

Collaborating Foundation Models for Domain Generalized Semantic Segmentation

Yasser Benigmim^{1,2} Subhankar Roy³ Slim Essid¹ Vicky Kalogeiton² Stéphane Lathuilière¹
¹ LTCI, Télécom-Paris, Institut Polytechnique de Paris
² LIX, Ecole Polytechnique, CNRS, Institut Polytechnique de Paris, ³ University of Aberdeen
 yasser.benigmim@telecom-paris.fr

Abstract

Domain Generalized Semantic Segmentation (DGSS) deals with training a model on a labeled source domain with the aim of generalizing to unseen domains during inference. Existing DGSS methods typically effectuate robust features by means of Domain Randomization (DR). Such an approach is often limited as it can only account for style diversification and not content. In this work, we take an orthogonal approach to DGSS and propose to use an assembly of CoLLaborative FOundation models for Domain Generalized Semantic Segmentation (CLOUDS). In detail, CLOUDS is a framework that integrates Foundation Models of various kinds: (i) CLIP backbone for its robust feature representation, (ii) Diffusion Model to diversify the content, thereby covering various modes of the possible target distribution, and (iii) Segment Anything Model (SAM) for iteratively refining the predictions of the segmentation model. Extensive experiments show that our CLOUDS excels in adapting from synthetic to real DGSS benchmarks and under varying weather conditions, notably outperforming prior methods by 5.6% and 6.7% on averaged mIoU, respectively. Our code is available at <https://github.com/yasserben/CLOUDS>

1. Introduction

Deep Neural Networks have showcased remarkable ability in scene understanding tasks like Semantic Segmentation (SS) [7, 70, 72], when the training and test distribution are the same. This dependency reveals a significant vulnerability: their performance substantially diminishes when encountering domain shifts [66], highlighting a fundamental challenge in generalizing these networks to unseen domains [43, 46, 61]. To address this, Domain Generalized Semantic Segmentation (DGSS) has gained prominence [28, 31, 33, 82]. DGSS aims to develop models that leverage a source-annotated dataset while remaining effec-




Figure 1. Performance over time by various methods on the GTA \rightarrow {Cityscapes, BDD, Mapillary} benchmark. Recent open-vocabulary approaches, like FC-CLIP, are shown to excel in zero-shot learning and surpass traditional domain generalization methods trained in closed-set scenarios, challenging the relevance of the DGSS setting. CLOUDS, by harnessing multiple foundation models, demonstrates its ability to effectively utilize the source dataset, thereby outperforming both conventional DGSS and open-vocabulary methods.

tive across unseen domains, thus overcoming the limitations of traditional DNNs in handling unseen environments.

Recently, the advent of large-scale pretrained models, often referred to as Foundation Models (FMs) [2, 35, 56, 59, 67], have brought a paradigm shift in computer vision tasks. The FMs comprise of contrastively trained image classification models (e.g., CLIP [56], ALIGN [30]), text-conditioned generative models (e.g., Stable Diffusion [59], DALL-E [57]), vision transformer-based segmentation model trained on mammoth dataset (e.g., Segment Anything Model (SAM) [35]), to name a few. Of particular interest to SS, SAM is a promptable segmentation model that accepts as input an image and geometric prompts (points, scribbles, boxes, masks) and outputs class-agnostic masks. Owing to its pre-training on billion-scale dataset [35], SAM has demonstrated excellent performance on varied out-of-distribution tasks namely, medical imaging



Figure 2. **Qualitative comparison at inference:** (a) Input image, (b) SHADE, a traditional style diversification DGSS method, (c) SAM, a foundation model that predicts precise class-agnostic maps, (d) GroundingSAM that leverages SAM and text-prompts to output semantic maps, and (e) Our proposed CLOUDS that leverages an assembly of foundation models to predict high-quality semantic maps.

[22, 44, 71], crater detection [19], and so on. In summary, large-scale pre-training holds promise for DGSS as robustness is greatly improved when a model is trained on large datasets that cover various possible distributions [79].

While the FMs present an excellent solution for addressing DGSS, they have not been adopted so far in the literature. A great majority of DGSS methods rely on the technique of Domain Randomization (DR) [55, 68, 74, 82], where the goal is to diversify the labeled source domain images by photometric and geometric image transformations. Examples of such transformations include style diversification [82, 83], adversarial style augmentation [83], and difficulty-aware photometric augmentations [31], among others [33, 78]. Albeit effective to some extent, these techniques can not bring content augmentation, resulting in sub-optimal performance (see Fig. 2b). Moreover, adopting SAM for the task of DGSS is not straightforward as it outputs class-agnostic masks, making it unsuitable for DGSS (see Fig. 2c). While there are recent works, such as Grounding-SAM [35, 41], that equip SAM with semantic predictions, it misses out on several important objects (e.g., “poles”) and *stuff* classes (e.g., “tree”, “sidewalk”) (Fig. 2d). The lack of semantic awareness can be attributed to ambiguous text prompts or poor vision and text alignment. These results indicate that DGSS is far from solved and accommodating FMs in DGSS is an open question.

Driven by the motivation of narrowing down the research gap, in this work we propose to amalgamate an *assembly* of FMs into a cohesive system to address DGSS. In detail, we propose to exploit: (i) strong representations of contrastive FMs, (ii) diverse content augmentation using generative FMs, and (iii) near accurate class-agnostic mask prediction of SAM and turn them into semantic predictions. Concretely, for the segmentation network we use CLIP [56] as a backbone, serving as a robust feature extractor. To introduce content diversification, we generate synthetic images using a combination of **Large Language Models** (LLMs) [67] and text-to-image **Diffusion Model** (DM) [59]. The LLM generates task-aware diverse textual prompts which then guide the DM in producing photorealistic synthetic images. Since the generated images are unlabeled, we employ self-training strategy [24, 80, 84, 85] where we use the pseudo labels from the teacher network

to train the student model. However, as the pseudo labels from the teacher can be noisy, we leverage **SAM**’s excellent class-agnostic mask predictions to refine the pseudo labels. As shown in Fig. 2e, the pseudo-label refinement greatly assists in improving the reliability of self-training.

