

Style-Aware Normalized Loss for Improving Arbitrary Style Transfer

Jiaxin Cheng^{*†}, Ayush Jaiswal[‡], Yue Wu[‡], Pradeep Natarajan[‡], Prem Natarajan[‡]
†USC Information Sciences Institute, ‡Amazon Alexa Natural Understanding

chengjia@isi.edu, {ayujaisw, wuayue, natarap, premknat}@amazon.com



Figure 1: Popular Arbitrary Style Transfer methods suffer from under-stylization and over-stylization due to imbalanced style transferability during training. Our new balanced loss mitigates these under- and over-stylization issues.

Abstract

Neural Style Transfer (NST) has quickly evolved from single-style to infinite-style models, also known as Arbitrary Style Transfer (AST). Although appealing results have been widely reported in literature, our empirical studies on four well-known AST approaches (GoogleMagenta [14], AdaIN [19], LinearTransfer [29], and SANet [37]) show that more than 50% of the time, AST stylized images are not acceptable to human users, typically due to under- or over-stylization. We systematically study the cause of this imbalanced style transferability (IST) and propose a simple yet effective solution to mitigate this issue. Our studies show that the IST issue is related to the conventional AST style loss, and reveal that the root cause is the equal weightage of training samples irrespective of the properties of their corresponding style images, which biases the model towards certain styles. Through investigation of the theoretical bounds of the AST style loss, we propose a new loss that largely overcomes IST. Theoretical analysis and experimental results validate the effectiveness of our loss, with over 80% relative improvement in style deception rate and 98% relatively higher preference in human evaluation.

^{*}This work was completed during his internship at Amazon.

1. Introduction

Neural style transfer (NST) refers to the generation of a pastiche image P from two images C and S via a neural network, where P shares the content with C but is in the style of S . While the original NST approach of Gatys [13] optimizes the transfer model for each pair of C and S , the field has rapidly evolved in recent years to develop models that support arbitrary styles out-of-the-box. NST models can, hence, be classified based on their stylization capacity into models trained for (1) a single combination of C and S [13, 23, 28, 32, 39], (2) one S [21, 27, 47, 48], (3) multiple fixed S [2, 9, 24, 30, 42, 55], and (4) infinite (arbitrary) S [4, 14, 15, 17, 19, 20, 25, 29, 31, 37, 43, 44]. Intuitively, the category (4) of arbitrary style transfer (AST) is the most advantageous as it is agnostic to S , allowing trained models to be adopted for diverse novel styles without re-training.

Although superior in concept, current AST models are plagued by the issue of imbalanced style transferability (IST), where the stylization intensity of model outputs varies largely across styles S . More importantly, besides the nice results shown in previous works [14, 19, 29, 37], a large number of stylized images suffer under-stylization (e.g., only the dominant color is transferred) or over-stylization (i.e., content is barely visible) for various S , making them

visually undesirable (see samples in Figure 1). This is validated by our user-study described later in Section 3.2, with more than 50% of stylized images found to be unacceptable, irrespective of the used AST model. Hence, we are still far from the AST goal — successfully transferring style from an *arbitrary* image to another. This urges us to systematically study the underlying reasons for IST and find potential solutions to further boost AST performance in order to generate better stylized images for diverse styles.

In this paper, we make the following contributions. Firstly, we systematically study the IST problem in AST and discover that the AST style loss is problematic as it fails to reflect human evaluation scores. Secondly, we investigate the AST style loss function, and locate the core reason for IST to be the way sample-wise style loss is aggregated into a batch loss. Thirdly, we derive the theoretical expectation of a sample-wise style loss as well as its bounds, and use it to propose a new style loss that enables more balanced training across styles. Finally, we conduct extensive AST benchmarking experiments as well as human evaluation to validate the effectiveness of the proposed solution. Results show that IST issue is indeed greatly mitigated for all tested AST approaches by incorporating the proposed style loss.

The rest of the paper is organized as follows. Section 2 briefly reviews related AST works. Section 3 discusses two AST style loss related studies and shows that IST is related to the loss. Section 4 identifies style-agnostic sample weighting in training loss aggregation as the real culprit, derive our new style-aware loss, and validate its effectiveness by repeating the aforementioned two studies. Section 5 provides further results of application of the proposed loss to four well-known AST approaches and shows that the IST issue is largely overcome for all the approaches. Finally, Section 6 provides concluding remarks.

2. Related Work

Arbitrary style transfer methods can be classified as either non-parametric [10, 11, 15, 26, 31, 50, 51, 54] or parametric [14, 19, 20, 25, 29, 37, 43, 44, 46, 49]. Non-parametric methods find similar patches between content and style images, and transfer style based on matched patches. Early methods popularly performed texture synthesis [10, 11, 26, 51]. However, Neural Style Transfer (NST) methods have become mainstream since their inception in [13]. Improvements in the NST framework include multi-level whitening and coloring [31] on VGG [45] features, and feature reshuffling based on patch-based feature similarity [15].

Parametric AST [14, 19, 20, 25, 29, 37, 43, 44, 46, 49, 52] involves optimizing a target function that reflects visual similarity of (1) content between content and stylized images, and (2) style between style and stylized images. This class of AST methods [14, 19, 20, 29, 37, 43] typically uses a Gram matrix-based VGG perceptual loss [13, 21] with a

Table 1: The four studied AST methods with model links.

AST Method	Net Architecture w/ Unique Feature
GoogleMagenta [14]	ConvNet w/ meta-learned instance norm
AdaIN [19]	Enc.&Dec. w/ adaptive instance norm
LinearTransfer [29]	Enc.&Dec. w/ linear transform matrix
SANet [37]	Enc.&Dec. w/ style attention

few modifications to the training procedure. AdaIN [19] applies an adaptive instance normalization layer on image features. Meta-learning is incorporated for learning style representation in GoogleMagenta [14] and [43]. LinearTransfer [29] learns a linear transform function from content and style features for stylization. Attention mechanism has also been integrated with content and style feature fusion in SANet [37] and other works [46, 52]. Light-weight efficient style transfer has been explored through instance normalization and dynamic convolution [20]. Gram matrix-based loss is also widely used in arbitrary video style transfer [1, 3, 18, 29, 40, 41]. Methods without Gram matrix-based losses use adversarial [25] or reconstruction [44] objectives. This paper studies parametric AST methods involving Gram-matrix based losses as listed in Table 1.

