

# SHUNIT: Style Harmonization for Unpaired Image-to-Image Translation

Seokbeom Song<sup>1</sup>, Suhyeon Lee<sup>1</sup>, Hongje Seong<sup>1</sup>, Kyoungwon Min<sup>2</sup>, Euntai Kim<sup>1\*</sup>

<sup>1</sup>Yonsei University, Seoul, Korea

<sup>2</sup>Korea Electronics Technology Institute, Seongnam, Korea  
{lgs5751, hyeon93, hjseong, etkim}@yonsei.ac.kr, minkw@keti.re.kr

## Abstract

We propose a novel solution for unpaired image-to-image (I2I) translation. To translate complex images with a wide range of objects to a different domain, recent approaches often use the object annotations to perform per-class source-to-target style mapping. However, there remains a point for us to exploit in the I2I. An object in each class consists of multiple components, and all the sub-object components have different characteristics. For example, a car in CAR class consists of a car body, tires, windows and head and tail lamps, etc., and they should be handled separately for realistic I2I translation. The simplest solution to the problem will be to use more detailed annotations with sub-object component annotations than the simple object annotations, but it is not possible. The key idea of this paper is to bypass the sub-object component annotations by leveraging the original style of the input image because the original style will include the information about the characteristics of the sub-object components. Specifically, for each pixel, we use not only the per-class style gap between the source and target domains but also the pixel's original style to determine the target style of a pixel. To this end, we present Style Harmonization for unpaired I2I translation (SHUNIT). Our SHUNIT generates a new style by harmonizing the target domain style retrieved from a class memory and an original source image style. Instead of direct source-to-target style mapping, we aim for source *and* target styles harmonization. We validate our method with extensive experiments and achieve state-of-the-art performance on the latest benchmark sets. The source code is available online: <https://github.com/bluejangbaljang/SHUNIT>.

## Introduction

Unpaired image-to-image (I2I) translation aims to learn source-to-target style mapping, where source and target images are unpaired. It can be applied to data augmentation (Antoniou, Storkey, and Edwards 2017; Mariani et al. 2018; Huang et al. 2018a; Xie et al. 2020), domain adaptation (Hoffman et al. 2018; Murez et al. 2018) and various image editing applications, such as style transfer (Gatys, Ecker, and Bethge 2016; Huang and Belongie 2017; Ulyanov, Vedaldi, and Lempitsky 2017), colorization (Zhang, Isola,

and Efros 2016; Zhang et al. 2017), and image inpainting (Iizuka, Simo-Serra, and Ishikawa 2017; Pathak et al. 2016).

In the I2I translation, the biggest problem is how to deal with the style variations among objects or classes. In other words, when a global style gap is applied to an entire image as in Fig. 1a, the I2I translation often results in unrealistic images because each class has different the style gaps between the source domain and target domain. Recent advanced methods (Shen et al. 2019; Bhattacharjee et al. 2020; Jeong et al. 2021; Kim et al. 2022) addressed the problem by leveraging additional object annotations. They simplify the task into class-level I2I translation and then perform per-class source-to-target style mapping. This enables the networks to explicitly estimate class-wise target styles, but it has a critical limitation: *An object in each class consists of multiple components, and all the sub-object components might also have different characteristics.* Let us consider the example given in Fig. 1b.

Understandably, when a road image taken on a sunny day is translated into a night image, head and tail lamps in a car should be brighter while the rest of the components, such as car body, tires, and windows, should be darker than before. Therefore, each component in a car should be handled separately for realistic I2I translation. However, if the previous approaches are applied to perform per-class source-to-target style mapping, they will translate all head and tail lamps into white lights, making unrealistic images, as shown in Fig. 1b. Here, one might think that this issue can be addressed by annotating more detailed sub-object components than the simple object categories, but it is actually impossible. A brief example is as follows. A car consists of body, window, and tires. A tire consists of wheel and gum. In this way, sub-object components can be divided endlessly.

To solve the above limitation of the previous class-level I2I methods, we present Style Harmonization for unpaired I2I translation (SHUNIT). The key idea of SHUNIT is to bypass the sub-object component annotations by leveraging the original style of the input image because the original style will include the information about the characteristics of the sub-object components. Thus, instead of mapping source-to-target style directly, SHUNIT harmonizes the source *and* target styles to realize realistic and practical I2I translation. As illustrated in Fig. 1c, SHUNIT uses not only the per-

\*Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

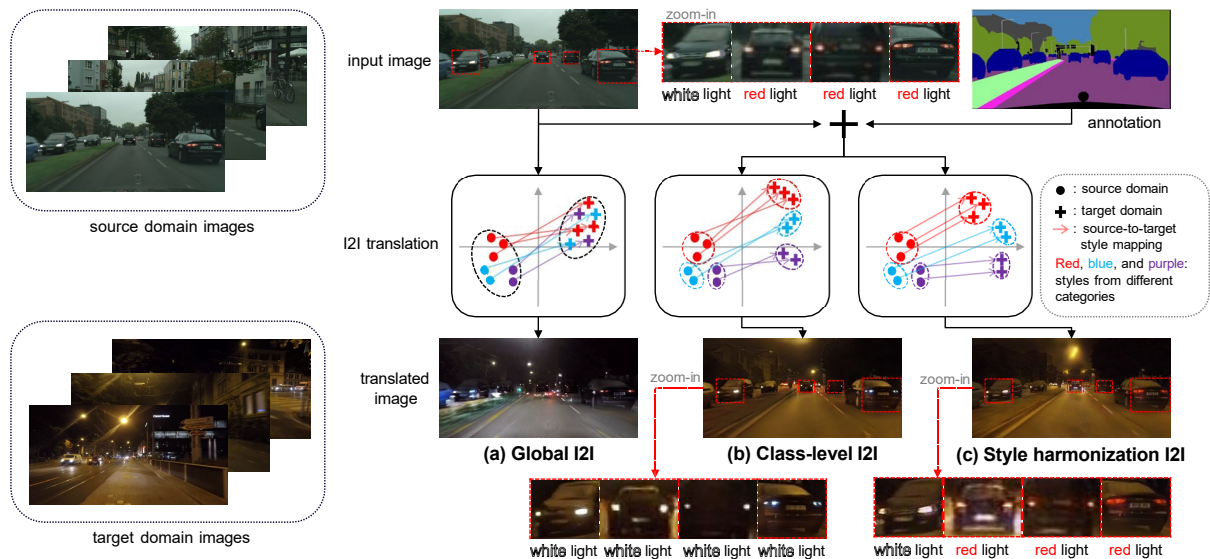


Figure 1: Illustration of the concepts in unpaired I2I translation. The results are obtained on Cityscapes  $\rightarrow$  ACDC (night) setting. In the image, many head and tail lamps should be bright at night. (a) Global I2I (Huang et al. 2018b) converts all classes to bright because it translates the image with a single source-to-target style mapping function. (b) Class-level I2I (Jeong et al. 2021) leverages additional annotations to address the problem of (a) and performs per-class source-to-target style mapping. It can effectively deal with multiple classes in an image, but loses the original style: All white lights and red lights become white lights. (c) Style harmonization I2I also performs class-wise style mapping, while adaptively preserving the original styles.

class style gap between the source and target domains but also the pixel’s original style to determine the target style of a pixel. To achieve this, we disentangle the target style into class-aware memory style and image-specific component style. The class-aware memory style is stored in a style memory, and image-specific component style is taken from the original input image.

