

Content-Style Decoupling for Unsupervised Makeup Transfer without Generating Pseudo Ground Truth

Zhaoyang Sun¹ Shengwu Xiong^{1,2,3,4} Yaxiong Chen¹ Yi Rong^{1†}
¹Wuhan University of Technology ²Wuhan Huaxia Institute of Technology
³Shanghai AI Laboratory ⁴Qiongtai Normal University
zhaoyangsun0304@outlook.com, {xiongsw, chen yaxiong, yrong}@whut.edu.cn

Abstract

The absence of real targets to guide the model training is one of the main problems with the makeup transfer task. Most existing methods tackle this problem by synthesizing pseudo ground truths (PGTs). However, the generated PGTs are often sub-optimal and their imprecision will eventually lead to performance degradation. To alleviate this issue, in this paper, we propose a novel Content-Style Decoupled Makeup Transfer (CSD-MT) method, which works in a purely unsupervised manner and thus eliminates the negative effects of generating PGTs. Specifically, based on the frequency characteristics analysis, we assume that the low-frequency (LF) component of a face image is more associated with its makeup style information, while the high-frequency (HF) component is more related to its content details. This assumption allows CSD-MT to decouple the content and makeup style information in each face image through the frequency decomposition. After that, CSD-MT realizes makeup transfer by maximizing the consistency of these two types of information between the transferred result and input images, respectively. Two newly designed loss functions are also introduced to further improve the transfer performance. Extensive quantitative and qualitative analyses show the effectiveness of our CSD-MT method. Our code is available at <https://github.com/Snowfallingplum/CSD-MT>.

1. Introduction

Given a pair of source and reference face images, the main goal of makeup transfer is to generate an image that simultaneously satisfies the following conditions: (1) Containing the makeup styles transferred from the reference image, such as lipstick, eye shadow and powder blush. (2) Preserving the content details of the source image, including iden-

tity, facial structure and background. This technique has been widely studied and is attracting increasing attentions from the computer vision and artificial intelligence communities, due to its great economic potential in the fields of e-commerce, entertainment and beauty industries. However, considering the diversity and complexity of different makeup styles, makeup transfer remains a challenging task.

One of the major problems with the makeup transfer task is its unsupervised nature, which means that there is no real transferred image that can be used as a target ground truth to guide the model training. To address this problem, most existing methods [3, 7, 9, 13, 17, 21, 23, 27, 30, 32] propose to synthesize a pseudo ground truth (PGT) image, from each input source-reference image pair, as an alternative supervision target. After that, the model parameters are optimized by minimizing the difference between each generated transferred output and its corresponding PGT (see Figure 1(a)).

As illustrated in Figure 2, according to the PGT generation strategy used, previous makeup transfer approaches can be roughly divided into two categories: (1) **Histogram-matching-based methods** [3, 9, 13, 15, 32] attempt to align the color distribution of each facial region (e.g., lip, eye and face areas) in the source image with that of the same region in the reference face. However, the PGT produced by this strategy discards all spatial information of makeup styles, and usually suffers from the large color difference (e.g., eye shadow and powder blusher) between the source and reference images. (2) **Geometric-distortion-based methods** [7, 21] synthesize the PGT by warping the reference face so that its shape (typically represented by some facial landmarks) is matched to that of the source one. But such process often introduces undesired artifacts and also leads to the loss of source content information. As a result, these low-quality PGTs will consequently degrade the transfer performance of all above-mentioned methods. Although a recent effort [27] has been made to create more effective PGTs through a hybrid strategy, the generated PGTs are still sub-optimal and their imprecision will severely affect the final transferred results.

[†]Corresponding author.

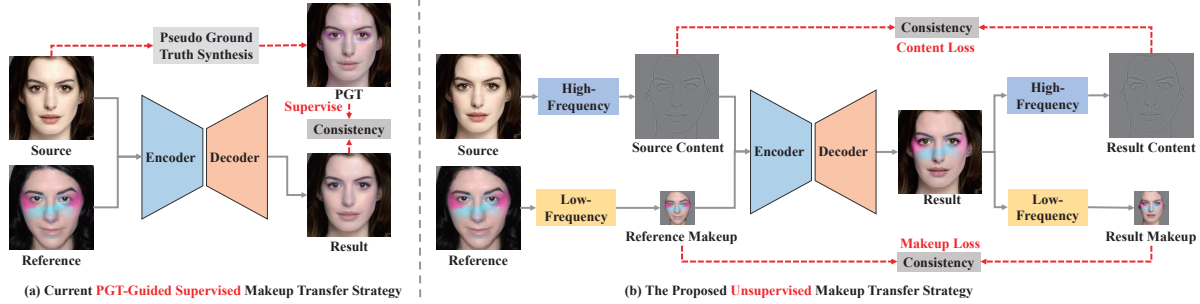


Figure 1. The comparison of different training strategies.




