

Generalization by Adaptation: Diffusion-Based Domain Extension for Domain-Generalized Semantic Segmentation

Joshua Niemeijer^{1*} Manuel Schwonberg^{2*} Jan-Aike Termöhlen^{3*}

Nico M. Schmidt² Tim Fingscheidt³

¹DLR ²CARIAD SE ³Technische Universität Braunschweig

joshua.niemeijer@dlr.de {manuel.schwonberg, nico.schmidt}@cariad.technology
{j.termoehlen, t.fingscheidt}@tu-bs.de

Abstract

When models, e.g., for semantic segmentation, are applied to images that are vastly different from training data, the performance will drop significantly. Domain adaptation methods try to overcome this issue, but need samples from the target domain. However, this might not always be feasible for various reasons and therefore domain generalization methods are useful as they do not require any target data. We present a new diffusion-based domain extension (DIDEX) method and employ a diffusion model to generate a pseudo-target domain with diverse text prompts. In contrast to existing methods, this allows to control the style and content of the generated images and to introduce a high diversity. In a second step, we train a generalizing model by adapting towards this pseudo-target domain. We outperform previous approaches by a large margin across various datasets and architectures without using any real data. For the generalization from GTA5, we improve state-of-the-art mIoU performance by 3.8 % absolute on average and for SYNTHIA by 11.8 % absolute, marking a big step for the generalization performance on these benchmarks. Code is available at <https://github.com/JNiemeijer/DIDEX>

1. Introduction

In recent years, the success of deep learning has led to significant advancements in the field of computer vision, e.g., for semantic segmentation. For this task, the usage of synthetic data is particularly interesting as manual data labeling is time- and cost-intensive. Synthetic data can be valuable for training as well as validation, since it allows the simulation of rare and dangerous events. However, it is still challenging to deploy models trained on synthetic data in real-world settings with varying data distri-

* indicates equal contribution

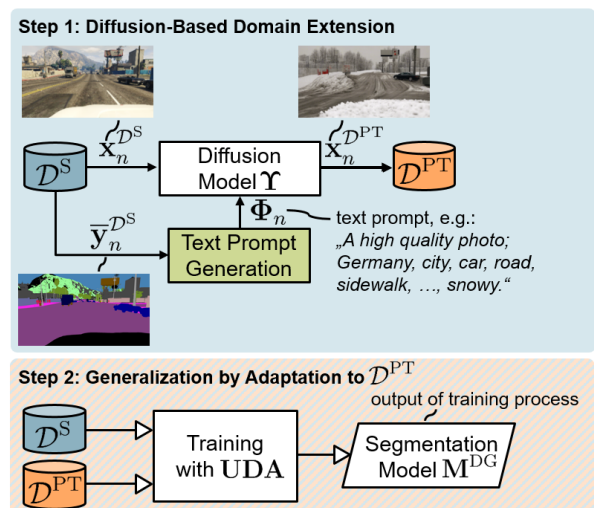


Figure 1. **Overview** showing a **block diagram** of our diffusion-based domain extension (DIDEX) method (step 1), followed by the **generalization-by-adaptation process** to the newly generated pseudo-target domain \mathcal{D}^{PT} (step 2).

butions due to large domain shifts towards real-world data [29, 48]. One approach to overcome this problem is by unsupervised domain adaptation (UDA) [42], where unlabeled data from a real target domain is available to adapt to [3, 13, 22, 27, 30, 31, 46, 51, 55]. Some approaches also perform this task source-free [22] or in a continual manner [21, 46]. Recently, vision transformer models [9, 52] caused a significant increase in performance and reducing the domain gap with DAFormer [14] being the initial work. All these methods rely on access to target domain data to adapt to this particular domain. In practice, this data might not always be available due to various reasons. Data collection can be difficult, because, e.g., adverse weather conditions such as fog, rain, and snow do not constantly occur, and sometimes the target domain cannot be anticipated at all. Consequently, the field of domain generalization (DG)

emerged where no data from the target domain is available at all and the task is to generalize from only a single, usually synthetic, source domain to unseen and unknown target domains [2, 18, 23, 32, 34, 58]. Style transfer methods such as AdaIN [19] are widely used in UDA, especially in combination with other techniques [13, 49, 50]. For domain generalization these methods cannot be used since no target domain guidance is available for the style transfer. For this reason, the majority of DG methods performs style randomization or augmentation to alter the visual appearance of the source domain [18, 20, 23, 35, 44, 58]. However, these methods reveal two major issues. First, many of them require additional real data from auxiliary domains [18, 20, 23, 35], which disrupts the idea of domain generalization. Second, the style randomization is difficult to control and often leads to limited diversity, as mostly only textures are changed.

In this work, we propose a novel method that tackles the domain generalization problem by diffusion-based domain extension (DIDEX). Diffusion models have demonstrated remarkable capabilities in capturing complex distributions and semantics and generating realistic samples with high quality [7, 12, 38, 43]. Our method allows not only to generate realistic pseudo-target images from synthetic images but also to alter important parameters such as location, time, weather conditions and semantic content via text prompts. We use these capabilities to generate a realistic looking pseudo-target domain as shown in Figure 1. Diffusion models have severe limitations in the semantic consistency of the generated output [60] w.r.t. the input image. To address these limitations, we draw inspiration from the field of *unsupervised* domain adaptation (UDA) and define the diffusion data as the target domain for our adaptation process. The key contributions of this paper are as follows:

- We propose a new method leveraging diffusion models for domain transfer, enhancing our model’s ability to generalize across domains without accessing real data.
- By utilizing UDA techniques, we overcome the semantic inconsistency limitation of diffusion models, ensuring improved domain generalization performance.
- We report results on common benchmarks, where our method outperforms the state-of-the-art domain generalization methods by a large margin.

2. Related Work

In this section, we will give a brief overview on the related methods for unsupervised domain adaptation, domain generalization, and diffusion models.

2.1. Unsupervised Domain Adaptation (UDA)

UDA methods are clustered into input, feature, output and hybrid adaptation approaches [42]. Our method has

similarities to the idea of UDA input space adaptation where the aim is to style transfer the visual appearance between the source and target domain. Often a CycleGAN [64] is utilized for a semantically consistent style transfer, usually in conjunction with other methods [8, 11, 13, 26, 54]. In the feature space, contrastive learning methods such as SePiCo [51], UCDA [59], ProDA [61], CPSL [24] and CONFETI [25] obtained significant improvements. DAFormer [14], HRDA [15] and MIC [17] are transformer-based approaches improving the state-of-the-art performance significantly by using a slightly modified `SegFormer` [52] architecture and employing methods such as rare class sampling, feature distance, teacher-student training, self-training, multi-resolution training and masked image consistency. Benigim et al. [1] proposed a data augmentation technique using a finetuned text-to-image diffusion model for one-shot unsupervised domain adaptation.

2.2. Domain Generalization (DG)

Domain generalization methods for semantic segmentation try to overcome the domain gap by either constraining learned feature distributions or randomizing distributions by augmentation or extension during training.

