

Diffusion Features to Bridge Domain Gap for Semantic Segmentation

Yuxiang Ji*

*Institute of Artificial Intelligence
Xiamen University
Xiamen, China
yuxiangji@stu.xmu.edu.cn*

Boyong He*

*Institute of Artificial Intelligence
Xiamen University
Xiamen, China
boyonghe@stu.xmu.edu.cn*

Chenyuan Qu

*School of Computer Science
University of Birmingham
Birmingham, United Kingdom
cxq134@student.bham.ac.uk*

Zhuoyue Tan

*Institute of Artificial Intelligence
Xiamen University
Xiamen, China
tanzhuoyue@stu.xmu.edu.cn*

Chuan Qin

*Department of Infrastructure Engineering
The University of Melbourne
Melbourne, Australia
qinqc1@student.unimelb.edu.au*

Liaoni Wu[†]

*Institute of Artificial Intelligence
Xiamen University
Xiamen, China
wuliaoni@xmu.edu.cn*

Abstract—Pre-trained diffusion models have demonstrated remarkable proficiency in synthesizing images across a wide range of scenarios with customizable prompts, indicating their effective capacity to capture universal features. Motivated by this, our study delves into the utilization of the implicit knowledge embedded within diffusion models to address challenges in cross-domain semantic segmentation. This paper investigates the approach that leverages the sampling and fusion techniques to harness the features of diffusion models efficiently. We propose Diffusion Feature Fusion (DIFF) as a backbone use for extracting and integrating effective semantic representations through the diffusion process. By leveraging the strength of text-to-image generation capability, we introduce a new training framework designed to implicitly learn posterior knowledge from it. Through rigorous evaluation in the contexts of domain generalization semantic segmentation, we establish that our methodology surpasses preceding approaches in mitigating discrepancies across distinct domains and attains the state-of-the-art (SOTA) benchmark. The implementation code is released at <https://github.com/YuxiangJi/DIFF>.

Index Terms—Diffusion model, domain generalization, semantic segmentation

I. INTRODUCTION

The paradigm of training segmentation models on large-scale datasets has demonstrated significant successes; nevertheless, the obstacles associated with the acquisition of data specific to niche scenarios, continue to pose significant challenges. Synthetic data, while complementing some missing data, usually suffers from the issue of domain gaps. This issue arises because models trained on limited synthetic data tend to decline in accuracy when applied in real-world settings, attributable to domain shifts in the test data [1], [2]. Research has shown that one of the important factors is the representation discrepancy caused when the perspective, background, style, or imaging conditions are changed to the unseen domain [3]–[5]. Taking this issue, the study of Domain Generalization (DG) focuses on reducing the domain variance performance and improving the model robustness across unseen domains.

Recently, the stunning performance of diffusion models on various tasks of image generation has attracted a great deal of research attention. The pre-trained text-to-image diffusion model (e.g., Stable Diffusion [6]) possesses the capability to synthesis images of remarkable realism and high quality across diverse styles, scenes, and categories, contingent upon the customized prompts given. This indicates that the diffusion model learns generic visual features while being able to disentangle the representations of image features according to conditional text inputs. Several studies also validate the ability of pre-trained diffusion models on representational and perceptual tasks [7]–[12]. Drawing inspiration from the implicit universal knowledge embedded within pre-trained diffusion models, it leads us to think: *how to utilize such knowledge to reduce the domain discrepancy in semantic segmentation?*

We therefore introduce **DIFFusion Feature Fusion (DIFF)**, a module based on the pre-trained Stable Diffusion model, to collect and integrate the feature sets from the whole diffusion process. Given an image from arbitrary domains, robust features modeled in the diffusion embedding space could be extracted from the Stable Diffusion. Different from similar studies [13], [14] that use single-step denoising for extracting features, we consider the diffusion trajectory as a more meaningful feature. The feature sets from the whole multi-step diffusion process will be fused by a convolutional fusion block and aligned to the visual embedding from the standard backbones, e.g. ResNet [15], Vision Transformer [16]. To further utilize the conditional generation capability of the pre-trained diffusion model and to address the absence of corresponding annotation text as conditional input when doing prediction, we introduce a special implicit posterior knowledge learning framework for supervised learning. By leveraging vision-language joint modeling, the implicit posterior knowledge from the conditional generation capability of diffusion models could be learned to maintain generalization when confronted with unseen data.