In summary, in this work we make the following contributions: (i) We highlight for the first time in DGSS the importance of content diversification, which is more effective than traditional DGSS methods relying on style diversification; (ii) We propose CLOUDS, a system of collaborative Foundation models for DGSS. We run extensive experiments on several DGSS benchmarks and demonstrate its effectiveness. As shown in Fig. 1, our FM-based DGSS method outperforms the traditional DGSS and open-vocabulary segmentation models by a non-trivial margin.

2. Related Works

Domain Generalization (DG). DG focuses on training models robust against domain shift [47, 69]. In the field of Domain Generalized Semantic Segmentation (DGSS), methods can be classified into two main categories. The first involves integrating tailor-made modules and transformations to explicitly eliminate domain-specific features [12, 50, 51, 53], such as IBN-Net [50] that uses instance normalization blocks and ISW [12] that applies whitening transformations. The second category relies on Domain Randomization (DR) by diversifying image styles [31, 52, 65, 81, 82]. SHADE [82] generates new styles from basis styles of the source domain, and MoDify [31] leverages difficulty-aware photometric augmentations. HRDA [26], designed for Unsupervised Domain Adaptation [25], shows remarkable performance in DGSS, combining the VIT-based model [17] and multiple training strategies for SS. While our approach aligns with DR, it differs in two aspects: (i) CLOUDS emphasizes data generation over stylization for greater diversification, and (ii) the use of generated data enables the introduction of a self-training strategy which, to our knowledge, has never been used for DGSS.

Foundation Models for Segmentation. Foundation models [2, 35, 56, 59] have recently garnered significant interest, and their application across a wide range of

downstream tasks constitutes an active area of research.

Diffusion Models (DM), a key type among generative models [4, 14, 59, 67], stand out in producing photorealistic images, thereby enhancing synthetic dataset generation and boosting vision task performance [1, 49, 63]. Yet, this approach remains underexplored in DGSS, a gap our paper aims to address. Diffusion models have also been effectively used for discriminative tasks [20, 32, 73]. In particular, the U-Net architecture [40] of these models, known for its rich and robust image representation, is particularly well suited for pixel-level prediction tasks like semantic segmentation. ODISE [73], for instance, adopts this approach for open-vocabulary semantic segmentation. The study closely related to ours [20] investigates both DGSS and test-time domain adaptation using a diffusion-based backbone.

CLIP also offers robust feature representations and image-text alignment, which have been instrumental in various vision applications [10, 11, 16, 20, 73, 77]. CLIP has emerged as a crucial component in open-vocabulary segmentation models which segment unseen classes via the introduction of a masked self-distillation mechanism [16] or a novel cost aggregation layer [11]. Drawing inspiration from this success, our work employs a CNN-based [37] CLIP backbone, as recommended in [77].

Segment Anything Model (SAM) [35], a prominent vision foundation model, is trained for promptable segmentation tasks. SAM excels in producing high-quality masks for any segmentation prompt. Various studies [6, 35, 38, 41] have proposed solutions to address SAM’s primary limitation of generating class-agnostic outputs. One notable approach is GroundingSAM [35, 41], which synergizes an open-set object detector with SAM to yield labeled masks. In contrast, our work proposes leveraging SAM to enhance pseudo labels derived from a generated dataset.

Large Language Models. LLMs have also impacted computer vision in many ways [3, 45, 75]. For instance, CuPL [54] leverages the knowledge of GPT-3 [4] to generate rich text descriptions which are prompted to CLIP and improve performance on the zero-shot image classification task. In CLOUDS, we use an LLM to increase diversity in the textual prompts conditioning the generation of images using a diffusion model, offering a notably more cost-effective approach than collecting and curating real data.

3. Collaborating Foundation Models

The goal of Domain Generalized Semantic Segmentation (DGSS) is to train a segmentation model $g \circ f$, where f corresponds to the feature extractor and g the decoder, on a source domain \mathcal{S} such that it can generalize to any unseen domain \mathcal{T} that follows a different distribution. We assume having access only to a labeled source domain

$\mathcal{S} = \{(\mathbf{x}_i^S, \mathbf{y}_i^S)\}_{i=1}^{N^S}$ of size N^S where $\mathbf{x}_i^S \in \mathbb{R}^{H \cdot W \cdot 3}$ represents an RGB source image and $\mathbf{y}_i^S \in \{0, 1\}^{H \cdot W \cdot C}$ the corresponding one-hot encoded ground-truth label. The proposed approach shown in Fig. 3, leverages the power of established foundation models to advance on DGSS.

Overview: CLOUDS uses a CNN-based CLIP backbone, harnessing its capability to extract robust features (Section 3.1). To enable Domain Randomization (DR) with content diversity, we employ a text-to-image diffusion model, which generates photo-realistic images while being conditioned on textual prompts which are also generated using a Large Language Model to broaden the diversity of generated images used for self-training (Section 3.2). Lastly, the predicted pseudo labels are refined with the help of the Segment-Anything Model (SAM) in a student-teacher fashion (Section 3.3). This collaborative strategy aims to fortify the model’s generalizability across diverse domains.

3.1. CLIP-based Student-Teacher Network

The segmentation model is composed of an encoder E and a decoder D . For the encoder E , we use a CNN-based CLIP instead of a ViT-based one, as it has been demonstrated in [77] that it produces better semantic features, and performs better on high-resolution images for semantic segmentation. For the decoder D , we employ the one from Mask2Former [9], a refined version of MaskFormer [8]. This model redefines the task of semantic segmentation by treating it as mask classification. It achieves this by dissociating the division of the image into regions from their subsequent classification, thereby stepping aside from the conventional per-pixel classification in semantic segmentation. The decoder comprises two components: a pixel decoder, and a transformer decoder, which perform mask prediction and classification. More details about the architecture can be found in [9].

