

Causal-Tune: Mining Causal Factors from Vision Foundation Models for Domain Generalized Semantic Segmentation

Yin Zhang^{1,2}, Yongqiang Zhang^{3*}, Yaoyue Zheng^{4,2}, Bogdan Raducanu², Dan Liu¹

¹School of Instrument Science and Engineering, Harbin Institute of Technology, Harbin, China

²Computer Vision Center, Universitat Autònoma de Barcelona, Barcelona, Spain

³College of Computer Science, Inner Mongolia University, Inner Mongolia, China

⁴Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University, Shaanxi, China

{yin.zhang.hit, yaoyz105}@gmail.com, zhangyongqiang@imu.edu.cn, bogdan@cvc.uab.es, liudan@hit.edu.cn

Abstract

Fine-tuning Vision Foundation Models (VFMs) with a small number of parameters has shown remarkable performance in Domain Generalized Semantic Segmentation (DGSS). Most existing works either train lightweight adapters or refine intermediate features to achieve better generalization on unseen domains. However, they both overlook the fact that long-term pre-trained VFMs often exhibit artifacts, which hinder the utilization of valuable representations and ultimately degrade DGSS performance. Inspired by causal mechanisms, we observe that these artifacts are associated with non-causal factors, which usually reside in the low- and high-frequency components of the VFM spectrum. In this paper, we explicitly examine the causal and non-causal factors of features within VFMs for DGSS, and propose a simple yet effective method to identify and disentangle them, enabling more robust domain generalization. Specifically, we propose Causal-Tune, a novel fine-tuning strategy designed to extract causal factors and suppress non-causal ones from the features of VFMs. First, we extract the frequency spectrum of features from each layer using the Discrete Cosine Transform (DCT). A Gaussian band-pass filter is then applied to separate the spectrum into causal and non-causal components. To further refine the causal components, we introduce a set of causal-aware learnable tokens that operate in the frequency domain, while the non-causal components are discarded. Finally, refined features are transformed back into the spatial domain via inverse DCT and passed to the next layer. Extensive experiments conducted on various cross-domain tasks demonstrate the effectiveness of Causal-Tune. In particular, our method achieves superior performance under adverse weather conditions, improving +4.8% mIoU over the baseline in snow conditions.

Code — <https://github.com/zhangyin1996/Causal-Tune>

Introduction

Recently, Vision Foundation Models (VFMs), such as DINOv2 (Oquab et al. 2023), CLIP (Radford et al. 2021), EVA02 (Fang et al. 2023, 2024), and SAM (Kirillov et al. 2023), have shown strong performance on various computer vision tasks. Specifically, using VFMs as backbones and

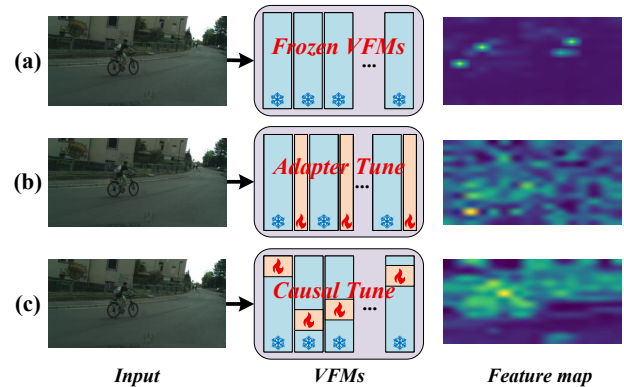


Figure 1: Visualization of DINOv2 feature maps. (a) Features extracted from the frozen DINOv2 contain noticeable artifacts. (b) Artifacts persist after applying existing adapter-based fine-tuning methods. (c) Our proposed Causal-Tune effectively suppresses these artifacts and guides the model to focus on domain-invariant causal factors.

fine-tuning them for Domain Generalized Semantic Segmentation (DGSS) (Wei et al. 2024; Zhao et al. 2025; Zhang and Tan 2025) has led to significant improvements, even surpassing traditional CNN-based methods trained from scratch (He et al. 2016; Ma et al. 2018; Zheng et al. 2025).

The predominant approach for fine-tuning VFMs is Parameter-Efficient Fine-Tuning (PEFT) (Han et al. 2024), where VFMs are kept frozen and only lightweight adapters are fine-tuned. Despite their superior performance compared to CNN-based models, transformer-based VFMs are often found to exhibit artifacts in extracted features (Darcet et al. 2023; Wang et al. 2025). As illustrated in Figure 1(a), feature extraction with frozen DINOv2 reveals such artifacts. This phenomenon primarily arises from long-term pretraining on large-scale datasets. Although such training makes them highly expressive, it also introduces considerable redundancy in the extracted features, *i.e.*, artifacts. Notably, as shown in Figure 1(b), adapter-based fine-tuning methods can improve feature map quality, yet artifacts still persist. Based on these observations, we hypothesize that these PEFT methods tend to indiscriminately fine-tune features

*Corresponding Author.




Figure 2: Left: Causal factors and non-causal factors (contain explicit and implicit non-causal factors). Right: Visualization of images adding various non-causal factors actively after DCT, high- and low-frequency filtering (H&LF), and inverse DCT.

from layers that may contain such artifacts. As a result, redundant features are not effectively suppressed, hindering the utilization of valuable representations from the VFM and ultimately degrading DGSS performance.

