

# Adversarial Image Composition with Auxiliary Illumination

Fangneng Zhan<sup>1,2</sup>[0000–0003–1502–6847], Shijian Lu<sup>1</sup>[0000–0002–6766–2506],  
Changgong Zhang<sup>2</sup>, Feiying Ma<sup>2</sup>, and Xuansong Xie<sup>2</sup>

<sup>1</sup> Nanyang Technological University

<sup>2</sup> Alibaba DAMO Academy

{fnzhan, shijian.lu}@ntu.edu.sg, {changgong.zcg, feiying.mfy}@alibaba-inc.com,  
xingtong.xxs@taobao.com

**Abstract.** Dealing with the inconsistency between a foreground object and a background image is a challenging task in high-fidelity image composition. State-of-the-art methods strive to harmonize the composed image by adapting the style of foreground objects to be compatible with the background image, whereas the potential shadow of foreground objects within the composed image which is critical to the composition realism is largely neglected. In this paper, we propose an Adversarial Image Composition Net (AIC-Net) that achieves realistic image composition by considering potential shadows that the foreground object projects in the composed image. A novel branched generation mechanism is proposed, which disentangles the generation of shadows and the transfer of foreground styles for optimal accomplishment of the two tasks simultaneously. A differentiable spatial transformation module is designed which bridges the local harmonization and the global harmonization to achieve their joint optimization effectively. Extensive experiments on pedestrian and car composition tasks show that the proposed AIC-Net achieves superior composition performance qualitatively and quantitatively.

## 1 Introduction

With the advances of deep neural networks (DNNs), image composition has been attracting increasing attention as a typical approach for image synthesis, augmented reality, etc. With a foreground object and a background image, a direct combination tends to introduce unrealistic artifacts in the composed image with various inconsistency in colors, illuminations, texture, etc. Several methods have been proposed for harmonizing the foreground object and the background image, including DNN-based image translation [1], local statistic transfer [2], etc. On the other hand, existing methods typically simplify the harmonization by solving the image consistency within the foreground object region only but neglecting that realistic image composition goes far beyond that.

Specifically, a typical visual effect that is often associated with objects in scenes is their projected shadows due to the 3D nature of the world. While embedding a foreground object into a background image for image composition,



**Fig. 1.** Given foreground objects and directly composed local images as shows in columns 1 and 2, the proposed AIC-Net disentangles the shadow generation and style transfer to produce realistic local harmonization as shown in column 3. With our spatial transformation module, the locally harmonized images in column 3 can be adapted and composed into the globally harmonized images as shown in column 4.

the shadow effects of the foreground object are critical which can significantly improve the realism of the composed image if handled properly. Due to various constraints in 2D images, such as complicated illuminations, losing of 3D information, etc., traditional computer vision methods are highly susceptible to failure for the task of high-fidelity shadow generation.

In this work, we design a novel Adversarial Image Composition Net (AIC-Net) that aims to achieve realistic image composition with adaptive shadow effects as illustrated in Fig. 1. AIC-Net adopts generative adversarial ideas to harmonize the foreground object within the background image and generates realistic shadow effects by learning from real shadows in natural images. It consists of local harmonization and global harmonization which are inter-connected and can be optimized jointly. We design a novel branched generation mechanism for high-fidelity local harmonization which disentangles the shadow generation and the style transfer of the foreground object for optimal accomplishment of two tasks simultaneously. An auxiliary illumination model is introduced to infer the illumination conditions of the background image such as lighting directions which help generate realistic shadows of the foreground object greatly. In addition, a differentiable spatial transformer is introduced to bridge the local region and the global background for joint optimization of the local and global harmonization. Two discriminators are employed where a local discriminator guides to learn style transfer and shadow generation of the foreground object and a global discriminator guide to learn global image harmonization.

The contributions of this work are three folds. First, we design an innovative AIC-Net that is capable of generating high-fidelity shadow of the foreground object adaptively. Second, we design a novel branched generation mechanism

that disentangles the style transfer and shadow generation of the foreground object and accomplishes the two tasks optimally and simultaneously. Third, we design spatial transformer module (STM) that achieves joint optimization of local harmonization and global harmonization effectively.

## 2 Related Work

### 2.1 Image Composition & Harmonization

Image composition aims to generate new images by embedding foreground objects into background images [3–13]. GP-GAN [3] composes high-resolution images by leveraging Poisson blending [14]. [4] presents a spatial transformer GAN by inserting STNs into the generator. [15] describes a Compositional GAN that introduces a self-consistent composition-decomposition network. [16] proposes a model to distinguish natural photographs from automatically generated composite images. [17] proposes a generative adversarial network (GAN) architecture for automatic image compositing with consideration of geometric, color and boundary consistency.

Image harmonization deals with the inconsistency between a foreground region and a background image for composition. Several methods have been proposed to adapt the appearance of the foreground object for image harmonization. For example, [18] matches contrast, texture, and blur by manipulating the scales of a pyramid decomposition of images. [19] studies image statistics that determine the realism of the composite images. [20] proposes a gradient domain to achieve composition which can prevent color bleeding without changing the boundary location. A number of deep-learning based image harmonization methods have also proposed in recent years. For example, [21] utilizes visual semantics to guide the process of replacement to generate realistic, artifact-free images with diverse styles. [2] achieves both spatial and inter-scale statistical consistency by carefully determining the local statistics to be transferred. [1] proposes an end-to-end deep convolutional neural network for image harmonization, which can capture both the context and semantic information of the composite images during harmonization. [17] combines a transformation network, a refinement network, a segmentation network and a pair of discriminator networks to achieve image composition via adversarial learning. Besides image harmonization, some work [22, 16] focuses on the realism assessment of the composed images. [23] handles image harmonization through domain verification discriminator with the insight that the foreground needs to be translated to the same domain as background.