Training procedure: We initially train our model on the labeled source dataset \mathcal{S} , while keeping the backbone f frozen as this offers: Firstly, it leads to faster and more efficient training. Secondly, it leverages pre-established, robust feature representations which ensure a better generalization while not overfitting on the source domain. Lastly, it preserves the image-text-aligned representations.

Thus, we only optimize the model’s decoder and employ the loss function used in Mask2Former, which is composed of two parts: a mask loss $\mathcal{L}_{\text{mask}}$ and a classification loss \mathcal{L}_{cls} . The mask loss $\mathcal{L}_{\text{mask}}$ is defined as a linear combination of binary cross-entropy and dice losses. The total supervised loss for the source domain, denoted as \mathcal{L}_S , is formulated as:

$$\mathcal{L}_S = \mathcal{L}_{\text{mask}} + \lambda_{\text{cls}} \mathcal{L}_{\text{cls}} \quad , \quad (1)$$

where λ_{cls} is a factor balancing the two loss components.

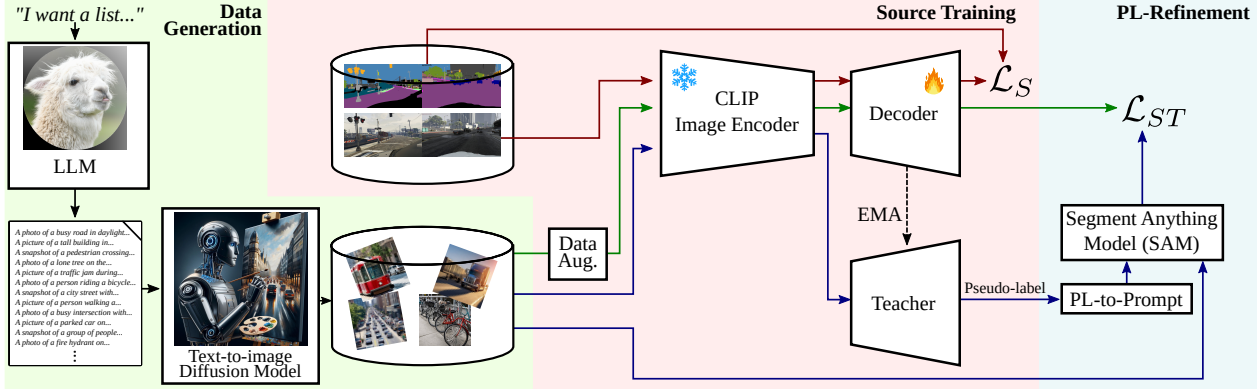


Figure 3. **Training pipeline of CLOUDS:** The model integrates a CLIP image encoder with a MaskFormer decoder (Sec. 3.1). Our domain randomization strategy is based on a data generalization module (Sec. 3.2) that combines a Large Language Model (LLM) with a text-to-image diffusion model to generate a varied dataset, representative of potential target datasets. This data is then employed in a Self-Training framework (Sec. 3.3), where initial pseudo labels (PL) prompt the Segment Anything Model (SAM) for refined pseudo labels, thereby fortifying the decoder’s robustness.

Self-Training: The generated data are leveraged via a self-training strategy that trains the model using its own predictions, referred to as *pseudo labels*. The pseudo labels are continuously updated during the training process, allowing them to evolve and become more accurate as the model’s predictions improve. Specifically, our approach employs a student-teacher framework to ensure training stability. The encoder remains frozen, which means our student-teacher framework only involves the decoder. The decoder D (Section 3.1) assumes the student’s role. Concurrently, the teacher network D_T is derived from the Exponential Moving Average (EMA) of the student’s parameters, as in [26]. The update rule for the teacher’s parameters θ_T is:

$$\theta'_T \leftarrow \alpha\theta_T + (1 - \alpha)\theta \quad , \quad (2)$$

where θ represents the student’s parameters, and α is the factor controlling the teacher’s momentum. Following [26], data augmentation is applied to the student branch to enhance the robustness of the student decoder. Meanwhile, the teacher branch receives the original images, facilitating the generation of more reliable pseudo labels.

3.2. Domain Randomization with Generative FMs

We use text-conditioned DMs for generating images and leverage LLMs for creating the text-prompts.

Diffusion models: To achieve domain randomization with high content diversity, we leverage a pretrained text-conditional diffusion model, as it excels in generating photorealistic images. Specifically, we employ the latent diffusion model [59] Stable Diffusion trained on LAION-5B [64]. We provide the model with textual prompts describing scenes relevant to potential future target domains.

In our experiments, we focus on prompts depicting urban street scenes, aligning with our benchmark scenarios. Note, this is in line with real-world scenarios, where we may not know the exact test environment, but we know the generic deployment setup (e.g. indoor, outdoor, or driving settings). To obtain diverse prompts without effort, we rely on a LLM.

LLM for diversifying prompts: Given their semantic diversity capabilities, we leverage LLMs (Llama-2 [67]), to enrich prompts for image generation. Specifically, we employ an LLM to both create a wide range of synonyms for predefined class names and increase the semantic compositionality of the scene by varying descriptions of environmental factors such as lighting and weather conditions. Our goal is to generate prompts formatted as ‘a photo of X in Z’, where ‘X’ represents any class name from our source dataset or its synonym, and ‘Z’ encompasses any contextual information describing the environment. This leverages prior knowledge of LLMs, which possess an understanding of the contexts in which objects typically appear. Thus, we employ the LLM to enable the creation of synthetic target data, mirroring plausible future target environments.