Recent advances in causal theory suggest that domain generalization can benefit from eliminating non-causal factors and mining the invariance of causal ones, as only causal factors tend to remain stable across domains (Lv et al. 2022; Zhang et al. 2022; Xu et al. 2023; Liu et al. 2024; Zhang et al. 2025). However, existing efforts in DGSS mainly focus on removing *explicit* non-causal factors, like style-related information. They often overlook the fact that *implicit* non-causal factors may also adversely affect performance (Xu et al. 2023). As illustrated on the left of Figure 2, commonly used DGSS datasets contain two types of non-causal factors: (i) explicit ones, such as rain, snow, fog, and night; and (ii) implicit ones, such as brightness, blur, noise, and reflections. Prior works have largely ignored these implicit factors, which may help explain the remaining artifacts observed in Figure 1(b). These observations naturally motivate our idea of removing both types of non-causal factors to purify VFM features and enhance generalization.

In this paper, we investigate two key questions and propose a method to explicitly identify and disentangle causal and non-causal factors in DGSS, to improve our understanding and enhance the performance on DGSS tasks. **(1) How can causal and non-causal factors be effectively identified during the fine-tuning of VFMs?** Frequency-based methods have long been considered effective for domain generalization. However, most works adopt the Fast Fourier Transform (FFT) (Nussbaumer 1981) or Haar Wavelet Transform (HWT) (Porwik and Lisowska 2004), which often fail to identify implicitly non-causal factors and thus lead to sub-optimal generalization performance. From a causal perspective, the frequency spectrum derived from Discrete Cosine Transform (DCT) (Ahmed, Natarajan, and Rao 2006) has been shown to better separate causal and non-causal factors (Xu et al. 2023). Therefore, we alternatively utilize DCT to convert features from the spatial domain to the frequency domain before fine-tuning, to explicitly identify both causal and non-causal factors. **(2) How can causal and non-causal factors be disentangled within the frequency domain?** We actively inject various non-causal factors into images and then apply DCT, high and low frequency filtering (H&LF), and inverse DCT to transform images from the frequency

domain back to the spatial domain. The results are visualized on the right side of Figure 2. We observe that non-causal factors tend to concentrate in both the high- and low-frequency components (second row). In contrast, residual components (last row) preserve structural and textural patterns that are relatively invariant across domains and are thus likely to represent causal factors.

To address the above challenges, we propose a simple but effective method: **Causal-Tune**. First, we extract the frequency spectrum of features from each layer using the DCT. A Gaussian band-pass filter is then applied to decompose the spectrum into causal and non-causal components. Then, the non-causal components are discarded, and only the causal components are retained for fine-tuning. A set of causal-aware learnable tokens is introduced to refine the causal components in the frequency domain. Finally, we use iDCT to transform the refined features back to the spatial domain and pass them to the next layer. In summary, the main contributions of this paper are as follows:

- We explore Parameter-Efficient Fine-Tuning of VFMs for domain generalization semantic segmentation tasks from a causal perspective.
- We propose Causal-Tune, a novel fine-tuning strategy to mine causal factors and remove non-causal factors from the features of VFMs by a band-pass filter and then we utilize a set of learnable tokens to refine the causal factors from the frequency domain.
- Extensive experiments conducted on various cross-domain tasks show the effectiveness of Causal-Tune. In particular, Causal-Tune achieves superior performance in adverse weather conditions, improving +4.8% mIoU over the baseline in ‘Snow’ condition.

Related Work

Domain Generalized Semantic Segmentation with Parameter-Efficient Fine-Tuning

The goal of Domain Generalized Semantic Segmentation (DGSS) is to enhance a model’s generalization ability on unseen domains by training only in a source domain. Recently, VFMs have demonstrated strong representation abilities thanks to large-scale pretraining, and DGSS methods based on Parameter-Efficient Fine-Tuning (PEFT) have attracted increasing attention. Rein (Wei et al. 2024) is the

first work to leverage PEFT for DGSS. They propose an embedded fine-tuning mechanism that inserts a small number of trainable parameters to refine feature maps between layers within the VFM. Subsequent PEFT-based DGSS works, such as SET (Yi et al. 2024), design spectrally decomposed tokens to learn style-invariant features, while FADA (Bi et al. 2024) proposes a frequency-adapted learning scheme to decouple style information from domain-invariant content. However, both of the above methods overlook an important fact: long-term pretrained VFMs often exhibit artifacts, and indiscriminately fine-tuning them can negatively impact performance on downstream tasks. We observe that artifacts are often associated with non-causal factors and aim to investigate this issue from a causal perspective.

Frequency-based Domain Generalization

Many works applying frequency domain transformation methods for domain generalization have demonstrated that frequency-based information is robust to domain shifts. For instance, FDA (Yang and Soatto 2020) utilizes the FFT to swap the amplitude spectrum between images, achieving strong performance in DGSS. FACT (Xu et al. 2021) addresses domain generalization by mixing the amplitude spectra of image pairs from different source domains. More recently, fine-tuning VFMs in the frequency domain for DGSS has attracted growing interest. SET (Yi et al. 2024) transforms the features of VFMs into the frequency domain using FFT, and learns style-invariant features through a set of learnable tokens. Similarly, FADA (Bi et al. 2024) leverages the HWT to extract style information from the frequency domain. However, they overlook the fact that long-term pre-trained VFMs often exhibit artifacts, which hinder the utilization of valuable representations and ultimately degrade DGSS performance. Different from previous works, we explore the relationship between the frequency domain and causal factors in DGSS, aiming to enhance generalization performance by removing non-causal factors.