Illumination estimation is a classic computer vision problem and it is critical for achieving realistic lighting and shadows in image harmonization. Several deep learning based methods have been proposed to recover illumination from 2-dimensional RGB images [24–28].

### 2.2 Generative Adversarial Networks

GANs [29] have achieved great success in image generation from either existing images or random noises. Instead of manually selecting features and parameters,

GAN generators learn an optimal mapping from random noise or existing images to synthesized images, while GAN discriminators differentiate the synthesized images from real ones via adversarial learning. Several GAN-based image synthesis methods have been reported in recent years. [30] introduces Laplacian pyramids that improve the quality of GAN-synthesized images greatly. [31] proposes an end-to-end trainable network for inserting an object instance mask into the semantic label map of an image. Other systems attempt to learning more reasonable potential features [32], exploring new training approach [33], visualizing GANs at the unit, object and scene level [34], etc. To expand the applicability, [13] generate shadow for augmented reality by modeling the mapping relation between the virtual object shadow and the real-world.

GAN-based image-to-image translation has been widely explored due to its wide applicability. CycleGAN [35] proposes a cycle-consistent adversarial network for realistic image-to-image translation. [36] investigates conditional adversarial networks for image-to-image translation. [37] adopts unsupervised learning to improve the realism of synthetic images using unlabelled real data. Other GANs [38–41] also achieve great performance in image-to-image translation.

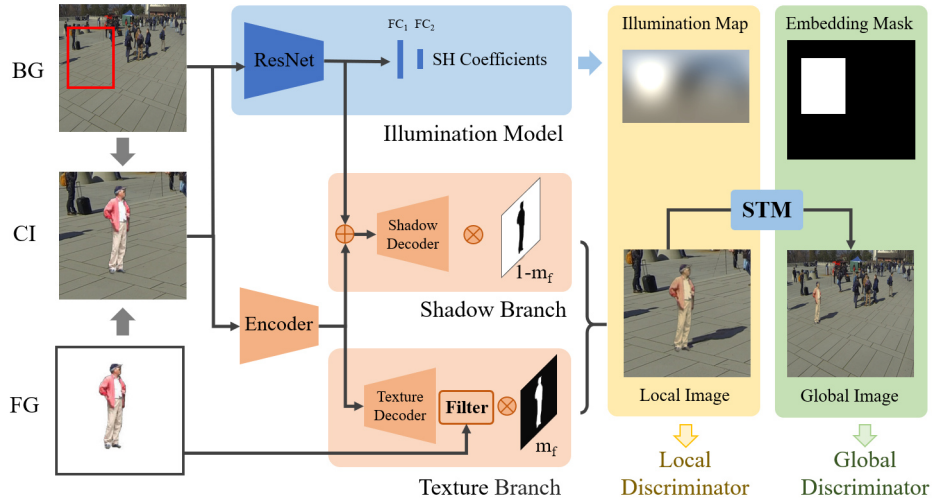
### 3 Proposed Method

The proposed AIC-Net consists of local harmonization and global harmonization which are inter-connected and can be optimized jointly as illustrated in Fig. 2. Detailed network structures and training strategies will be presented in the following subsections.

#### 3.1 Local Harmonization

The local harmonization is achieved by a local discriminator and a branched generation network which consists of a shadow branch and a texture branch as illustrated in Fig. 2. Given a background image (BG), a meaningful region within it is first selected for composition with a foreground object (FG). The selection of the meaningful region will be described in more details in *Experiment Setting*. We thus obtain a directly composed image (CI)  $x$  by embedding the foreground object into the selected region in the background image, as well as a foreground mask  $m_f$  by setting the foreground object pixels to 1 and the rest pixels in  $x$  to 0 as illustrated in Fig. 2

In local harmonization, the generator aims to transfer the style of the composed image and generate realistic object shadows concurrently. Although the two processes can be jointly handled through image-to-image translation, they have very different translation objectives and learn with distinct image-to-image mapping mechanisms. The concurrent learning of both degrades the translation severely as illustrated in ‘W/o Branched Generation’ in Fig. 3. We design a novel branched generation strategy to disentangle style transfer and shadow generation for optimal image harmonization. As shown in Fig. 2, the generator has two branches that share the same Encoder. With a background image and



**Fig. 2.** The structure of the proposed AIC-Net: With a background image  $BG$  (with specified composition regions) and foreground object  $FG$  to be composed, a directly composed images  $CI$  can be derived directly. A shadow branch and a texture branch will learn to generate shadow and transfer the style of  $CI$  which produces a harmonized *Local Image*. A spatial transformer module (STM) then learns to compose the *Local Image* with the background image to produce a *Global Image*. An *Illumination Model* will predict SH coefficients which can be used to reconstruct an *Illumination Map* (‘FC’ denote fully connected layer). The *Local Image* will be concatenated with the *Illumination Map* to be distinguished by the *Local Discriminator*. The *Global Image* will be concatenated with an *Embedding Mask* to be distinguished by the *Global Discriminator*.