Pan *et al.* proposed IBN-Net [32] as one of the first DG approaches and employed instance-batch normalization (IBN), which is more robust w.r.t. appearance changes. Further approaches utilize instance whitening to decorrelate features from different layer channels and thereby removing style-related information. Pan *et al.* [33] proposed an approach which can switch between different whitening and normalization techniques depending on the task. RobustNet [5], DURL [53], and SAN+SAW [34] all proposed improvements such as a better guidance of the whitening process by a so-called sensitivity-aware prior module [53] or semantic-aware whitening and normalization [34].

Domain randomization can be separated into approaches that require additional real-world data (auxiliary domains) and those which do not need it. Yue *et al.* [58] proposed one of the first domain randomization methods combining domain randomization and pyramid consistency (DRPC), enforcing semantic feature consistency between differently stylized images with transferred textures from ImageNet. Similarly, Huang *et al.* [18] with FSDR employ style randomization in the frequency space, Peng *et al.* [35] use painting for their style randomization and WildNet by Lee *et al.* [23] applies style randomization in conjunction with contrastive and consistency learning. Kim *et al.* [20] utilize style randomization with internet sampled images and self-training similar to UDA approaches. In contrast to these methods, our proposed method can be seen as a more structured distribution randomization without the need for real-world auxiliary domains.

Only few works perform style randomization without




Figure 2. **Simplified representation of different domain distributions** and the effects of domain adaptation and generalization methods. Shown in (a) are the initial distributions for a source domain \mathcal{D}^S and two target domains \mathcal{D}^{T_1} and \mathcal{D}^{T_2} . In (b), the effect of domain adaptation from \mathcal{D}^S to a specific target domain \mathcal{D}^{T_1} is shown. In (c), the effect of commonly used domain generalization methods is shown that broaden the source distribution, e.g., via data augmentation, but in a less directed manner. In (d), we show the effect of our proposed method, where we perform a domain extension and generate a pseudo-target domain \mathcal{D}^{PT} that also covers the distributions of multiple target domains (\mathcal{D}^{T_1} and \mathcal{D}^{T_2}). For the purpose of clarity, only $K=2$ target domains are shown here. In the remainder of this paper, by default, we employ $K=4$ target domains for the evaluation.

additional real data. AugFormer by Schwonberg *et al.* [41] showed that simple augmentations can significantly increase domain generalization. Similarly, Sun *et al.* [44] proposed a strong style randomization in the CIELAB color space. Zhong *et al.* [63] proposed to alter the channel-wise mean and standard deviation in an adversarial manner to generate stylized images which are hard to segment for the model. SHADE by Zhao *et al.* [62] introduces a style hallucination module which is based on farthest point sampling and generates new diverse samples by a linear combination of the basic styles. Bi *et al.* [2] proposed CMFormer which is the first DG approach tailored towards vision transformers. It fuses low and original resolution in a so-called content enhanced mask attention. In contrast, our approach is independent from the architecture. Gong *et al.* [10] proposed PromptFormer as the first DG approach utilizing diffusion models. However, their method fundamentally differs from ours since they employ a diffusion backbone for domain-invariant pre-training, where scene and category text prompts help the model to disentangle domain-variant and invariant information. They also employ a consistency loss to enforce same predictions under different input prompts.

2.3. Diffusion Models

Diffusion models for image synthesis recently emerged and outperform the long established standard using GANs [7]. Several improvements were proposed to accelerate the image generation, namely denoising diffusion probabilistic models (DDPMs) [12] and denoising diffusion implicit models (DDIMs) [43]. Rombach *et al.* proposed latent diffusion models (LDMs) sometimes referred to as stable diffusion models which reduce the training and inference time and enable text-to-image and image-to-image genera-

tion. Based on this, several extensions and improvements were proposed such as stable diffusion XL (SDXL) [36] or ControlNet [60], which offers the possibility to constrain the generation process with different additional inputs such as depth or semantics.

3. Proposed Method

In this section, we will describe the approach on the text prompt generation of our new diffusion-based domain extension (DIDEX) method. Subsequently, we will describe the generalization-by-adaptation approach.

3.1. Diffusion-Based Domain Extension

The general approach and text prompt generation will be introduced in the following in detail.

3.1.1 General Approach

We define an input image as $\mathbf{x}_n \in \mathbb{G}^{H \times W \times C}$, with \mathbb{G} denoting the set of integer color intensity values, H and W image height and width in pixels, and $C=3$ the number of color channels, respectively. Images in a dataset with N images are indexed $n \in \{1, 2, \dots, N\}$. A semantic segmentation network M maps \mathbf{x}_n to an output map $\mathbf{y}_n = (y_{n,i,s}) \in \mathbb{I}^{H \times W \times S}$ with posterior probabilities $y_{n,i,s} = P(s|i, \mathbf{x}_n)$ for each class $s \in \mathcal{S}$ at pixel index $i \in \mathcal{I} = \{1, 2, \dots, H \cdot W\}$. Here, $\mathcal{S} = \{1, 2, \dots, S\}$ denotes the set of S classes and $\mathbb{I} = [0, 1]$. Superscripts ‘‘S’’ and ‘‘T’’ on \mathbf{x}_n and \mathbf{y}_n denote the domain from which the variables stem, with, e.g., \mathcal{D}^S being the source domain and \mathcal{D}^T being the target domain. With domain generalization, the aim is to generalize to multiple unseen domains so the target domains \mathcal{D}^{T_k} are indexed with $k \in \mathcal{K} = \{1, 2, \dots, K\}$, where K is the total amount of target domains that the method

will be evaluated on. The crucial advantage of diffusion models is the prompt-based control of the generated image. That enables us to define the content and style of the generated images and create a pseudo-target domain \mathcal{D}^{PT} . Two different strategies for the generation of pseudo-target domains are possible. First, a specialized pseudo-target domain could be created, if some information regarding the target domains are available, such as location, weather, etc. In our approach, however, we aim for a domain generalization without any assumptions about the target domains and generate \mathcal{D}^{PT} with high diversity to cover a wide range of possible target domain distributions as shown in Figure 2 (d). We formulate the objective to generate a pseudo-target domain, which is the union of all relevant target domains as

$$\mathcal{D}^{\text{PT}} \approx \bigcup_{k \in \mathcal{K}} \mathcal{D}^{\text{T}_k}, \quad (1)$$

where \mathcal{D}^{T_k} denotes the relevant target domains. In our domain generalization setting we only have a single source domain distribution available to obtain \mathcal{D}^{PT} . We use the diffusion model to extend the initial source distribution into the more diverse and wider pseudo-target distribution. A single image of \mathcal{D}^{PT} can be obtained by:

$$\mathbf{x}_n^{\text{PT}} = \Upsilon(\mathbf{x}_n, \Phi_n) \quad (2)$$

where Υ denotes the diffusion model and Φ_n the text prompt. The generated pseudo-target domain provides the crucial advantage that we have a target domain available that can be used for adapting the segmentation network as shown in Figure 2 (d). Compared to other DG approaches, this is a novel and unique feature of our method and enables us to obtain generalization by adaptation.

