Domain Enhanced Arbitrary Image Style Transfer via Contrastive Learning

Yuxin Zhang
NLPR, Institute of Automation, CAS
School of Artificial Intelligence, UCAS
China
zhangyuxin2020@ia.ac.cn

Fan Tang*
School of Artificial Intelligence,
Jilin University
China
tangfan@jlu.edu.cn

Weiming Dong*
NLPR, Institute of Automation, CAS
School of Artificial Intelligence, UCAS
China
weiming.dong@ia.ac.cn

Haibin Huang
Kuaishou Technology
China
huanghaibin03@kuaishou.com

Chongyang Ma
Kuaishou Technology
China
chongyangma@kuaishou.com

Tong-Yee Lee
National Cheng-Kung University
Taiwan
tonylee@ncku.edu.tw

Changsheng Xu
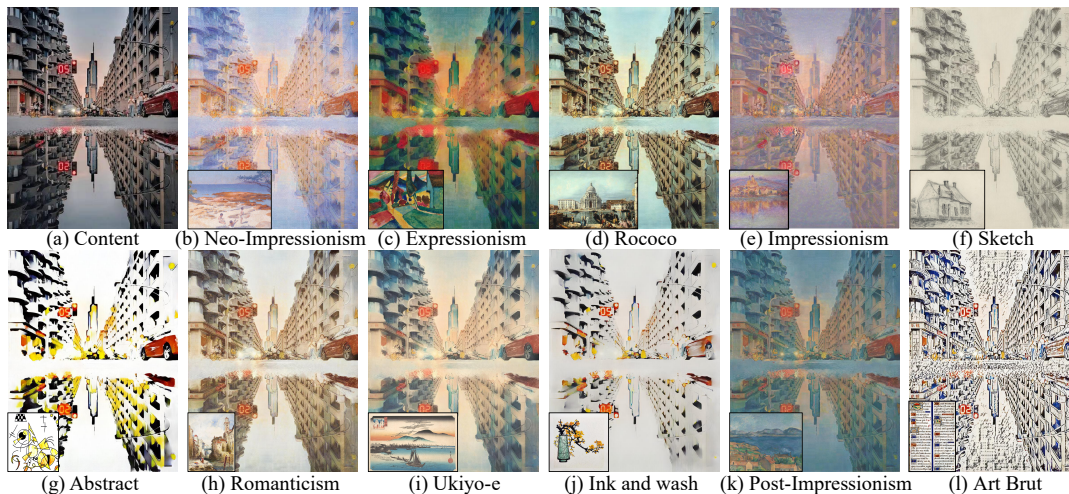NLPR, Institute of Automation, CAS
csxu@nlpr.ia.ac.cn

Figure 1: Style transfer results using our method, which can robustly and effectively handle various painting styles. The input content image is in the top-left corner and the style reference is shown as the inset for each result. Our method can faithfully capture the style of each painting and generate a result with a unique artistic visual appearance. Style image credits: {(b) Henri Edmond Cross, (c) Vasily Kandinsky, (d) Michele Marieschi, (e) Claude Monet, (h) Richard Parkes Bonington, (i) Utagawa Hiroshige, (k) Paul Cezanne}/The Art Institute of Chicago (CC0), (f) Vincent van Gogh/National Gallery of Art (CC0).

ABSTRACT

In this work, we tackle the challenging problem of arbitrary image style transfer using a novel style feature representation learning method. A suitable style representation, as a key component in image stylization tasks, is essential to achieve satisfactory results. Existing deep neural network based approaches achieve reasonable results with the guidance from second-order statistics such as Gram matrix of content features. However, they do not leverage sufficient style information, which results in artifacts such as local distortions and style inconsistency. To address these issues, we propose to learn style representation directly from image features instead of their second-order statistics, by analyzing the similarities and differences between multiple styles and considering the style distribution. Specifically, we present Contrastive Arbitrary Style Transfer (CAST), which is a new style representation learning and style transfer method via contrastive learning. Our framework

* Fan Tang and Weiming Dong are co-corresponding authors.

consists of three key components, i.e., a multi-layer style projector for style code encoding, a domain enhancement module for effective learning of style distribution, and a generative network for image style transfer. We conduct qualitative and quantitative evaluations comprehensively to demonstrate that our approach achieves significantly better results compared to those obtained via state-of-the-art methods. Code and models are available at https://github.com/zyxElsa/CAST_pytorch.

CCS CONCEPTS

• **Computing methodologies → Image processing**.

KEYWORDS

Arbitrary style transfer, contrastive learning, style encoding

1 INTRODUCTION

If a picture is worth a thousand words, then an artwork tells the whole story. Art styles, which describe the way the artwork looks, are the manner in which the artist portrays his or her subject matter and how the artist expresses his or her vision. Style is determined by the characteristics that describe the artwork, such as the way the artist employs form, color, and composition. Artistic style transfer, as an efficient way to create a new painting by combining the content of a natural images and the style of an existing painting image, is a major research topic in computer graphics and computer vision [Jing et al. 2020b; Liao et al. 2017], with style representation as the most important issue.

Since Gatys et al. [2016] proposed to use Gram matrix as artistic style representation, high-quality visual results are generated by advanced neural style transfer networks. Despite the remarkable progress made in the field of arbitrary image style transfer, the second-order feature statistics (Gram matrix or mean/variance) style representation has restricted the further development and application. As shown in Figure 1, the appearances of different artwork styles vary considerably in terms of not only the colors and local textures but also the layouts and compositions. Figures 2d and 2e show the results of two recently proposed state-of-the-art style transfer approaches. We obverse that aligning the distributions of neural activation between images using second-order statistics results in difficulty to capture the color distribution or the special layouts, or imitate specific detailed brush effects of different styles.

