

Textual Query-Driven Mask Transformer for Domain Generalized Segmentation

Byeonghyun Pak* Byeongju Woo* Sunghwan Kim*
Dae-hwan Kim Hoseong Kim[†]

Agency for Defense Development (ADD)
{byeonghyun_pak, woobj903, ssshwan, dhkim7, hoseongkim}@add.re.kr

Abstract. In this paper, we introduce a method to tackle Domain Generalized Semantic Segmentation (DGSS) by utilizing domain-invariant semantic knowledge from text embeddings of vision-language models. We employ the text embeddings as object queries within a transformer-based segmentation framework (textual object queries). These queries are regarded as a domain-invariant basis for pixel grouping in DGSS. To leverage the power of textual object queries, we introduce a novel framework named the **t**extual **Q**uery-**d**riven **M**ask transformer (**tqdm**). Our **tqdm** aims to (1) generate textual object queries that maximally encode domain-invariant semantics and (2) enhance the semantic clarity of dense visual features. Additionally, we suggest three regularization losses to improve the efficacy of **tqdm** by aligning between visual and textual features. By utilizing our method, the model can comprehend inherent semantic information for classes of interest, enabling it to generalize to extreme domains (*e.g.*, sketch style). Our **tqdm** achieves 68.9 mIoU on GTA5→Cityscapes, outperforming the prior state-of-the-art by 2.5 mIoU. The project page is at <https://byeonghyunpak.github.io/tqdm>.

Keywords: Domain Generalized Semantic Segmentation · Leveraging Vision-Language Models · Transformer-Based Segmentation

1 Introduction

Developing a model that generalizes robustly to unseen domains has been a long-standing goal in the field of machine perception. In this context, Domain Generalized Semantic Segmentation (DGSS) aims to build models that can effectively operate across diverse target domains, trained solely on a single source domain. This area has made notable improvements with a wide range of approaches, including normalization and whitening [7, 42–44], domain randomization [12, 23, 24, 59, 63], and utilizing the inherent robustness of transformers [10, 50].

Meanwhile, the recent advent of Vision-Language Models (VLMs) (*e.g.*, CLIP [46]) has introduced new possibilities and applications in various vision tasks,

*Equal contribution [†]Corresponding author

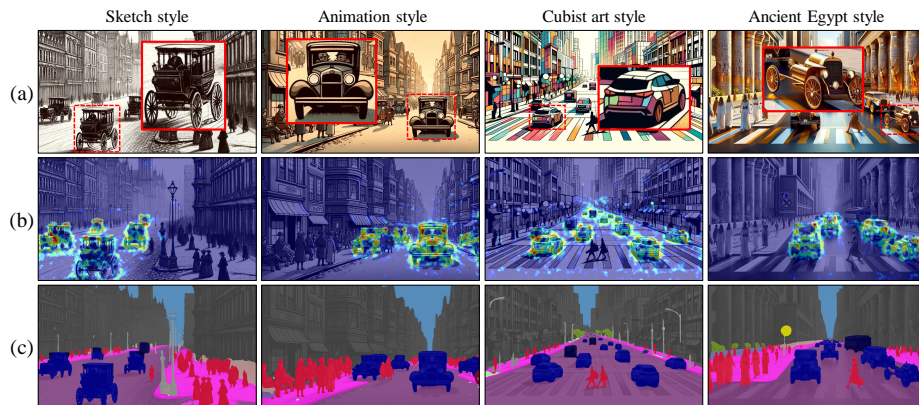


Fig. 1: (a) A collection of driving scene images with diverse styles generated by ChatGPT¹. (b) The image-text similarity maps of a pre-trained VLM (*i.e.*, EVA02-CLIP [51]) on diverse domains. The text embedding of ‘car’ is consistently well-aligned with the corresponding class regions of images across various domains. (c) The segmentation results predicted by our proposed **tqdm**. Note that our model can generalize to extreme domains (*e.g.*, sketch style) and effectively identify the cars in multiple forms that are not present in the source domain (*i.e.*, GTA5 [48]).

thanks to their rich semantic representations [68]. One notable strength of VLMs is their ability to generalize across varied domain shifts [40, 46]. This capability has inspired the development of methods for domain generalization in image classification [6, 19, 36, 55]. Furthermore, there have been efforts to incorporate the robust visual representation of VLMs in DGSS [12, 20]. However, existing methods in DGSS have not yet explored direct language-driven approaches to utilize textual representations from VLMs for domain-generalized recognition. Note that the contrastive learning objective in VLMs aligns a text caption (*e.g.*, “a photo of a car”) with images from a wide range of domains in a joint space [46]. This alignment enables the text embeddings to capture domain-invariant semantic knowledge [19], as demonstrated in Fig. 1

In this paper, we introduce a language-driven DGSS method that harnesses domain-invariant semantics from textual representations in VLMs. Our proposed method can make accurate predictions even under extreme domain shifts, as it comprehends the inherent semantics of targeted classes. In Fig. 1c, our model exhibits accurate segmentation results on driving scene images with diverse styles generated by ChatGPT¹. Note that our model can effectively identify cars in multiple forms that are not present in the source domain (*i.e.*, GTA5 [48]).

The key idea of our method is to utilize text embeddings of classes of interest from VLMs as object queries, referred to as *textual object queries*. Given that an object query can be considered as a mask embedding vector to group regions belonging to the same class [4], textual object queries generate robust mask predictions for classes of interest across diverse domains, thanks to their

¹<https://chat.openai.com>

domain-invariant semantics (Sec. 3). Building on this insight, we propose a textual query-driven mask transformer (**tqdm**) to leverage textual object queries for DGSS. Our design philosophy lies in (1) generating the queries that maximally encode domain-invariant semantic knowledge and (2) enhancing their adaptability in dense predictions by improving the semantic clarity of pixel features (Sec. 4.1). Additionally, we discuss three regularization losses to preserve the robust vision-language alignment of pre-trained VLMs, thereby improving the effectiveness of our method (Sec. 4.2).

Our contributions can be summarized into three aspects:

- To the best of our knowledge, we are the first to introduce a direct language-driven DGSS approach using text embeddings from VLMs to enable domain-invariant recognition, effectively handling extreme domain shifts.
- We propose a novel query-based segmentation framework named **tqdm** that leverages textual object queries, along with three regularization losses to support this framework.
- Our **tqdm** demonstrates the state-of-the-art performance across multiple DGSS benchmarks; *e.g.*, **tqdm** achieves 68.9 mIoU on GTA5→Cityscapes, which improves the prior state-of-the-art method by 2.5 mIoU.

