

Style-Hallucinated Dual Consistency Learning for Domain Generalized Semantic Segmentation

Yuyang Zhao¹, Zhun Zhong², Na Zhao¹, Nicu Sebe², and Gim Hee Lee¹

¹ National University of Singapore

² University of Trento

<https://github.com/HeliosZhao/SHADE>

Abstract. In this paper, we study the task of synthetic-to-real domain generalized semantic segmentation, which aims to learn a model that is robust to unseen real-world scenes using only synthetic data. The large domain shift between synthetic and real-world data, including the limited source environmental variations and the large distribution gap between synthetic and real-world data, significantly hinders the model performance on unseen real-world scenes. In this work, we propose the **Style-HALLUCINATED DUAL** consistency learning (**SHADE**) framework to handle such domain shift. Specifically, SHADE is constructed based on two consistency constraints, Style Consistency (SC) and Retrospection Consistency (RC). SC enriches the source situations and encourages the model to learn consistent representation across style-diversified samples. RC leverages real-world knowledge to prevent the model from overfitting to synthetic data and thus largely keeps the representation consistent between the synthetic and real-world models. Furthermore, we present a novel style hallucination module (SHM) to generate style-diversified samples that are essential to consistency learning. SHM selects basis styles from the source distribution, enabling the model to dynamically generate diverse and realistic samples during training. Experiments show that our SHADE yields significant improvement and outperforms state-of-the-art methods by 5.05% and 8.35% on the average mIoU of three real-world datasets on single- and multi-source settings, respectively.

1 Introduction

Semantic segmentation, *i.e.*, predicting a semantic category for each pixel, plays a crucial role in autonomous driving. With modern deep neural networks [19,20] and abundant annotated data [5,6], fully-supervised methods [2,21,39] have achieved remarkable success on many public datasets. Nonetheless, annotating each pixel for a high-resolution image is expensive and time-consuming. For example, it takes more than 1.5 hours to annotate a 1024×2048 driving scene [5], and the time is even doubled for scenes under adverse weather [32] and poor lighting conditions [31]. To alleviate the heavy annotation cost, unsupervised domain adaptive semantic segmentation (UDA-Seg) [35,46,47] has been introduced to learn models that can perform well on a target set, given the labeled (source) data and unlabeled (target) data. Considering that synthetic data can

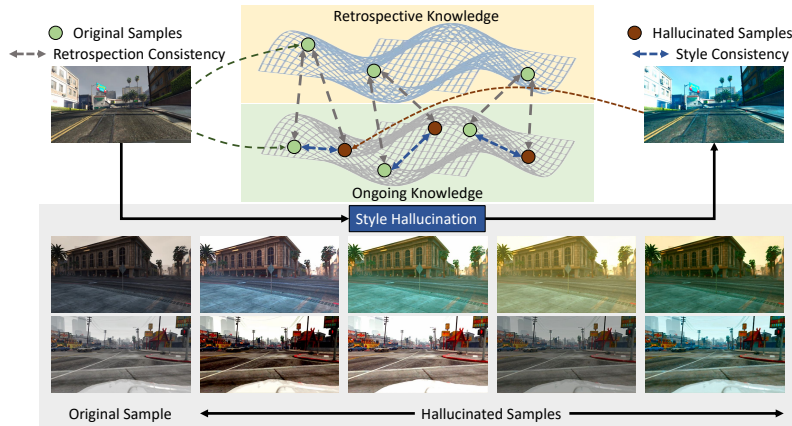


Fig. 1. Illustration of dual consistency constraints and examples of style hallucination. We generate hallucinated samples (brown circle) from the style hallucination module and then utilize the paired samples and real-world (retrospective) knowledge to learn style consistency (blue dash line) and retrospection consistency (gray dash line).

be automatically generated and annotated by a pre-designed engine [29,30], existing works of UDA-Seg commonly use synthetic data as the source domain and real-world data as the target domain.

Despite the success of UDA-Seg, the limitations are still severe. First, the adapted model achieves remarkable performance on the target data, but still degrades when facing unseen out-of-distribution samples. Second, it is almost impractical to collect diverse enough real-world data that can cover all the conditions. For example, a model trained with Singapore scenes cannot well address the snowy road in Switzerland. To mitigate such limitations, a more practical but challenging setting, domain generalized semantic segmentation (DG-Seg) [4,41,43], is introduced in the community. DG-Seg only leverages annotated source data to train a robust model that can cope with different unseen conditions. Similar to UDA-Seg, synthetic data are commonly used as the source data to release expensive annotation cost in DG-Seg. In this paper, we aim to solve the synthetic-to-real DG task in semantic segmentation.

The main challenge for synthetic-to-real DG-Seg is to cope with the significant domain shift between source and unseen target domains, which can be roughly divided into two aspects. First, the environmental variations in source data are very limited compared to those of unseen target data. Second, there exists large distribution gap between synthetic and real-world data, *e.g.*, image styles and characteristics of objects. To learn the domain-invariant model that can address the domain shift, previous works design tailor-made modules [4,26] to remove domain-specific information, or leverage extra real-world data to transfer synthetic data [14,41] to real-world styles for narrowing the distribution gap. However, the removal of domain-specific information [4,26] is not complete and

explicit due to the lack of real-world information; the real-world style transfer [14,41] heavily relies on extra data, which are not always available in practice, and ignores the invariant representation within the source domain. Taking the above into account, in this paper, we aim to explicitly learn domain-invariant representation with the stylized samples in the source domain and bridge the gap between synthetic and real-world data without using extra real-world data.