The goal of the style memory is to obtain class-wise source-to-target style gaps and it is motivated by (Jeong et al. 2021). Compared to the memory in (Jeong et al. 2021), our style memory differs in two aspects. First, the output from the style memory was used alone as a target style in (Jeong et al. 2021), but the output from the memory is adaptively aggregated (=harmonized) in this paper with the style of the original input image to make a target style. Second, the memory was simply updated in (Jeong et al. 2021), whereas our style memory is jointly trained and optimized with the other parts of SHUNIT. Specifically, the class-aware memory in (Jeong et al. 2021) was not trained but simply was updated using the input features during the training, memorizing the style features from the target domain. The gradient was not propagated to the style memory. Thus, the style memory in (Jeong et al. 2021) cannot update their parameters based on the error of memory. In SHUNIT, however, we overcome this problem by enabling the memory to learn through backpropagation. To this end, we train the style memory from randomly initialized parameters and introduce style contrastive loss to constrain the memory to learn class-wise style representations. The backpropagation forces the style memory to reduce the final loss jointly and effectively along with the other parts of SHUNIT. To demonstrate the

superiority of our SHUNIT, we conduct extensive experiments on the latest benchmark sets and achieve state-of-the-art performance.

Overall, the contributions of our work are summarized as follows:

- We present a novel challenge in I2I translation: an object might have various styles.
- We propose a new I2I method, style harmonization, that leverages two distinct styles: class-aware memory style and image-specific component style. To the best of our knowledge, the style harmonization is the first method to estimate the target style in multiple perspectives for unpaired I2I translation.
- We achieve new state-of-the-art performance on latest benchmarks and provide extensive experimental results with analysis.

## Related Work

**Image-to-image translation.** The goal of I2I is to learn source-to-target style mapping. For I2I translation, pix2pix (Isola et al. 2017) proposes a general solution using conditional generative adversarial networks (Mirza and Osindero 2014). However, it has a significant limitation: paired training data should be used for training networks. CycleGAN (Zhu et al. 2017) successfully addresses this problem with a cycle consistency loss. The loss allows us to train the networks with unpaired training data by supervising the reconstructed original image only. Based on CycleGAN, many approaches (Kim et al. 2017; Choi et al. 2018)

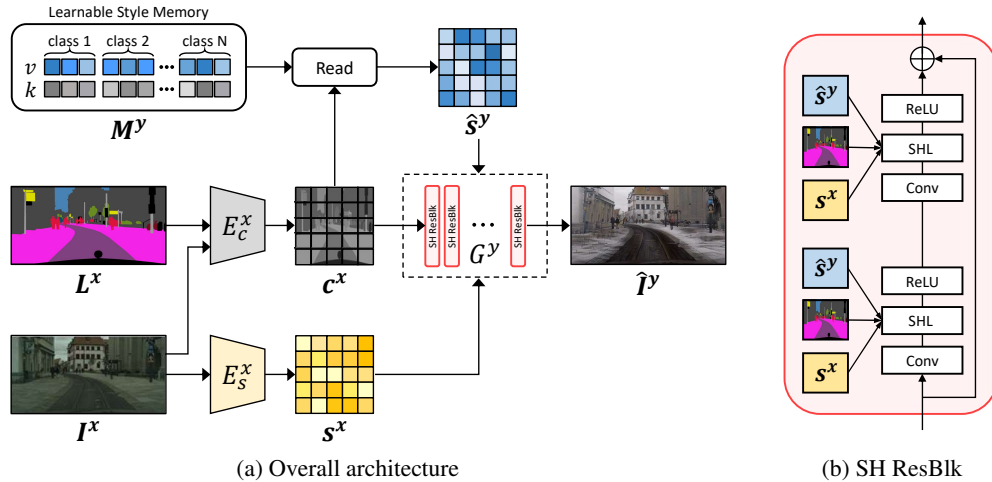


Figure 2: An overview of SHUNIT. (a) From a pair of image  $I^x$  and label  $L^x$  in source domain  $\mathcal{X}$ , the two encoders (*i.e.*,  $E_c^x$  and  $E_s^x$ ) extract the content  $c^x$  and component style  $s^x$ , respectively. The content retrieves the memory style  $\hat{s}^y$ . The three features  $c^x$ ,  $s^x$ , and  $\hat{s}^y$  are fed to the target generator  $G^y$ , which consists of several Style Harmonization Residual Blocks (SH ResBlk). The generator  $G^y$  outputs the translated image  $\hat{I}^y$ . (b) The input of the SH ResBlk is the content  $c^x$  for the first layer, and the output of the previous block is used as input for the remainders. Each SH ResBlk includes two Style Harmonization Layers (SHL) that transfer target styles, *i.e.*,  $s^x$  and  $\hat{s}^y$ .

have been proposed to tackle I2I translation take in an unpaired manner. UNIT (Liu, Breuel, and Kautz 2017) proposes another unpaired I2I translation solution by mapping two images in different domains into the same latent code in a shared-latent space. MUNIT (Huang et al. 2018b) and DRIT (Lee et al. 2018) introduce a disentangled representation to achieve diverse and multi-modal I2I translation from unpaired data. Basically, they perform global I2I translation which focuses on mapping a global style on all pixels in an image. Although they work well on object-centric images, they bring severe artifacts for complex images, such as multiple objects being presented or large domain gap scenarios, as illustrated in Fig. 1a. To complement the problem, recent approaches leverage additional object annotations and perform class-level I2I translation.

**Class-level image-to-image translation.** Recent several approaches (Mo, Cho, and Shin 2019; Shen et al. 2019; Bhattacharjee et al. 2020; Jeong et al. 2021; Kim et al. 2022) propose class-level image-to-image translation solutions with object annotations. Specifically, INIT (Shen et al. 2019) generates instance-wise target domain images. DUNIT (Bhattacharjee et al. 2020) additionally employs an object detection network and jointly trains it with the I2I translation network. MGUIT (Jeong et al. 2021) proposes an approach to store and read class-wise style representations with key-value memory networks (Miller et al. 2016). This approach, however, cannot directly supervise the memory with objective functions for I2I translation. Instaformer (Kim et al. 2022) proposes a transformer-based (Vaswani et al. 2017; Dosovitskiy et al. 2021) architecture that mixes instance-aware content and style representations. The existing methods that leverage object annotations learn the direct class-wise source-to-target style map-

ping, as shown in Fig. 1b. This effectively simplifies the I2I translation problem into per-class I2I translation, but they overlook an important point that not all pixels in the same class should be translated with the same style. Our approach, style harmonization, addresses this problem by introducing the component style that facilitates preserving the original style of the source image, as illustrated in Fig. 1c.

