

Generalizing Vision-Language Models with Dedicated Prompt Guidance

Xinyao Li¹, Yinjie Min², Hongbo Chen¹, Zhekai Du¹, Fengling Li³, Jingjing Li^{1*}

¹School of Computer Science and Engineering, University of Electronic Science and Technology of China

²School of Statistics and Data Science, Nankai University

³University of Technology Sydney

xinyao326@outlook.com, nk.yjmin@gmail.com, jayus8643@gmail.com, zhekaid@uestc.edu.cn, fenglingli2023@gmail.com, lijn117@yeah.net

Abstract

Fine-tuning large pretrained vision-language models (VLMs) has emerged as a prevalent paradigm for downstream adaptation, yet it faces a critical trade-off between domain specificity and domain generalization (DG) ability. Current methods typically fine-tune a universal model on the entire dataset, which potentially compromises the ability to generalize to unseen domains. To fill this gap, we provide a theoretical understanding of the generalization ability for VLM fine-tuning, which reveals that training multiple parameter-efficient expert models on partitioned source domains leads to better generalization than fine-tuning a universal model. Inspired by this finding, we propose a two-step domain-expert-Guided DG (GuiDG) framework. GuiDG first employs prompt tuning to obtain source *domain experts*, then introduces a Cross-Modal Attention module to guide the fine-tuning of the vision encoder via adaptive expert integration. To better evaluate few-shot DG, we construct ImageNet-DG from ImageNet and its variants. Extensive experiments on standard DG benchmarks and ImageNet-DG demonstrate that GuiDG improves upon state-of-the-art fine-tuning methods while maintaining efficiency.

Code, Appendix — <https://github.com/TL-UESTC/GuiDG>

Introduction

Domain Generalization (DG) (Zhou et al. 2022a) aims to learn general knowledge from multiple source domains that is applicable to unseen target distributions. While traditional approaches focus on extracting domain-invariant features (Li et al. 2018b), the emergence of general-purpose vision-language models (VLMs) like CLIP (Radford et al. 2021) has fundamentally changed this landscape. Thanks to extensive pretraining, these models demonstrate remarkable zero-shot generalization capability to novel objects and domains.

However, adapting CLIP to downstream tasks presents a challenge: **endowing the model with domain-specific knowledge while preserving its zero-shot generalization ability** (Wortsman et al. 2022; Lai et al. 2023). Current methods tackle this by balancing specialization and generalization during fine-tuning. One line of work employs weight ensemble techniques, combining pretrained and fine-tuned

ID Acc	1. Fine-tune	Methods	2. Generalize	OOD Acc
66.7% ↓↓	✗	CLIP zero-shot	✗	57.2% ↓↓
70.9% ↑	✓	ERM fine-tune	✗	59.4% ↓
72.6% ↑↑	✓	ERM + GuiDG	✓	61.0% ↑↑

Figure 1: Illustration of the specialization - generalization balance. ERM fine-tuning fits to source knowledge at the cost of generalization ability, while our GuiDG achieves consistent improvements on both seen and unseen domains.

model weights (Wortsman et al. 2022) to mitigate over-fitting. Another explores careful training paradigms to prevent catastrophic forgetting of pretrained knowledge (Lai et al. 2023).

Despite their effectiveness, these approaches are limited by their design of sharing a universal model between source and target data. Such claim is supported by real-world examples (sampled from the ImageNet) in Figure 1, where the model is overly fitted to limited source data but lacks adaptability to unknown distributions. A holistic source model cannot guarantee consistent performance across potential target domains. The source fine-tuning may even harm the zero-shot ability in pretrained CLIP. To address such challenge, we first derive a novel upper bound for DG risks that reveals two key insights. (1) Training with partitioned source data and reduced hypothesis space achieves lower generalization risks in fine-tuning. This motivates us to train multiple parameter-efficient prompts to serve as source *domain experts* rather than fine-tuning a universal model. (2) An ensemble of dedicated source models provides more robust generalization than a universal model. Proper combination of experts can dynamically handle unseen target distributions.

Based on the theoretical insights, we design a two-step framework termed domain-expert-Guided DG (GuiDG). In Step 1, we leverage prompt tuning (Zhou et al. 2022c) to learn domain experts dedicated to model each source domain. Prompt tuning updates less than 1% of the total parameters in VLMs, providing a natural way to reduce the hypothesis space while capturing domain-specific knowledge. As shown in the left of Figure 1, GuiDG discriminates better on source data than naive fine-tuning, which aligns with the benefits of

*Corresponding author.

reduced hypothesis space. Step 2 focuses on the combination of the domain experts. We introduce a lightweight Cross-Modal Attention (CMAtn) module that generates weights to determine the contribution of experts. By assigning larger weights to more compatible experts, CMAtn guides the vision encoder to learn better representations, which in turn enables CMAtn to learn better weighting strategies.

Only domain experts are learnable in Step 1. In Step 2, the experts are frozen while the vision encoder and CMAtn module are jointly optimized. Such design maintains computation efficiency while promoting cross-modal information exchange. As illustrated in the right part of Figure 1, GuiDG generalizes well to unseen target domain while retaining zero-shot ability in CLIP. The contributions of this work include:

- We conduct theoretical analysis for fine-tuning VLM and derive a novel upper bound for generalization risks. We reveal that, contrary to end-to-end fine-tuning, properly trained and aggregated domain-specific models can achieve better generalization than a single universal model.
- Guided by our theoretical findings, we propose GuiDG, a two-step framework that first learns parameter-efficient domain experts on partitioned source data, then employs a Cross-Modal Attention module to adaptively integrate these experts during VLM fine-tuning.
- We develop ImageNet-DG, a DG benchmark derived from ImageNet and its variants to evaluate few-shot DG. Extensive experiments demonstrate the consistent performance gains and parameter efficiency of GuiDG.