2 Related Work

Vision-Language Models (VLMs). VLMs [9, 21, 35, 45, 46, 52], which are trained on extensive web-based image-caption datasets, have gained attention for their rich semantic understanding. CLIP [46] employs contrastive language-image pre-training, and several studies [40, 46] have explored its robustness to natural distribution shifts. EVA02-CLIP [13, 14] further enhances CLIP by exploring robust dense visual features through masked image modeling [16].

The advanced capabilities of VLMs enable more challenging segmentation tasks. For example, open-vocabulary segmentation [25, 30, 56, 67] attempts to segment an image by arbitrary categories described in texts. Although these works use text embeddings for segmentation tasks similar to our approach, our work distinctly differs in the problem of interest and the solution. We aim to build models that generalize well to unseen domains, whereas the prior works primarily focus on adapting to unseen classes.

Domain Generalized Semantic Segmentation (DGSS). DGSS aims to develop a robust segmentation model that can generalize robustly across various unseen domains. Prior works have concentrated on learning domain-invariant representations through approaches such as normalization and whitening [7, 42–44], and domain randomization [23, 24, 59, 63, 64]. Normalization and whitening remove domain-specific features to focus on learning domain-invariant features. For instance, RobustNet [7] selectively whitens features sensitive to photometric changes. Domain randomization seeks to diversify source domain images by augmenting them into various domain styles. DRPC [59] ensures consistency among multiple stylized images derived from a single content image. TLDR [23]

explores domain randomization while focusing on learning textures. The incorporation of vision transformers [10, 50] has further enhanced DGSS by utilizing the robustness of attention mechanisms. However, these methods largely correspond to visual pattern recognition and have limited ability in understanding high-level semantic concepts inherent to each class.

Recent studies [12, 20, 54] have attempted to utilize VLMs in DGSS. Rein [54] introduces an efficient fine-tuning method that preserves the generalization capability of large-scale vision models, including CLIP [46] and EVA02-CLIP [13, 14]. FAMix [12] employs language as the source of style augmentation, along with a minimal fine-tuning method for the vision backbone of VLMs. VLTseg [20] fine-tunes the vision encoder of VLMs, aligning dense visual features with text embeddings via auxiliary loss. Despite these advancements, the existing approaches either do not utilize language information [54] or use it primarily as an auxiliary tool to support training pipelines [12, 20]. In contrast, this paper introduces a direct language-driven approach to fully harness domain-invariant semantic information embedded in the textual features of VLMs.

Object query design. Recent studies [4, 5, 29, 58, 60] have explored query-based frameworks that utilize a transformer decoder [53] for segmentation tasks, inspired by DETR [2]. In these frameworks, an object query serves to group pixels belonging to the same semantic region (*e.g.*, object or class) by representing the region as a latent vector. Given the critical role of object queries, existing studies have focused on their design and optimization strategies [1, 4, 22, 26, 32, 33, 61, 62]. Mask2former [4] guides query optimization via masked cross-attention for restricting the query to focus on predicted segments. ECENet [33] generates object queries from predicted masks to ensure that the queries represent explicit semantic information. Furthermore, several works [26, 61, 62] have introduced additional supervision into object queries to enhance training stability, while others [1, 22] have designed conditional object queries for cross-modal tasks.

While these studies have demonstrated that purpose-specific object queries contribute to performance, convergence, and functionality in dense prediction tasks [27], research on object queries for DGSS remains unexplored. Our work aims to address DGSS by designing domain-invariant object queries and developing a decoder framework to improve the adaptability of these object queries.

3 Textual Object Query

Our observation is that utilizing text embeddings from VLMs as object queries within a query-based segmentation framework enables domain-invariant recognition. This recognition leads to the effective grouping of dense visual features across diverse domains.

Recent studies [19, 36] have suggested that the text embedding of a class captures core semantic concepts to represent the class across different visual domains, *i.e.*, domain invariant semantics. This capability stems from web-scale contrastive learning [46], which aligns the text embedding for a class of interest with image features of the corresponding class from a wide variety of domains.




Fig. 2: Effectiveness of textual object query. (a) In all DGSS benchmarks, textual object queries (\mathbf{q}_{text}) outperforms randomly initialized object queries (\mathbf{q}_{rand}). (b) We visualize the mask predictions corresponding to a class (*i.e.*, ‘bicycle’), derived from \mathbf{q}_{text} on the left and \mathbf{q}_{rand} on the right, respectively. \mathbf{q}_{rand} yields a degraded result, whereas \mathbf{q}_{text} produces a robust one on a unseen domain.

Given that the text embeddings have the potential to align with dense visual features [38, 47, 65], one can leverage the textual information for domain-generalized dense predictions. We find that the text embedding of a class name is well-aligned with the visual features of the class region, even under extreme domain shifts (see Appendix A).

To leverage domain-invariant semantics in textual features from VLMs, we propose utilizing *textual object queries*. Generally, object queries in transformer-based segmentation frameworks are conceptualized as mask embedding vectors, representing regions likely to be an ‘object’ or a ‘class’ [4]. In semantic segmentation, the queries are optimized to represent the semantic information for classes of interest. Our key insight is that designing object queries with generalized semantic information for classes of interest leads to domain-invariant recognition. We implement textual object queries using the text embeddings of targeted classes from the text encoder of VLMs.

We conduct a motivating experiment to demonstrate the capability of textual object queries to generalize to unseen domains. We design a simple architecture comprising an encoder and textual object queries (\mathbf{q}_{text}), and compare it with an architecture that employs conventional, randomly initialized object queries (\mathbf{q}_{rand}). The details of this experiment are described in Appendix B. In Fig. 2a, \mathbf{q}_{text} outperforms \mathbf{q}_{rand} in all unseen target domains. Fig. 2b clearly supports this observation: \mathbf{q}_{rand} yields a degraded mask prediction for a ‘bicycle,’ whereas \mathbf{q}_{text} produces a robust result on a unseen domain. **query driven mask transformer (tqdm)** to leverage the power of textual object queries.

4 Method

In this section, we propose **textual query-driven mask transformer (tqdm)**, a segmentation decoder that comprehends domain-invariant semantic knowledge by leveraging *textual object queries* as pixel grouping basis (Sec. 4.1). Furthermore, we discuss three regularization losses that aim to maintain robust vision-language alignment, thereby enhancing the efficacy of **tqdm** (Sec. 4.2).