* Contribute equally to the work.

† Corresponding author.




Figure 1: Overview of proposed diffusion feature fusion (DIFF) module and implicit posterior knowledge learning (IPKL) training pipeline. (a) In the conditional branch, we extract the categories and reference masks from the semantic segmentation annotations and use them as conditions, which are input into the DIFF module along with the input image. Through the DIFF module, we obtain features enhanced with conditional information for supervised training. (b) In the unconditional branch, we only use the image as input to the DIFF module and employ the prediction results from the conditional branch as a teacher for consistency learning.

II. METHOD

By given a text-image paired dataset, the original generation-purpose text-to-image diffusion model is set up to model a process that gradually removes noise from a standard noise distribution to an image distribution based on text prompt, or in other words, it possesses the conditional distribution $p(\mathcal{X}|\mathcal{C})$. As for standard segmentation tasks, let $\{\mathcal{X}, \mathcal{Y}\}$ be the set of images and mask labels, the target is to learn $p(\mathcal{Y}|\mathcal{X})$. Considering text as a higher-level understanding, it is straightforward to leverage the prior knowledge $p(\mathcal{X}|\mathcal{C})$ of pre-trained text-to-image diffusion models to help learn a more generalized posterior distribution for segmentation models.

Therefore, like some recent perceptual works based on diffusion models [7], [8], [11]–[14], we rely on the critical denoising component U-Net [17] for a more generalized encoding. Unlike the usual generative tasks where the noise predicted by U-Net is gradually removed, we perform an inverse process, that the predicted noise is gradually added to the input image. Different from similar studies [7], [13] that use single-step denoising for extracting features, we consider the whole multi-step diffusion trajectory as a more meaningful feature. In order to retain more information from pre-trained diffusion model without disruption and to ensure that the features can be utilized by the segmentation head model, we freeze the entire U-Net and sample as many effective features as possible for fusion, as detailed in Sec. II-A. On the other hand, considering that standard segmentation tasks do not have text prompts as input, we introduce a dual-stream network for learning posterior knowledge to perform conventional segmentation predictions, as detailed in Sec. II-B. The overview of the

proposed pipeline is shown in Fig. 1, where the training is performed on two branches and the prediction is only performed on the unconditional branch. The trainable modules of the two branches share weights.

A. Diffusion Feature Fusion

In alignment with the architectural design and pre-training scheme of diffusion model, we devise two distinct feature sets from U-Net aimed at facilitating the integration of visual and text semantic understanding. Specifically, we extract the intermediate features $\{\mathcal{F}_{t,l}^{\text{inter}}\} \in \mathbb{R}^{d_l \times w_l \times h_l}$ from each layer l and each step t within the U-Net decoder as the visual representation, and the cross-attention maps $\{\mathcal{F}_{t,l}^{\text{cross}}\} \in \mathbb{R}^{d_l \times w_l \times h_l}$ as the interaction representations between visual and text content. However, directly taking out the features of the whole process layer by layer would lead to too numerous and unwieldy for practical utilization, which is also the reason that many works apply a one-step diffusion pipeline and hand-select the block, step pairs by grid search. Inspired by DiffHyperfeature [11], here we propose Diffusion Feature Fusion (DIFF) to fuse features extracted from different layers with different steps by an aggregating convolutional block, including a series of convolution, normalization, and activation layers.

Formally, given intermediate features $\{\mathcal{F}_{t,l}^{\text{inter}}\}$ and cross-attention maps $\{\mathcal{F}_{t,l}^{\text{cross}}\}$ from the diffusion process of the input images, the final features for segmenting $\mathcal{F}_{\text{diff}}$ takes the form

$$\mathcal{F}_{\text{diff}} = F_{\text{conv}}(\oplus_{t,l}[\mathcal{F}_{t,l}^{\text{inter}}, \mathcal{F}_{t,l}^{\text{cross}}]), \quad (1)$$

where $\oplus_{t,l}$ means the concatenation across t and l .