To obtain such prompts, we prompt the LLM with the following text: *I want a list of prompts that can be used by an image generation model to generate synthetic images [...] The prompt should strictly follow this template: “a photo of X in Z” where X contains one or multiple class names within C [...]. Can you provide 100 diverse and simple prompts.* Where \mathbf{C} is replaced by the list of all the class names of our source dataset.

3.3. Pseudo Label refinement using FMs

The pseudo labels generated by the teacher decoder demonstrate prediction inaccuracies. For their refinement, we use




Figure 4. Pseudo Label refinement with SAM. We extract binary masks from the predicted segmentation. After labeling connected components and filtering noisy ones, we select random points within each binary mask. These points, along with the corresponding RGB images, are then used to prompt SAM, enabling the generation of more accurate segmentation maps

SAM [35], capable of generating multiple detailed masks for any image within a zero-shot framework. SAM, originally trained for promptable segmentation, is adapted to our task through the use of point-based prompting (Fig. 4).

Specifically, we adopt the following methodology: Given a generated image x , we sequentially derive its pseudo labels using the CLIP encoder and the teacher decoder $p = D_T(E(x))$. These pseudo labels are assembled into C class-specific binary masks m_i , $1 \leq i \leq C$, each representing one of the predefined classes. For each class i , the mask m_i might encompass multiple objects, as they are not differentiated by instances. To resolve this, we apply the Hoshen–Kopelman algorithm, known for its efficiency in labeling connected components in masks [23]. Consequently, for every class i segmented in the image, we obtain M_i distinct binary masks. Additionally, to further refine the pseudo labels, a filtering process is implemented to remove regions falling below a set size threshold.

Then, we randomly select points within each binary mask, utilizing these points to prompt SAM. For each set of points, SAM provides an enhanced binary mask. The final enhanced pseudo label map is obtained by aggregating all the predicted objects across all classes. SAM may predict overlapping masks for points from different classes. In such instances, the intersecting pixels are categorized as an “unlabeled” class to prevent the student model’s training with erroneous pseudo labels.

The Student model is then trained using these refined pseudo labels. The self-training loss \mathcal{L}_{st} for generated images is the same as the source loss \mathcal{L}_S of the source dataset.

4. Experiments

4.1. Experimental Setups

Datasets: To evaluate the model’s adaptability to new environments, we train on a single source domain and test on multiple unseen domains. We assess two scenarios: (a) Synthetic-to-Real domain generalization, where we rely on

GTA [58], which comprises 24,966 images at a resolution of 1914×1052, and SYNTHIA [60] which contains 9,400 images at a resolution of 1280×760. In the real-world domain, we use Cityscapes [13], featuring 2,975 training images and 500 validation images at a resolution of 2048×1024. We also consider BDD [76] which involves 1,000 validation images at a resolution of 1280×720, and Mapillary [48] including 2,000 validation images across diverse resolutions. For brevity, we denote ‘C’ for Cityscapes, ‘B’ for BDD100K, and ‘M’ for Mapillary. (b) Clear-to-adverse weather scenarios, where we incorporate the ACDC dataset [62] containing 406 validation images of 1920×1080 resolution.

Implementation details: The network is trained for 40K iterations on GTA alone for initial reliable pseudo labels and an additional 40K iterations on both GTA and the generated dataset. On SYNTHIA, a smaller dataset, each step consists of 30K iterations. For Cityscapes, an even smaller dataset, we use 10K iterations for each step. The learning rate is $lr = 1 \times 10^{-4}$ for Cityscapes and $lr = 1 \times 10^{-5}$ for GTA5 and SYNTHIA and the batch size is 8. For other hyperparameters, we follow Mask2Former [9], using the same loss functions with AdamW [34] optimizer with a decay of 0.05. We use a random scaling in the range of [0.5,2.0] and random cropping of 768×768 size, and then we apply data augmentation (random flipping, color jittering). Our assessment involves ResNet-50 [21] and ResNet-101 [21] encoders pretrained on both ImageNet [15] and WIT [56] from OpenAI.¹ We also use a ConvNext-Large encoder[42, 56] pretrained on the LAION-2B [64] from OpenCLIP [29]. For the decoder, we use the one from Mask2Former [9].

Evaluation protocol: We measure model performance using the mean Intersection over Union (mIoU). For GTA, mIoU is computed across 19 classes; for SYNTHIA, we report mIoU for 16 shared classes with Cityscapes, BDD, and Mapillary. We consistently evaluate the model’s mIoU using the final trained model across all experiments.

4.2. State-of-the-art comparison

We compare CLOUDS against both DGSS methods and methods based on FMs. We first describe the compared methods and then report and analyze the mIoU results in two settings: training on GTA and SYNTHIA, when the target domains are Cityscapes, BDD100K, and Mapillary. Similar to prior work in DGSS [31, 33, 82], we assess the performance using ResNet-50, ResNet-101, and MiT-B5. We also extend our evaluation to include ConvNext-Large.

Compared Methods: For *traditional DGSS*, we include all leading DGSS methods: IBN-Net [50], ISW [12], SHADE [82], TLDR [33] and MoDify [31]. We also com-