Figure 2. The PGTs and transferred results generated by different categories of makeup transfer methods.

To eliminate these negative effects of the PGT, in this paper, we propose a new Content-Style Decoupled Makeup Transfer (CSD-MT) method, which works in a purely unsupervised manner without generating any PGT. To achieve this, one important observation is that the main differences of the same face image before and after makeup are concentrated in its low-frequency (LF) component, while the high-frequency (HF) component remains almost unchanged. Therefore, we can assume that *the LF component of a face image is more associated with its makeup style information, and the HF component is more related to its content details*. With this assumption, CSD-MT first performs frequency decomposition on each input and output image to decouple their contents and makeup styles. Then, for model training, CSD-MT simultaneously maximizes the content and makeup consistencies of the transferred result with the source and reference images based on their HF and LF components, respectively, as shown in Figure 1(b). Additionally, we introduce two novel loss functions to enhance the transfer of the spatial and color information in makeup. The effectiveness of our proposed CSD-MT method is evaluated on three publicly available datasets, covering various makeup styles as well as different pose and expression variations. Our main contributions are summarized as follows:

- We propose a novel Content-Style Decoupled Makeup Transfer (CSD-MT) method, which works in a purely un-

supervised manner. Based on frequency decomposition, CSD-MT for the first time eliminates the requirement of generating pseudo ground truth.

- Two newly designed loss functions, namely the self-augmented reconstructive loss and the color contrastive loss, are introduced to facilitate a better transfer of the spatial and color information in makeup.
- Extensive quantitative and qualitative comparisons on three datasets indicate that CSD-MT outperforms seven state-of-the-art makeup transfer methods. In addition, the ablation study validates the superiority of our proposed unsupervised learning strategy over PGT-guided training.

2. Related Works

2.1. Makeup Transfer

During the past decade, makeup transfer has attracted increasing attention from the computer vision community. According to the PGT generation strategy, the previous methods can be roughly divided into two categories. (1) **For histogram-matching-based methods**, BeautyGAN [13] pioneers a histogram matching loss and designs a dual input/output GAN to perform makeup transfer and removal simultaneously. To handle misaligned head poses and facial expressions, SCGAN [3] encodes component-wise makeup regions into spatially-invariant style codes, while PSGAN

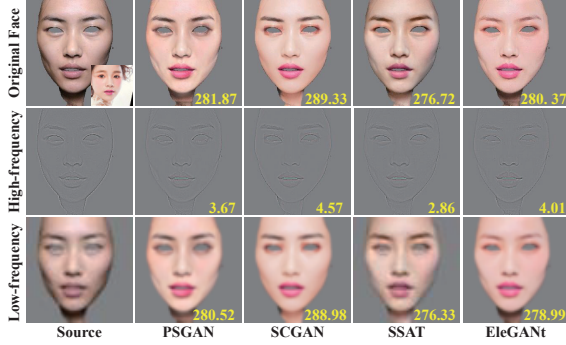


Figure 3. Visualization of the frequency components decomposed from the source image and the transferred results. The low-frequency components are resized for better visualization. The mean square errors of the different components between source images and transferred results are marked in the lower left corner.

[9, 15] utilizes an attention mechanism to adaptively deform the makeup feature maps based on source images. CPM [17] bypasses semantic alignment by converting the images to UV [5] space, where each pixel corresponds to a specific semantic point on the face. RamGAN [26] and SpMT [32] explore local attention to eliminate potential associations between different makeup components. (2) **For geometric-distortion-based methods**, PairedCycleGAN [2] trains a style discriminator to measure the makeup consistency between the results and the reference images. For instances of extreme makeup, LADN [7] employs multiple overlapping local makeup style discriminators. To ensure color fidelity, FAT [23] and SSAT [21] utilize cross-attention [22] to calculate semantic correspondence between the two input images. In addition, the recent EleGANt [27] achieves a more effective PGT using a hybrid strategy that integrates the advantages of these two PGTs by dynamically assigning different weights.

Different from all the above methods, whose transfer performance is heavily influenced by the quality of PGTs, the goal of our CSD-MT is to investigate a PGT-free makeup transfer approach to eliminate the negative effects of generating PGTs.

2.2. Frequency Decomposition

Frequency decomposition has shown its effectiveness in various tasks, including classification [6, 25], image synthesis [4, 11], and image translation [12, 14, 28]. For instance, LRR [6] utilizes the Laplacian pyramid [1] to refine the boundary details of semantic segmentation. LapSRN [11] consists of multiple generators that progressively reconstruct the HF residuals of high-resolution images. WTC² [28] employs wavelet transform to accelerate the stylization process of high-resolution images. For makeup transfer, we observe that the main differences of the same face image

before and after makeup are concentrated in its LF component, while the HF component remains almost unchanged. Therefore, unlike the methods mentioned above, our goal of frequency decomposition is to decouple the content information and makeup style from a face image.