## Proposed Method

### Definition and Overview

Let  $\mathcal{X}$  and  $\mathcal{Y}$  be the visual source and target domains, respectively. Given an image and the corresponding label (=bounding box or segmentation mask) in  $\mathcal{X}$  domain, our framework generates a new image in  $\mathcal{Y}$  domain while remaining the semantic information in the given image. We assume that each domain consists of images and labels denoted by  $(I^x, L^x) \in \mathcal{X}$  and  $(I^y, L^y) \in \mathcal{Y}$ , and both domains have the same set of  $N$  classes. Our framework contains the source encoder  $E^x = \{E_c^x, E_s^x\}$ , target generator  $G^y$ , and target style memory  $M^y$  for source-to-target mapping, and the target encoder  $E^y = \{E_c^y, E_s^y\}$ , source generator  $G^x$ , source style memory  $M^x$  for target-to-source mapping. For convenience, we will only describe the source-to-target direction, and the overview of our framework is depicted in Fig. 2.

Following the previous studies (Huang et al. 2018b; Lee et al. 2018), we assume that an image can be disentangled into domain-invariant content and domain-specific style. For this, we basically follow the MUNIT (Huang et al. 2018b) architecture. The content encoder  $E_c^x$  consists of several strided convolutional layers and residual blocks (He et al. 2016), and all the convolutional layers are followed by Instance Normalization (Ulyanov, Vedaldi, and Lempitsky 2016). The content encoder extracts the domain-invariant




Figure 3: Read operation in style memory. The content feature  $c^x$  is separated by the label  $L^x$ . For each class, the memory read is independently performed. After class-wise memory read, the memory style  $\hat{s}^y$  is obtained by gathering the retrieved class-wise values into the original locations.  $\otimes$  denotes matrix multiplication.

content feature  $c^x$  from the image  $I^x$  and label  $L^x$ . The style encoder  $E_s^x$  also consists of several strided convolutional layers and residual blocks, and it extracts the *component style* feature  $s^x$  from the image  $I^x$ . The *memory style*  $\hat{s}^y$  is read by retrieving from the learnable style memory  $M^y$  to the content feature  $c^x$ . The generator  $G^y$  consists of several style harmonization layers and residual blocks, and it produces the translated image  $\hat{I}^y$  from  $c^x$ ,  $s^x$ , and  $\hat{s}^y$ .

### Style Harmonization for Unpaired Image-to-Image Translation (SHUNIT)

The important point of SHUNIT is that two styles are employed to determine the target style: One is image-specific *component style*, and the other one is class-aware *memory style*. We focus on extracting two distinct styles accurately and then harmonizing them. In what follows, we describe the detail of each step.

**Component style.** The style encoder  $E_s^x$  takes the image  $I^x$  as input and extracts the style feature  $s^x$  of size  $H \times W \times C$ , where  $H$ ,  $W$ , and  $C$  are the height, width, and number of channels of the feature, respectively. Here,  $s^x$  explicitly represents the style of the input image, thus we use it as image-specific *component style*. The component style is used together with the memory style in the style harmonization layer to reduce the artifacts of the generated target image.

**Memory style.** The component style is not sufficient to handle complex scenes with multiple objects. Therefore, we exploit the class-specific style that leverages an object annotation. To this end, we construct the class-wise style memory retrieved by the content feature. The target style memory  $M^y$  consists of  $N$  class memories to store class-wise style representations of the target domain  $\mathcal{Y}$ . Each class memory  $M_n^y$  has  $U$  key-value pairs  $(k^y, v^y)$  (Jeong et al. 2021), considering that various styles exist in one class (e.g., different styles of headlamp and tires in CAR class). The key  $k^y$  is used for matching with the content feature and the value  $v^y$  has the class-aware style representations. The key and value

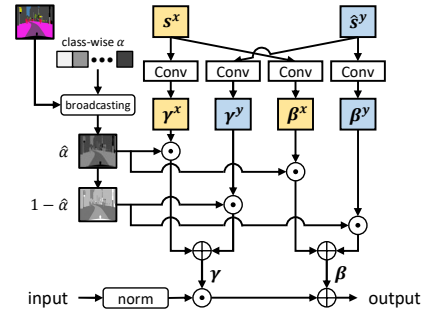


Figure 4: Style harmonization layer. From the component style  $s^x$  and memory style  $\hat{s}^y$ , scale ( $\gamma^x$  and  $\gamma^y$ ) and shift ( $\beta^x$  and  $\beta^y$ ) factors are extracted via four convolutional layers. The alpha mask  $\hat{\alpha}$  is obtained by broadcasting the class-wise alpha  $\alpha$  with the semantic label  $L^x$ . We weighted sum the image and memory styles with the alpha mask, and transfer it to the input feature.  $\odot$  denotes Hadamard product.

are learnable vectors, each one of size  $1 \times 1 \times C$ .

Fig. 3 shows a detailed implementation of the process of reading the corresponding *memory style*  $\hat{s}^y$  from the memory  $M^y$ . With the semantic label, we separate the content feature  $c^x$  into  $\{c_1^x, c_2^x, \dots, c_N^x\}$ , where  $c_n^x$  denotes source content feature for the  $n$ -th class. Let  $c_{n,i}^x$  be the  $i$ -th pixel of the  $n$ -th class source content feature and  $(k_{n,j}^y, v_{n,j}^y)$  be the  $j$ -th key-value pair of the  $n$ -th class target style memory  $M_n^y$ . In this work, we aim to read the target memory style  $\hat{s}_{n,i}^y$  corresponding to the source content  $c_{n,i}^x$  using the similarity between the source content and key of target memory. To this end, we calculate the similarity  $w_{n,i,j}^x$  between  $c_{n,i}^x$  and  $k_{n,j}^y$  as:

$$w_{n,i,j}^x = \frac{\exp(d(c_{n,i}^x, k_{n,j}^y))}{\sum_{u=1}^U \exp(d(c_{n,i}^x, k_{n,u}^y))} \quad (1)$$

where  $d(\cdot, \cdot)$  is the cosine similarity. We then read the memory style  $\hat{s}_{n,i}^y$  corresponding to  $c_{n,i}^x$  by calculating the weighted sum of values in  $n$ -th class style memory:

$$\hat{s}_{n,i}^y = \sum_{j=1}^U w_{n,i,j}^x v_{n,j}^y. \quad (2)$$

The same process is applied for content features of other classes, finally extracting spatially varying target memory style features  $\hat{s}^y$  of size  $H \times W \times C$ .

