

# Test-Time Modification: Inverse Domain Transformation for Robust Perception

Arpit Jadon<sup>1,\*</sup> Joshua Niemeijer<sup>1,\*</sup> Yuki M. Asano<sup>2</sup>

<sup>1</sup>German Aerospace Center Braunschweig <sup>2</sup>University of Technology Nuremberg

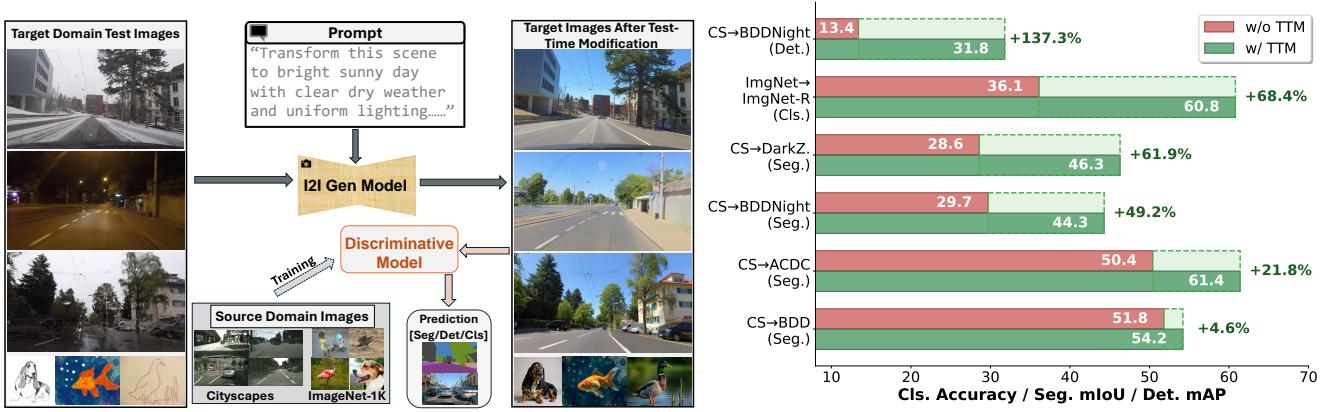


Figure 1. **Left:** Our proposed Test-time view modification via inverse domain transfer using a strong image-to-image (I2I) Foundation Model. **Right:** Subset of reported performance improvements on domain generalization benchmarks.

## Abstract

Generative foundation models contain broad visual knowledge and can produce diverse image variations, making them particularly promising for advancing domain generalization tasks. While they can be used for training data augmentation, synthesizing comprehensive target-domain variations remains slow, expensive, and incomplete. We propose an alternative: using diffusion models at test time to map target images back to the source distribution where the downstream model was trained. This approach requires only a source domain description, preserves the task model, and eliminates large-scale synthetic data generation. We demonstrate consistent improvements across segmentation, detection, and classification tasks under challenging environmental shifts in real-to-real domain generalization scenarios with unknown target distributions. Our analysis spans multiple generative and downstream models, including an ensemble variant for enhanced robustness. The method achieves substantial relative gains: 137% on BDD100K-Night, 68% on ImageNet-R, and 62% on Dark-Zurich.

## 1. Introduction

Despite major advances in vision architectures and training strategies, model performance still relies strongly on the quality and coverage of the training data. Predictions tend to be reliable when test images follow the same distribution as the training set, but even moderate distribution shifts can lead to a clear drop in accuracy. To tackle this, the field of domain generalization aims to train models that remain robust to unseen target domains without assuming access to data or annotations from those domains. Most existing approaches fall into two categories (Fig. 2b,c).

The first one “Unspecific Image Augmentation” modifies the training data in broad, unspecific ways: typically through heavy augmentations such as blur, noise, or color jitter, with the hope that the model will generalize to new domains. The second approach, “Generative Training Image Augmentation” leverages recent diffusion-based generative models to synthesize additional training data. Using text prompts, these methods attempt to approximate potential target domains and generate corresponding images. However, estimating all relevant target domains in advance is difficult, and the diversity of synthesized data remains limited.

In this work, we introduce a third direction. Instead of

\* Equal contribution

expanding the training distribution, we use modern generative models to compute an *inverse transformation at test time*. The goal is to map a target-domain image back toward the source-domain distribution on which the discriminative model was trained. This idea is related to test-time normalization techniques such as adjusting BatchNorm statistics, but diffusion-based transformations allow for far richer and more expressive changes. Because the discriminative model performs best on its original training distribution, this inverse mapping improves prediction quality without modifying the model itself. Importantly, computing this inverse transformation requires only a description of the source domain, which is readily available. In contrast, expanding the training distribution with synthetic data would require anticipating all relevant target domains, which is a far harder problem.

We emphasize that our approach is not intended as a run-time optimization for deployment scenarios, as it requires inference through large generative models that introduce significant computational overhead. Rather, we present this work as a proof of principle that demonstrates how the vast world knowledge encoded in foundation models can be effectively harnessed to improve traditional computer vision tasks. Our method showcases a paradigm where generative models serve as powerful domain translators, unlocking their learned representations to bridge distribution gaps without any task-specific training. This demonstration opens new research directions for understanding and utilizing the complementary strengths of generative and discriminative models, even if practical deployment would require further innovations in efficiency.

We evaluate our inverse domain transformation on domain generalization benchmarks for semantic segmentation, detection, and classification. Our test-time modification (TTM) pipeline improves performance across all settings without finetuning either the discriminative or generative models, offering a simple, task-agnostic approach to domain generalization. Our main contributions are as follows:

- A formalisation of and recipe for Inverse Domain Transformation for view modification of test data
- Quantitative evaluation and analysis of Test-Time Modification (TTM) as a new paradigm that requires no retraining of generative or discriminative models
- State-of-the-art results for various pretrained models on benchmarks such as BDD [39], DarkZurich [27], ACDC [29], ImageNet-R [10]

## 2. Related Work

Domain Generalization is the task of making a discriminative model robust to data from distributions unseen during training. Given a discriminative model trained on a so-called source domain, it is tested on data from unseen tar-

get domains that exhibit substantial domain shifts with respect to, for instance, lighting, weather, or environmental conditions. We present existing approaches to domain generalization and works related to our inverse transformation approach.