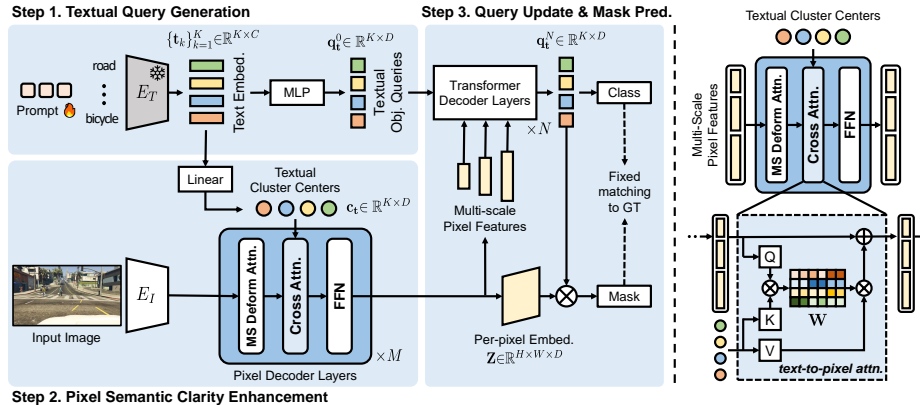


Fig. 3: Overall pipeline of tqdm . (Step 1) We generate initial textual object queries \mathbf{q}_t^0 from the K class text embeddings $\{\mathbf{t}_k\}_{k=1}^K$. (Step 2) To improve the segmentation capabilities of these queries, we incorporate text-to-pixel attention within the pixel decoder. This process enhances the semantic clarity of pixel features, while reconstructing high-resolution per-pixel embeddings \mathbf{Z} . (Step 3) The transformer decoder refines these queries for the final prediction. Each prediction output is then assigned to its corresponding ground truth (GT) through fixed matching, ensuring that each query consistently represents the semantic information of one class.

Preliminary. We employ the image encoder E_I and text encoder E_T , both of which are initialized with a pre-trained Vision-Language Model (VLM) (e.g., CLIP [46]). We fully fine-tune E_I to learn enhanced dense visual representations, whereas E_T is kept frozen to preserve robust textual representations. E_I extracts multi-scale pixel features from an image and feeds them into our tqdm decoder. Additionally, E_I outputs visual embeddings \mathbf{x} , which are projected in a joint vision-language space.

4.1 Textual Query-Driven Mask Transformer

Our proposed framework, tqdm , aims to leverage textual object queries for DGSS in three key steps. Initially, *textual query generation* focuses on generating textual object queries that maximally encode domain-invariant semantic knowledge. Subsequently, *pixel semantic clarity enhancement* aims to improve the segmentation capability of textual object queries by incorporating text-to-pixel attention within a pixel decoder. Lastly, following the practices of mask transformer [4, 5], a transformer decoder refines the object queries for the final mask prediction. The overall pipeline of tqdm is demonstrated in Fig. 3.

Textual query generation. To generate textual object queries for DGSS, we prioritize two key aspects: the queries should (1) preserve domain-invariant semantic information for robust prediction and (2) adapt to the segmentation task to ensure promising performance. To meet the first requirement, we maintain E_T frozen to preserve its original language representations. For the second requirement, we employ learnable prompts [66] to adapt textual features from E_T .

Specifically, E_T generates a text embedding $\mathbf{t}_k \in \mathbb{R}^C$ for each class label name embedding $\{\text{class}_k\}$ with a learnable prompt \mathbf{p} :

$$\mathbf{t}_k = E_T([\mathbf{p}, \{\text{class}_k\}]), \quad (1)$$

where $1 \leq k \leq K$ for total K classes, and C denotes the channel dimension. Finally, we obtain initial textual object queries $\mathbf{q}_t^0 \in \mathbb{R}^{K \times D}$ from the text embeddings $\mathbf{t} = \{\mathbf{t}_k\}_{k=1}^K \in \mathbb{R}^{K \times C}$ through multi-layer perceptron (MLP). Note that D is the dimension of the query vectors in tqdm .

Pixel semantic clarity enhancement. To improve the segmentation capabilities of textual object queries, we incorporate a *text-to-pixel attention* mechanism that enhances the semantic clarity of each pixel feature (refer to ‘‘Cross Attn.’’ in Fig. 3). This mechanism ensures that pixel features are clearly represented in terms of domain-invariant semantics, allowing them to be effectively grouped by textual object queries.

Initially, we derive textual cluster centers $\mathbf{c}_t \in \mathbb{R}^{K \times D}$ by compressing the channel dimension of the text embeddings $\mathbf{t} \in \mathbb{R}^{K \times C}$ with a linear layer. Then, within a pixel decoder layer, a text-to-pixel attention block utilizes multi-scale pixel features $\mathbf{z} \in \mathbb{R}^{L \times D}$ as query tokens \mathbf{Q}_z with a linear projection. Here, L denotes the length of pixel features. The textual clustering centers \mathbf{c}_t are projected into key tokens \mathbf{K}_t and value tokens \mathbf{V}_t . The attention weights $\mathbf{W} \in \mathbb{R}^{L \times K}$ and the enhanced pixel features are calculated as follows:

$$\mathbf{W} = \text{softmax}(\mathbf{Q}_z \mathbf{K}_t^\top), \quad (2)$$

$$\mathbf{z} \leftarrow \mathbf{z} + \mathbf{W} \mathbf{V}_t. \quad (3)$$

We consider this text-to-pixel attention mechanism as a method for updating the pixel features toward K textual clustering centers. The attention weight \mathbf{W} calculates similarity scores between the L pixel features and the K textual clustering centers, where \mathbf{K}_t serves as the clustering centers. Then, in Eq. (3), we refine the pixel features \mathbf{z} with $\mathbf{W} \mathbf{V}_t$, aiming to align more closely with the K textual clustering centers. This approach promotes the grouping of regions belonging to the same classes, thereby improving their semantic clarity.

Query update and mask prediction. The final step involves updating textual object queries and predicting segmentation masks. The transformer decoder, including masked attention [4] with N layers, progressively refines the initial textual object queries \mathbf{q}_t^0 into \mathbf{q}_t^N by integrating pixel features from the pixel decoder. The refined textual object queries $\mathbf{q}_t^N \in \mathbb{R}^{K \times D}$ then predict K masks via dot product with per-pixel embeddings $\mathbf{Z} \in \mathbb{R}^{H \times W \times D}$ followed by sigmoid activation, where H and W are the spatial resolutions. These queries are then classified by a linear classifier with softmax activation to produce a set of class probabilities. We optimize tqdm using the segmentation loss \mathcal{L}_{seg} , following [4]:

$$\mathcal{L}_{\text{seg}} = \lambda_{\text{bce}} \mathcal{L}_{\text{bce}} + \lambda_{\text{dice}} \mathcal{L}_{\text{dice}} + \lambda_{\text{cls}} \mathcal{L}_{\text{cls}}, \quad (4)$$

where the binary cross-entropy loss \mathcal{L}_{bce} and the dice loss [37] $\mathcal{L}_{\text{dice}}$ optimize the predicted masks, and the categorical cross-entropy loss \mathcal{L}_{cls} optimizes the class




Fig. 4: Three regularization losses to enhance the efficacy of tqdm. (a) Language regularization prevents the learnable prompts from distorting the semantic meaning of text embeddings. (b) Vision-language regularization aims to align visual and textual features at the pixel-level. (c) Vision regularization maintains the ability of the vision encoder to align with textual information at the image-level.

prediction of queries. The loss weights λ_{bce} , λ_{dice} and λ_{cls} are set to the same values as those in [4]. To assign each query to a specific class, we adopt fixed matching instead of bipartite matching (see Tab. 2c for ablation). This matching ensures that each query solely represents the semantic information of one class.