Semantic consistency: We observed that semantic or structural consistency of the image-to-image generation is often not given for complex urban street scenes, i.e., pedestrians or cars can disappear and the semantic content of the pixels is changed (see Figure 4). Directly using the synthetic labels for the pseudo-target domain is therefore not possible. However, we assume that a certain structural and semantic consistency is also beneficial for the adaptation in step two and for this reason we employ depth-guided Stable Diffusion 2.0 [38] and ControlNet [60], which offer a higher structural and semantic consistency. Both extend the image generation to

$$\mathbf{x}_n^{\text{PT}} = \Upsilon(\mathbf{x}_n, \mathbf{C}(\mathbf{x}_n), \Phi_n), \quad (3)$$

where \mathbf{C} is a constraint based on the image \mathbf{x}_n , e.g., by canny edge extraction, depth or semantic prediction.

3.1.2 Text Prompt Generation

Text prompts are crucial for the class distribution of the pseudo-target domain, since the prompts determine the style

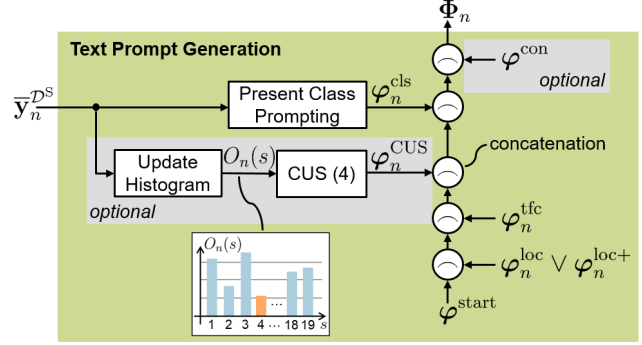


Figure 3. Block diagram of the **text prompt generation** from Figure 1 providing Φ_n . The operation in the circles denotes a concatenation. Concatenation of φ^{CUS} and φ^{con} is optional. Class-uniform sampling (CUS) is described in (4). The rarest class in the example histogram $O_n(s)$ is marked in orange.

and content of the generated images. To the best of our knowledge, there are no existing strategies how to design text prompts for diffusion models in the context of complex urban street scenes. For this reason, we developed a new, systematic and modular text prompt generation that is visualized in Figure 3. Our text prompt strategy aims at varying the image content and style at different levels, introducing as much diversity as possible into the pseudo-target domain. To design the building blocks of the prompts, we considered major causes of real-to-real domain shifts as a guidance.

The start of each prompt is $\varphi_n^{\text{start}} = \text{"A high quality photo;"}$ We observed that particularly emphasizing photorealism with prompts such as "real" or "photorealistic" did not improve the realism of the output. The location shift between different countries or continents is important in automated driving so we include a set of different countries and regions to include their style and content in \mathcal{D}^{PT} . We concatenate a geographic location to the start prompt by choosing either a location $\varphi_n^{\text{loc}} \in \{\text{"Europe"}, \text{"Germany"}\}$ or a location from an extended location set $\varphi_n^{\text{loc+}} \in \{\text{"Europe"}, \text{"Germany"}, \text{"USA"}, \text{"China"}, \text{"India"}\}$. We are aware that this set of locations is not globally comprehensive and may have a significant impact on the domain generalization performance. We leave an extensive study with more regions of the world for future work. Subsequently we concatenate a traffic location $\varphi_n^{\text{tfc}} \in \{\text{"_"}\}$ (blank), "highway", "city". The blank denotes that no traffic location is provided.

Present class prompting: We utilize the GTA5 segmentation label map $\bar{y}_n^{\mathcal{D}^{\text{S}}}$ to identify and to concatenate all classes that are present in the current image. This serves as an additional guidance on which classes $\varphi_n^{\text{cls}} \subseteq \mathcal{S}$ the diffusion model should include in the output.

Class-uniform sampling (CUS): Optionally, we intro-

duce a class-uniform sampling (CUS) to mitigate the unbalanced class distribution by tracking the number of times a class occurred in the synthetic source images and concatenating the least often seen class to the prompt. For each image with index n , we utilize the segmentation labels from the respective source image $\bar{y}_n^{\mathcal{D}^S}$ to update a histogram with the class occurrences $O_n(s)$. We then concatenate the rarest class name to the prompt:

$$\varphi_n^{\text{CUS}} = \text{class}(\arg \min_{s \in \mathcal{S}}(O_n(s))). \quad (4)$$

After concatenating the rarest class to our prompt, we subsequently increase the counter for this class. This leads to a more uniform class distribution in our domain extension.

Image conditions: Adverse weather or visibility conditions are a challenge for autonomous vehicles and we therefore also represent these shifts in our pseudo-target domain by adding one of the following words to the prompt:

$$\varphi_n^{\text{con}} \in \{ \text{"rain"}, \text{"fog/mist"}, \text{"snowy"}, \text{"sunny"}, \\ \text{"overcast"}, \text{"stormy"}, \text{"overexposure"}, \\ \text{"underexposure"}, \text{"evening"}, \text{"morning"}, \\ \text{"night/darkness"}, \text{"backlighting"}, \\ \text{"artificial lighting"}, \text{"harsh light"}, \\ \text{"dappled light"}, \text{"sun flare"}, \text{"hazy/haze"}, \\ \text{"spring"}, \text{"autuum"}, \text{"winter"}, \text{"summer"} \}$$

By concatenating all the different building blocks, we obtain our text prompt according to Figure 3 as:

$$\Phi_n = \varphi^{\text{start}} \frown (\varphi_n^{\text{loc}} \vee \varphi_n^{\text{loc}+}) \frown \varphi^{\text{tfc}} \\ \frown \varphi^{\text{CUS}} \frown \varphi^{\text{cls}} \frown \varphi^{\text{con}}, \quad (5)$$

with \frown symbolizing concatenation and φ^{CUS} and φ^{con} being optional. We denote the base prompt without the optional strings as Φ_n^{base} . A sample prompt can become, e.g., “A high quality photo; Europe, highway, road, car, building, vegetation, winter”. If there are multiple options within one of the building blocks, we select the prompt randomly. We generate \mathcal{D}^{PT} before the adaptation step.

3.2. Generalization by Adaptation

Even with strong semantic guidance by depth or the semantic ground truth, the stable diffusion models generate partially highly inconsistent outputs. However, the area of unsupervised domain adaptation provides powerful methods which are capable to handle datasets without labels effectively. We utilize these methods to adapt the segmentation model towards the pseudo-target domain that was generated in a 1st step (Section 3.1). For UDA methods, we usually obtain an adapted segmentation model by $\text{UDA}(\mathcal{D}^S, \mathcal{D}^T) \rightarrow \mathbf{M}^{\text{UDA}}$, where \mathcal{D}^T is usually a real data

target domain without labels and \mathbf{M} our domain-adapted model. Our method changes this setting slightly but notably to $\text{UDA}(\mathcal{D}^S, \mathcal{D}^{\text{PT}}) \rightarrow \mathbf{M}^{\text{DG}}$, where model \mathbf{M} is adapted to our pseudo-target domain \mathcal{D}^{PT} and therefore also generalizes without accessing any real data. In our case, UDA strategies create this supervision signal by aligning the distributions of the source \mathcal{D}^S and pseudo-target domain \mathcal{D}^{PT} .