¹<https://github.com/openai/CLIP>

	Method	Encoder	Pre-training	Training	Target Domains (mIoU in %)			Avg.
					Cityscapes	BDD100K	Mapillary	
Domain Generalization	IBN-Net [50]	ResNet-50	ImageNet	GTA5	33.9	32.3	37.8	34.6
	ISW [12]				36.6	35.2	40.3	37.4
	SHADE [82]				44.7	39.3	43.3	42.4
	TLDR [33]				46.5	42.6	46.2	45.1
	MoDify [31]				45.7	40.1	46.2	44.0
	CLOUDS (Ours)				54.6	46.7	58.6	53.3
	IBN-Net [50]	ResNet-101	ImageNet	GTA5	37.4	34.2	36.8	36.1
	ISW [12]				37.2	33.4	35.6	35.4
	SHADE [82]				46.6	43.7	45.5	45.3
	TLDR [33]				47.6	44.9	48.8	47.1
	MoDify [31]				48.8	44.2	47.5	46.8
	HRDA * [26]				39.6	38.7	42.2	40.1
	CLOUDS (Ours)				50.6	44.8	56.6	50.7
	CLOUDS (Ours)				55.7	49.3	59.0	54.7
HRDA [26]	MiT-B5	ImageNet	GTA5	57.4	49.1	61.1	55.9	
CLOUDS (Ours)				58.1	53.8	62.3	58.1	
Foundation Models	PTDiffSeg [20]	Diffusion	LAION-5B	GTA5	52.0	–	–	–
	Grounding-SAM [35, 41]	ViT-H	SA-1B	–	43.5	39.4	48.4	43.8
	CAT-Seg [11]	ViT-G/14		COCO-Stuff [5]	45.0	47.6	51.8	48.2
	ODISE [73]	Diffusion		COCO Panoptic [39]	53.8	53.6	59.1	55.5
	FC-CLIP [77]	ConvNeXt-L	LAION-2B	COCO Panoptic [39]	56.2	54.2	60.6	57.0
	FC-CLIP * [77]	ConvNeXt-L		GTA5	53.6	47.6	57.4	52.9
	CLOUDS (Ours)	ConvNext-L		GTA5	60.2	57.4	67.0	61.5

Table 1. Comparison with state of the art on GTA \rightarrow {Cityscapes, BDD, Mapillary} with leading DGSS methods and foundation models. * denotes experiment obtained using the official code. † denotes Web Image-Text dataset used to train CLIP model from OpenAI.

pare against HRDA [26], initially designed for Unsupervised Domain Adaptation and adjusted for DGSS in [25].

For *Foundation models*, we compare against three categories: (i) SAM-based [35] (GroundingSAM [35, 41]), (ii) CLIP-based[56] (FC-CLIP[77], CATSeg [11]), and (iii) Diffusion-based[59] (ODISE[73], PTDiffSeg [20]) methods. (i) For SAM-based methods, GroundingSAM merges SAM’s segmentation power with the GroundingDINO open-set object detector that uses text prompts to predict bounding boxes. SAM encodes these boxes to predict their corresponding semantic mask. (ii) For CLIP-based methods, FC-CLIP uses a frozen CNN-based backbone to extract semantic features for both mask generation and CLIP classification. CAT-Seg uses cost aggregation to finetune CLIP image embeddings for better generalization and robustness to unseen domains. (iii) For Diffusion-based methods, PTDiffSeg [20] is tailored for DGSS and uses the feature representations of a diffusion model for semantic segmentation. We present results only on Cityscapes, as the paper does not report results on other datasets and the code is not available. While ODISE (55.5%) falls into this category, we note that

during inference, to improve its performance it uses CLIP.

GTA as source domain: Table 1 reports the results when benchmarking GTA \rightarrow {C, B, M}. Overall, we observe that CLOUDS outperforms both all DGSS methods and FM-based methods for all types of encoders and pre-training setups. For traditional DGSS methods, when using ResNet-50, CLOUDS achieves 53.3%, i.e. +8.1% compared to the second competitor MoDify, whereas with ResNet-101, it reaches 54.7%, improving the benchmark by +7.6%. When using ConvNext-L, CLOUDS achieves 61.5%, outperforming the previously strongest DGSS model HRDA (55.9%) by +5.5%. Additionally, CLOUDS notably outperforms HRDA by 10% when using ResNet-101 as backbone and by 3% when using MiT-B5 [72]. These results show that our method efficiently leverages several FMs, leading to improved features and enhanced training.

Furthermore, CLOUDS consistently outperforms all methods based on FMs, on all three datasets, achieving an average mIoU of 61.5%. GroundingSAM alongside open-vocabulary methods (ODISE, FC-CLIP, and CAT-Seg) face

a significant limitation due to their reliance on text prompts, which leads to ambiguity when it comes to segmenting classes that can hardly be described with words only (e.g., classes "pole" and "terrain"). Text prompts alone fall short in fully capturing all visual elements, particularly in complex scenes like autonomous driving, as the same class may cover diverse objects with different characteristics. Instead, CLOUDS leverages image data by training on the source domain and harnessing visual cues that are essential for accurately segmenting these complex classes.

Method	Encoder	C	B	M	Avg
DRPC [78]	ResNet-50	35.7	31.5	32.7	33.3
SAN-SAW [53]		38.9	35.2	34.5	36.2
MoDify [31]		38.9	33.7	36.2	36.3
TLDR [33]		41.9	34.4	36.8	37.7
CLOUDS (Ours)		46.1	37.6	48.1	43.9
DRPC [78]	ResNet-101	37.6	34.4	34.1	35.3
GTR [52]		39.7	35.3	36.4	37.1
FSDR [27]		40.8	37.4	39.6	39.3
SAN-SAW [53]		40.9	36.0	37.3	38.0
TLDR [33]		42.6	35.5	37.5	38.5
HRDA * [26]		34.9	25.0	34.0	31.3
MoDify [31]		43.4	39.5	42.3	41.7
CLOUDS (Ours)		49.1	40.3	50.1	46.5
HRDA * [26]	MiT-B5	39.6	32.6	40.0	37.4
CLOUDS (Ours)		42.2	38.3	43.6	41.4
FC-CLIP * [77]	ConvNext-L	38.0	29.9	39.0	35.6
CLOUDS (Ours)		53.4	47.0	55.8	52.1

Table 2. Comparison with state-of-the-art methods for DGSS on Synthia \rightarrow {Cityscapes (C), BDD (B), Mapillary (M)}. * denotes experiment obtained using the official code

SYNTHIA as source domain: In Table 2, we address the setting SYNTHIA \rightarrow {C, B, M}. CLOUDS achieves 43.9% using ResNet-50 and outperforms the previously leading method (TLDR) by +6.2%. When using ResNet-101, CLOUDS achieves 46.5% and improves the previously best method (MoDify) by +4.8%. Moreover, CLOUDS outperforms HRDA with MiT-B5 backbone by +4%. When using ConvNext-L we obtain the best results, i.e. 52.1% mIoU, outperforming FC-CLIP by a large margin.