3. Methodology

3.1. Problem Statement

Let \mathcal{X} and \mathcal{Y} denote the non-makeup source domain and the makeup reference domain, respectively. In general, the image samples in \mathcal{X} and \mathcal{Y} are unpaired, which means that the source and reference images are collected from different persons with distinct identity information, and each reference face showcases a unique makeup style. Given a pair of source and reference images $\{(x, y) | x \in \mathcal{X}, y \in \mathcal{Y}\}$ as input, the main goal of makeup transfer is to generate a transferred result \hat{x} , which maximally preserves the content information in x and contains the same makeup style as y . Such task can be considered as a cross-domain image-to-image translation problem with specific conditions, while its unsupervised nature makes it even more challenging.

3.2. Content and Makeup Style Decoupling

In order to solve the makeup transfer task without generating PGT, we attempt to seek the content and makeup style supervision signals from the input images (x, y) themselves to accurately control the corresponding information in \hat{x} . We approach this purpose by investigating the frequency characteristics of these two types of information. To do this, we randomly select 500 pairs of test images from the MT dataset [13], and perform frequency decomposition on each source image and its corresponding transferred results generated by different methods. More specifically, given an arbitrary image $x \in \mathbb{R}^{H \times W \times 3}$ where H and W denote its height and width, we first remove its background region x_{bg} through a face parsing technique [29]. After that, by applying a fixed Gaussian kernel on the remaining foreground face image x_{fg} , we calculate a low-pass prediction $x_l \in \mathbb{R}^{\frac{H}{d} \times \frac{W}{d} \times 3}$, where d represents a downsampling factor. Based on this, the high-frequency residual $x_h \in \mathbb{R}^{H \times W \times 3}$ is finally obtained by $x_h = x_{fg} - up(x_l)$, where $up(\cdot)$ is a bilinear interpolation upsampling operation.

We measure the mean squared errors (MSE) of these decomposed LF and HF components between source and transferred images in Figure 3. It can be observed that the MSE values calculated on the LF components are much larger than those obtained on the HF components. This suggests that the main differences of the same face images before and after makeup are primarily concentrated in their LF components, while the HF components remain almost unchanged. Additionally, the visual comparisons displayed in Figure 3 also support this claim. Therefore, we can assume

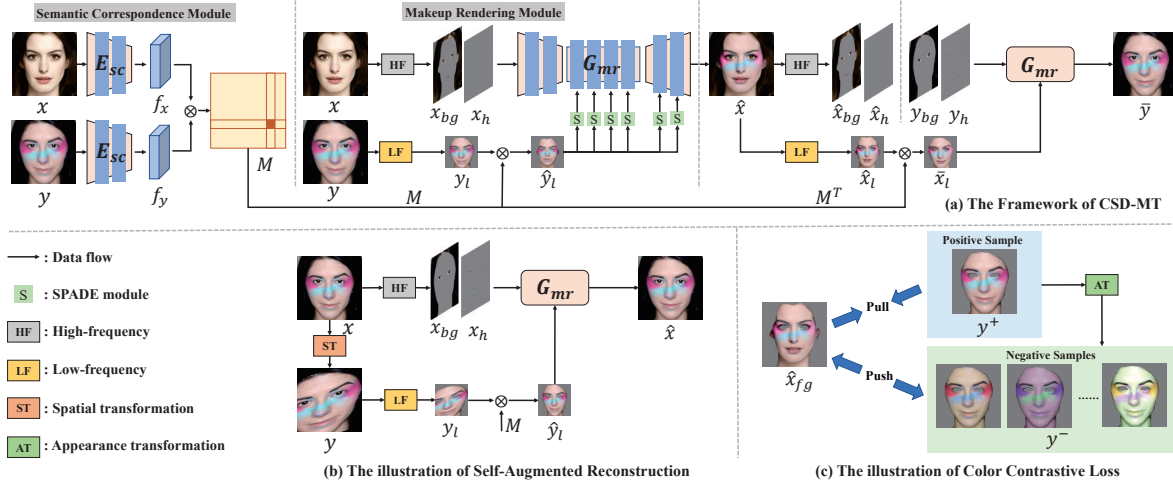


Figure 4. Illustration of the proposed CSD-MT framework. (a) Given a source image x and a reference image y , the semantic correspondence module first constructs a pixel-wise correlation matrix M between them. Next, by performing face parsing and frequency decomposition, the makeup rendering module G_{mr} obtains the background area x_{bg} and the HF component x_h that contain the content information of x , as well as the LF component y_l comprising the makeup style of y . Then, each pixel in \hat{y}_l aggregates the information from the corresponding pixels in y_l according to the correlation matrix M . Finally, the final transferred result $\hat{x} = G_{mr}([x_{bg}, x_h], \hat{y}_l)$ is generated using x_{bg} , x_h and \hat{y}_l . Furthermore, we introduce a self-augmented reconstructive loss (b) and a color contrastive loss (c) to enhance the transfer of the spatial and color information in makeup, respectively.

that the LF component of a face image is more associated with its makeup style information, and the HF component is more related to its content details. Such assumption allows us to appropriately decouple the content information and makeup style contained in each face image by using the frequency decomposition process described above.