Different from the previous key-value memory networks (Jeong et al. 2021) that learn the memory via updating mechanism, we learn the memory through backpropagation. The updating mechanism is used to directly store the external input features. However, it has a critical drawback: The memory cannot be trained with the network jointly with the same objective function because the gradient should be stopped at the updated memory. To solve the problem, we discard the update mechanism and learn the memory with the loss functions presented in Eq. (7). The effectiveness of our memory learning strategy is validated in the experiments section.

**Style harmonization layer.** Our goal is harmonizing the source *and* target styles instead of mapping source-to-target style directly. To this end, we propose the style harmonization layer to adaptively aggregate the component style and memory style. The style harmonization layer consists of several convolution layers and class-wise alpha parameters, and it is illustrated in Fig. 4. Here we use three conditional inputs: memory style, component style, and label. Convolutional layers are used to compute pixel-wise scale  $\gamma$  and shift  $\beta$  factors from the two styles. Following (Jiang et al. 2020; Park et al. 2019; Zhu et al. 2020; Ling et al. 2021), we transfer the harmonized target style by scaling and shifting the normalized input feature with the computed factors (*i.e.*,  $\gamma$  and  $\beta$ ). In the layer, we additionally set class-wise alpha parameters  $(\alpha_1, \alpha_2, \dots, \alpha_N)$ . It is used to decide which style has more influence for each class in the generated images. If the alpha value is large, the component style  $s^x$  has more influence, and vice versa.

Let  $f$ , of size  $H \times W \times C$ , be the input feature of the current style harmonization layer in the generator  $G^y$ . With the style harmonizing scale  $\gamma$  and shift  $\beta$  factors, the feature is denormalized by

$$\gamma_{c,h,w} \frac{(f_{c,h,w} - \mu_c)}{\sigma_c} + \beta_{c,h,w} \quad (3)$$

where  $\mu_c$  and  $\sigma_c$  are the mean and standard deviation of the input feature  $f$  at the channel  $c$ , respectively. The modulation parameters  $\gamma_{c,h,w}$  and  $\beta_{c,h,w}$  are obtained from  $\gamma_{c,h,w}^x$ ,  $\gamma_{c,h,w}^y$ ,  $\beta_{c,h,w}^x$ , and  $\beta_{c,h,w}^y$ , which are the scale and shift factors of the component style  $s^x$  and memory style,  $s^y$  respectively, and they are computed by

$$\begin{aligned} \gamma_{c,h,w} &= \hat{\alpha}_{h,w} \gamma_{c,h,w}^x + (1 - \hat{\alpha}_{h,w}) \gamma_{c,h,w}^y, \\ \beta_{c,h,w} &= \hat{\alpha}_{h,w} \beta_{c,h,w}^x + (1 - \hat{\alpha}_{h,w}) \beta_{c,h,w}^y, \end{aligned} \quad (4)$$

where  $\hat{\alpha}$  denotes the alpha mask. It is obtained by broadcasting the class-wise alpha parameters to their corresponding semantic regions of the label  $L^x$ . We experimentally demonstrate that our style harmonization layer adaptively controls the style of each object well, and the results are given in the experiments section.

## Loss Functions

We leverage standard loss functions used in MUNIT (Huang et al. 2018b) to generate proper target domain images. It includes self-reconstruction  $L_{self}$  (Zhu et al. 2017), cycle consistency  $L_{cycle}$  (Zhu et al. 2017), perceptual  $L_{perc}$  (Johnson, Alahi, and Fei-Fei 2016) and adversarial loss  $L_{adv}$  (Goodfellow et al. 2014). The detailed explanations of those loss functions are given in the supplementary material.

In this paper, we propose two advanced loss functions to facilitate style harmonization: content contrastive loss and style contrastive loss. It is used with the aforementioned standard loss functions jointly. In what follows, we introduce the proposed two loss functions.

**Content contrastive loss.** To extract domain invariant content features from the content encoder, MUNIT (Huang et al. 2018b) simply reduces the L1 distances between  $c^x$  and

$c^y$ . We replace this with contrastive representation learning to improve discrimination within a class. For a content feature  $\hat{c}_i^y$  at pixel  $i$ , which is the content feature extracted from translated target image, we set the positive sample to  $c_i^x$  and we set the remaining features at the other pixels as negative samples. The content contrastive loss is defined with the form of InfoNCE (Oord, Li, and Vinyals 2018) as:

$$\mathcal{L}_{content} = - \sum_{i=1}^{HW} \log \left( \frac{\exp((c_i^x \cdot \hat{c}_i^y)/\tau)}{\sum_{j=1}^{HW} \exp((c_j^x \cdot \hat{c}_i^y)/\tau)} \right) \quad (5)$$

where  $\tau$  is a temperature parameter. In this equation, the features in the same class at the pixel  $i$  can be considered as negative samples. This encourages the content encoder to extract more diverse style representations from the style memory within the same class. It is also applied to the target-to-source pipeline with  $\hat{c}^x$  and  $c^y$

**Style constrative loss.** We propose the style contrastive loss to allow the style memory to learn class-wise style representations. Similar to the content contrastive loss, for a style feature  $\hat{s}_i^x$  at pixel  $i$ , which is the memory style of target-to-source mapping, we set the positive sample to the source component style  $s_i^x$  and we set the remaining features at the other pixels as negative samples. The style constrative loss is defined as follows:

$$\mathcal{L}_{style} = - \sum_{i=1}^{HW} \log \left( \frac{\exp((s_i^x \cdot \hat{s}_i^x)/\tau)}{\sum_{j=1}^{HW} \exp((s_j^x \cdot \hat{s}_i^x)/\tau)} \right) \quad (6)$$

This loss directly supervises the memory style  $\hat{s}^x$  for the translated target image. This improves the stability of cycle consistency learning for reconstructing  $I^x$  from  $\hat{I}^y$ . It is also applied to the target styles, *i.e.*,  $\hat{s}^y$  and  $s^y$

Finally, all loss functions are summarized as follows:

$$\begin{aligned} \min_{(E^x, E^y, G^x, G^y, D^x, D^y)} \max_{(D^x, D^y)} \mathcal{L}(E^x, E^y, G^x, G^y, D^x, D^y) = \\ \lambda_{self} \mathcal{L}_{self} + \lambda_{cycle} \mathcal{L}_{cycle} + \lambda_{perc} \mathcal{L}_{perc} \\ + \lambda_{adv} \mathcal{L}_{adv} + \lambda_{content} \mathcal{L}_{content} + \lambda_{style} \mathcal{L}_{style} \end{aligned} \quad (7)$$

where  $D^x$  and  $D^y$  denote the multi-scale discriminators (Wang et al. 2018) for each visual domain,  $\mathcal{X}$  and  $\mathcal{Y}$ . The details of  $\mathcal{L}_{self}$ ,  $\mathcal{L}_{cycle}$ ,  $\mathcal{L}_{perc}$ , and  $\mathcal{L}_{adv}$  are described in the supplementary material.