In this paper, we revisit the core problem for neural style transfer, that is, the proper artistic style representation. The widely used second-order statistics as a global style descriptor can distinguish styles to some extent, but they are not the optimal way to represent styles. By second-order statistics, arbitrary stylization formulates styles through artificially designed image features and loss functions in a heuristic manner. In other words, the network learns

(a) Content   (b) Style   (c) Ours   (d) AdaAttN   (e) IEST

Figure 2: Compared with AdaAttnN [Liu et al. 2021b] and IEST [Chen et al. 2021a] which rely on second-order statistics, our method can faithfully transfer styles while ensuring structural consistency with the content images. Style image (1st row) credit: Claude Monet/AIC (CC0).

to fit the second-order statistics of the style image and generated image, instead of the style itself. Exploring the relationship and distribution of styles directly from artistic images instead of using pre-defined style representations is worthwhile.

Toward this end, we propose to improve arbitrary style transfer with a novel style representation by contrastive learning-based optimization. Our key insight is that a person without artistic knowledge has difficulty defining the style if only one artistic image is given, but identifying the difference between different styles is relatively easy. Specifically, we present a novel **C**ontrastive **A**rbitrary **S**tyle **T**ransfer (**CAST**) framework for image style representation and style transfer. CAST consists of a backbone based on an encoder-transformation-decoder structure, a **m**ulti-layer **s**tyle **p**rojector (**MSP**) module, and a **d**omain **e**nhancement (**DE**) module. We introduce contrastive learning to consider the positive and negative relationships between styles, and we use DE to learn the distribution of overall art image domains. To capture the style features at various scales, our MSP module projects the features of each layer of the style image to the corresponding style encoding space. Our contribution can be summarized as follows:

- We propose an MSP module for style encoding and a novel CAST model for encoder-transformation-decoder-based arbitrary style transfer without using the second-order statistics as style representations.
- We introduce contrastive learning and domain enhancement by considering the relationships between positive and negative examples as well as the global distribution of styles, which solves the problem that existing style transfer models cannot fully utilize a large amount of style information.
- Experiments show that our method achieves state-of-the-art style transfer results in terms of visual quality. A challenging subjective survey was conducted, as inspired by the Turing test, to show that output of CAST could mislead participants from telling the fake painting images from real ones.

2 RELATED WORK

Image style transfer. Traditional style transfer methods such as stroke-based rendering [Fišer et al. 2016] and image filtering [Wang et al. 2004] typically use low-level hand-crafted features. Gatys
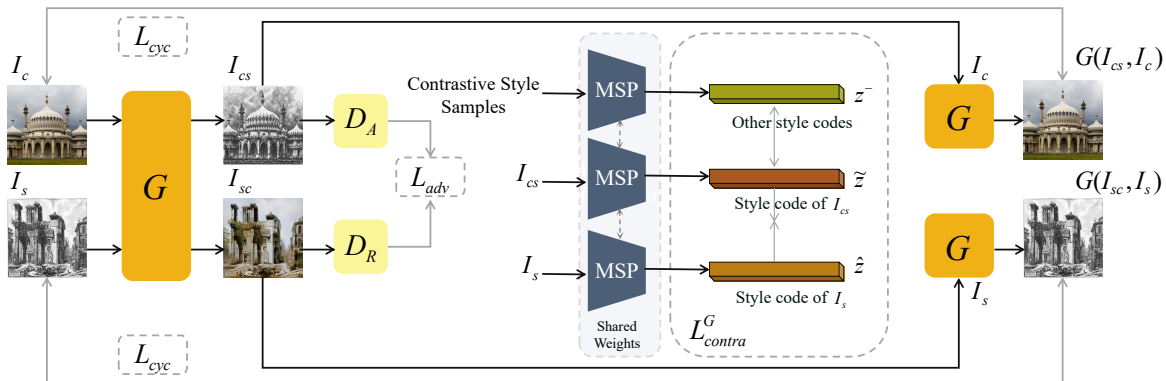
Figure 3: CAST consists of an encoder-transformation-decoder-based generator G, a multi-layer style projector (MSP) module, and a domain enhancement module. We first generate images I_{cs} and I_{sc} from the content image I_c and the style image I_s using the generator. Then, I_{cs} and I_s are fed into the MSP module to generate the corresponding style code $\tilde{z}$ and $\hat{z}$, which will be used as positive samples in the style contrastive learning process. The style codes z^- of other artistic images in the style bank will be used as negative samples. We compute a contrastive style loss $\mathcal{L}_{contra}^{G}$ based on these style codes. DE module is based on the adversarial loss $\mathcal{L}_{adv}$ and the cycle consistency loss $\mathcal{L}_{cyc}$. Style image credit: Giovanni Battista Piranesi/AIC (CC0).

et al. [2016] and the follow-up variants [Gatys et al. 2017; Kolkin et al. 2019] demonstrate that the statistical distribution of features extracted from pre-trained deep convolutional neural networks can capture style patterns effectively. Although the results are remarkable, these methods formulate the task as a complex optimization problem, which leads to high computational cost. Some recent approaches rely on a learnable neural network to match the statistical information in feature space for efficiency. Per-style-per-model methods [Gao et al. 2020; Johnson et al. 2016; Puy and Pérez 2019] train a specific network for each individual style. Multiple-style-per-model methods [Chen et al. 2017; Dumoulin et al. 2017; Ulyanov et al. 2016; Zhang and Dana 2018] represent multiple styles using one single model.