To this end, we introduce a novel dual consistency learning framework that can jointly achieve the above two goals. As shown in Fig. 1, we introduce two consistency constraints, *style consistency* (SC) and *retrospection consistency* (RC), which explicitly encourage the model to learn style invariant representation and to be less overfitting to the synthetic data respectively. Specifically, we first diversify the samples in SC by style variation, and then enforce the consistency between them. To obtain the diverse samples, we adopt the style features, *i.e.*, channel-wise mean and standard deviation, to generate new data. Compared with directly transferring the whole image (*e.g.*, CycleGAN [48]), changing style features can maintain the pixel alignment to the utmost extent. This forces the model to learn pixel-level consistency between paired samples. In addition, our RC lies in the guidance of real-world information. We leverage ImageNet [6] pre-trained model which is available in all the DG-Seg models. The pre-trained model contains general knowledge of classifying some “things” class, *e.g.*, bicycle, bus and car. Furthermore, the features of these classes in the pre-trained model reflect the representation in the context of real-world. Consequently, such features can serve as the guidance for the ongoing model to retrospect what the real-world looks like and to lead the model less overfitting to the synthetic data.

Style diversifying is crucial for the success of dual consistency learning. Previous works [33,44] commonly mix or swap styles within the source domain, which will generate more samples of the dominant styles (*e.g.*, daytime in GTAV [29]). Nevertheless, it is not the best way since the target styles may be quite different from the dominant styles. To fully take the advantage of all the source styles, we propose style hallucination module (SHM), which leverages C basis styles to represent the style space of C dimension and thus to generate new styles. Ideally, the basis styles should be linearly independent so the linear combination of basis styles can represent all the source styles. However, many unrealistic styles that impair the model training are generated when we directly take C orthogonal unit vectors as the basis. To reconcile diversity and realism, we use farthest point sampling (FPS) [28] to select C styles from all the source styles as basis styles. Such basis styles contain many rare styles since rare styles are commonly far away from the dominant ones. With these basis styles that represent the style space in a better way, we utilize linear combination to generate new styles. To summarize, we propose the **Style-HAllucinated Dual consistENcy** learning (**SHADE**) framework for domain generalization in the context of semantic segmentation, and our contributions are as follows:

- We propose the dual consistency constraints for domain generalized semantic segmentation, which learn the style invariant representation among diversi-

fied styles and narrow the domain gap between synthetic and real domains by retrospecting the real-world knowledge.

- We propose the style hallucination module to generate new and diverse styles with the representative basis styles, enhancing the dual consistency learning.
- Experiments on single-source and multi-source DG benchmark demonstrate the effectiveness of the proposed method and we achieve new state-of-the-art performance.

2 Related Work

Domain Generalization. To tackle the performance degradation in the out-of-distribution conditions, domain generalization [14,41,4,23,44] is introduced to learn a robust model with one or multiple source domains, aiming to perform well on unseen domains. Considering the expensive annotation cost in semantic segmentation, synthetic data are commonly adopted as the source domain in recent DG-Seg works. The large domain shift and the restricting conditions in the training data greatly limit the performance on the unseen real-world test data. To address these problems, one main stream of previous works [14,41] focuses on diversifying training data with real-world templates and learning the invariant representation from all the domains. Another main stream aims at directly learning the explicit domain-invariant features within the source domain [4,26,33]. IBN-Net [26] and ISW [4] leverage tailor-made instance normalization block and whitening transformation to reduce the impact of domain-specific information. Different from previous methods, we generate new styles only with the synthetic source domain and learn the invariant representation across styles.

Consistency learning (CL). CL is adopted by many computer vision tasks and settings. One main stream is leveraging CL to exploit the unlabeled samples in semi-supervised learning [8,34] and unsupervised domain adaptation [25,49,45] since the consistency between two view samples can be used for pseudo-labeling. CL is also applied to address the corruptions and perturbations [12,36] by maximizing the similarity between different augmentations. In addition, CL is also used in self-supervised learning [3,10] as the contrastive loss to utilize totally unlabeled data. We introduce CL into domain generalization, leading the model robust to various styles. We also leverage consistency with real-world knowledge to narrow the domain gap between synthetic and real-world data.

Style Variation. Style features are widely explored in style transfer [7,16], which aims at changing the image style but maintaining the content. Inspired by this, recent domain generalization methods leverage the style features to generate diverse data of different styles to improve the generalization ability. Swapping [33,42] and mixing [44] existing styles within the source domains is an effective way and generating new styles [38] by specially designed modules can also make sense. We also only leverage the styles within the source domain but take the relatively rare styles in the source domains into account, thus generating more diverse samples to improve generalization ability.

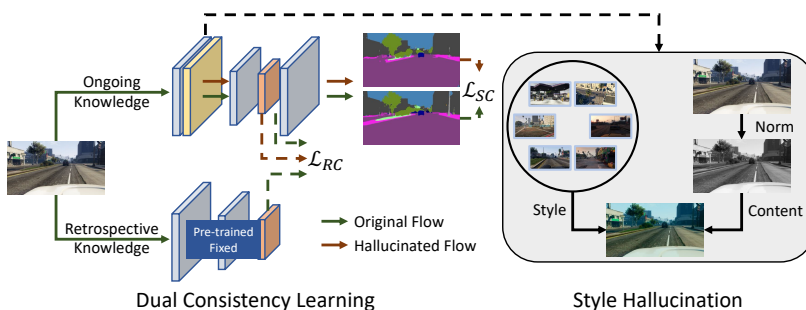


Fig. 2. Overview of the proposed style-hallucinated dual consistency learning (SHADE) framework. Training images are fed into the segmentation model (ongoing knowledge) and the ImageNet pre-trained model (retrospective knowledge). The style hallucination module is inserted into a certain layer of the segmentation model to generate stylized samples. Finally, the model is optimized by the dual consistency losses: style consistency \mathcal{L}_{SC} and retrospection consistency \mathcal{L}_{RC} . Note that the cross entropy loss is also used and we omit it here for brevity.

