

Deep Image Harmonization by Bridging the Reality Gap

Junyan Cao[#]

Joy_C1@sjtu.edu.cn

Wenyan Cong[#]

wycong@utexas.edu

Li Niu^{*}

ustcnewly@sjtu.edu.cn

Jianfu Zhang

c.sis@sjtu.edu.cn

Liqing Zhang

zhang-lq@cs.sjtu.edu.cn

MoE Key Lab of Artificial Intelligence

Shanghai Jiao Tong University

Shanghai, China

Abstract

Image harmonization has been significantly advanced with large-scale harmonization dataset. However, the current way to build dataset is still labor-intensive, which adversely affects the extendability of dataset. To address this problem, we propose to construct rendered harmonization dataset with fewer human efforts to augment the existing real-world dataset. To leverage both real-world images and rendered images, we propose a cross-domain harmonization network to bridge the domain gap between two domains. Moreover, we also employ well-designed style classifiers and losses to facilitate cross-domain knowledge transfer. Extensive experiments demonstrate the potential of using rendered images for image harmonization and the effectiveness of our proposed network.

1 Introduction

Image composition [13] combines foreground from one image and background from another image into one composite image. However, due to the appearance discrepancy caused by different capture conditions (*e.g.* weather, season, time of the day) of foreground and background, the quality of composite image might be greatly degraded. Image harmonization aims to diminish the discrepancy by adjusting the appearance of the composite foreground according to the background. Recently, many deep learning based harmonization methods [4, 6, 15, 16, 28, 38, 45] have achieved promising results.

Deep learning based methods require a large-scale training set that contains pairs of composite images and corresponding harmonized results. However, it is nontrivial and infeasible to collect abundant training pairs by manually harmonizing composite images. So inversely, [15] and [6] take real images as the ground-truth and adjust foreground appearance to generate synthetic pairs of composite images and ground-truth images. Though the above data

[#]Equal contribution. ^{*}Corresponding author.

© 2022. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

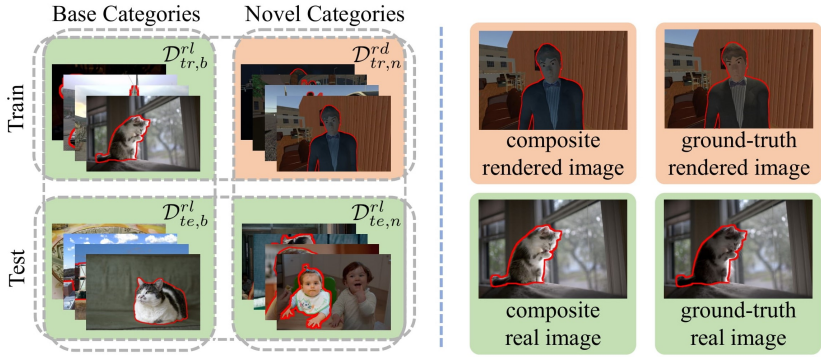


Figure 1: Take “human” as an example novel category. Left sub-figure: illustration of train- ing and test images from base categories and novel categories. Right sub-figure: image pairs from rendered image domain and real image domain. Images in orange (*resp.*, green) are from rendered (*resp.*, real) image domain. The foregrounds are outlined in red.

acquisition is feasible in implementation, it still requires manual labor for foreground seg- mentation and unqualified composite image filtering [4], which hinders the extendibility of harmonization dataset. For example, the foreground category of test composite images may be out of the scope of training set and we observe that cross-category harmonization suf- fers from a performance drop (see Section 5.2). However, extending the dataset to include novel foreground categories requires heavy human efforts. In the remainder of this paper, the category of a composite image means the category of its foreground to be harmonized.

In this work, we propose to construct image harmonization dataset using 3D rendering techniques, which can escape from heavy manual efforts and supplement existing real-world harmonization dataset. Specifically, we utilize 3D rendering software *Unity3D* to generate rendered 2D scenes, and we leverage its plugin *UniStorm* to generate a group of images with different capture conditions for each 2D scene. Since it is effortless to obtain accurate foreground masks in *Unity3D*, we could exchange the foregrounds within images from the same 2D scene, then obtain abundant pairs of composite rendered images and ground-truth rendered images.

Suppose that a real-world (“real” for short) harmonization dataset has real images from base categories, we can enrich this dataset using rendered images from novel categories. Then, as in Figure 1, the image harmonization model could be trained on the combination of rendered images and real images, and tested on real images from both base and novel categories. This is a practical setting when our goal is enhancing the performance on novel categories which are out of the scope of real dataset.

In this work, we contribute a Rendered Image Harmonization dataset *RdHarmony*, which totally contains 6 categories (“human”, “bottle”, “cake”, “motorcycle”, “cow”, and “couch”). In the main paper, we mainly take “human” as an example novel category because human harmonization is rather challenging and of paramount interest for many applications (*e.g.*, augmented reality [47, 49]). The experiments using the other five categories as novel cate- gories are left to the Supplementary.