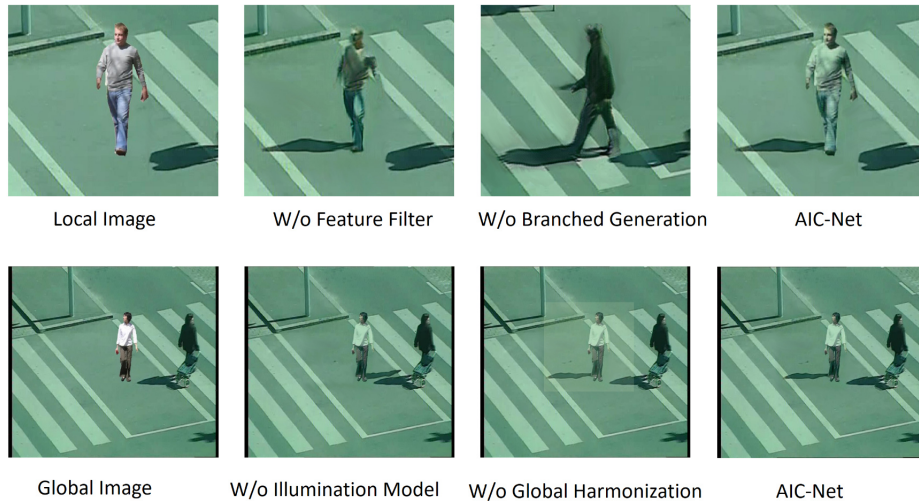
a composed image as the input, the decoder in shadow branch strives to generate shadows in the composed image while the decoder in the texture branch strives to transfer the style of the composed image to be compatible with the background image.

**Shadow Branch** To avoid synthesizing shadows with false directions as shown in ‘W/o Illumination Model’ in Fig. 3, we need to infer the global illumination information from the background images. In this work, we adopt a deep learning based method for illumination estimation. Specially, the ground truth of illumination map which is a panorama is represented by spherical harmonics (SH) coefficients [42], then a model with ResNet [43] as the backbone is adopt to regress the SH coefficients from the background images with limited field of view. With the predicted SH coefficients, an illumination map  $M_s$  can reconstructed as follows:

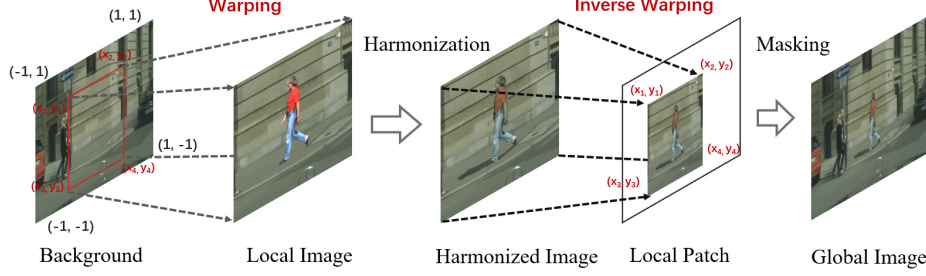
$$M(s) = \sum_{i=1}^M SH_i * y_i(s) \quad (1)$$

where  $s$  is the spatial direction,  $M$  is the number of SH coefficients,  $y_i$  and  $SH_i$  are the  $i$ -th SH basis function and SH coefficients, respectively. We pre-train the illumination model on the Laval Indoor dataset [25] and [44] and fix the model weights while training AIC-Net.

The illumination model is able to infer the illumination from cues such as shadow existed in the background. Thus, the feature extracted by the illumination model contains rich information of shadow direction, size and intensity. In the shadow branch as shown in Fig. 2, we concatenate the feature extracted by the illumination model and the feature extracted by Encoder to integrate the shadow and illumination information into the shadow generation. The shadow decoder then strives to generate shadows with realistic shadow directions and styles in the local images. A binary mask ( $1-m_f$ ) is applied to the output for integration with the texture branch. As the shadow branch may translate other regions beyond the shadow areas undesirably, we feed real images into the shadow branch to keep the non-shadow areas be the same with the input. The shadow branch will thus concentrate on the translation of potential shadow areas for realistic shadow generation. To drive the effective learning of shadow branch, the reconstructed illumination map will be combined with the locally harmonized image to be distinguished by the local discriminator.



**Fig. 3.** Illustration of the proposed guided feature filter, branched generation, illumination model and global harmonization: The introduction of the branched generation and guided feature filter helps to suppress various artifacts and generate realistic shadows and foreground object in the local image. The illumination model helps to generate shadow with high-fidelity shadow styles and directions. The global harmonization helps to suppress artifacts and style discrepancy between the local and global images.



**Fig. 4.** Global harmonization: With a selected region in the background image, a homography matrix  $H$  can be estimated to transform the selected region to the local image. After the local harmonization, the harmonized local image can be transformed back to the background image by an inverse warping  $H^{-1}$ . A masking operation is followed to obtain the global image.

**Texture Branch** The texture branch aims to translate the texture of the foreground region to be harmonious with the background image. As a normal GAN generator tends to over-translate the foreground region which often degrades the foreground features as illustrated in ‘W/o Feature Filter’ in Fig. 3, we introduce a guided feature filter to preserve the ‘content features’ while performing style transfer. Different from the guided image filter [45] which is conducted in image level, guided feature filter is conducted in the feature space to achieve flexible style transfer. With the foreground object as the content guidance ( $C$ ) and the feature decoded by the Texture Decoder as the style guidance ( $S$ ), an output feature ( $T$ ) with transferred style but preserved contents can be obtained through a linear model as follows:

$$T_i = a_k C_i + b_k, \forall i \in \omega_k \quad (2)$$

where  $i$  is a pixel index and  $\omega_k$  is a local square window centered at pixel  $k$ ,  $a_k$  and  $b_k$  are linear coefficients to be optimized. The minimization of the reconstruction error between  $C$  and  $T$  can be achieved by the following cost function:

$$E(a_k, b_k) = \sum_{i \in \omega_k} ((a_k S_i + b_k - C_i)^2 + \epsilon a_k^2) \quad (3)$$

where  $\epsilon$  is a regularization parameter that prevents  $a_k$  from being too large. It can be solved via linear regression:

$$a_k = \frac{\frac{1}{|\omega|} \sum_{i \in \omega_k} S_i - \mu_k \bar{C}_k}{\bar{\sigma}_k + \epsilon} \quad (4)$$

$$b_k = \bar{R}_k - a_k \mu_k \quad (5)$$

where  $\mu_k$  and  $\sigma_k^2$  are the mean and variance of  $S$  in  $\omega_k$ ,  $|\omega|$  is the number of pixels in  $\omega_k$ , and  $\bar{C}_k = \frac{1}{|\omega|} \sum_{i \in \omega_k} C_i$  is the mean of  $C$  in  $\omega_k$ . Then we can obtain the

transferred feature  $T$ . A following convolution layer will finally generate images with transferred style. A binary mask  $m_f$  is applied to the output for further integration with the shadow branch. By denoting the outputs of the shadow branch and the texture branch by  $x_s$  and  $x_t$ , the locally harmonized image  $x_h$  can be obtained as follows:

$$x_h = x_t * m_f + x_s * (1 - m_f) \quad (6)$$

### 3.2 Global Harmonization

The harmonized local region should be consistent with global image regarding the image style especially the boundary of the local region in the global image, otherwise the global image will tend to present clearly artifacts as shown in ‘W/o Global Harmonization’ of Fig. 3 which should be addressed as well for high-fidelity global harmonization. We adopt a global discriminator to ensure the consistency in global harmonization. With real background image as the training reference, the global discriminator learns to distinguish the harmonized global image (Fake Global in Fig. 2) and real global images (Real Global in Fig. 2), driving the generator for generating harmonious shadows and an attention network for eliminating the style discrepancy and boundary artifacts. A spatial transformer module (STM) is designed to bridge the local image and global image to achieve joint optimization.

**Spatial Transformer Module:** The spatial transformer is a differentiable module that applies a homography to the local image as illustrated in Fig. 4. We denote the coordinates of the four vertices of the background image by  $[(-1, 1), (1, 1), (1, -1), (-1, -1)]$ , and the coordinate of the four vertices of the selected region by  $[(x_1, y_1), (x_2, y_2), (x_3, y_3), (x_4, y_4)]$  ( $-1 < x_i < 1, -1 < y_i < 1, i = 1, 2, 3, 4$ ). With four paired vertices, the homography  $H$  can be computed which applied STM to warp the selected region to the local image. Once the local image is harmonized, it can be warped back to the local region of the background image by using the inverse homography  $H_{-1}$  as illustrated in Fig. 4. A warping mask with the selected region as 1 and other regions as 0 can be obtained simultaneously. A globally harmonized image can thus be obtained by masking the local patch into the background image. Since the STM and the masking operation are both differentiable, the local harmonization and the global harmonization can be bridged for end-to-end joint optimization.

### 3.3 Adversarial Training

With a GAN-based structure, the network is trained through adversarial learning between a generator and two discriminators. For clarity purpose, we denote the  $BG, CI$  in Fig. 2 as  $X$  and  $x$ , the generator and the two branches *Shadow Branch* and *Texture Branch* as  $G, G_s$  and  $G_t$ , **Global Discriminator** and **Local Discriminator** as  $D_G$  and  $D_L$ . The foreground mask is denoted as  $m_f$ , the real global and local images as  $Y$  and  $y$ . The **Embedding Mask** in real global images and harmonized global images as  $m_Y$  and  $m_X$ .



For the local harmonization, the shadow and texture branches strive to generate shadow and translate foreground that are denoted by  $G_s(X, x)$  and  $G_t(X, x)$ , respectively. The local image  $x_h$  can thus be denoted by:

$$x_h = G_s(X, x) * (1 - m_f) + G_t(X, x) * m_f \quad (7)$$

The local discriminator tries to distinguish the local image while the generator tries to mislead the local discriminator. As we adopt Wasserstein GAN [33] objective for training, the local adversarial loss of the generator and local discriminator can be defined by:

$$L_{D_L} = E[D_L(x_h)] - E[D_L(y)] \quad (8)$$

$$L_{G_L} = -E[D_L(x_h)] \quad (9)$$

To ensure the irrelevant region keep unchanged and preserve more details of the original image, an identity loss is defined as follows:

$$L_{S_{idt}} = E[G(Y, y) - y] \quad (10)$$

For the global harmonization, the global discriminator will distinguish the pair of global image  $X_h$  and embedding mask  $m_X$  which drives the generator to generate consistent texture and suppress artifacts around the local image boundary. The global adversarial loss of the generator and the global discriminator can thus be defined by:

$$L_{D_G} = E[D_G(X_h, m_X)] - E[D_G(Y, m_Y)] \quad (11)$$

$$L_{G_G} = -E[D_G(X_h, m_X)] \quad (12)$$

The overall objective function for the generator is:

$$L_G = L_{G_L} + \lambda_G L_{G_G} + \lambda_{G_{idt}} L_{S_{idt}} \quad (13)$$

where  $\lambda_G, \lambda_{G_{idt}}$  denote weights of global adversarial loss and generator identity loss respectively. The overall objective function for the discriminators is:

$$L_D = L_{D_L} + \lambda_{D_G} L_{D_G} \quad (14)$$

where  $\lambda_{D_G}$  denotes the weight of global adversarial loss for the discriminator.  $L_D$  and  $L_G$  drive the learning of the whole model alternately.