3 Methodology

Preliminary. In synthetic-to-real domain generalized semantic segmentation, one or multiple labeled source domains $\mathcal{S} = \{x_S^i, y_S^i\}_{i=1}^{N_S}$, where N_S is the number of source domains, are used to train a segmentation model, which is deployed to unseen real-world target domains \mathcal{T} directly. In general, the source and target domains share the same label space $Y_S, Y_T \in (1, N_C)$ but belong to different distributions. The goal of this task is to improve the generalization ability of model in unseen domains using only the source data.

Overview. To solve the above challenging problem, we propose the **Style-Hallucinated Dual consistency learning (SHADE)** framework, which is equipped with dual consistency constraints and a Style Hallucination Module (SHM). The dual consistency constraints effectively learn the domain-invariant representation and reduce the gap between the real and synthetic data. SHM enriches the training samples by dynamically generating diverse styles, which catalyzes the advantage of dual consistency learning. The overall framework is shown in Fig. 2.

3.1 Dual Consistency Constraints

In SHADE, we introduce two consistency learning constraints: (1) Style Consistency (SC) that aims at learning the consistent predictions across stylized samples. (2) Retrospection Consistency (RC) that focuses on narrowing the distribution shift between synthetic and real-world data in the feature-level.

Style Consistency (SC). Different from traditional cross entropy constraint focusing on the high-level semantic information, logit pairing [12,17] has demonstrated its effectiveness in learning adversarial samples by highlighting the most

invariant information. Inspired by this, we propose SC to ameliorate the style shift with logit pairing. To fulfill this, SC requires the style-diversified samples \tilde{x}_S which share the same semantic content with the source samples x_S but of different styles. Due to the pixel-level segmentation requirement, it is better to obtain the style-diversified samples \tilde{x}_S by non-geometry style variation, *e.g.*, MixStyle [44], CrossNorm [33] and the proposed style hallucination in Sec. 3.2. Formally, we minimize the Jensen-Shannon Divergence (JSD) between the posterior probability p of the semantically aligned \tilde{x}_S and x_S :

$$\begin{aligned}\mathcal{L}_{SC}(x_S, \tilde{x}_S) &= JSD(p(x_S); p(\tilde{x}_S)) \\ &= \frac{1}{2} (D[p(x_S)||Q] + D[p(\tilde{x}_S)||Q]),\end{aligned}\tag{1}$$

where $Q = (p(x_S) + p(\tilde{x}_S))/2$ is the average information of the original and style-diversified samples. D denotes the KL Divergence between the posterior probability $p \in \{p(x_S), p(\tilde{x}_S)\}$ and Q . JSD highlights the invariant pixel-level semantic information across two styles, impelling the model to be stable and insensitive to varied styles.

Retrospection Consistency (RC). Backbones of semantic segmentation models commonly start from ImageNet [6] pre-trained weights since the pre-trained backbones have learned general representation of many “things” categories in the real-world, including some classes in the self-driving scenes, *e.g.*, bicycle, bus, and car. However, the model learns more task-specific information and fit to the synthetic data when training with synthetic semantic segmentation data. As the model is required to deploy on unseen real-world scenarios, it is important to obtain some knowledge for the real-world objects. Previous works [41,14] leverage the abundant and even carefully selected images from real-world data, which may not be readily fulfilled in applications. Since ImageNet pre-trained weights are available to every segmentation model as the initialization, we propose RC to leverage such knowledge to lead the model less overfitting to the synthetic data and to retrospect the real-world knowledge lying in initialization. As the ImageNet pre-trained model is trained by image classification, only “things” classes are learned instead of “stuff” classes like road and wall. Consequently, RC is implemented as the feature-level distance minimization on those pixels of “things” classes. In addition, the style-diversified samples \tilde{x}_S in style consistency is also used in RC, which can lead the generated samples close to the real-world style. We define RC as:

$$\begin{aligned}\mathcal{L}_{RC}(x_S, \tilde{x}_S) &= \frac{1}{\sum_m M_{things}^{(m)}} \sum_m M_{things}^{(m)} \cdot \left(\left(f(x_S; \theta_S)^{(m)} - f(x_S; \theta_{IN})^{(m)} \right)^2 \right. \\ &\quad \left. + \left(f(\tilde{x}_S; \theta_S)^{(m)} - f(x_S; \theta_{IN})^{(m)} \right)^2 \right),\end{aligned}\tag{2}$$

where M_{things} denotes the mask of “things” classes. $f(x_S; \theta_S)$ and $f(\tilde{x}_S; \theta_S)$ denote the bottleneck feature of original sample x_S and style-diversified sample \tilde{x}_S respectively, which are obtained from the ongoing segmentation model θ_S .

$f(x_S; \theta_{IN})$ denotes the bottleneck feature of original sample x_S in the retrospective ImageNet pre-trained model θ_{IN} .