4. Experimental Setup

In the following, we introduce the employed datasets, metrics, and diffusion and segmentation network architectures. Afterwards, we introduce the utilized UDA methods.

4.1. Datasets and Metrics

Datasets: We follow the common domain generalization standard setting [10, 23, 34, 35, 44] and employ the two synthetic datasets SYNTHIA (SYN) [39] and GTA5 [37] as our source domains. The datasets comprise 9400 and 24966 images, respectively, which we denote as $\mathcal{D}_{\text{train}}^{\text{SYN}}$ and $\mathcal{D}_{\text{train}}^{\text{GTA5}}$. To evaluate the domain generalization capabilities, we use Cityscapes (CS) [6], BDD100k (BDD) [57], Mapillary Vistas (MV) [28] as our real-world target domains with 500, 1000 and 2000 validation images, respectively. The training sets of these datasets remain unused in our generalization method. We evaluate all experiments on the validation sets \mathcal{D}_{val} of our target domains and compute the domain generalization (DG) mean over these sets, as is common practice in domain generalization [18, 34, 58]. As an addition to this domain generalization benchmark, we also include ACDC [40] with 406 validation images in some of our experiments, which contains adverse weather conditions. It is excluded from the DG mean, because most other approaches do not report ACDC performance.

Metrics: For evaluation we employ the mean intersection over union (mIoU) of $S = 19$ segmentation classes [6, 37, 40] for GTA5 trained models and $S = 16$ classes for models trained on SYNTHIA, as it is common practice [22].

4.2. Network Architectures

Diffusion models: For the data generation process, we employ Stable Diffusion 2.0 (SD2.0) [38] as our standard diffusion model. We also compare against ControlNet [60], which provides several options to constrain the output generation in our ablation studies. To create data, we utilize the image-to-image prompting strategy described in Section 3. The ablation studies w.r.t. the data generation process that are described in Section 5 analyze the different prompting elements.

Segmentation models: As the network architectures we selected the common standards in the area of domain generalization. The DeepLabV2 [4] architecture with a ResNet-101 is used to evaluate the generalization on a

Table 1. **Domain generalization performance** (mIoU (%)) of several methods employing two different encoder networks. **Training** was performed on the synthetic **GTA5** ($\mathcal{D}^S = \mathcal{D}_{\text{train}}^{\text{GTA5}}$) dataset. **Evaluation** is performed on various real-world **validation sets** ($\mathcal{D}^T = \mathcal{D}_{\text{val}}$). Prior work results are either cited from [23] (marked with $^\circ$) or from the respective paper (marked with *).

Enc.	DG Method	No real data	mIoU (%) on			
			$\mathcal{D}_{\text{val}}^{\text{CS}}$	$\mathcal{D}_{\text{val}}^{\text{BDD}}$	$\mathcal{D}_{\text{val}}^{\text{MV}}$	DG mean
ResNet	Baseline	✓	36.1	36.6	43.8	38.8
	IBN-Net $^\circ$ [32]	✓	37.7	36.7	36.8	37.1
	RobustNet $^\circ$ [5]	✓	37.3	38.7	38.1	38.0
	DRPC* [58]	✗	42.5	38.7	38.1	39.8
	SW* [33]	✓	36.1	36.6	32.6	35.1
	FSDR* [18]	✗	44.8	41.2	43.4	43.1
	SAN+SAW* [34]	✓	45.3	41.2	40.8	42.4
	WEDGE* [20]	✗	45.2	41.1	48.1	44.8
	GTR* [35]	✗	43.7	39.6	39.1	40.8
	SHADE* [62]	✓	46.7	43.7	45.5	45.3
	WildNet $^\circ$ [23]	✗	45.8	41.7	47.1	44.9
	RICA* [44]	✓	48.0	45.2	46.3	46.5
	DIDEX with MIC [17]	✓	52.4	40.9	49.2	47.5
	Transformer	Baseline	✓	46.6	45.6	50.1
ReVT* [45]		✓	50.0	48.0	52.8	50.3
DAFormer* [14]		✓	52.7	47.9	54.7	51.7
HRDA* [16]		✓	57.4	49.1	61.2	55.9
CMFormer* [2]		✓	55.3	49.9	60.1	55.1
PromptFormer* [10]		✓	52.0	-	-	-
DIDEX with MIC [17]		✓	62.0	54.3	63.0	59.7

CNN-based architecture and for the recently emerged vision transformer we use the DAFormer [14] network.

4.3. Employed UDA Methods

We utilize five different state-of-the-art domain adaptation approaches for our generalization by adaptation step, namely DACS [47], DAFormer [14], HRDA [15], MIC [17] and SePiCo [51]. As described in Section 3, we utilize these UDA methods to adapt from the synthetic source domain to our pseudo-target domain. Previous experiments have shown that the adaptation of UDA methods to other domain shifts can diminish the performance [40]. In this context, *it has to be noted that we do not finetune any hyperparameters of the employed UDA methods*. All the methods are employed as provided by the respective authors.

5. Evaluation and Discussion

In this section, we will first compare our approach to state-of-the-art domain generalization methods. Afterwards, we investigate the impact of different prompt-

Table 2. **Domain generalization performance** (mIoU (%)) of several methods employing two different encoder networks. **Training** was performed on the synthetic **SYNTHIA** ($\mathcal{D}^S = \mathcal{D}_{\text{train}}^{\text{SYN}}$) dataset. **Evaluation** is performed on real-world **validation sets** ($\mathcal{D}^T = \mathcal{D}_{\text{val}}$). Prior work results are either cited from [23] (marked with $^\circ$) or from the respective paper (marked with *).

Enc.	DG Method	No real data	mIoU (%) on			
			$\mathcal{D}_{\text{val}}^{\text{CS}}$	$\mathcal{D}_{\text{val}}^{\text{BDD}}$	$\mathcal{D}_{\text{val}}^{\text{MV}}$	DG mean
ResNet	Baseline	✓	34.3	27.8	38.0	33.4
	IBN-Net $^\circ$ [32]	✓	34.2	32.6	36.2	34.3
	DRPC* [58]	✗	37.6	34.4	34.1	35.4
	SW* [33]	✓	36.1	36.6	32.6	35.1
	FSDR* [18]	✗	40.8	37.4	39.6	39.3
	SAN+SAW* [34]	✓	40.9	36.0	37.3	38.1
	WEDGE* [20]	✗	40.9	38.1	43.1	40.7
	GTR* [35]	✗	39.7	35.3	36.4	37.1
	RICA* [44]	✓	45.0	36.3	41.6	41.0
	DIDEX with MIC [17]	✓	53.1	41.8	50.3	48.4
	Transformer	Baseline	✓	41.4	36.2	42.4
ReVT* [45]		✓	46.3	40.3	44.8	43.8
CMFormer* [2]		✓	44.6	33.4	43.3	40.4
PromptFormer* [10]		✓	49.3	-	-	-
DIDEX with MIC [17]		✓	59.8	47.4	59.5	55.6

ing strategies, UDA approaches, semantic consistency constraints, and the quantity of generated images.