**Constraining the distributions.** Early methods constrain input distributions via normalization in latent space, using batch-norm [31] or instance normalization [21] with test-time parameters from target images. Subsequent works introduce whitening [22] to decorrelate channel representations and reduce style information, with various improvements for guided whitening [4, 24, 38]. Our approach instead leverages generative models’ world knowledge to map target distributions back to source, providing more robust transformations that handle aleatoric uncertainty.

**Data augmentation and randomization.** Extending source distributions to include target domain information has been widely explored, from basic augmentations (Gaussian blur, color jitter) to domain randomization [42] using random style transformations. Several works [12, 23] leverage auxiliary domains (paintings, web images) for style randomization, while Huang et al. [11] perform randomization in frequency space. Zhong et al. [45] proposed an adversarial style augmentation and Zhao et al. [44] a style hallucination module for extending the source domain.

**Prior knowledge for source domain extension.** Generative models like Stable Diffusion [26] enable synthetic target domain creation using prior knowledge. Niemeijer et al. [18] generate pseudo-target domains with weather and traffic variations for autonomous driving, then apply unsupervised domain adaptation [30]. Benigmim et al. [1] utilize the prior knowledge contained in LLMs to create meaningful prompts for the generation of relevant synthetic data. ControlNet variants [15, 16, 19, 43] enable supervised training on synthetic data by conditioning on both labels and target domain descriptions. However, limited prior knowledge of prevents comprehensive coverage of all relevant domains that might be encountered at test-time. In contrast, we argue that using these generative models is more efficient when done in reverse: we describe the source domain and transform target images back to the source distribution.

**Inverse transformations.** Prior inverse transformation methods use style transfer via Fourier transforms [32] or GANs [17], but have required *target domain access* to determine style parameters. The approach of Yu et al. [40] is closest to ours. They proposed an inverse transformation to the source domain based on a source domain trained unconditional diffusion model. We instead propose a new paradigm that leverages world knowledge from generative models like Flux1, Kontext [14] and Qwen-Image-Edit-2509 [34] to perform arbitrary transformations using only text descriptions. Beyond style transfer, these models can interpret image content and even reduce aleatoric uncer-

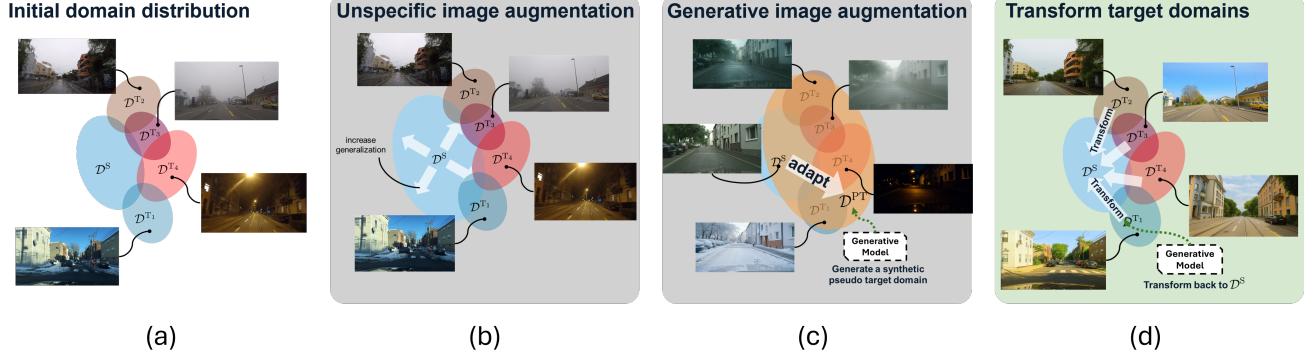


Figure 2. (a) Initial distributions:  $\mathcal{D}^S$  source domain,  $\mathcal{D}^{T_k}$  different target domain distributions. (b) untargeted extension of the source domain by traditional augmentations. (c) use of generative models to create pseudo target domain that compasses all relevant  $\mathcal{D}^{T_k}$ . (d) use of generative models to transform images from  $\mathcal{D}^{T_k}$  back to the distribution of  $\mathcal{D}^S$ .

tainty, as we describe in Sec.3.1.

### 3. Methods

**Preliminaries.** Consider a model  $f_\theta$  (for classification, detection, or segmentation) with parameters  $\theta$ , pretrained on a source domain distribution  $\mathcal{D}^S$ . Our goal is to generalize this model to unseen target domains  $\mathcal{D}^{T_k}$ , where  $k \in \{1, \dots, K\}$  represents  $K$  potential target domains.

We denote images as  $x_n \in \mathbb{R}^{H \times W \times 3}$  for RGB inputs, where  $n \in \{1, \dots, N\}$  indexes a dataset of  $N$  images. The model  $f_\theta$  produces predictions  $y_n$  for input  $x_n$ , with  $y_{n,c}$  representing the posterior probability for class  $c \in \mathcal{C}$ .

Generative models like Stable Diffusion [26, 35] are trained on large-scale web data containing diverse image distributions, enabling them to capture extensive world knowledge. Recent works [13, 18, 36] have exploited this capability to augment training data distributions. In these approaches, a diffusion generative model  $G$  is guided by text prompt  $t_k$  to create synthetic images  $x_n^{PT}$  that approximate target domain characteristics:

$$x_n^{PT} = G(x_n^S, t_k) \quad (1)$$

These methods estimate relevant target domains by encoding prior knowledge about  $\mathcal{D}^{T_k}$  distributions into the text prompt  $t_k$ , producing a pseudo-target domain  $\mathcal{D}^{PT}$ .

However, this approach assumes we possess sufficient prior knowledge to generate distributions for all unseen domains  $\mathcal{D}^{T_k}$ . We argue this assumption is impractical: (1) the number of relevant target domains may be unbounded, and (2) symbolic text descriptions may fail to capture the full complexity of these distributions. Therefore, we propose to invert this paradigm.

**Test-Time Image Modification.** Instead of generating pseudo-target domains during training, we propose using generative models for image editing (e.g., Flux.1 Kon-

text [14] or Qwen-Image-Edit [34]) to transform target domain images back to the source domain at test time:

$$x_n^{PS} = G(x_n^T, t^S) \quad (2)$$

where  $t^S$  describes the source domain distribution.