Figure 2: Segmenting prediction on the unseen data of existing SOTA domain generalization (DG) semantic segmentation methods (DAFormer [18], ReVT [19], CMFormer [20]) and our method.

B. Implicit Posterior Knowledge Learning

In order to learn posterior knowledge $p(\mathcal{Y}|\mathcal{X})$ for standard semantic segmentation from the prior knowledge $p(\mathcal{X}|\mathcal{C})$ of pre-trained diffusion models, we still have two unresolved issues: (1) The first issue is that original text-to-image diffusion models are set up to model the relationship between the overall textual description and the image, which leaves a gap to the pixel-level fine-grained segmentation. (2) The second issue is that there are no reference guided input provided when predicting in the standard semantic segmentation tasks. To address these two issues, we adopt path controlled diffusion in Sec. II-B1 and unconditional consistency learning in Sec. II-B2.

1) Path Controlled Diffusion

To ensure the extracted features correspond more precisely with the textual prompts, we adopt a training-free method following MultiDiffusion [21]. By given segmentation data $\{\mathcal{X}, \mathcal{Y}\}$ during training, where \mathcal{X} are images and \mathcal{Y} are pixel-level semantic annotations, we decompose the annotations \mathcal{Y} into multiple groups of masks $\mathcal{M} \in \{0, 1\}^{\text{cls} \times w \times h}$ and their corresponding category descriptions $\mathcal{C} \in \{\text{sky}, \text{vegetation}, \text{road}, \text{people}, \text{car}, \dots\}^{\text{cls}}$, where cls represents the number of all categories. Then we feed the category descriptions \mathcal{C} as text prompts into the diffusion model and fuse the predicted latent variables by the corresponding masks \mathcal{M} at each diffusion step, formulate as

$$I_{t+1}(I_t, \mathcal{M}, \mathcal{C}) = \sum_{i=1}^{\text{cls}} \frac{\mathcal{M}_i}{\sum_{j=1}^{\text{cls}} \mathcal{M}_j} [\mathcal{M}_i \odot (\Phi(I_t, \mathcal{C}_i))], \quad (2)$$

where Φ represents the pre-trained text-to-image diffusion model and I represent the latent variables.

Then by feeding such conditional features extracted by DIFF to the segment head D , we could train the entire model using cross-entropy, just like in conventional segmentation models as

$$\mathcal{L}_{\text{condit}} = \text{CE}(D(\mathcal{F}_{\text{diff}}^{\text{con}}(\mathcal{X}, \mathcal{M}, \mathcal{C})), \mathcal{Y}). \quad (3)$$

2) Unconditional Consistency Learning

To address the issue of not having corresponding text prompts during prediction, we set up an additional unconditional branch

as shown in Fig. 1. In this branch, the diffusion features $\mathcal{F}_{\text{diff}}^{\text{uncon}}$ are acquired with DIFF through an unconditional process $p(\mathcal{F}|\mathcal{X}, \mathcal{C} = \emptyset)$. On this basis, we simply employ a L2 loss for consistency learning between the outputs from two branches, forcing them to produce nearly identical outputs, as

$$\mathcal{L}_{\text{consis}} = \|D(\mathcal{F}_{\text{diff}}^{\text{con}}(\mathcal{X}, \mathcal{M}, \mathcal{C})), D(\mathcal{F}_{\text{diff}}^{\text{uncon}}(\mathcal{X}))\|_2. \quad (4)$$

In this way, the implicit posterior knowledge learned from the conditional generative model could be gradually distilled onto the unconditional input branch for predicting. The complete learning objective is the combination of conditional segmentation loss $\mathcal{L}_{\text{condit}}$ and consistency loss $\mathcal{L}_{\text{consis}}$ as

$$\mathcal{L}_{\text{final}} = \lambda_1 \mathcal{L}_{\text{condit}} + \lambda_2 \mathcal{L}_{\text{consis}}. \quad (5)$$

Here, λ_1 and λ_2 are two hyper-parameters to control the weights of two objectives, which are both set to 1 in our experiments.

