

## SSH: A Self-Supervised Framework for Image Harmonization

Yifan Jiang<sup>†</sup>, He Zhang<sup>‡</sup>, Jianming Zhang<sup>‡</sup>, Yilin Wang<sup>‡</sup>, Zhe Lin<sup>‡</sup>, Kalyan Sunkavalli<sup>‡</sup>,  
Simon Chen<sup>‡</sup>, Sohrab Amirghodsi<sup>‡</sup>, Sarah Kong<sup>‡</sup>, Zhangyang Wang<sup>†</sup>

<sup>†</sup>The University of Texas at Austin    <sup>‡</sup>Adobe Inc.

<sup>†</sup>{yifanjiang97, atlaswang}@utexas.edu,

<sup>‡</sup>{hezhan, jianmzha, yilwang, zlin, sunkaval, sichen, tamirgho, sakong}@adobe.com



Figure 1: A visual comparison between the image harmonization results of the direct composition (DC), the current state-of-the-art methods, and the proposed SSH. Best viewed in color and zoomed in.

### Abstract

Image harmonization aims to improve the quality of image compositing by matching the “appearance” (e.g., color tone, brightness and contrast) between foreground and background images. However, collecting large-scale annotated datasets for this task requires complex professional retouching. Instead, we propose a novel Self-Supervised Harmonization framework (SSH) that can be trained using just “free” natural images without being edited. We reformulate the image harmonization problem from a representation fusion perspective, which separately processes the foreground and background examples, to address the background occlusion issue. This framework design allows for a dual data augmentation method, where diverse [foreground, background, pseudo GT] triplets can be generated by cropping an image with perturbations using 3D color lookup tables (LUTs). In addition, we build a real-world harmonization dataset as carefully created by expert users, for evaluation and benchmarking purposes. Our results show that the proposed self-supervised method outperforms previous state-of-the-art methods in terms of reference metrics, visual quality, and subject user study. Code and dataset are available at <https://github.com/VITA-Group/SSHharmonization>.

### 1. Introduction

Image harmonization is a crucial step in image compositing that aims at adjusting (harmonizing) the appear-

ance—e.g., the color, saturation, brightness and contrast—of a foreground object to better match the background image so that the resulting composite is more realistic. For example, a subject captured under sunlight looks different from one on a cloudy day and its appearance needs to be edited when composited into a cloudy scene.

Previous approaches tackle this issue by transferring the statistic information between the foreground and background regions, including color [19, 36] and texture [31]. More recently, [33, 5, 6] train deep neural networks to address the image harmonization problem, necessitating the large-scale dataset of input-harmonized composite training pairs. However, collecting a large-scale high-quality harmonization dataset, in general, requires tedious professional expert retouching. Instead, existing methods [33, 5, 6] bypass this by selecting foreground objects in existing images, perturb their color to *simulate* an unharmonized composite, and train the network to regress the original input image, as manifested in Fig. 2 left. While these approaches [33, 5, 6] are effective to an extent, they have several limitations:

- **Limited ground truth paired data.** Collecting high-quality paired harmonization data is time-consuming and laborious. Even in the constrained case presented above, it requires an accurate mask of the foreground object in each image, shown as Fig. 2 left.
- **Background occlusion.** Due to synthesizing by naively compositing, existing methods cannot make effective use