3. Analysis of the AST Style Loss

3.1. AST Training Loss

A number of loss functions have been proposed recently for AST training, *e.g.*, discriminator-based adversarial losses [25] and reconstruction-based loss terms [44]. However, the original [13] NST loss \mathcal{L}_{NST} is still the one that is the most popularly employed [14, 19, 20, 29, 37, 43]. As described in Equation (1), it is composed of two terms: $\mathcal{L}_{\text{NST}_c}(C, P)$ for learning content from C and $\mathcal{L}_{\text{NST}_s}(S, P)$ for deriving style from S , with a trade-off factor β .

$$\mathcal{L}_{\text{NST}} = \mathcal{L}_{\text{NST}_c}(C, P) + \beta \mathcal{L}_{\text{NST}_s}(S, P) \quad (1)$$

One typically needs an ImageNet [8] pretrained VGG [45] network \mathcal{F} for extracting features from C , S , and P . Next, the content loss is calculated by comparing the features of P and C , and the style term is calculated by comparing the Gram matrices \mathcal{G} of the features of P and S , as \mathcal{G} is known [13, 32] to be effective in deriving style information. In practice, the style and content terms are calculated for features from several layers and aggregated using a weighted sum (weights w^l are typically set as ones) across layers. The following equations summarize the loss calculation, where MSE is the mean squared error.

$$\mathcal{L}_{\text{NST}_c}^l(C, P) = \text{MSE}(\mathcal{F}^l(C), \mathcal{F}^l(P)) \quad (2)$$

$$\mathcal{L}_{\text{NST}_s}^l(S, P) = \text{MSE}(\mathcal{G} \circ \mathcal{F}^l(S), \mathcal{G} \circ \mathcal{F}^l(P)) \quad (3)$$

$$\mathcal{L}_{\text{NST}_c}(C, P) = \sum_{l \in \mathbb{L}_{\text{NST}_c}} w_c^l \cdot \mathcal{L}_{\text{NST}_c}^l(C, P) \quad (4)$$

$$\mathcal{L}_{\text{NST}_s}(S, P) = \sum_{l \in \mathbb{L}_{\text{NST}_s}} w_s^l \cdot \mathcal{L}_{\text{NST}_s}^l(S, P) \quad (5)$$

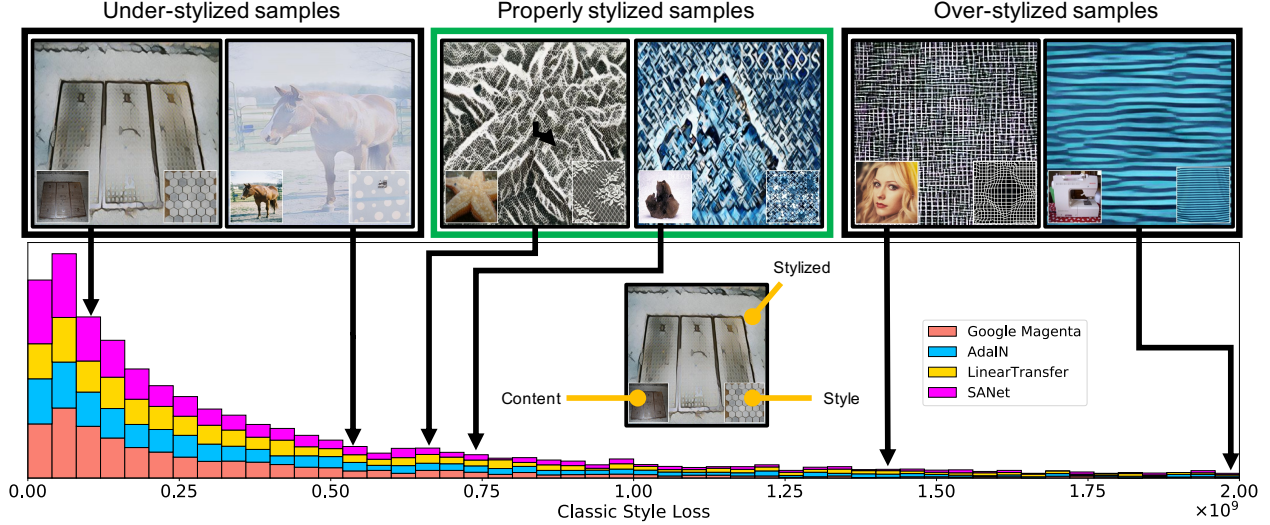


Figure 2: Distribution of classic Gram matrix-based style losses for four AST methods [14, 19, 29, 37]. Smaller loss does not guarantee better style transfer (left two images) while high quality transferred images can have larger style losses (middle two images), with over-stylized images counter-intuitively attaining the highest losses. Zoom in for a better view.

3.2. Analysis

The IST issue could be intuitively attributed to the “naturally higher” difficulty of transferring certain styles compared to others. In order to study IST systematically, we calculate content and style losses for 20,000 randomly sampled ImageNet [8] images stylized with images from the Describable Textures Dataset (DTD) [7] and pretrained models listed in Table 1. Ideally, under- and over-stylized P are those with low content and low style losses, respectively. On analyzing the samples, however, we find that the former relationship is valid but the latter is not — over-stylized P typically attain (sometimes significantly) higher style losses than understylized samples. In the following studies, we examine the distribution of the style loss and its correlation with visual perception of stylization quality.

Study I: AST Style Loss Distribution. We compute the empirical distribution of style losses for the models listed in Table 1 and inspect stylized samples belonging to different sections of the distribution – low, moderate, and high style losses. We use a VGG-16 model pretrained on ImageNet as the feature extractor, and calculate the style loss (see Equation (5)) using layers \mathcal{F}^l in the conventional style layer set [21] $\mathbb{L}_{AST_s} = \{\mathcal{F}_{b1}^{r2}, \mathcal{F}_{b2}^{r2}, \mathcal{F}_{b3}^{r3}, \mathcal{F}_{b4}^{r4}\}$, where \mathcal{F}_{bi}^{rj} denotes the j -th ReLU layer in i -th convolutional block of VGG-16. The AST style loss is thus restated as below.

$$\mathcal{L}_{AST_s}^l(S, P) = \text{MSE}(\mathcal{G} \circ \mathcal{F}^l(S), \mathcal{G} \circ \mathcal{F}^l(P)) \quad (6)$$

$$\mathcal{L}_{AST_s}(S, P) = \sum_{l \in \mathbb{L}_{AST_s}} w_s^l \cdot \mathcal{L}_{AST_s}^l(S, P) \quad (7)$$

Figure 2 summarizes our findings. Despite large differences among the tested methods, (1) their \mathcal{L}_{AST_s} distributions are similar, and (2) \mathcal{L}_{AST_s} does not reflect stylization quality:

under-stylized samples attain lower loss values than over-stylized ones. Similar conclusions can be drawn for VGG-19-based \mathcal{L}_{AST_s} , as shown in the supplementary material.