Causal Mechanism in Domain Generalization

Causal mechanism suggests that domain generalization can benefit from mining invariant causal factors while removing non-causal factors, inspiring new methodologies for improving the generalization ability of models. Specifically, (Lv et al. 2022) proposes a representation learning method to extract causal factors from inputs, thereby boosting generalization performance. (Zhang et al. 2022) introduces a causal intervention reasoning module to learn invariant feature representations for multiple adverse weather generalization. UFR (Liu et al. 2024) constructs a structural causal model to analyze the limitations of domain generalization and designs a causal attention learning module to address these challenges. The most similar work to ours is MAD (Xu et al. 2023), which removes implicit non-causal factors by a multi-view adversarial discriminator. However, the main difference between MAD and our work is that MAD relies on a data-augmentation strategy, whereas our method directly fine-tunes the features of VFMs after the explicit disentangle of causal and non-causal factors.

Proposed Method

Preliminaries

Given a frozen VFM with N layers (L_1, L_2, \dots, L_N), the output feature of the i -th layer L_i is denoted as f_i . The feature propagation from layer L_i to L_{i+1} can be expressed as:

$$f_{i+1} = L_{i+1}(f_i). \quad (1)$$

The core idea of fine-tuning VFMs in DGSS is to refine the features between layers to improve the generalization on unseen domains, which can be formulated as:

$$\hat{f}_{i+1} = L_{i+1}(f_i + \Delta f_i). \quad (2)$$

As f_i remains frozen, the goal is to propose a method to generate refined features Δf_i . To this end, we design a novel fine-tuning method, which formulates the refinement as $\Delta f_i = \text{CausalTune}(f_i)$. The proposed method leverages causal mechanism to enhance feature representations by identifying and preserving causal factors while filtering out non-causal ones.

Causal & Non-causal Factors Filter

As discussed above, the spectrum obtained by the Discrete Cosine Transform (DCT) provides a more effective representation for distinguishing between causal and non-causal factors. In particular, the extremely high- and low-frequency components of the spectrum are more likely to encode non-causal information, while the remaining parts tend to capture more causal factors. Based on this observation, we adopt a simple yet effective approach by applying a band-pass filter to separate causal and non-causal components. Specifically, as illustrated in Figure 3(a), for the output features $f_i(h, w)$ extracted from layer L_i of the frozen VFM, we first use the DCT to transform $f_i(h, w)$ into the frequency domain. This process is mathematically formulated as:

$$F_i^{DCT}(u, v) = \alpha(u)\alpha(v) \sum_{h=0}^{H-1} \sum_{w=0}^{W-1} f_i(h, w) \cdot \cos\left(\frac{\pi(2h+1)u}{2H}\right) \cdot \cos\left(\frac{\pi(2w+1)v}{2W}\right),$$

$$\alpha(u) = \begin{cases} \sqrt{\frac{1}{H}}, u = 0 \\ \sqrt{\frac{2}{H}}, u > 0 \end{cases} \quad \alpha(v) = \begin{cases} \sqrt{\frac{1}{W}}, v = 0 \\ \sqrt{\frac{2}{W}}, v > 0 \end{cases} \quad (3)$$

where $F_i^{DCT}(u, v)$ denotes the frequency domain features after DCT. H and W are the height and width of the features. $\alpha(u)$ and $\alpha(v)$ are the normalization factors that ensure the energy conservation of DCT and inverse Discrete Cosine Transform (iDCT).

Then, we apply a Gaussian band-pass filter to separate causal and non-causal factors (features):

$$G(u, v) = \exp\left(-\frac{u^2 + v^2}{2R_H^2}\right) - \exp\left(-\frac{u^2 + v^2}{2R_L^2}\right), \quad (4)$$

where R_H and R_L denote the cutoff frequency for high- and low-frequency, respectively. As a result, the causal features