Arbitrary style transfer methods [Deng et al. 2022, 2020; Li et al. 2017; Svoboda et al. 2020; Wu et al. 2021a] build more flexible feed-forward architectures to handle an arbitrary style using a unified model. AdaIN [Huang and Belongie 2017] and DIN [Jing et al. 2020a] directly align the overall statistics of content features with the statistics of style features and adopt conditional instance normalization. However, dynamic generation of affine parameters in the instance normalization layer may cause distortion artifacts. Instead, several methods follow the encoder-decoder manner, where feature transformation and/or fusion is introduced into an auto-encoder-based framework. For example, Li et al. [2019] learn a cross-domain feature linear transformation matrix (LST) to enable universal style transfer and generate the desired stylization results by decoding from the transformed features. Park et al. [2019] introduce SANet to flexibly match the semantically nearest style features onto the content features. Deng et al. [2021] propose MCCNet to fuse exemplar style features and input content features by multi-channel correlation for efficient style transfer. An et al. [2021] propose reversible neural flows and an unbiased feature transfer module (ArtFlow) to prevent content leak during universal style transfer. Liu et al. [2021b] present an adaptive attention normalization module (AdaAttN) to consider both shallow and deep features for attention

score calculation. GAN-based methods [Kotovenko et al. 2019a,b; Sanakoyeu et al. 2018a; Svoboda et al. 2020; Zhu et al. 2017] have been successfully used in collection style transfer, which considers style images in a collection as a domain [Chen et al. 2021b; Lin et al. 2021; Xu et al. 2021].

Contrastive learning. Contrastive learning has been used in many applications, such as image dehazing [Wu et al. 2021b], context prediction [Santa Cruz et al. 2019], geometric prediction [Liu et al. 2019] and image translation. Contrastive learning is introduced in image translation to preserve the content of the input [Han et al. 2021] and reduce mode collapse [Jeong and Shin 2021; Kang and Park 2020; Liu et al. 2021a]. CUT [Park et al. 2020] proposes patch-wise contrastive learning by cropping input and output images into patches and maximizing the mutual information between patches. Following CUT, TUNIT [Baek et al. 2021] adopts contrastive learning on images with similar semantic structures. However, the semantic similarity assumption does not hold for arbitrary style transfer tasks, which leads the learned style representations to a significant performance drop. IEST [Chen et al. 2021a] applies contrastive learning to image style transfer based on feature statistics (mean and standard deviation) as style priors. The contrastive loss is calculated only within the generated results. Contrastive learning in IEST is an auxiliary method to associate stylized images sharing the same style, and the ability comes from the feature statistics from pre-trained VGG. Differently, we introduce contrastive learning for style representation by proposing a novel framework that uses visual features comprehensively to represent style for the task of arbitrary image style transfer.

3 METHOD

As shown in Figure 3, our framework consists of three key components: (1) a multi-layer style projector which is trained to project features of artistic image into style code; (2) a contrastive style learning module which is applied to guide both the training of the

multi-layer style projector and the style image generation; and (3) a domain enhancement scheme to further help learn the distribution of artistic image domain. All these components are used for learning style representations to measure the difference between the input artistic images and generated results and thus, they could be applied to different kinds of arbitrary style transfer networks.

3.1 Multi-layer Style Projector

Our goal is to develop an arbitrary style transfer framework that can capture and transfer the local stroke characteristics and overall appearance of an artistic image to a natural image. A key component is to find a suitable style representation which can be used to distinguish different styles and further guide the generation of style images. To this end, we design an MSP module, which includes a style feature extractor and a multi-layer projector. Instead of using features from a specific layer or a fusion of multiple layers, our MSP projects features of different layers into separate latent style spaces to encode local and global style cues.

Specifically, we adapt VGG-19 [Simonyan and Zisserman 2014] and finetune the VGG-19 model pre-trained on ImageNet with a collection of 18,000 artistic images in 30 categories. We then select M layers of feature maps in VGG-19 as input to our multi-layer projector (we use layers of ReLU1_2, ReLU2_2, ReLU3_3, and ReLU4_3 in all experiments). We use max pooling and average pooling to capture the mean and peak value of features. The multi-layer projector consists of pooling, convolution, and several multilayer perceptron layers, and it projects the style features into a set of K-dimensional latent style code, as shown in Figure 4.

After training, MSP can encode an artistic image into a set of latent style code $\{z_i | i \in [1, M], z_i \in \mathbb{R}^K\}$, which can be plugged into an existing style transfer network (i.e., replacing the mean and variance in AdaIN [Huang and Belongie 2017]) as the guidance for stylization. Next, we will describe how to jointly train MSP and style transfer networks with a contrastive learning strategy.

3.2 Contrastive Style Learning

As demonstrated above, the style code $\{z_1, z_2, ..., z_M\}$ of an image can be used as the target for MSP training and the guidance for the style transfer network. However, we lack the ground-truth style code for supervised training. Therefore, we adopt contrastive learning and design a new contrastive style loss as an implicit measurement for network training.

When train the MSP module, an image I and its augmented version I^+ (random resizing, cropping, and rotations) are fed into a M-layer style feature extractor, which is the pre-trained VGG-19 network. The extracted style features are then sent to the multi-layer projector, which is an M-layer neural network and maps the style features to a set of K-dimensional vectors $\{z\}$. The contrastive representation learns the visual styles of images by maximizing the mutual information between I and I^+ in contrast to other artistic images within the dataset considered as negative samples $\{I^-\}$. Specifically, the images I, I^+, and N negative samples are respectively mapped into M groups of K-dimensional vectors z, $z^+ \in \mathbb{R}^K$ and $\{z^- \in \mathbb{R}^K\}$. The vectors are normalized to prevent collapsing. Following [Van den Oord et al. 2018], we define the
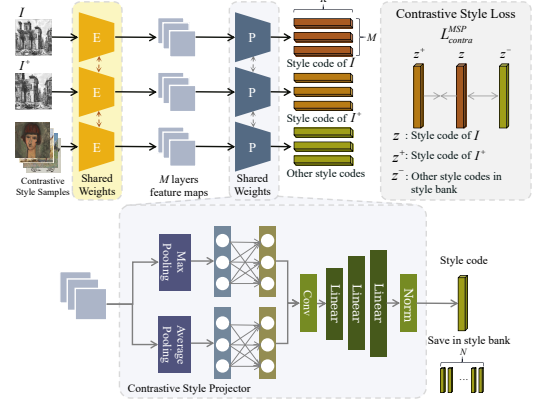