## Experiments

In this section, we present extensive experimental results and analysis. To demonstrate the superiority of our method, we compare our SHUNIT with state-of-the-art I2I translation methods. The implementation details of our method are provided in the supplementary material.

### Datasets

We evaluate our SHUNIT on three I2I translation scenarios: Cityscapes (Cordts et al. 2016)  $\rightarrow$  ACDC (Sakaridis, Dai, and Van Gool 2021) and INIT (Shen et al. 2019), and KITTI (Geiger et al. 2013)  $\rightarrow$  Cityscapes (Cordts et al. 2016). In all scenarios, INIT (Shen et al. 2019), DUNIT (Bhattacharjee et al. 2020), MGUIT (Jeong et al. 2021), InstaFormer (Kim et al. 2022), and our method use semantic labels provided in each dataset.



Figure 5: Qualitative comparison on Cityscapes (clear)  $\rightarrow$  ACDC (snow/rain/fog/night). From the given clear image (first column), we generate four adverse condition images using (Zhu et al. 2017; Huang et al. 2018b; Jeong et al. 2021) and SHUNIT. In the last column, we show a sample of the real image for each adverse condition.

	clear $\rightarrow$ snow		clear $\rightarrow$ rain		clear $\rightarrow$ fog		clear $\rightarrow$ night	
	cFID $\downarrow$	mIoU $\uparrow$	cFID $\downarrow$	mIoU $\uparrow$	cFID $\downarrow$	mIoU $\uparrow$	cFID $\downarrow$	mIoU $\uparrow$
CycleGAN (Zhu et al. 2017)	21.88	23.68	20.16	35.96	31.72	13.73	15.26	31.33
UNIT (Liu, Breuel, and Kautz 2017)	13.89	31.24	16.25	39.39	29.36	28.70	12.28	35.29
MUNIT (Huang et al. 2018b)	13.79	33.83	12.62	44.20	29.34	27.44	12.56	<b>37.43</b>
TSIT (Jiang et al. 2020)	10.47	<u>38.08</u>	14.16	<u>46.40</u>	25.16	<u>36.68</u>	<b>11.62</b>	<u>35.92</u>
MGUIT (Jeong et al. 2021)	<u>8.75</u>	33.33	<u>10.76</u>	42.60	<u>24.36</u>	10.22	15.83	31.36
SHUNIT (ours)	<b>6.62</b>	<b>45.15</b>	<b>8.47</b>	<b>48.84</b>	<b>6.53</b>	<b>38.96</b>	14.08	33.66

Table 1: Quantitative comparison on Cityscapes  $\rightarrow$  ACDC. We measure class-wise FID (lower is better) and mIoU (higher is better). For brevity, class-wise FID is written as cFID.

	clear $\rightarrow$ snow	clear $\rightarrow$ rain	clear $\rightarrow$ fog	clear $\rightarrow$ night
AdaptSegNet (Tsai et al. 2018)	35.3	49.0	31.8	29.7
ADVENT (Vu et al. 2019)	32.1	44.3	32.9	31.7
BDL (Li, Yuan, and Vasconcelos 2019)	36.4	<u>49.7</u>	37.7	<u>33.8</u>
CLAN (Luo et al. 2019)	37.7	44.0	39.0	31.6
FDA (Yang and Soatto 2020)	<b>46.9</b>	<b>53.3</b>	<b>39.5</b>	<b>37.1</b>
SIM (Wang et al. 2020)	33.3	44.5	36.6	28.0
MRNet (Zheng and Yang 2021)	38.7	45.4	38.8	27.9
SHUNIT (ours)	<u>45.2</u>	48.8	<u>39.0</u>	33.7

Table 2: Quantitative Comparison on domain adaptation for semantic segmentation. We report mIoU for Cityscapes  $\rightarrow$  ACDC.

**Cityscapes  $\rightarrow$  ACDC** Cityscapes (Cordts et al. 2016) is one of the most popular urban scene dataset. ACDC (Sakaridis, Dai, and Van Gool 2021) is the latest dataset with multiple adverse condition images and consists of four conditions of street scenes: snow, rain, fog, and night. ACDC dataset provides images with corresponding dense pixel-level semantic annotations, and it has 19 classes the same as Cityscapes dataset for all adverse conditions. Following (Sakaridis, Dai, and Van Gool 2021), we leverage Cityscapes dataset as a clear condition and translate it to the adverse conditions (*i.e.*, snow, rain, fog, and night) in ACDC dataset. Therefore, this scenario is challenging because not

only the weather conditions, but also layouts, such as camera model, view, and angle, are different. To train the networks, 2975, 400, 400, 400, and 400 images are used for clear, snow, rain, fog, and night conditions, respectively. For a fair comparison on this benchmark, we reproduce existing state-of-the-art methods (Zhu et al. 2017; Liu, Breuel, and Kautz 2017; Huang et al. 2018b; Jiang et al. 2020; Jeong et al. 2021) in our system. For fair comparison, we set the number of key-value pairs for style memory to be the same as our setting and use segmentation mask for reproducing (Jeong et al. 2021).

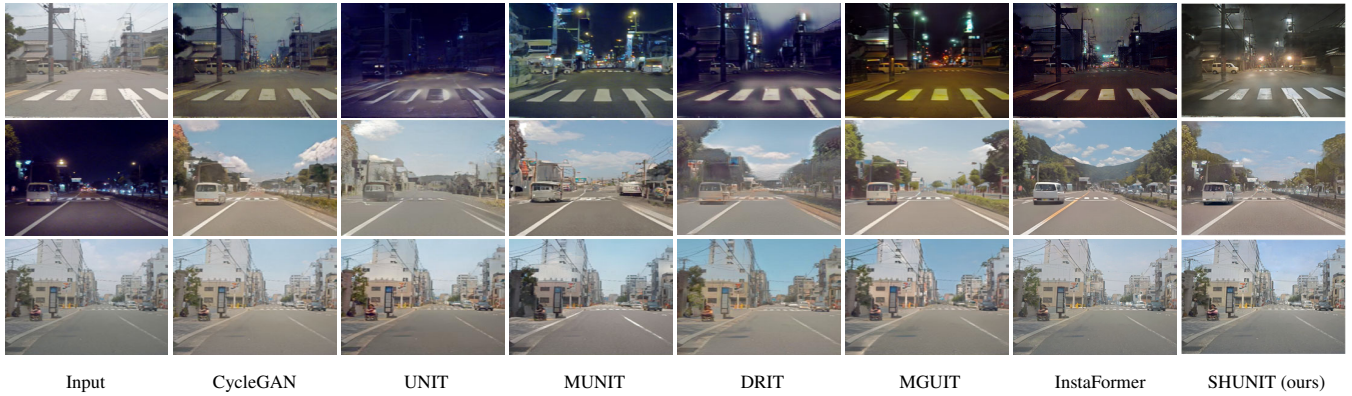


Figure 6: Qualitative comparison on INIT dataset. (Top to bottom) sunny  $\rightarrow$  night, night  $\rightarrow$  sunny, cloudy  $\rightarrow$  sunny results. Our method preserves object details and looks more realistic.