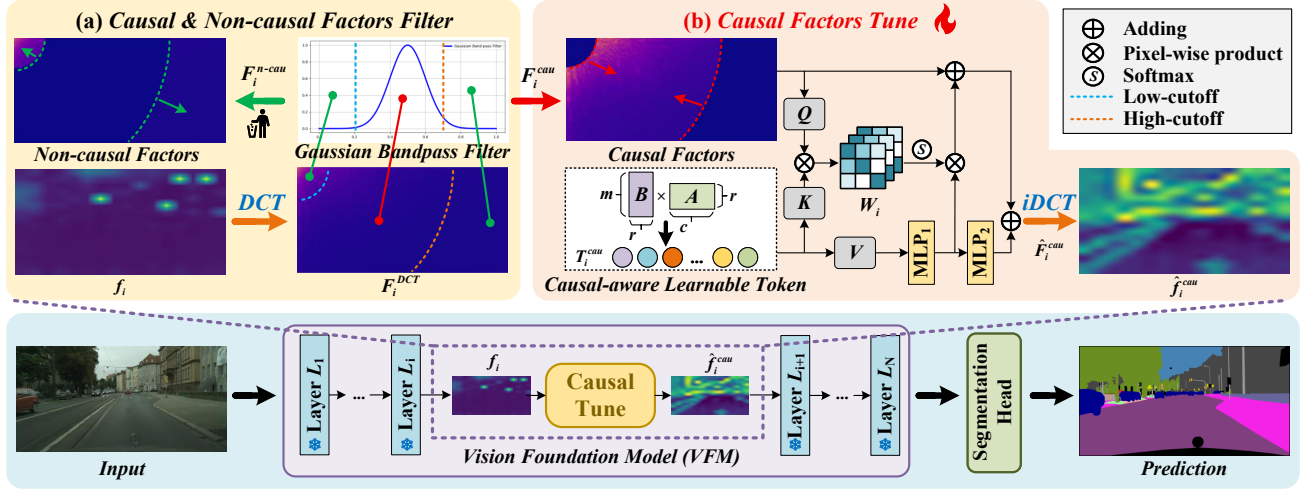


Figure 3: The pipeline of our proposed Causal-Tune. (a) The output feature f_i of layer L_i are first transformed to the frequency domain feature F_i^{DCT} using DCT, and then a Gaussian band-pass filter separates it into causal factors F_i^{cau} (red) and non-causal factors F_i^{n-cau} (green). Only the causal factors are used for subsequent fine-tuning, while the non-causal factors are discarded. (b) A series of causal-aware learnable tokens T_i^{cau} interacts with the causal factors F_i^{cau} through an attention mechanism to refine them. The refined causal factors \hat{F}_i^{cau} are then transformed back to the spatial domain \hat{f}_i^{cau} using the iDCT.

F_i^{cau} and non-causal features F_i^{n-cau} can be obtained from the frequency domain by using a Gaussian band-pass filter:

$$\begin{aligned} F_i^{cau} &= F_i^{DCT}(u, v) \cdot G(u, v), \\ F_i^{n-cau} &= F_i^{DCT}(u, v) \cdot (1 - G(u, v)). \end{aligned} \quad (5)$$

As shown in Figure 3(a), the components below the low cut-off frequency R_L and above the high cut-off frequency R_H are considered as non-causal factors F_i^{n-cau} (green), while the remaining components are considered as causal factors F_i^{cau} (red). The objective of our method is to mine causal factors from the VFMs, as they contain domain-invariant information, whereas non-causal factors primarily encode domain-specific information. Thus, the non-causal factors F_i^{n-cau} are discarded, and only the causal factors F_i^{cau} are retained for fine-tuning.

Causal Factors Tune

Building on awareness of domain-invariant causal factors, we introduce a learnable token T_i^{cau} to fine-tune these causal factors within each VFM layer. As shown in Figure 3(b), the learnable token of layer L_i can be represented as:

$$T_i^{cau} = B_i A_i, \quad T_i^{cau} \in \mathbb{R}^{m \times c}, \quad (6)$$

where $B_i \in \mathbb{R}^{m \times r}$ and $A_i \in \mathbb{R}^{r \times c}$, with m representing the length of the sequence T_i^{cau} , and c denoting the dimensionality of the feature f_i .

To allow each learnable token to capture task-specific knowledge and improve segmentation performance, we utilize an attention-based mechanism to enhance the semantic clarity of causal factors (Pak et al. 2024; Yi et al. 2024). Specifically, we consider causal factors F_i^{cau} as *Query*,

causal-aware learnable token T_i^{cau} as *Key* and *Value*. The attention weights W_i are computed through a pixel-wise product between F_i^{cau} and T_i^{cau} . Then, the token T_i^{cau} (*Value*) are projected into the feature space of F_i^{cau} using a multi-layer perceptron MLP_1 . Additionally, a residual connection is incorporated to better align the token T_i^{cau} with the corresponding causal factors F_i^{cau} . Finally, an additional MLP layer, denoted as MLP_2 , is applied for flexible feature refinement. Residual connections are incorporated to mitigate vanishing gradient issues. The refined causal factors \hat{F}_i^{cau} in the frequency domain are computed as follows:

$$W_i = \text{Softmax} \left(\frac{F_i^{cau} \times T_i^{cau \top}}{\sqrt{c}} \right), \quad (7)$$

$$\tilde{F}_i^{cau} \leftarrow F_i^{cau} + W_i \times MLP_1(T_i^{cau}), \quad (8)$$

$$\hat{F}_i^{cau} = F_i^{cau} + MLP_2(\tilde{F}_i^{cau}), \quad (9)$$

where the softmax activation function is used to normalize the attention weights. Furthermore, an inverse Discrete Cosine Transform (iDCT) is required to convert the refined causal features from the frequency domain back to the spatial domain. The detailed process is described as follows:

$$\begin{aligned} \hat{f}_i^{cau}(h, w) &= \alpha(u)\alpha(v) \sum_{u=0}^{H-1} \sum_{v=0}^{W-1} \hat{F}_i^{cau}(u, v) \cdot \\ &\quad \cos \left(\frac{\pi(2h+1)u}{2H} \right) \cdot \cos \left(\frac{\pi(2w+1)v}{2W} \right) \\ \alpha(u) &= \begin{cases} \sqrt{\frac{1}{H}}, & u = 0 \\ \sqrt{\frac{2}{H}}, & u > 0 \end{cases} \quad \alpha(v) = \begin{cases} \sqrt{\frac{1}{W}}, & v = 0 \\ \sqrt{\frac{2}{W}}, & v > 0 \end{cases} \end{aligned} \quad (10)$$

where $\hat{f}_i^{cau}(h, w)$ denotes refined causal features in the spatial domain.

