optimized using a loss function composed of the following parts. First, the decoder adversarial loss:

$$L_{gen} = \mathbb{E}_{x_i \sim \mathbb{P}} \left[ \left( C\left(x_i\right) - \mathbb{E}_{x_f \sim \mathbb{Q}} C\left(x_f\right) + 1 \right)^2 \right] \\ + \mathbb{E}_{x_f \sim \mathbb{Q}} \left[ \left( \mathbb{E}_{x_i \sim \mathbb{P}} C\left(x_i\right) - C\left(x_f\right) + 1 \right)^2 \right],$$
(7)

where  $\mathbb{P}$  is the distribution of the real data and  $\mathbb{Q}$  is the distribution of the generated (fake) data and C is the discriminator.

In order to enforce the stylization preserve the content part of the latent codes while recombining the style part of the latent codes to represent the target style class, we use so called transfer latent cycle loss:

$$L_{z_{transf}} = f[E(D(T(z_i, z_t)))_C - (z_i)_C] + f[E(D(T(z_i, z_t)))_S - (z_t)_S].$$
(8)

Further, to make the main decoder learn to reconstruct the original inputs, we employ the classical reconstruction loss as we did for auxiliary decoder as well:

$$L_{idt} = f[D(T(z_i, z_i)) - x_i] + f[D(T(z_t, z_t)) - x_t].$$
 (9)

The above put together composes the main decoder loss  $L_D$ , where the reconstruction loss  $L_{idt}$  is weighted by the same hyperparameter  $\lambda$  used also in  $L_{\tilde{D}}$ .

$$L_D = L_{gen} + L_{z_{transf}} + \lambda L_{idt} \tag{10}$$

**Discriminator.** The discriminator is a convolutional network receiving two images concatenated over the channel dimension and producing an  $N \times N$  map of predictions. The first image is the one to discriminate, whereas the second one serves as conditioning for the style class. The output prediction is ideally 1 if the two inputs come from the same style class and 0 otherwise. The discriminator loss is defined as:

$$L_{C} = \mathbb{E}_{x_{i} \sim \mathbb{P}} \left[ \left( C\left(x_{i}\right) - \mathbb{E}_{x_{f} \sim \mathbb{Q}} C\left(x_{f}\right) - 1 \right)^{2} \right] + \mathbb{E}_{x_{f} \sim \mathbb{Q}} \left[ \left( \mathbb{E}_{x_{i} \sim \mathbb{P}} C\left(x_{i}\right) - C\left(x_{f}\right) - 1 \right)^{2} \right].$$
(11)

## 3.2. Two-stage Peer-regularized Feature Recombination (TPFR)

The TPFR module draws inspiration from PeerNets [34] and Graph Attention Layer (GAT) [37] to perform style transfer in latent space, taking advantage of the separation of content and style information (enforced by Equations 2 and 3). Peer-regularized feature recombination is done in two stages, as explained in the following paragraphs.

**Style recombination.** It receives  $z_i = [(z_i)_C, (z_i)_S]$  and  $z_t = [(z_t)_C, (z_t)_S]$  as an input and computes the k-Nearest-Neighbors (k-NN) between  $(z_i)_C$  and  $(z_t)_C$  using the Euclidean distance to induce the graph of peers.

Attention coefficients over the graph nodes are computed and used to recombine the style portion of  $(z_{out})_S$  as convex combination of its nearest neighbors representations. The content portion of the latent code remains instead unchanged, resulting in:  $z_{out} = [(z_i)_C, (z_{out})_S]$ .

Given a pixel  $(\mathbf{z_m})_{\mathbf{C}}$  of feature map m, its k-NN graph in the space of d-dimensional feature maps of all pixels of all peer feature maps  $n_k$  is considered. The new value of the style part  $(z)_S$  for the pixel is expressed as:

$$(\tilde{\mathbf{z}}_p^m)_S = \sum_{k=1}^K \alpha_{mn_k pq_k} (\mathbf{z}_{q_k}^{n_k})_S, \qquad (12)$$

$$\alpha_{mn_kpq_k} = \frac{\text{LReLU}(\exp(a((\mathbf{z}_p^m)_C, (\mathbf{z}_{q_k}^{n_k})_C))))}{\sum_{k'=1}^{K} \text{LReLU}(\exp(a((\mathbf{z}_p^m)_C, (\mathbf{z}_{q_{k'}}^{n_{k'}})_C))))}$$
(13)

where  $a(\cdot, \cdot)$  denotes a fully connected layer mapping from 2*d*-dimensional input to scalar output, and  $\alpha_{mn_kpq_k}$  are attention scores measuring the importance of the  $q_k$ th pixel of feature map n to the output pth pixel  $\tilde{\mathbf{x}}_p^m$  of feature map m. The resulting style component of the input feature map  $\tilde{\mathbf{X}}^m$  is the weighted pixel-wise average of its peer pixel features defined over the style input image.