	sunny $\rightarrow$ night		night $\rightarrow$ sunny		sunny $\rightarrow$ rainy		sunny $\rightarrow$ cloudy		cloudy $\rightarrow$ sunny		Average	
	CIS $\uparrow$	IS $\uparrow$	CIS $\uparrow$	IS $\uparrow$	CIS $\uparrow$	IS $\uparrow$	CIS $\uparrow$	IS $\uparrow$	CIS $\uparrow$	IS $\uparrow$	CIS $\uparrow$	IS $\uparrow$
CycleGAN (Zhu et al. 2017)	0.014	1.026	0.012	1.023	0.011	1.073	0.014	1.097	0.090	1.033	0.025	1.057
UNIT (Liu, Breuel, and Kautz 2017)	0.082	1.030	0.027	1.024	0.097	1.075	0.081	1.134	0.219	1.046	0.087	1.055
MUNIT (Huang et al. 2018b)	1.159	1.278	1.036	1.051	1.012	1.146	1.008	1.095	1.026	1.321	1.032	1.166
DRIT (Lee et al. 2018)	1.058	1.224	1.024	1.099	1.007	1.207	1.025	1.104	1.046	1.321	1.031	1.164
INIT (Shen et al. 2019)	1.060	1.118	1.045	1.080	1.036	1.152	1.040	1.142	1.016	1.460	1.043	1.179
DUNIT (Bhattacharjee et al. 2020)	1.166	1.259	1.083	1.108	1.029	1.225	1.033	1.149	1.077	1.472	1.079	1.223
MGUIT (Jeong et al. 2021)	1.176	1.271	<u>1.115</u>	<u>1.130</u>	1.092	1.213	1.052	1.218	<u>1.136</u>	<u>1.489</u>	1.112	1.254
Instaformer (Kim et al. 2022)	<u>1.200</u>	<u>1.404</u>	<u>1.115</u>	1.127	<b>1.158</b>	<u>1.394</u>	<b>1.130</b>	<u>1.257</u>	<b>1.141</b>	<b>1.585</b>	1.149	1.353
SHUNIT (ours)	<b>1.205</b>	<b>1.503</b>	<b>1.308</b>	<b>1.585</b>	<u>1.136</u>	<b>1.609</b>	<u>1.111</u>	<b>1.405</b>	1.085	1.315	<b>1.169</b>	<b>1.483</b>

Table 3: Quantitative Comparison on INIT dataset. We measure CIS and IS (higher is better).

**INIT** INIT (Shen et al. 2019) is a public benchmark set for I2I translation. It contains street scenes images including 4 weather categories (*i.e.*, sunny, night, rainy, and cloudy) with the corresponding bounding box labels. Following (Shen et al. 2019), we split the 155K images into 85% for training and 15% for testing. We conduct five translation experiments: sunny  $\leftrightarrow$  night, sunny  $\leftrightarrow$  cloudy, sunny  $\rightarrow$  rainy. In this dataset, we directly copied the results of the existing methods from (Shen et al. 2019; Bhattacharjee et al. 2020; Jeong et al. 2021; Kim et al. 2022). Similarly, for fair comparison with MGUIT, the number of key-value pairs in style memory is set equally.

**KITTI  $\rightarrow$  Cityscapes** KITTI is a public benchmark set for object detection. It contains 7481 images with bounding boxes annotations for training and 7518 images for testing. Following the previous I2I translation methods (Bhattacharjee et al. 2020; Jeong et al. 2021; Kim et al. 2022), we select the common 4 object classes (person, car, truck, bicycle) for evaluation.

### Qualitative Comparison

Fig. 5 shows qualitative results on Cityscapes  $\rightarrow$  ACDC. Since our I2I translation setting, Cityscapes  $\rightarrow$  ACDC, is very challenging as discussed in the datasets section, existing methods cannot generate realistic images in several scenarios. Specifically, CycleGAN (Zhu et al. 2017) often

destroys the semantic layout. MUNIT (Huang et al. 2018b) translates images with a global style, thus it also often generates artifacts, as shown in the snow, rain, and fog images. MGUIT (Jeong et al. 2021) also includes artifacts in the car even though leveraging memory style. It shows the limitation of the updating mechanism for training memory style, and the limitation is clearly depicted in the challenging scenario. In contrast to them, our SHUNIT accurately generates images in the target domains without losing the original style in the input image. In the supplementary material, we further provide the results of UNIT (Liu, Breuel, and Kautz 2017) and TSIT (Jiang et al. 2020).

As shown in Fig. 6, which depicts qualitative results on INIT dataset, our method generates high-quality images in various scenarios. In the night  $\rightarrow$  sunny scenario (second row), InstaFormer (Kim et al. 2022) translates the color of the road lane to yellow. On the other hand, our method keeps the color of the lane as white and generates a sunny scene by harmonizing the target domain style retrieved from a style memory and an image style.

### Quantitative Comparison

The quantitative results on Cityscapes  $\rightarrow$  ACDC are presented in Table 1. To quantify the per-class image-to-image translation quality, we measure class-wise FID (Shim et al. 2022). We further measure mIoU on ACDC test set. The mIoU metric is used to validate the results on the practical

	Pers.	Car	Truc.	Bic.	mAP
DT (Inoue et al. 2018)	28.5	40.7	25.9	29.7	31.2
DAF (Huang et al. 2018b)	39.2	40.2	25.7	48.9	38.5
DARL (Kim et al. 2019)	46.4	58.7	27.0	49.1	45.3
DAOD (Rodriguez and Mikolajczyk 2019)	47.3	59.1	28.3	49.6	46.1
DUNIT (Bhattacharjee et al. 2020)	<u>60.7</u>	65.1	32.7	<u>57.7</u>	54.1
MGUIT (Jeong et al. 2021)	58.3	68.2	33.4	<b>58.4</b>	54.6
InstaFormer (Kim et al. 2022)	<b>61.8</b>	<u>69.5</u>	<u>35.3</u>	55.3	<u>55.5</u>
SHUNIT (ours)	56.3	<b>74.4</b>	<b>51.9</b>	53.2	<b>59.0</b>

Table 4: Quantitative Comparison on domain adaptation for object detection. We report per-class AP for KITTI  $\rightarrow$  Cityscapes.