This approach offers a key advantage: rather than requiring  $K$  prompts  $t_{1\dots K}$  to describe all possible unknown target domains, we need only a *single prompt*  $t^S$  describing the known source domain. This dramatically reduces complexity from describing multiple unknown distributions to characterizing one well-understood distribution.

At test time, we apply the discriminative model  $f_\theta$  to the pseudo-source-domain image:

$$y_n^{PS} = f_\theta(x_n^{PS}) = f_\theta(G(x_n^T, t^S)) \quad (3)$$

Since  $f_\theta$  was trained on the source domain distribution, predictions on  $x_n^{PS}$  should be more accurate than those obtained directly from  $x_n^T$ . Crucially, this approach requires no additional training or fine-tuning of the discriminative model.

One potential challenge is maintaining semantic consistency during the inverse transformation. The generative model  $G$  may inadvertently introduce or remove semantic content—a known limitation affecting even source-to-target transformation methods. To address this, we ensemble predictions from both the original and pseudo-source domain images:

$$y_n^T = 0.5f_\theta(x_n^{PS}) + 0.5f_\theta(x_n^T) \quad (4)$$

This fusion strategy grounds the final prediction in the original image while benefiting from the domain-aligned features of the transformed image. Note that this fusion step is task-specific—we apply it for semantic segmentation but not for object detection or image classification.

**Generating Effective Source Domain Prompts.** To obtain an effective source domain prompt  $t^S$ , we employ the two-stage pipeline illustrated in Figure 3.

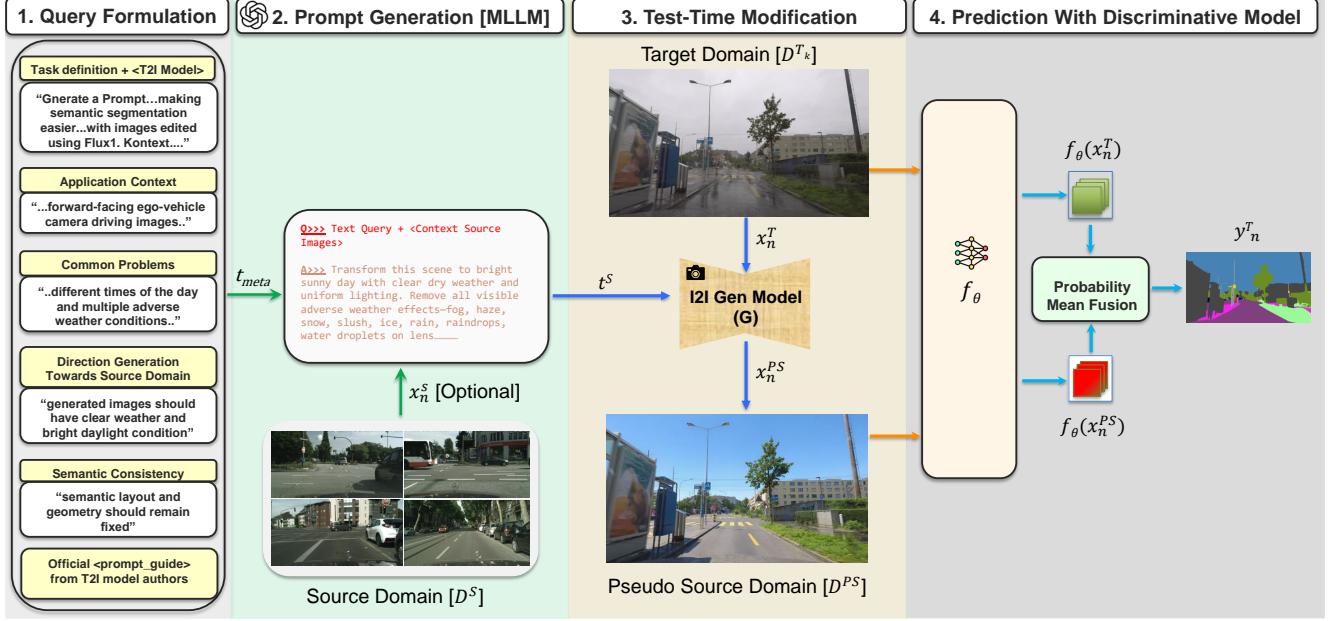


Figure 3. Overview of our test-time modification pipeline via inverse domain transfer. The process begins with formulating a meta-prompt that describes the desired transformation. An MLLM processes this meta-prompt to generate an effective prompt  $t^S$  for the I2I generation model. During test time, the I2I model transforms target domain images to the source domain. Finally, the pre-trained discriminative model processes both images, with task-specific fusion strategies applied as needed.

**Step 1: Meta-Prompt Formulation.** We construct a comprehensive meta-prompt: a human-written text, denoted as  $t_{\text{meta}}$ , designed to generate another prompt. This meta-prompt instructs a multimodal large language model (MLLM), such as GPT-5 [20], to generate a prompt suitable for inverse domain transformation using an Image-to-Image (I2I) generation model. The meta-prompt  $t_{\text{meta}}$  integrates the following components:

- **Task specification:** The downstream task (e.g., semantic segmentation)
- **Model information:** The specific I2I model being used
- **Objective:** The intended goal (e.g., improving segmentation on test images)
- **Domain context:** Application-specific information (e.g., “road driving scenes”)
- **Expected challenges:** Common domain shifts such as weather conditions or lighting variations
- **Transformation requirements:** Instructions to preserve semantic layout while adapting appearance
- **Model guidelines:** Best practices from the I2I model’s documentation

While specifying expected challenges requires minimal prior knowledge, this information could alternatively be inferred at test time using MLLMs. Optionally, we can provide a small subset of  $m$  source domain images  $\{x_1^S, \dots, x_m^S\}$  as visual context, though this is typically unnecessary for well-known datasets.