## 4 Experiment

### 4.1 Dataset

**PRID 2011** [46] was created to evaluate person re-identification methods. The dataset consists of images extracted from multiple person trajectories recorded from two surveillance cameras.

**WILDTRACK** [47] is multi-camera detection and tracking dataset of pedestrians in the wild, where the cameras’ fields of view has large part overlap.

**Penn-Fudan** [48] is an image database that was used for pedestrian detection evaluation as reported in [1]. The images are taken from scenes around campus and urban streets. The objects we are interested in these images are pedestrians. Each image will have at least one pedestrian.

**UA-DETRAC** [49] dataset consists of 100 challenging video sequences captured from over 140 thousand frames of real-world traffic scenes. More than 1.2 million labeled bounding boxes of vehicles with information of occlusion, weather, vehicle category and truncation are annotated.

**Cityscapes** [50] dataset focuses on semantic understanding of urban street scenes. It has street scenes from 50 cities and provides semantic segmentation of 30 classes. 3475 images are annotated with bounding boxes of persons.

**ShapeNet** [51] is a large repository of 3D object models which index more than 3 million models of which 220 thousand are classified into 3,135 categories.

## 4.2 Experiment Setting

In the pedestrian harmonization experiment, 2000 background images are collected from [46] and [47] and 500 pedestrians are cropped from [48] as the foreground objects. In the car harmonization experiment, we collect the background images from [49] and [50] which contain images captured under different illumination conditions and viewpoints. Random horizontal flipping is applied to the background images to augment the training set. For the foreground images, 1000 3D car models from [51] are utilized to render 3000 car images in different views.

The local region to compose with the foreground object is derived by a simple mechanism. As the bounding boxes of the foreground objects are know, we first select a box region around an existed bounding box and the one with no overlap with other bounding boxes is selected as the local region.

The proposed method is implemented using the Tensorflow framework. The optimizer is Adadelta which employs adaptive learning rate. The model is trained in 100 thousand iterations with a batch size of 4. In addition, the network is trained on a workstation with an Intel Core i7-7700K CPU and a NVIDIA GeForce GTX 1080 Ti graphics card with 12GB memory. All input images are resized to  $256 \times 256$  in training.

## 4.3 Experiment Analysis

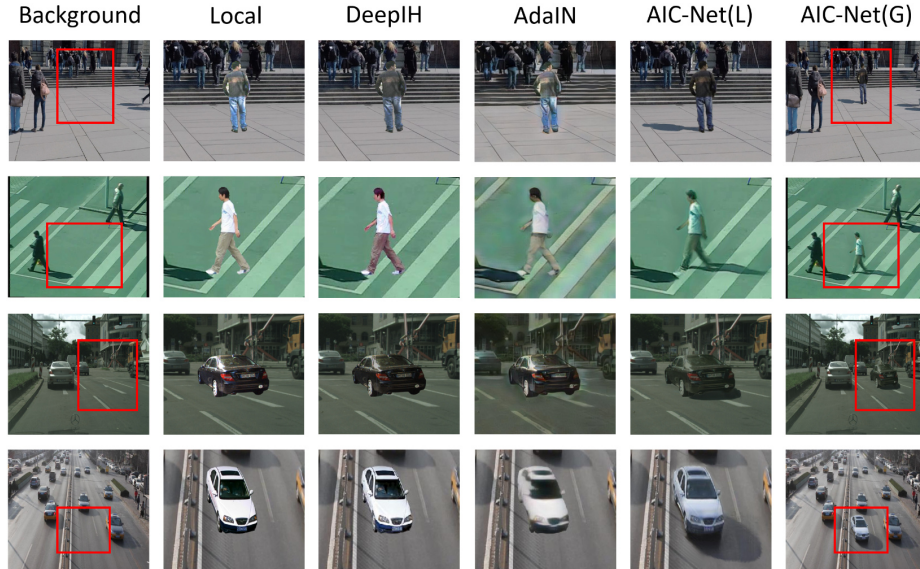
**Quantitative Analysis:** The evaluation is performed through Fréchet Inception Distance (FID) [52], Manipulation Score (MS) [17] and Amazon Mechanical Turk (AMT) study. For the evaluation with FID, the harmonized images are fed to the Inception network [53] trained on ImageNet [54], and the features from the layer before the last fully-connected layer are used to calculate the Fréchet Inception Distance. The Manipulation Score [17] is generated by a state-of-the-art manipulation detection algorithm [55] where a higher score indicates a higher

**Table 1.** Comparison of harmonization methods with different metrics: ‘Car’ and ‘Ped’ denote the results of car and pedestrian synthesis respectively. ‘Local’ and ‘Global’ denote harmonized local and global images. ‘MS’ metric denotes the manipulation scores. For AMT user study, the number in each cell is the percentage of the harmonized images that are deemed as the most realistic by Turkers.