Method	Trained on $C.$ \rightarrow ACDC					Trained on $C.$		Trained on $G.$		
	Night	Snow	Fog	Rain	Avg.	$\rightarrow B.$	$\rightarrow M.$	$\rightarrow C.$	$\rightarrow B.$	$\rightarrow M.$
<i>Resnet-based:</i>										
IBN (Pan et al. 2018)	21.2	49.6	63.8	50.4	46.3	48.56	57.04	33.85	32.30	37.75
Itenorm (Huang et al. 2019)	23.8	49.9	63.3	50.1	46.9	49.23	56.26	31.81	32.70	33.88
IW (Pan et al. 2019)	21.8	47.6	62.4	52.4	46.1	48.49	55.82	29.91	27.48	29.71
ISW (Choi et al. 2021)	24.3	49.8	64.3	56.0	48.6	50.73	58.64	36.58	35.20	40.33
DIRL (Xu et al. 2022)	-	-	-	-	-	51.80	-	41.04	39.15	41.60
<i>Transformer-based:</i>										
HGFormer (Ding et al. 2023)	52.7	68.6	69.9	72.0	65.8	53.40	66.90	-	-	-
CMFormer (Bi, You, and Gevers 2024)	33.7	64.3	77.8	67.6	60.9	59.27	71.10	55.31	49.91	60.09
<i>VFM-based:</i>										
Rein (Wei et al. 2024)	55.9	70.6	79.5	72.5	69.6	63.54	74.03	66.40	60.40	66.10
SET (Yi et al. 2024)	<u>57.3</u>	<u>73.6</u>	80.1	74.8	<u>71.5</u>	65.07	75.67	<u>68.06</u>	61.64	67.68
FADA (Bi et al. 2024)	57.4	<u>73.5</u>	<u>80.2</u>	<u>75.0</u>	<u>71.5</u>	<u>65.12</u>	<u>75.86</u>	68.23	61.94	<u>68.09</u>
Ours	56.2	75.4	81.3	75.2	72.0	66.28	76.05	66.22	<u>61.80</u>	68.21
<i>Comparison with Baseline (Rein):</i>										
	+0.3	+4.8	+1.8	+2.7	+2.4	+2.74	+2.02	-0.18	+1.40	+2.11

Table 1: DGSS performance of our proposed Causal-Tune. $G.$, $C.$, $B.$ and $M.$ denotes GTA5, Cityscapes, BDD100K and Mapillary datasets, respectively. ‘-’ indicates that the results are not reported in the original paper and no source code is available. Note that in order to be consistent with previous methods, only report one decimal official results under $C. \rightarrow$ ACDC setting. **Blod** and underline results represent best and second-best performance, respectively.

Experiments

Benchmarks and Training Setting

Datasets and Metric. We evaluate the performance of Causal-Tune on five driving-scene datasets that share 19 categories for DGSS. **Cityscapes** (Cordts et al. 2016) is a widely used dataset for autonomous driving, which contains 2975 training images and 500 validation images with a resolution of 2048×1024 . **GTA5** (Richter et al. 2016) provides 24966 images with 1914×1052 resolution obtained from the game engine. **ACDC** (Sakaridis, Dai, and Van Gool 2021) comprises 406 images for validation, which covers four adverse conditions, *i.e.*, night, snow, fog, and rain. **BDD100K** (Yu et al. 2018) is a large-scale dataset that consists of 70k training images and 10k validation images. **Mapillary** (Neuhold et al. 2017) has 18000 and 2000 images for training and validation, respectively. Following previous works (Wei et al. 2024; Yi et al. 2024), we adopt the mean Intersection of Union (mIoU) as the evaluation metric.

Implementation details. Following previous VFM-based fine-tuning methods on DGSS, we choose DINOv2 (Oquab et al. 2023) and Mask2Former (Cheng et al. 2022) as the default VFM and segmentation head. AdamW is used as the optimizer with a base learning rate of $1e-4$. We set $R_L = 0.2$ and $R_H = 0.7$ in Gaussian band-pass filter. The model is trained 40000 iterations with a batch size of 4 and all the resolution of input images is 512×512 . Our method is implemented based on *MMSegmentation* (Contributors 2020) with a single RTX3090 GPU, occupying 14 GB memory.

Comparison with SOTA Methods

We use Rein (Wei et al. 2024) as baseline and following by Rein, we compare our method with existing DGSS on

three setups: i) Cityscapes ($C.$) \rightarrow ACDC, ii) Cityscapes ($C.$) \rightarrow BDD100K ($B.$), Mapillary ($M.$) and iii) GTA5 ($G.$) \rightarrow Cityscapes ($C.$), BDD100K ($B.$), Mapillary ($M.$). In addition, we also compare with three types of DGSS methods: 1) ResNet-based, 2) Transformer-based, and 3) VFM-based methods, to verify the effectiveness of our Causal-Tune. All of our experiments are implemented and averaged by three independent repetitions.

As shown in Table 1, Causal-Tune outperforms all Resnet-based and Transformer-based methods. In addition, by elaborately designing the causal factors tuning, our method achieves the best results in 7 out of 10 cases and the second-best results in 1 case compared to VFM-based methods. The detailed results are as follows:

i) Cityscapes ($C.$) \rightarrow ACDC. As illustrated in the second column of Table 1, our method has an obvious advantage in performance owing to its Causal & Non-causal Factors Filter and Causal Factor Tune strategy. It achieves the best results on ‘Snow’, ‘Fog’, and ‘Rain’ weather conditions, outperforming SET (Yi et al. 2024) and FADA (Bi et al. 2024) by (1.8% & 1.9%), (1.2% & 1.1%), and (0.4% & 0.2%) respectively. Furthermore, compared to baseline method Rein (Wei et al. 2024) (in the last row of Table 1), Causal-Tune brings a 2.4% improvement (72.0% vs. 69.6%) on average, where the most notable improvement is on ‘Snow’ condition (from 70.6% to 75.4%, +4.8%). Moreover, our method achieves a large margin improvement (from 72.5% to 75.2%, +2.7%) on ‘Rain’ condition, an impressive improvement of +1.8% is obtained on ‘Fog’ condition, and a satisfactory result on ‘Night’ condition (56.2% vs. 55.9%). The above comparison clearly demonstrates the effectiveness of our proposed method on adverse weather conditions.