Related Work

Vision-Language Models (VLMs) (Radford et al. 2021) are derived from web-scale image-text pairs with contrastive learning. VLMs generally feature a vision and text encoder to handle image and text inputs. By comparing vision and text representations, VLMs can make robust zero-shot inference (Li et al. 2025c). Efforts have been made to adapt VLMs to downstream applications. Prompt-tuning methods (Zhou et al. 2022c; Khattak et al. 2023a) learn prompt embeddings to achieve parameter-efficient adaptation in a few-shot style. As a more practical scenario, some methods propose to distill the pretrained VLM to a smaller model for client-side deployment or fine-tuning (Addepalli et al. 2024; Li et al. 2024b). VLMs’ strong zero-shot ability has attracted researches on transfer learning tasks (Li et al. 2025b). Some works choose to integrate domain information in prompts (Ge et al. 2023; Du et al. 2024), while others refine the representation space to match target distribution (Li et al. 2025a, 2024a).

Domain Generalization (DG) aims to learn general knowledge from multiple source domains generalizes to unseen domains. Adversarial-based methods (Li et al. 2018c; Deng et al. 2020) extract domain-invariant features from source domains via a min-max game between a feature extractor and domain discriminator. Augment-based methods refine source images to improve model generalization ability (Islam et al. 2024; Zhao et al. 2024). Zhou et al. zhou2021domain, zhou2024mixstyle mix-up images to synthesize novel domains that enhance model generalization. Inspired by meta-learning, some methods try to close domain

shift by splitting source data into meta-train and meta-test subsets (Li et al. 2018a; Khoee, Yu, and Feldt 2024). With recent advances in VLMs, Chen et al. (Chen et al. 2024) propose to solve hybrid DG tasks with perturbation distillation of VLMs. Cheng et al. (Cheng et al. 2024) utilize pretrained large language models to disentangle text prompts of VLMs for domain-agnostic visual features. Addepalli et al. (Addepalli et al. 2024) solve white- and black-box DG settings by aligning vision and text representations before distillation. While there are attempts on ensemble-based DG (Zhong et al. 2022; Bai et al. 2024), we are the first to theoretically support the benefits of such design, and propose a grounded GuiDG framework to reveal the generalization abilities of properly integrated domain-specific knowledge.

Robust Fine-Tuning focuses on preserving the generalizability of pretrained models while incorporating task-specific knowledge. For VLMs, prompt-tuning-based methods preserve base knowledge and learn new information (Zhang et al. 2024), or learn instance-conditioned prompts to prevent over-fitting (Zhou et al. 2022b). Unsupervised fine-tuning methods (Liang et al. 2024; Tanwisuth et al. 2023) are based on zero-shot inference results and apply regularization terms to prevent forgetting. On the fine-tuning of large backbone networks, weight ensemble (Wortsman et al. 2022; Cha et al. 2021) methods observe that mixing-up several weights in model optimization trajectory improves generalization performance. Adjusting learning rate also proves critical to preventing catastrophic forgetting (Wortsman et al. 2022; Lai et al. 2023). MIRO (Cha et al. 2022) proposes to constraint newly learned feature representations by oracle representations. This paper leverages domain specifics for generalization, which is compatible with existing fine-tuning methods.

Method

Preliminaries

Problem definition. This work investigates domain generalizable fine-tuning for C -class classification. We have labeled source data partitioned into source domains $D^S = \{D_i^S\}_{i=1}^d$, where d is the number of source domains and $D_i^S = \{(x_j^i, y_j^i)\}_{j=1}^{n_i^S}$ is the i th source domain with n_i^S labeled samples. The unseen target domain is denoted as $D^T = \{x_i^t\}_{i=1}^{n_t}$. The goal is to train a function f^S on D^S that minimizes prediction error on any unseen D^T .

Preliminaries on CLIP. We investigate the generalizable fine-tuning of CLIP (Radford et al. 2021). CLIP features a vision encoder E_v that extract vision representations from input images x : $I = E_v(x)$. For each class, one can use a general description t_c , e.g., A photo of a [CLASS $_c$], for zero-shot classification task, where CLASS $_c$ is the category name of the c th possible class. The text encoder E_t takes the sentence t_c as the input then generates text representations by $T_c = E_t(t_c)$. By computing cosine similarities ($\cos \langle \cdot, \cdot \rangle$) between the vision representation of image x and text representations $\{T_c\}_{c=1}^C$ of all classes, we can obtain the probability that x belongs to class c by:

$$P(y = c | x) = \frac{\exp(\cos \langle I, T_c \rangle / \tau)}{\sum_{i=1}^C \exp(\cos \langle I, T_i \rangle / \tau)}, \quad (1)$$

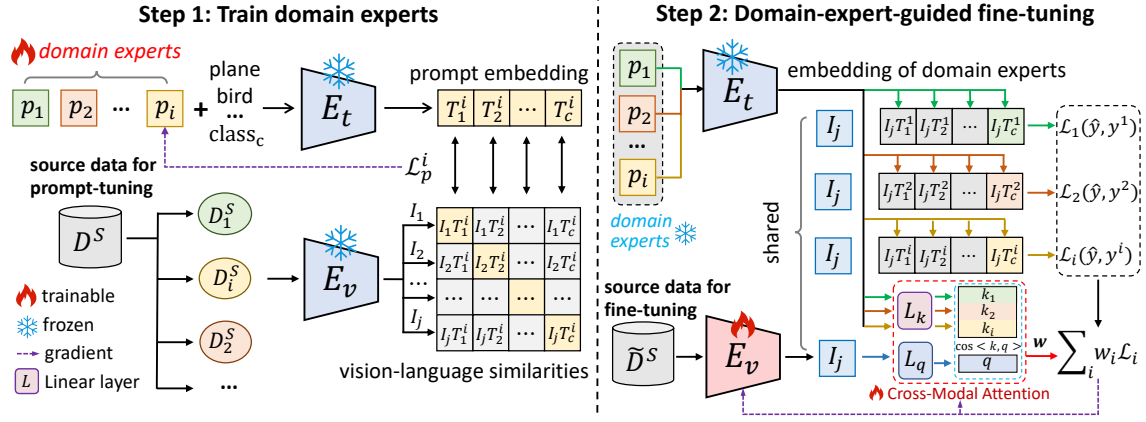


Figure 2: The two-step GuiDG framework. In Step 1, we split source data according to their domain characteristics. On each domain, a domain expert is learned with off-the-shelf prompt tuning methods. In Step 2, all domain experts are frozen. A Cross-Modal Attention (CMA) module decides ensemble weights from vision and text representations. These weights aggregate the knowledge in domain experts to guide the fine-tuning of the vision encoder, and assemble predictions for inference.

where τ is the temperature hyperparameter.