Comparison with zero-shot Domain Adaptation: Zero-shot Domain Adaptation assumes having access to the annotated source domain and a domain description of the target using natural language, in the form of a text prompt. The task of prompt-driven zero-shot domain adaptation was introduced by PØDA [18], a method that harnesses CLIP’s robust feature representations. It adjusts source feature statistics to align with the target domain, guided by natural language domain prompts. To ensure a fair comparison with PØDA [18], we employed the same ResNet-50 backbone

Method	Encoder	Night	Snow	Rain
CLIPStyler [36]	ResNet-50	21.3	41.0	38.7
PØDA [18]		25.0	43.9	42.3
CLOUDS (Ours)		29.4	52.1	49.7
CLOUDS (Ours)	ResNet-101	33.0	51.3	53.4
	ConvNext-L	45.1	65.3	64.4

Table 3. Comparison with zero-shot Domain Adaptation methods. evaluation on ACDC dataset.

pretrained on CLIP. The results in Table 3 demonstrate that our method consistently outperforms PØDA, achieving improvements of +4.4%, +8.2% and +7.4% across three scenarios: day \rightarrow night, clear \rightarrow snow and clear \rightarrow rain, respectively. We also provide results using ResNet-101 pretrained on CLIP and ConvNext-L on LAION-2B.

4.3. Ablation studies

We ablate the various components of CLOUDS. We focus our evaluation on the setting GTA \rightarrow {C, B, M}.

Backbone	CLIP	{LLM, Diffusion}	SAM	Avg
ResNet-50	✓			50.0
	✓	✓		50.7
	✓	✓	✓	53.3
ResNet-101	✓			51.9
	✓	✓		53.3
	✓	✓	✓	54.7
ConvNext-L	✓			58.5
	✓	✓		58.6
	✓	✓	✓	61.5

Table 4. Ablation of the key components of CLOUDS on GTA \rightarrow {Cityscapes, BDD, Mapillary} using three different backbones.

Effect of each Foundation model: We investigate the impact of integrating various FMs on the performance of DGSS. Table 4 reports the results, i.e., average mIoU across three diverse datasets: C, B, and M. Notably, employing CLIP as a standalone feature extractor with various backbones, including ResNet-50, ResNet-101 and ConvNext-L, yields state-of-the-art results, reaching an average mIoU of 50.0%, 51.9% and 58.5%, respectively. This impressive performance underscores CLIP’s rich feature representation, hence leading to robustness in handling unseen environments. Including the diffusion model conditioned by the LLM-generated prompt brings marginal improvements with ConvNext-L and ResNet-50 backbones. This can be attributed to the noise in the pseudo labels predicted by the Teacher model, which comes from the domain shift in the generated images produced by the text-to-image model,




Figure 5. Effect of generated dataset size on mIoU. Experiments performed on GTA \rightarrow {Cityscapes, BDD, Mapillary}.

thus posing difficulties in self-training. Adding SAM leads to the largest performance gains, i.e. approx. +3% in mIoU for all models. This affirms the critical role of SAM in refining pseudo labels and enhancing their accuracy before using them in supervision. This gain is evident in all backbones, highlighting SAM’s effectiveness in addressing the challenges of self-training.

Effect of Scaling: Here, we examine the impact of the generated dataset size on the performance. Figure 5 shows that increasing the size of the generated dataset improves the mIoU significantly (up to +3%) up approximately 5K images, where a plateau is reached. Therefore, in our experiments, we use a dataset size of 5000 images.




Figure 6. Ablation Study on SAM Prompting: This study examines the effects of varying the number of sampled points and the threshold used to separate connected components.

Prompting SAM: We examine how various prompts used in SAM affect the performance, Fig. 6 reports the results. The left barplot displays the impact of varying the number of points used to prompt SAM for predicting masks. Increasing them leads to a decrease in mIoU. This is most likely due to the noise in pseudo labels and the possibility of points being placed on different objects on the image, resulting in less accurate masks, which leads to noisy pseudo labels. The right barplot illustrates the relationship between

the threshold size used in the Hoshen–Kopelman algorithm (which separates the connected components of the binary mask) and the performance of our model. A higher threshold helps eliminate small regions that might be incorrectly identified as objects belonging to a certain class, thereby enhancing the precision of the pseudo labels. Therefore, the model’s performance improves.

Backbone	Cityscapes	BDD100K	Mapillary	Avg.
Trainable	58.6	53.0	62.8	58.1
Frozen	60.2	57.4	67.0	61.5

Table 5. Impact of fine-tuning CLIP on the performance.

Fine-tuning the CLIP encoder: Here, we examine the impact of fine-tuning CLIP on the performance. Table 5 reports the results. We observe that keeping the backbone frozen during training, our method effectively retains the encoder’s ability to generalize and perform well in unseen environments. Contrarily, finetuning the CLIP encoder resulted in poorer performance, likely due to overfitting on the source domain, which consequently led to diminished generalization capabilities on the target domains.

5. Conclusion

In this work, we proposed CLOUDS, a novel approach to Domain Generalized Semantic Segmentation that uniquely integrates various Foundation Models in a collaborative manner. By leveraging the robust features of CLIP to unseen domains and the content diversification capabilities of diffusion models, enriched by LLM-generated text prompts, and employing the ability of SAM to enhance pseudo labels for self-training, we effectively address key challenges in DGSS. Our approach introduces a self-training strategy using generated data. Through extensive experimentation, CLOUDS demonstrates superior performance over traditional DGSS and open-vocabulary segmentation models on various benchmarks, marking a significant advancement in the field. This work not only bridges a crucial research gap but also establishes a new standard for robustness and adaptability in semantic segmentation across diverse domains in the era of Foundation Models.