**Step 2: Prompt Generation.** The MLLM processes the meta-prompt to generate the final source domain prompt:

$$t^S = \text{MLLM}(t_{\text{meta}}, \{x_1^S, \dots, x_m^S\}) \quad (5)$$

The resulting  $t^S$  enables transformation of target domain images into semantically consistent pseudo-source domain images. For example, in our driving scene experiments, we use the following prompt that removes various domain shift artifacts while preserving scene structure:

Transform this scene to bright sunny day with clear dry weather and uniform lighting. Remove all visible adverse weather effects [...] while maintaining original object positions, scale, scene composition, camera angle, and framing exactly as in the input.

We provide complete prompt variations and ablation studies in the appendix.

### 3.1. Intuition: Reducing Aleatoric Uncertainty

Our approach not only addresses epistemic uncertainty (model uncertainty) but also reduces aleatoric uncertainty inherent in the data. For a discriminative model  $f_\theta(x)$ , the predictive variance decomposes as:

$$\text{Var}[y \mid x] = \underbrace{\text{Var}_\theta[f_\theta(x)]}_{\text{epistemic}} + \underbrace{\mathbb{E}_\theta[\text{Var}(y \mid x, \theta)]}_{\text{aleatoric}} \quad (6)$$

While epistemic uncertainty can be reduced through additional training data, aleatoric uncertainty arises from the input itself—due to occlusions, low light, noise, or adverse weather—and cannot be reduced by collecting more training samples:

$$\text{Var}(y | x, \theta) = \text{Var}(y | x, \theta^+) \quad (7)$$

where  $\theta^+$  represents parameters learned from additional training data.

#### Inverse Transformation as Aleatoric Reduction.

Modern generative models, trained on diverse image corpora, possess broad prior knowledge similar to human visual experience. By transforming a corrupted target image  $x^T$  to a cleaner pseudo-source image  $x^{PS}$ , we effectively reduce the input-dependent uncertainty:

$$\text{Var}(y | x^{PS}, \theta) < \text{Var}(y | x^T, \theta) \quad (8)$$

This allows our test-time modification to mitigate effects such as snow cover, raindrops, or low-light conditions (see Figures 4 and 7). Importantly, this reduction in aleatoric uncertainty would not be achievable by synthesizing pseudo-target domains during training, as aleatoric uncertainty is fundamentally tied to the input rather than the model.

## 4. Experiments

**Datasets** Across all downstream tasks, we use a total of six target datasets for evaluation and two datasets for training our source domain models. For semantic segmentation, we use the clear weather Cityscapes [6] as our source dataset. Whereas we evaluate on adverse weather and nighttime datasets - ACDC [29], BDD100K [39], BDD100K-Night [28, 39], and Dark Zurich [27]. For the object detection task, we again use Cityscapes to train our source model and evaluate on a randomly sampled subset of 501 BDD100K nighttime images with object detection annotations. We refer to this validation set as BDD100K-Night-Det. For image classification, we use the ImageNet-1K [7] dataset as our source domain and ImageNet-R [10] dataset as our target set. The latter contains difficult, rendered image versions of a subset of 200 categories. Refer to the appendix for more information on the various datasets.

**Implementation Details** For our semantic segmentation experiments, we utilize both convolutional neural networks (CNNs) and transformer-based models of varying sizes. CNN based models include DeepLabV3+ [2] and OCRNet [41] whereas transformer based models include Segformer [37] and Mask2Former [3]. All these models are pre-trained on the Cityscapes dataset and are directly evaluated on the validation set of the target dataset. For object detection evaluation, we use the Faster-RCNN [25] and Mask-RCNN [9] models. For the image classification task, we

use the ResNet-50 and ResNet-152 models [8]. For image generation, we utilize two different image-to-image (I2I) generation models as part of TTM, including a commercial model, Flux.1 Kontext Max and an open-source model Qwen-Image-Edit-2509 [QIE-2509]. We also do an ablation study on additional I2I generation models. More details on different models are provided in the appendix.

### 4.1. Semantic Segmentation

We evaluate on the domain generalized semantic segmentation task in the autonomous driving context, focusing on weather and daytime shifts, as they cause the most significant domain gap in autonomous driving applications.

Table 1 shows a performance comparison of our Test-Time Modification (TTM) method on the domain generalized semantic segmentation task. As shown in the table, using TTM yields significant improvements in mIoU across various datasets and segmentation models. This performance jump is even more significant with smaller models, such as DeepLabV3+, OCRNet, and Segformer [MiT-B1]. This demonstrates that weaker models can be made more robust by using TTM. An interesting observation that supports this reasoning is that the weakest DeepLabV3+ model, when used with TTM, outperforms a much stronger Segformer [MiT-B5] model for the CS  $\rightarrow$  ACDC benchmark. The same trend persists with CS  $\rightarrow$  DarkZurich and CS  $\rightarrow$  BDD100K-Night. Also, for CS  $\rightarrow$  BDD100K, Flux1. Kontext Max outperforms QIE-2509 likely because it is more robust to geographical shifts, which is the primary change in BDD100K w.r.t Cityscapes. Qualitative results for domain generalized semantic segmentation on different target datasets are shown in Fig. 4.

**Ablation** We perform an ablation study with QIE-2509 in Table 2, where we generate 10 different Dark Zurich validation sets (for 10 random seeds with the same prompt) using TTM and report the average mIoU across these generations with standard deviation. We try to study the scale of randomization in these I2I models. Considering a small validation set (50 images), the deviation is reasonable, and we are still able to make significant improvements. We perform another ablation in Table 3 where we use five different I2I generation models and compare the segmentation performance on the BDD100K-Night dataset. We use the same prompt for all generative models, obtained during the TTM process. Flux. 1 Kontext Max performs the best when used for TTA. We show a comparison of the generation quality of test-modified images between different generation models in Fig. 5. As can be seen, all models are capable of performing night-to-day transfer with good fidelity. However, Flux.1 Kontext Max maintains the best semantic consistency among all models, as highlighted in the figure. In another ablation study, we show the effect of correct

Table 1. Semantic segmentation performance (mIoU in %) on different domain generalization benchmarks with and without (base model) using image to image generation model based test-time modification (TTM). CS refers to Cityscapes and Geo. shift refers to geographic shift.