Theoretical Formulation

Assume that each source domain follows a distribution $P_i = P_{X,i} \times P_{Y|X}$, where $P_{X,i}$ represents the marginal distribution of features and $P_{Y|X}$ denotes the conditional distribution of labels given features. The entire source domains follow a mixture distribution $P = P_X \times P_{Y|X}$, where the marginal feature distribution P_X is defined as $P_X = \sum_{i=1}^d \pi_i P_{X,i}$. Similarly, the target domain follows distribution $P' = P'_X \times P_{Y|X}$ with $P'_X = \sum_{i=1}^d \pi'_i P_{X,i}$. Here $\pi_i, \pi'_i \geq 0$ holds for all $i = 1, \dots, d$, and satisfies $\sum_{i=1}^d \pi_i = \sum_{i=1}^d \pi'_i = 1$. Let $n = \sum_{i=1}^d n_i^S$ be the total number of source samples. We assume that $n_i^S = \pi_i n$ for all $i = 1, \dots, d$ (Albuquerque et al. 2019). For any hypothesis $h \in \mathcal{H}$ and distribution P_i , we define its risk as $\mathcal{E}_i(h) = \mathbb{E}_{x,y \sim P_i} \mathcal{L}(y, h(x))$, where $\mathcal{L}(\cdot, \cdot)$ is a loss function bounded by c_L . Our goal is to find $h \in \mathcal{H}$ that minimizes the risk $\mathcal{E}'(h) = \mathbb{E}_{x,y \sim P'} \mathcal{L}(y, h(x))$ on the target distribution P' . The classical approach seeks to train a universal predictor f by empirical risk minimization (ERM) (Gulrajani and Lopez-Paz 2020) across all domains:

$$\hat{f} = \arg \min_{f \in \mathcal{H}} \sum_{i=1}^d \pi_i \hat{\mathcal{E}}_i(f), \quad (2)$$

where $\hat{\mathcal{E}}_i(h) = \frac{1}{n_i^S} \sum_{j=1}^{n_i^S} \mathcal{L}(y_j^i, h(x_j^i))$ denotes the empirical risk on domain i . Instead, we propose a two-stage approach. First, for each domain i with hypothesis space \mathcal{H}_i , we find a domain-specific predictor $\hat{f}_i = \arg \min_{f \in \mathcal{H}_i} \hat{\mathcal{E}}_i(f)$. Then, we aggregate these predictors $\{\hat{f}_i\}_{i=1}^d$ using algorithm \mathcal{A} and an additional independent dataset $\tilde{D}^S = \{\tilde{D}_i^S\}_{i=1}^d$, where $\tilde{D}_i^S = \{(\tilde{x}_j^i, \tilde{y}_j^i)\}_{j=1}^{\tilde{n}_i^S}$ is also independently drawn from P_i . Let $m = \sum_{i=1}^d \tilde{n}_i^S$ denote the total number of samples in \tilde{D}^S with $\tilde{n}_i^S = \pi_i m$. The aggregated predictor

$\tilde{f} = \mathcal{A}(\hat{f}_1, \dots, \hat{f}_d; \tilde{D}^S)$ belongs to hypothesis space $\tilde{\mathcal{H}}$. For comparison, we redefine \hat{f} as the minimizer in \mathcal{H} of the empirical loss on $D^S \cup \tilde{D}^S$.

Theorem 1. Assume hypothesis space \mathcal{H} , $\tilde{\mathcal{H}}$ and \mathcal{H}_i have VC-dimension d_0 , \tilde{d} and d_i respectively. There exists constant $C > 0$, such that for any $\delta \in (0, 1)$ with probability at least $1 - 3d\delta$, the following inequality hold:

$$\begin{aligned} \mathcal{E}'(\tilde{f}) - \sum_{i=1}^d \pi'_i \hat{\mathcal{E}}_i(\hat{f}_i) &\leq \left(\sum_{i=1}^d \pi'_i / \sqrt{\pi_i} \right) \sqrt{\frac{c_L \log(1/\delta)}{2m}} \\ &+ C \sqrt{\frac{\tilde{d} \log(m) + \log(1/\delta)}{m}} \\ &+ C \sum_{i=1}^d \pi'_i \sqrt{\frac{d_i \log(n_i^S) + \log(1/\delta)}{n_i^S}}, \end{aligned} \quad (3)$$

and denote $N = n + m$, with probability at least $1 - d\delta$ the following inequality hold:

$$\mathcal{E}'(\hat{f}) - \sum_{i=1}^d \pi'_i \hat{\mathcal{E}}_i(\hat{f}_i) \leq C \sqrt{\frac{d_0 \log(N) + \log(1/\delta)}{N}}. \quad (4)$$

Denote the right-hand side of Eq. (3) and Eq. (4) as $\text{Upp}(\mathcal{E}'(\tilde{f}), \delta)$ and $\text{Upp}(\mathcal{E}'(\hat{f}), \delta)$, we have corollary:

Corollary 2. Assume $m = n$, $c_\pi = \sum_{i=1}^d \pi'_i / \sqrt{\pi_i}$ and $\sum_{i=1}^d \pi'_i \sqrt{2d_i} / \sqrt{\pi_i} \leq c(\delta) \sqrt{d_0}$, where for specified $\delta \in (0, 1)$, $c(\delta) = \inf_{1 \leq i \leq d} \sqrt{\frac{\log(2n) + (1/d_0) \log(1/\delta)}{\log(n_i^S) + (1/d_i) \log(3/\delta)}}$. We have

$$\text{Upp}(\mathcal{E}'(\tilde{f}), \delta/3) \leq \text{Upp}(\mathcal{E}'(\hat{f}), \delta) + \varepsilon, \quad (5)$$

where

$$\varepsilon = c_\pi \sqrt{\frac{c_L \log(3/\delta)}{N}} + C \sqrt{\frac{2\tilde{d} \log(N/2) + 2 \log(3/\delta)}{N}}.$$

Remark 3. When \mathcal{H} is a parameterized neural network space, assume the number of parameters as $n(\mathcal{H})$. The VC-dimension of \mathcal{H} is approximately $n(\mathcal{H}) \log\{n(\mathcal{H})\}$ (Bartlett and Maass 2003). When $\tilde{d} \ll d_0$, which is equivalent to $n(\tilde{\mathcal{H}}) \ll n(\mathcal{H})$, ε in inequality (5) is a very small item compared with $\text{Upp}(\mathcal{E}'(\hat{f}), \delta)$ as d_0 is usually large. Corollary 2 suggests that we should seek ways to make $\sum_{i=1}^d \pi'_i \sqrt{2n(\mathcal{H}_i) \log\{n(\mathcal{H}_i)\}} / \sqrt{\pi_i} < c(\delta) \sqrt{n(\mathcal{H}) \log\{n(\mathcal{H})\}}$, thereby ensuring the upper bound of ensemble risk is much smaller than that of a universal model, i.e., $\text{Upp}(\mathcal{E}'(\tilde{f}), \delta/3) < \text{Upp}(\mathcal{E}'(\hat{f}), \delta)$.