Mem.	Comp.	$\alpha$	clear $\rightarrow$ snow		clear $\rightarrow$ rain		$\mathcal{L}_{content}$	$\mathcal{L}_{style}$	clear $\rightarrow$ snow		clear $\rightarrow$ rain	
			cFID $\downarrow$	mIoU $\uparrow$	cFID $\downarrow$	mIoU $\uparrow$			cFID $\downarrow$	mIoU $\uparrow$	cFID $\downarrow$	mIoU $\uparrow$
✓			12.93	26.33	<b>7.83</b>	39.06	✓		11.06	32.43	13.23	38.39
✓	✓		8.72	37.20	8.47	38.97		✓	12.38	30.89	9.33	39.47
✓	✓	✓	<b>6.62</b>	<b>44.35</b>	8.47	<b>42.77</b>	✓	✓	<b>6.62</b>	<b>44.35</b>	<b>8.47</b>	<b>42.77</b>

(a) Style ablation study on SHL.

(b) Ablation study on loss functions.

			clear $\rightarrow$ snow		clear $\rightarrow$ rain					clear $\rightarrow$ snow		clear $\rightarrow$ rain		
			cFID $\downarrow$	mIoU $\uparrow$	cFID $\downarrow$	mIoU $\uparrow$				cFID $\downarrow$	mIoU $\uparrow$	cFID $\downarrow$	mIoU $\uparrow$	
Updating	10.55	36.70	12.05	38.21	L1	12.21	38.98	<b>7.72</b>	39.78	w/o	6.73	44.11	9.03	38.72
Backprop.	<b>6.62</b>	<b>44.35</b>	<b>8.47</b>	<b>42.77</b>	Contrastive	<b>6.62</b>	<b>44.35</b>	8.47	<b>42.77</b>	w/	<b>6.62</b>	<b>44.35</b>	<b>8.47</b>	<b>42.77</b>

(c) Style memory training strategies.

(d) L1 vs. Contrastive loss for style and content losses.

(e) Experimental results of label input for content encoder.

Table 5: Ablation Study. We report class-wise FID and mIoU in two scenarios: clear  $\rightarrow$  {snow, rain}.

problem, semantic segmentation. We generate a training set by Cityscapes  $\rightarrow$  ACDC and then train DeepLabV2 (Chen et al. 2017) on it. The mIoU score is obtained with the trained DeepLabV2 by evaluating on ACDC test set. As shown in Table 1, we surpass the state-of-the-art I2I translation methods by a significant margin in most scenarios, demonstrating the superiority of our style harmonization for unpaired I2I translation. Table 2 shows the quantitative results of domain adaptation for semantic segmentation with the state-of-the-art methods. Despite we trained DeepLabV2 with only a simple cross-entropy loss, our SHUNIT achieves comparable performance with the domain adaptation methods that were trained DeepLabV2 with additional loss functions and several techniques for boosting the performance of mIoU score on the target domain.

Table 3 shows another quantitative results on the testing split of INIT (Shen et al. 2019). To directly compare our method with the public results, we evaluate our method with Inception Score (IS) (Salimans et al. 2016) and Conditional Inception Score (CIS) (Huang et al. 2018b). As shown in Table 3, we achieve the best performance in most scenarios.

We further evaluate our method on domain adaptation benchmark following DUNIT (Bhattacharjee et al. 2020). We use Faster-RCNN (Ren et al. 2015) trained on the source domain as a detector. As shown in Table 4, we achieve the state-of-the-art performance.

## Ablation Study

In this section, we study the effectiveness of each component in our method. We validate on Cityscapes (clear)  $\rightarrow$  ACDC (snow/rain) scenarios and use ACDC validation set for both class-wise FID and mIoU<sup>1</sup>.

**Ablation study on style harmonization layer.** We ablate the memory style, component style, and class-wise  $\alpha$  in the style harmonization layer, and they are denoted as ‘‘Mem.’’, ‘‘Comp.’’, and ‘‘ $\alpha$ ’’ in Table 5a, respectively. As shown in the table, the memory style-only is far behind the full model. With component style, we can achieve performance improvement on clear  $\rightarrow$  snow while decreasing on clear  $\rightarrow$  rain. We obtain significant improvement on most scenarios with class-wise  $\alpha$ . The results demonstrate that the existing approach, which only leverages the memory style, is not sufficient for I2I translation, and we successfully address the problem by adaptively harmonizing two styles.

**Ablation study on content and style losses.** We study the effectiveness of the proposed two loss functions,  $\mathcal{L}_{style}$  and  $\mathcal{L}_{content}$ , by ablating them step-by-step, and the results are given in Table 5b. As shown in the table, our model is effective when two losses are used jointly.

<sup>1</sup>mIoU on test set should be evaluated on the online server (Sakaridis, Dai, and Van Gool 2021) and it has a limit on the number of submissions. Therefore, we use validation set for ablation study.



**Style memory training strategy.** As described in the proposed method section, we opt for backpropagation to train the style memory rather than updating the mechanism used in (Jeong et al. 2021). The results are shown in Table 5c. We surpass the existing updating method by a large margin.

**L1 vs. Contrastive loss.** Table 5d shows the efficacy of our contrastive-based approach by replacing  $L_{content}$  and  $L_{style}$  with L1 losses as used in MUNIT (Huang et al. 2018b). The L1 losses are designed to reduce L1 distances within positive pairs without consideration of negative pairs. As we discussed in loss functions section, our method effectively encourages extracting more diverse style representations, leading to performance improvement.

**Label input for content encoder.** Table 5e shows that label input is not significant but always leads to performance improvements. Therefore, we have no reason to omit the label input.

## Limitations

Since our framework leverages the component style of the source image, the generated image’s quality relies on the source image’s quality. If the source image has a very bright colors, SHUNIT often generates a relatively bright night image. In Fig. 5, SHUNIT struggles to generate geometrically distinct lights from the source image. Due to the above problems, SHUNIT cannot achieve the best performance on clear  $\rightarrow$  night scenario in Table. 1. We believe that these problems can be alleviated by leveraging geometric information such as depth or camera pose. Additionally, our approach can give limited benefit to some I2I translation scenarios, such as dog  $\rightarrow$  cat, because these tasks need to change the content; however, we tackle the unpaired I2I translation task under the condition that the content will not be changed, and only the style will be changed.

## Conclusion

We present a new perspective of the target style: It can be disentangled into class-aware and image-specific styles. Furthermore, our SHUNIT effectively harmonizes the two styles, and its superiority is demonstrated through extensive experiments. We believe that our proposal has the potential to break new ground in style-based image editing applications such as style transfer, colorization, and image inpainting.

## Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2021-0-00800, Development of Driving Environment Data Transformation and Data Verification Technology for the Mutual Utilization of Self-driving Learning Data for Different Vehicles).