However, models trained on rendered images (source domain) would suffer from a per- formance drop when applied to real images (target domain) (see Section 5.2), owing to considerably different data distributions between two domains, which is typically known as

domain gap [44, 52]. Hence, we propose a novel Cross-domain **h**armonization Network named **CharmNet** to align two domains when training with a combination of real images and rendered images. Our CharmNet has three stages. The *domain-specific encoding stage* projects data from both domains to the same shared domain, where the projected features of two domains are further aligned with an adversarial loss. Then, the *domain-invariant encoding-decoding stage* encourages information sharing between real and rendered images. The *domain-specific decoding stage* projects the domain-invariant features back to different domains by enforcing the features to reconstruct the ground-truth images in each domain.

Moreover, by naming different capture conditions as different “styles”, we can easily obtain the style labels of rendered images using UniStorm. Therefore, we can explicitly exploit such information to facilitate cross-domain knowledge transfer. We assume that the input features before the second stage are inharmonious and the output features after the second stage are harmonious, so that the harmonization process is shared across two domains and more useful information can be transferred from rendered image domain to real image domain. For rendered images, their style labels before and after harmonization are readily available, which could be used to supervise style classification. For real images, their style labels before and after harmonization are unknown, but we assume that the style distribution should become more concentrated after harmonization because the foreground style is adapted to the background style. Hence, we propose a novel Style Aggregation (SA) loss to enforce the style distribution to be more concentrated after harmonization. To verify the effectiveness of our proposed network, we conduct comprehensive experiments on real dataset iHarmony4 and our contributed rendered dataset.

Our contributions are summarized as follows: 1) We investigate on the cross-domain and cross-category issues in image harmonization; 2) We have contributed and released the first large-scale rendered image harmonization dataset RdHarmony; 3) We propose the first cross-domain image harmonization network CharmNet with novel network architecture and style aggregation loss; 4) Extensive experiments on real and rendered datasets demonstrate the potential of using rendered images for image harmonization.

2 Related Works

2.1 Image Harmonization

Image harmonization has drawn increasing attention in the field of computer vision. Early works [21, 25, 35, 36, 37, 39, 41, 43, 50] mainly focused on manipulating the low-level image statistics, such as matching color distributions [36, 37], applying gradient-domain compositing [21, 35, 43], and mapping multi-scale statistics [41]. Recently, deep learning based methods [4, 6, 12, 22, 45] spring up rapidly. [38, 45] both leveraged auxiliary semantic features. [9] released the first large-scale image harmonization dataset iHarmony4 and introduced a domain verification discriminator pulling close the foreground domain and background domain. [2, 17] explored various attention mechanisms. [12, 15] harmonized composite images by harmonizing reflectance and illumination separately. [28] proposed to explicitly formulates background style and adaptively applied it to the foreground, which is extended in [16] by extracting the local background style relevant to the foreground. [8] proposed to combine pixel-to-pixel transformation and color-to-color transformation in a unified framework. Different from existing methods only using real images, we are the first to use both real images and rendered images to train a cross-domain image harmonization model.

2.2 Domain Adaptation

Domain adaptation strives to reduce the distribution mismatch and make the model trained on source domain generalize well to target domain. In our task, real images and rendered images are treated as two domains. Domain adaptation has been studied widely in classification [12, 13, 14], detection [15, 16, 17, 18], segmentation [19, 20, 21, 22], person re-identification [8, 9, 23], and pose estimation [24, 25]. In the above applications, the label/output spaces (*e.g.*, class label, segmentation mask) are the same between two domains. However, in image harmonization, the input and output spaces are the same in each domain but different across domains. Therefore, previous domain adaptation methods cannot be directly applied to our task. To the best of our knowledge, we are the first to address such a challenging task in the realm of domain adaptation.

Another possible domain adaptation approach is to translate images from the source domain to the target domain using Image-to-Image (I2I) translation methods [6, 26, 27, 28, 29, 30, 31, 32, 33], so that the translated images can be used for any downstream tasks. Nevertheless, during experiments, we observe that I2I translation methods have poor performance when translating rendered images to real images because they severely distort the illumination statistics. Therefore, I2I translation methods are ill-suited for our task.

3 Dataset Construction

We construct a rendered image harmonization dataset RdHarmony that contains rendered images with foregrounds of “human” category and other 5 object categories, *i.e.*, “bottle”, “cake”, “motorcycle”, “cow”, and “couch”.

The whole process can be divided into two steps. The first step is *ground-truth rendered image generation*. We place various 3D characters of each novel category in different 3D scenes using Unity3D, and vary the camera viewpoints to generate 2D scenes. Meanwhile, we utilize UniStorm to control the weather and time. Note that UniStorm is a weather system plugin for Unity3D, which can simulate real-world natural phenomena by setting multiple attributes (*e.g.*, light intensity, illumination color). With UniStorm, each 2D scene could produce a group of images with different capture conditions (*i.e.*, weather, time), which are referred to as different styles. In this work, we focus on natural lighting conditions because most images in iHarmony4 [9] are captured in natural lighting conditions. We select 3 representative weathers (Clear, Partly Cloudy, Cloudy) from UniStorm with pre-defined values. And we split time-of-the-day into 4 distinct phases (sunrise&sunset, noon, night, other-times) by sampling every 5 minutes from 5:00 am to 21:00 pm. Based on the combinations of weather and time-of-the-day, we define 10 representative styles, including the night style as well as styles of Clear/Partly Cloudy/Cloudy weather at sunrise&sunset/noon/other-times. *Note that each style is not a discrete style, but covers a range of illumination intensity and direction.*