Figure 4: Overview of our MSP module, which includes a VGG-19 based style feature extractor E and a multi-layer projector P. P maps the extracted features to style codes $\{z\}$ which are saved in the style memory bank to compute the contrastive style loss $\mathcal{L}_{contra}^{MSP}$. Image credits: {Giovanni Battista Piranesi, Amedeo Modigliani}/AIC (CC0).

contrastive loss function to train our MSP module as:

$$\mathcal{L}_{contra}^{MSP} = -\sum_{i=1}^{M} \log \frac{\exp(z_i \cdot z_i^+ / \tau)}{\exp(z_i \cdot z_i^+ / \tau) + \sum_{j=1}^{N} \exp(z_i \cdot z_{i_j}^- / \tau)}, \quad (1)$$

where $\cdot$ denotes the dot product of two vectors, and τ is a temperature scaling factor and is set to be 0.07 in all of our experiments. Meanwhile, we maintain a large dictionary of 4096 negative examples using a memory bank architecture following MOCO [He et al. 2020]. It is worth noting that we calculate the contrastive loss between *images*, as opposed to CUT [Park et al. 2020] which adopts contrastive learning by cropping images into patches and maximizing the mutual information between *patches*.

The contrastive representation also provides proper guidance for the generator G to transfer styles between images. We adopt the same form of contrastive loss as used for learning MSP in Eq. (1), but compute the loss using the contrastive representations of the output image I_{cs} and the reference style image I_s, then I_{cs} will have a style similar to I_s:

$$\mathcal{L}_{contra}^{G} = -\sum_{i=1}^{M} \log \frac{\exp(\tilde{z}_i \cdot \hat{z}_i / \tau)}{\exp(\tilde{z}_i \cdot \hat{z}_i / \tau) + \sum_{j=1}^{N} \exp(\tilde{z}_i \cdot z_{i_j}^- / \tau)}, \quad (2)$$

where $\tilde{z}$ and $\hat{z}$ denote the contrastive representation of I_{cs} and I_s, respectively. The negative examples are sampled from the same dictionary used for training of the MSP module. Notably, we take the specific generated and reference images as positive examples and utilize contrastive loss as guidance to transfer styles, which is an one-on-one process. Differently, the contrastive loss in IEST [Chen et al. 2021a] is calculated only within generated results and it takes a set of images as positive examples, which may reduce the style consistency with the given reference (see Figure 2e).

3.3 Domain Enhancement

We introduce DE with adversarial loss to enable the network to learn the style distribution Recent style transfer models employ

GAN [Goodfellow et al. 2014] to align the distribution of generated images with specific artistic images [Chen et al. 2021b; Lin et al. 2021]. The adversarial loss can enhance the holistic style of the stylization results, while it strongly relies on the distribution of datasets. Even with the specific artistic style loss, the generation process is often not robust enough to be artifact-free.

Differently from these previous methods, we divide the images in the training set into realistic domain and artistic domain, and we use two discriminators D_R and D_A to enhance them respectively (see Figure 3). During the training process, we first randomly select an image from the realistic domain as the content image I_c and another image from the artistic domain as the style image I_s. I_c and I_s are used as the real samples of D_R and D_A, respectively. The generated image $I_{cs} = G(I_c, I_s)$ is used as the fake sample of D_A. We exchange the content and style images to generate an image $I_{sc} = G(I_s, I_c)$ as the fake sample of D_R. The adversarial loss is:

$$\begin{aligned}\mathcal{L}_{adv} =&\mathbb{E}[\log D_R(I_c)] + \mathbb{E}[\log(1 - D_R(I_{cs}))]\\ &+ \mathbb{E}[\log D_A(I_s)] + \mathbb{E}[\log(1 - D_A(I_{sc}))],\end{aligned} \quad (3)$$

To maintain the content information of the content image in the process of style transfer between the two domains, we also add a cycle consistency loss:

$$\mathcal{L}_{cyc} = \mathbb{E}[\|I_c - G(I_{cs}, I_c)\|_1] + \mathbb{E}[\|I_s - G(I_{sc}, I_s)\|_1]. \quad (4)$$

3.4 Network Training

Our full objective function for training of the generator G and discriminators D_R and D_A is formulated as:

$$\mathcal{L}(G, D_R, D_A) = \lambda_1 \mathcal{L}_{adv} + \lambda_2 \mathcal{L}_{cyc} + \lambda_3 \mathcal{L}^G_{contra}, \quad (5)$$

where λ_1, λ_2, and λ_3 are weights to balance different loss terms. We set $\lambda_1 = 1$, $\lambda_2 = 2$, and $\lambda_3 = 0.2$ in all of our experiments.

Implementation details. We collect 100,000 artistic images in different styles from WikiArt [Phillips and Mackintosh 2011] and randomly sample 20,000 images as our artistic dataset. We averagely sample 20,000 images from Places365 [Zhou et al. 2018] as realistic image dataset. We train and evaluate our framework on those artistic and realistic images. In the training phase, all images are loaded with 256×256 resolution. The number of feature map layers M is set to be 4. The dimension K of style latent code is set to 512, 1024, 2048, and 2048 for the four different layers, respectively. We use Adam [Kingma and Ba 2014] as optimizer with $\beta_1 = 0.5$, $\beta_2 = 0.999$, and a batch size of 4. The initial learning rate is set to 1×10^{-4} and linear decayed linear for total 8×10^5 iterations. The training process takes about 18 hours on one NVIDIA GeForce RTX3090. We choose the same backbone as AdaIN [Huang and Belongie 2017] in our experiments for simplicity. The results of using other backbones are shown in the supplementary materials.

4 EXPERIMENTS

We compare CAST with several state-of-the-art style transfer methods, including NST [Gatys et al. 2016], AdaIN [Huang and Belongie 2017], LST [Li et al. 2019], SANet [Park and Lee 2019], ArtFlow [An et al. 2021], MCCNet [Deng et al. 2021], AdaAttN [Liu et al. 2021b], and IEST [Chen et al. 2021a]. All the baselines are trained using publicly available implementations with default configurations. The comparison of inference speed is shown in Table 1.