With carefully constructed \mathcal{H}_i that satisfies condition in Corollary 2, when $\tilde{\mathcal{E}}_i(\hat{f}_i) \leq \tilde{\mathcal{E}}_i(\hat{f})$ (each \hat{f}_i is trained to minimize empirical risk of its domain with a much smaller hypothesis space, leading to lower empirical risk compared to the universal model \hat{f} that needs to compromise across all domains), $\mathcal{E}'(\tilde{f})$ has a tighter upper bound than $\mathcal{E}'(\hat{f})$ on P' . The proofs on Theorem 1 and Corollary 2, as well as an illustrative toy example on Remark 3 are in Appendix.

Learning Domain Experts

Current methods (Wortsman et al. 2022) fully fine-tune a universal model on all source data. However, Remark 3 indicates that an ensemble of parameter-efficient domain-specific models brings better generalization ability. Motivated by such finding, we propose to learn **dedicated** function \hat{f}_i from each carefully designed \mathcal{H}_i to ensure $\sum_{i=1}^d \pi'_i \sqrt{2\tilde{d}_i} / \sqrt{\pi_i} \leq c(\delta) \sqrt{\tilde{d}_0}$. More details are in Appendix. The learned experts incorporate more specific domain knowledge into the default text description, e.g., *A [real/art/...] photo of a [CLASS]*.

We learn the domain experts in a few-shot style (Zhou et al. 2022c), while any prompt-tuning method is viable. On domain D_i^S , we construct the learnable prompt for class j as:

$$t_j^i = [p_{i1}] [p_{i2}] \dots [p_{im}] [\text{CLASS}_j], \quad (6)$$

where $\mathbf{p}_i = [p_{i1}] [p_{i2}] \dots [p_{im}]$ is domain expert, m is length of expert, and $[\text{CLASS}_j]$ is embedding of class j . The text embedding of class j in domain i is obtained: $T_j^i = E_t(t_j^i)$. The probability of image x^i belonging to class j is obtained:

$$P_j(\hat{y} | x^i, t_{1:C}^i) = \frac{\exp(\cos \langle E_v(x^i), T_j^i \rangle / \tau)}{\sum_{c=1}^C \exp(\cos \langle E_v(x^i), T_c^i \rangle / \tau)}. \quad (7)$$

Standard cross-entropy loss is used to learn \mathbf{p}_i :

$$\mathcal{L}_p^i = - \sum_{j=1}^{n_i^S} \sum_{c=1}^C [y_j^i = c] \cdot \log P_c(\hat{y} | x_j^i, t_{1:C}^i), \quad (8)$$

where $[\cdot]$ is indicator function. Both E_v and E_t are frozen during the optimization of Eq. (8). The i th domain expert is obtained by optimizing Eq. (8): $\mathbf{p}_i = \arg \min_{\mathbf{p}_i} \mathcal{L}_p^i$.

Domain-Expert-Guided Fine-Tuning

With trained domain experts \mathbf{p}_i to serve as dedicated domain predictors \hat{f}_i , we proceed to learn algorithm \mathcal{A} that aggregates the domain knowledge in $\{\hat{f}_i\}_{i=1}^d$ to guide the fine-tuning of CLIP. As shown in Step 2 of Figure 2, we design a

Cross-Modal Attention (CMAtn) module to approximate the aggregation algorithm \mathcal{A} . CMAtn is composed of a linear layer L_q that transforms vision features into query embeddings: $q(x) = L_q(E_v(x))$, and a linear layer L_k to transform text features into key embeddings: $k_i = L_k(E_t(t_i))$. We then compute the normalized cosine similarities between a query embedding and key embeddings of all domain experts to obtain ensemble weights:

$$\mathbf{w}(x) = \text{Softmax}(\cos \langle q(x), [k_1, k_2, \dots, k_d] \rangle), \quad (9)$$

where $\mathbf{w}(x) = (w_1(x), w_2(x), \dots, w_d(x))$ is the weight vector, and $[\cdot]$ is the concatenation operation. CMAtn learns to assign larger weights to more compatible experts, and smaller weights to irrelevant experts. The training data in Step 2 include data from all available source domains, therefore CMAtn can learn various weighting scenarios that are generalizable to unseen target domains. For an image input, different classification results are obtained from each domain expert via Eq. (7). Instead of computing their weighted average before optimization, we propose to weight the training losses. We observe that such design brings more direct and effective optimization for both the image encoder and CMAtn module. As instructed in Theoretical Formulation, we train CMAtn on \tilde{D}^S with m samples that are *i.i.d* with D^S . Combining Eq. (7), the training loss for Step 2 is defined:

$$\mathcal{L}_f = \sum_{i=1}^d \sum_{j=1}^{\tilde{n}_i^S} w_i(x_j^i) \left(- \sum_{c=1}^C [y_j^i = c] \log P_c(\hat{y} | x_j^i, t_{1:C}^i) \right). \quad (10)$$

During the optimization of Eq. (10), only the CMAtn and the vision encoder are trainable.