Discussion. Our retrospection consistency is inspired by the Feature Distance (FD) in DAFormer [13] but with different motivation and implementation. *First*, DAFormer focuses on unsupervised domain adaptation in semantic segmentation with unlabeled real-world data available, so it only focuses on fitting to the specific target domain (CityScapes [5]) rather than addressing unseen domain shift. *Second*, FD in DAFormer is used to better classify those similar classes (*e.g.*, bus and train) with the classification knowledge from ImageNet. Since we have no idea about the target distribution in DG-Seg, RC in our framework serves as an important guidance for the real-world knowledge, especially for more complex real-world scenes (*e.g.*, BDD100K [40] and Mapillary [24]). In addition, RC can also enhance the learning of real-world styles by taking the style-diversified samples into account.

3.2 Style Hallucination

Background. Style transfer [1,16] and domain generalization [33,44] methods show that the channel-wise mean and standard deviation can represent the non-content style of the image, which plays an important role in the domain shift. The style features can be readily used by AdaIN [16] which can transfer the image to an arbitrary style while remaining the content:

$$\text{AdaIN}(x, y) = \sigma(y) \left(\frac{x - \mu(x)}{\sigma(x)} \right) + \mu(y), \quad (3)$$

where x and y denotes the feature maps providing the content and style respectively. $\mu(\ast)$ and $\sigma(\ast)$ denotes the channel-wise mean and standard deviation. In domain generalization, as only one or multiple source domains are accessible, previous works modify AdaIN by replacing style features with other source styles. Those styles can be directly obtained from other samples [33] or can be generated by mixing other styles with its own styles [44].

Style Hallucination Module (SHM). Unlike image classification benchmarks where images are commonly of the same style in one dataset, even one semantic segmentation dataset contains various styles, *e.g.*, daytime, nighttime and twilight. This is why existing style variation methods [33,44] can work in single-source domain generalized segmentation. However, the ways of these methods in generating extra styles are sub-optimal, since they just randomly swap or mix source styles without considering the frequency and diversity of the source styles. As a result, more samples of the dominant style (*e.g.*, daytime) will be generated, yet the generated distribution may be quite different from the real-world one. Since we have no idea about the real-world target set, it is better to diversify the source samples as much as possible. We next introduce the Style Hallucination Module (SHM) for generating diverse source samples.

Definition 1: A basis B of a vector space V over a field F is a linearly independent subset of V that spans V . When the field is the reals \mathbb{R} , the resulting basis vectors are n -tuples of reals that span n -dimensional Euclidean space \mathbb{R}^n [9].

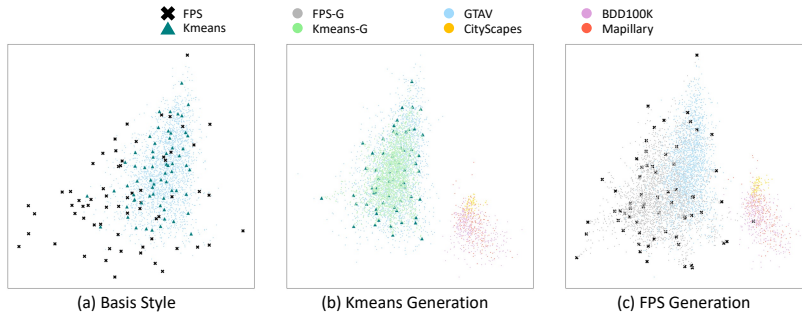


Fig. 3. Visualization of distributions of different domains. (a) Comparison of two kinds of basis styles; (b) generated style with Kmeans basis style; (c) generated style with FPS basis style. (Zoom in for details.)

According to Definition 1, style space can be viewed as a subspace of C -dimensional vector space, and thus all possible styles can be represented by the basis vectors. However, if we directly take C linearly independent vectors as the basis, *e.g.*, orthogonal unit vectors, many unrealistic styles are generated since the realistic styles are only in a small subspace, and such generated styles can impair the model training. To reconcile the diversity and realism, we use farthest point sampling (FPS) [28] to select C styles from all the source styles as basis styles. FPS is widely used for point cloud downsampling, which can iteratively choose C points from all the points, such that the chosen points are the most distant points with respect to the remaining points. Despite not strictly linearly independent, basis styles obtained by FPS can represent the style space to the utmost extent, and also contain many rare styles since rare styles are commonly far away from dominant ones. In addition, we recalculate the dynamic basis styles every k epochs instead of fixing basis style, as the style space is changing along with the model training. To generate new styles, we sample the combination weight $W = [w_1, \dots, w_C]$ from Dirichlet distribution $B([\alpha_1, \dots, \alpha_C])$ with the concentration parameters $[\alpha_1, \dots, \alpha_C]$ all set to $1/C$. The basis styles are then linearly combined by W :

$$\mu_{HS} = W \cdot \mu_{base}, \quad \sigma_{HS} = W \cdot \sigma_{base}, \quad (4)$$

where $\mu_{base} \in \mathbb{R}^{C \times C}$ and $\sigma_{base} \in \mathbb{R}^{C \times C}$ are the C basis styles. With the generated styles, style hallucinated samples \tilde{x}_S can be obtained by:

$$\tilde{x}_S = \sigma_{HS} \left(\frac{x_S - \mu(x_S)}{\sigma(x_S)} \right) + \mu_{HS}. \quad (5)$$

Discussion. Selecting representative basis styles is crucial for SHM. FPS is adopted in our method as it can cover the rare styles to the utmost extent. Another way is taking the Kmeans [22] clustering centers as the basis. As shown in Fig. 3(a), FPS samples (black cross) spread out more than Kmeans centers (teal

triangle), and can cover almost all possible source styles (lightskyblue point). When using the basis styles for style generation, styles obtained from Kmeans centers (Fig. 3(b)) are still within the source distributions and even ignore some possible rare styles. In contrast, FPS basis styles can generate more diverse styles (Fig. 3(c)), and even generate some styles close to the real-world ones (yellow, pink and orange point). Tab. 4 further demonstrates the effectiveness of FPS basis styles and shows that the Kmeans basis styles are even worse than directly swapping and mixing source styles.