3.3. The Proposed CSD-MT Method

To highlight that the improvement of transfer performance mainly comes from our proposed unsupervised learning strategy, we design a relative concise architecture for the proposed CSD-MT method. As illustrated in Figure 4, the generator \mathcal{G} of CSD-MT contains a semantic correspondence module and a makeup rendering module, which are presented in detail in the following subsections.

Semantic Correspondence Module. Generally, due to the differences in head pose and facial expression, the same facial parts in the input source image x and reference image y often appear at different spatial locations [9, 21, 27], and such semantic misalignment will eventually lead to performance degradation. To alleviate this problem, our semantic correspondence module constructs a pixel-wise correlation matrix M by calculating the cosine similarity as:

$$M(i, j) = \frac{f_x(i)^T f_y(j)}{\|f_x(i)\|_2 \|f_y(j)\|_2}. \quad (1)$$

Here, $f_x = E_{sc}(x)$, $f_y = E_{sc}(y)$ denote the semantic features extracted by a convolutional encoder $E_{sc}(\cdot)$. Both f_x

and f_y have the same spatial resolution as the LF component of input images, i.e., $\frac{H}{d} \times \frac{W}{d}$. $f(i)$ represents the feature vector of the i -th pixel in f and $M(i, j)$ indicates the element at the (i, j) -th location of M . We consider the correlation matrix M as a deformation mapping function, and use it to achieve semantic alignment between the source and reference images in our makeup rendering module.

Makeup Rendering Module. By performing face parsing and frequency decomposition on the input images, we can obtain the background area x_{bg} and the HF component x_h that contain the content information of the source image x , as well as the LF component y_l comprising the makeup style of the reference image y . Then, the correlation matrix M in Eq. (1) is used to spatially deform y_l as follow:

$$\hat{y}_l(i) = \sum_j Softmax(M(i, j)/\tau) \cdot y_l(j), \quad (2)$$

where $Softmax(\cdot)$ denotes a softmax computation along the column dimension, which normalizes the element values in each row of M , and $\tau > 0$ is a temperature parameter. Based on M , each pixel in \hat{y}_l aggregates the information from the corresponding pixels in y_l according to the semantic correspondence between x and y . Therefore, the deformed \hat{y}_l is semantically aligned with the source image.

Finally, the makeup rendering module generates the final transferred result \hat{x} based on x_{bg} , x_h and \hat{y}_l as:

$$\hat{x} = G_{mr}([x_{bg}, x_h], \hat{y}_l). \quad (3)$$

Here, $[\cdot, \cdot]$ denotes a channel-wise concatenation operation. $G_{mr}(\cdot, \cdot)$ is an encoder-decoder network implemented with




Figure 5. Qualitative comparison with several state-of-the-art methods on different makeup styles. The proposed CSD-MT produces the most precise transferred results with desired makeup information and high-quality content details. Please zoom in for better comparison.

the U-Net structure. In $G_{mr}(\cdot, \cdot)$, the conditional makeup information is introduced through the SPADE modules [18], whose modulation parameters are generated from \hat{y}_l .

3.4. Training Objectives

Transfer Loss. Similar to the input images, the transferred result \hat{x} produced by our generator \mathcal{G} can also be decomposed into the LF component \hat{x}_l and the HF component \hat{x}_h . By employing the transposed correlation matrix M^T as in Eq. (2), we re-deform \hat{x}_l into \bar{x}_l to make it semantically aligned with the reference image. According to our assumption, \bar{x}_l is expected to be consistent with y_l , such that the makeup style can be faithfully transferred. And meanwhile, \hat{x}_h is required to preserve the content information of the source image and thus should be consistent with x_h . Therefore, the following transfer loss is defined to simultaneously promote the makeup and content consistencies:

$$L_{trans} = L_{makeup} + \alpha L_{cont}, \quad (4)$$

$$L_{makeup} = \|\bar{x}_l - y_l\|_1, L_{cont} = GP(\hat{x}_h, x_h),$$

where $\alpha > 0$ balances the importance of the two terms. L_{makeup} is defined as the L1 distance between \bar{x}_l and y_l . L_{cont} calculates the Gradient Profile loss [19] $GP(\cdot, \cdot)$ between \hat{x}_h and x_h , which is computed in the image gradients space and thus is more powerful in distilling HF details.