Train and Inference

As shown in Figure 2, GuiDG consists of two training steps. In Step 1, we train d independent domain experts $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_d$ by optimizing Eq. (8). In Step 2, CMAtn and the vision encoder are trained by minimizing Eq. (10). Step 2 is compatible with existing regularization terms (detailed in Appendix) for robust fine-tuning, therefore we have:

$$\theta_{E_v}, \theta_{L_q}, \theta_{L_k} = \arg \min_{\theta_{E_v}, \theta_{L_q}, \theta_{L_k}} \mathcal{L}_f + \alpha \mathcal{L}_r, \quad (11)$$

where $\theta_{E_v}, \theta_{L_q}, \theta_{L_k}$ are parameters in the vision encoder and CMAtn, \mathcal{L}_r is off-the-shelf regularization loss for fine-tuning, and α controls the regularization effects.

During inference, the VLM generates d sets of logits from each domain expert given target data x^t . CMAtn assigns proper weights w_i for each output. Assume hidden variable $\mathcal{I}(x^t)$ indicating the index of \hat{f}_i that best predicts y^t . For class c , by conditioning on $\{\hat{f}_i\}_{i=1}^d$ we have $\mathbb{P}(y^t = c | x^t) = \sum_{i=1}^d \mathbb{P}(\mathcal{I}(x^t) = i) \mathbb{P}(y^t = c | x^t, \mathcal{I}(x^t) = i)$. In our design, weight $w_i(x^t)$ is an estimator of $\mathbb{P}(\mathcal{I}(x^t) = i)$ and $\cos \langle E_v(x^t), T_c^i \rangle / \tau$ is an estimator approximately proportion to $\mathbb{P}(y^t = c | x^t, \mathcal{I}(x^t) = i)$. Thus, the weighted average of outputs serve as the final inference results \hat{y}^t :

$$\hat{y}^t = \arg \max_c \sum_{i=1}^d w_i(x^t) \cdot \cos \langle E_v(x^t), T_c^i \rangle / \tau. \quad (12)$$

Method	OfficeHome					DomainNet						
	Art	Clp	Prod	RW	Avg.	clp	inf	pnt	qdr	rel	skt	Avg.
CLIP-zeroshot	82.9	67.8	89.0	89.8	82.4	70.1	46.4	61.7	13.7	82.9	62.6	56.2
Full data												
MIRO (Cha et al. 2022)	83.6	75.7	89.7	90.2	84.8	79.7	43.5	67.4	24.6	79.2	68.4	60.5
WiSE-FT (Wortsman et al. 2022)	85.2	76.2	92.9	91.0	86.3	76.8	49.5	69.4	20.1	81.7	67.2	60.8
VL2V-SD (Addepalli et al. 2024)	87.3	78.6	92.0	91.7	87.4	80.0	49.0	71.1	23.3	82.1	71.4	62.8
ERM (WF)*	84.9	71.2	92.4	92.0	85.1 \pm 0.2	74.5	49.5	69.6	16.1	84.5	66.7	60.2 \pm 0.1
ERM (WF) + GuiDG*	85.9	71.7	92.6	92.3	85.6 \pm 0.2	76.0	53.1	70.9	17.3	84.8	68.2	61.7 \pm 0.2
UEO (WF)* (Liang et al. 2024)	85.6	72.8	93.2	92.5	86.0 \pm 0.3	75.5	50.8	70.2	16.9	84.6	66.8	60.8 \pm 0.1
UEO (WF) + GuiDG*	86.8	73.6	92.9	93.4	86.7 \pm 0.3	76.5	52.9	71.1	17.8	85.0	68.9	62.0 \pm 0.2
CLIPood* (Shu et al. 2023)	87.8	73.8	92.7	92.9	86.8 \pm 0.2	77.6	54.6	72.7	20.8	85.2	69.7	63.4 \pm 0.1
CLIPood + GuiDG*	89.1	74.6	92.1	93.6	87.4 \pm 0.2	77.7	54.3	72.5	20.4	85.3	69.9	63.4 \pm 0.2
16-shot												
ERM (WF)*	86.0	69.5	92.3	93.0	85.2 \pm 0.3	73.9	49.8	68.3	16.3	84.6	65.6	59.8 \pm 0.2
ERM (WF) + GuiDG*	86.4	70.0	93.1	91.6	85.3 \pm 0.3	74.2	51.6	68.5	15.0	84.9	66.8	60.2 \pm 0.2
UEO (WF)* (Liang et al. 2024)	85.4	68.1	92.7	92.8	84.8 \pm 0.3	74.1	52.9	68.3	14.9	85.1	66.9	60.4 \pm 0.2
UEO (WF) + GuiDG*	87.2	71.1	93.2	92.6	86.0 \pm 0.2	75.4	52.9	69.9	17.6	85.4	68.3	61.6 \pm 0.2
CLIPood* (Shu et al. 2023)	86.2	71.7	92.7	93.2	86.0 \pm 0.3	77.3	53.0	70.9	19.3	85.1	68.2	62.3 \pm 0.1
CLIPood + GuiDG*	87.6	72.9	93.4	93.4	86.8 \pm 0.4	76.8	53.7	71.7	20.3	85.2	68.9	62.8 \pm 0.2
8-shot												
ERM (WF)*	82.9	65.3	90.5	92.2	82.7 \pm 0.4	71.5	46.3	66.2	13.9	83.8	63.9	57.6 \pm 0.2
ERM (WF) + GuiDG*	85.8	69.7	93.0	92.1	85.2 \pm 0.3	73.9	51.5	68.4	14.9	85.0	66.2	60.0 \pm 0.2
UEO (WF)* (Liang et al. 2024)	85.1	67.4	92.2	92.3	84.3 \pm 0.2	74.3	51.6	68.7	15.1	84.9	66.5	60.2 \pm 0.2
UEO (WF) + GuiDG*	86.2	72.1	92.3	92.7	85.8 \pm 0.3	76.1	52.7	70.2	17.9	85.3	67.9	61.7 \pm 0.3
CLIPood* (Shu et al. 2023)	86.8	69.9	92.7	93.9	85.8 \pm 0.3	76.9	51.7	70.5	18.9	84.7	68.0	61.8 \pm 0.2
CLIPood + GuiDG*	86.6	73.5	92.4	92.9	86.4 \pm 0.3	76.1	53.1	71.0	18.8	85.2	68.5	62.1 \pm 0.2

* Results based on our own runs.

Table 1: DG results of GuiDG. Best results are in bold. Most significant improvements by incorporating GuiDG are underlined.