Method	DeepLabV3+ [ResNet-101-d8]	OCRNet [HRNetV2-W48]	Segformer [MiT-B1]	Segformer [MiT-B5]	Mask2Former [Swin-L]	Avg mIoU	$\Delta\uparrow$ (%)
<b>Clear-to-Adverse-Weather: CS <math>\rightarrow</math> ACDC</b>							
Base Model	42.0	45.5	46.9	57.1	60.5	50.4	
+ Flux1. Kontext Max (TTM)	56.5	56.0	55.7	61.3	64.3	58.8	
+ QIE-2509 (TTM)	58.3	58.9	58.9	64.6	66.4	<b>61.4</b>	<b>+21.8%</b>
<b>Day-to-Nighttime: CS <math>\rightarrow</math> DarkZurich</b>							
Base Model	19.8	22.9	23.2	36.8	40.6	28.6	
+ Flux1. Kontext Max (TTM)	41.2	40.9	43.1	47.9	51.7	45.0	
+ QIE-2509 (TTM)	44.8	43.9	44.9	46.7	51.3	<b>46.3</b>	<b>+61.8%</b>
<b>Day-to-Nighttime + Geo Shift: CS <math>\rightarrow</math> BDD100K-Night</b>							
Base Model	23.8	24.4	25.8	34.0	40.6	29.7	
+ Flux1. Kontext Max (TTM)	42.3	37.6	41.9	44.7	49.1	43.1	
+ QIE-2509 (TTM)	42.8	43.0	41.8	46.6	47.1	<b>44.3</b>	<b>+49.1%</b>
<b>Clear-to-Adverse-Weather + Geo Shift: CS <math>\rightarrow</math> BDD100K</b>							
Base Model	48.8	49.7	47.4	55.2	58.0	51.8	
+ Flux1. Kontext Max (TTM)	53.2	52.8	50.5	55.8	58.7	<b>54.2</b>	<b>+4.6%</b>
+ QIE-2509 (TTM)	50.7	50.5	48.7	53.7	55.0	51.7	

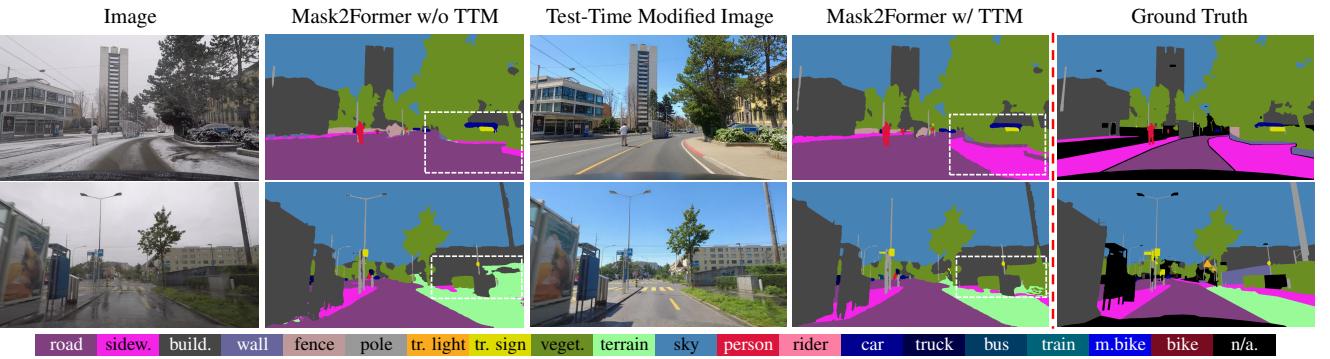


Figure 4. Qualitative performance comparison for domain generalized semantic segmentation on ACDC dataset - with and without using TTM. The same, unmodified Mask2Former trained on Cityscapes is used as the segmentation model for inference. Dashed boxes show how TTM can reduce aleatoric uncertainty of the modified image by incorporating world knowledge, leading to improved predictions.

Table 2. CS  $\rightarrow$  DarkZurich: mIoU across 10 seeds (mean  $\pm$  standard deviation) for 10 different versions of DarkZurich generated using QIE-2509 with the same text prompt. The second row reports average mIoU with standard deviation.

	DeepLabV3+	OCRNet	MiT-B1	MiT-B5	Mask2Former
w/o TTM	19.8	22.9	23.2	36.8	40.6
w/ TTM	$43.7 \pm 1.2$	$43.1 \pm 1.3$	$42.9 \pm 1.6$	$47.5 \pm 1.3$	<b><math>50.5 \pm 1.2</math></b>

prompting in Fig. 6 by generating images using QIE-2509. Handcrafted short prompts, such as "remove snow from the image," often do not yield the desired output. Metaprompting an MLLM, such as GPT-5, is a better strategy; however, without a precise query, an MLLM can generate weak prompts. As we see in Fig. 6c, the car on the road also gets removed, resulting in a significant layout mismatch. In contrast, our query formulation from the TTM process, as

Table 3. CS  $\rightarrow$  BDD100K-Night: segmentation performance comparison for different I2I generative models with Mask2Former segmentation model.

	w/o TTM	Flux1. Kont Dev [14]	Nano- Banana [5]	SeedEdit 3.0 [33]	QIE- 2509	Flux1. Kont Max
mIoU	40.6	42.3	43.5	45.3	47.1	<b>49.1</b>

described in Section 3, results in complete snow removal while preserving the semantic layout.