5.1. Comparison with State of the Art

We compare the generalization performance of our method to state-of-the-art methods in Table 1 ($\mathcal{D}^S = \mathcal{D}_{\text{train}}^{\text{GTA5}}$) and Table 2 ($\mathcal{D}^S = \mathcal{D}_{\text{train}}^{\text{SYN}}$). For GTA5-trained models, for both the ResNet-based and the Transformer-based backbones, we clearly achieve a new state-of-the-art (SOTA) performance for the DG mean. For the ResNet-based models, we improve significantly on Cityscapes and Mapillary. For SYNTHIA-trained models (Table 2) our method provides an even larger improvement. With a ResNet-based backbone we outperform other approaches by 7.4% absolute for the DG mean and by 11.8% abs. with a transformer backbone. On Mapillary Vistas, our method provides 14.7% abs. mIoU improvement compared with the best performing prior method [45]. With 59.8% mIoU on Cityscapes, DIDEX performs competitive with the UDA method DAFormer [14] with 60.9% mIoU *without using any real data during training*.

5.2. Influence of Prompting Strategy

Table 3 shows the results of the resulting domain generalization when varying the prompts w.r.t. location, environment condition, and class-uniform sampling. Overall,

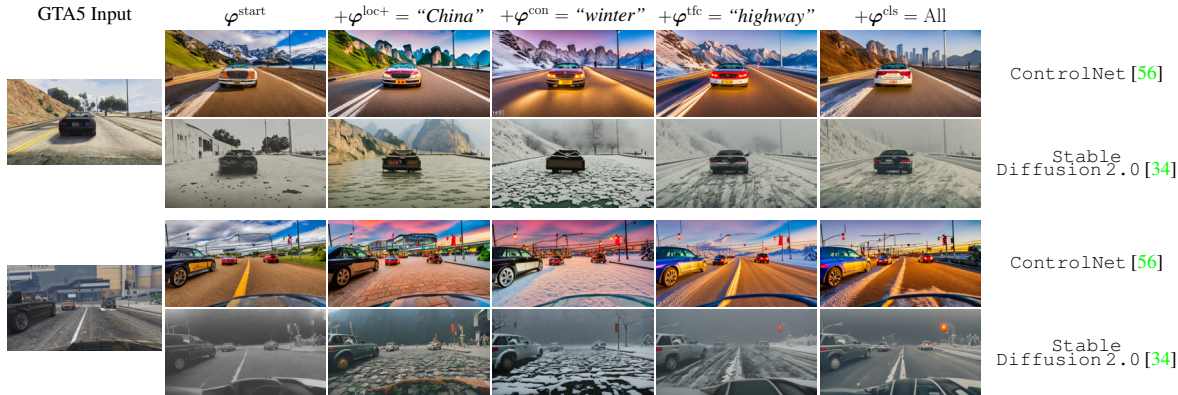


Figure 4. **Visualization of the prompt ablation study** showing the impact of the prompt on the images of \mathcal{D}^{PT} for ControlNet [60] and SD 2.0 [38]. Each prompt has a corresponding impact on the content of the generated images. However, semantic inconsistencies can be seen such as trees turning into mountains or buildings which disappear. More examples are provided in the supplement.

Table 3. **Influence of the different text prompt building blocks** on the generalization performance (mIoU(%)). **Training** was performed on the synthetic **GTA5** ($\mathcal{D}^{\text{S}} = \mathcal{D}_{\text{train}}^{\text{GTA5}}$) dataset. **Evaluation** is performed on various real-world **validation sets** ($\mathcal{D}^{\text{T}} = \mathcal{D}_{\text{val}}$). The adaptation was performed with the HRDA method [15] and a Transformer-based encoder.

Base Prompt (Φ^{base})	Add. Location ($\varphi^{\text{loc}} \rightarrow \varphi^{\text{loc+}}$)	+ Conditions ($+\varphi^{\text{con}}$)	+ CUS ($+\varphi^{\text{CUS}}$)	mIoU (%) on				
				$\mathcal{D}_{\text{val}}^{\text{CS}}$	$\mathcal{D}_{\text{val}}^{\text{BDD}}$	$\mathcal{D}_{\text{val}}^{\text{MV}}$	$\mathcal{D}_{\text{val}}^{\text{ACDC}}$	DG mean
✓	–	–	–	58.5	52.2	62.9	46.9	57.9
✓	✓	–	–	58.7	52.5	63.4	46.7	58.2
✓	–	✓	–	59.4	52.7	62.7	46.8	58.3
✓	–	–	✓	61.2	52.5	63.7	48.8	59.1
✓	✓	–	✓	58.6	51.8	62.8	45.2	57.7
✓	–	✓	✓	60.1	53.7	63.5	46.6	59.1
✓	✓	✓	✓	58.8	52.7	63.2	47.4	58.3

the most important part of the text prompts seems to be the class-uniform sampling ($+\varphi^{\text{CUS}}$) as it results in the highest domain generalization mIoU. But increasing the variation in each of the dimensions seems to have a positive effect, although the increase in performance is rather small for $\varphi^{\text{loc}} \rightarrow \varphi^{\text{loc+}}$ and $+\varphi^{\text{con}}$. Increasing the variety of locations has the biggest effect on the performance on the Mapillary Vistas dataset [28], which consists of images from various locations all over the world.

However, increasing the number of conditions does not increase performance on MV or ACDC which are very diverse w.r.t. the conditions. This might be related to insufficiently realistic generation of such conditions by the diffusion model, or related to the fact that the HRDA [15] method has difficulties to adapt to these conditions. Finally, we observe that using all text prompts does not improve over only class-uniform sampling.

5.3. Influence of UDA Approaches

Table 4 shows the influence of the UDA approach that is used for the generalization by adaptation step to the pseudo-target domain. Recent UDA methods, such as

DAFormer [14] and MIC [17] adapt better to the pseudo-target domain and thus generalize better across domains. It should be noted that even comparably simple UDA methods such as DACS [47] obtain a high domain generalization and outperform previous SOTA methods for SYNTHIA as the source dataset. However, SePiCo [51] with a ResNet-101 backbone performs worse than the other methods which might be caused by the lack of hyperparameter optimization.

5.4. Influence of Image Quantity & Consistency

In Figure 5, we show results of randomly sampled subsets of either Cityscapes (purple) or our pseudo-target domain (orange) used in the UDA approach. For each step, we applied DAFormer [14] for the adaptation from GTA5 to the subset. We can observe that the model adapted to Cityscapes is on par with the pseudo-target domain model on Cityscapes up until a subset size of 29 UDA images, but gets increasingly better with more sampled images. This divergence of performance does, however, not occur for the other domains (BDD, MV, ACDC), where the models are mostly on par with each other. This indicates that the in-

Table 4. **Influence of DIDEX combined different UDA methods** on the generalization performance (mIoU(%)). **Training** was performed on the synthetic **GTA5** ($\mathcal{D}^S = \mathcal{D}_{\text{train}}^{\text{GTA5}}$, upper part) or **SYNTIA** ($\mathcal{D}^S = \mathcal{D}_{\text{train}}^{\text{SYN}}$, lower part) dataset. **Evaluation** is performed on various real-world **validation sets** ($\mathcal{D}^T = \mathcal{D}_{\text{val}}$). Text prompts comprised the base prompt Φ^{base} and the CUS φ^{CUS} . * indicates additional usage of the masked image consistency loss [17].