Experiments

Setup

Benchmark. We conduct experiments on standard domain generalization benchmarks. All experiments are repeated 5 times with different seeds and means were reported. **Office-Home** (Venkateswara et al. 2017) includes 4 domains with 65 categories of office items. **VLCS** (Torralla and Efros 2011) includes 4 domains with 5 classes. **PACS** (Li et al. 2017) provides 4 art-style domains with 7 classes. **DomainNet** (Peng et al. 2019) contains 0.6 million samples from 6 domains and 345 categories. **TerraIncognita (TI)** (Beery, Van Horn, and Perona 2018) includes 10 classes of animal pictures taken in 4 different locations. We notice existing DG benchmarks include limited classes and instances, hindering sufficient evaluation of modern models like CLIP. Therefore, we sample data from ImageNet (Deng et al. 2009) and its variants (Hendrycks et al. 2021b,a; Wang et al. 2019; Recht et al. 2019) to construct a new subset **ImageNet-DG**. We also experiment on single-source DG, where the model is fine-tuned on ImageNet and generalizes to ImageNet variants.

Implementation details. We use CLIP ViT-B/16 (Radford et al. 2021) on all experiments. The temperature in Eq. (7), Eq. (12) are set to 0.01 as in original CLIP. The query transformer L_q in CMAtn is a linear layer of shape (d_f, d_f) where d_f is dimension of the feature representations in CLIP. L_k is a linear layer of shape $(C, 1)$ that transforms prompt embedding of C classes to 1. These linear layers in CMAtn introduce 1M parameters, ensuring our method’s parameter efficiency. On all datasets, we follow leave-one-out paradigm to test one unseen target domain with others as source do-

Method	PACS	VLCS	TI	Avg.
CLIP (Radford et al. 2021)	96.2	81.8	33.8	70.6
MIRO (Cha et al. 2022)	95.6	82.2	54.3	77.4
SWAD (Cha et al. 2021)	91.4	79.1	42.9	71.1
WiSE-FT (Wortsman et al. 2022)	97.3	82.9	54.5	78.2
RISE (Huang et al. 2023)	93.3	80.6	49.6	74.5
VL2V-SD (Addepalli et al. 2024)	96.7	83.3	58.5	79.5
DSPL (Cheng et al. 2024)	97.5	86.4	57.1	80.3
ERM (WF)*	96.7 \pm 0.4	83.3 \pm 0.2	51.0 \pm 0.4	77.0
ERM (WF) + GuiDG*	97.6 \pm 0.2	83.8 \pm 0.2	52.9 \pm 0.5	78.1
UEO (WF)* (Liang et al. 2024)	96.9 \pm 0.2	81.3 \pm 0.3	51.5 \pm 0.3	76.6
UEO (WF) + GuiDG*	97.6 \pm 0.2	84.2 \pm 0.2	52.3 \pm 0.4	78.0
CLIPood* (Shu et al. 2023)	97.3 \pm 0.1	84.2 \pm 0.2	60.0 \pm 0.6	80.5
CLIPood + GuiDG*	97.3 \pm 0.1	84.3 \pm 0.2	60.9 \pm 0.4	80.8

* Results based on our own runs.

Table 2: DG results on PACS, VLCS and TI. Best results are in bold. Most significant improvements are underlined.

mains each time. On VLCS, PACS and TerraIncognita with fewer classes, all data are used. On OfficeHome, DomainNet and ImageNet-DG we adopt few-shot fine-tuning protocol to evaluate the few-shot generalization ability. To satisfy the data independence requirements in section Theoretical Formulation, we randomly split each domain to construct disjoint D^S for phase 1 and \tilde{D}^S for phase 2. We adopt AdamW (Loshchilov 2017) with learning rate 5e-6 on all tasks. The regularization loss \mathcal{L}_r and α in Eq. (11) are set according to original baseline methods, with more details in Appendix.

Main Results

Standard benchmarks. Table 1 and Table 2 include results on the five standard DG benchmarks in DomainBed (Gulra-

Method	ImageNet-DG source = I, S, V2				Single-source ImageNet source = I					
	A	I	R	Avg.	A	I	R	S	V2	Avg.
CLIP-zeroshot (Radford et al. 2021)	47.0	66.7	73.9	62.6	47.8	66.7	74.0	46.1	60.8	59.1
16-shot										
PromptSRC [†] (Khattak et al. 2023b)	51.0	72.1	79.1	67.4	50.9	71.3	77.8	49.6	64.4	62.8
Apex [†] (Yang, Ko, and Yun 2023)	50.6	72.5	78.7	67.3	50.7	72.0	76.8	48.5	64.7	62.5
ERM (WF)* (Gulrajani and Lopez-Paz 2020)	48.9	70.7	78.9	66.2 \pm 0.1	49.1	70.9	75.8	48.4	64.1	61.7 \pm 0.1
ERM (WF) + GuiDG*	50.2	72.8	80.8	67.9 \pm 0.2	51.1	72.6	77.9	49.7	65.3	63.3 \pm 0.1
UEO (WF)* (Liang et al. 2024)	50.5	69.8	77.2	65.8 \pm 0.2	48.6	71.9	75.9	48.6	65.1	62.0 \pm 0.1
UEO (WF) + GuiDG*	50.8	72.9	80.7	68.1 \pm 0.2	51.0	73.5	78.1	50.1	66.2	63.8 \pm 0.2
CLIPood* (Shu et al. 2023)	49.8	71.8	81.1	67.6 \pm 0.1	50.4	71.6	77.2	49.3	64.9	62.7 \pm 0.1
CLIPood + GuiDG*	50.9	73.0	81.8	68.6 \pm 0.1	49.6	73.4	77.7	50.4	66.3	63.5 \pm 0.1
8-shot										
ERM (WF)* (Gulrajani and Lopez-Paz 2020)	49.0	69.8	77.4	65.4 \pm 0.1	47.9	70.3	75.9	48.2	63.3	61.1 \pm 0.1
ERM (WF) + GuiDG*	50.7	71.3	79.4	67.1 \pm 0.1	50.9	71.9	78.1	49.5	64.8	63.0 \pm 0.1
UEO (WF)* (Liang et al. 2024)	50.5	70.0	77.0	65.8 \pm 0.1	48.0	70.3	76.2	48.5	63.4	61.3 \pm 0.2
UEO (WF) + GuiDG*	50.4	71.7	80.1	67.4 \pm 0.1	50.5	72.8	78.0	49.8	65.6	63.3 \pm 0.2
CLIPood* (Shu et al. 2023)	50.0	71.6	80.5	67.4 \pm 0.2	45.5	71.7	75.2	47.9	64.8	61.0 \pm 0.1
CLIPood + GuiDG*	50.3	71.8	81.2	67.8 \pm 0.2	49.5	73.0	77.3	50.2	65.7	63.1 \pm 0.1

* Results based on our own runs. [†] Prompt-tuning baselines.