**Content recombination.** Once the style latent code is recombined, an analogous process is repeated to transform the content latent code according to the new style information. In this case, it starts off with inputs  $z_{out} = [(z_i)_C, (z_{out})_S]$  and  $z_t = [(z_t)_C, (z_t)_S]$  and the *k*-NN graph is computed given the style latent codes  $(z_{out})_S, (z_t)_S$ . This graph is used together with Equation 12 in order to compute the attention coefficients and recombine the content latent code as  $(z_{final})_C$ .

The output of the TPFR module therefore is a new latent code  $z_{final} = [(z_{final})_C, (z_{out})_S]$  which recombines both style and content part of the latent code.

#### 4. Experimental setup and Results

The proposed approach is compared against the state-ofthe-art on extensive qualitative evaluations and, to support the choice of architecture and loss functions, ablation studies are performed to demonstrate roles of the various components and how they influence the final result. Only qualitative comparisons are provided as no standard quantitative evaluations are defined for NST algorithms.

#### 4.1. Training

A modification of the dataset collected by [31] is used for training the model. It is composed of a collection of thirteen different painters representing different target styles and selection of relevant classes from the Places365 dataset [42] providing the real photo images. This gives us 624,077 real photos and 4,430 paintings in total. Our network can be trained end-to-end alternating optimization steps for the auxiliary decoder, the main decoder, and the discriminator. The loss used for training is defined as:

$$L = L_C + L_D + L_{\tilde{D}},\tag{14}$$

where *C* is the discriminator, *D* the main decoder, and *D* is the auxiliary decoder (see Section 3.1). Using ADAM [18] as the optimization scheme, with an initial learning rate of 4e-4 and a batch size of 2, training is performed for a total of 200 epochs. After 50 epochs, the learning rate is decayed linearly to zero. Please note that choice of the batch size can be arbitrary, and we choose 2 only due to limited computing resources. In each epoch, we randomly visit 6144 of the real photos. The weighting of the reconstruction identity loss  $\lambda = 25.0$  and the margin for the metric learning  $\mu = 1.0$  during all of our experiments. The training images are cropped and resized to  $256 \times 256$  resolution. Note that during testing, our method can operate on images of arbitrary size.

#### 4.2. Style transfer

A set of test images from Sanakoyeu [31] is stylized and compared against competing methods in Figure 4 (inputs of size  $768 \times 768$  px) to demonstrate arbitrary stylization of a content image given several different styles. It is important to note that, as opposed to majority of the competing methods, our network does not require retraining for each style and allows therefore also for transfer of previously unseen styles, e.g. Pissarro in row 5 of Figure 4, which is a painter that has not been present in the training set.

Qualitative results<sup>2</sup> that were done on color images of size  $512 \times 512$ , using previously unseen paintings from painters that are in the training set, are shown in Figure 5(b). It is worth noticing that our approach can deal also with very abstract painting styles, such as the one of Jackson Pollock (Figure 5(b), row 3, style 1). It motivates the claim that our model generalizes well to many different styles, allowing zero-shot style transfer.

**Zero-shot style transfer.** In order to support our claims regarding zero-shot style transfer, we have collected samples of a few painters that were not seen during training. In particular, we collected paintings from Salvador Dali, Camille Pissarro, Henri Matisse, Katigawa Utamaro and Rembrandt. The evaluation presented in Figure 5(a) shows that our approach is able to account for fine details in painting style of Camille Pissarro (row 1, style 1), as well as create larger flat regions to recombine style that is own to Katigawa Utamaro (row 6, style 2). We find the results of zero-shot style transfer using the aforementioned painters very encouraging.