Acknowledgements. This work has been supported by the French National Research Agency (ANR) with the ANR-20-CE23-0027 project and was granted access to the HPC resources of IDRIS under the allocation AD011013071 made by GENCI. We would like to thank T.Delatolas and I.Marouf for proofreading.

References

- [1] Yasser Benigim, Subhankar Roy, Slim Essid, Vicky Kalogeiton, and Stéphane Lathuilière. One-shot unsupervised domain adaptation with personalized diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 698–708, 2023. [3](#)
- [2] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. [1](#), [2](#)
- [3] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. [3](#)
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. [3](#)
- [5] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Cocostuff: Thing and stuff classes in context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1209–1218, 2018. [6](#)
- [6] Jiaqi Chen, Zeyu Yang, and Li Zhang. Semantic segment anything. <https://github.com/fudan-zvg/Semantic-Segment-Anything>, 2023. [3](#)
- [7] 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*, 40(4):834–848, 2017. [1](#)
- [8] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. [3](#)
- [9] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. [3](#), [5](#)
- [10] Junhyeong Cho, Gilhyun Nam, Sungyeon Kim, Hunmin Yang, and Suha Kwak. Promptstyler: Prompt-driven style generation for source-free domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15702–15712, 2023. [3](#)
- [11] Seokju Cho, Heeseong Shin, Sunghwan Hong, Seungjun An, Seungjun Lee, Anurag Arnab, Paul Hongsuck Seo, and Seungryong Kim. Cat-seg: Cost aggregation for open-vocabulary semantic segmentation. *arXiv preprint arXiv:2303.11797*, 2023. [3](#), [6](#)
- [12] 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 *CVPR*, pages 11580–11590, 2021. [2](#), [5](#), [6](#)
- [13] 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 *CVPR*, pages 3213–3223, 2016. [5](#)
- [14] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1):53–65, 2018. [3](#)
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [5](#)
- [16] Xiaoyi Dong, Jianmin Bao, Yinglin Zheng, Ting Zhang, Dongdong Chen, Hao Yang, Ming Zeng, Weiming Zhang, Lu Yuan, Dong Chen, et al. Maskclip: Masked self-distillation advances contrastive language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10995–11005, 2023. [3](#)
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [2](#)
- [18] Mohammad Fahes, Tuan-Hung Vu, Andrei Bursuc, Patrick Pérez, and Raoul de Charette. Poda: Prompt-driven zero-shot domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18623–18633, 2023. [7](#)
- [19] Iraklis Giannakis, Anshuman Bhardwaj, Lydia Sam, and Georgios Leontidis. Deep learning universal crater detection using segment anything model (sam). *arXiv preprint arXiv:2304.07764*, 2023. [2](#)
- [20] Rui Gong, Martin Danelljan, Han Sun, Julio Delgado Mangas, and Luc Van Gool. Prompting diffusion representations for cross-domain semantic segmentation. *arXiv preprint arXiv:2307.02138*, 2023. [3](#), [6](#)
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [5](#)
- [22] Sheng He, Rina Bao, Jingpeng Li, P Ellen Grant, and Yangming Ou. Accuracy of segment-anything model (sam) in medical image segmentation tasks. *arXiv preprint arXiv:2304.09324*, 2023. [2](#)
- [23] Joseph Hoshen and Raoul Kopelman. Percolation and cluster distribution. i. cluster multiple labeling technique and critical concentration algorithm. *Physical Review B*, 14:3438–3445, 1976. [5](#)
- [24] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9924–9935, 2022. [2](#)

- [25] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Hrda: Context-aware high-resolution domain-adaptive semantic segmentation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXX*, pages 372–391. Springer, 2022. **2, 6**
- [26] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Domain adaptive and generalizable network architectures and training strategies for semantic image segmentation. *T-PAMI*, 2023. **2, 4, 6, 7**
- [27] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Fsd: Frequency space domain randomization for domain generalization. In *CVPR*, pages 6891–6902, 2021. **7**
- [28] Lei Huang, Yi Zhou, Fan Zhu, Li Liu, and Ling Shao. Iterative normalization: Beyond standardization towards efficient whitening. In *CVPR*, pages 4874–4883, 2019. **1**
- [29] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, et al. Openclip. *Zenodo*, 4:5, 2021. **5**
- [30] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. **1**
- [31] Xueying Jiang, Jiaying Huang, Sheng Jin, and Shijian Lu. Domain generalization via balancing training difficulty and model capability. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18993–19003, 2023. **1, 2, 5, 6, 7**
- [32] Laurynas Karazija, Iro Laina, Andrea Vedaldi, and Christian Rupprecht. Diffusion models for zero-shot open-vocabulary segmentation. *arXiv preprint arXiv:2306.09316*, 2023. **3**
- [33] Sunghwan Kim, Dae-hwan Kim, and Hoseong Kim. Texture learning domain randomization for domain generalized segmentation. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. **1, 2, 5, 6, 7**
- [34] Diederik P Kingma and Jimmy Ba. A method for stochastic optimization. In *International conference on learning representations (ICLR)*, page 6, 2015. **5**
- [35] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. **1, 2, 3, 5, 6**
- [36] Gihyun Kwon and Jong Chul Ye. Clipstyler: Image style transfer with a single text condition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18062–18071, 2022. **7**
- [37] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. **3**
- [38] Feng Li, Hao Zhang, Peize Sun, Xueyan Zou, Shilong Liu, Jianwei Yang, Chunyuan Li, Lei Zhang, and Jianfeng Gao. Semantic-sam: Segment and recognize anything at any granularity. *arXiv preprint arXiv:2307.04767*, 2023. **3**
- [39] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. **6**
- [40] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. **3**
- [41] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. **2, 3, 6**
- [42] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022. **5**
- [43] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2507–2516, 2019. **1**
- [44] Jun Ma and Bo Wang. Segment anything in medical images. *arXiv preprint arXiv:2304.12306*, 2023. **2**
- [45] Liliane Momeni, Mathilde Caron, Arsha Nagrani, Andrew Zisserman, and Cordelia Schmid. Verbs in action: Improving verb understanding in video-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15579–15591, 2023. **3**
- [46] Saeid Motiian, Quinn Jones, Seyed Iranmanesh, and Gianfranco Doretto. Few-shot adversarial domain adaptation. *Advances in neural information processing systems*, 30, 2017. **1**
- [47] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International conference on machine learning*, pages 10–18. PMLR, 2013. **2**
- [48] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Buló, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, pages 4990–4999, 2017. **5**
- [49] Joshua Niemeijer, Manuel Schwonberg, Jan-Aike Termöhlen, Nico M. Schmidt, and Tim Fingscheidt. Generalization by adaptation: Diffusion-based domain extension for domain-generalized semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2830–2840, 2024. **3**
- [50] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 464–479, 2018. **2, 5, 6**
- [51] Xingang Pan, Xiaohang Zhan, Jianping Shi, Xiaoou Tang,