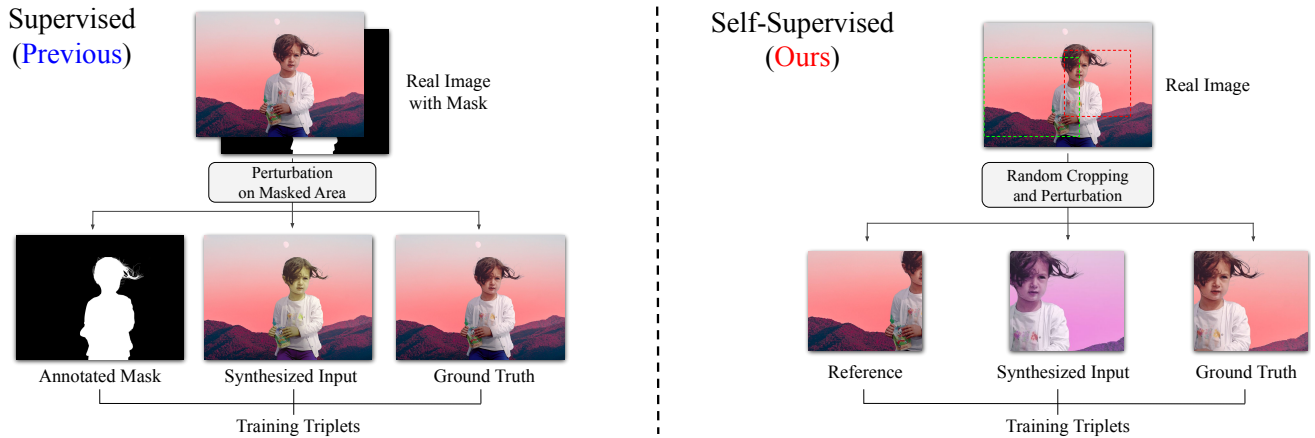


Figure 2: Comparison between previous supervised methods [5, 6, 33] and the proposed SSH. Unlike previous methods that demand annotated masks in the training process, our self-supervised framework requires NO mask during training. **The annotated mask is only needed at testing for actually making the composed image.**

of the background context for harmonization. For example, when the foreground object occupies a large part of the image, their performance commonly degrade.

- **Limited harmonization variability.** Current methods only consider simple low-dimensional color transfer functions to generate training (and even testing) data. That does not generalize well to real-world scenarios with drastically complex appearance discrepancies.

To tackle these limitations, we propose a new Self-Supervised image Harmonization framework, dubbed SSH. Different from previous approaches [33, 5, 6] which directly take the compositing image as the input (as shown in Fig. 2 left), the proposed SSH method attempts to reformulate the harmonization problem from a representation fusion perspective. The proposed framework separately extracts the “content” and “appearance” representation from foreground and background images, and then aggregates these representations to synthesize the harmonized output. Based on this form, we introduce a novel *dual data augmentation* engine to generate various synthesized data that can be directly taken as [foreground, background, pseudo GT] triplets to support self-supervised training (shown in Fig. 2 right). Meanwhile, we propose to adopt the 3D lookup table (LUT) to replace the traditional color transfer augmentation, which generates diverse visual examples on-the-fly. To this end, the proposed approach has no requirement for any foreground mask during training, and also allows us to leverage the *entire* background image to generate high-quality harmonization results.

Previous methods evaluate their performance synthesized data [5] or real-world data without ground truth [33]). In view of this gap, we build a new real-world, high-quality benchmark of harmonized composite images, that

are retouched by professional Photoshop users. This collected dataset contains 216 composite images whose foregrounds include both the human portraits and general objects, while their backgrounds cover diverse environments such as mountains, rivers, buildings, sky and more (details are described in Sec. 3.3). Experiments demonstrate that SSH significantly outperforms state-of-the-art harmonization methods on the realistic data.

We summarize our contributions below:

- We propose the first **self-supervised** harmonization framework that needs neither human-annotated mask nor professionally created images for training.
- We develop a novel *dual data augmentation* scheme, empowered by leveraging more complex 3D LUTs, to simulate more diverse and realistic training data on-the-fly.
- We collect the first-of-its-kind real-world benchmark set, containing 216 high-quality composite images that are professionally curated, to evaluate state-of-the-art image harmonization methods. Our method also significantly outperform existing approaches.

## 2. Related Works

**Image Harmonization:** Traditional image harmonization methods mainly target at better adjusting the low-level appearance statistics, such as color distribution [27, 29] and multi-scale features [14, 26, 32]. Besides traditional approaches, several recent works try to adopt learning-based method on harmonization task to better understand the context information between foreground and background images. Zhu *et al.* [42] propose to use a discriminative model that can distinguish between natural images and composite images. Tsai *et al.* firstly adopt segmentation mask

as semantic information to train an end-to-end deep learning method. Cun *et al.* [6] propose the channel-wise and spatial-wise attention mechanism that further improve the visual quality of harmonized results. The most recent approaches DoveNet [5] treats the image harmonization as a domain adaptation problem and successfully use adversarial learning to achieve notable performance. However, these methods [33, 6, 5] mainly rely on the annotated segmentation mask to synthesize training pairs. Different from those methods, our method tackles the issue of expensive human annotation thanks to the superiority of self-supervision.

**Self-Supervised Learning:** Self-supervised learning has been popular in high-level tasks by learning under pretext tasks [11, 24] or contrasting augmenting views [4, 12]. For low-level tasks, self-supervision is often implemented by self-generating synthetic data pairs, such as blur kernel in deblurring [23, 17], or bicubic interpolation in super resolution [7]. Recently, [18] applies self-supervised learning on image denoising without accessing clean reference images. However, most of these approaches meet severe challenges when tested on real-world data [3, 30, 15]. To the best of our knowledge, SSH is the first self-supervised learning method on the image harmonization task.