Table 3: DG results on ImageNet-DG and single-source DG on ImageNet and its variants. The ‘source’ row indicates the source domain(s) used for fine-tuning. Best results are in bold. Most significant improvements by incorporating GuiDG are underlined.

Step 1		Step 2		results
multiple experts	i.i.d data	no weights	learnable weights	Avg.
		✓		64.9
✓		✓		67.0
✓			✓	67.3
	✓	✓		66.8
✓	✓	✓		67.6
✓	✓		✓	67.9

Table 4: Ablation study on ImageNet-DG.

jani and Lopez-Paz 2020). On OfficeHome and DomainNet we additionally provide few-shot results and comparisons. On all tasks, we adopt the prompt-tuning procedure in (Zhou et al. 2022c) for training domain experts with $m = 16$. For regularization losses in Eq. (11), we experiment with the entropy loss in UEO (Liang et al. 2024), the Margin Metric Softmax loss in CLIPood (Shu et al. 2023), and ERM (Gulrajani and Lopez-Paz 2020) (without regularization). To achieve competitive results, we incorporate WiSE-FT (WF) (Wortsman et al. 2022) when implementing ERM and UEO. The Beta Moving Average in CLIPood achieves similar effects. We provide domain-wise DG results in Table 1 and average over all domains in Table 2. We can observe that GuiDG provides steady improvements over the baseline methods on all tasks, achieving new state-of-the-art. Specifically, in few-shot settings, the significant reduce in fine-tuning data harms the generalization effects of fine-tuning methods. GuiDG mitigates such performance drop adaptive expert integration.

Evaluation on ImageNet. Table 3 presents DG results on ImageNet and its variants. On ImageNet-DG, the source domains are ImageNet (train split), ImageNet-S and ImageNet-V2, and the target domains are ImageNet-A, ImageNet (eval-




Figure 3: Bars and lines are relative accuracies (average accuracy subtracted) and weights of domain experts. The right-most bar in each group shows the gains by prompt ensemble.

uation split) and ImageNet-R. We ensure that in each task, the label space among source and target domains are identical. We can observe significant boosts brought by GuiDG on all baseline methods. The 8-shot performance with GuiDG even surpasses 16-shot performance without GuiDG, supporting the efficacy and data-efficiency of leveraging multiple dedicated domain models. We also compare with prompt-tuning methods (Khattak et al. 2023b; Yang, Ko, and Yun 2023). The results show the superiority of our method on both ID (domain I) and OOD generalization tasks. To evaluate GuiDG without domain labels, we perform single-source DG on ImageNet by randomly splitting the training set of ImageNet into 4 pseudo-domains. In such case, the domain experts tend to be homogeneous, but GuiDG still performs better than competing baselines. More discussions are in Appendix.

Analytical Experiments

Ablation study. Table 4 presents ablation study of GuiDG. We evaluate the effectiveness of key designs in Step 1 and 2 on baseline method 16-shot ERM (WF). The results indicate

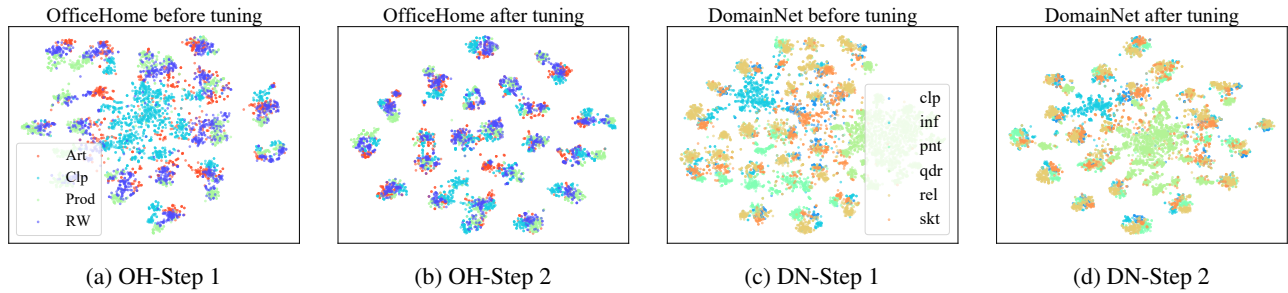
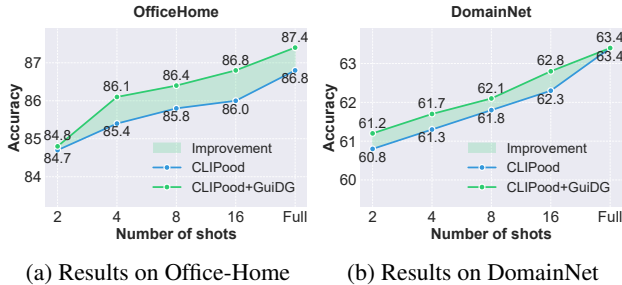


Figure 4: Vision features of source and target data before and after fine-tuning. (a),(b) Results obtained on domain ‘Art’ of OfficeHome. (c),(d) Results obtained on domain ‘clp’ of DomainNet.



Dataset	CMAttn	Experts	Ratio
OfficeHome (4)	1.050M	0.025M	0.7%
DomainNet (6)	1.052M	0.041M	0.7%

Table 5: Parameter analysis of GuiDG. In the parentheses are the number of domains. Ratio refers to the percentage of additional parameters by integrating GuiDG.

Figure 5: Few-shot results, averaged over all target domains.

each component contributes positively. The most significant performance drop emerges if only one prompt is trained instead of multiple domain experts (‘multiple experts’). The results support our design to ‘divide and conquer’ the large source domain by smaller expert models. CMAttn further guides the fine-tuning process in Step 2 by combining appropriate domain experts (learnable weights), achieving better generalization than simple averaging (no weights). By utilizing i.i.d data between Step 1 and 2, the requirements of our theory are satisfied and improvements are observed.