Trained on $C.$ \rightarrow ACDC					
Method	Night	Snow	Fog	Rain	Avg.
FFT	57.1	71.1	80.6	69.2	<u>69.5</u>
HWT	52.8	<u>72.2</u>	79.4	<u>72.4</u>	69.2
DCT (Ours)	<u>56.2</u>	75.4	81.3	75.2	72.0

Trained on $C.$			Trained on $G.$		
Method	$\rightarrow B.$	$\rightarrow M.$	$\rightarrow C.$	$\rightarrow B.$	$\rightarrow M.$
FFT	65.07	71.00	65.22	60.56	64.77
HWT	64.43	<u>71.52</u>	<u>65.77</u>	<u>61.08</u>	<u>66.66</u>
DCT (Ours)	66.28	76.05	66.22	61.80	68.21

Table 2: Analysis of different frequency domain transformation methods.

Frequency Filter			Trained on $C.$				
LF	HF	$L\&HF$	Night	Snow	Fog	Rain	Avg.
-	-	-	55.9	70.6	79.5	72.5	69.6
✓	-	-	57.3	72.1	73.2	71.1	68.4
-	✓	-	54.8	<u>73.7</u>	<u>80.5</u>	<u>74.1</u>	<u>70.4</u>
-	-	✓	<u>56.2</u>	75.4	81.3	75.2	72.0

Table 3: Ablation studies on different filtering methods for DCT frequency spectrum. LF is low-frequency filtering, HF is high-frequency filtering and $L\&HF$ denotes both low and high-frequency filtering, *i.e.*, a band-pass filter.

ii) **Cityscapes ($C.$) \rightarrow BDD100K ($B.$), Mapillary ($M.$).** As illustrated in the third column of Table 1, Causal-Tune achieves a new SOTA performance under $C. \rightarrow B.$ (66.28%) and $C. \rightarrow M.$ (76.05%) setting, respectively.

iii) **GTA5 ($G.$) \rightarrow Cityscapes ($C.$), BDD100K ($B.$), Mapillary ($M.$).** As shown in the fourth column of Table 1, compared with the SOTA method FADA (Bi et al. 2024), our method obtains the best result on $G. \rightarrow M.$ setting (68.21%). Moreover, our method can still achieve competitive results under $G. \rightarrow B.$ (61.80% vs. 61.94%) setting.

However, our method achieves a worse performance under $G. \rightarrow C.$ setting. We think the reason is that GTA5 (Richter et al. 2016) is a synthetic dataset while Cityscapes (Cordts et al. 2016) is obtained from real scenes. Our method may not be able to mine causal factors effectively under *synthetic-to-real* generalization.

In summary, our proposed method not only obtains an excellent performance in *real-to-real* setting ($C. \rightarrow$ ACDC and $C. \rightarrow B., M.$) but also maintains a satisfactory performance in *synthetic-to-real* setting ($G. \rightarrow C., B., M.$).

Ablation Studies and Analysis

Analysis of different frequency domain transformation methods. As mentioned earlier, the frequency spectrum derived from the DCT has been shown to better separate causal and non-causal factors compared to methods based on FFT and HWT. In this section, we explore the effects of different frequency domain transformation methods on Causal-

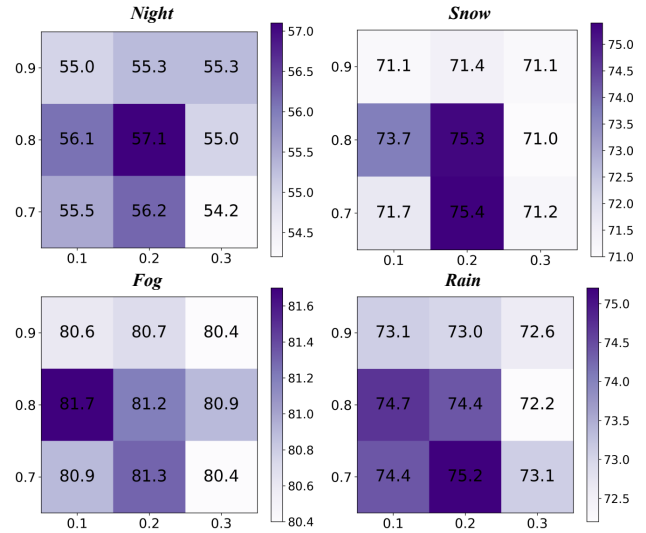


Figure 4: Accuracy matrix of cutoff frequency analysis under $C. \rightarrow$ ACDC generalization. The horizontal and vertical axis are low- and high-cutoff frequency, respectively.

Tune. From Table 2 we can observe that HWT cannot reveal causal and non-causal factors well, while FFT is simply better at removing the non-causal factor of ‘Night’. However, our proposed method outperforms FFT and HWT on most of the DGSS setups, which demonstrates that Causal-Tune could mine causal factors and remove non-causal factors better by utilizing DCT.

Different filtering method on DCT frequency spectrum.