**Comparing Image Harmonization and Style Transfer:** Style transfer can be traced back to the seminal work of Gatys *et al.* [10], followed by many improved methods aiming at improving either transfer efficiency [13, 16] or scalability [39, 37, 9]. Nevertheless, its powerful ability on abstracting the texture feature makes it in general unsuitable for harmonizing the realistic *photography images*. A particularly relevant route to us is the photorealistic style transfer [22, 20, 38], which adds a photorealism regularization term on standard style transfer, resulting in a visually pleasing and photorealistic output. Although their goal is very similar to ours, existing works in this vein either need semantic segmentation masks to indicate specific regions [22], or require two input images to share a similar layout [20, 38] (*e.g.*, building-to-building). Different from all of them, SSH can be adopted on arbitrary photography images without needing the segmentation mask or assuming similar layouts. We compare with the state-of-the-art photorealistic style transfer method  $WCT^2$  [38] in our experiments to better illustrate the differences.

### 3. Method

An overview of our proposed self-supervised harmonization framework (SSH) is shown in Fig. 3. Our main goal is to avoid the expensive human annotation that would be otherwise required for the harmonization task and define *pseudo* ground truths that can instead serve as proxies for this task. To this end, our method utilizes a content network and a reference network to extract representation about an input image’s content and appearance respectively. The rep-

resentations are extracted in different image crops. Then these features are concatenated and fed into a fusion network which aims at reconstructing the harmonized image. In addition, processing the foreground and background images individually (rather than in a composite image [5]) allows us maximize information and avoid the background occlusion problem (Fig. 6).

We observe that the different crops from one image tend to share the same appearance (color, lighting condition, and contrast) since they are captured from the same environment (lighting, weather condition) and camera setting. Thus the different crops and their different appearance version with proper perturbations can serve as a pseudo triplet (the foreground, background and ground truth) for the content and reference networks. Inspired by this observation, we propose a dual data augmentation scheme, consisting of a content augmentation and an appearance augmentation. We will first introduce the dual data augmentation strategy, then proceed to explaining the details of our self-supervised framework and the data collection procedure.

#### 3.1. Dual Data Augmentation

The goal of dual data augmentation is to provide pseudo training triplets that contain various appearance and also mimic real testing scenarios. In each iteration, we perform both content and appearance augmentations. The content augmentation samples two different crops with some overlapping region. Meanwhile, the appearance augmentation applies multiple 3D color lookup table (LUT) for the given image to obtain its corresponding stylized images.

##### 3.1.1 Content Augmentation

To simulate real testing scenarios where the foreground and background images are totally different, we apply content augmentation in the data synthesis process. The content augmentation adopts a simple yet effective multi-cropping method to generate different crops of one original image. The cropping size ranges from a local region to a global region, and thus can mimic diverse environments which reduces the gap between synthetic data and real testing data. The bottom left image in Fig. 3 shows a typical example where we can obtain different looks from an image by this cropping method.

##### 3.1.2 Appearance Augmentation

Synthesizing data with various appearance is a common step in training image harmonization models [5, 33, 6]. Existing approaches either choose to adopt a single color transfer method [33], or simply extend it to several different color transfer approaches [5]. However, in real harmonization scenarios, the appearance mismatch between the foreground and background images can be significantly more