	Enc.	DIDEX + ...	mIoU (%) on				
			$\mathcal{D}_{\text{val}}^{\text{CS}}$	$\mathcal{D}_{\text{val}}^{\text{BDD}}$	$\mathcal{D}_{\text{val}}^{\text{MV}}$	$\mathcal{D}_{\text{val}}^{\text{ACDC}}$	DG mean
\mathcal{D}^S : GTA5	ResNet	DACS* [47]	46.9	40.0	45.2	31.9	44.0
		SePiCo [51]	44.9	36.4	38.8	30.4	40.0
		DAFormer* [14]	50.4	41.8	47.1	33.9	46.4
		MIC [17]	52.4	40.9	49.2	36.1	47.5
	Transformer	DACS* [47]	52.0	49.0	53.7	41.9	51.8
		SePiCo [51]	57.4	49.7	56.4	44.5	54.5
		DAFormer* [14]	57.7	56.6	60.7	46.4	58.3
		MIC [17]	62.0	54.3	63.0	50.1	59.7
\mathcal{D}^S : SYNTIA	ResNet	DACS* [47]	47.5	38.6	41.5	30.1	42.5
		SePiCo [51]	43.3	32.6	40.6	27.7	38.8
		DAFormer* [14]	49.8	40.0	45.5	33.7	45.1
		MIC [17]	53.1	41.8	50.3	33.3	48.4
	Transformer	DACS* [47]	52.1	38.4	48.0	36.4	46.2
		SePiCo [51]	54.4	45.5	52.3	37.7	50.7
		DAFormer* [14]	53.3	44.5	52.3	38.6	50.0
		MIC [17]	59.8	47.4	59.5	43.5	55.6

crease in performance on Cityscapes is caused by a specialization of the model to the target domain, but does not increase its generalization capabilities relative to the model adapted to the pseudo-target domain. It further indicates that the images in the pseudo-target domain are of a similar "functional quality" for generalizing to other real-world domains as the real-world Cityscapes data.

In Table 5, the influence of different ways to constrain the image generation process is shown. We observe that ControlNet [60] constrained with semantics provides a small improvement over SD2.0 [38] and can be a promising future research direction.

6. Conclusions

We introduced a novel diffusion-based domain extension (DIDEX) method for domain generalization which utilizes the generative capabilities of diffusion models. We projected the problem of domain generalization to the problem of domain adaptation which enabled us to utilize powerful adaptation methods for domain generalization. Our diffusion-based domain extension outperforms previous state-of-the-art methods by a large margin across datasets and architectures; with GTA5 as the source dataset




Figure 5. **Influence of the # of UDA images** (target domain) on the generalization performance (mIoU (%)). **Training** was performed on the synthetic **GTA5** ($\mathcal{D}^S = \mathcal{D}_{\text{train}}^{\text{GTA5}}$) dataset. **Evaluation** is performed on various real-world **validation sets** ($\mathcal{D}^T = \mathcal{D}_{\text{val}}$). DAFormer [14] was used for UDA and text prompts comprised only the base prompt Φ^{base} and the CUS block φ^{CUS} .

Table 5. **Influence of the diffusion model Υ architecture and constraints $C(x_n)$** on the generalization performance (mIoU (%)). **Training** was performed on the synthetic **GTA5** ($\mathcal{D}^S = \mathcal{D}_{\text{train}}^{\text{GTA5}}$) dataset. **Evaluation** is performed on various real-world **validation sets** ($\mathcal{D}^T = \mathcal{D}_{\text{val}}$). UDA was performed with HRDA [15] for SD2.0, MIC [17] for the others, and a full prompt Ω_n .

Diffusion Model Υ	$C(x_n)$	mIoU (%) on				
		$\mathcal{D}_{\text{val}}^{\text{CS}}$	$\mathcal{D}_{\text{val}}^{\text{BDD}}$	$\mathcal{D}_{\text{val}}^{\text{MV}}$	$\mathcal{D}_{\text{val}}^{\text{ACDC}}$	DG mean
SD2.0 [38]	Depth	58.8	52.7	63.2	47.4	58.3
ControlNet [60]	Depth	58.1	54.0	63.7	46.7	58.6
ControlNet [60]	Seg.	58.9	53.9	64.1	49.0	59.0

by 3.8% abs. mIoU and with SYNTIA even by 11.8% abs. mIoU on average. Our text prompt ablation study has shown that information about present classes is beneficial for the pseudo-target domain. A remarkable result is that the functional quality of the diffusion generated data for the purpose of domain generalization is comparable to real Cityscapes data, highlighting the potential of using diffusion models for domain generalization.

Acknowledgment

This work was supported by the German Federal Ministry for Economic Affairs and Climate Action within the project "KI Wissen – Entwicklung von Methoden für die Einbindung von Wissen in maschinelles Lernen" and the SynthBAD project.