Metrics	CycleGAN		DeepIH		DeepPH		AdaIN		AIC-Net	
	Car	Ped	Car	Ped	Car	Ped	Car	Ped	Car	Ped
<b>AMT(Local)</b>	8%	4%	12%	11%	20%	23%	18%	17%	<b>42%</b>	<b>45%</b>
<b>AMT(Global)</b>	3%	1%	16%	15%	25%	24%	21%	22%	<b>41%</b>	<b>42%</b>
<b>FID(Local)</b>	115.1	131.6	136.2	151.4	99.1	103.3	88.7	97.2	<b>84.4</b>	<b>97.1</b>
<b>FID(Global)</b>	102.4	126.1	121.8	139.3	92.4	101.9	84.6	92.8	<b>72.2</b>	<b>89.6</b>
<b>MS(Local)</b>	0.69	0.70	0.75	0.74	0.72	0.73	0.69	0.71	<b>0.63</b>	<b>0.64</b>
<b>MS(Global)</b>	0.71	0.70	0.73	0.76	0.73	0.74	0.70	0.74	<b>0.62</b>	<b>0.62</b>

possibility that the image is manipulated. For the AMT study which involves 10 users, we present a set of images (5 images by 5 methods) each time to the user to select the most realistic one. With 20 sets of presented images, an AMT score is derived by averaging the percentage of most realistic images as judged by the 10 users. ‘Local’ and ‘Global’ denote the results of locally and globally harmonized images respectively.

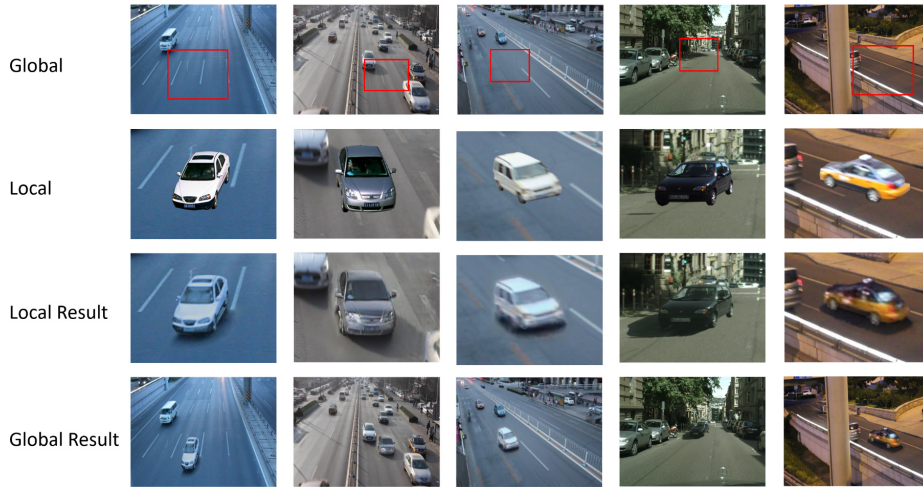
Table 1 compares the quality and realism of images processed by different harmonization methods. For the local harmonization, AIC-Net outperforms all compared methods with all metrics clearly as it harmonizes the foreground region and generates relevant shadow which improves the image realism significantly. The using of guided filter helps to preserve texture of high quality, thus obtaining the best FID scores. The spatial transformer module achieves the joint optimization of local and global harmonization, leading to best Manipulation score. CycleGAN can also synthesize shadows but the foreground region tends to be translated undesirably as it does not disentangle the shadow generation and style transfer processed. As DeepIH requires paired training images with semantic annotations, it performs terribly bad when paired training images are not available. As style transfer methods, DeepPH and AdaIN harmonize the foreground region successfully but all of them ignore the shadow generation completely. For the global harmonization, AIC-Net outperforms other methods by large margins as its joint optimization of local and global harmonization ensures consistent shadow and smooth local boundary. The lack of global harmonization leads to clear artifacts in the global image harmonized by CycleGAN. For DeepIH, DeepPH and AdaIN, there are not many artifacts but the lack of shadow and guided filter degrades the image realism and quality greatly. Normally, the global harmonized images present better score than local harmonized images as only small part of global image needs to be harmonized.



**Fig. 5.** Comparing AIC-Net with other harmonization methods: Images in col. 1 are background images with selected regions. Col. 2 shows directly composed local images. AIC-Net (L) and AIC-Net (G) show images after local and global harmonization.

**Qualitative Analysis:** We compare the images harmonized by different methods in Fig. 5, where AIC-Net (L) and AIC-Net (G) denote harmonization of local images and global images by our proposed AIC-Net. As Fig. 5 shows, DeepIH [1] can only achieve a slight translation and cannot generate any potential shadow, thus resulting in unrealistic harmonized images. The branched generation in AIC-Net disentangles the style transfer and shadow generation, leading to much more realistic image harmonization. The global discriminator in AIC-Net works better by driving the generator to remove the boundary artifacts. AdaIN [56] transfer the style of the foreground region to be harmonious with the background, but they do not handle shadow generation which leads to unrealistic composition. Besides, AdaIN is unable to preserve the texture quality of harmonized image due to the missing of guided filter.

To validate that the proposed model is able to generate various shadow adaptively, we tested the trained model in different environments as shown in Fig. 6. The first two rows show the backgrounds and local images respectively, the last row shows locally harmonized images and globally harmonized images produced by our model. As shown in the first three columns of Fig. 6, the proposed model generates shadows conditioned on the existed shadow cases in the backgrounds. For background images without clearly shadow case as shown in the last two columns of Fig. 6, the proposed model is still able to synthesize realistic shadows conditioned on the feature extracted by the illumination model.



**Fig. 6.** Shadow generation under different conditions by AIC-Net: The first row shows the background image with the local region for foreground embedding. The second row show the direct composition. Rows 3-4 show the local harmonization and global harmonization by AIC-Net, respectively. AIC-Net can generate various shadows adaptively conditioned on illumination, viewpoint and weather of the background image, etc.