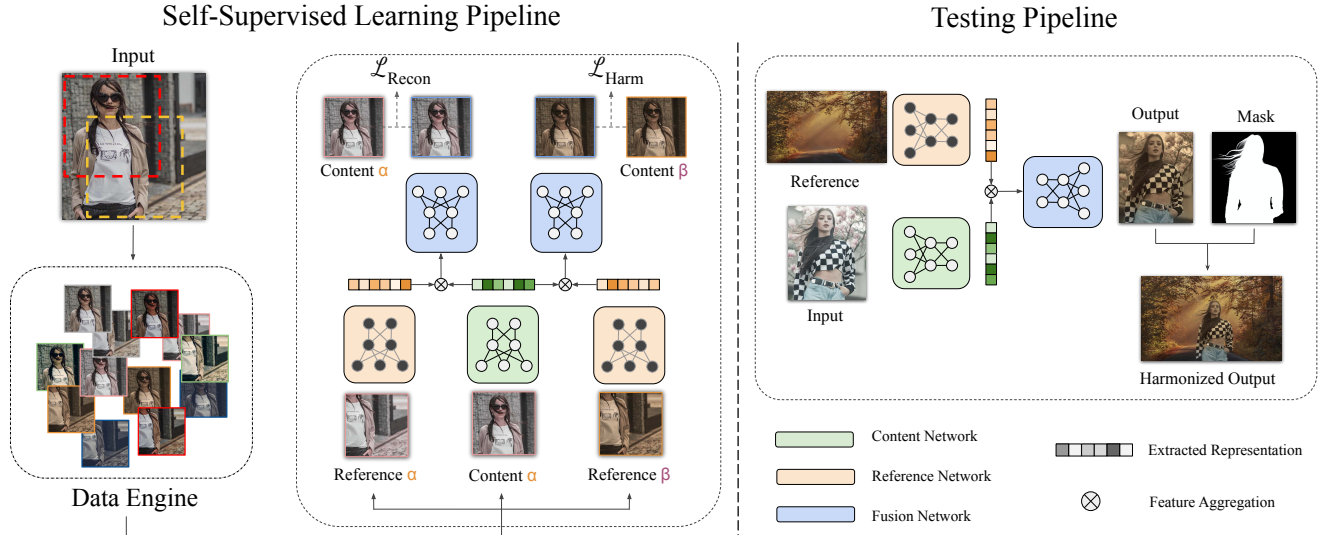


Figure 3: **Details of Training and Testing Stage.** The left figure illustrate the main pipeline of our self-supervised framework. We firstly use the proposed dual data augmentation engine to generate *Content  $\alpha$*   *$\beta$*  and *Reference  $\alpha$*   *$\beta$* , serving as the input of content network  $G_C$  and reference network  $G_S$  respectively ( $\alpha$  and  $\beta$  represents two different appearances after applying different 3D LUTs). After that the training pipeline learns to synthesize content  $\beta$  from content  $\alpha$  when the reference  $\beta$  is given, and also learns to reconstruct content  $\alpha$  itself when reference  $\alpha$  is given. The translation and reconstruction process result in both the harmonization loss and the reconstruction loss. The right figure describe our testing stage. Noted that, **the human annotated mask is needed only in the testing stage for necessary composition.**

complex and include contrast, brightness, and saturation differences. To tackle this issue, we propose to use a 3D color lookup table as the basic transformation approach.

A **3D Color Lookup Table (LUT)** maps one color space to another, and is widely used in film post-production industry. It is essentially a 3D-to-3D mapping that can transform any RGB color of an input image to any other RGB color. It can also represents functions like contrast enhancement where the tonal range of the input image is manipulated.

A LUT has a number of advantages including: 1) Unlike simple color transfer functions, it can represent complex appearance adjustments; for example, the bottom left image in Fig. 3 shows that the LUT can provide a non-linear transformation where the appearance of different parts of the image are transformed differently (e.g., the skin can be transformed to red color while the T-shirt remains white color), 2) Given one input image, there exist hundreds of LUTs that can be applied to generate its stylized versions and dramatically enrich the training data, and 3) LUT processing is real-time so that it can be applied as an on-the-fly data augmentation strategy.

### 3.2. Self-Supervised Framework

The proposed self-supervised framework SSH takes the foreground as content image  $C$  and the background as reference image  $R$ . We adopt a reference network  $G_r$  to capture

the appearance representation (color, brightness, and contrast etc.) from reference image  $R$ , and a content network  $G_c$  to capture the content representation (structure, texture etc.) from content input  $C$ . Then the fusion network  $F$  aggregates the appearance and content representation and learns to synthesize the output  $C'$ , so that the output  $C'$  matches the appearance of reference image  $R$  and preserves the content of content image  $C$ . We formulate the process as following:

$$C' = F(G_c(C), G_r(R)) \quad (1)$$

Using the proposed dual data augmentation, we can generate pseudo triplets [foreground, background, and ground truth] from one image, and adopt it for training SSH. As shown in the left part of Fig. 3, we generate two image with their appearance perturbed by two different 3D LUTs separately (denoted by  $\alpha$  and  $\beta$ ). Therefore, content crops (denoted as  $C_\alpha, C_\beta$ ) and reference crops (denoted as  $R_\alpha, R_\beta$ ) corresponding to them can be obtained.

Here  $C_\alpha$  contain similar appearance information with  $R_\alpha$  and the same content information as  $C_\beta$ . Then the network is expect to map  $C_\alpha$  to  $C_\beta$  when taking  $R_\beta$  as the reference, and reconstruct the  $C_\alpha$  to itself when taking  $R_\alpha$  as the reference. This special design simulate real testing scenarios where the foreground and background does not share the same content while the output is expected to have same






Figure 4: **The Main Pipeline of Dataset Collection.** The detailed pipeline of how the composite image is retouched by professional user is shown above. It includes both the object annotation and appearance matching (brightness, color/saturation, contrast adjustment, and local polish).