**Ablation study.** There are several key components in our solution which make arbitrary style transfer with a single model and end-to-end training possible. The effect of suppressing each of them during the training is examined, and results for the various models are compared, highlighting the importance of each component in Figure 6. Clearly, the *auxiliary decoder* used during training is a centerpiece of the whole approach, as it prevents degenerate solutions. We observe that training encoder directly with the main decoder end-to-end does not work (Fig. 6 *NoAux*). *Separation of the* 

<sup>&</sup>lt;sup>2</sup>More results are shown in the supplementary material.



Figure 4. Qualitative comparison with respect to other state-of-the-art methods. It should be noted that most of the compared methods had to train a new model for each style. While providing competitive results, our method performs arbitrary style transfer using a single model.

*latent code* into content and style part allows for the introduced two-stage style transfer and is important to account for changes in shape of the objects for styles like, e.g. Picasso (Fig. 6 NoSep). Two-stage recombination provides better generalization to variety of styles (Fig. 6 NoTS). Performing only exchange of style based on content features completely fails in some cases (e.g. row 1 in Fig. 6). Next, *metric learning* on the style latent space enforces its better clustering and enhances some important details in the stylized image (Fig. 6 NoML). Last but not least, the combined *local and global* style latent code is important in order to be able to account for changes in edges and brushstrokes appropriately (Fig. 6 NoGlob).

# 5. Conclusions

We propose a novel model for neural style transfer which mitigates various limitations of current state-of-theart methods and that can be used also in the challenging zero-shot transfer setting. This is thanks to a Two-Stage Peer-Regularization Layer using graph convolutions to recombine the style component of the latent representation and with a metric learning loss enforcing separation of different styles combined with cycle consistency in feature space. An auxiliary decoder is also introduced to prevent degenerate solutions and to enforce enough variability of the generated samples. The result is a state-of-the-art method that can be trained end-to-end without the need of a pre-trained model to compute the perceptual loss, therefore lifting recent concerns regarding the reliability of such features for NST. More importantly the proposed method requires only a single encoder and a single decoder to perform transfer among arbitrary styles, contrary to many competing methods requiring a decoder (and possibly an encoder) for each input and target pair. This makes our method more applicable to real-world image generation scenarios where users define their own styles.



(a) Zero-shot style transfer.

(b) Styles seen during training.

Figure 5. Qualitative evaluation of our method for previously unseen styles (left) and for styles seen during training (right). It can be observed that the generated images are consistent with the provided target style (inferred from a single sample only), showing the good generalization capabilities of the approach.



Figure 6. Ablation study evaluating different architecture choices for our approach. Detail of style transfer for row 2 is shown in the rightmost column. *Ours* refers to the final approach, *NoAux* makes no use of the auxiliary decoder during training, *NoSep* ignores the separation of content and style in the latent space during feature exchange in TPFR module and exchanges the whole latent code at once, *NoTS* recombines only style features based on content and keeps the original content features, *NoML* does not use metric learning and *NoGlob* does not separate the style latent code into a global and local part.

# References

- Asha Anoosheh, Eirikur Agustsson, Radu Timofte, and Luc Van Gool. Combogan: Unrestrained scalability for image domain translation. *CoRR*, abs/1712.06909, 2017. 1, 3
- M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Process. Mag.*, 34(4):18–42, 2017.
- [3] Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, and Gang Hua. Stylebank: An explicit representation for neural image style transfer. *CoRR*, abs/1703.09210, 2017. 3
- [4] Tian Qi Chen and Mark Schmidt. Fast patch-based style transfer of arbitrary style. *CoRR*, abs/1612.04337, 2016. 3
- [5] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *CoRR*, abs/1610.07629, 2016. 3
- [6] A. A. Efros and T. K. Leung. Texture synthesis by nonparametric sampling. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 1033–1038 vol.2, Sep. 1999. 2
- [7] Jeffrey De Fauw, Sander Dieleman, and Karen Simonyan. Hierarchical autoregressive image models with auxiliary decoders. *CoRR*, abs/1903.04933, 2019. 3
- [8] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. A neural algorithm of artistic style. *CoRR*, abs/1508.06576, 2015. 1, 2
- [9] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*, 2019. 1, 2
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *NIPS*. 2014. 3
- [11] Xun Huang and Serge J. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. *CoRR*, abs/1703.06868, 2017. 2, 3
- [12] Xun Huang, Ming-Yu Liu, Serge J. Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. *CoRR*, abs/1804.04732, 2018. 1, 3
- [13] Peter J. Huber. Robust estimation of a location parameter. Ann. Math. Statist., 35(1):73–101, 03 1964. 4
- [14] Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, and Mingli Song. Neural style transfer: A review. *CoRR*, abs/1705.04058, 2017. 2, 3
- [15] Justin Johnson, Alexandre Alahi, and Fei-Fei Li. Perceptual losses for real-time style transfer and super-resolution. *CoRR*, abs/1603.08155, 2016. 2, 3
- [16] Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard GAN. CoRR, abs/1807.00734, 2018. 3
- [17] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *CoRR*, abs/1812.04948, 2018.

- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. 6
- [19] Dmytro Kotovenko, Artsiom Sanakoyeu, Sabine Lang, and Bjorn Ommer. Content and style disentanglement for artistic style transfer. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4422–4431, 2019. 3, 4
- [20] Dmytro Kotovenko, Artsiom Sanakoyeu, Pingchuan Ma, Sabine Lang, and Bjorn Ommer. A content transformation block for image style transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10032–10041, 2019. 2, 3
- [21] Chuan Li and Michael Wand. Combining markov random fields and convolutional neural networks for image synthesis. *CoRR*, abs/1601.04589, 2016. 3
- [22] Chuan Li and Michael Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. *CoRR*, abs/1604.04382, 2016. 3
- [23] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Diversified texture synthesis with feed-forward networks. *CoRR*, abs/1703.01664, 2017. 3
- [24] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. *CoRR*, abs/1705.08086, 2017. 3
- [25] Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. Demystifying neural style transfer. *CoRR*, abs/1701.01036, 2017. 3
- [26] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. *CoRR*, abs/1703.00848, 2017. 3
- [27] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. arXiv preprint arXiv:1905.01723, 2019. 3
- [28] Xiao-Jiao Mao, Chunhua Shen, and Yu-Bin Yang. Image restoration using convolutional auto-encoders with symmetric skip connections. *CoRR*, 2016. 3
- [29] Jonathan Masci, Ueli Meier, Dan Cireşan, and Jürgen Schmidhuber. Stacked Convolutional Auto-Encoders for Hierarchical Feature Extraction, pages 52–59. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011. 3
- [30] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. Film: Visual reasoning with a general conditioning layer. *CoRR*, abs/1709.07871, 2017.
   2
- [31] Artsiom Sanakoyeu, Dmytro Kotovenko, Sabine Lang, and Björn Ommer. A style-aware content loss for real-time HD style transfer. *CoRR*, abs/1807.10201, 2018. 3, 4, 6
- [32] Lu Sheng, Ziyi Lin, Jing Shao, and Xiaogang Wang. Avatarnet: Multi-scale zero-shot style transfer by feature decoration. In *CVPR*, 2018. 2, 3
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 2
- [34] Jan Svoboda, Jonathan Masci, Federico Monti, Michael M. Bronstein, and Leonidas J. Guibas. Peernets: Exploiting peer wisdom against adversarial attacks. *CoRR*, abs/1806.00088, 2018. 5

- [35] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S. Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. *CoRR*, abs/1603.03417, 2016. 2, 3
- [36] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. *CoRR*, abs/1701.02096, 2017. 3
- [37] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Alejandro Romero, Pietro Lió, and Yoshua Bengio. Graph attention networks. *CoRR*, abs/1710.10903, 2018. 5
- [38] Hao Wang, Xiaodan Liang, Hao Zhang, Dit-Yan Yeung, and Eric P. Xing. Zm-net: Real-time zero-shot image manipulation network. *CoRR*, abs/1703.07255, 2017. 3
- [39] Li-Yi Wie, Sylvain Lefebvre, Vivek Kwatra, and Greg Turk. State of the Art in Example-based Texture Synthesis. In M. Pauly and G. Greiner, editors, *Eurographics 2009 - State of the Art Reports*, 2009. 2
- [40] Hang Zhang and Kristin J. Dana. Multi-style generative network for real-time transfer. *CoRR*, abs/1703.06953, 2017.
   3
- [41] Junbo Jake Zhao, Michal Mathieu, Ross Goroshin, and Yann LeCun. Stacked what-where auto-encoders. *CoRR*, 2015. 3
- [42] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis* and machine intelligence, 40(6):1452–1464, 2017. 6
- [43] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycleconsistent adversarial networks. *CoRR*, abs/1703.10593, 2017. 1, 2, 3