4.1 Qualitative Evaluation

We first present qualitative results of our method against the selected state-of-the-art methods in Figure 5. The comparison shows the superiority of CAST in terms of visual quality. NST is likely to encounter the issue of unpleasant local minimum (e.g., the 1st, 3rd and 8th rows). AdaIN often fails to generate sharp details and introduces undesired patterns that do not exist in style images (e.g., the 1st, 3rd, 6th and 8th rows). LST tends to transfer low-level style patterns like colors but the local details of strokes are often ignored (e.g., the 2nd-5th rows). SANet often generates repetitive patterns in the stylized images (e.g., the 2nd, 5th, 6th and 8th rows). ArtFlow sometimes generates unexpected colors or patterns in relatively smooth regions in some cases (e.g., the 1st-5th and 8th rows). MCCNet can effectively preserve the input content but may fail to capture the stroke details and often generates haloing artifacts around object contours (e.g., the 2nd, 4th, 6th-8th rows). AdaAttN cannot well capture some stroke patterns (e.g., the 1st, 2nd, 5th and 6th rows) and fails to transfer important colors of the style references to the results (the 3rd and 8th rows). Although the generated visual effects of IEST are of high quality, the usage of second-order statistics as style representation causes color distortion (e.g., the 1st row in Figure 2e and the 4th row in Figure 5) and cannot capture the detailed stylized patterns (e.g., the regions of sky in the 5th and 7th rows in Figure 5). In particular, these state-of-the-art methods cannot capture the *leaving blank* characteristic of Chinese painting style in the 1st row of Figure 5 and fail to generate results with a clean background.

In comparison, CAST achieves the best stylization performance that balances characteristics of style patterns and content structures. Instead of using second-order statistics as a global style descriptor, we use an MSP module for style encoding with the help of a DE module for effective learning of style distribution. Thus, CAST can flexibly represent vivid local stroke characteristics and the overall appearance while still preserving the content structure. For instance, as shown in Figures 1j and 2c (the 2nd row), CAST successfully captures the large portion of empty regions in the style images, and it generates a stylization results which have salient objects in the center and blank space around. The whole vague appearance of the style image (drawn by Claude Monet) in Figures 1e and 2c (the 1st row) is also effectively transferred to the content images.

4.2 Quantitative Evaluation

We use the content loss [Li et al. 2017], LPIPS [Chen et al. 2021a], and deception rate [Sanakoyeu et al. 2018b] and conduct two user studies to evaluate our method quantitatively. The two user studies are online surveys that cover art/computer science students/professors and civil servants.

For content loss and LPIPS, we use a pre-trained VGG-19 and compute the average perceptual distances between the content image and the stylized image. The statistics are shown in Table 1. For deception rate, we train a VGG-19 network to classify 10 styles on WikiArt. Then, the deception rate is calculated as the percentage of stylized images that are predicted by the pre-trained network as the correct target styles. We report the deception rate for the proposed CAST and the baseline models in the 2nd column of Table 1. As observed, CAST achieves the highest accuracy and surpasses other

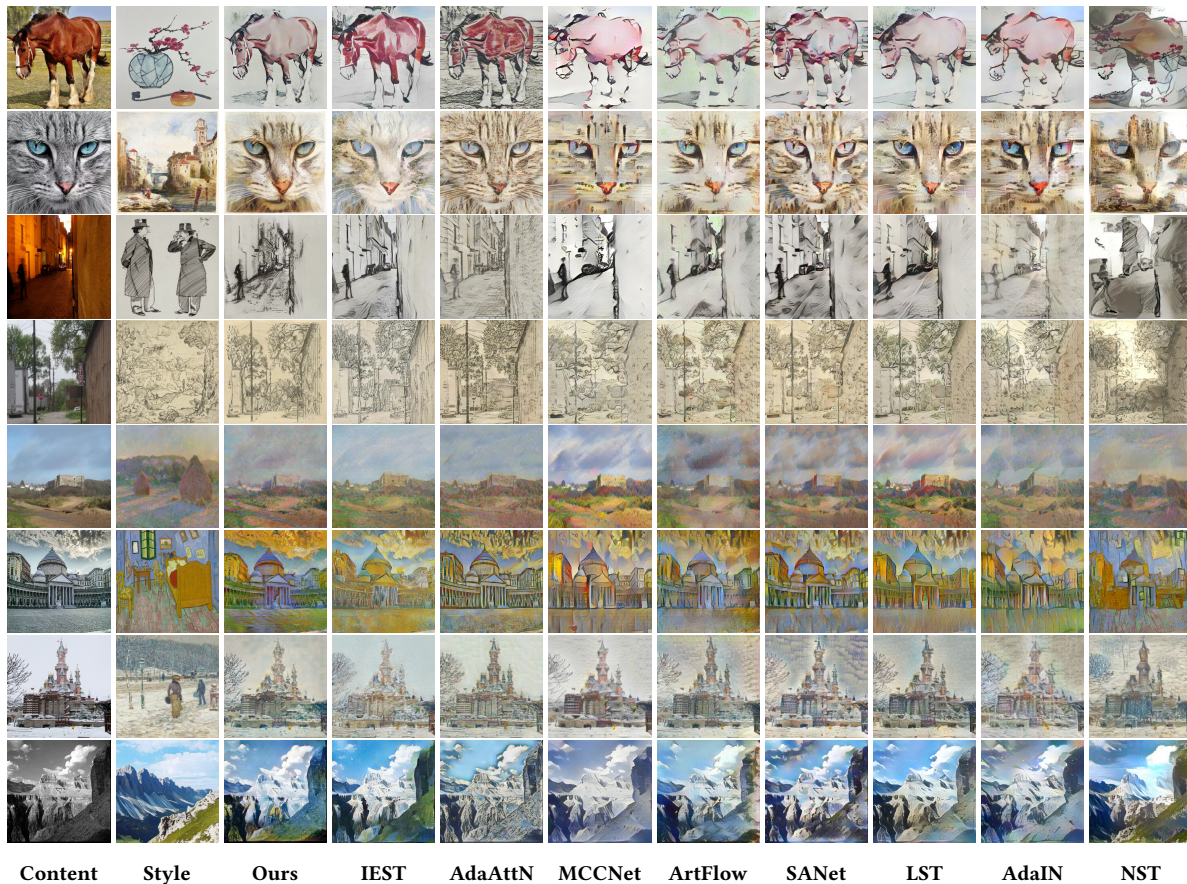| Content | Style | Ours | IEST | AdaAttN | MCCNet | ArtFlow | SANet | LST | AdaIN | NST |