the appearance as the background and the same content as the foreground. The mapping process and reconstruction process use the following harmonization loss  $L_{harm}$  and reconstruction loss  $L_{harm}$ :

$$C'_\beta = F(G_c(C_\alpha), G_s(R_\beta)) \quad (2)$$

$$C'_\alpha = F(G_c(C_\alpha), G_s(R_\alpha)) \quad (3)$$

$$L_{harm} = \|C'_\beta - C_\beta\|^2 \quad (4)$$

$$L_{recon} = \|C'_\alpha - C_\alpha\|^2 \quad (5)$$

Since the reference network  $G_r$  needs to capture the appearance information from given image, it is expected to extract similar representation when receiving different crops with the same appearance. Meanwhile, the content network  $G_c$  is expected to capture the same feature given the same crops with different appearance. Thus we design another disentanglement loss  $L_{dis}$  formulated as following:

$$L_{dis} = \|G_c(C_\alpha) - G_c(C_\beta)\|^2 + \|G_s(C_\alpha) - G_s(R_\alpha)\|^2 \quad (6)$$

Combining with harmonization loss and reconstruction loss defined above, the overall loss function for training SSH is thus written as:

$$Loss = L_{harm} + w_1 * L_{recon} + w_2 * L_{dis} \quad (7)$$

Where  $w_1$  and  $w_2$  is set to 0.4 and 0.05 in experiments. Therefore, the proposed SSH can translate a image to a synthesized one that the synthesized output can match the appearance of another images. As shown in the right part of Fig. 3, only in the testing stage, we adopt the object mask to composite synthesized output and background/reference image to generate the harmonized output

### 3.3. Dataset Preparation

#### 3.3.1 Training Data and LUT Collection

Since SSH requires only self-supervision (rather than the labeled object masks or harmonization ground truth from professionals user required by previous methods [33, 6, 5]), we are able to collect a larger-scale unlabeled training set of images with diverse semantics, environmental lighting conditions and contents. Our unlabeled training set consists of 81917 images from several datasets [21, 2, 41] and the Internet<sup>12</sup>, containing mountain, river, sky, general object, human portrait under diverse lighting conditions. In addition, we also collect 100 3D color lookup tables from the Internet. We randomly select two LUTs from the collection in each training iteration, causing  $100 \times 100$  possible combinations that enable diversity during training.

#### 3.3.2 Real-World Image Harmonization Benchmark

Although there exist several testing datasets [33, 5] for evaluating the harmonization performance, all of them are generated by perturbing the foreground objects with simple color-transfer methods [29, 35, 8, 28]. Moreover, these methods [5, 33] apply the same data synthesis procedure for both training stage and testing stage. For learning-based approaches, the evaluation results will be inevitably biased and cannot faithfully probe their real-world generalization.

To bridge the gap between current evaluation protocols with real image harmonization demands, we propose a **Real-world HarMonization** dataset for evaluation, named **RealHM**. However, collecting a well-annotated real-world testing dataset is not a trivial task, as several elements need to be adjusted together and the best result is beyond one image. In particular, at least three main steps are required to generate high-quality harmonized composition, as follows: 1) A challenging compositing pair need to be picked where directly compositing these two images will cause severe visual disharmony. 2) A high-quality mask to cut out the foreground objects is needed, where the hard boundary and soft boundary (hair/fur region) are treated differently. 3) Matching the appearance (color, brightness, saturation, contrast) of the foreground with the background by using PhotoShop tools, where some particular local region needs further adjustment. A typical example is shown in Fig. 4.

Finally, we collect 216 high-quality, high-resolution foreground/background pairs with corresponding harmonized outputs, where the foregrounds include both the human portraits and general objects and the backgrounds cover diverse environments such as mountain, river, buildings, sky and else.

<sup>1</sup><https://unsplash.com>

<sup>2</sup><https://www.flickr.com/>



Figure 5: **Comparison with the State-of-the-art methods.** The first column “DC” represents the direct compositing results of foreground and background images, second column “GT” represents the ground truth harmonized results annotated by human expert users. The rest five column shows the output of comparing methods and our proposed SSH methods. More visual results will be shown in the supplementary. Best viewed zoomed in.

## 4. Experiments

### 4.1. Implementation

SSH is first trained from scratch for 70 epochs with the learning rate of  $2e-4$ , followed by another 30 epochs with the learning rate linearly decayed to 0. We adopt a scale jittering range of [256, 320] and then crop a  $224 \times 224$  image in the training stage. We use the Adam optimizer and the batch size is set to be 64. The whole training process takes 20 hours on 8 Nvidia 2080Ti GPUs. We implement it with PyTorch [25] framework. We use  $256 \times 256$  resolution for full-reference metrics evaluation (PSNR, MSE, SSIM, and LPIPS) and use the original resolution when showing the visual comparison for better view. We follow [1] to generate high-resolution visual examples using color mapping function.

### 4.2. Comparing with the State-of-the-art methods

In this section we compare the performance of our proposed method SSH with the current state-of-the-art methods. We conduct both the quantitative and qualitative evaluation, including visual quality comparison, referenced image quality assessment (IQA) and human subjective test.