Study II: Classic AST Style Loss versus Human Score.

Study I has revealed a counter-intuitive lack of relationship between style transfer quality and conventional AST style loss. In this study, we further investigate this issue by conducting a human study to assess the correlation between \mathcal{L}_{AST_s} and human perception. Specifically, we requested five volunteers to manually annotate AST samples by partitioning the samples produced in Study I into five random disjoint subsets. Each sample was presented as a tuple of (S, P) and had to be annotated as “Good” (-1), “OK” (0), or “Bad” (1), in decreasing order of stylization quality. The annotators were not given additional instructions and were told to classify the samples based on their own perception.

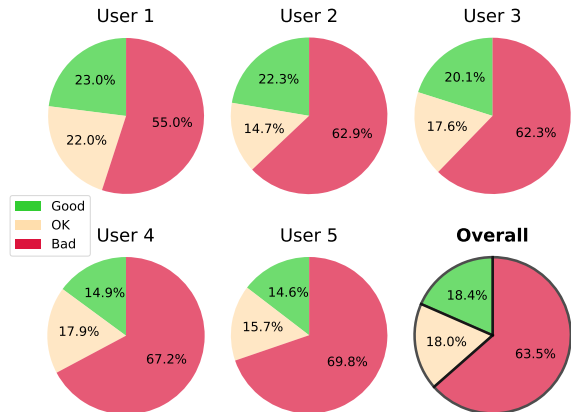


Figure 3: Statistics of human perception of stylization quality as assessed in Study II.

Figure 3 shows the statistics of the collected annotations. We compute Pearson correlation between the human scores and corresponding style losses. Table 2 summarizes the results, showing that the conventional style loss not only fails to reflect human perception but is also negatively correlated. Hence, $\mathcal{L}_{\text{AST}_s}$ (as defined in Equation (7)) is inappropriate for the AST task. Furthermore, the negative correlation indicates that this style loss penalizes over-stylized samples more than under-stylized ones — contrary to what one would expect for a good AST style loss.

Table 2: Pearson correlation between the classic AST style loss ($\mathcal{L}_{\text{AST}_s}$) and human score (h). \mathcal{F}_{bi}^{rj} indicates the the j -th ReLU layer in i -th convolutional block of VGG-16.

Volunteer	Correlation Score $\rho(h, \mathcal{L}_s)$				
	$\mathcal{L}_s^{\mathcal{F}_{b1}^{r2}}$	$\mathcal{L}_s^{\mathcal{F}_{b2}^{r2}}$	$\mathcal{L}_s^{\mathcal{F}_{b3}^{r3}}$	$\mathcal{L}_s^{\mathcal{F}_{b4}^{r4}}$	$\mathcal{L}_{\text{AST}_s}$
1	-0.166	-0.153	-0.186	-0.114	-0.167
2	-0.163	-0.141	-0.139	-0.074	-0.150
3	-0.184	-0.161	-0.147	-0.177	-0.163
4	-0.194	-0.147	-0.127	0.006	-0.138
5	-0.214	-0.224	-0.192	-0.070	-0.219
Average	-0.184	-0.165	-0.158	-0.088	-0.167

4. A New Balanced AST Style Loss

The conventional AST style loss defined in Equation (7) reuses the classic NST style loss, which has been shown to work in practice in several previous works [14, 19, 29, 37, 43]. However, results of studies I and II reveal that this AST style loss is problematic: trained models work partially but suffer from Imbalanced Style Transferability (IST), *i.e.*, under- or over-stylization for various styles with loss values that do not reflect stylization quality. In this section, we first identify the core issue in the AST style loss by viewing AST from a multi-task learning point-of-view. We then propose a simple yet effective solution to mitigate the problem.

4.1. Identifying the Real Problem with AST Loss

The conventional AST style loss (also the classic NST style loss) used in studies I and II is a sample-wise loss. However, in order to ascertain the cause of the aforementioned issue, it is important to inspect how it is used in training — it needs to be aggregated into a batch-wise loss as

$$\mathcal{L}_{\text{AST}_s}^{\text{Batch}} = \sum_{k \in \{1, \dots, B\}} \frac{1}{B} \cdot \mathcal{L}_{\text{AST}_s}(S_k, P_k) \quad (8)$$

where B is the batch size. While it is typical to average sample-wise losses into batch-losses in this fashion, this protocol is not suitable for AST. This is because the AST learning setup resembles multi-task learning, where each

batch has B tasks — one for each input style. The overall multi-task loss can be written as

$$\mathcal{L}_{\text{AST}_s}^{\text{Multitask}} = \sum_{k \in \{1, \dots, B\}} \lambda_k \cdot \mathcal{L}_{\text{AST}_s}(S_k, P_k) \quad (9)$$

where λ_k is typically a task-specific contribution factor. Comparing Equations (8) and (9), it is clear that the AST style loss is a special case of the multi-task loss when $\lambda_k = 1/B$ for each k . However, this equal-task-weight setting is known to be problematic in multi-task learning unless all task losses are within similar dynamic ranges [6, 22].

In case of AST, style losses for different style images can differ by more than 1,000 times for both randomly initialized and fully trained AST models. Consequently, styles with small or large dynamic loss ranges are under- or over-stylized, respectively. Although $\lambda_k = 1/B$ works for some style images, generating nice stylization results for them, this setting is unsuitable for the general AST problem and is the root cause of the discrepancy between stylization quality and loss values. Therefore, we should neither simply aggregate the sample-wise losses to form a batch-wise loss nor directly compare losses from different styles.

4.2. A New Balanced AST Style Loss

The multi-task view discussed in the previous section implies that the IST problem could be resolved by assigning each style transfer task in a batch the “right” task weight. Hence, we seek to formulate a balanced AST style loss as

$$\hat{\mathcal{L}}_{\text{AST}_s}^l(S, P) = \frac{\mathcal{L}_{\text{AST}_s}^l(S, P)}{\mathcal{V}^l(S, P)} \quad (10)$$

where $\mathcal{V}^l(S, P)$ is the appropriate task-dependent normalization term that needs to be determined (where $\lambda_k = 1/\mathcal{V}^l(S, P)$). An intuitive approach to achieve this is to adopt automatic task loss-weight tuning methods from the multi-task literature [6, 16, 22, 34]. However, these methods require estimation of statistics (*e.g.*, gradient norms) for all tasks in multiple iterations (if not continuously), which is infeasible for AST as the tasks change across batches and the number of (S, P) combinations is potentially infinite. Therefore, weight tuning approaches are not suitable for AST. Furthermore, the choice of $\mathcal{V}^l(S, P)$ is limited to something that could be computed without historical data.