The second step is *composite rendered image generation*. For each 2D scene, we treat one 3D character as foreground and obtain the foreground mask effortlessly using Unity3D. By exchanging the foregrounds of ground-truth rendered images in the same 2D scene with each other, we can produce composite rendered images with mixed styles. We define the style label of the ground-truth rendered image as a 10-dim one-hot vector. Inspired by [34], we define a soft style label for the composite rendered image, which mixes the foreground and background style labels according to the foreground and background area ratio.

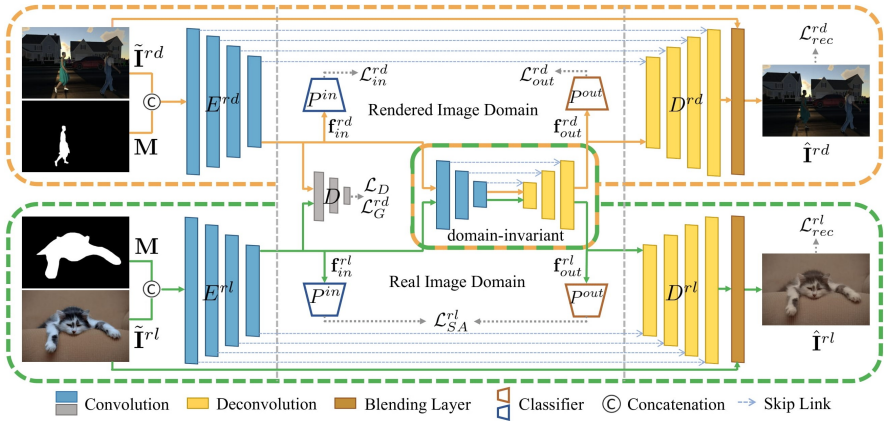


Figure 2: The network architecture of CharmNet, which contains two three-stage generators, a domain discriminator D , and two style classifiers P^{in} and P^{out} . The data flow of real (*resp.*, rendered) images is marked with green (*resp.*, orange) lines.

For “human” category, we generate 15,000 ground-truth rendered images with 1,500 different 2D scenes, and produce 135,000 rendered training pairs. For each object category (“bottle”, “cake”, “motorcycle”, “cow”, and “couch”), we create 9,000 rendered training pairs. Totally, our RdHarmony has 180,000 pairs of composite rendered images and ground-truth rendered images. More details about RdHarmony are left to the Supplementary.

4 Our Method

In this work, we focus on training an image harmonization network using a combination of rendered image pairs and real image pairs. In the training stage, we have access to rendered (rd) image pairs $\{\tilde{\mathbf{I}}^{rd}, \mathbf{I}^{rd}\}$ from novel categories and real (rl) image pairs $\{\tilde{\mathbf{I}}^{rl}, \mathbf{I}^{rl}\}$ from base categories, where $\tilde{\mathbf{I}}^{rd}$ (*resp.*, $\tilde{\mathbf{I}}^{rl}$) denotes the composite rendered (*resp.*, real) image and \mathbf{I}^{rd} (*resp.*, \mathbf{I}^{rl}) denotes the corresponding ground-truth rendered (*resp.*, real) image. We also have access to the binary foreground mask \mathbf{M} , which is not distinguished between domains. Given a composite rendered (*resp.*, real) image $\tilde{\mathbf{I}}^{rd}$ (*resp.*, $\tilde{\mathbf{I}}^{rl}$) and mask \mathbf{M} , the goal of image harmonization is to reconstruct \mathbf{I}^{rd} (*resp.*, \mathbf{I}^{rl}) with the harmonized rendered (*resp.*, real) image $\hat{\mathbf{I}}^{rd}$ (*resp.*, $\hat{\mathbf{I}}^{rl}$). In the testing stage, we aim to harmonize composite real images from both base and novel categories.

4.1 Cross-domain Harmonization Network

As stated in Section 2.2, after harmonization, the composite rendered (*resp.*, real) image is converted to harmonized rendered (*resp.*, real) image. They share neither input space nor output space, so we design a three-stage network for cross-domain harmonization, including a domain-specific encoding stage, a domain-invariant encoding-decoding stage, and a domain-specific decoding stage.

The pipeline of our network is shown in Figure 2. The generator G adopts the improved backbone iDIH proposed in [63], where an image blending layer is added to the UNet-like

architecture. Since both input space and output space are different across two domains, we split the first (*resp.*, last) 4 layers in the encoder (*resp.*, decoder) into E^{rd} (*resp.*, D^{rd}) and E^{rl} (*resp.*, D^{rl}), where E^{rd} (*resp.*, E^{rl}) and D^{rd} (*resp.*, D^{rl}) are the domain-specific encoder and decoder for the rendered (*resp.*, real) image domain.