Cycle Consistency Loss. Inspired by CycleGAN [31], we feed y_{bg}, y_h and \bar{x}_l into the makeup rendering network in Eq. (3) once again. We expect the produced transferred result $\bar{y} = G_{mr}([y_{bg}, y_h], \bar{x}_l)$ can be as close as possible to the original reference image y , which can be formulated as:

$$L_{cycle} = \|\bar{y} - y\|_1. \quad (5)$$

Adversarial Loss. To make the transferred results more realistic, we construct a multi-scale discriminator \mathcal{D} [24] to distinguish the face images generated by CSD-MT from the reference images containing real makeup information. Based on objective function of LSGAN [16], our adversarial loss L_{adv} is defined as follow:

$$L_{adv} = \mathbb{E}_y[(\mathcal{D}(y) - 1)^2] + \mathbb{E}_{\hat{x}}[(\mathcal{D}(\hat{x}))^2]. \quad (6)$$

Self-Augmented Reconstructive Loss. To further enhance the robustness to different head poses and facial expressions, we develop a self-augmented reconstruction process that facilitates a better transfer of the spatial information in makeup, as shown in Figure 4(b). Specifically, given an image x with makeup, we impose a random spatial transformation $T_s(\cdot)$ on it and obtain a transformed image $y = T_s(x)$. This image has the same makeup style as x , but the spatial information in the original x (e.g., facial structure and person identity) is completely destroyed. Considering x as the

Methods	MT			Wild-MT			LADN		
	FID	Self-Aug PSNR/SSIM		FID	Self-Aug PSNR/SSIM		FID	Self-Aug PSNR/SSIM	
		Crop	Rotate		Crop	Rotate		Crop	Rotate
BeautyGAN [13]	48.24	21.70/0.848	21.58/0.847	99.62	21.67/0.875	21.33/0.862	69.50	19.73/0.719	20.39/0.731
PSGAN [9]	45.02	23.46/0.886	22.68/0.875	89.92	22.56/0.879	21.89/0.870	57.80	22.97/0.824	22.16/0.815
SCGAN [3]	39.20	23.92/0.887	24.05/0.888	79.54	24.02/0.902	24.24/0.905	51.39	22.42/0.839	22.44/0.838
SpMT [32]	46.10	24.36/0.891	23.79/0.887	77.04	24.30/0.909	23.77/0.904	48.18	23.31/0.817	22.89/0.815
LADN [7]	73.91	19.27/0.759	19.02/0.755	104.91	19.22/0.790	19.09/0.786	65.87	18.61/0.727	18.39/0.723
SSAT [21]	38.01	24.01/0.894	23.93/0.893	70.53	23.57/0.905	23.80/0.907	53.84	22.81/0.848	22.95/0.846
EleGANt [27]	54.06	25.15/0.885	24.60/0.875	86.19	24.82/0.893	24.49/0.877	61.40	24.64/0.838	24.31/0.848
CSD-MT (Ours)	37.56	27.28/0.920	26.68/0.915	60.82	27.83/0.934	26.70/0.923	40.87	25.19/0.868	25.23/0.868

Table 1. Quantitative comparison of FID and Self-augmented PSNR/SSIM on the MT, Wild-MT and LADN datasets.

source image and y as the reference image, the generator \mathcal{G} of CSD-MT takes (x, y) as input and outputs a transferred result \hat{x} . Based on this process, \hat{x} should faithfully reconstruct x , since it contains the same content and makeup style (distilled from y) information as in x . Therefore, we define the following self-augmented reconstructive loss:

$$L_{aug} = \|\hat{x} - x\|_1. \quad (7)$$

Color Contrastive Loss. To promote the color fidelity, as shown in Figure 4(c), we propose a color contrastive loss which can be formulated as follow:

$$L_{cts} = -\log\left(1 - \frac{\ell(\hat{x}_{fg}, y^+)}{\sum_{i=1}^N \ell(\hat{x}_{fg}, y_i^-)}\right), \quad (8)$$

where \hat{x}_{fg} is the foreground face area separated from the transferred result \hat{x} . For this anchor, y^+ and y^- denote the positive and negative samples, respectively. In our implementation, the face area y_{fg} of the input reference image is used as the only positive sample. And each negative sample y_i^- is obtained by performing a random appearance transformation $T_a(\cdot)$ on y_{fg} , i.e., $y_i^- = T_a(y_{fg})$, N is the total number of negative samples. Based on the perceptual loss [10], the distance function $\ell(\cdot, \cdot)$ is defined as:

$$\ell(x, y) = \sum_l \|\text{Gram}(\phi_l(x)) - \text{Gram}(\phi_l(y))\|_1, \quad (9)$$

where $\phi_l(\cdot)$ represents the feature map extracted from the l -th layer of the pre-trained VGG19 [20] model. $\text{Gram}(\cdot)$ calculates the gram matrix of a feature map. By minimizing the color contrastive loss in Eq. (8), \hat{x}_{fg} and y_{fg} with similar color distributions are pulled closer, while \hat{x}_{fg} and y_i^- with different color distributions are pushed away.