#### 4.2.1 Visual Quality Comparison

We firstly compare our SSH framework with current state-of-the-art methods in terms of visual quality, shown in Fig. 5. The first column shows the results by directly compositing foreground and background images. The second column represents the annotated harmonized results retouched by the human expert user with professional editing skill. The third to seventh column shows the output generated by:  $WCT^2$  [38], DIH[33],  $S^2AM$  [6], DoveNet[5], and our proposed SSH methods. From the results in Fig. 5, photorealistic style transfer method [38] generates the worst results due to the foreground and background do not share similar layout, violating the prerequisite of its suitable scenarios. The results of DIH show subtle adjustment compared to the inputs since they only use one color transfer method to synthesize training data. The results of  $S^2AM$  and DoveNet either show incorrect color or disharmony contrast. In contrast, SSH successfully not only learns to extract correct color but also avoids contrast/brightness mismatching. More results will be shown in the supplementary.

#### 4.2.2 Full-referenced metrics

We next evaluate the performance of these methods on our **RealHM** benchmark. Specifically, the human annotated harmonized output is set to be the ground truth label in each example, and the direct compositing results is

Methods	PSNR $\uparrow$	MSE $\downarrow$	SSIM $\uparrow$	LPIPS $\downarrow$
DC	25.91	409.54	0.9385	0.049
WCT <sup>2</sup> [38]	22.13	446.85	0.8559	0.096
DIH [33]	23.96	433.52	0.8661	0.082
S <sup>2</sup> AM [6]	26.77	283.27	0.9366	0.053
DoveNet [5]	27.41	214.11	0.9416	0.044
SSH	<b>27.91</b>	<b>206.85</b>	<b>0.9479</b>	<b>0.039</b>

Table 1: **Comparing with the State-of-the-art methods.** We compare the the proposed method SSH with others with referenced-metric including PSNR, MSE, SSIM, and perceptual metrics. Here  $\downarrow$  represents the lower the better and  $\uparrow$  represents the higher the better. Our method outperforms previous methods under all these four metrics.

Methods	Score $\uparrow$
WCT <sup>2</sup> [38]	0.821
DIH [33]	1.201
S <sup>2</sup> AM [6]	1.744
DoveNet [5]	1.256
SSH	<b>2.295</b>

Table 2: **Human Subjects Evaluation.** The higher score indicates the better result.

Methods	PSNR $\uparrow$	MSE $\downarrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Single-Crops	24.80	376.74	0.9203	0.069
Multi-Crops (ours)	<b>27.91</b>	<b>206.85</b>	<b>0.9479</b>	<b>0.039</b>

Table 3: **Evaluation of Content Augmentation.** The first row shows the results of SSH with single-cropping augmentation and second row shows the results of SSH with random multi-cropping augmentation (the proposed one).

considered as a baseline score here. Following previous work [6], we adopt referenced metrics, including PSNR, Mean Squared Error (MSE), SSIM [34], and LPIPS [40]. The results are shown in Table 1. To be more detailed, WCT<sup>2</sup> [38] shows the worst performance since its applicable scenarios requires the two images to share the similar layout (*e.g.*, building-to-building). Besides, S<sup>2</sup>AM [6] outperforms DIH [33] due to its dual attention module. DoveNet [5] is slightly better than S<sup>2</sup>AM thanks to its domain verification discriminator. However, these learning-based method all requires human annotated object mask. Different from those approaches, SSH reaches the best performance among all these four metrics without the necessity of any labels, demonstrating the effectiveness of the proposed self-supervised framework. We also include a typical example in Fig. 6 to show that the back ground occlusion issue can be well addressed by the proposed SSH framework.

### 4.2.3 Human Subjects Evaluation

We conduct a human subjective review to compare the performance of SSH with other methods. We randomly select

Methods	PSNR $\uparrow$	MSE $\downarrow$	SSIM $\uparrow$	LPIPS $\downarrow$
color transfer	27.01	311.96	0.9451	0.043
Saturation	26.68	337.11	0.9432	.0460
3D LUT (ours)	<b>27.91</b>	<b>206.85</b>	<b>0.9479</b>	<b>0.039</b>

Table 4: **Evaluation of Appearance Augmentation.** We evaluate different ‘‘appearance’’ augmentation strategy including color transfer, random saturation change, and our proposed 3D color lookup table (LUT) augmentation.