3.3 Training Objective

The overall training objective is the combination of pixel-level cross entropy loss and the proposed two consistency constraints:

$$\begin{aligned} \mathcal{L} = & \frac{1}{2} (\mathcal{L}_{CE}(x_S, y_S) + \mathcal{L}_{CE}(\tilde{x}_S, y_S)) \\ & + \lambda_{SC} \mathcal{L}_{SC}(x_S, \tilde{x}_S) + \lambda_{RC} \mathcal{L}_{RC}(x_S, \tilde{x}_S), \end{aligned} \quad (6)$$

where λ_{SC} and λ_{RC} are the weights for style consistency and retrospection consistency, respectively.

4 Experiments

4.1 Experimental Setup

Datasets. Two synthetic datasets (GTAV [29] and SYNTHIA [30]) are taken as the source domains. GTAV [29] contains 24,966 images with the size of 1914×1052 , splitting into 12,403 training, 6,382 validation, and 6,181 testing images. SYNTHIA [30] contains 9,400 images of 960×720 , where 6,580 images are used for training and 2,820 images for validation. We evaluate the model on the validation sets of the three real-world datasets. CityScapes [5] contains 500 validation images of 2048×1024 . BDD-100K [40] and Mapillary [24] contain 1,000 1280×720 images and 2,000 1920×1080 images for validation, respectively.

Implementation Details. Following [4], we use DeepLabV3+ [2] as the segmentation model. The segmentation model is equipped with two backbones, ResNet-50 and ResNet-101 [11]. The SHM is inserted after the first Conv-BN-ReLU layer (layer0). We re-select the basis styles with the interval $k = 3$. We set $\lambda_{SC} = 10$ and $\lambda_{RC} = 1$. Models are optimized by the SGD optimizer with the learning rate 0.01, momentum 0.9 and weight decay 5×10^{-4} . The polynomial decay [18] with the power of 0.9 is used as the learning rate scheduler. All models are trained with the batch size of 8 for 40K iterations. During training, four widely used data augmentation techniques are adopted, including color jittering, Gaussian blur, random flipping, and random cropping of 768×768 .

Protocols. We conduct experiments on both single-source DG and multi-source DG. For *single-source DG*, to conduct a fair comparison with [4] and [14], we train the model with GTAV training data (12,403 images) when using ResNet-50

Table 1. Comparison with state-of-the-art methods on single-source DG with ResNet-50 and ResNet-101 as backbone, respectively. “Extra Data” denotes using extra real-world data during training. [§] denotes selecting best checkpoint for each target dataset.

Net	Methods (GTAV)	Extra Data	CityScapes	BDD100K	Mapillary	Mean
ResNet-50	Baseline	x	28.95	25.14	28.18	27.42
	SW [27]	x	29.91	27.48	29.71	29.03
	IterNorm [15]	x	31.81	32.70	33.88	32.79
	IBN-Net [26]	x	33.85	32.30	37.75	34.63
	DRPC [41] [§]	✓	37.42	32.14	34.12	34.56
	ISW [4]	x	36.58	35.20	40.33	37.37
	Ours	x	44.65	39.28	43.34	42.42
ResNet-101	Baseline	x	32.97	30.77	30.68	31.47
	IBN-Net [26]	x	37.37	34.21	36.81	36.13
	ISW [4]	x	37.20	33.36	35.57	35.38
	DRPC [41] [§]	✓	42.53	38.72	38.05	39.77
	FSDR [14] [§]	✓	44.80	41.20	43.40	43.13
	Ours	x	46.66	43.66	45.50	45.27

backbone (same with [4]), and with the whole GTAV datasets (24,966 images) when using ResNet-101 backbone (same with [14]). For *multi-source DG*, we follow [4] to train the model with the training set of GTAV (12,403 images) and SYNTHIA (6,580 images) using ResNet-50 backbone. **Note that** several state-of-the-art works, *e.g.*, [14] and [41], select the *best checkpoint* for each target dataset respectively, which is impractical since we cannot estimate the model performance on unseen domains in real-world applications. Instead, we directly use the *last checkpoint* to evaluate the three target datasets, which is more in line with the practical purpose of domain generalization.

Baseline. The baseline in each protocol is the model trained with the corresponding source training data by cross entropy loss function.

Evaluation Metric. We use the 19 shared semantic categories for training and evaluation. The mean intersection-over-union (mIoU) of the 19 categories on the three real-world datasets is adopted as the evaluation metric.

4.2 Comparison with State-of-the-art Methods

Single-source DG. In Tab. 1, we compare SHADE with state-of-the-art methods under single source setting, including SW [27], IterNorm [15], IBN-Net [26], ISW [4], DRPC [41] and FSDR [14]. *First*, we compare models that are trained with GTAV training set, using ResNet-50 backbone. SHADE achieves an average mIoU of 42.42% on the three real-world target datasets, yielding an improvement of 15.00% mIoU over the baseline and outperforming the previous best method (ISW) by 5.05%. *Second*, we further compare methods under the training protocol of DRPC [41] and FSDR [14], using ResNet-101 backbone and taking the

Table 2. Comparison with state-of-the-art methods on multi-source DG. All models use ResNet-50 backbone and are trained with training sets of GTAV and SYNTHIA.