III. EXPERIMENTS

A. Experimental Setup

Following the setup of prior research on domain generalization for semantic segmentation [13], [18], [22], the model trained on a source domain dataset will be evaluated on a series of unknown target datasets. Based on two synthetic datasets GTAV [23] and Synthia [24], three real-world clear datasets Cityscapes (CS) [25], BDD-100K (BDD) [26], and Mapillary Vistas (MV) [27], and two real-world adverse weather datasets ACDC [28] and Dark Zurich (DZ) [29], we consider two practically significant generalization tasks: *synthetic-to-real* and *real-to-adverse*. *Synthetic-to-real*: We use two synthetic datasets GTAV and Synthia as source domain, and five real-world datasets CS, BDD, MV, ACDC, and DZ as target domain. *Real-to-Adverse*: We use clear dataset CS as source domain, and two adverse weather datasets ACDC and DZ as target domain.

For implementation, our DIFF module is based on the released Stable-Diffusion v2-1 [6] checkpoint, and we use the same decoder head and training settings as in DAFormer [18].

Table I: Synthetic-to-real DG results comparing to SOTA methods on GTAV source domain. Training is performed on synthetic dataset GTAV [23]. Evaluation is performed on five real-world datasets with 19 categories.

DG Method	Backbone	mIoU (%) on						
		CS [25]	BDD [26]	MV [27]	Avg3	ACDC [28]	DZ [29]	Avg5
IBN-Net [30]	ResNet-101	37.37	32.29	33.84	33.15	-	-	-
FSDR [22]	ResNet-101	44.80	41.20	43.40	43.13	24.77	9.66	32.77
SAN-SAW [31]	ResNet-101	45.33	41.18	40.77	42.23	-	-	-
ReVT [19]	MiT-B5	49.96	48.01	53.06	50.34	41.15	21.99	42.84
DAFormer [18]	MiT-B5	52.65	47.89	54.66	51.73	38.25	17.45	42.18
CMFormer [20]	SWin-B	55.31	49.91	60.09	55.10	41.34	22.58	45.85
PromptDiff [13]	Diffusion	52.00	-	-	-	-	-	-
Ours	Diffusion	58.01	53.60	59.85	57.15	46.32	34.27	50.41
		+2.70	+3.69	-0.24	+2.05	+4.98	+11.69	+4.56

Table II: Synthetic-to-real DG results comparing to SOTA methods on Synthia source domain. Training is performed on synthetic dataset Synthia [24]. Evaluation is performed on five real-world datasets with 16 categories.

DG Method	Backbone	mIoU (%) on						
		CS [25]	BDD [26]	MV [27]	Avg3	ACDC [28]	DZ [29]	Avg5
IBN-Net [30]	ResNet-101	32.04	30.57	32.16	31.59	-	-	-
FSDR [22]	ResNet-101	40.80	39.60	37.40	39.30	-	-	-
SAN-SAW [31]	ResNet-101	38.92	35.24	34.52	36.23	-	-	-
DAFormer [18]	MiT-B5	44.08	33.20	42.99	40.09	26.62	14.14	32.21
ReVT [19]	MiT-B5	46.28	40.30	44.76	43.78	35.75	20.10	37.44
CMFormer [20]	SWin-B	44.59	33.44	43.25	40.43	34.50	19.57	35.07
PromptDiff [13]	Diffusion	49.10	-	-	-	-	-	-
Ours	Diffusion	49.31	42.20	49.47	46.99	36.27	23.39	40.13
		+0.21	+1.90	+4.71	+3.21	+0.52	+3.29	+2.69

Remark:

In tables, the best results are highlighted in **bold**, while the second best is underlined.

B. Results

Under the *synthetic-to-real* setting as shown in Tab. I, our method is demonstrated to outperform the previous SOTA DG methods based on ResNet-101 [15], MiT-B5 [32], SWin-B [33], and diffusion backbones. Specifically, our method shows higher improvements (4.98% and 11.69%) on the two adverse weather datasets ACDC and DZ, that the data with larger domain discrepancies. Comparing to the PromptDiff [13] which is also a diffusion-based method, our performance surpasses it by 6.01% on GTAV \rightarrow CS. We also extend the experiments on the synthia source domain in Tab. II for further validation.