4.2 Regularization

Our textual query-driven approach is based on a strong alignment between visual and textual features. To maintain this alignment, we propose three regularization strategies: (1) language regularization that prevents the learnable prompts from distorting the semantic meaning of text embeddings, (2) vision-language regularization that ensures the pixel-level alignment between visual and textual features, and (3) vision regularization that preserves the textual alignment capability of the vision encoder from a pre-trained VLM. The regularization losses are demonstrated in Fig. 4.

Language regularization. We use learnable text prompts [66] to adapt text embeddings for DGSS. During training, the prompts can distort the semantic meaning of text embeddings. To address this issue, we introduce a language regularization loss, which ensures semantic consistency between the text embeddings derived from the learnable prompt \mathbf{p} and those derived from the fixed prompt \mathbf{P}_0 . These embeddings are denoted as $\mathbf{t} \in \mathbb{R}^{K \times C}$ and $\mathbf{T}_0 \in \mathbb{R}^{K \times C}$, respectively.

$$\mathcal{L}_{\text{reg}}^L = \text{Cross-Entropy}(\text{Softmax}(\hat{\mathbf{t}}\hat{\mathbf{T}}_0^\top), \mathbf{I}_K), \quad (5)$$

where $\hat{\mathbf{t}}$ and $\hat{\mathbf{T}}_0$ are the ℓ_2 normalized versions of \mathbf{t} and \mathbf{T}_0 along the channel dimension, respectively, and \mathbf{I}_K is a K -dimensional identity matrix. We use “a clean origami of a [class].” as \mathbf{P}_0 , which is effective for segmentation [31].

Vision-language regularization. For improved segmentation capability of textual object queries, we need to preserve joint vision-language alignment at the pixel-level. We incorporate an auxiliary segmentation loss for a pixel-text score map [20, 47]. This score map is defined by the cosine-similarity between the visual embeddings \mathbf{x} and the text embeddings \mathbf{t} , computed as $\mathbf{S} = \hat{\mathbf{x}}\hat{\mathbf{t}}^\top \in \mathbb{R}^{hw \times K}$, where

$\hat{\mathbf{x}}$ and $\hat{\mathbf{t}}$ are the ℓ_2 normalized versions of \mathbf{x} and \mathbf{t} along the channel dimension, respectively. The score map \mathbf{S} is optimized with a per-pixel cross-entropy loss:

$$\mathcal{L}_{\text{reg}}^{\text{VL}} = \text{Cross-Entropy}(\text{Softmax}(\mathbf{S}/\tau), \mathbf{y}), \quad (6)$$

where τ is a temperature coefficient [17], and \mathbf{y} denotes the ground-truth labels.

Vision regularization. In vision transformers [11], a [class] token captures the global representation of an image [11]. CLIP [46] aligns this [class] token with text embedding of a corresponding caption. Therefore, we consider the token as having the preminent capacity for textual alignment within visual features. To this end, we propose a vision regularization loss $\mathcal{L}_{\text{reg}}^{\text{V}}$ to ensure that the visual backbone preserves its textual alignment at the image-level while learning dense pixel features. Specifically, $\mathcal{L}_{\text{reg}}^{\text{V}}$ enforces the consistency between the [class] token of the training model \mathbf{x}^{CLS} and that of the initial visual backbone $\mathbf{x}_0^{\text{CLS}}$ from a pre-trained VLM:

$$\mathcal{L}_{\text{reg}}^{\text{V}} = \|\mathbf{x}^{\text{CLS}} - \mathbf{x}_0^{\text{CLS}}\|_2. \quad (7)$$

Full objective. The full training objective consists of the segmentation loss \mathcal{L}_{seg} and the regularization loss $\mathcal{L}_{\text{reg}} = \mathcal{L}_{\text{reg}}^{\text{L}} + \mathcal{L}_{\text{reg}}^{\text{VL}} + \mathcal{L}_{\text{reg}}^{\text{V}}$:




$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{seg}} + \mathcal{L}_{\text{reg}}. \quad (8)$$








5 Experiments

5.1 Implementation Details

Datasets. We evaluate the performance of `tqdm` under both synthetic-to-real and real-to-real settings. As synthetic datasets, GTA5 [48] provides 24,966 images at a resolution of 1914×1052 , split into 12,403 for training, 6,382 for validation, and 6,181 for testing. SYNTHIA [49] comprises 6,580 images for training and 2,820 for validation, each at a resolution of 1280×760 . As real-world datasets, Cityscapes [8] includes 2,975 images for training and 500 images for validation, with images at 2048×1024 resolution. BDD100K [57] consists of 7,000 training and 1,000 validation images, each at 1280×720 resolution. Mapillary [39] offers 18,000 training images and 2,000 validation images, with resolutions varying across the dataset. For simplicity, we abbreviate GTA5, SYNTHIA, Cityscapes, BDD100K, and Mapillary as G, S, C, B, and M, respectively.

Network architecture. We employ vision transformer-based backbones, initialized with either CLIP [46] or EVA02-CLIP [51]. The CLIP model incorporates a Vision Transformer-base (ViT-B) backbone [11] with a patch size of 16, while the EVA02-CLIP model utilizes the EVA02-large (EVA02-L) backbone with a patch size of 14. For the pixel decoder, we adopt a multi-scale deformable attention transformer [69], which includes $M=6$ layers, and integrate our text-to-pixel attention layer within it. Regarding the transformer decoder, we follow the default settings outlined in [4], which consist of $N=9$ layers with masked attention.

Table 1: Comparison of mIoU (%) higher is better) for synthetic-to-real setting ($G \rightarrow \{C, B, M\}$) and real-to-real setting ($C \rightarrow \{B, M\}$). , , and  denote initialization with CLIP [46], EVA02-CLIP [51], and DINOv2 [41] pre-training, respectively. The best and second-best results are **highlighted** and underlined, respectively. Our method is marked in blue. The results denoted with † are both trained and tested with an input resolution of 1024×1024 .