Methods (G+S)	CityScapes	BDD100K	Mapillary	Mean
Baseline	35.46	25.09	31.94	30.83
IBN-Net [26]	35.55	32.18	38.09	35.27
ISW [4]	37.69	34.09	38.49	36.76
Ours	47.43	40.30	47.60	45.11

Table 3. Ablation studies on loss functions. All models use ResNet-50 backbone and are trained with GTAV training set. SHM: our style hallucination module; EMA: using exponential moving average model instead of ImageNet pre-trained model.

SHM	\mathcal{L}_{SC}	\mathcal{L}_{RC}	EMA	CityScapes	BDD100K	Mapillary	Mean
X	X	X	X	28.95	25.14	28.18	27.42
✓	X	X	X	38.68	32.40	35.96	35.68
✓	✓	X	X	42.66	35.92	40.42	39.67
✓	X	✓	X	41.43	37.65	41.77	40.29
✓	✓	X	✓	42.38	38.04	42.34	40.92
✓	✓	✓	X	44.65	39.28	43.34	42.42

whole set of GTAV (24,966 images) as the training data. Note that DRPC [41] and FSDR [14] utilize extra real-world data from ImageNet [6] or even driving scenes. Moreover, they select the best checkpoint for each target dataset, which is impractical in the real-world applications. Even so, we achieve the best results on all three datasets, **46.66% on CityScapes, 43.66% on BDD100K, 45.50% on Mapillary**, outperforming FSDR [14] by 2.14% in the average mIoU. These results show that we produce new state of the art in domain generalized semantic segmentation.

Multi-source DG. To further verify the effectiveness of SHADE, we compare SHADE with IBN-Net [26] and ISW [4] under the multi-source setting. We use ResNet-50 as the backbone and take the training set of GTAV and SYNTHIA as the source domains. As shown in Tab. 2, SHADE gains an improvement of 14.28% in average mIoU over the baseline, and outperforms ISW and IBN-Net by 8.35% and 9.84% respectively. The significant improvement over ISW and IBN-Net is mainly benefited from the various samples. With richer source samples, our SHM can generate more informative and diverse styles, which can effectively facilitate the dual consistency learning.

4.3 Ablation Studies

To investigate the effectiveness of each component in SHADE, we conduct ablation studies in Tab. 3.

Table 4. Comparison of different style variation methods. All models use ResNet-50 backbone and are trained with GTAV training set.

Methods (GTAV)	CityScapes	BDD100K	Mapillary	Mean
Baseline	28.95	25.14	28.18	27.42
Random Style	37.99	37.63	38.06	37.89
MixStyle [44]	43.14	37.94	42.22	41.10
CrossNorm [33]	43.13	37.20	41.83	40.72
Kmeans Basis	40.50	37.62	39.46	39.19
Ours	44.65	39.28	43.34	42.42

Effectiveness of Style Hallucination Module (SHM). SHM is the basis of SHADE. When using SHM only, we directly apply cross entropy loss on the style hallucinated samples. As shown in the second row of Tab. 3, our SHM can largely improve the model performance even without using the proposed dual consistency learning. This demonstrates the importance of training the model with diverse samples and the effectiveness of our SHM.

Effectiveness of Style Consistency (SC). SC is the consistency constraint that leads the model to learn style invariant representation. In Tab. 3, compared with only applying SHM, SC yields an improvement of 3.99% in average mIoU, demonstrating the superiority of the proposed logit pairing over cross entropy loss in learning style invariant model.

Effectiveness of Retrospection Consistency (RC). *First*, RC serves as an important guidance for narrowing the domain gap between synthetic and real data. Applying RC on top of SHM can yield an improvement of 4.61%, while removing RC will degrade the performance of SHADE by 2.75% in mIoU. *Second*, we conduct experiments to verify that the effectiveness of RC lies in the real-world knowledge instead of feature-level distance minimization of the paired samples. As directly minimizing the feature-level absolute distance of paired samples will lead to sub-optimal results (lead all the features close to zero), we replace the ImageNet pre-trained model in RC by exponential moving average (EMA) model. Comparing the fifth row and the sixth row in Tab. 3, EMA model only gains 1.25% improvement while RC improves the SC model by 2.75%. The results verify the significance of the retrospective knowledge in RC.

4.4 Further Evaluation

In this section, we compare other style variation methods with SHM and evaluate two important factors influencing SHM, *i.e.*, basis style selection interval k and the insert location of SHM.