We start with deriving the theoretical upper- and lower-bounds for the classic AST layerwise style loss (Equation (6)) as shown in Equations (11) and (12), respectively:

$$\sup\{\mathcal{L}_{\text{AST}_s}^l(S, P)\} = \frac{\|\mathcal{G} \circ \mathcal{F}^l(S)\|^2 + \|\mathcal{G} \circ \mathcal{F}^l(P)\|^2}{N^l} \quad (11)$$

$$\inf\{\mathcal{L}_{\text{AST}_s}^l(S, P)\} = \frac{(\|\mathcal{G} \circ \mathcal{F}^l(S)\| - \|\mathcal{G} \circ \mathcal{F}^l(P)\|)^2}{N^l} \quad (12)$$

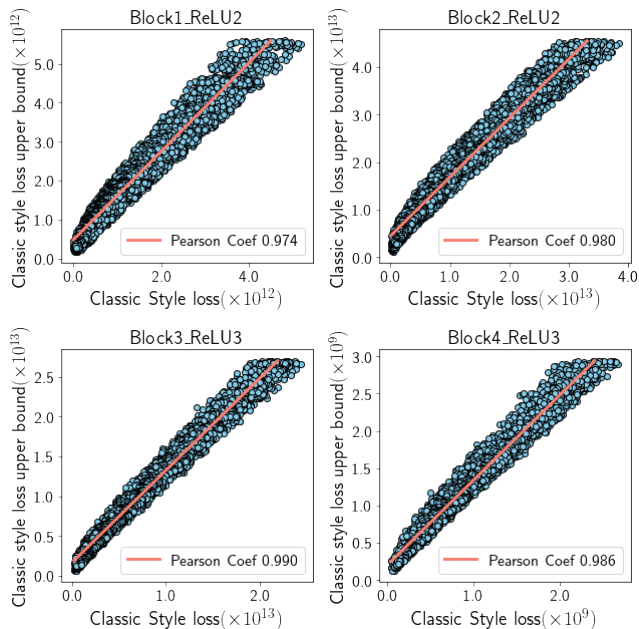


Figure 4: Relationship between the classic AST style loss $\mathcal{L}_{AST_s}^l(S, P)$ and $\sup\{\mathcal{L}_{AST_s}^l(S, P)\}$. The subplots correspond to the four VGG-16 style layers used in analysis.

where N^l is a constant that is equal to the product of spatial dimensions of the feature tensor at layer l . The detailed derivations can be found in the supplementary material.

In order to mitigate Imbalanced Style Transferability, we propose a new style-balanced loss $\hat{\mathcal{L}}_{AST_s}^l$ by normalizing the style loss of each AST task with its supremum as:

$$\hat{\mathcal{L}}_{AST_s}^l(S, P) = \frac{\mathcal{L}_{AST_s}^l(S, P)}{\sup\{\mathcal{L}_{AST_s}^l(S, P)\}} \quad (13)$$

4.3. Analysis and Validation of Effectiveness

We conduct three studies to analyze and validate the correctness and effectiveness of our new loss (Equation (13)).

Study III: Relationship between $\mathcal{L}_{AST_s}^l(S, P)$ and $\sup\{\mathcal{L}_{AST_s}^l(S, P)\}$. It is important to establish this relationship to ensure that $\sup\{\mathcal{L}_{AST_s}^l(S, P)\}$ is a suitable normalization term. Specifically, the relationship has to be close to linear to balance all the training tasks by ensuring that all the training tasks have the same upper bound of 1.

We randomly sample 200,000 pairs of images from the Painter by Numbers (PBN) dataset [36]. For each pair, we compute the classic layerwise style loss $\mathcal{L}_{AST_s}^l(S, P)$ using the VGG-16 style layer set [21] (Equation (7)), and its upper-bound (Equation (11)). Figure 4 provides a scatter plot of the two terms, where each dot is a sample and the red line is the linear fit of all samples, showing that $\mathcal{L}_{AST_s}^l(S, P)$ and $\sup\{\mathcal{L}_{AST_s}^l(S, P)\}$ are strongly correlated.

Study IV: Distribution of the New Balanced AST Style

Table 3: Pearson correlation between the new AST style loss ($\hat{\mathcal{L}}_{AST_s}$) and human scores (h). \mathcal{F}_{bi}^{rj} indicates the the j -th ReLU layer in i -th convolutional block of VGG-16.

Volunteer	Correlation Score $\rho(h, \hat{\mathcal{L}}_s)$				
	$\hat{\mathcal{L}}_s^{\mathcal{F}_{b1}^{r2}}$	$\hat{\mathcal{L}}_s^{\mathcal{F}_{b2}^{r2}}$	$\hat{\mathcal{L}}_s^{\mathcal{F}_{b3}^{r3}}$	$\hat{\mathcal{L}}_s^{\mathcal{F}_{b4}^{r3}}$	$\hat{\mathcal{L}}_{AST_s}$
1	-0.086	0.095	0.188	0.121	0.152
2	0.002	0.001	0.098	0.122	0.124
3	-0.087	0.106	0.217	0.217	0.228
4	-0.067	-0.089	0.053	0.210	0.149
5	-0.141	-0.118	0.112	0.208	0.171
Average	-0.076	-0.001	0.134	0.176	0.165

Loss. As previously noted in Study I, under-stylized samples typically attain lower loss values than over-stylized ones in the classic AST style loss. Here we verify whether our new AST style loss fixes this issue. We reuse the data from Study I for the four tested AST approaches (see Table 1), and compute the corresponding new style loss distributions as shown in Figure 5. Results show that over-stylized samples now attain lower loss values than under-stylized ones under the new AST style loss $\hat{\mathcal{L}}_{AST_s}$.

Study V: New AST Style Loss vs. Human Score. We investigate the relationship between our new balanced AST style loss $\hat{\mathcal{L}}_{AST_s}$ and human scores for the samples generated in Study II. These results are presented in Table 3. Unlike the negative correlation between the classic AST style loss and human scores (see Table 2), the new balanced AST style loss is positively correlated with human scores. Despite the treatment of ‘‘OK’’ annotations as zero scores in human study analysis, this is a strong indication that the proposed new AST style loss aligns much better with human scores than the classic AST style loss.

5. Experimental Evaluation

In this section, we provide qualitative and quantitative evaluation of the proposed new loss and its benefits in comparison with the conventional AST style loss.