Since our method freezes most of the network parameters, the learning capacity in the source domain could be limited. Under the *real-to-adverse* setting as shown in Tab. III, in settings with smaller domain differences (i.e., CS \rightarrow ACDC), our method is inferior (-1.30%) to those based on a transformer backbone, but it shows significant improvement (+4.38%) in settings with larger domain differences (i.e., CS \rightarrow DZ).

Fig. 2 provides some visual examples of semantic segmentation results from our method comparing to the previous SOTA DG methods. It shows that our method improves the generalization on the unseen data to a large extent, especially on the *road*, *sidewalk* categories.

C. Ablation Study

We evaluate the effects of proposed diffusion feature fusion (DIFF) and implicit posterior knowledge learning (IPKL) with different consistency losses on the overall results in Tab. IV.

Table III: Clear-to-adverse DG results comparing to SOTA methods on Cityscapes source domain. Training is performed on clear dataset Cityscapes [25]. Evaluation is performed on two adverse weather datasets with 19 categories.

DG Method	Backbone	mIoU (%) on		
		ACDC [28]	DZ [29]	Avg2
FSDR [22]	ResNet-101	47.18	22.60	34.89
SAN-SAW [31]	ResNet-101	49.00	24.80	36.90
DAFormer [18]	MiT-B5	55.15	28.28	41.72
CMFormer [20]	SWin-B	60.21	33.90	<u>47.05</u>
PromptDiff [13]	Diffusion	58.60	<u>34.00</u>	46.30
Ours	Diffusion	<u>58.91</u>	38.38	48.65
		-1.30	+4.38	+1.60

Table IV: Ablation study on the proposed components DIFF and IPKL with different consistency losses.

DIFF	IPKL	$\mathcal{L}_{\text{consis}}$	mIoU (%) on	
			GTAV [23]	\rightarrow CS [25]
-	-	-	49.72	(baseline)
✓	-	-	56.28	(+6.56)
✓	✓	w/o. $\mathcal{L}_{\text{consis}}$	44.82	(-4.90)
✓	✓	KL Loss	57.81	(+8.09)
✓	✓	L2 Loss	58.01	(+8.29)




Figure 3: Heatmap (on *sidewalk*) and segmentation results comparison between (b) w/o. IPKL; (c) w/. IPKL and w/o. reference input; (d) w/. IPKL and w/. reference input.

The results indicate that both DIFF and IPKL could improve the domain generalization performance of the model, while only introducing the IPKL training framework without using consistency loss leads to significant performance degradation. As for the choice of consistency loss function, both L2 loss and KL loss contribute to improving consistency learning, and the results of L2 loss are better.

As shown in Fig. 3, by introducing IPKL, the model could predict results (c) on the unseen data without a guided reference input similar to those (d) with a guided reference input. Both of them are significantly superior to the results without IPKL (b).

IV. CONCLUSION

This paper delves into the potential of representations from pre-trained diffusion models in the challenging context of domain generalization for semantic segmentation. By introducing Diffusion Feature Fusion (DIFF) and implicit posterior knowledge learning (IPKL), the network could uniquely model cross-domain features by leveraging the rich prior knowledge of pre-trained diffusion models. Extensive experiments on multiple settings demonstrated the superior performance of our method compared to the existing domain generalization semantic segmentation methods.