Method	Backbone	<i>synthetic-to-real</i>				<i>real-to-real</i>		
		G→C	G→B	G→M	Avg.	C→B	C→M	Avg.
SAN-SAW [44]	RN101	45.33	41.18	40.77	42.43	54.73	61.27	58.00
WildNet [24]	RN101	45.79	41.73	47.08	44.87	47.01	50.94	48.98
SHADE [63]	RN101	46.66	43.66	45.50	45.27	50.95	60.67	55.81
TLDR [23]	RN101	47.58	44.88	48.80	47.09	-	-	-
FAMix [12] 	RN101	49.47	46.40	51.97	49.28	-	-	-
SHADE [64]	MiT-B5	53.27	<u>48.19</u>	54.99	52.15	-	-	-
IBAFFormer [50]	MiT-B5	<u>56.34</u>	49.76	<u>58.26</u>	<u>54.79</u>	-	-	-
VLTSeg [20] 	ViT-B	47.50	45.70	54.30	49.17	-	-	-
tqdm (ours) 	ViT-B	57.50	47.66	59.76	54.97	50.54	65.74	58.14
HGFormer [10]	Swin-L	-	-	-	-	61.50	72.10	66.80
VLTSeg [20] 	EVA02-L	65.60	58.40	<u>66.50</u>	63.50	64.40 [†]	76.40[†]	<u>70.40[†]</u>
Rein [54] 	EVA02-L	65.30	60.50	64.90	63.60	64.10	69.50	66.80
Rein [54] 	ViT-L	<u>66.40</u>	<u>60.40</u>	66.10	<u>64.30</u>	65.00	72.30	68.65
tqdm (ours) 	EVA02-L	68.88	59.18	70.10	66.05	<u>64.72</u>	<u>76.15</u>	70.44

The number of textual object queries is set to 19 to align with the number of classes in the Cityscapes dataset [8]. Additionally, the length of the learnable prompt \mathbf{p} is set to 8.

Training. We use the same training configuration for both the CLIP and EVA02-CLIP models. All experiments are conducted using a crop size of 512×512 , a batch size of 16, and 20k training iterations. Following [4, 20, 54], we adopt an AdamW [34] optimizer. We set the learning rate at 1×10^{-5} for the synthetic-to-real setting and 1×10^{-4} for the real-to-real setting, with the backbone learning rate reduced by a factor of 0.1. Linear warm-up [15] is applied over $t_{\text{warm}}=1.5k$ iterations, followed by a linear decay. We apply standard augmentations for segmentation tasks, including random scaling, random cropping, random flipping, and color jittering. Additionally, we adopt rare class sampling, following [18].

5.2 Comparison with Previous Methods

We compare our **tqdm** with existing DGSS methods. We conduct experiments in two settings: synthetic-to-real ($G \rightarrow \{C, B, M\}$) and real-to-real ($C \rightarrow \{B, M\}$), both for the CLIP and EVA02-CLIP models. Tab. 1 shows that our **tqdm** generally outperforms the existing methods and achieves state-of-the-art results in both the synthetic-to-real and real-to-real settings. In particular, our approach with the EVA02-CLIP model improves the G→C benchmark by 2.48 mIoU. More synthetic-to-real setting (*i.e.*, $S \rightarrow \{C, B, M\}$) results are shown in Appendix C.

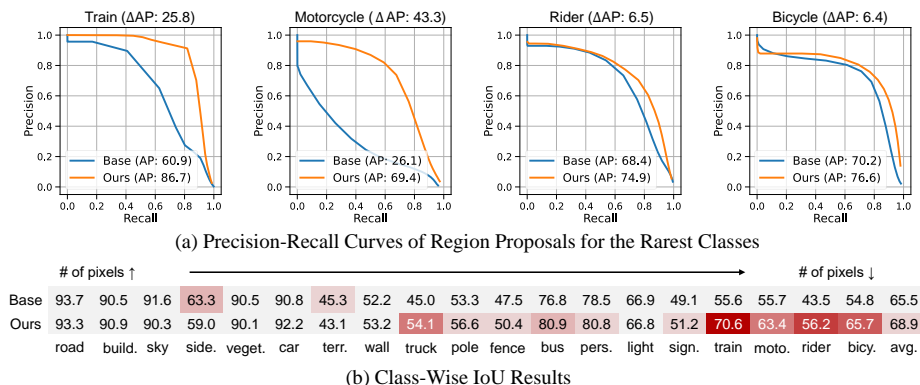


Fig. 5: Precision-recall curves of region proposals for the rarest classes and class-wise IoU results. (a) For the rarest classes (*i.e.*, ‘train,’ ‘motorcycle,’ ‘rider,’ and ‘bicycle’), our \mathbf{tqdm} with textual object query produces more robust region proposals than the baseline with randomly initialized object query. (b) This enhanced robustness leads to the superior DGSS performances of our \mathbf{tqdm} for these classes. The red colors visualize the differences in IoU between the baseline and \mathbf{tqdm} .

5.3 In-Depth Analysis

We design experiments to investigate the factors contributing to the performance improvements achieved by our \mathbf{tqdm} . Our analysis focuses on two key aspects: (1) robustness of object query representations for classes of interest, and (2) semantic coherence² of pixel features across domains. We compare our \mathbf{tqdm} model with the baseline Mask2Former [4] model, both initialized with EVA02-CLIP, and trained on GTA5. The baseline adopts randomly initialized object queries and lacks text-to-pixel attention blocks in its pixel decoder. For a fair comparison, the baseline employs K object queries with fixed matching, the same as \mathbf{tqdm} .

Robustness of object query representations. One notable aspect of \mathbf{tqdm} is its inherent robustness in object query representations, as we use text embeddings from VLMs as a basis for these queries. Given that the role of the initial object query (\mathbf{q}_t^0) involves localizing region proposals [4] and aggregating pixel information within these proposals to obtain the final object query (\mathbf{q}_t^N), developing a robust initial object query is crucial for overcoming domain shifts. Thus, we investigate whether the queries derived from text embeddings produce more robust region proposals compared to the randomly initialized queries.

To quantify this robustness, we demonstrate the precision-recall curves of region proposal predictions on $G \rightarrow C$. As shown in Fig. 5a, for the classes which are rare in the source dataset (*e.g.*, ‘train,’ and ‘motorcycle’) [18], \mathbf{tqdm} outperforms the baseline in Average-Precision (AP). While the baseline with randomly initialized queries is prone to overfitting for rare classes [28], our \mathbf{tqdm} is effective for these classes by encoding domain-invariant representations through the

²Semantic coherence is a property of vision models in which semantically similar regions in images exhibit similar pixel representations [3, 38].