Comparison of different style variation methods. We compare SHM with random style, MixStyle [44], CrossNorm [33] and style hallucination with Kmeans basis in Tab. 4. Random style utilizes the randomly sampled new styles from the standard normal distribution to form new samples. MixStyle [44] generates new

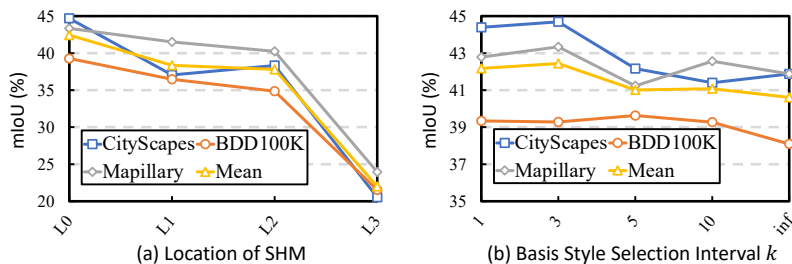


Fig. 4. Parameter analysis on the location of SHM and the basis style selection interval.

styles by mixing the original style with the random shuffled style within a mini-batch, and CrossNorm [33] swaps the original style with another style within the shuffled mini-batch. SHM and Kmeans basis both use the linear combination of basis style to generate new styles, but the basis styles of SHM are selected by FPS [28] while those of Kmeans basis are obtained by Kmeans clustering centers. We can make four observations from Tab. 4. **First**, we cannot make full use of dual consistency to achieve significant performance with the unrealistic random styles since standard normal distribution cannot represent the source nor the target domains. **Second**, despite the use of realistic source styles, random utilization of MixStyle and CrossNorm leads to the generation of more samples from the dominant styles that may be different from the real-world target styles. When using MixStyle and CrossNorm, the model achieves an average mIoU of 41.10% and 40.72%, respectively. **Third**, as shown in Fig. 3(b), Kmeans basis suffers from the similar but more severe dominant style issue in style generation. As a result, rare styles are discarded and thus the model achieves poorer performance than the above two. **Fourth**, SHM selects basis styles with FPS, and thus the selection can cover the source distribution to a large extent, especially those rare styles. With such basis styles, SHM generates styles from all the source distributions, and some generated styles are even close to the target domains (Fig. 3(c)). Consequently, combining SHM with dual consistency learning, SHADE can reap the benefit of the source data and outperforms other methods on all three target datasets.

Location of SHM. We investigate the impact of inserting SHM in different locations in Fig. 4(a). “L0” denotes inserting SHM after the first Conv-BN-ReLU layer (layer0) and “L1” to “L3” denote inserting SHM after the corresponding (1-3) ResNet layer. As shown in Fig. 4(a), “L0” achieves the best result while the performance of “L1” and “L2” drops a little. However, the model suffers from drastic performance degradation when inserting SHM after layer3. The reasons are two-fold. First, the channel-wise mean and standard deviation represent more style information in the shallow layers of deep neural networks while they contain more semantic information in deep layers [16, 7]. Second, the residual connections in ResNet will lead the ResNet activations of deep layers to have large peaks and

Table 5. Comparison with state-of-the-art methods trained on CityScapes.

Methods	GTAV	SYNTHIA	BDD100K	Mapillary
IBN-Net [26]	45.06	26.14	48.56	57.04
ISW [4]	45.00	26.20	50.73	58.64
Ours	48.61	27.62	50.95	60.67

small entropy, which makes the style features biased to a few dominant patterns instead of the global style [37].

Basis style selection interval. The distribution of source styles is varied along with the model training. To better represent the style space, we re-select the basis styles with the interval of k epochs. The abscissa of Fig. 4(b) denotes the selection interval k and “inf” denotes only selecting the basis style once in the beginning of training. As shown in Fig. 4(b), the model achieves consistent and good performance with frequent re-selection ($k \leq 3$) while the performance degrades with the increase of selection interval, and the average mIoU is lower than 41% when only selecting once. Taking both the performance and computational cost into consideration, we set $k = 3$ in SHADE.

4.5 Real-to-Others Domain Generalization

To further demonstrate the effectiveness of SHADE, we leverage CityScapes [5] as the source domain and generalize to real (BDD100K [40] and Mapillary [24]) and synthetic (GTAV [29] and SYNTHIA [30]) domains. As shown in Tab. 5, SHADE consistently outperforms ISW [4] and IBN-Net [26] on both real and synthetic datasets. These results verify the versatility of our method.

5 Conclusion

In this paper, we propose a novel framework (SHADE) for synthetic-to-real domain generalized semantic segmentation. SHADE leverages two consistency constraints to learn the domain-invariant representation by seeking consistent representation across styles and the guidance of retrospective knowledge. In addition, the style hallucination module (SHM) is equipped into our framework, which can effectively catalyze the dual consistency learning by generating diverse and realistic source samples. Experiments on three real-world dataset show that SHADE achieves state-of-the-art performance under both single- and multi-source domain generalization settings with different backbones.

Acknowledgements. This research/project is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-RP-2020-016), the Tier 2 grant MOE-T2EP20120-0011 from the Singapore Ministry of Education, and the EU H2020 project AI4Media (No. 951911).