## 4.2. Object Detection

A performance comparison for the domain-generalized object detection task on Cityscapes  $\rightarrow$  BDD100 K Night-Det is shown in Table 4. We report the mean Average Precision (mAP) metric at 0.5 threshold. We find that TTM yields no

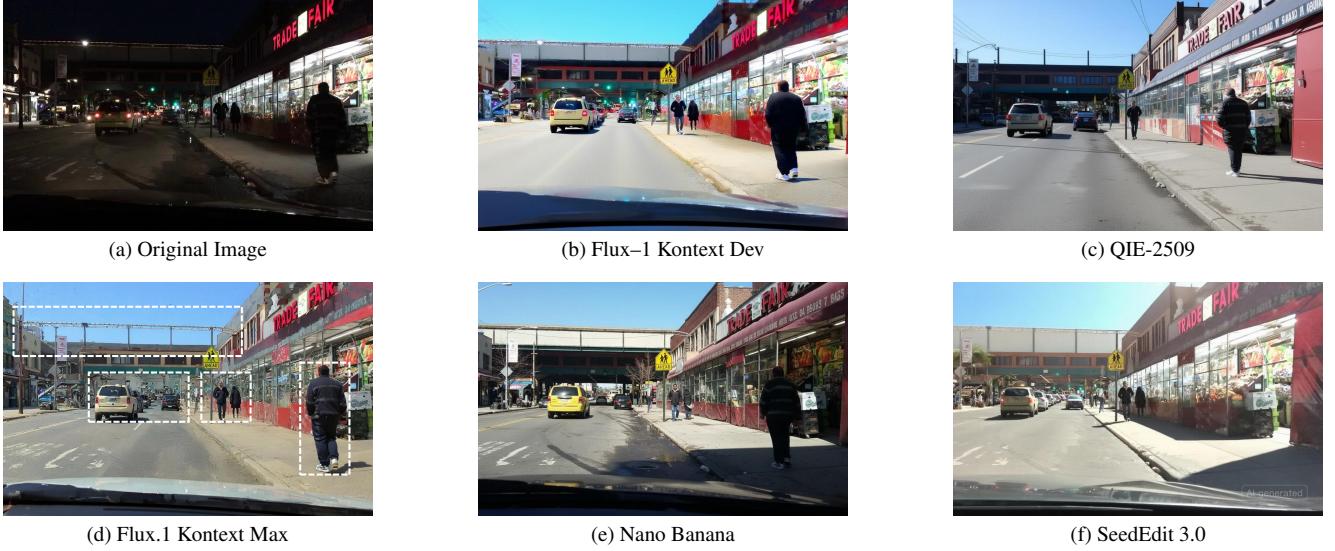


Figure 5. Qualitative comparison of the generation quality of test-modified images among different I2I generation models on BDD100K-Night. Dashed boxes highlight key layout components that are not respected by some generations.

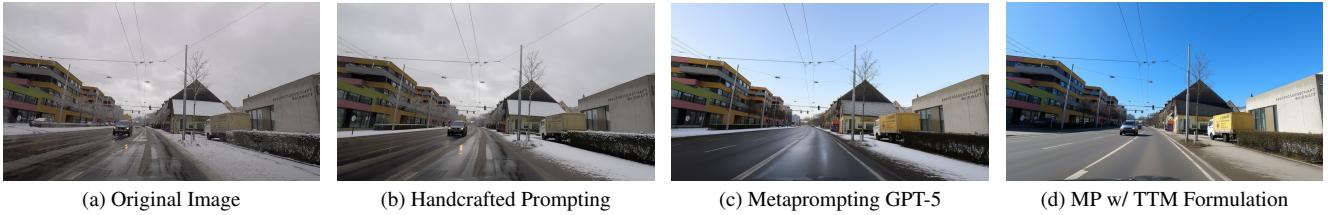


Figure 6. Comparing prompting strategies for inverse domain transfer using QIE-2509. MP in (d) refers to Metaprompting GPT-5



Figure 7. Qualitative performance comparison for domain generalized object detection on Cityscapes  $\rightarrow$  BDD100KNight-Det using Faster-RCNN - without and with using TTM. Three different examples from BDD100KNight-Det dataset. (Zoom recommended)

table gains by doubling and even tripling the mAP score for Faster-RCNN and Mask R-CNN, respectively. Qualitative results are shown in Fig. 7. As is visually evident from the

figure, detecting objects becomes much easier after TTM, as challenging nighttime conditions are turned into easier daytime ones.

### 4.3. Image Classification

For image classification task, apart from the base model (w/o TTM), we also compare against other models introduced in [10] that are trained on ImageNet-1K while using additional strategies to reduce the domain gap like pre-training on ImageNet-21K (10x the size of ImageNet-1K), different augmentation techniques, using self-attention for better spatial feature learning, and adversarial training. However, we only take raw ResNet-50 and ResNet-152 models pre-trained on ImageNet-1K as candidates for applying TTM without using any of the additional domain reduction techniques. Here, due to the large validation set size, for image generation, we use a 4-bit quantized variant of QIE-2509 with 8-step Lightning LoRA weights for faster inference. The improvement should be even higher if we use the non-quantized complete QIE-2509 model.

Quantitative results with our TTM method are shown in Table 5. As can be seen, we are able to get the highest

Table 4. Object detection performance (mAP@50) on Cityscapes→BDD100KNight-Det in a domain generalization setting.

Model	Method	Per-Class mAP@50 (%)							Overall	
		person	rider	car	truck	bus	motor.	bicycle	mAP@50	Δ↑ (%)
Faster R-CNN	Without TTM	16.1	22.8	18.6	4.3	1.0	3.3	41.1	13.4	
	+ Flux1. Kontext Max (TTM)	36.9	29.7	51.4	29.7	12.3	30.6	36.3	<b>28.4</b>	+111.9%
	+ QIE-2509 (TTM)	5.0	22.8	23.7	10.9	22.0	27.7	34.9	18.4	
Mask R-CNN	Without TTM	15.9	22.8	13.6	1.3	0.0	2.5	25.6	10.2	
	+ Flux1. Kontext Max (TTM)	35.1	40.0	51.6	31.7	9.3	25.2	61.1	<b>31.8</b>	+211.8%
	+ QIE-2509 (TTM)	4.5	22.8	24.5	12.1	14.7	36.6	36.6	19.0	