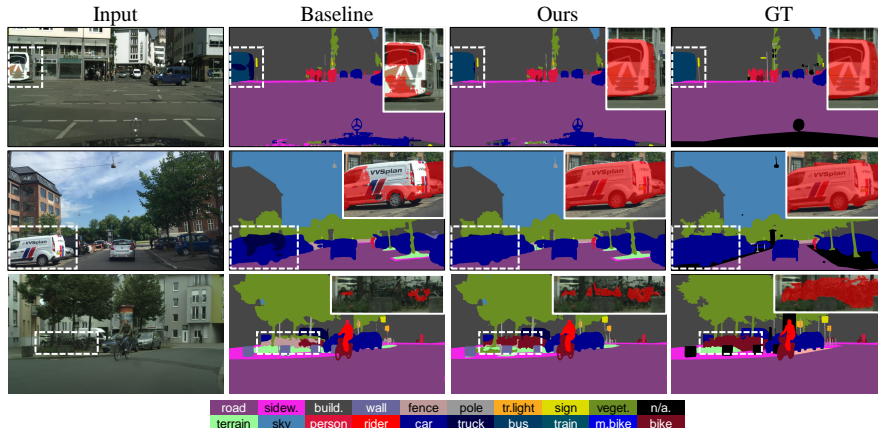


Fig. 6: Qualitative results and region proposals by object queries. Our tqdm provides better prediction results across unseen domains (*i.e.*, C and M) compared to the baseline. We highlight the region proposals generated by object queries with solid lines and the corresponding prediction results with dashed lines. In contrast to the randomly initialized queries of the baseline, our textual object queries lead to more robust region proposals for these classes, resulting in improved predictions.

utilization of language information. Indeed, the class-wise IoU results in Fig. 5b demonstrate notable performance gains across multiple classes, with more pronounced gains in rarer ones. These results affirm that the robust region proposals of tqdm , derived from textual object queries, contribute to enhanced final predictions. Further experimental details and results are provided in Appendix D.

Our qualitative results for unseen domains (*i.e.*, C and M), as shown in Fig. 6, also support the idea that robust region proposals lead to enhanced prediction results. The tqdm model provides better prediction results with high-quality region proposals, while the baseline produces degraded region proposals, resulting in inferior predictions (refer to white boxes). We provide more qualitative comparisons with other DGSS methods in Appendix E.

Semantic coherence on pixel features. The other notable aspect of tqdm is its semantic coherence [3, 38] for pixel features across source and target domains. This property is achieved by incorporating text-to-pixel attention within the pixel decoder to enhance the semantic clarity of pixel features. To visualize this property, we start by selecting a pixel embedding from a source domain image (indicated by \times in Fig. 7). We then measure the cosine similarities across all pixel embeddings for the source image itself and for unseen domain images, both before ($m=0$) and after ($m=M$) processing by the pixel decoder. Here, m is the index of pixel decoder layers. Finally, we plot these similarity measurements as heatmaps to visualize the results.

In the source domain image, both the tqdm and baseline models exhibit comparable levels of semantic coherence before and after processing by the pixel decoder (see Fig. 7a). In contrast, for the unseen domain images, the tqdm model demonstrates significantly better semantic coherence for the class of interest (*i.e.*,

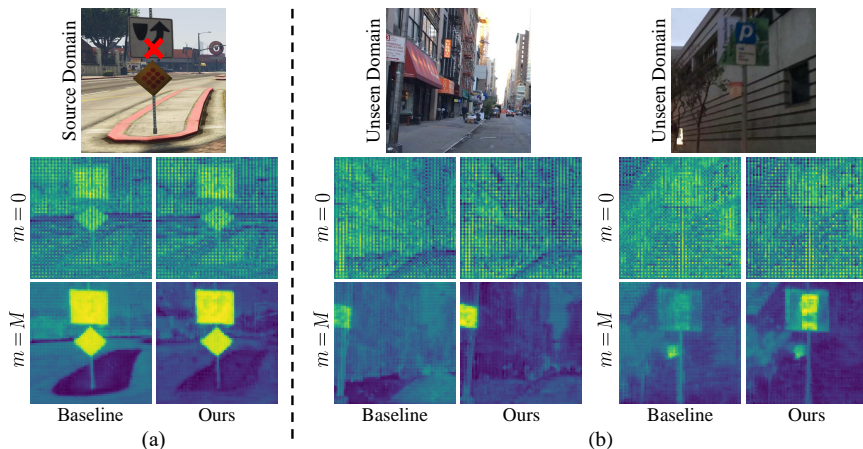


Fig. 7: Visualization of semantic coherence on pixel features. We compare the semantic coherence of pixel features across source and target domains between `tqdm` and the baseline. We select a pixel embedding from a source domain image (indicated by \times). Then, we measure the cosine similarities across all pixel embeddings (a) for the source image itself and (b) for unseen domain images, both before ($m=0$) and after ($m=M$) processing by the pixel decoder. For the unseen domain images, the `tqdm` model demonstrates significantly better semantic coherence for the class of interest (*i.e.*, ‘traffic sign’) compared to the baseline after processing by the pixel decoder.

‘traffic sign’) after processing by the pixel decoder, compared to the baseline (see Fig. 7b). These results imply that text-to-pixel attention enhances the semantic clarity of pixel features, consequently leading to the refined pixel features being more effectively grouped by textual object queries, even in unseen domains.

5.4 Ablation Studies

In our ablation experiments, we train the EVA-CLIP model on GTA5 and evaluate it on Cityscapes, BDD100K, and Mapillary.

Key components. We investigate how the key components contribute to the overall performance of our method. We first verify the effectiveness of textual object queries (\mathbf{q}_t). In Tab. 2a, the model with \mathbf{q}_t (row 2) shows better results compared to the model with randomly initialized object query (row 1), with an average gain of 0.69 mIoU. Then, we evaluate the contribution of the text-to-pixel attention block (\mathbf{A}_{t2p}), which complements the segmentation capacity of the queries. \mathbf{A}_{t2p} further improves the performance by 1.12 mIoU on average, enhancing the semantic clarity of per-pixel embeddings (row 2 and 3).

Regularization losses. We analyze how each regularization loss contributes to the overall performance. Tab. 2b presents the results under configurations without the language regularization loss ($\mathcal{L}_{\text{reg}}^L$), vision-language regularization loss ($\mathcal{L}_{\text{reg}}^{\text{VL}}$), and vision regularization loss ($\mathcal{L}_{\text{reg}}^V$), respectively. The performance degrades when any of these three regularization losses is excluded, which implies

Table 2: Ablation experiments. We use the EVA02-L model. The models are trained on GTA5, and evaluated on Cityscapes, BDD100K and Mapillary. The best results are **highlighted**, and the default setting is marked in blue.