Methods	PSNR $\uparrow$	MSE $\downarrow$	SSIM $\uparrow$	LPIPS $\downarrow$
w/o recon loss.	24.05	475.45	0.9095	0.0780
w/o disentangle loss	24.94	372.21	0.9217	0.0643
SSH	<b>27.91</b>	<b>206.85</b>	<b>0.9479</b>	<b>0.039</b>

Table 5: **Evaluation of Loss Function Design Choices.** The first row shows the results generated by SSH without reconstruction loss and the second row shows the results without disentanglement loss

15 foreground and background pairs from the RealHM benchmark. Each image is first processed by five methods WCT<sup>2</sup> [38], DIH [33], S<sup>2</sup>AM[6], DoveNet[5], and SSH), and then displayed on a screen for comparison. The human annotated ground truth is also provided as a reference. We then ask 50 subjects to independently score the visual quality, considering the following factors: 1) whether the image contains color/saturation disharmony; 2) whether the foreground and the background has different illumination; and 3) whether the composite images show texture distortions/artifacts. The score of visual quality ranges from 0 to 4 (worst to best quality). As shown in the Table 2, the photorealistic style transfer method WCT<sup>2</sup> [38] shows the worst performance since it is not suitable for image harmonization task, which is consistent with the observation from full reference evaluation. The learning-based image harmonization methods [33, 6, 5] show comparably good performance, while the proposed method SSH achieves best score.





Figure 6: **Background occlusion problem.** “BG” and “DC” represent the background image and the direct compositing result, respectively. Since DoveNet [5] only takes the direct compositing image as the input, it fails to capture the true appearance of the background when the foreground object is too large, while SSH successfully generates a harmonized output.

### 4.3. Ablation Study

#### 4.3.1 Effectiveness of Dual Data Augmentation

As the superiority of SSN is built upon the strong dual data augmentation, we evaluate each component starting from the content augmentation and appearance augmentation. The content augmentation act as a crucial role in our framework as it generate two different crops of the same image, simulating the real scenarios where the foreground and background images are totally different. We ablate it by replacing the multi-cropping method with a single-cropping method and then make the content and reference network receive the same crop as the input. As shown in Table 3, we observe that without adopting multi-cropping strategy, the performance largely degrades due to the pseudo label is exactly same as the input in the training stage, making the model easily minimizing the loss without actually learning the representation of the appearance.

Furthermore, we study the effectiveness of 3D LUT by comparing it with other appearance augmentation such as color transfer and random saturation adjustment. As shown in Table 4, the proposed 3D LUT augmentation strategy outperforms both the color transfer and random saturation change. This is because 3D LUT provides a stronger appearance change thanks to its diverse color mapping way. Also, 3D LUT enable local appearance change instead of a simple global translation.

#### 4.3.2 Evaluation of Loss function

To study the effectiveness of reconstruction loss and disentanglement loss, we conducted the ablation experiments by removing them separately. As shown in Table 5, either removing the reconstruction loss or the disentanglement loss will cause performance degradation, demonstrating the effectiveness of these loss. We also find that the disentanglement loss can help stabilize the performance our method, the visual examples will be shown in the supplementary.




Figure 7: **Locality-aware harmonization.** “DC” denotes direct compositing results.

### 4.4. Locality-aware Harmonization

Since the proposed self-supervised framework takes the full reference images as the inputs without annotated mask, the missing locality information brought by the annotated mask becomes a probable concern of current approach. In practice, we find that this issue can be well addressed by a proper cropping strategy during the inference stage. We take a spatially-variant colorful image as a typical background example, shown in Fig. 7. The background contains two different appearance in different location, top right and bottom left. However, the appearance of the harmonized output is expected to be close to the region where the foreground object is placed. To avoid the harmonized output being affected by the misleading environment, we show that simply adopt the cropping method can well preserve the locality information. By separately considering the top right and bottom left cropping box as the background, the proposed methods successfully extract correct “appearance” representation and generate reasonable and pleasing visual results, demonstrating the effectiveness of SSH in this challenging scenario.

## 5. Conclusion

We propose a self-supervised framework for image harmonization, named SSH. The proposed method does not require any human annotated labels in the training phase thus reduce the tedious effort of collecting a large-scale high-quality human annotated dataset retouched by professional users. Furthermore, we propose a dual data augmentation which include both the content data augmentation and appearance data augmentation to not only provide stable pseudo labels but also enrich the diversity of training data. Besides, we built a real harmonization benchmark that fills the gap in real testing scenarios. Our method outperforms all previous methods in a variety of metrics.