Overall Loss. In summary, the overall loss function for the generator \mathcal{G} and discriminator \mathcal{D} of the proposed CSD-MT method is defined as:

$$\min_{\mathcal{G}} \max_{\mathcal{D}} L = \lambda_{trans} L_{trans} + \lambda_{cycle} L_{cycle} + \lambda_{adv} L_{adv} + \lambda_{aug} L_{aug} + \lambda_{cts} L_{cts}. \quad (10)$$

Methods	Simple	Complex	Extreme
SCGAN [3]	5.4%	1.7%	1.1%
SpMT [32]	11.1%	2.6%	1.5%
SSAT [21]	13.3%	9.8%	7.4%
EleGANt [27]	14.3%	15.4%	13.2%
CSD-MT (Ours)	55.9%	70.4%	76.7%

Table 2. The ratio selected as best (%) on different types of makeup styles. We classify "Simple", "Complex" and "Extreme" makeup based on our subjective experience.

4. Experiments

4.1. Datasets

MT Dataset [13] contains 1,115 non-makeup and 2,719 makeup images, which are mostly well-aligned and have plenty of makeup styles. We split the training and testing sets by following the strategy in [9, 13].

Wild-MT Dataset [9] consists of 369 non-makeup and 403 makeup images. Most of them contain large variations in head pose and facial expression.

LADN Dataset [7] has 333 non-makeup and 302 makeup images, including 155 extreme makeup images with great variances on makeup color, style and region coverage.

4.2. Baselines

We compare our proposed CSD-MT approach with seven state-of-the-art makeup transfer methods, including four histogram-matching-based methods (BeautyGAN [13], PSGAN [9], SCGAN [3], and SpMT [32]), two geometric-distortion-based methods (LADN [7], SSAT [21]), and one hybrid method (EleGANt [27]).

All these methods are trained by only using the training set of the MT dataset. And their performance and generalization ability are evaluated on the test set of the MT dataset, as well as on the Wild-MT and LADN datasets. See supplementary materials for the training details of CSD-MT.




Figure 6. Transferred results produced by the CSD-MT models trained with different learning strategies.




Figure 7. Reconstruction results obtained by rendering the de-makeup source image using a randomly rotated reference image.

4.3. Qualitative Comparison

Figure 5 displays the transferred results of all competing methods on various makeup styles. It can be seen that the histogram-matching-based methods all fail to work when there are large color differences between the source and reference images. For geometric-distortion-based methods, LADN generates unrealistic results with undesired artifacts, while SSAT cannot effectively transfer the makeup details such as the eye shadow and lipstick. The hybrid method EleGANt achieves better results than other baselines, but still struggles with transferring extreme makeup styles that distributed throughout the entire facial area. In contrast, for all types of makeup styles, our unsupervised CSD-MT method produces the most precise transferred results with desired makeup information and high-quality content details.

4.4. Quantitative Comparison

Fréchet Inception Distance (FID). Following [15], we calculate the FID score [8] (lower is better) between the reference images and the transferred results generated by different methods, which are reported in Table 1. The lowest FID scores achieved by our CSD-MT method indicate that its outputs are more realistic.

Methods	Self-Aug PSNR/SSIM	
	Crop	Rotate
CSD-MT with PSGAN's PGT	22.22/0.881	22.37/0.883
CSD-MT with SSAT's PGT	23.23/0.893	23.37/0.891
PSGAN with L_{aug} and L_{cts}	23.17/0.883	22.42/0.874
SSAT with L_{aug} and L_{cts}	23.78/0.894	23.64/0.891
(Setting A) $L_{adv} + L_{cont}$	19.97/0.866	20.01/0.867
(Setting B) $A + L_{makeup}$	23.01/0.872	22.77/0.871
(Setting C) $B + L_{cycle}$	24.92/0.883	24.61/0.882
(Setting D) $C + L_{aug}$	26.45/0.914	25.61/0.907
(Setting E) $D + L_{cts}$	27.28/0.920	26.68/0.915

Table 3. Quantitative results of ablation studies on the MT dataset.




Figure 8. Transferred results produced by the CSD-MT models trained with different loss functions.