REFERENCES

- [1] J. Wang, C. Lan, C. Liu, Y. Ouyang, T. Qin, W. Lu, Y. Chen, W. Zeng, and P. Yu, "Generalizing to unseen domains: A survey on domain generalization," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [2] K. Zhou, Z. Liu, Y. Qiao, T. Xiang, and C. C. Loy, "Domain generalization: A survey," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, vol. 45, no. 04, pp. 4396–4415, 2023.
- [3] H. Nam, H. Lee, J. Park, W. Yoon, and D. Yoo, "Reducing domain gap by reducing style bias," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 8690–8699.
- [4] L. Gao, L. Zhang, and Q. Zhang, "Addressing domain gap via content invariant representation for semantic segmentation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 9, 2021, pp. 7528–7536.
- [5] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *International conference on machine learning*. PMLR, 2015, pp. 1180–1189.
- [6] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10684–10695.
- [7] W. Zhao, Y. Rao, Z. Liu, B. Liu, J. Zhou, and J. Lu, "Unleashing text-to-image diffusion models for visual perception," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 5729–5739.
- [8] D. Baranchuk, A. Voynov, I. Rubachev, V. Khrlukov, and A. Babenko, "Label-efficient semantic segmentation with diffusion models," in *International Conference on Learning Representations*, 2021.
- [9] S. Mukhopadhyay, M. Gwilliam, V. Agarwal, N. Padmanabhan, A. Swaminathan, S. Hegde, T. Zhou, and A. Shrivastava, "Diffusion Models Beat GANs on Image Classification," Jul. 2023.
- [10] X. Chen, Z. Liu, S. Xie, and K. He, "Deconstructing denoising diffusion models for self-supervised learning," *arXiv preprint arXiv:2401.14404*, 2024.
- [11] G. Luo, L. Dunlap, D. H. Park, A. Holynski, and T. Darrell, "Diffusion hyperfeatures: Searching through time and space for semantic correspondence," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [12] L. Tang, M. Jia, Q. Wang, C. P. Phoo, and B. Hariharan, "Emergent correspondence from image diffusion," *Advances in Neural Information Processing Systems*, vol. 36, pp. 1363–1389, 2023.
- [13] R. Gong, M. Danelljan, H. Sun, J. D. Mangas, and L. Van Gool, "Prompting diffusion representations for cross-domain semantic segmentation," *arXiv preprint arXiv:2307.02138*, 2023.
- [14] J. Xu, S. Liu, A. Vahdat, W. Byeon, X. Wang, and S. De Mello, "Open-vocabulary panoptic segmentation with text-to-image diffusion models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 2955–2966.
- [15] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [16] A. DOSOVITSKIY, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [17] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*. Springer, 2015, pp. 234–241.
- [18] L. Hoyer, D. Dai, and L. Van Gool, "Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 9924–9935.
- [19] J.-A. Termöhlen, T. Bartels, and T. Fingscheidt, "A re-parameterized vision transformer (revt) for domain-generalized semantic segmentation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 4376–4385.
- [20] Q. Bi, S. You, and T. Gevers, "Learning content-enhanced mask transformer for domain generalized urban-scene segmentation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 2, 2024, pp. 819–827.
- [21] O. Bar-Tal, L. Yariv, Y. Lipman, and T. Dekel, "Multidiffusion: Fusing diffusion paths for controlled image generation," in *International Conference on Machine Learning*. PMLR, 2023, pp. 1737–1752.
- [22] J. Huang, D. Guan, A. Xiao, and S. Lu, "Fsd: Frequency space domain randomization for domain generalization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6891–6902.
- [23] S. R. Richter, V. Vineet, S. Roth, and V. Koltun, "Playing for data: Ground truth from computer games," in *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II 14*. Springer, 2016, pp. 102–118.
- [24] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez, "The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3234–3243.
- [25] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3213–3223.
- [26] F. Yu, H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, and T. Darrell, "Bdd100k: A diverse driving dataset for heterogeneous multitask learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2636–2645.
- [27] G. Neuhold, T. Ollmann, S. Rota Bulo, and P. Kontschieder, "The mapillary vistas dataset for semantic understanding of street scenes," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 4990–4999.
- [28] C. Sakaridis, D. Dai, and L. Van Gool, "Acdd: The adverse conditions dataset with correspondences for semantic driving scene understanding," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10765–10775.
- [29] C. Sakaridis, D. Dai, and L. V. 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*, 2019, pp. 7374–7383.
- [30] X. Pan, P. Luo, J. Shi, and X. Tang, "Two at once: Enhancing learning and generalization capacities via ibn-net," in *Proceedings of the european conference on computer vision (ECCV)*, 2018, pp. 464–479.
- [31] D. Peng, Y. Lei, M. Hayat, Y. Guo, and W. Li, "Semantic-aware domain generalized segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 2594–2605.
- [32] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo, "Segformer: Simple and efficient design for semantic segmentation with transformers," *Advances in neural information processing systems*, vol. 34, pp. 12 077–12 090, 2021.
- [33] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10012–10022.