5.1. Experiment Settings

We reuse the pretrained models listed in Table 1 and train additional models with either the classic or the new style loss in order to validate the generalizability of the improvements due to the latter. These models are listed in Table 4. We use content images from MS-COCO [33] and style images from Painter by Numbers [36] to train the models. The same four layers ($\mathbb{L}_{AST_s} = \{\mathcal{F}_{b1}^{r2}, \mathcal{F}_{b2}^{r2}, \mathcal{F}_{b3}^{r3}, \mathcal{F}_{b4}^{r3}\}$) of the ImageNet-trained VGG-16 model that were used in Studies I–V were employed here to calculate style losses for training. The \mathcal{F}_{b3}^{r3} layer was reused to calculate content losses, following previous works [13, 21]. In order to conduct fair

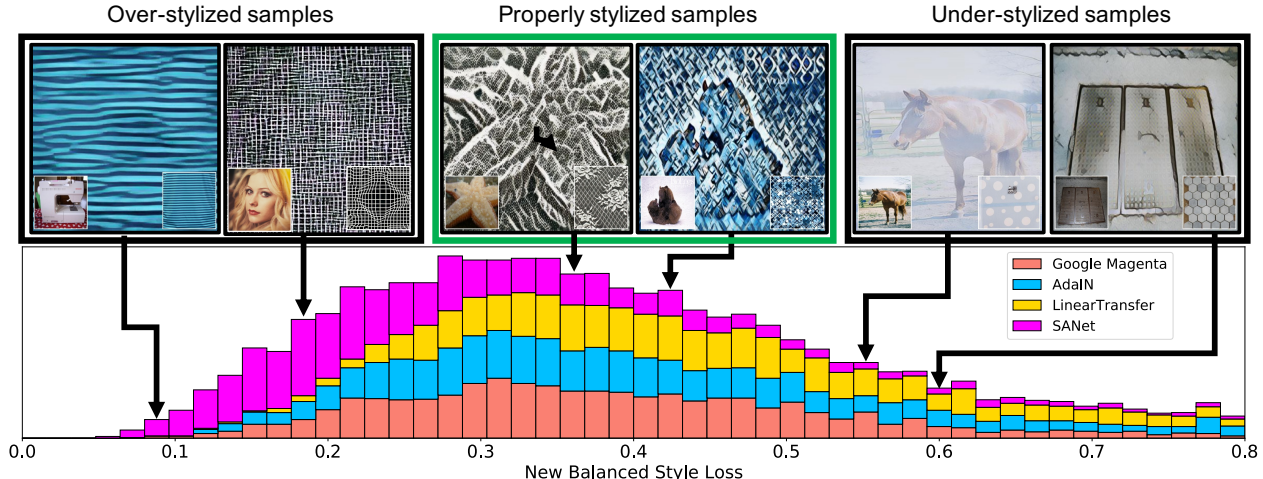


Figure 5: Distribution of our style-balanced loss for four AST methods [14, 19, 29, 37]. Small loss values indicate over-stylization while large values correspond to under-stylization, with properly-stylized samples in the middle, as expected.

Table 4: The AST models used in experimental evaluation. “3P” indicates publicly available pretrained models, while “1P” shows the settings used in our experiments. We train all models using VGG-16-based losses for consistency.

Model Name	AST Method	AST Style Loss	1P/3P
GoogleMagenta	GoogleMagenta	Classic + VGG-16	3P
OurGM	GoogleMagenta	Classic + VGG-16	1P
OurBalGM	GoogleMagenta	Balanced + VGG-16	1P
AdaIN	AdaIN	Classic + VGG-19	3P
OurAI	AdaIN	Classic + VGG-16	1P
OurBalAI	AdaIN	Balanced + VGG-16	1P
LinearTransfer	LinearTransfer	Classic + VGG-19	3P
OurLT	LinearTransfer	Classic + VGG-16	1P
OurBalLT	LinearTransfer	Balanced + VGG-16	1P
SANet	SANet	Classic + VGG-19	3P
OurSAN	SANet	Classic + VGG-16	1P
OurBalSAN	SANet	Balanced + VGG-16	1P

comparison of models, we pick style vs. content trade-off weights β in the overall loss (Equation (1)) that ensure that the magnitudes of the style and content losses are similar. We use the same optimizers that were used by the authors of the four models as reported in literature. We use “1P” models for the classic AST style loss related comparisons. We do not backpropagate gradient through the denominator to obtain more stable training and better results.

5.2. Qualitative Evaluation

Comparison with Classic Loss. Figure 6 shows a few visual examples of under- and over-stylization that are mitigated by training models with our new balanced style loss. As evident, our loss is effective in both cases and generalizes across all tested AST models, providing a better trade-

off between content and style. Results of under-stylized samples show that our loss helps in capturing both global and low-level texture-related style information where models trained with the classic loss only contain style-color. Results also show that our loss can preserve more content in cases of over-stylization. While content is completely unrecognizable in over-stylized images due to the classic loss, our loss produces both visible content and proper stylization. More results are presented in supplementary material. **Comparison with Style Interpolation.** Style interpolation is a common method for fusing styles by combining style features of different images before decoding [14, 19, 29, 43]. Typically, an interpolation coefficient is employed to control the contributions of different styles. Since over-stylization is mainly about transferring too much style, one plausible remedy is to apply style interpolation between the style and content (*i.e.* treating the content image as a new style). However, as shown in Figure 7, it is not very effective. In contrast, stylization using our balanced loss (Figure 7(c)) provides superior results in better preserving content while properly transferring style. Last but not least, finding a “good” interpolation coefficient is not a trivial task, since it will be different for different styles.

5.3. Quantitative Evaluation

AST Style Loss Comparison. We compute the classic and the new balanced AST style losses for all the testing samples for each AST model. Results are presented in Table 5. As evident, our models trained with the new balanced loss always attain significantly lower overall loss values. This is because the new loss allows more tasks ((S, C) combinations) to be trained fairly, achieving lower classic style losses for the same pairs of content and style images.