References

- [1] Yasser Benigmim, Subhankar Roy, Slim Essid, Vicky Kalogeiton, and Stéphane Lathuilière. One-Shot Unsupervised Domain Adaptation With Personalized Diffusion Models. In *Proc. of CVPR - Workshops*, pages 698–708, Vancouver, BC, Canada, June 2023. [2](#)
- [2] Qi Bi, Shaodi You, and Theo Gevers. Learning Content-enhanced Mask Transformer for Domain Generalized Urban-Scene Segmentation. *arXiv:2307.00371*, pages 1–18, 2023. [2](#), [3](#)
- [3] Jan-Aike Bolte, Markus Kamp, Antonia Breuer, Silviu Homocanu, Peter Schlicht, Fabian Hüger, Daniel Lipinski, and Tim Fingscheidt. Unsupervised Domain Adaptation to Improve Image Segmentation Quality Both in the Source and Target Domain. In *Proc. of CVPR - Workshops*, pages 1404–1413, Long Beach, CA, USA, June 2019. [1](#)
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. DeepLab: Semantic Image Segmentation With Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 40(4):834–848, Apr. 2017. [5](#)
- [5] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T. Kim, Seungryong Kim, and Jaegul Choo. RobustNet: Improving Domain Generalization in Urban-Scene Segmentation via Instance Selective Whitening. In *Proc. of CVPR*, pages 11580–11590, virtual, June 2021. [2](#)
- [6] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *Proc. of CVPR*, pages 3213–3223, Las Vegas, NV, USA, June 2016. [5](#), [12](#)
- [7] Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In *Proc. of NeurIPS*, pages 8780–8794, virtual, Dec. 2021. [2](#), [3](#)
- [8] Jiahua Dong, Yang Cong, Gan Sun, Zhen Fang, and Zhengming Ding. Where and How to Transfer: Knowledge Aggregation-Induced Transferability Perception for Unsupervised Domain Adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pages 1–17, 2021. [2](#)
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *Proc. of ICLR*, pages 1–21, virtual, May 2021. [1](#)
- [10] Rui Gong, Martin Danelljan, Han Sun, Julio Delgado Mangas, and Luc Van Gool. Prompting Diffusion Representations for Cross-Domain Semantic Segmentation. *arXiv:2307.02138*, pages 1–17, 2023. [3](#), [5](#)
- [11] Rui Gong, Wen Li, Yuhua Chen, and Luc Van Gool. DLOW: Domain Flow for Adaptation and Generalization. In *Proc. of CVPR*, pages 2477–2486, Long Beach, CA, USA, June 2019. [2](#)
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *Proc. of NeurIPS*, pages 6840–6851, virtual, Dec. 2020. [2](#), [3](#)
- [13] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Philip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. CyCADA: Cycle-Consistent Adversarial Domain Adaptation. In *Proc. of ICML*, pages 1989–1998, Stockholm, Sweden, July 2018. [1](#), [2](#)
- [14] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. DAFormer: Improving Network Architectures and Training Strategies for Domain-Adaptive Semantic Segmentation. In *Proc. of CVPR*, pages 9924–9935, New Orleans, LA, USA, June 2022. [1](#), [2](#), [6](#), [7](#), [8](#)
- [15] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. HRDA: Context-Aware High-Resolution Domain-Adaptive Semantic Segmentation. In *Proc. of ECCV*, pages 372–391, Tel Aviv, Israel, Oct. 2022. [2](#), [6](#), [7](#), [8](#)
- [16] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Domain Adaptive and Generalizable Network Architectures and Training Strategies for Semantic Image Segmentation. *arXiv:2304.13615*, pages 1–15, Apr. 2023.
- [17] Lukas Hoyer, Dengxin Dai, Haoran Wang, and Luc Van Gool. MIC: Masked Image Consistency for Context-Enhanced Domain Adaptation. In *Proc. of CVPR*, pages 11721–11732, Vancouver, BC, Canada, June 2023. [2](#), [6](#), [7](#), [8](#)
- [18] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. FSDR: Frequency Space Domain Randomization for Domain Generalization. In *Proc. of CVPR*, pages 6891–6902, virtual, June 2021. [2](#), [5](#)
- [19] Xun Huang and Serge Belongie. Arbitrary Style Transfer in Real-Time with Adaptive Instance Normalization. In *Proc. of ICCV*, pages 1501–1510, Venice, Italy, Oct. 2017. [2](#)
- [20] Namyup Kim, Taeyoung Son, Cuiling Lan, Wenjun Zeng, and Suha Kwak. WEDGE: Web-Image Assisted Domain Generalization for Semantic Segmentation. *arXiv:2109.14196*, pages 1–14, Sept. 2021. [2](#)
- [21] Marvin Klingner, Mouadh Ayache, and Tim Fingscheidt. Continual BatchNorm Adaptation (CBNA) for Semantic Segmentation. *IEEE Transactions on Intelligent Transportation Systems*, 23(11):20899–20911, 2022. [1](#)
- [22] Marvin Klingner, Jan-Aike Termöhlen, Jacob Ritterbach, and Tim Fingscheidt. Unsupervised BatchNorm Adaptation (UBNA): A Domain Adaptation Method for Semantic Segmentation Without Using Source Domain Representations. In *Proc. of WACV - Workshops*, pages 210–220, Waikoloa, HI, USA, Jan. 2022. [1](#), [5](#)
- [23] Suhyeon Lee, Hongje Seong, Seongwon Lee, and Euntai Kim. WildNet: Learning Domain Generalized Semantic Segmentation from the Wild. In *Proc. of CVPR*, pages 9936–9946, New Orleans, LA, USA, June 2022. [2](#), [5](#), [6](#)
- [24] Ruihuang Li, Shuai Li, Chenhang He, Yabin Zhang, Xu Jia, and Lei Zhang. Class-Balanced Pixel-Level Self-Labeling for Domain Adaptive Semantic Segmentation. In *Proc. of CVPR*, pages 11593–11603, New Orleans, LA, USA, June 2022. [2](#)
- [25] Tianyu Li, Subhankar Roy, Huayi Zhou, Hongtao Lu, and Stéphane Lathuilière. Contrast, Stylize and Adapt: Unsuper-