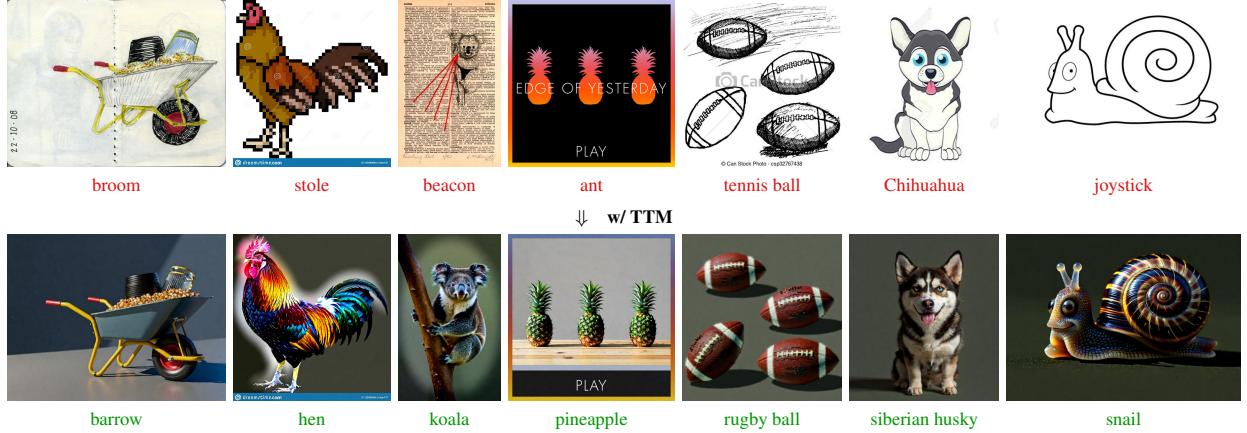


Figure 8. Performance comparison on domain generalized image classification with ResNet-50. Top row: predictions on original ImageNet-R images (without TTM). Bottom row: predictions on the same images after TTM. **Green** = correct, **red** = incorrect.

Table 5. Image classification performance comparison on ImageNet-R dataset. Top-1 accuracies are reported in %.

Method	ImageNet-R (%)	Δ↑ (%)
ResNet-50	36.1	
+ ImageNet-21K <i>Pretraining</i>	37.2	+3.0%
+ CBAM ( <i>Self-Attention</i> )	36.8	+1.9%
+ $\ell_\infty$ Adversarial Training	31.4	-13.0%
+ Speckle Noise	37.9	+5.0%
+ Style Transfer Augmentation	41.5	+15.0%
+ AugMix	41.1	+13.8%
+ DeepAugment	42.2	+16.9%
+ DeepAugment + AugMix	46.8	+29.6%
<b>ResNet-50 + TTM (Ours)</b>	<b>60.8</b>	<b>+68.5%</b>
ResNet-152	41.3	
<b>ResNet-152 + TTM (Ours)</b>	<b>63.5</b>	<b>+53.8%</b>

top-1 accuracy with ResNet-50 using TTM on ImageNet-R without employing any additional training strategies. Moreover, even with the larger ResNet-152 model, we are able to significantly improve the top-1 accuracy. Additionally, following a similar trend to domain generalized semantic segmentation, applying TTM with ResNet-50 even surpasses the performance of a larger ResNet-152 model (from 41.3 to 60.8). Qualitative results are shown in Fig. 8. As visible

from the figure, the transformation of different renditions to real-life-like objects reduces the domain gap w.r.t ImageNet thereby helping to achieve better image classification performance for a fixed discriminative model.

## 5. Conclusion

In this paper, we propose an efficient approach to improve the perception robustness during inference time called Test-Time Modification (TTM). During TTM, we perform an inverse domain transformation of the target domain images back to the source domain by implicitly distilling the world knowledge of state-of-the-art Image-to-Image generation models. To demonstrate the effectiveness of our approach, we experimentally focus on the domain generalization setting for semantic segmentation, object detection, and image classification tasks, utilizing challenging target data distributions across various weather and daytime conditions, as well as rendering shifts. The significant improvements in our quantitative and qualitative results across multiple downstream tasks, datasets, and generation models demonstrate that test-time modification using large, information-rich Image-to-Image generation models can significantly help in creating more robust perception pipelines for challenging environments.

## 6. Acknowledgment

The research is funded by the German Federal Ministry for Economic Affairs and Energy within the project “NXT GEN AI METHODS – Generative Methoden für Perzeption, Prädiktion und Planung”. We thank the consortium and also acknowledge the Gauss Centre for Supercomputing for providing computing time on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC).

## References

[1] Yasser Benigmim, Subhankar Roy, Slim Essid, Vicky Kalogeiton, and Stéphane Lathuilière. Collaborating foundation models for domain generalized semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3108–3119, 2024. [2](#)

[2] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. [5](#)

[3] 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. [5](#)

[4] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T. Kim, Seungryong Kim, and Jaegul Choo. RobustNet: Improving Domain Generalization in Urban-Scene Segmentation via Instance Selective Whitening. In *Proc. of CVPR*, pages 11580–11590, virtual, 2021. [2](#)

[5] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blissein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025. [6](#)

[6] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. [5](#)

[7] 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](#)

[8] 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](#)

[9] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [5](#)

[10] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kada-vath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8340–8349, 2021. [2, 5, 7](#)

[11] Jiaxing Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. FSDR: Frequency Space Domain Randomization for Domain Generalization. In *Proc. of CVPR*, pages 6891–6902, virtual, 2021. [2](#)

[12] Namyup Kim, Taeyoung Son, Cuiling Lan, Wenjun Zeng, and Suha Kwak. WEDGE: Web-Image Assisted Domain Generalization for Semantic Segmentation. *arXiv:2109.14196*, pages 1–14, 2021. [2](#)

[13] Orest Kupyn and Christian Rupprecht. Dataset enhancement with instance-level augmentations. In *European Conference on Computer Vision*, pages 384–402. Springer, 2024. [3](#)

[14] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey, Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini, Axel Sauer, and Luke Smith. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space, 2025. [2, 3, 6](#)

[15] Ming Li, Taojiannan Yang, Huafeng Kuang, Jie Wu, Zhaoning Wang, Xuefeng Xiao, and Chen Chen. Controlnet++: Improving conditional controls with efficient consistency feedback. In *European Conference on Computer Vision*, 2024. [2](#)

[16] Yumeng Li, Margret Keuper, Dan Zhang, and Anna Khoreva. Adversarial supervision makes layout-to-image diffusion models thrive. In *The Twelfth International Conference on Learning Representations*, 2024. [2](#)

[17] Annika Mütze, Matthias Rottmann, and Hanno Gottschalk. Semi-supervised domain adaptation with cyclegan guided by a downstream task loss, 2022. [2](#)

[18] 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. [2, 3](#)