Deception Rate. Sanakoyeu *et al.* [42] introduced Decep-

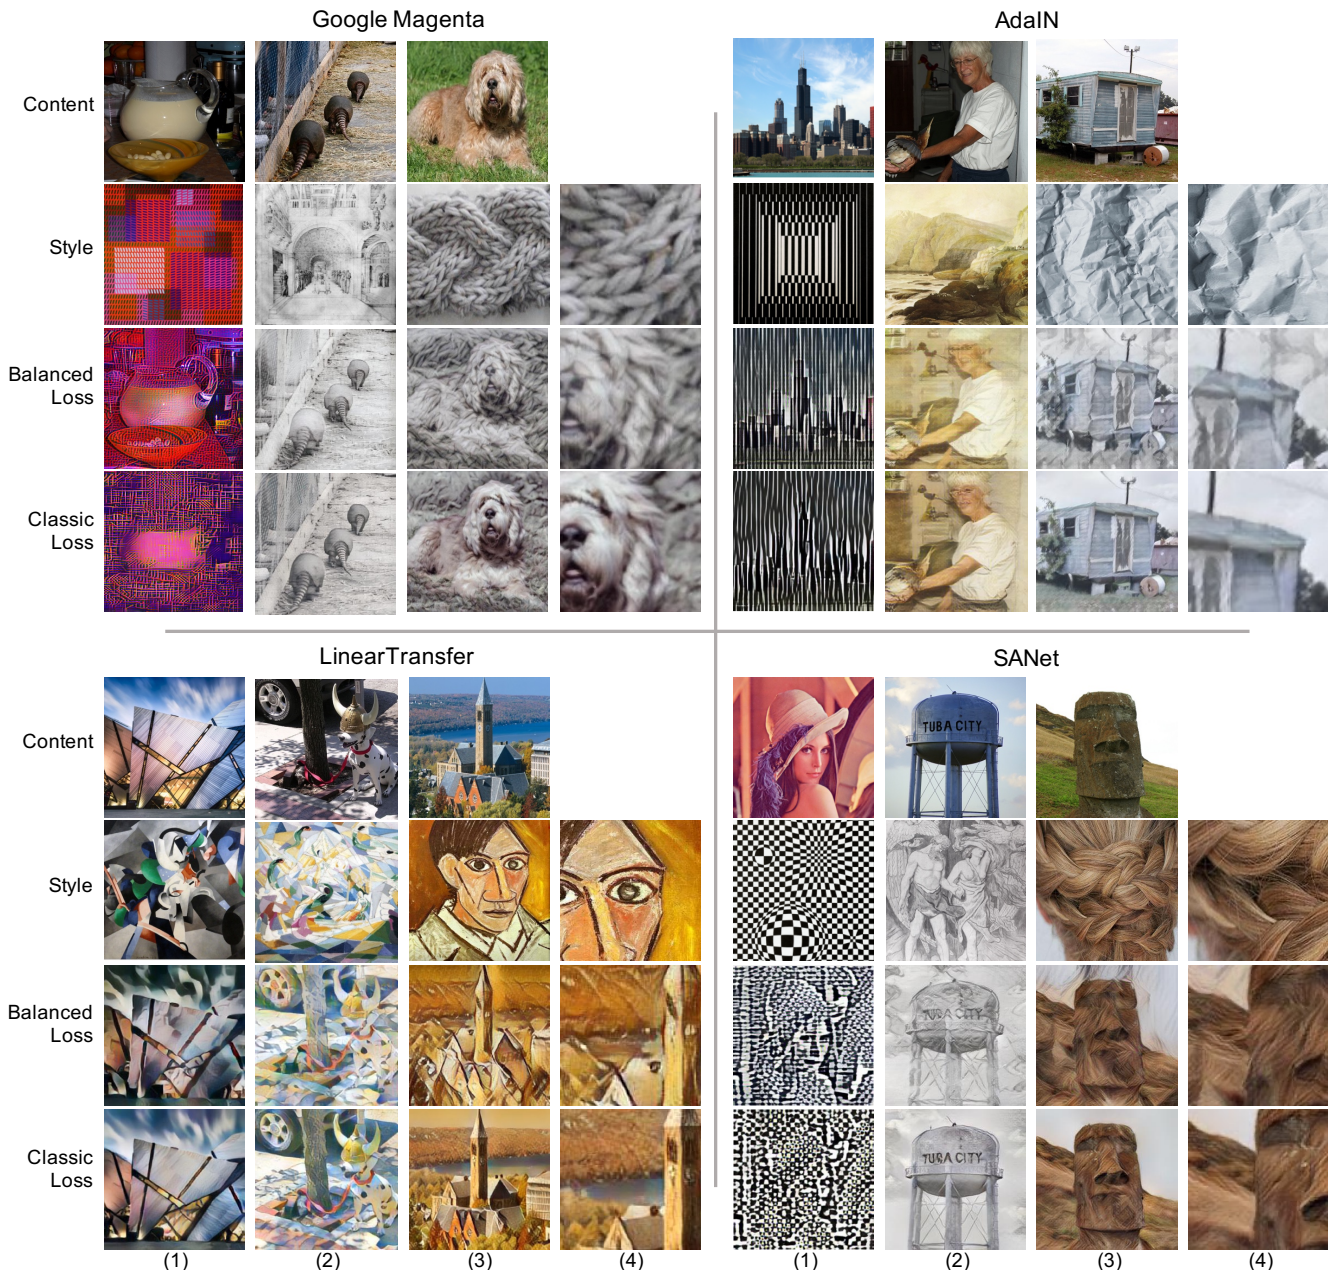


Figure 6: Improvements in cases that were under- or over-stylized due to the classic loss. In the former case, local and global texture of stylized images are closer to style images when our new loss is used, e.g., columns (3) and (4) for Google Magenta. In the latter case, original content is more visible in stylized images when our loss is used, e.g., column (1) of AdaIN.

tion Rate as a metric for evaluating style transfer quality. It is defined as the success rate of stylized images at deceiving an artist classification model, such that the same artist is predicted for both the style image and the stylized image.

We generate 5,000 stylized images for each AST model with content images from the ImageNet test set and style images from the Painter by Numbers (PBN) test set in order to ensure that the images and styles used in this experiment do not overlap with the training dataset. Further-

more, we exclude style images from artists who were seen in the training data or have less than 30 paintings in the test set. This results in 1,798 style images (paintings) from 34 artists. Content-style pairs were randomly sampled to generate the stylized images for evaluation.

We use the winning solution of the PBN challenge (<https://github.com/inejc/painters>) to compute the deception rate. First, we use the model to generate 2,048-dimensional features for all the style images. Next, for each stylized im-

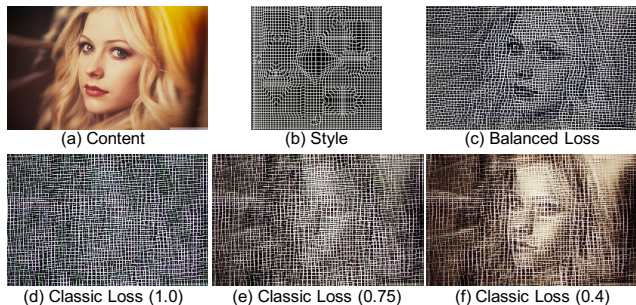


Figure 7: Our new balanced style loss is more effective in mitigating the over-stylization issue than the classic style interpolation approach using the GoogleMagenta solution. The style interpolation coefficient is presented in braces.