## References

- Antoniou, A.; Storkey, A.; and Edwards, H. 2017. Data augmentation generative adversarial networks. *arXiv preprint arXiv:1711.04340*.
- Bhattacharjee, D.; Kim, S.; Vizier, G.; and Salzmann, M. 2020. Dunit: Detection-based unsupervised image-to-image translation. In *CVPR*, 4787–4796.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2017. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4): 834–848.
- Choi, Y.; Choi, M.; Kim, M.; Ha, J.-W.; Kim, S.; and Choo, J. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, 8789–8797.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 3213–3223.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Gatys, L. A.; Ecker, A. S.; and Bethge, M. 2016. Image style transfer using convolutional neural networks. In *CVPR*, 2414–2423.
- Geiger, A.; Lenz, P.; Stiller, C.; and Urtasun, R. 2013. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11): 1231–1237.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *NeurIPS*, 27.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hoffman, J.; Tzeng, E.; Park, T.; Zhu, J.-Y.; Isola, P.; Saenko, K.; Efros, A.; and Darrell, T. 2018. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 1989–1998. PMLR.
- Huang, S.-W.; Lin, C.-T.; Chen, S.-P.; Wu, Y.-Y.; Hsu, P.-H.; and Lai, S.-H. 2018a. Auggan: Cross domain adaptation with gan-based data augmentation. In *ECCV*, 718–731.
- Huang, X.; and Belongie, S. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 1501–1510.
- Huang, X.; Liu, M.-Y.; Belongie, S.; and Kautz, J. 2018b. Multimodal unsupervised image-to-image translation. In *ECCV*, 172–189.
- Iizuka, S.; Simo-Serra, E.; and Ishikawa, H. 2017. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4): 1–14.
- Inoue, N.; Furuta, R.; Yamasaki, T.; and Aizawa, K. 2018. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5001–5009.

- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *CVPR*, 1125–1134.
- Jeong, S.; Kim, Y.; Lee, E.; and Sohn, K. 2021. Memory-guided unsupervised image-to-image translation. In *CVPR*, 6558–6567.
- Jiang, L.; Zhang, C.; Huang, M.; Liu, C.; Shi, J.; and Loy, C. C. 2020. Tsit: A simple and versatile framework for image-to-image translation. In *ECCV*, 206–222. Springer.
- Johnson, J.; Alahi, A.; and Fei-Fei, L. 2016. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 694–711. Springer.
- Kim, S.; Baek, J.; Park, J.; Kim, G.; and Kim, S. 2022. InstaFormer: Instance-Aware Image-to-Image Translation with Transformer. In *CVPR*.
- Kim, T.; Cha, M.; Kim, H.; Lee, J. K.; and Kim, J. 2017. Learning to discover cross-domain relations with generative adversarial networks. In *ICML*, 1857–1865. PMLR.
- Kim, T.; Jeong, M.; Kim, S.; Choi, S.; and Kim, C. 2019. Diversify and match: A domain adaptive representation learning paradigm for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12456–12465.
- Lee, H.-Y.; Tseng, H.-Y.; Huang, J.-B.; Singh, M.; and Yang, M.-H. 2018. Diverse image-to-image translation via disentangled representations. In *ECCV*, 35–51.
- Li, Y.; Yuan, L.; and Vasconcelos, N. 2019. Bidirectional learning for domain adaptation of semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6936–6945.
- Ling, J.; Xue, H.; Song, L.; Xie, R.; and Gu, X. 2021. Region-aware adaptive instance normalization for image harmonization. In *CVPR*, 9361–9370.
- Liu, M.-Y.; Breuel, T.; and Kautz, J. 2017. Unsupervised image-to-image translation networks. *NeurIPS*, 30.
- Luo, Y.; Zheng, L.; Guan, T.; Yu, J.; and Yang, Y. 2019. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2507–2516.
- Mariani, G.; Scheidegger, F.; Istrate, R.; Bekas, C.; and Malossi, C. 2018. Bagan: Data augmentation with balancing gan. *arXiv preprint arXiv:1803.09655*.
- Miller, A. H.; Fisch, A.; Dodge, J.; Karimi, A.-H.; Bordes, A.; and Weston, J. 2016. Key-Value Memory Networks for Directly Reading Documents. In *EMNLP*.
- Mirza, M.; and Osindero, S. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
- Mo, S.; Cho, M.; and Shin, J. 2019. InstaGAN: Instance-aware Image-to-Image Translation. In *ICLR*.
- Murez, Z.; Kolouri, S.; Kriegman, D.; Ramamoorthi, R.; and Kim, K. 2018. Image to image translation for domain adaptation. In *CVPR*, 4500–4509.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Park, T.; Liu, M.-Y.; Wang, T.-C.; and Zhu, J.-Y. 2019. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2337–2346.
- Pathak, D.; Krahenbuhl, P.; Donahue, J.; Darrell, T.; and Efros, A. A. 2016. Context encoders: Feature learning by inpainting. In *CVPR*, 2536–2544.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Rodriguez, A. L.; and Mikolajczyk, K. 2019. Domain adaptation for object detection via style consistency. *arXiv preprint arXiv:1911.10033*.
- Sakaridis, C.; Dai, D.; and Van Gool, L. 2021. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *ICCV*, 10765–10775.
- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. *NeurIPS*, 29.
- Shen, Z.; Huang, M.; Shi, J.; Xue, X.; and Huang, T. S. 2019. Towards instance-level image-to-image translation. In *CVPR*, 3683–3692.
- Shim, S.-H.; Hyun, S.; Bae, D.; and Heo, J.-P. 2022. Local Attention Pyramid for Scene Image Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7774–7782.
- Tsai, Y.-H.; Hung, W.-C.; Schulter, S.; Sohn, K.; Yang, M.-H.; and Chandraker, M. 2018. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7472–7481.
- Ulyanov, D.; Vedaldi, A.; and Lempitsky, V. 2016. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*.
- Ulyanov, D.; Vedaldi, A.; and Lempitsky, V. 2017. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *CVPR*, 6924–6932.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *NeurIPS*, 30.
- Vu, T.-H.; Jain, H.; Bucher, M.; Cord, M.; and Pérez, P. 2019. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2517–2526.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. In *CVPR*, 8798–8807.
- Wang, Z.; Yu, M.; Wei, Y.; Feris, R.; Xiong, J.; Hwu, W.-m.; Huang, T. S.; and Shi, H. 2020. Differential treatment for stuff and things: A simple unsupervised domain adaptation method for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12635–12644.

- Xie, X.; Chen, J.; Li, Y.; Shen, L.; Ma, K.; and Zheng, Y. 2020. Self-supervised cyclegan for object-preserving image-to-image domain adaptation. In *ECCV*, 498–513. Springer.
- Yang, Y.; and Soatto, S. 2020. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4085–4095.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *ECCV*, 649–666. Springer.
- Zhang, R.; Zhu, J.-Y.; Isola, P.; Geng, X.; Lin, A. S.; Yu, T.; and Efros, A. A. 2017. Real-time user-guided image colorization with learned deep priors. In *SIGGRAPH*.
- Zheng, Z.; and Yang, Y. 2021. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *International Journal of Computer Vision*, 129(4): 1106–1120.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2223–2232.
- Zhu, P.; Abdal, R.; Qin, Y.; and Wonka, P. 2020. Sean: Image synthesis with semantic region-adaptive normalization. In *CVPR*, 5104–5113.