Figure 5: Qualitative comparisons with several state-of-the-art style transfer methods, including IEST [Chen et al. 2021a], AdaAttN [Liu et al. 2021b], MCCNet [Deng et al. 2021], ArtFlow [An et al. 2021], SANet [Park and Lee 2019], LST [Li et al. 2019], AdaIN [Huang and Belongie 2017], and NST [Gatys et al. 2016]. Content image credits: *Horse* **Pixabay/Pexels (CC0),** *Cat* **Pixabay/Pexels (CC0). Style image credits (from the the 2nd row to the 7th row): {Richard Parkes Bonington, Philip William May, Michel Ange Corneille, Claude Monet, Vincent van Gogh, Childe Hassam}/The Art Institute of Chicago (CC0).**

methods by a large margin. As a reference, the mean accuracy of the network on real images of the artists from WikiArt is 78%.

User Study I. We compare CAST with eight state-of-the-art style transfer methods to evaluate which method generates results that are most favored by humans. For each participant, 50 content-style pairs are randomly selected and the stylized results of CAST and one of the other methods are displayed in a random order. Then, we ask the participant to choose the image that learns the most characteristics from the style image. Participants were told that the consistency of content and style was the primary metrics. The style is subjective and the effectiveness of training also depends on their understanding ability. Finally, we collect 3,400 votes from 68 participants. We report the percentage of votes for each method in the 3rd column of Table 1. CAST obtains significantly higher preferences in categories of Sketch, Chinese painting, and Impressionism.

User Study II. We design a novel user study to evaluate the stylized images quantitatively, which is called the Stylized Authenticity Detection (SAD). For each question, we show participants ten artworks of similar styles, including two to four stylized fake painting and ask them to select the synthetic ones. Within each single question, the stylized paintings are generated by the same method. Each participant finished 25 questions. Finally, we collect 2125 groups of results from 85 participants and use the average precision and recall as the measurement for how likely the results will be recognized as synthetics. Table 1 shows the statistics. The paintings generated by CAST have the lowest chance to be decided by people as fake paintings. We also notice that the precision and recall of CAST is less than 50%, which means that users could not tell the real ones from the fakes and tend to select more real paintings as synthetics when doing the testing.

4.3 Ablation Study

Contrastive style loss. We remove the contrastive style loss from Eq. (5) to train the model. As shown in Figures. 6e and 6i, the model without our contrastive style loss cannot capture the color and the stroke characteristics of the style image compared with the full CAST model. The sharp pencil lines of the style image in the 1st

(a) Content   **(b) Style**   **(c) AdaIN**   **(d) w/o DE**   **(e) w/o $\mathcal{L}^G_{cont}$**   **(f) mix-DE**   **(g) one-DE**   **(h) $1/2$ cycle**   **(i) full CAST**
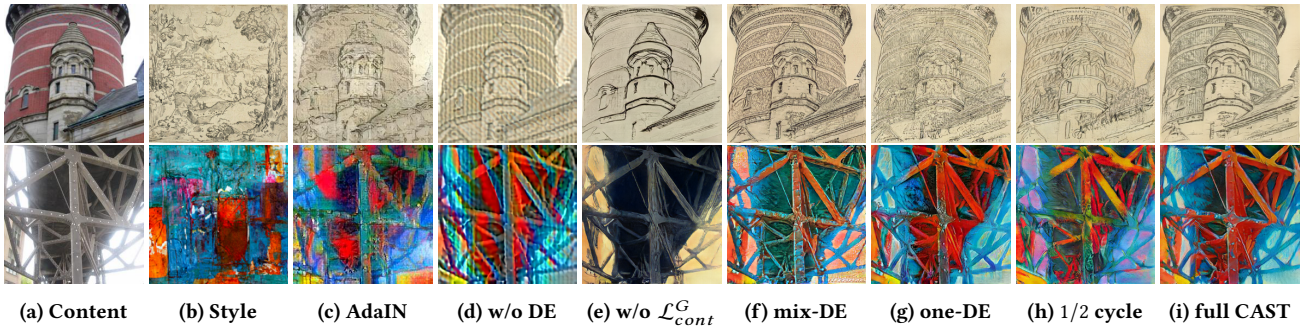
Figure 6: Ablation study results. From left to right: (a) content image; (b) style image; (c) AdaIN with perceptual loss; (d) CAST without DE; (e) AdaIN with DE, and cycle consistency loss; (f) CAST using mixed DE; (g) CAST using one DE without the realistic domain; (h) CAST trained with asymmetric cycle consistent loss by only reconstructing the realistic images; and (i) full CAST model. Style image credits: Michel Ange Corneille/The Art Institute of Chicago (CC0), Steve Johnson/Pexels (CC0).

Table 1: Statistics of inference speed and quantitative comparison with state-of-the-art methods. The results of user study I represent the average percentage of cases in which the result of the corresponding method is preferred compared with ours. The results of user study II show the accuracy and recall of being selected as fake paintings by the participants. The best results are in bold while the second best results are marked with <u>underline</u>.