**Ablation Study:** We conducted an ablation study to evaluate the effectiveness of different technical designs in AIC-Net. As shown in Table 2, four variants of AIC-Net are designed including: 1) AIC-Net (WF) denotes AIC-Net without the guided feature filter; 2) AIC-Net (WB) denotes AIC-Net without using the branched generation mechanism; 3) AIC-Net (WS) denotes AIC-Net without incorporating the spatial transformer module; 4) AIC-Net(WI) denotes AIC-Net without illumination model. The evaluations are performed with three metrics including FID, MS and AMT study that involves 10 users. For AMT user study, a set of images (5 generated by AIC-Net and its four variants) are presented to each of 10 users who will choose the most realistic image. We presented 25 sets of harmonized images and the AMT score is derived by averaging the percentages of the most realistic images as judged by the 10 users. The scores of car and pedestrian experiments are averaged for each metric to obtain the final scores.

As Table 2 shows, AIC-Net (WF) obtains clearly lower scores with all metrics as compared with the standard AIC-Net which shows the importance of guided feature filter in generating realistic image harmonization. AIC-Net (WB) also obtains lower scores than AIC-Net as the branched generation in AIC-Net disentangles the shadow generation and style transfer to produce more realistic local harmonization as shown in Fig. 3. For AIC-Net (WS), the scores for the global images are much lower than the other models, demonstrating the importance of the spatial transformer module which bridges the local and global harmonization to produce globally realistic harmonization. On the other hand, AIC-Net (WS) achieves similar scores which can be expected as both models

**Table 2.** We conduct ablation study with three metrics (FID, AMT, MS) to evaluate the quality of harmonized local and global images. AIC-Net (WF), AIC-Net (WB), AIC-Net (WS) and AIC-Net(WI) denote AIC-Net without the guided feature filter, without branched generation mechanism, without spatial transformer module and without illumination model respectively. For AMT user study, the percentage numbers represent how often the harmonized images in each category are classified as *most realistic* by Turkers. We average the FID, AMT and MS scores of car and pedestrian images as the final scores.

Methods	FID		AMT		MS	
	Local	Global	Local	Global	Local	Global
<b>AIC-Net(WF)</b>	115.5	106.1	14%	19%	0.72	0.70
<b>AIC-Net(WB)</b>	109.2	93.7	19%	20%	0.70	0.69
<b>AIC-Net(WS)</b>	98.5	92.1	21%	16%	0.73	0.71
<b>AIC-Net(WI)</b>	94.2	91.8	17%	18%	0.67	0.65
<b>AIC-Net</b>	90.8	81.8	29%	27%	0.63	<b>0.62</b>

are almost the same in local harmonization. AIC-Net (WI) present clearly drop of AMT score compared with standard AIC-Net, indicating the significance of illumination model for the generation of realistic shadow.

## 5 Conclusion

This paper presents AIC-Net, an end-to-end trainable network that harmonizes foreground objects and background images with adaptive and realistic shadow effects. A branched generation mechanism is designed to disentangle style transfer and shadow generation to produce high-fidelity local harmonization. An auxiliary illumination model is collaborated in shadow branch to achieve adaptive shadow generation. A spatial transformer module is introduced to bridge the local and global harmonization for joint optimization. The quantitative and qualitative experiments demonstrate the superior harmonization performance and wide adaptability of the proposed method. We will continue to study AIC-Net for image harmonization with shadow generation within more complicated illumination environments.

## References

1. Tsai, Y.H., Shen, X., Lin, Z., Sunkavalli, K., Lu, X., Yang, M.H.: Deep image harmonization. In: CVPR. (2017)
2. Luan, F., Paris, S., Shechtman, E., Bala, K.: Deep painterly harmonization. arXiv preprint arXiv:1804.03189 (2018)
3. Wu, H., Zheng, S., Zhang, J., Huang, K.: Gp-gan: Towards realistic high-resolution image blending. arXiv:1703.07195 (2017)
4. Lin, C.H., Yumer, E., Wang, O., Shechtman, E., Lucey, S.: St-gan: Spatial transformer generative adversarial networks for image compositing. In: CVPR. (2018)