(a) Key Components						(b) Regularization Losses								
q_t	A_{t2p}	C	B	M	Avg.	\mathcal{L}_{reg}^L	\mathcal{L}_{reg}^{VL}	\mathcal{L}_{reg}^V	C	B	M	Avg.		
1	-	-	66.24	57.79	68.68	64.24	1	-	✓	✓	69.34	58.01	68.98	65.44
2	✓	-	67.66	58.09	69.03	64.93	2	✓	-	✓	65.51	55.40	68.16	63.02
3	✓	✓	68.88	59.18	70.10	66.05	3	✓	✓	-	69.21	57.44	69.27	65.31
							4	✓	✓	✓	68.88	59.18	70.10	66.05

(c) Matching Assignment Choice					(d) Text Prompt Choice				
Matching	C	B	M	Avg.	Prompt	C	B	M	Avg.
1 bipartite	66.67	56.79	69.11	64.19	1 fixed template	68.48	57.81	69.30	65.20
2 fixed	68.88	59.18	70.10	66.05	2 learnable	68.88	59.18	70.10	66.05

the importance of maintaining robust vision-language alignment. Particularly, \mathcal{L}_{reg}^{VL} significantly contributes to the performance by 3.03 mIoU on average (row 2 and 4). This result underlines that the vision-language alignment at the pixel-level is essential for enhancing the efficacy of the textual query-driven framework.

Matching assignment choice. We adopt fixed matching to ensure that each query represents the semantics of a single class. We compare the fixed matching approach with conventional bipartite matching [4, 5]. The model with fixed matching outperforms one with bipartite matching by 1.86 mIoU on average.

Text prompt choice. We validate the effectiveness of prompt tuning. In Tab. 2d, the learnable prompt tuning yields an average of 0.85 mIoU over one that adopts a fixed template prompt (*i.e.*, “a clean origami of a [class].” [31]).

6 Conclusion

Fine-grained visual features for classes can vary across different domains. This variation presents a challenge for developing models that genuinely understand fundamental semantic knowledge of the classes and generalize well between the domains. To address this challenge in DGSS, we propose utilizing text embeddings from VLMs as object queries within a transformer-based segmentation framework, *i.e.*, textual object queries. Moreover, we introduce a novel framework called **tqdm** to fully harness the power of textual object queries. Our **tqdm** is designed to (1) generate textual object queries that fully capture domain-invariant semantic information and (2) improve their adaptability in dense predictions through enhancing pixel semantic clarity. Additionally, we suggest three regularization losses to preserve the robust vision-language alignment of pre-trained VLMs. Our comprehensive experiments demonstrate the effectiveness of textual object queries in recognizing domain-invariant semantic information in DGSS. Notably, **tqdm** achieves the state-of-the-art performance on multiple DGSS benchmarks, *e.g.*, 68.9 mIoU on GTA5→Cityscapes, outperforming the prior state-of-the-art method by 2.5 mIoU.

Acknowledgements

We sincerely thank Chanyong Lee and Eunjin Koh for their constructive discussions and support. We also appreciate Junyoung Kim, Chaehyeon Lim and Minkyu Song for providing insightful feedback. This work was supported by the Agency for Defense Development (ADD) grant funded by the Korea government (279002001).

References

1. Cai, Z., Kwon, G., Ravichandran, A., Bas, E., Tu, Z., Bhotika, R., Soatto, S.: X-DETR: A versatile architecture for instance-wise vision-language tasks. In: ECCV (2022)
2. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV. Springer (2020)
3. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV (2021)
4. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: CVPR (2022)
5. Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. NeurIPS (2021)
6. Cho, J., Nam, G., Kim, S., Yang, H., Kwak, S.: PromptStyler: Prompt-driven style generation for source-free domain generalization. In: ICCV (2023)
7. Choi, S., Jung, S., Yun, H., Kim, J.T., Kim, S., Choo, J.: RobustNet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In: CVPR (2021)
8. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016)
9. Desai, K., Johnson, J.: VirTex: Learning visual representations from textual annotations. In: CVPR (2021)
10. Ding, J., Xue, N., Xia, G.S., Schiele, B., Dai, D.: HGFormer: Hierarchical grouping transformer for domain generalized semantic segmentation. In: CVPR (2023)
11. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
12. Fahes, M., Vu, T.H., Bursuc, A., Pérez, P., de Charette, R.: A simple recipe for language-guided domain generalized segmentation. arXiv preprint arXiv:2311.17922 (2023)
13. Fang, Y., Sun, Q., Wang, X., Huang, T., Wang, X., Cao, Y.: EVA-02: A visual representation for neon genesis. arXiv preprint arXiv:2303.11331 (2023)
14. Fang, Y., Wang, W., Xie, B., Sun, Q., Wu, L., Wang, X., Huang, T., Wang, X., Cao, Y.: EVA: Exploring the limits of masked visual representation learning at scale. In: CVPR (2023)
15. Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, large minibatch sgd: Training ImageNet in 1 hour. arXiv preprint arXiv:1706.02677 (2017)

16. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: CVPR (2022)
17. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020)
18. Hoyer, L., Dai, D., Van Gool, L.: DAFormer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In: CVPR (2022)
19. Huang, Z., Zhou, A., Ling, Z., Cai, M., Wang, H., Lee, Y.J.: A sentence speaks a thousand images: Domain generalization through distilling CLIP with language guidance. In: ICCV (2023)
20. Hümmer, C., Schwonberg, M., Zhong, L., Cao, H., Knoll, A., Gottschalk, H.: VLT-Seg: Simple transfer of CLIP-based vision-language representations for domain generalized semantic segmentation. arXiv preprint arXiv:2312.02021 (2023)
21. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: ICML (2021)
22. Kamath, A., Singh, M., LeCun, Y., Synnaeve, G., Misra, I., Carion, N.: MDETR – Modulated detection for end-to-end multi-modal understanding. In: ICCV (2021)
23. Kim, S., Kim, D.h., Kim, H.: Texture learning domain randomization for domain generalized segmentation. ICCV (2023)
24. Lee, S., Seong, H., Lee, S., Kim, E.: WildNet: Learning domain generalized semantic segmentation from the wild. In: CVPR (2022)
25. Li, B., Weinberger, K.Q., Belongie, S., Koltun, V., Ranftl, R.: Language-driven semantic segmentation. In: ICLR (2022)
26. Li, F., Zhang, H., Liu, S., Guo, J., Ni, L.M., Zhang, L.: DN-DETR: Accelerate detr training by introducing query denoising. In: CVPR (2022)
27. Li, X., Ding, H., Zhang, W., Yuan, H., Pang, J., Cheng, G., Chen, K., Liu, Z., Loy, C.C.: Transformer-based visual segmentation: A survey. arXiv preprint arXiv:2304.09854 (2023)
28. Li, Z., Kamnitsas, K., Glocker, B.: Analyzing overfitting under class imbalance in neural networks for image segmentation. IEEE transactions on medical imaging **40**(3), 1065–1077 (2020)
29. Li, Z., Wang, W., Xie, E., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P., Lu, T.: Panoptic segformer: Delving deeper into panoptic segmentation with transformers. In: CVPR (2022)
30. Liang, F., Wu, B., Dai, X., Li, K., Zhao, Y., Zhang, H., Zhang, P., Vajda, P., Marculescu, D.: Open-vocabulary semantic segmentation with mask-adapted CLIP. In: CVPR (2023)
31. Lin, Y., Chen, M., Wang, W., Wu, B., Li, K., Lin, B., Liu, H., He, X.: CLIP is also an efficient segmenter: A text-driven approach for weakly supervised semantic segmentation. In: CVPR (2023)
32. Liu, S., Li, F., Zhang, H., Yang, X., Qi, X., Su, H., Zhu, J., Zhang, L.: DAB-DETR: Dynamic anchor boxes are better queries for detr. ICLR (2022)
33. Liu, Y., Liu, C., Han, K., Tang, Q., Qin, Z.: Boosting semantic segmentation from the perspective of explicit class embeddings. In: ICCV (2023)
34. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. ICLR (2019)
35. Lu, J., Batra, D., Parikh, D., Lee, S.: ViLBERT: Pretraining task-agnostic vision-linguistic representations for vision-and-language tasks. NeurIPS (2019)
36. Mangla, P., Chandhok, S., Aggarwal, M., Balasubramanian, V.N., Krishnamurthy, B.: INDIGO: intrinsic multimodality for domain generalization. arXiv preprint arXiv:2206.05912 (2022)

37. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 3DV (2016)
38. Mukhoti, J., Lin, T.Y., Poursaeed, O., Wang, R., Shah, A., Torr, P.H., Lim, S.N.: Open vocabulary semantic segmentation with patch aligned contrastive learning. In: CVPR (2023)
39. Neuhold, G., Ollmann, T., Rota Bulò, S., Kotschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: ICCV (2017)
40. Nguyen, T., Ilharco, G., Wortsman, M., Oh, S., Schmidt, L.: Quality not quantity: On the interaction between dataset design and robustness of CLIP. NeurIPS (2022)
41. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: DINOv2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)
42. Pan, X., Luo, P., Shi, J., Tang, X.: Two at once: Enhancing learning and generalization capacities via ibn-net. In: ECCV (2018)
43. Pan, X., Zhan, X., Shi, J., Tang, X., Luo, P.: Switchable whitening for deep representation learning. In: ICCV (2019)
44. Peng, D., Lei, Y., Hayat, M., Guo, Y., Li, W.: Semantic-aware domain generalized segmentation. In: CVPR (2022)
45. Pham, H., Dai, Z., Ghiasi, G., Kawaguchi, K., Liu, H., Yu, A.W., Yu, J., Chen, Y.T., Luong, M.T., Wu, Y., et al.: Combined scaling for zero-shot transfer learning. Neurocomputing (2023)
46. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
47. Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., Zhou, J., Lu, J.: DenseCLIP: Language-guided dense prediction with context-aware prompting. In: CVPR (2022)
48. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: ECCV (2016)
49. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: CVPR (2016)
50. Sun, Q., Chen, H., Zheng, M., Wu, Z., Felsberg, M., Tang, Y.: IBAFormer: Intra-batch Attention Transformer for Domain Generalized Semantic Segmentation. arXiv preprint arXiv:2309.06282 (2023)
51. Sun, Q., Fang, Y., Wu, L., Wang, X., Cao, Y.: Eva-clip: Improved training techniques for clip at scale. arXiv preprint arXiv:2303.15389 (2023)
52. Tan, H., Bansal, M.: LXMERT: Learning cross-modality encoder representations from transformers. arXiv preprint arXiv:1908.07490 (2019)
53. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. NeurIPS (2017)
54. Wei, Z., Chen, L., Jin, Y., Ma, X., Liu, T., Lin, P., Wang, B., Chen, H., Zheng, J.: Stronger, fewer, & superior: Harnessing vision foundation models for domain generalized semantic segmentation. arXiv preprint arXiv:2312.04265 (2023)
55. Wortsman, M., Ilharco, G., Kim, J.W., Li, M., Kornblith, S., Roelofs, R., Lopes, R.G., Hajishirzi, H., Farhadi, A., Namkoong, H., et al.: Robust fine-tuning of zero-shot models. In: CVPR (2022)
56. Xu, M., Zhang, Z., Wei, F., Lin, Y., Cao, Y., Hu, H., Bai, X.: A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In: ECCV (2022)

57. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: BDD100k: A diverse driving dataset for heterogeneous multitask learning. In: CVPR (2020)
58. Yu, Q., Wang, H., Qiao, S., Collins, M., Zhu, Y., Adam, H., Yuille, A., Chen, L.C.: kMaX-DeepLab: k-means mask transformer. In: ECCV (2022)
59. Yue, X., Zhang, Y., Zhao, S., Sangiovanni-Vincentelli, A., Keutzer, K., Gong, B.: Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In: ICCV (2019)
60. Zhang, B., Tian, Z., Tang, Q., Chu, X., Wei, X., Shen, C., et al.: Segvit: Semantic segmentation with plain vision transformers. NeurIPS (2022)
61. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: DINO: DETR with improved denoising anchor boxes for end-to-end object detection. arXiv preprint arXiv:2203.03605 (2022)
62. Zhang, H., Li, F., Xu, H., Huang, S., Liu, S., Ni, L.M., Zhang, L.: MP-Former: Mask-piloted transformer for image segmentation. In: CVPR (2023)
63. Zhao, Y., Zhong, Z., Zhao, N., Sebe, N., Lee, G.H.: Style-hallucinated dual consistency learning for domain generalized semantic segmentation. In: ECCV (2022)
64. Zhao, Y., Zhong, Z., Zhao, N., Sebe, N., Lee, G.H.: Style-hallucinated dual consistency learning: A unified framework for visual domain generalization. IJCV (2023)
65. Zhou, C., Loy, C.C., Dai, B.: Extract free dense labels from CLIP. In: ECCV (2022)
66. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. IJCV (2022)
67. Zhou, Z., Lei, Y., Zhang, B., Liu, L., Liu, Y.: ZegCLIP: Towards adapting clip for zero-shot semantic segmentation. In: CVPR (2023)
68. Zhu, C., Chen, L.: A survey on open-vocabulary detection and segmentation: Past, present, and future. arXiv preprint arXiv:2307.09220 (2023)
69. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: Deformable transformers for end-to-end object detection. ICLR (2020)