- vised Contrastive Learning Framework for Domain Adaptive Semantic Segmentation. In *Proc. of CVPR*, pages 4868–4878, Vancouver, BC, Canada, June 2023. [2](#)
- [26] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional Learning for Domain Adaptation of Semantic Segmentation. In *Proc. of CVPR*, pages 6936–6945, Long Beach, CA, USA, June 2019. [2](#)
- [27] Robert A. Marsden, Alexander Bartler, Mario Döbler, and Bin Yang. Contrastive Learning and Self-Training for Unsupervised Domain Adaptation in Semantic Segmentation. In *Proc. of IJCNN*, pages 1–8, Padua, Italy, July 2022. [1](#)
- [28] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kontschieder. The Mapillary Vistas Dataset for Semantic Understanding of Street Scenes. In *Proc. of ICCV*, pages 4990–4999, Venice, Italy, Oct. 2017. [5](#), [7](#), [12](#)
- [29] Joshua Niemeijer, Sudhanshu Mittal, and Thomas Brox. Synthetic Dataset Acquisition for a Specific Target Domain. In *Proc. of ICCV - Workshops*, pages 4055–4064, Paris, France, Oct. 2023. [1](#)
- [30] Joshua Niemeijer and Jörg Peter Schäfer. Domain Adaptation and Generalization: A Low-Complexity Approach. In *Proc. of CoRL*, Auckland, NZ, Dec. 2022. [1](#)
- [31] Joshua Niemeijer and Jörg P. Schäfer. Combining Semantic Self-Supervision and Self-Training for Domain Adaptation in Semantic Segmentation. In *Proc. of IV - Workshops*, pages 364–371, Nagoya, Japan, July 2021. [1](#)
- [32] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at Once: Enhancing Learning and Generalization Capacities via IBN-Net. In *Proc. of ECCV*, pages 464–479, Munich, Germany, Sept. 2018. [2](#)
- [33] Xingang Pan, Xiaohang Zhan, Jianping Shi, Xiaoou Tang, and Ping Luo. Switchable Whitening for Deep Representation Learning. In *Proc. of ICCV*, pages 1863–1871, Seoul, Korea, Oct. 2019. [2](#)
- [34] Duo Peng, Yinjie Lei, Munawar Hayat, Yulan Guo, and Wen Li. Semantic-Aware Domain Generalized Segmentation. In *Proc. of CVPR*, pages 2594–2605, New Orleans, LA, USA, June 2022. [2](#), [5](#)
- [35] Duo Peng, Yinjie Lei, Lingqiao Liu, Pingping Zhang, and Jun Liu. Global and Local Texture Randomization for Synthetic-to-Real Semantic Segmentation. *IEEE Transactions on Image Processing*, 30:6594–6608, 2021. [2](#), [5](#)
- [36] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. *arXiv:2307.01952*, pages 1–21, July 2023. [3](#)
- [37] Stephan Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for Data: Ground Truth from Computer Games. In *Proc. of ECCV*, pages 102–118, Amsterdam, Netherlands, Oct. 2016. [5](#)
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proc. of CVPR*, pages 10684–10695, New Orleans, LA, USA, June 2022. [2](#), [4](#), [5](#), [7](#), [8](#), [12](#), [15](#), [16](#)
- [39] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M. Lopez. The SYNTHIA Dataset: A Large Collection of Synthetic Images for Semantic Segmentation of Urban Scenes. In *Proc. of CVPR*, pages 3234–3243, Las Vegas, NV, USA, June 2016. [5](#)
- [40] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The Adverse Conditions Dataset with Correspondences for Semantic Driving Scene Understanding. In *Proc. of ICCV*, pages 10765–10775, virtual, Oct. 2021. [5](#), [6](#), [12](#)
- [41] Manuel Schwonberg, Fadoua El Bouazati, Nico M Schmidt, and Hanno Gottschalk. Augmentation-Based Domain Generalization for Semantic Segmentation. In *Proc. of IV - Workshops*, pages 1–8, Anchorage, AK, USA, June 2023. [3](#)
- [42] Manuel Schwonberg, Joshua Niemeijer, Jan-Aike Termöhlen, Jörg P. Schäfer, Nico M. Schmidt, Hanno Gottschalk, and Tim Fingscheidt. Survey on Unsupervised Domain Adaptation for Semantic Segmentation for Visual Perception in Automated Driving. *IEEE Access*, 11:54296–54336, May 2023. [1](#), [2](#)
- [43] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. In *Proc. of ICLR*, pages 1–20, virtual, May 2021. [2](#), [3](#)
- [44] Qiyu Sun, Pavlo Melnyk, Michael Felsberg, and Yang Tang. Augment Features Beyond Color for Domain Generalized Segmentation. *arXiv:2307.01703*, pages 1–10, July 2023. [2](#), [3](#), [5](#)
- [45] Jan-Aike Termöhlen, Timo Bartels, and Tim Fingscheidt. A Re-Parameterized Vision Transformer (ReVT) for Domain-Generalized Semantic Segmentation. In *Proc. of ICCV - Workshops*, pages 4376–4385, Paris, France, Oct. 2023. [6](#)
- [46] Jan-Aike Termöhlen, Marvin Klingner, Leon J. Brettin, Nico M. Schmidt, and Tim Fingscheidt. Continual Unsupervised Domain Adaptation for Semantic Segmentation by Online Frequency Domain Style Transfer. In *Proc. of ITSC*, pages 2881–2888, virtual, Sept. 2021. [1](#)
- [47] Wilhelm Tranheden, Viktor Olsson, Juliano Pinto, and Lennart Svensson. DACS: Domain Adaptation via Cross-Domain Mixed Sampling. In *Proc. of WACV*, pages 1379–1389, virtual, Jan. 2021. [6](#), [7](#)
- [48] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to Adapt Structured Output Space for Semantic Segmentation. In *Proc. of CVPR*, pages 7472–7481, Salt Lake City, UT, USA, June 2018. [1](#)
- [49] Zhonghao Wang, Mo Yu, Yunchao Wei, Rogerio Feris, Jinjun Xiong, Wen-Mei Hwu, Thomas S. Huang, and Honghui Shi. Differential Treatment for Stuff and Things: A Simple Unsupervised Domain Adaptation Method for Semantic Segmentation. In *Proc. of CVPR*, pages 12635–12644, Seattle, WA, USA, June 2020. [2](#)
- [50] Zuxuan Wu, Xintong Han, Yen-Liang Lin, Mustafa Gokhan Uzunbas, Tom Goldstein, Ser Nam Lim, and Larry S. Davis. DCAN: Dual Channel-Wise Alignment Networks for Unsupervised Scene Adaptation. In *Proc. of ECCV*, pages 535–552, Munich, Germany, Sept. 2018. [2](#)
- [51] Binhui Xie, Shuang Li, Mingjia Li, Chi Harold Liu, Gao Huang, and Guoren Wang. SePiCo: Semantic-Guided Pixel Contrast for Domain Adaptive Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 45(7), Jan. 2023. [1](#), [2](#), [6](#), [7](#)

- [52] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. In *Proc. of NeurIPS*, pages 12077–12090, virtual, Dec. 2021. [1](#), [2](#)
- [53] Qi Xu, Liang Yao, Zhengkai Jiang, Guannan Jiang, Wenqing Chu, Wenhui Han, Wei Zhang, Chengjie Wang, and Ying Tai. DURL: Domain-Invariant Representation Learning for Generalizable Semantic Segmentation. In *Proc. of AAAI*, pages 2884–2892, virtual, June 2022. [2](#)
- [54] Yonghao Xu, Fengxiang He, Bo Du, Dacheng Tao, and Liangpei Zhang. Self-Ensembling GAN for Cross-Domain Semantic Segmentation. *IEEE Transactions on Multimedia*, Dec. 2022. [2](#)
- [55] Yanchao Yang and Stefano Soatto. FDA: Fourier Domain Adaptation for Semantic Segmentation. In *Proc. of CVPR*, pages 4085–4095, Seattle, WA, USA, June 2020. [1](#)
- [56] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. BDD100K: A Diverse Driving Dataset for Heterogeneous Multitask Learning. In *Proc. of CVPR*, pages 1–14, Seattle, WA, USA, June 2020. [12](#)
- [57] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. BDD100K: A Diverse Driving Video Database With Scalable Annotation Tooling. *arXiv*, (1805.04687), Aug. 2018. [5](#)
- [58] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain Randomization and Pyramid Consistency: Simulation-to-Real Generalization Without Accessing Target Domain Data. In *Proc. of ICCV*, pages 2100–2110, Seoul, Korea, Oct. 2019. [2](#), [5](#)
- [59] Feihu Zhang, Vladlen Koltun, Philip Torr, René Ranftl, and Stephan R. Richter. Unsupervised Contrastive Domain Adaptation for Semantic Segmentation. *arXiv:2204.08399*, pages 1–12, Apr. 2022. [2](#)
- [60] Lvmin Zhang and Maneesh Agrawala. Adding Conditional Control to Text-to-Image Diffusion Models. *arXiv:2302.05543*, pages 1–33, Feb. 2023. [2](#), [3](#), [4](#), [5](#), [7](#), [8](#), [12](#), [13](#), [14](#)
- [61] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. Prototypical Pseudo Label Denoising and Target Structure Learning for Domain Adaptive Semantic Segmentation. In *Proc. of CVPR*, pages 12414–12424, virtual, 2021. [2](#)
- [62] Yuyang Zhao, Zhun Zhong, Na Zhao, Nicu Sebe, and Gim Hee Lee. Style-Hallucinated Dual Consistency Learning for Domain Generalized Semantic Segmentation. In *Proc. of ECCV*, pages 535–552, Tel Aviv, Israel, 2022. [3](#)
- [63] Zhun Zhong, Yuyang Zhao, Gim Hee Lee, and Nicu Sebe. Adversarial Style Augmentation for Domain Generalized Urban-Scene Segmentation. In *Proc. of NeurIPS*, pages 338–350, New Orleans, LA, USA, Dec. 2022. [3](#)
- [64] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proc. of ICCV*, pages 2223–2232, Venice, Italy, Oct. 2017. [2](#)