To further verify the effectiveness of the DCT in our method, we conduct an ablation study on different filtering methods for the DCT frequency spectrum. Specifically, compared to baseline (Wei et al. 2024) (the first row in Table 3), when we only remove the low-frequency components (LF) in the DCT spectrum, the performance improves under ‘Night’ and ‘Snow’ condition, while it degrades under other conditions (the second row in Table 3). Moreover, from the third row of Table 3, removing high-frequency (HF) could enhance the generalization of the model under ‘Snow’, ‘Fog’ and ‘Rain’ conditions. In our method, we remove both low- and high-frequency components in the DCT spectrum, *i.e.*, use a band-pass filter, we achieve the best average performance in four conditions (72.0% in the last row).

Based on this experiment, we can observe that non-causal factors like ‘Fog’ and ‘Rain’ are mostly concentrated in the high-frequency of the DCT spectrum, whereas non-causal factors of ‘Night’ are mostly concentrated in the low-frequency of the DCT spectrum. Moreover, ‘Snow’ exists in both low- and high-frequency and a band-pass filter could remove these non-factors more effectively.

Analysis of the cutoff frequency parameters for the band-pass filter.

To verify the impact of the band-pass filter cutoff frequencies in our method on the DGSS task, we conduct experiments under $C. \rightarrow$ ACDC setting with different cutoff frequencies and compute the corresponding ac-




Figure 5: of feature map from different VFMs. We show the feature maps of frozen DINOv2 (Oquab et al. 2023), EVA02 (Fang et al. 2024), CLIP (Radford et al. 2021) and after fine-tuning by our method.




Figure 6: Visualization segmentation results under $C. \rightarrow ACDC$ generalization. Causal-Tune is compared with ISW (Choi et al. 2021) (ResNet-based), HGFormer (Bi, You, and Gevers 2024) (Transformer-based) and FADA (Bi et al. 2024) (VFM-based).

curacy matrix. As shown in Figure 4, when low-cutoff frequency $R_L \leq 0.2$ and high-cutoff frequency $R_H \leq 0.8$, the model generalizes better under four adverse weather conditions. However, as both cutoff frequencies increase ($R_L > 0.2$, $R_H > 0.8$), the performance of the model degrades. Thus, we set $R_L = 0.2$, $R_H = 0.7$ as the default.

Visualization Results

Feature map visualization of different VFMs. As mentioned before, the artifacts in the feature of VFMs may be caused by implicit non-causal factors that are not fully eliminated during the fine-tuning. In Figure 5, we observe that the artifacts in EVA02 are significantly reduced after fine-tuning by our method. Moreover, the artifacts in DINOv2 and CLIP are also removed. This demonstrates that removing non-causal factors can reduce artifacts in VFMs and allow models to focus on domain-invariant causal factors.

Visualization of segmentation results. To further show the effectiveness of our Causal-Tune, we present some visualization results on $C. \rightarrow ACDC$ setting. As shown in Fig-

ure 6, our method shows better segmentation results not only the ResNet-based method ISW (Choi et al. 2021) and the Transformer-based method HGFormer (Bi, You, and Gevers 2024), but also VFM-based method FADA (Bi et al. 2024).

Conclusion

In this paper, we propose a novel VFM-based PEFT method named Causal-Tune for DGSS from a causal perspective. In order to explicitly mine the causal and non-causal factors of features within VFMs for DGSS, our method incorporates two key components: i) Causal & Non-causal Factors Filter to disentangle causal factors and non-causal factors within the frequency domain, and ii) Causal Factors Tune, which utilizes a set of causal-aware learnable tokens to refine the causal factors to enhance the ability of generalization. Extensive experiments conducted on various cross-domain tasks show the effectiveness of our method. However, the generalization of the model is somewhat sensitive to the high- and low-cutoff frequency of the band-pass filter, and we plan to design a dynamic cutoff frequency to allow the model to adapt to various conditions in the future.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62206077), the Inner Mongolia Natural Science Foundation for Distinguished Young Scholars (No. 2025JQ009), and the Inner Mongolia Talent Development Project for Outstanding Young Talents. Yaoyue Zheng acknowledges the China Scholarship Council (CSC) No.202406280387. This work is supported by Grant PID2022-143257NB-I00 funded by MICIU/AEI/10.13039/501100011033 and ERDF/EU, the Departament de Recerca i Universitats from Generalitat de Catalunya with reference 2021SGR01499, and the Generalitat de Catalunya CERCA Program.