Self-Augmented PSNR/SSIM. A major challenge in the performance evaluation for makeup transfer tasks is the lack of ground truth images. We utilize a similar self-augmented reconstruction process mentioned in section 3.4 to address this issue. As shown in Figure 7, given a makeup sample, we randomly crop or rotate it to obtain a pseudo reference image whose content information has been corrupted. We also generate a de-makeup pseudo source image using the makeup removal function of SSAT. With these pseudo inputs, we treat the original makeup image as the ground-truth, and compute the PSNR and SSIM (higher is better) for model evaluation. Both the quantitative and qualitative results in Table 1 and Figure 7 show that CSD-MT outperforms other state-of-the-art methods, demonstrating its effectiveness in generating high-quality results.



Figure 9. Robustness of CSD-MT in various complex scenarios.

User Study. We also conduct a user study to quantitatively evaluate the performance of different models. We randomly select 20 pairs of images with different types of makeup styles and generate the transferred results using different methods. Then, totally 23 participants are asked to choose the most satisfactory result based on the image quality and makeup similarity. For a fair comparison, the transferred results are shown simultaneously in a random order. The results of user study are shown in Table 2.

4.5. Ablation Study

PGT-guided vs. Unsupervised. To intuitively compare these two strategies, we keep the network architecture of CSD-MT unchanged and train this model by using the PGT-guided training process as in PSGAN and SSAT, respectively. From Table 3 and Figure 6, we can see that the outputs produced by the resulted models are very similar to those of PSGAN and SSAT. This suggests that the transfer performance is heavily influenced by the synthesized PGTs rather than the network structure. It can be also found that the original CSD-MT significantly outperforms these two models, demonstrating the superiority of our proposed unsupervised strategy over PGT-guided training. Additionally, we integrate the losses L_{aug} and L_{cts} into PSGAN and SSAT. As shown in Table 3, the influence of L_{aug} and L_{cts} on previous methods is slight, indicating that our unsupervised strategy fits better with these two losses.

Loss Functions. As shown in Table 3, to analyze the effect of different losses defined in section 3.4, we gradually add each loss into a basic setting ($L_{adv} + L_{cont}$), resulting in 5 different loss combinations (Setting A-E). The quantitative and qualitative results of the CSD-MT models trained with these settings are displayed in Table 3 and Figure 8, respectively. We can observe that the model trained with only $L_{adv} + L_{cont}$ can already preserve the content details effectively. This is mainly attributed to the content-style decoupling operation and the content consistency loss L_{cont} . By adding L_{makeup} and L_{cycle} , the makeup style information is transferred but some complex details are still miss-



Figure 10. Generalization of CSD-MT in unseen anime makeup.

ing. Further equipped with L_{aug} facilitates the transfer of spatial information, so the makeup can appear at the correct location on the face. Finally, L_{cts} ensures the color fidelity. **Robustness.** As shown in Figure 9, capturing the semantic correspondence between the source and reference images makes our CSD-MT model insensitive to age, pose and expression variations. In addition, separating the foreground and background areas through the face parsing allows our method to be unaffected by image occlusions.

Generalization. To evaluate the generalization ability of our method, we collect some anime makeup examples which have a significant domain gap with the training samples in the MT dataset and have never been encountered by the model before. The results are displayed in Figure 10.

Limitation. In CSD-MT, we assume that the HF component is more closely associated with the content details of face images. This assumption may result in CSD-MT being unable to handle cases with patterns.

5. Conclusion

In this paper, we propose an unsupervised makeup transfer method called CSD-MT to eliminate the negative effects of generating PGTs. Inspired by the observed frequency characteristics, CSD-MT decouples the content and makeup style information through frequency decomposition and realizes makeup transfer by maximizing the consistency of these two types of information between the transferred result and input images, respectively. Experiments demonstrate that our CSD-MT method significantly outperforms existing state-of-the-art methods in quantitative and qualitative analyses.

Acknowledgments. This work was supported in part by the National Key Research and Development Program of China (Grant No.2022ZD0160604), the Project of Sanya Yazhou Bay Science and Technology City (Grant No.SCKJ-JYRC-2022-76 and SKJC-2022-PTDX-031), the Young Scientists Fund of the National Natural Science Foundation of China (Grant No.62306219), and the CAAI-Huawei MindSpore Open Fund (Grant No.CAAIXSILJJ-2022-004A).