- and Ping Luo. Switchable whitening for deep representation learning. In *ICCV*, pages 1863–1871, 2019. [2](#)
- [52] Duo Peng, Yinjie Lei, Lingqiao Liu, Pingping Zhang, and Jun Liu. Global and local texture randomization for synthetic-to-real semantic segmentation. *IEEE TIP*, pages 6594–6608, 2021. [2](#), [7](#)
- [53] Duo Peng, Yinjie Lei, Munawar Hayat, Yulan Guo, and Wen Li. Semantic-aware domain generalized segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2594–2605, 2022. [2](#), [7](#)
- [54] Sarah Pratt, Ian Covert, Rosanne Liu, and Ali Farhadi. What does a platypus look like? generating customized prompts for zero-shot image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15691–15701, 2023. [3](#)
- [55] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *CVPR*, pages 12556–12565, 2020. [2](#)
- [56] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [1](#), [2](#), [5](#), [6](#)
- [57] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. [1](#)
- [58] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, pages 102–118. Springer, 2016. [5](#)
- [59] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [1](#), [2](#), [3](#), [4](#), [6](#)
- [60] 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 *CVPR*, pages 3234–3243, 2016. [5](#)
- [61] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Computer Vision—ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part IV 11*, pages 213–226. Springer, 2010. [1](#)
- [62] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Acde: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10765–10775, 2021. [5](#)
- [63] Mert Bulent Sariyildiz, Kartteek Alahari, Diane Larlus, and Yannis Kalantidis. Fake it till you make it: Learning transferable representations from synthetic imagenet clones. In *CVPR 2023—IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. [3](#)
- [64] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. [4](#), [5](#)
- [65] Zhiqiang Tang, Yunhe Gao, Yi Zhu, Zhi Zhang, Mu Li, and Dimitris N Metaxas. Selfnorm and crossnorm for out-of-distribution robustness. 2020. [2](#)
- [66] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR 2011*, pages 1521–1528. IEEE, 2011. [1](#)
- [67] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. [1](#), [2](#), [3](#), [4](#)
- [68] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. *Advances in neural information processing systems*, 31, 2018. [2](#)
- [69] Riccardo Volpi, Diane Larlus, and Grégory Rogez. Continual adaptation of visual representations via domain randomization and meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4443–4453, 2021. [2](#)
- [70] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2020. [1](#)
- [71] Junde Wu, Rao Fu, Huihui Fang, Yuanpei Liu, Zhaowei Wang, Yanwu Xu, Yueming Jin, and Tal Arbel. Medical sam adapter: Adapting segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.12620*, 2023. [2](#)
- [72] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. SegFormer: Simple and efficient design for semantic segmentation with transformers. *NeurIPS*, pages 12077–12090, 2021. [1](#), [6](#)
- [73] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023. [3](#), [6](#)
- [74] Qinwei Xu, Ruipeng Zhang, Ya Zhang, Yanfeng Wang, and Qi Tian. A fourier-based framework for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14383–14392, 2021. [2](#)
- [75] Yue Yang, Artemis Panagopoulou, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, and Mark Yatskar. Language in a bottle: Language model guided concept bottlenecks

- for interpretable image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19187–19197, 2023. 3
- [76] 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 *CVPR*, pages 2636–2645, 2020. 5
- [77] Qihang Yu, Ju He, Xueqing Deng, Xiaohui Shen, and Liang-Chieh Chen. Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip. *arXiv preprint arXiv:2308.02487*, 2023. 3, 6, 7
- [78] 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 *ICCV*, pages 2100–2110, 2019. 2, 7
- [79] Giacomo Zara, Alessandro Conti, Subhankar Roy, Stéphane Lathuilière, Paolo Rota, and Elisa Ricci. The unreasonable effectiveness of large language-vision models for source-free video domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10307–10317, 2023. 2
- [80] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12414–12424, 2021. 2
- [81] Yuyang Zhao, Zhun Zhong, Zhiming Luo, Gim Hee Lee, and Nicu Sebe. Source-free open compound domain adaptation in semantic segmentation. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(10):7019–7032, 2022. 2
- [82] Yuyang Zhao, Zhun Zhong, Na Zhao, Nicu Sebe, and Gim Hee Lee. Style-hallucinated dual consistency learning for domain generalized semantic segmentation. In *ECCV*, pages 535–552. Springer, 2022. 1, 2, 5, 6
- [83] Zhun Zhong, Yuyang Zhao, Gim Hee Lee, and Nicu Sebe. Adversarial style augmentation for domain generalized urban-scene segmentation. *Advances in Neural Information Processing Systems*, 35:338–350, 2022. 2
- [84] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305, 2018. 2
- [85] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5982–5991, 2019. 2