## References

- [1] Mahmoud Afifi and Michael S Brown. Deep white-balance editing. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 1397–1406, 2020. 6
- [2] Vladimir Bychkovsky, Sylvain Paris, Eric Chan, and Frédo Durand. Learning photographic global tonal adjustment with a database of input/output image pairs. In *CVPR 2011*, pages 97–104. IEEE, 2011. 5
- [3] Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3086–3095, 2019. 3
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 3
- [5] Wenyang Cong, Jianfu Zhang, Li Niu, Liu Liu, Zhixin Ling, Weiyuan Li, and Liqing Zhang. Dovenet: Deep image harmonization via domain verification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8394–8403, 2020. 1, 2, 3, 5, 6, 7, 8
- [6] Xiaodong Cun and Chi-Man Pun. Improving the harmony of the composite image by spatial-separated attention module. *IEEE Transactions on Image Processing*, 29:4759–4771, 2020. 1, 2, 3, 5, 6, 7
- [7] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. 3
- [8] Ulrich Fecker, Marcus Barkowsky, and André Kaup. Histogram-based prefiltering for luminance and chrominance compensation of multiview video. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(9):1258–1267, 2008. 5
- [9] Wei Gao, Yijun Li, Yihang Yin, and Ming-Hsuan Yang. Fast video multi-style transfer. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, March 2020. 3
- [10] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. 3
- [11] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. 3
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 3
- [13] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. 3
- [14] Jiaya Jia, Jian Sun, Chi-Keung Tang, and Heung-Yeung Shum. Drag-and-drop pasting. *ACM Transactions on graphics (TOG)*, 25(3):631–637, 2006. 2
- [15] Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. Enlightengan: Deep light enhancement without paired supervision. *IEEE Transactions on Image Processing*, 30:2340–2349, 2021. 3
- [16] 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. 3
- [17] Orest Kupyn, Tetiana Martyniuk, Junru Wu, and Zhangyang Wang. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8878–8887, 2019. 3
- [18] Samuli Laine, Tero Karras, Jaakko Lehtinen, and Timo Aila. High-quality self-supervised deep image denoising. In *Advances in Neural Information Processing Systems*, pages 6970–6980, 2019. 3
- [19] Jean-Francois Lalonde and Alexei A Efros. Using color compatibility for assessing image realism. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007. 1
- [20] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 453–468, 2018. 3
- [21] 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. 5
- [22] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4990–4998, 2017. 3
- [23] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3883–3891, 2017. 3
- [24] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016. 3
- [25] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019. 6
- [26] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. In *ACM SIGGRAPH 2003 Papers*, pages 313–318. 2003. 2

- [27] Francois Pitie, Anil C Kokaram, and Rozenn Dahyot. N-dimensional probability density function transfer and its application to color transfer. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1434–1439. IEEE, 2005. 2
- [28] François Pitié, Anil C Kokaram, and Rozenn Dahyot. Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding*, 107(1-2):123–137, 2007. 5
- [29] Erik Reinhard, Michael Adhikhmin, Bruce Gooch, and Peter Shirley. Color transfer between images. *IEEE Computer graphics and applications*, 21(5):34–41, 2001. 2, 5
- [30] Jaesung Rim et al. *Real-World Blur Dataset for Learning and Benchmarking Deblurring Algorithms*. PhD thesis, DG-IST, 2020. 3
- [31] Kalyan Sunkavalli, Micah K Johnson, Wojciech Matusik, and Hanspeter Pfister. Multi-scale image harmonization. *ACM Transactions on Graphics (TOG)*, 29(4):1–10, 2010. 1
- [32] Michael W Tao, Micah K Johnson, and Sylvain Paris. Error-tolerant image compositing. In *European Conference on Computer Vision*, pages 31–44. Springer, 2010. 2
- [33] Yi-Hsuan Tsai, Xiaohui Shen, Zhe Lin, Kalyan Sunkavalli, Xin Lu, and Ming-Hsuan Yang. Deep image harmonization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3789–3797, 2017. 1, 2, 3, 5, 6, 7
- [34] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 7
- [35] Xuezhong Xiao and Lizhuang Ma. Color transfer in correlated color space. In *Proceedings of the 2006 ACM international conference on Virtual reality continuum and its applications*, pages 305–309, 2006. 5
- [36] Su Xue, Aseem Agarwala, Julie Dorsey, and Holly Rushmeier. Understanding and improving the realism of image composites. *ACM Transactions on graphics (TOG)*, 31(4):1–10, 2012. 1
- [37] Shuai Yang, Zhangyang Wang, Zhaowen Wang, Ning Xu, Jiaying Liu, and Zongming Guo. Controllable artistic text style transfer via shape-matching gan. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4442–4451, 2019. 3
- [38] 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. 1, 3, 6, 7
- [39] Hang Zhang and Kristin Dana. Multi-style generative network for real-time transfer. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 3
- [40] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 7
- [41] Hao Zhou, Torsten Sattler, and David W Jacobs. Evaluating local features for day-night matching. In *European Conference on Computer Vision*, pages 724–736. Springer, 2016. 5
- [42] Jun-Yan Zhu, Philipp Krahenbuhl, Eli Shechtman, and Alexei A Efros. Learning a discriminative model for the perception of realism in composite images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3943–3951, 2015. 2