References

1. Chen, H., Zhao, L., Zhang, H., Wang, Z., Zuo, Z., Li, A., Xing, W., Lu, D.: Diverse image style transfer via invertible cross-space mapping. In: ICCV (2021) [7](#)
2. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: ECCV (2018) [1](#), [9](#)
3. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2020) [4](#)
4. Choi, S., Jung, S., Yun, H., Kim, J.T., Kim, S., Choo, J.: Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In: CVPR (2021) [2](#), [4](#), [9](#), [10](#), [11](#), [14](#)
5. 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) [1](#), [7](#), [9](#), [14](#)
6. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009) [1](#), [3](#), [6](#), [11](#)
7. Dumoulin, V., Shlens, J., Kudlur, M.: A learned representation for artistic style. In: ICLR (2017) [4](#), [13](#)
8. French, G., Laine, S., Aila, T., Mackiewicz, M., Finlayson, G.: Semi-supervised semantic segmentation needs strong, varied perturbations. In: BMVC (2020) [4](#)
9. Halmos, P.R.: Finite-dimensional vector spaces. Springer (1987) [7](#)
10. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020) [4](#)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) [9](#)
12. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: AugMix: A simple data processing method to improve robustness and uncertainty. In: ICLR (2020) [4](#), [5](#)
13. Hoyer, L., Dai, D., Van Gool, L.: Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In: CVPR (2022) [7](#)
14. Huang, J., Guan, D., Xiao, A., Lu, S.: Fsd: Frequency space domain randomization for domain generalization. In: CVPR (2021) [2](#), [3](#), [4](#), [6](#), [9](#), [10](#), [11](#)
15. Huang, L., Zhou, Y., Zhu, F., Liu, L., Shao, L.: Iterative normalization: Beyond standardization towards efficient whitening. In: CVPR (2019) [10](#)
16. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV (2017) [4](#), [7](#), [13](#)
17. Kamnitsas, H., Kirr, A., Goodfellow, I.: Adversarial logit pairing. In: ICML (2018) [5](#)
18. Liu, W., Rabinovich, A., Berg, A.C.: Parsenet: Looking wider to see better. In: CoRR (2015) [9](#)
19. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021) [1](#)
20. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. arXiv preprint arXiv:2201.03545 (2022) [1](#)
21. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR (2015) [1](#)

22. MacQueen, J.: Some methods for classification and analysis of multivariate observations. In: Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability (1967) 8
23. Meng, R., Li, X., Chen, W., Yang, S., Song, J., Wang, X., Zhang, L., Song, M., Xie, D., Pu, S.: Attention diversification for domain generalization. In: ECCV (2022) 4
24. Neuhold, G., Ollmann, T., Rota Bulò, S., Kotschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: ICCV (2017) 7, 9, 14
25. Pan, F., Shin, I., Rameau, F., Lee, S., Kweon, I.S.: Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In: CVPR (2020) 4
26. Pan, X., Luo, P., Shi, J., Tang, X.: Two at once: Enhancing learning and generalization capacities via ibn-net. In: ECCV (2018) 2, 4, 10, 11, 14
27. Pan, X., Zhan, X., Shi, J., Tang, X., Luo, P.: Switchable whitening for deep representation learning. In: ICCV (2019) 10
28. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: NeurIPS (2017) 3, 8, 13
29. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: ECCV (2016) 2, 3, 9, 14
30. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: CVPR (2016) 2, 9, 14
31. Sakaridis, C., Dai, D., Gool, L.V.: Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In: ICCV (2019) 1
32. Sakaridis, C., Dai, D., Van Gool, L.: Acdd: The adverse conditions dataset with correspondences for semantic driving scene understanding. In: ICCV (2021) 1
33. Tang, Z., Gao, Y., Zhu, Y., Zhang, Z., Li, M., Metaxas, D.: Selfnorm and crossnorm for out-of-distribution robustness. ICCV (2021) 3, 4, 6, 7, 12, 13
34. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: NeurIPS (2017) 4
35. Tsai, Y.H., Hung, W.C., Schulter, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: CVPR (2018) 1
36. Wang, H., Xiao, C., Kossaifi, J., Yu, Z., Anandkumar, A., Wang, Z.: Augmax: Adversarial composition of random augmentations for robust training. In: NeurIPS (2021) 4
37. Wang, P., Li, Y., Vasconcelos, N.: Rethinking and improving the robustness of image style transfer. In: CVPR (2021) 14
38. Wang, Z., Luo, Y., Qiu, R., Huang, Z., Baktashmotlagh, M.: Learning to diversify for single domain generalization. In: ICCV (2021) 4
39. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. NeurIPS (2021) 1
40. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In: CVPR (2020) 7, 9, 14
41. Yue, X., Zhang, Y., Zhao, S., Sangiovanni-Vincentelli, A., Keutzer, K., Gong, B.: Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In: ICCV (2019) 2, 3, 4, 6, 10, 11

42. Zhao, Y., Zhong, Z., Luo, Z., Lee, G.H., Sebe, N.: Source-free open compound domain adaptation in semantic segmentation. *IEEE TCSVT* (2022) [4](#)
43. Zhong, Z., Zhao, Y., Lee, G.H., Sebe, N.: Adversarial style augmentation for domain generalized urban-scene segmentation. *arXiv preprint arXiv:2207.04892* (2022) [2](#)
44. Zhou, K., Yang, Y., Qiao, Y., Xiang, T.: Domain generalization with mixstyle. In: *ICLR* (2021) [3](#), [4](#), [6](#), [7](#), [12](#)
45. Zhou, Q., Feng, Z., Gu, Q., Cheng, G., Lu, X., Shi, J., Ma, L.: Uncertainty-aware consistency regularization for cross-domain semantic segmentation. *Computer Vision and Image Understanding* (2022) [4](#)
46. Zhou, Q., Feng, Z., Gu, Q., Pang, J., Cheng, G., Lu, X., Shi, J., Ma, L.: Context-aware mixup for domain adaptive semantic segmentation. *arXiv preprint arXiv:2108.03557* (2021) [1](#)
47. Zhou, Q., Zhuang, C., Yi, R., Lu, X., Ma, L.: Domain adaptive semantic segmentation with regional contrastive consistency regularization. In: *ICME* (2022) [1](#)
48. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *ICCV* (2017) [3](#)
49. Zou, Y., Yu, Z., Vijaya Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: *ECCV* (2018) [4](#)