After the domain-specific encoding stage, we assume that images from different domains are projected into the same feature space. Inspired by [14, 15], we employ a domain discriminator D with adversarial loss [46] to further align two domains, which is given by

$$\mathcal{L}_D = \mathbb{E}[\max(0, 1 - D(\mathbf{f}_{in}^{rl}))] + \mathbb{E}[\max(0, 1 + D(\mathbf{f}_{in}^{rd}))], \quad \mathcal{L}_G^{rd} = -\mathbb{E}[D(\mathbf{f}_{in}^{rd})], \quad (1)$$

where \mathbf{f}_{in}^{rd} (*resp.*, \mathbf{f}_{in}^{rl}) denotes the feature extracted by the domain-specific encoder E^{rd} (*resp.*, E^{rl}) in the rendered (*resp.*, real) image domain. By playing a minimax game, the discriminator D encourages the encoders E^{rd} and E^{rl} to project rendered and real images to the same feature space.

Then, through the domain-invariant encoder-decoder, the knowledge of harmonization is transferred from rendered image domain to real image domain. Finally, the third stage projects the domain-invariant features back to the input domain, by enforcing the harmonized output to approach the ground-truth image in each domain:

$$\mathcal{L}_{rec}^{rd} = \|\hat{\mathbf{I}}^{rd} - \mathbf{I}^{rd}\|_1, \quad \mathcal{L}_{rec}^{rl} = \|\hat{\mathbf{I}}^{rl} - \mathbf{I}^{rl}\|_1. \quad (2)$$

4.2 Style Classifiers and Style Aggregation Loss

When a composite rendered image passes through our three-stage network, we assume that the input feature before the second stage is inharmonious with mixed styles, while the output feature after the second stage is harmonious with a unified style. With this assumption, the main harmonization process is accomplished in the second stage, so that more useful knowledge can be transferred from rendered images to real images. Therefore, we employ two auxiliary style classifiers positioned before and after the second stage.

Because rendered images are associated with ground-truth style labels, we adopt standard cross-entropy classification losses for the inharmonious input feature \mathbf{f}_{in}^{rd} and harmonious output feature \mathbf{f}_{out}^{rd} , which is given by

$$\mathcal{L}_{in}^{rd} = -\sum_{k=1}^K \tilde{y}_k^{rd} \log P_k^{in}(\mathbf{f}_{in}^{rd}), \quad \mathcal{L}_{out}^{rd} = -\sum_{k=1}^K y_k^{rd} \log P_k^{out}(\mathbf{f}_{out}^{rd}), \quad (3)$$

where K is the number of styles ($K = 10$ as mentioned in Section 3). \tilde{y}_k^{rd} (*resp.*, y_k^{rd}) is the k -th entry of style label vector for the composite (*resp.*, ground-truth) rendered image $\tilde{\mathbf{I}}^{rd}$ (*resp.*, \mathbf{I}^{rd}). The style label vectors of rendered images have been introduced in Section 3. P_k^{in} (*resp.*, P_k^{out}) is the style classifier before (*resp.*, after) the second stage. $P_k^{in}(\cdot)$ or $P_k^{out}(\cdot)$ is the k -th entry of predicted style distribution.

For real images without ground-truth style labels, standard classification loss cannot be used. Since the foreground style should be harmonized to the same as background style during harmonization, we design a novel style aggregation (SA) loss composed of two loss terms based on the following two assumptions. Firstly, \mathbf{f}_{out}^{rl} is likely to have one style only if \mathbf{f}_{in}^{rl} has this style. Thus, we adopt the weighted classification loss:

$$\mathcal{L}_W^{rl} = -\sum_{k=1}^K P_k^{in}(\mathbf{f}_{in}^{rl}) \log P_k^{out}(\mathbf{f}_{out}^{rl}). \quad (4)$$

Intuitively, if $P_k^{in}(\mathbf{f}_{in}^{rl})$ is small, $P_k^{out}(\mathbf{f}_{out}^{rl})$ would also be penalized. Because rendered images and real images are projected to the same feature space after the first stage, P^{in} and P^{out} are shared across two domains.

Secondly, the style distribution of \mathbf{f}_{out}^{rl} should be more concentrated than that of \mathbf{f}_{in}^{rl} . Since entropy depends on concentration and lower entropy implies higher concentration, we propose an entropy reduction loss to guarantee that the style distribution becomes more concentrated after harmonization:

$$\mathcal{L}_{ER}^{rl} = \max(0, m + \sum_{k=1}^K P_k^{in}(\mathbf{f}_{in}^{rl}) \log P_k^{in}(\mathbf{f}_{in}^{rl}) - \sum_{k=1}^K P_k^{out}(\mathbf{f}_{out}^{rl}) \log P_k^{out}(\mathbf{f}_{out}^{rl})), \quad (5)$$

where m is a margin. We wrap up the above two loss terms as a style aggregation loss $\mathcal{L}_{SA}^{rl} = \mathcal{L}_W^{rl} + \mathcal{L}_{ER}^{rl}$. Therefore, the total loss function for training the generators is

$$\mathcal{L}_G = \mathcal{L}_{rec}^{rd} + \mathcal{L}_{rec}^{rl} + \lambda_{adv} \mathcal{L}_G^{rd} + \lambda_{sty}^{rd} (\mathcal{L}_{in}^{rd} + \mathcal{L}_{out}^{rd}) + \lambda_{sty}^{rl} \mathcal{L}_{SA}^{rl}, \quad (6)$$

where λ_{adv} , λ_{sty}^{rd} , and λ_{sty}^{rl} are trade-off parameters. More implementation details can be found in the Supplementary.