Table 5: Loss comparison between models trained with the classic loss and those trained with our new loss. Models trained with our loss always attain lower overall loss values.

Model Name	Classic Loss	New Balanced Loss
GoogleMagenta	4.69×10^8	0.36
OurGM	3.40×10^8	0.33
OurBalGM	3.35×10^8	0.22
AdaIN	7.05×10^8	0.33
OurAI	6.62×10^8	0.43
OurBalAI	6.58×10^8	0.31
LinearTransfer	6.11×10^8	0.33
OurLT	6.78×10^8	0.47
OurBalLT	4.27×10^8	0.25
SANet	5.00×10^8	0.28
OurSAN	5.25×10^8	0.41
OurBalSAN	4.03×10^8	0.21

Table 6: Deception Rate (%) for models trained with classic and our balanced style losses, shown in columns ‘‘Classic Loss’’ and ‘‘Bal. Loss’’, respectively. ‘‘Imp.’’ and ‘‘RImp.’’ show absolute and relative improvements, respectively.

AST Method	Classic Loss	Bal. Loss	Imp.	RImp.
GoogleMagenta	16.84%	35.38%	18.54%	110%
AdaIN	10.78%	22.18%	11.40%	106%
LinearTransfer	18.58%	39.18%	20.60%	111%
SANet	18.08%	33.16%	15.08%	83%

age, we extract its features using the same model and find its nearest style image using L_2 distance. For a successful deception, the artist of the nearest neighbor must match that of the style image used for generating the stylized sample.

Table 6 summarizes the results, showing large boosts in deception rate from adopting our new loss. The improvements range from a massive 83% minimum relative improvement for SANet to 111% (or 2.1x) for LinearTransfer.

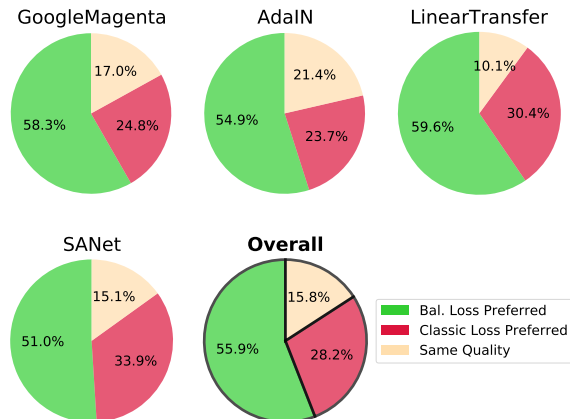


Figure 8: Human evaluation results. On all methods, users prefer images stylized with models trained with our loss.

Human Evaluation. We further quantify improvement in stylization quality due to our new loss through human evaluation. We randomly sample 5,000 content-style pairs with content images from the ImageNet test set and style images each from the test sets of PBN [36] and Describable Textures Dataset (DTD) [7]. For each pair and for each of the four AST methods (Table 1), we generate two stylized images: one using the model trained with the classic loss and another with that trained using our new loss. Random subsets of images are then voted on by annotators, who are requested to select one of the two stylized images for each pair based on their preference of stylization quality. In total, 2,200 annotations were collected, and show that 28.2% votes prefer stylized images from models trained with the classic loss while a much larger 55.9% prefer those trained with our new balanced loss, with 15.8% finding them of similar quality. Figure 8 provides the breakdown for each model, showing that models trained with our loss generate results that are more preferred by human perception.

6. Conclusion

In this work, we systematically studied the discrepancy between the classic AST style loss and human perception of stylization quality. We identified the root cause of the issue as the style-agnostic aggregation of sample-wise losses during training and derived theoretical bounds of the style loss to design a new style-balanced loss with style-aware normalization. We showed that unlike the classic loss, our new loss is positively correlated with human perception. Finally, experimental results show up to 111% and 98% relative improvements in Deception Rate and human preference, respectively. Future work can adopt our new loss in related problems, *e.g.*, video [1, 3, 5, 12, 18, 29, 40, 41] and photo [29, 35, 38, 53] stylization, texture synthesis [10, 11, 26, 51], *etc.* Future work can also derive tighter bounds for the style loss to improve style-aware normalization.