[19] Joshua Niemeijer, Jan Ehrhardt, Heinz Handels, and Hristina Uzunova. Uncertainty-aware controlnet: Bridging domain gaps with synthetic image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 4184–4193, 2025. [2](#)

[20] OpenAI. Gpt-5: Openai’s multimodal large language model. <https://openai.com/index/introducing-gpt-5/>, 2025. [4](#)

[21] Xingang Pan, Ping Luo, Jianping Shi, and Xiaou Tang. Two at Once: Enhancing Learning and Generalization Capacities via IBN-Net. In *Proc. of ECCV*, pages 464–479, Munich, Germany, 2018. [2](#)

[22] Xingang Pan, Xiaohang Zhan, Jianping Shi, Xiaou Tang, and Ping Luo. Switchable Whitening for Deep Representation Learning. In *Proc. of ICCV*, pages 1863–1871, Seoul, Korea, 2019. [2](#)

[23] Duo Peng, Yinjie Lei, Lingqiao Liu, Pingping Zhang, and Jun Liu. Global and Local Texture Randomization for Synthetic-to-Real Semantic Segmentation. *IEEE Transactions on Image Processing*, 30:6594–6608, 2021. 2

[24] Duo Peng, Yinjie Lei, Munawar Hayat, Yulan Guo, and Wen Li. Semantic-Aware Domain Generalized Segmentation. In *Proc. of CVPR*, pages 2594–2605, New Orleans, LA, USA, 2022. 2

[25] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 5

[26] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022. 2, 3

[27] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7374–7383, 2019. 2, 5

[28] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Map-guided curriculum domain adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):3139–3153, 2020. 5

[29] Christos Sakaridis, Haoran Wang, Ke Li, René Zurbrügg, Arpit Jadon, Wim Abbeloos, Daniel Olmeda Reino, Luc Van Gool, and Dengxin Dai. Acdc: The adverse conditions dataset with correspondences for robust semantic driving scene perception. *arXiv preprint arXiv:2104.13395*, 2021. 2, 5

[30] Manuel Schwonberg, Joshua Niemeijer, Jan-Aike Termöhlen, Jörg P. Schäfer, Nico M. Schmidt, Hanno Gottschalk, and Tim Fingscheidt. Survey on Unsupervised Domain Adaptation for Semantic Segmentation for Visual Perception in Automated Driving. *IEEE Access*, 11: 54296–54336, 2023. 2

[31] Mattia Segu, Alessio Tonioni, and Federico Tombari. Batch normalization embeddings for deep domain generalization. *Pattern Recognition*, 135:109115, 2023. 2

[32] Jan-Aike Termöhlen, Marvin Klingner, Leon J. Brettin, Nico M. Schmidt, and Tim Fingscheidt. Continual unsupervised domain adaptation for semantic segmentation by online frequency domain style transfer. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 2881–2888, 2021. 2

[33] Peng Wang, Yichun Shi, Xiaochen Lian, Zhonghua Zhai, Xin Xia, Xuefeng Xiao, Weilin Huang, and Jianchao Yang. Seededit 3.0: Fast and high-quality generative image editing. *arXiv preprint arXiv:2506.05083*, 2025. 6

[34] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, Yuxiang Chen, Zecheng Tang, Zekai Zhang, Zhengyi Wang, An Yang, Bowen Yu, Chen Cheng, Dayiheng Liu, Deqing Li, Hang Zhang, Hao Meng, Hu Wei, Jingyuan Ni, Kai Chen, Kuan Cao, Liang Peng, Lin Qu, Minggang Wu, Peng Wang, Shuting Yu, Tingkun Wen, Wensen Feng, Xiaoxiao Xu, Yi Wang, Yichang Zhang, Yongqiang Zhu, Yujia Wu, Yuxuan Cai, and Zenan Liu. Qwen-image technical report, 2025. 2, 3

[35] Weijia Wu, Yuzhong Zhao, Hao Chen, Yuchao Gu, Rui Zhao, Yefei He, Hong Zhou, Mike Zheng Shou, and Chunhua Shen. Datasetdm: Synthesizing data with perception annotations using diffusion models. *Advances in Neural Information Processing Systems*, 36:54683–54695, 2023. 3

[36] Weijia Wu, Yuzhong Zhao, Mike Zheng Shou, Hong Zhou, and Chunhua Shen. Diffumask: Synthesizing images with pixel-level annotations for semantic segmentation using diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1206–1217, 2023. 3

[37] Enze Xie, Wenhui Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems*, 34: 12077–12090, 2021. 5

[38] Qi Xu, Liang Yao, Zhengkai Jiang, Guannan Jiang, Wenzheng Chu, Wenhui Han, Wei Zhang, Chengjie Wang, and Ying Tai. DIRL: Domain-Invariant Representation Learning for Generalizable Semantic Segmentation. In *Proc. of AAAI*, pages 2884–2892, virtual, 2022. 2

[39] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020. 2, 5

[40] Runpeng Yu, Songhua Liu, Xingyi Yang, and Xinchao Wang. Distribution shift inversion for out-of-distribution prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3592–3602, 2023. 2

[41] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Object-contextual representations for semantic segmentation. In *European conference on computer vision*, pages 173–190. Springer, 2020. 5

[42] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain Randomization and Pyramid Consistency: Simulation-to-Real Generalization Without Accessing Target Domain Data. In *Proc. of ICCV*, pages 2100–2110, Seoul, Korea, 2019. 2

[43] Lvmin Zhang and Maneesh Agrawala. Adding Conditional Control to Text-to-Image Diffusion Models. *arXiv:2302.05543*, pages 1–33, 2023. 2

[44] Yuyang Zhao, Zhun Zhong, Na Zhao, Nicu Sebe, and Gim Hee Lee. Style-hallucinated dual consistency learning for domain generalized semantic segmentation. In *Computer Vision – ECCV 2022*, pages 535–552, Cham, 2022. Springer Nature Switzerland. 2

[45] Zhun Zhong, Yuyang Zhao, Gim Hee Lee, and Nicu Sebe. Adversarial style augmentation for domain generalized urban-scene segmentation. In *Advances in Neural Information Processing Systems*, pages 338–350. Curran Associates, Inc., 2022. 2