5 Experiments

5.1 Dataset

Rendered Image Dataset RdHarmony contains 180,000 pairs of composite rendered images and ground-truth rendered images. Note that we sample 65,000 pairs with “human” foregrounds from maximal 135,000 pairs to train the harmonization network, aiming to reduce data redundancy and improve training efficiency.

Real Image Dataset iHarmony4 [2] contains 73,146 pairs of composite real images and ground-truth real images. In the main experiments, we treat “human” as an example novel category and other categories as base categories. The training set is split into novel training set (18,718 pairs) from “human” category and base training set (47,024 pairs) from other categories. Similarly, the test set is split into novel test set (1670 pairs) and base test set (5734 pairs). Experiments of the other 5 object categories could be found in the Supplementary.

Formally, we denote the rendered training set from novel category as $\mathcal{D}_{tr,n}^{rd}$ and the real training set from base (*resp.*, novel) categories as $\mathcal{D}_{tr,b}^{rl}$ (*resp.*, $\mathcal{D}_{tr,n}^{rl}$). We denote real test set from base (*resp.*, novel) categories as $\mathcal{D}_{te,b}^{rl}$ (*resp.*, $\mathcal{D}_{te,n}^{rl}$). By default, we merge $\mathcal{D}_{tr,n}^{rd}$ and $\mathcal{D}_{tr,b}^{rl}$ as the whole training set, and evaluate on both $\mathcal{D}_{te,b}^{rl}$ and $\mathcal{D}_{te,n}^{rl}$ (see Figure 1).

5.2 Preliminary Results

We evaluate the performance of cross-category harmonization and cross-domain harmonization. We train iDIH on $\mathcal{D}_{tr,n}^{rl}$, $\mathcal{D}_{tr,b}^{rl}$ and $\mathcal{D}_{tr,n}^{rd}$ separately and test on $\mathcal{D}_{te,n}^{rl}$. As the size of $\mathcal{D}_{tr,n}^{rl}$ is $N = 18718$, we sample N training images from $\mathcal{D}_{tr,b}^{rl}$ (*i.e.*, $\mathcal{D}_{tr,b}^{rl}(sub)$) and $\mathcal{D}_{tr,n}^{rd}$ (*i.e.*, $\mathcal{D}_{tr,n}^{rd}(sub)$) for fair comparison. From Table 1, we observe that although cross-category harmonization could harmonize the composite images (column 3 *v.s.* column 1), its performance is much worse than within-category harmonization (column 3 *v.s.* column 2). We also observed that cross-domain harmonization hurts the performance badly, which demonstrates the large domain gap between rendered images and real images (column 4 *v.s.* column 2).

#	1	2	3	4
Training data	-	$\mathcal{D}_{tr,n}^r$	$\mathcal{D}_{tr,b}^r(sub)$	$\mathcal{D}_{tr,n}^{rd}(sub)$
fMSE↓	931.74	386.36	486.54	1100.86
PSNR↑	33.29	35.71	34.64	31.25

Table 1: Results of models trained on $\mathcal{D}_{tr,n}^r$, $\mathcal{D}_{tr,b}^r(sub)$, and $\mathcal{D}_{tr,n}^{rd}(sub)$ and tested on $\mathcal{D}_{te,n}^r$. “-” denotes the metrics directly tested on the composite real images.

#	Training data	Method	$\mathcal{D}_{te,n}^r$			$\mathcal{D}_{te,b}^r$		
			MSE↓	fMSE↓	PSNR↑	MSE↓	fMSE↓	PSNR↑
1	-	Input Composite	155.74	931.74	33.29	177.34	1505.92	31.15
2	$\mathcal{D}_{tr,b}^r$	DoveNet [9]	69.06	458.15	35.01	65.87	674.57	33.94
3		Hao <i>et al.</i> [10]	52.42	407.26	35.43	58.21	582.79	34.54
4		iDIH [68]	46.67	417.71	35.21	53.03	566.96	34.77
5		RainNet [43]	58.40	525.36	34.97	59.54	701.07	34.44
6		Guo <i>et al.</i> [45]	49.92	416.21	35.58	52.31	577.16	35.01
7		two-stage training	43.06	383.39	35.67	49.56	533.10	34.98
8	$\mathcal{D}_{tr,n}^r$ & $\mathcal{D}_{tr,b}^r$	dataset fusion	38.31	368.50	35.64	44.15	485.94	35.27
9		UNIT [29]	55.09	458.72	34.74	56.21	607.14	34.38
10		CycleGAN [54]	51.82	474.78	34.64	56.30	592.33	34.52
11		CUT [54]	48.59	427.25	35.14	52.81	572.53	34.72
12		CharmNet	30.83	296.40	36.60	39.41	432.19	35.83
13	$\mathcal{D}_{tr,n}^r$ & $\mathcal{D}_{tr,b}^r$	iDIH [68]	27.71	259.17	37.12	36.17	422.98	36.04

Table 2: Results of models trained on various training data and tested on $\mathcal{D}_{te,n}^r/\mathcal{D}_{te,b}^r$. “-” denotes metrics directly tested on composite real images. Note that row 13 using real training images $\mathcal{D}_{tr,n}^r$ from novel category serves as the upper bound.