References

- [1] Dongdong Chen, Jing Liao, Lu Yuan, Nenghai Yu, and Gang Hua. Coherent online video style transfer. In *IEEE International Conference on Computer Vision*, pages 1105–1114, 2017.
- [2] Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, and Gang Hua. Stylebank: An explicit representation for neural image style transfer. In *IEEE conference on Computer Vision and Pattern Recognition*, pages 1897–1906, 2017.
- [3] Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, and Gang Hua. Stereoscopic neural style transfer. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6654–6663, 2018.
- [4] Tian Qi Chen and Mark Schmidt. Fast patch-based style transfer of arbitrary style. In *NIPS Workshop on Constructive Machine Learning*, 2016.
- [5] Xinghao Chen, Yiman Zhang, Yunhe Wang, Han Shu, Chun-jing Xu, and Chang Xu. Optical flow distillation: Towards efficient and stable video style transfer. In *European Conference on Computer Vision*, pages 614–630. Springer, 2020.
- [6] Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, pages 794–803. PMLR, 2018.
- [7] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *IEEE conference on Computer Vision and Pattern Recognition*, 2014.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE conference on Computer Vision and Pattern Recognition*, pages 248–255. Ieee, 2009.
- [9] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *arXiv preprint arXiv:1610.07629*, 2016.
- [10] Alexei A Efros and William T Freeman. Image quilting for texture synthesis and transfer. In *Conference on Computer Graphics and Interactive Techniques*, pages 341–346, 2001.
- [11] Alexei A Efros and Thomas K Leung. Texture synthesis by non-parametric sampling. In *IEEE International Conference on Computer Vision*, volume 2, pages 1033–1038. IEEE, 1999.
- [12] Wei Gao, Yijun Li, Yihang Yin, and Ming-Hsuan Yang. Fast video multi-style transfer. In *IEEE Winter Conference on Applications of Computer Vision*, pages 3222–3230, 2020.
- [13] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *IEEE conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016.
- [14] Golnaz Ghiasi, Honglak Lee, Manjunath Kudlur, Vincent Dumoulin, and Jonathon Shlens. Exploring the structure of a real-time, arbitrary neural artistic stylization network. In *British Machine Vision Conference*, 2017.
- [15] Shuyang Gu, Congliang Chen, Jing Liao, and Lu Yuan. Arbitrary style transfer with deep feature reshuffle. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8222–8231, 2018.
- [16] Michelle Guo, Albert Haque, De-An Huang, Serena Yeung, and Li Fei-Fei. Dynamic task prioritization for multi-task learning. In *European Conference on Computer Vision*, pages 270–287, 2018.
- [17] Zhiyuan Hu, Jia Jia, Bei Liu, Yaohua Bu, and Jianlong Fu. Aesthetic-aware image style transfer. In *ACM International Conference on Multimedia*, pages 3320–3329, 2020.
- [18] Haozhi Huang, Hao Wang, Wenhan Luo, Lin Ma, Wenhao Jiang, Xiaolong Zhu, Zhifeng Li, and Wei Liu. Real-time neural style transfer for videos. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 783–791, 2017.
- [19] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE International Conference on Computer Vision*, pages 1501–1510, 2017.
- [20] Yongcheng Jing, Xiao Liu, Yukang Ding, Xinchao Wang, Errui Ding, Mingli Song, and Shilei Wen. Dynamic instance normalization for arbitrary style transfer. In *AAAI Conference on Artificial Intelligence*, volume 34, pages 4369–4376, 2020.
- [21] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711. Springer, 2016.
- [22] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *IEEE conference on Computer Vision and Pattern Recognition*, pages 7482–7491, 2018.
- [23] Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. Style transfer by relaxed optimal transport and self-similarity. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 10051–10060, 2019.
- [24] Dmytro Kotovenko, Artsiom Sanakoyeu, Sabine Lang, and Bjorn Ommer. Content and style disentanglement for artistic style transfer. In *IEEE International Conference on Computer Vision*, pages 4422–4431, 2019.
- [25] Dmytro Kotovenko, Artsiom Sanakoyeu, Pingchuan Ma, Sabine Lang, and Bjorn Ommer. A content transformation block for image style transfer. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 10032–10041, 2019.
- [26] Vivek Kwatra, Arno Schödl, Irfan Essa, Greg Turk, and Aaron Bobick. Graphcut textures: image and video synthesis using graph cuts. *ACM Transactions on Graphics*, 22(3):277–286, 2003.
- [27] Chuan Li and Michael Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. In *European Conference on Computer Vision*, pages 702–716. Springer, 2016.
- [28] Shaohua Li, Xinxing Xu, Liqiang Nie, and Tat-Seng Chua. Laplacian-steered neural style transfer. In *ACM International Conference on Multimedia*, pages 1716–1724, 2017.
- [29] Xueting Li, Sifei Liu, Jan Kautz, and Ming-Hsuan Yang. Learning linear transformations for fast image and video style transfer. In *IEEE conference on Computer Vision and Pattern Recognition*, 2019.
- [30] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Diversified texture synthesis with

- feed-forward networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3920–3928, 2017.
- [31] 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.
- [32] Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. Demystifying neural style transfer. In *International Joint Conference on Artificial Intelligence*, pages 2230–2236, 2017.
- [33] 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.
- [34] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1871–1880, 2019.
- [35] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4990–4998, 2017.
- [36] K Nichol. *Painter by numbers*, 2016.
- [37] Dae Young Park and Kwang Hee Lee. Arbitrary style transfer with style-attentional networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5880–5888, 2019.
- [38] Sebastian Penhouët and Paul Sanzenbacher. Automated deep photo style transfer. *arXiv preprint arXiv:1901.03915*, 2019.
- [39] Eric Risser, Pierre Wilnot, and Connelly Barnes. Stable and controllable neural texture synthesis and style transfer using histogram losses. *arXiv preprint arXiv:1701.08893*, 2017.
- [40] Manuel Ruder, Alexey Dosovitskiy, and Thomas Brox. Artistic style transfer for videos. In *German Conference on Pattern Recognition*, pages 26–36. Springer, 2016.
- [41] Manuel Ruder, Alexey Dosovitskiy, and Thomas Brox. Artistic style transfer for videos and spherical images. *International Journal of Computer Vision*, 126(11):1199–1219, 2018.
- [42] Artsiom Sanakoyeu, Dmytro Kotovenko, Sabine Lang, and Bjorn Ommer. A style-aware content loss for real-time hd style transfer. In *European Conference on Computer Vision*, pages 698–714, 2018.
- [43] Falong Shen, Shuicheng Yan, and Gang Zeng. Neural style transfer via meta networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8061–8069, 2018.
- [44] Lu Sheng, Ziyi Lin, Jing Shao, and Xiaogang Wang. Avatar-net: Multi-scale zero-shot style transfer by feature decoration. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8242–8250, 2018.
- [45] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [46] Jan Svoboda, Asha Anoopshah, Christian Osendorfer, and Jonathan Masci. Two-stage peer-regularized feature recombination for arbitrary image style transfer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13816–13825, 2020.
- [47] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. In *International Conference on Machine Learning*, volume 1, page 4, 2016.
- [48] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6924–6932, 2017.
- [49] Huan Wang, Yijun Li, Yuehai Wang, Haoji Hu, and Ming-Hsuan Yang. Collaborative distillation for ultra-resolution universal style transfer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1860–1869, 2020.
- [50] Zhizhong Wang, Lei Zhao, Haibo Chen, Lihong Qiu, Qihang Mo, Sihuan Lin, Wei Xing, and Dongming Lu. Diversified arbitrary style transfer via deep feature perturbation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7789–7798, 2020.
- [51] Li-Yi Wei and Marc Levoy. Fast texture synthesis using tree-structured vector quantization. In *Conference on Computer Graphics and Interactive Techniques*, pages 479–488, 2000.
- [52] Yuan Yao, Jianqiang Ren, Xuansong Xie, Weidong Liu, Yong-Jin Liu, and Jun Wang. Attention-aware multi-stroke style transfer. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1467–1475, 2019.
- [53] Jonghwa Yim, Jisung Yoo, Won-joon Do, Beomsu Kim, and Jihwan Choe. Filter style transfer between photos. In *European Conference on Computer Vision*, pages 103–119. Springer, 2020.
- [54] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. In *IEEE International Conference on Computer Vision*, pages 9036–9045, 2019.
- [55] Hang Zhang and Kristin Dana. Multi-style generative network for real-time transfer. In *European Conference on Computer Vision*, pages 0–0, 2018.