References

- Ahmed, N.; Natarajan, T.; and Rao, K. R. 2006. Discrete cosine transform. *IEEE transactions on Computers*, 100(1): 90–93.
- Bi, Q.; Yi, J.; Zheng, H.; Zhan, H.; Huang, Y.; Ji, W.; Li, Y.; and Zheng, Y. 2024. Learning frequency-adapted vision foundation model for domain generalized semantic segmentation. *Advances in Neural Information Processing Systems*, 37: 94047–94072.
- Bi, Q.; You, S.; and Gevers, T. 2024. Learning content-enhanced mask transformer for domain generalized urban-scene segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 819–827.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girdhar, R. 2022. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1290–1299.
- Choi, S.; Jung, S.; Yun, H.; Kim, J. T.; Kim, S.; and Choo, J. 2021. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11580–11590.
- Contributors, M. 2020. MMSegmentation: OpenMMLab Semantic Segmentation Toolbox and Benchmark. <https://github.com/open-mmlab/mms Segmentation>.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3213–3223.
- Darcet, T.; Oquab, M.; Mairal, J.; and Bojanowski, P. 2023. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*.
- Ding, J.; Xue, N.; Xia, G.-S.; Schiele, B.; and Dai, D. 2023. Hgformer: Hierarchical grouping transformer for domain generalized semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15413–15423.
- Fang, Y.; Sun, Q.; Wang, X.; Huang, T.; Wang, X.; and Cao, Y. 2024. Eva-02: A visual representation for neon genesis. *Image and Vision Computing*, 149: 105171.
- Fang, Y.; Wang, W.; Xie, B.; Sun, Q.; Wu, L.; Wang, X.; Huang, T.; Wang, X.; and Cao, Y. 2023. Eva: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19358–19369.
- Han, Z.; Gao, C.; Liu, J.; Zhang, J.; and Zhang, S. Q. 2024. Parameter-efficient fine-tuning for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huang, L.; Zhou, Y.; Zhu, F.; Liu, L.; and Shao, L. 2019. Iterative normalization: Beyond standardization towards efficient whitening. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4874–4883.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4015–4026.
- Liu, Y.; Zhou, S.; Liu, X.; Hao, C.; Fan, B.; and Tian, J. 2024. Unbiased faster r-cnn for single-source domain generalized object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 28838–28847.
- Lv, F.; Liang, J.; Li, S.; Zang, B.; Liu, C. H.; Wang, Z.; and Liu, D. 2022. Causality inspired representation learning for domain generalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8046–8056.
- Ma, N.; Zhang, X.; Zheng, H.-T.; and Sun, J. 2018. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European conference on computer vision (ECCV)*, 116–131.
- Neuhold, G.; Ollmann, T.; Rota Bulò, S.; and Kotschieder, P. 2017. The mapillary vistas dataset for semantic understanding of street scenes. In *Proceedings of the IEEE international conference on computer vision*, 4990–4999.
- Nussbaumer, H. J. 1981. The fast Fourier transform. In *Fast Fourier transform and convolution algorithms*, 80–111. Springer.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. DINOv2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Pak, B.; Woo, B.; Kim, S.; Kim, D.-h.; and Kim, H. 2024. Textual query-driven mask transformer for domain generalized segmentation. In *European Conference on Computer Vision*, 37–54. Springer.
- Pan, X.; Luo, P.; Shi, J.; and Tang, X. 2018. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Proceedings of the european conference on computer vision (ECCV)*, 464–479.

- Pan, X.; Zhan, X.; Shi, J.; Tang, X.; and Luo, P. 2019. Switchable whitening for deep representation learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1863–1871.
- Porwik, P.; and Lisowska, A. 2004. The Haar-wavelet transform in digital image processing: its status and achievements. *Machine graphics and vision*, 13(1/2): 79–98.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Richter, S. R.; Vineet, V.; Roth, S.; and Koltun, V. 2016. Playing for data: Ground truth from computer games. In *European conference on computer vision*, 102–118. Springer.
- Sakaridis, C.; Dai, D.; and Van Gool, L. 2021. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10765–10775.
- Wang, J.; Chen, M.; Karaev, N.; Vedaldi, A.; Rupprecht, C.; and Novotny, D. 2025. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5294–5306.
- Wei, Z.; Chen, L.; Jin, Y.; Ma, X.; Liu, T.; Ling, P.; Wang, B.; Chen, H.; and Zheng, J. 2024. Stronger fewer & superior: Harnessing vision foundation models for domain generalized semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 28619–28630.
- Xu, M.; Qin, L.; Chen, W.; Pu, S.; and Zhang, L. 2023. Multi-view adversarial discriminator: Mine the non-causal factors for object detection in unseen domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8103–8112.
- Xu, Q.; Yao, L.; Jiang, Z.; Jiang, G.; Chu, W.; Han, W.; Zhang, W.; Wang, C.; and Tai, Y. 2022. DirI: Domain-invariant representation learning for generalizable semantic segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 2884–2892.
- Xu, Q.; Zhang, R.; Zhang, Y.; Wang, Y.; and Tian, Q. 2021. A fourier-based framework for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14383–14392.
- Yang, Y.; and Soatto, S. 2020. Fda: Fourier domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4085–4095.
- Yi, J.; Bi, Q.; Zheng, H.; Zhan, H.; Ji, W.; Huang, Y.; Li, Y.; and Zheng, Y. 2024. Learning spectral-decomposed tokens for domain generalized semantic segmentation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 8159–8168.
- Yu, F.; Xian, W.; Chen, Y.; Liu, F.; Liao, M.; Madhavan, V.; Darrell, T.; et al. 2018. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2(5): 6.
- Zhang, H.; Xiao, L.; Cao, X.; and Foroosh, H. 2022. Multiple adverse weather conditions adaptation for object detection via causal intervention. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(3): 1742–1756.
- Zhang, T.; Liu, P.; Lu, Y.; Cai, M.; Zhang, Z.; Zhang, Z.; and Zhou, Q. 2025. CWNNet: Causal Wavelet Network for Low-Light Image Enhancement. *arXiv preprint arXiv:2507.10689*.
- Zhang, X.; and Tan, R. T. 2025. Mamba as a Bridge: Where Vision Foundation Models Meet Vision Language Models for Domain-Generalized Semantic Segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 14527–14537.
- Zhao, D.; Li, J.; Wang, S.; Wu, M.; Zang, Q.; Sebe, N.; and Zhong, Z. 2025. FisherTune: Fisher-Guided Robust Tuning of Vision Foundation Models for Domain Generalized Segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15043–15054.
- Zheng, Y.; Zhang, X.; Tian, Z.; and Du, S. 2025. Enhancing few-shot lifelong learning through fusion of cross-domain knowledge. *Information Fusion*, 115: 102730.