In summary, without real training images from novel category $\mathcal{D}_{tr,n}^r$, solely using $\mathcal{D}_{tr,b}^r$ or using $\mathcal{D}_{tr,n}^{rd}$ will both lead to sub-optimal performance, which motivates us to train with a combination of $\mathcal{D}_{tr,b}^r$ and $\mathcal{D}_{tr,n}^{rd}$.

5.3 Main Results

As shown in Table 2, we divide baselines into two groups according to their used training data. The first group contains recent image harmonization baselines (row 2 to row 6) [9, 10, 43, 45, 68] that are trained with only $\mathcal{D}_{tr,b}^r$. Among them, iDIH [68] achieves competitive performance. Although the recent method [9] achieves better performance, we opt for iDIH due to its simple network architecture. Thus, the rest of baselines (row 7 to row 12) and our method are all built upon iDIH. The second group contains domain adaptation models trained with both $\mathcal{D}_{tr,n}^r$ and $\mathcal{D}_{tr,b}^r$. Since there are no existing domain adaptation methods that could be directly applied to our task, we evaluate two straightforward methods, *i.e.*, two-stage training and dataset fusion, as well as Image-to-Image (I2I) translation methods [29, 54, 54] for comparison. Specifically, two-stage training means training with $\mathcal{D}_{tr,n}^{rd}$ and fine-tuning with $\mathcal{D}_{tr,b}^r$. Dataset fusion means directly training with a mixture of $\mathcal{D}_{tr,n}^{rd}$ and $\mathcal{D}_{tr,b}^r$. For I2I translation methods, we train UNIT [29], CycleGAN [54], and CUT [54] to translate from rendered image domain $\mathcal{D}_{tr,n}^{rd}$ to real image domain $\mathcal{D}_{tr,b}^r$. Then, we use translated $\mathcal{D}_{tr,n}^{rd}$ to augment $\mathcal{D}_{tr,b}^r$ as the new training set.

The results of the second baseline group are also reported in Table 2. Two-stage training only brings minor performance gain (row 7 *v.s.* row 4), probably because the statistics of real image domain are only considered during the second stage. The dataset fusion method improves the performance by a large margin (row 8 *v.s.* row 4), which indicates that RdHar-



Figure 3: Example results generated by different baselines and our method on $\mathcal{D}_{te,n}^r$. From left to right, we show the input composite real image, ground-truth real image, as well as the harmonized images generated by iDIH (row 4 in Table 2), CUT, dataset fusion, and our CharmNet. The foregrounds are outlined in red.

mony, though a rendered image dataset, is complementary to real image dataset by providing essential information of novel category. Besides, the performances of UNIT, CycleGAN, and CUT are significantly degraded and even worse than iDIH (row 4), due to the low quality of translated images and the severely distorted illumination statistics. Our CharmNet outperforms all the baselines within the second group by a large margin. Apart from novel category, our method also significantly enhances the performance on base categories, which indicates that cross-domain cross-category knowledge transfer is also very useful.

In addition, we report the result of iDIH trained with both $\mathcal{D}_{tr,n}^r$ and $\mathcal{D}_{tr,b}^r$ (row 13). This model serves as an upper bound for domain adaptation methods, assuming that there are abundant real training images from novel category. Despite the performance gap between row 12 and row 13, our method has shown great potential to bridge the gap between real images and rendered images. Since extending rendered image dataset costs significantly fewer manual efforts than extending real image dataset, we firmly believe that using both rendered images and real images for image harmonization is a promising research direction.

5.4 Qualitative Analyses

Given a composite real image from $\mathcal{D}_{te,n}^r$, the harmonized outputs generated by iDIH [88] (row 4 in Table 2), CUT [54], dataset fusion, and our CharmNet are shown in Figure 3. Compared with other baselines, CharmNet could generate more favorable results with consistent foreground and background. Besides, CharmNet concentrates more on adjusting the foreground style to the background style without corrupting the appearance of human foreground, making the harmonized results visually closer to the ground-truth real images. In the Supplementary, we additionally provide example results on $\mathcal{D}_{te,b}^r$ and compare the harmonization results of our method and upper bound (row 13 in Table 2).