References

- [1] Peter J Burt and Edward H Adelson. The laplacian pyramid as a compact image code. In *Readings in computer vision*, pages 671–679. Elsevier, 1987. 3
- [2] Huiwen Chang, Jingwan Lu, Fisher Yu, and Adam Finkelstein. Pairedcyclegan: Asymmetric style transfer for applying and removing makeup. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 40–48, 2018. 3
- [3] Han Deng, Chu Han, Hongmin Cai, Guoqiang Han, and Shengfeng He. Spatially-invariant style-codes controlled makeup transfer. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 6549–6557, 2021. 1, 2, 6
- [4] Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. *Advances in neural information processing systems*, 28, 2015. 3
- [5] Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. Joint 3d face reconstruction and dense alignment with position map regression network. In *Proceedings of the European conference on computer vision (ECCV)*, pages 534–551, 2018. 3
- [6] Golnaz Ghiasi and Charless C Fowlkes. Laplacian pyramid reconstruction and refinement for semantic segmentation. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, pages 519–534. Springer, 2016. 3
- [7] Qiao Gu, Guanzhi Wang, Mang Tik Chiu, Yu-Wing Tai, and Chi-Keung Tang. Ladv: Local adversarial disentangling network for facial makeup and de-makeup. In *Proceedings of the IEEE/CVF International conference on computer vision*, pages 10481–10490, 2019. 1, 3, 6
- [8] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 7
- [9] Wentao Jiang, Si Liu, Chen Gao, Jie Cao, Ran He, Jiashi Feng, and Shuicheng Yan. Psgan: Pose and expression robust spatial-aware gan for customizable makeup transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5194–5202, 2020. 1, 3, 4, 6
- [10] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 694–711. Springer, 2016. 6
- [11] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 624–632, 2017. 3
- [12] Biwen Lei, Xiefan Guo, Hongyu Yang, Miaomiao Cui, Xuansong Xie, and Di Huang. Abpn: Adaptive blend pyramid network for real-time local retouching of ultra high-resolution photo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2108–2117, 2022. 3
- [13] Tingting Li, Ruihe Qian, Chao Dong, Si Liu, Qiong Yan, Wenwu Zhu, and Liang Lin. Beautygan: Instance-level facial makeup transfer with deep generative adversarial network. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 645–653, 2018. 1, 2, 3, 6
- [14] Jie Liang, Hui Zeng, and Lei Zhang. High-resolution photo-realistic image translation in real-time: A laplacian pyramid translation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9392–9400, 2021. 3
- [15] Si Liu, Wentao Jiang, Chen Gao, Ran He, Jiashi Feng, Bo Li, and Shuicheng Yan. Psgan++: Robust detail-preserving makeup transfer and removal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):8538–8551, 2021. 1, 3, 7
- [16] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2017. 5
- [17] Thao Nguyen, Anh Tuan Tran, and Minh Hoai. Lipstick ain’t enough: beyond color matching for in-the-wild makeup transfer. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 13305–13314, 2021. 1, 3
- [18] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2337–2346, 2019. 5
- [19] M Saquib Sarfraz, Constantin Seibold, Haroon Khalid, and Rainer Stiefelhagen. Content and colour distillation for learning image translations with the spatial profile loss. *arXiv preprint arXiv:1908.00274*, 2019. 5
- [20] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6
- [21] Zhaoyang Sun, Yaxiong Chen, and Shengwu Xiong. Ssat: A symmetric semantic-aware transformer network for makeup transfer and removal. In *Proceedings of the AAAI Conference on artificial intelligence*, pages 2325–2334, 2022. 1, 3, 4, 6
- [22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [23] Zhaoyi Wan, Haoran Chen, Jie An, Wentao Jiang, Cong Yao, and Jiebo Luo. Facial attribute transformers for precise and robust makeup transfer. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1717–1726, 2022. 1, 3
- [24] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. 5

- [25] Travis Williams and Robert Li. Wavelet pooling for convolutional neural networks. In *International conference on learning representations*, 2018. 3
- [26] Jianfeng Xiang, Junliang Chen, Wenshuang Liu, Xianxu Hou, and Linlin Shen. Ramgan: region attentive morphing gan for region-level makeup transfer. In *European Conference on Computer Vision*, pages 719–735. Springer, 2022. 3
- [27] Chenyu Yang, Wanrong He, Yingqing Xu, and Yang Gao. Elegant: Exquisite and locally editable gan for makeup transfer. In *European Conference on Computer Vision*, pages 737–754. Springer, 2022. 1, 3, 4, 6
- [28] Jaejun Yoo, Youngjung Uh, Sanghyuk Chun, Byeongkyu Kang, and Jung-Woo Ha. Photorealistic style transfer via wavelet transforms. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9036–9045, 2019. 3
- [29] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 325–341, 2018. 3
- [30] Honglun Zhang, Wenqing Chen, Hao He, and Yaohui Jin. Disentangled makeup transfer with generative adversarial network. *arXiv preprint arXiv:1907.01144*, 2019. 1
- [31] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 5
- [32] Mingrui Zhu, Yun Yi, Nannan Wang, Xiaoyu Wang, and Xinbo Gao. Semi-parametric makeup transfer via semantic-aware correspondence. *arXiv preprint arXiv:2203.02286*, 2022. 1, 3, 6