Method		CAST	IEContraAST	AdaAttN	MCCNet	Artflow	SANet	LST	AdaIN	NST
Inference time (ms/image)↓		<u>11</u>	184	130	29	168	14	**7**	<u>11</u>	16863
Content loss↓		<u>0.148</u>	0.155	0.162	**0.117**	0.172	0.150	0.155	0.176	0.188
LPIPS↓		<u>0.245</u>	0.256	0.256	**0.234**	0.264	0.265	0.248	0.266	0.291
Deception Rate↑		**62.00%**	<u>56.42%</u>	50.70%	46.37%	43.79%	51.87%	48.29%	51.00%	37.70%
User Study I		-	30.25%	41.9%	38.3%	46.1%	44.7%	20.0%	29.8%	25.1%
User StudyII	Precision↓	**43.69%**	60.26%	58.64%	70.89%	<u>56.81%</u>	65.80%	65.49%	70.66%	60.18%
(SAD)	Recall↓	**41.19%**	58.67%	<u>58.16%</u>	72.76%	61.55%	62.55%	64.91%	75.93	64.76%

row become thick black lines. The textural oil painting strokes of the style image in the 2nd row become smooth blocks, and the vivid colors almost disappear. With the contrastive style loss, our full model can faithfully transfer the brushstrokes, textures, and colors from the input style image.

Domain enhancement. Our full CAST uses DE for realistic and artistic images separately. We train a simplified CAST model using one discriminator that mixes realistic and artistic images together (mix-DE). As shown in Figure 6f, the results generated by mix-DE model are acceptable, but the stroke details in the generated images are weaker than the ones by the full CAST model. This fact is due to the existence of a significant gap between the artistic and realistic image domains. We further abandon all images from realistic domain for ablation (one-DE). As shown in Figure 6g, the results generated by one-DE model lack details.

Cycle consistency loss. To better evaluate the improvement of the contrastive style loss on the style transfer task, we exclude the latent promotion of cycle consistency loss from network training. The reason is that the reconstruction process of artistic image may imply style information. We train CAST with an asymmetric cycle consistent loss, which only reconstructs the realistic images. The decoder of the style transfer network is unaffected by the reconstruction of the artistic image. As shown in Figure 6h, removing realistic image reconstruction will lead to slightly degraded stylization results.

5 CONCLUSION AND FUTURE WORK

In this work, we present a novel framework, namely CAST, for the task of arbitrary image style transfer. Instead of relying on second-order metrics such as Gram matrix or mean/variance of deep features, we use image features directly by introducing an MSP module for style encoding. We develop a contrastive loss function to leverage the available multi-style information in the existing collection of artwork and help train the MSP module and our generative style transfer network. We further propose a DE scheme to effectively model the distribution of realistic and artistic image domains. Extensive experimental results demonstrate that our proposed CAST method achieves superior arbitrary style transfer results compared with state-of-the-art approaches. In the future, we plan to improve the contrastive style learning process by considering artist and category information.

REFERENCES

Jie An, Siyu Huang, Yibing Song, Dejing Dou, Wei Liu, and Jiebo Luo. 2021. ArtFlow: Unbiased Image Style Transfer via Reversible Neural Flows. In *IEEE/CVF Conferences on Computer Vision and Pattern Recognition (CVPR)*. 862–871.

Kyungjune Baek, Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Hyunjung Shim. 2021. Rethinking the truly unsupervised image-to-image translation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*. 14154–14163.

Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, and Gang Hua. 2017. StyleBank: An explicit representation for neural image style transfer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1897–1906.

Haibo Chen, Lei Zhao, Zhizhong Wang, Zhang Hui Ming, Zhiwen Zuo, Ailin Li, Wei Xing, and Dongming Lu. 2021a. Artistic Style Transfer with Internal-external Learning and Contrastive Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Haibo Chen, Lei Zhao, Zhizhong Wang, Huiming Zhang, Zhiwen Zuo, Ailin Li, Wei Xing, and Dongming Lu. 2021b. DualAST: Dual Style-Learning Networks for Artistic Style Transfer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 872–881.

Yingying Deng, Fan Tang, Weiming Dong, Haibin Huang, Chongyang Ma, and Changsheng Xu. 2021. Arbitrary Video Style Transfer via Multi-Channel Correlation. In *AAAI Conference on Artificial Intelligence (AAAI)*. 1210–1217.

Yingying Deng, Fan Tang, Weiming Dong, Chongyang Ma, Xingjia Pan, Lei Wang, and Changsheng Xu. 2022. StyTr^2: Image Style Transfer with Transformers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yingying Deng, Fan Tang, Weiming Dong, Wen Sun, Feiyue Huang, and Changsheng Xu. 2020. Arbitrary style transfer via multi-adaptation network. In *ACM International Conference on Multimedia*. 2719–2727.

Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. 2017. A learned representation for artistic style. In *International Conference on Learning Representations*.

Jakub Fišer, Ondřej Jamriška, Michal Lukáč, Eli Shechtman, Paul Asente, Jingwan Lu, and Daniel Sýkora. 2016. StyLit: Illumination-Guided Example-Based Stylization of 3D Renderings. *ACM Transactions on Graphics* 35, 4, Article 92 (2016), 11 pages.

Wei Gao, Yijun Li, Yihang Yin, and Ming-Hsuan Yang. 2020. Fast video multi-style transfer. In *IEEE/CVF Winter Conference on Applications of Computer Vision*. 3222–3230.

Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2016. Image style transfer using convolutional neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2414–2423.

Leon A. Gatys, Alexander S. Ecker, Matthias Bethge, Aaron Hertzmann, and Eli Shechtman. 2017. Controlling Perceptual Factors in Neural Style Transfer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3730–3738.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems (NIPS)*.

Junlin Han, Mehrdad Shoeiby, Lars Petersson, and Mohammad Ali Armin. 2021. Dual Contrastive Learning for Unsupervised Image-to-Image Translation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 746–755.

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 9729–9738.

Xun Huang and Serge Belongie. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision (ICCV)*. 1501–1510.

Jongheon Jeong and Jinwoo Shin. 2021. Training GANs with Stronger Augmentations via Contrastive Discriminator. In *International Conference on Learning Representations*.