5.5 Ablation Studies

We perform ablation studies to support the design of our CharmNet. First, we ablate the adversarial loss and style-related losses in Eqn. 6. The results demonstrate that each term is helpful and their combination achieves further improvement with mutual collaboration. Since the main challenge in our problem is to bridge the huge domain gap between rendered images and real images, so the adversarial loss contributes most to the performance gain. The two loss terms in style aggregation (SA) loss work together to meet our expectation, which leads to further improvement. Next, we analyze the network design in terms of the number of unshared layers in the domain-specific encoding stage and domain-specific decoding stage (see Section 4.1), which shows that it is necessary to use a suitable number of unshared layers. The detailed results are left to the Supplementary.

5.6 Training Data Analyses

We first explore how the diversity of RdHarmony dataset influences the harmonization performance by varying the number of 3D scenes and styles. Then with all 3D scenes and 10 styles, we explore the performance variance when using different numbers of rendered training images. We also explore the performance when using the whole $\mathcal{D}_{tr,n}^{rd}$ and partial real training images from novel category. In other words, we can jointly use rendered images and real images for the novel category. Specifically, we can use rendered training images to augment the real training images of novel category, leading to consistently better performance. When real training images are insufficient, the performance gain brought by rendered training images is more notable. The detailed results are also left to the Supplementary.

5.7 Evaluation on Real Composite Images

In real-world applications, there is no ground-truth for a real-world composite image with human foregrounds, so it is infeasible to evaluate the model performance quantitatively. Following [9, 15], we select 46 images with human foregrounds from the real-world composite images [15] and conduct user study. Detailed results are left to the Supplementary.

6 Conclusion

In this paper, we have contributed a large-scale rendered image harmonization dataset RdHarmony, which simulates the process of capturing the same scene under different illumination conditions based on 3D rendering techniques. We have also presented a cross-domain image harmonization network CharmNet to bridge the gap between rendered images and real images. Our constructed dataset and proposed network have made a big step towards jointly using rendered images and real images for image harmonization.

Acknowledgement

The work was supported by the Shanghai Municipal Science and Technology Major/Key Project, China (2021SHZDZX0102, 20511100300) and National Natural Science Foundation of China (Grant No. 61902247).

References

- [1] Rita Chattopadhyay, Qian Sun, Wei Fan, Ian Davidson, Sethuraman Panchanathan, and Jieping Ye. Multisource domain adaptation and its application to early detection of fatigue. *ACM Trans. Knowl. Discov. Data*, 6(4), 2012.
- [2] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *CVPR*, 2018.
- [3] Yunjey Choi, Min-Je Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 2018.
- [4] Wenyan Cong, Jianfu Zhang, Li Niu, Liu Liu, Zhixin Ling, Weiyuan Li, and Liqing Zhang. DoveNet: Deep image harmonization via domain verification. In *CVPR*, 2020.
- [5] Wenyan Cong, Li Niu, Jianfu Zhang, Jing Liang, and Liqing Zhang. Bargainnet: Background-guided domain translation for image harmonization. In *ICME*, 2021.
- [6] Wenyan Cong, Xinhao Tao, Li Niu, Jing Liang, Xuesong Gao, Qihao Sun, and Liqing Zhang. High-resolution image harmonization via collaborative dual transformations. In *CVPR*, 2022.
- [7] Xiaodong Cun and Chi-Man Pun. Improving the harmony of the composite image by spatial-separated attention module. *IEEE Trans. Image Process.*, 29:4759–4771, 2020.
- [8] W. Deng, L. Zheng, Q. Ye, G. Kang, Y. Yang, and J. Jiao. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *CVPR*, 2018.
- [9] Yang Fu, Yunchao Wei, Guanshuo Wang, Yuqian Zhou, Honghui Shi, and Thomas S. Huang. Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification. In *ICCV*, 2019.
- [10] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015.
- [11] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *J. Mach. Learn. Res.*, 17(1):2096–2030, 2016.
- [12] W. Ge, S. Yang, and Y. Yu. Multi-evidence filtering and fusion for multi-label classification, object detection and semantic segmentation based on weakly supervised learning. In *CVPR*, 2018.
- [13] R. Gopalan, Ruonan Li, and R. Chellappa. Domain adaptation for object recognition: An unsupervised approach. In *ICCV*, 2011.
- [14] Zonghui Guo, Dongsheng Guo, Haiyong Zheng, Zhaorui Gu, Bing Zheng, and Junyu Dong. Image harmonization with transformer. In *ICCV*, 2021.
- [15] Zonghui Guo, Haiyong Zheng, Yufeng Jiang, Zhaorui Gu, and Bing Zheng. Intrinsic image harmonization. In *CVPR*, 2021.