5. Zhan, F., Lu, S., Xue, C.: Verisimilar image synthesis for accurate detection and recognition of texts in scenes. In: ECCV. (2018) 249–266
6. Zhan, F., Huang, J., Lu, S.: Adaptive composition gan towards realistic image synthesis. 1905.04693 (2019)
7. Zhan, F., Xue, C., Lu, S.: Ga-dan: Geometry-aware domain adaptation network for scene text detection and recognition. In: ICCV. (2019)
8. Zhan, F., Zhu, H., Lu, S.: Scene text synthesis for efficient and effective deep network training. arXiv:1901.09193 (2019)
9. Zhan, F., Zhu, H., Lu, S.: Spatial fusion gan for image synthesis. In: CVPR. (2019)
10. Zhan, F., Lu, S.: Esir: End-to-end scene text recognition via iterative image rectification. In: CVPR. (2019)
11. Zhan, F., Lu, S., Zhang, C., Ma, F., Xie, X.: Towards realistic 3d embedding via view alignment. arXiv preprint arXiv:2007.07066 (2020)
12. Zhan, F., Lu, S., Xiao, A.: Spatial-aware gan for unsupervised person re-identification. arXiv preprint arXiv:1911.11312 (2019)
13. Liu, D., Long, C., Zhang, H., Yu, H., Dong, X., Xiao, C.: Arshadowgan: Shadow generative adversarial network for augmented reality in single light scenes. In: CVPR. (2020)
14. Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. TOG **22** (2003)
15. Azadi, S., Pathak, D., Ebrahimi, S., Darrell, T.: Compositional gan: Learning conditional image composition. arXiv:1807.07560 (2018)
16. Zhu, J.Y., Krähenbühl, P., Shechtman, E., Efros, A.A.: Learning a discriminative model for the perception of realism in composite images. In: ICCV. (2015)
17. Chen, B.C., Kae, A.: Toward realistic image compositing with adversarial learning. In: CVPR. (2019)
18. Sunkavalli, K., Johnson, M.K., Matusik, W., Pfister, H.: Multi-scale image harmonization. ACM Transactions on Graphics (Proc. ACM SIGGRAPH) **29** (2010)
19. Xue, S., Agarwala, A., Dorsey, J., Rushmeier, H.: Understanding and improving the realism of image composites. TOG **31** (2012)
20. Tao, M.W., Johnson, M.K., Paris, S.: Error-tolerant image compositing. IJCV **103** (2013)
21. Tsai, Y.H., Shen, X., Lin, Z., Sunkavalli, K., Yang, M.H.: Sky is not the limit: Semantic-aware sky replacement. ACM Transactions on Graphics (Proc. SIGGRAPH) **35** (2016)
22. Efros, J.F.L.A.A.: Using color compatibility for assessing image realism. In: ICCV. (2007)
23. Cong, W., Zhang, J., Niu, L., Liu, L., Ling, Z., Li, W., Zhang, L.: Dovenet: Deep image harmonization via domain verification. In: CVPR. (2020)
24. Lalonde, J.F., Efros, A.A., Narasimhan, S.G.: Estimating the natural illumination conditions from a single outdoor image. IJCV (2012)
25. Gardner, M.A., Sunkavalli, K., Yumer, E., Shen, X., Gambaretto, E., Gagné, C., Lalonde, J.F.: Learning to predict indoor illumination from a single image. In: SIGGRAPH Asia. (2017)
26. Hold-Geoffroy, Y., Sunkavalli, K., Hadap, S., Gambaretto, E., Lalonde, J.F.: Deep outdoor illumination estimation. In: CVPR. (2017)
27. Gardner, M.A., Hold-Geoffroy, Y., Sunkavalli, K., Gagné, C., Lalonde, J.F.: Deep parametric indoor lighting estimation. In: ICCV. (2019)
28. Garon, M., Sunkavalli, K., Hadap, S., Carr, N., Lalonde, J.F.: Fast spatially-varying indoor lighting estimation. In: CVPR. (2019)

29. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. In: NIPS. (2014) 2672–2680
30. Denton, E., Chintala, S., Szlam, A., Fergus, R.: Deep generative image models using a laplacian pyramid of adversarial networks. In: NIPS. (2015)
31. Lee, D., Liu, S., Gu, J., Liu, M.Y., Yang, M.H., Kautz, J.: Context-aware synthesis and placement of object instances. In: NIPS. (2018)
32. Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., Abbeel, P.: InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In: NIPS. (2016)
33. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: ICML. (2017)
34. Bau, D., Zhu, J.Y., Strobel, H., Zhou, B., Tenenbaum, J.B., Freeman, W.T., Torralba, A.: Gan dissection: Visualizing and understanding generative adversarial networks. In: ICLR. (2019)
35. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV. (2017)
36. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR. (2017)
37. Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: CVPR. (2017)
38. Zhu, J.Y., Zhang, R., Pathak, D., Darrell, T., Efros, A.A., Wang, O., Shechtman, E.: Toward multimodal image-to-image translation. In: NIPS. (2017)
39. Azadi, S., Fisher, M., Kim, V., Wang, Z., Shechtman, E., Darrell, T.: Multi-content gan for few-shot font style transfer. In: CVPR. (2018)
40. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: CVPR. (2019)
41. Liu, M.Y., Huang, X., Mallya, A., Karras, T., Aila, T., Lehtinen, J., Kautz, J.: Few-shot unsupervised image-to-image translation. arXiv:1905.01723 (2019)
42. Green, R.: Spherical harmonic lighting: The gritty details. In: Game Developers Conference. (2003)
43. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. (2016)
44. Cheng, D., Shi, J., Chen, Y., Deng, X., Zhang, X.: Learning scene illumination by pairwise photos from rear and front mobile cameras. Computer Graphics Forum (2018)
45. He, K., Sun, J., Tang, X.: Guided image filtering. TPAMI (2013)
46. Hirzer, M., Beznai, C., Roth, P.M., Bischof, H.: Person re-identification by descriptive and discriminative classification. In: SCIA. (2011)
47. Chavdarova, T., Baqué, P., Bouquet, S., Maksai, A., Jose, C., Bagautdinov, T., Lettry, L., Fua, P., Gool, L.V., Fleuret, F.: Wildtrack: A multi-camera hd dataset for dense unscripted pedestrian detection. In: CVPR. (2018)
48. Wang, L., Shi, J., Song, G., fan Shen, I.: Object detection combining recognition and segmentation. In: ACCV. (2007)
49. Wen, L., Du, D., Cai, Z., Lei, Z., Chang, M.C., Qi, H., Lim, J., Yang, M.H., Lyu, S.: Ua-detrag: A new benchmark and protocol for multi-object detection and tracking. arXiv preprint arXiv:1511.04136 (2015)
50. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR. (2016)



51. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
52. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NIPS. (2017)
53. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR. (2015)
54. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. (2009)
55. Zhou, P., Han, X., Morariu, V.I., Davis, L.S.: Learning rich features for image manipulation detection. In: CVPR. (2018)
56. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV. (2017)