Yongcheng Jing, Xiao Liu, Yukang Ding, Xinchao Wang, Errui Ding, Mingli Song, and Shilei Wen. 2020a. Dynamic instance normalization for arbitrary style transfer. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 4369–4376.

Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, Yizhou Yu, and Mingli Song. 2020b. Neural Style Transfer: A Review. *IEEE Transactions on Visualization and Computer Graphics* 26, 11 (2020), 3365–3385.

Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*. Springer, 694–711.

Minguk Kang and Jaesik Park. 2020. ContraGAN: Contrastive Learning for Conditional Image Generation. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. 2019. Style Transfer by Relaxed Optimal Transport and Self-Similarity. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10043–10052.

Dmytro Kotovenko, Artsiom Sanakoyeu, Sabine Lang, and Bjorn Ommer. 2019a. Content and style disentanglement for artistic style transfer. In *IEEE/CVF International Conference on Computer Vision (ICCV)*. 4422–4431.

Dmytro Kotovenko, Artsiom Sanakoyeu, Pingchuan Ma, Sabine Lang, and Bjorn Ommer. 2019b. A content transformation block for image style transfer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10032–10041.

Xueting Li, Sifei Liu, Jan Kautz, and Ming-Hsuan Yang. 2019. Learning Linear Transformations for Fast Image and Video Style Transfer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 3804–3812.

Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. 2017. Universal style transfer via feature transforms. In *Advances Neural Information Processing Systems (NeurIPS)*. 386–396.

Jing Liao, Yuan Yao, Lu Yuan, Gang Hua, and Sing Bing Kang. 2017. Visual Attribute Transfer through Deep Image Analogy. *ACM Transactions on Graphics* 36, 4, Article 120 (2017), 15 pages.

Minxuan Lin, Fan Tang, Weiming Dong, Xiao Li, Changsheng Xu, and Chongyang Ma. 2021. Distribution Aligned Multimodal and Multi-Domain Image Stylization. *ACM Transactions on Multimedia Computing, Communications, and Applications* 17, 3, Article 96 (2021), 17 pages.

Rui Liu, Yixiao Ge, Ching Lam Choi, Xiaogang Wang, and Hongsheng Li. 2021a. DivCo: Diverse Conditional Image Synthesis via Contrastive Generative Adversarial Network. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 16372–16381.

Songhua Liu, Tianwei Lin, Dongliang He, Fu Li, Meiling Wang, Xin Li, Zhengxing Sun, Qian Li, and Errui Ding. 2021b. AdaAttN: Revisit attention mechanism in arbitrary neural style transfer. In *IEEE/CVF International Conference on Computer Vision (ICCV)*. 6649–6658.

Xialei Liu, Joost van de Weijer, and Andrew D. Bagdanov. 2019. Exploiting Unlabeled Data in CNNs by Self-Supervised Learning to Rank. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 8 (2019), 1862–1878.

Dae Young Park and Kwang Hee Lee. 2019. Arbitrary Style Transfer With Style-Attentional Networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 5880–5888.

Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. 2020. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision (ECCV)*. Springer, 319–345.

Fred Phillips and Brandy Mackintosh. 2011. Wiki Art Gallery, Inc.: A case for critical thinking. *Issues in Accounting Education* 26, 3 (2011), 593–608.

Gilles Puy and Patrick Pérez. 2019. A flexible convolutional solver for fast style transfers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 8963–8972.

Artsiom Sanakoyeu, Dmytro Kotovenko, Sabine Lang, and Bjorn Ommer. 2018a. A style-aware content loss for real-time HD style transfer. In *European Conference on Computer Vision (ECCV)*. 698–714.

Artsiom Sanakoyeu, Dmytro Kotovenko, Sabine Lang, and Björn Ommer. 2018b. A Style-Aware Content Loss for Real-Time HD Style Transfer. In *European Conference Computer Vision (ECCV)*. Springer International Publishing, Cham, 715–731.

Rodrigo Santa Cruz, Basura Fernando, Anoop Cherian, and Stephen Gould. 2019. Visual Permutation Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 12 (2019), 3100–3114.

Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*.

Jan Svoboda, Asha Anoosheh, Christian Osendorfer, and Jonathan Masci. 2020. Two-stage peer-regularized feature recombination for arbitrary image style transfer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13816–13825.

Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S Lempitsky. 2016. Texture networks: Feed-forward synthesis of textures and stylized images. In *International Conference on Machine Learning (ICML)*. 1349–1357.

Aaron Van den Oord, Yazhe Li, Oriol Vinyals, et al. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* 2, 3 (2018), 4.

Bin Wang, Wenping Wang, Huaiping Yang, and Jiaguang Sun. 2004. Efficient example-based painting and synthesis of 2D directional texture. *IEEE Transactions on Visualization and Computer Graphics* 10, 3 (2004), 266–277.

Haiyan Wu, Yanyun Qu, Shaohui Lin, Jian Zhou, Ruizhi Qiao, Zhizhong Zhang, Yuan Xie, and Lizhuang Ma. 2021b. Contrastive Learning for Compact Single Image Dehazing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10546–10555.

Xiaolei Wu, Zhihao Hu, Lu Sheng, and Dong Xu. 2021a. StyleFormer: Real-Time Arbitrary Style Transfer via Parametric Style Composition. In *IEEE/CVF International Conference on Computer Vision (ICCV)*. 14618–14627.

Wenju Xu, Chengjiang Long, Ruisheng Wang, and Guanghui Wang. 2021. DRB-GAN: A Dynamic ResBlock Generative Adversarial Network for Artistic Style Transfer. In *IEEE/CVF International Conference on Computer Vision (ICCV)*. 6383–6392.

Hang Zhang and Kristin Dana. 2018. Multi-style generative network for real-time transfer. In *European Conference on Computer Vision Workshops*. 349–365.

Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2018. Places: A 10 Million Image Database for Scene Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40, 6 (2018), 1452–1464.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*. 2223–2232.