- [16] Yucheng Hang, Bin Xia, Wenming Yang, and Qingmin Liao. Scs-co: Self-consistent style contrastive learning for image harmonization. In *CVPR*, 2022.
- [17] Guoqing Hao, Satoshi Iizuka, and Kazuhiro Fukui. Image harmonization with attention-based deep feature modulation. In *BMVC*, 2020.
- [18] Xun Huang, Ming-Yu Liu, Serge J. Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018.
- [19] N. Inoue, R. Furuta, T. Yamasaki, and K. Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *CVPR*, 2018.
- [20] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- [21] Jiaya Jia, Jian Sun, Chi-Keung Tang, and Heung-Yeung Shum. Drag-and-drop pasting. *ACM Transactions on Graphics*, 25(3):631–637, 2006.
- [22] Yifan Jiang, He Zhang, Jianming Zhang, Yilin Wang, Zhe Lin, Kalyan Sunkavalli, Simon Chen, Sohrab Amirghodsi, Sarah Kong, and Zhangyang Wang. Ssh: A self-supervised framework for image harmonization. In *ICCV*, 2021.
- [23] Chanyong Jung, Gihyun Kwon, and Jong Chul Ye. Exploring patch-wise semantic relation for contrastive learning in image-to-image translation tasks. In *CVPR*, 2022.
- [24] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Changick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In *ICCV*, 2019.
- [25] Jean-François Lalonde and Alexei A. Efros. Using color compatibility for assessing image realism. In *ICCV*, 2007.
- [26] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *ECCV*, 2018.
- [27] Hsin-Ying Lee, Hung-Yu Tseng, Qi Mao, Jia-Bin Huang, Yu-Ding Lu, Maneesh Singh, and Ming-Hsuan Yang. Drit++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision*, pages 1–16, 2020.
- [28] Jun Ling, Han Xue, Li Song, Rong Xie, and Xiao Gu. Region-aware adaptive instance normalization for image harmonization. In *CVPR*, 2021.
- [29] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, volume 30, 2017.
- [30] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *ICCV*, 2019.
- [31] Qi Mao, Hung-Yu Tseng, Hsin-Ying Lee, Jia-Bin Huang, Siwei Ma, and Ming-Hsuan Yang. Continuous and diverse image-to-image translation via signed attribute vectors. *International Journal of Computer Vision*, 130(2):517–549, 2022.

- [32] Chen Minghao, Shuai Zhao, Haifeng Liu, and Deng Cai. Adversarial-learned loss for domain adaptation. In *AAAI*, 2020.
- [33] Li Niu, Wenyan Cong, Liu Liu, Yan Hong, Bo Zhang, Jing Liang, and Liqing Zhang. Making images real again: A comprehensive survey on deep image composition. *arXiv preprint arXiv:2106.14490*, 2021.
- [34] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *ECCV*, 2020.
- [35] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *ACM Transactions on Graphics*, 22(3):313–318, 2003.
- [36] Francois Pitie, Anil C Kokaram, and Rozenn Dahyot. N-dimensional probability density function transfer and its application to color transfer. In *ICCV*, 2005.
- [37] Erik Reinhard, Michael Ashikhmin, Bruce Gooch, and Peter Shirley. Color transfer between images. *IEEE Computer Graphics and Applications*, 21(5):34–41, 2001.
- [38] Konstantin Sofiiuk, Polina Popenova, and Anton Konushin. Foreground-aware semantic representations for image harmonization. In *WACV*, 2021.
- [39] Olga Sorkine, Ran Gal, Tommer Leyvand, and Ying-Qing Xu. Color harmonization. *ACM Transactions on Graphics*, 25(3):624–630, 2006.
- [40] A. Spurr, J. Song, S. Park, and O. Hilliges. Cross-modal deep variational hand pose estimation. In *CVPR*, 2018.
- [41] Kalyan Sunkavalli, Micah K. Johnson, Wojciech Matusik, and Hanspeter Pfister. Multi-scale image harmonization. *ACM Transactions on Graphics*, 29(4):125:1–125:10, 2010.
- [42] Fuwen Tan, Crispin Bernier, Benjamin Cohen, Vicente Ordonez, and Connelly Barnes. Where and who? automatic semantic-aware person composition. In *WACV*, 2018.
- [43] Michael W. Tao, Micah K. Johnson, and Sylvain Paris. Error-tolerant image compositing. In *ECCV*, 2010.
- [44] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR*, 2011.
- [45] Yi-Hsuan Tsai, Xiaohui Shen, Zhe Lin, Kalyan Sunkavalli, Xin Lu, and Ming-Hsuan Yang. Deep image harmonization. In *CVPR*, 2017.
- [46] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017.
- [47] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Perez. ADVENT: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *CVPR*, 2019.
- [48] L. Wei, S. Zhang, W. Gao, and Q. Tian. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*, 2018.

- [49] Shuchen Weng, Wenbo Li, Dawei Li, Hongxia Jin, and Boxin Shi. MISC: Multi-condition injection and spatially-adaptive compositing for conditional person image synthesis. In *CVPR*, 2020.
- [50] Su Xue, Aseem Agarwala, Julie Dorsey, and Holly E. Rushmeier. Understanding and improving the realism of image composites. *ACM Transactions on Graphics*, 31(4): 84:1–84:10, 2012.
- [51] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. CutMix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019.
- [52] Y. Zhang, P. David, and B. Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *ICCV*, 2017.
- [53] Xingyi Zhou, Arjun Karpur, Chuang Gan, Linjie Luo, and Qixing Huang. Unsupervised domain adaptation for 3d keypoint estimation via view consistency. In *ECCV*, 2018.
- [54] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.
- [55] Yang Zou, Zhiding Yu, B. V. Kumar, and J. Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, 2018.