Weight analysis. Proper ensemble of domain experts in GuiDG ensures minimum generalization risks. Figure 3 investigates the compatibility between domain expert accuracies and their assigned weights, revealing the following insights. (1) CMAttn is generalizable to unknown domains. The assigned weights are reasonable and compatible with the target performance of experts without accessing target data. (2) The ensemble process always provides positive gains. As shown by the rightmost bars in each group, the performances by integrating multiple domain experts consistently *surpass* the best individual domain expert in the group (indicated by the shadowed parts in bars). (3) The weights take effect by reducing negative influences of experts. While it is hard to precisely match weights with every expert on unseen domains, CMAttn correctly assigns the lowest weights to the worst-performing experts in all cases to eliminate their drawbacks.

Feature visualization. We conduct t-SNE visualization (Van der Maaten and Hinton 2008) on vision features before and after the domain-expert-guided fine-tuning. As shown in Figure 4, before fine-tuning the features are chaotically distributed and cannot form compact locality structures, a

merit that good classification models possess (Li et al. 2019). After our fine-tuning process, the features are more distinguishable and form distinct class-wise feature groups. We observe that on the previously unseen domains, the model can still extract discriminative features. On the most challenging domain ‘qdr’ (Figure 4d), its features after tuning better align with other domains compared to Figure 4c.

Few-shot performances. To evaluate GuiDG with less fine-tuning data, we compare the performance of CLIPood and CLIPood+GuiDG under the setting of 2-, 4-, 8-, 16-shots and full data fine-tuning in Figure 5. The performances drop significantly with less training data, but we can still observe steady performance gains across all few-shot settings.

Parameter analysis. Table 5 analyzes additional parameters introduced by incorporating GuiDG. As the number of source domains increase, the additional parameters maintain at $\sim 1M$ in total. Such additional 1M parameters by incorporating GuiDG account for less than 1% of all tunable parameters in current baselines, which is reasonable.

Conclusion

This work investigates domain generalization of VLMs. Current methods train a universal model on all source domains for generalization, which is inevitably limited by the trade-off between model specificity and generalization ability. To address this, we show that ensemble of multiple smaller source expert models brings lower target risks while maintaining source specificity. Therefore, we design a domain-expert-guided DG framework that first learns prompt experts on source domains to encompass source knowledge. Secondly, a Cross-Modal Attention module is introduced to guide the tuning of VLMs with learnable weights. Experiments on standard DG benchmarks and a newly-proposed ImageNet-DG subset demonstrate the efficacy and efficiency of GuiDG.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 62572102, 52441801, and in part by the Fundamental Research Funds for the Central Universities (UESTC) under Grant ZYGX2024Z008.

References

- Addepalli, S.; Asokan, A. R.; Sharma, L.; and Babu, R. V. 2024. Leveraging Vision-Language Models for Improving Domain Generalization in Image Classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23922–23932.
- Albuquerque, I.; Monteiro, J.; Darvishi, M.; Falk, T. H.; and Mitliagkas, I. 2019. Generalizing to unseen domains via distribution matching. *arXiv preprint arXiv:1911.00804*.
- Bai, S.; Zhang, Y.; Zhou, W.; Luan, Z.; and Chen, B. 2024. Soft prompt generation for domain generalization. In *European Conference on Computer Vision*, 434–450. Springer.
- Bartlett, P. L.; and Maass, W. 2003. Vapnik–Chervonenkis dimension of neural nets. *The handbook of brain theory and neural networks*, 1188–1192.
- Beery, S.; Van Horn, G.; and Perona, P. 2018. Recognition in terra incognita. In *Proceedings of the European conference on computer vision (ECCV)*, 456–473.
- Cha, J.; Chun, S.; Lee, K.; Cho, H.-C.; Park, S.; Lee, Y.; and Park, S. 2021. Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems*, 34: 22405–22418.
- Cha, J.; Lee, K.; Park, S.; and Chun, S. 2022. Domain generalization by mutual-information regularization with pre-trained models. In *European conference on computer vision*, 440–457. Springer.
- Chen, Z.; Wang, W.; Zhao, Z.; Su, F.; Men, A.; and Meng, H. 2024. PracticalDG: Perturbation Distillation on Vision-Language Models for Hybrid Domain Generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23501–23511.
- Cheng, D.; Xu, Z.; Jiang, X.; Wang, N.; Li, D.; and Gao, X. 2024. Disentangled Prompt Representation for Domain Generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23595–23604.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Deng, Z.; Ding, F.; Dwork, C.; Hong, R.; Parmigiani, G.; Patil, P.; and Sur, P. 2020. Representation via representations: Domain generalization via adversarially learned invariant representations. *arXiv preprint arXiv:2006.11478*.
- Du, Z.; Li, X.; Li, F.; Lu, K.; Zhu, L.; and Li, J. 2024. Domain-agnostic mutual prompting for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23375–23384.
- Ge, C.; Huang, R.; Xie, M.; Lai, Z.; Song, S.; Li, S.; and Huang, G. 2023. Domain adaptation via prompt learning. *IEEE Transactions on Neural Networks and Learning Systems*.
- Gulrajani, I.; and Lopez-Paz, D. 2020. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; et al. 2021a. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8340–8349.
- Hendrycks, D.; Zhao, K.; Basart, S.; Steinhardt, J.; and Song, D. 2021b. Natural adversarial examples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15262–15271.
- Huang, Z.; Zhou, A.; Ling, Z.; Cai, M.; Wang, H.; and Lee, Y. J. 2023. A sentence speaks a thousand images: Domain generalization through distilling clip with language guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11685–11695.
- Islam, K.; Zaheer, M. Z.; Mahmood, A.; Nandakumar, K.; and Akhtar, N. 2024. Genmix: effective data augmentation with generative diffusion model image editing. *arXiv preprint arXiv:2412.02366*.
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2023a. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19113–19122.
- Khattak, M. U.; Wasim, S. T.; Naseer, M.; Khan, S.; Yang, M.-H.; and Khan, F. S. 2023b. Self-regulating prompts: Foundational model adaptation without forgetting. In *Proceedings of the IEEE/CVF international conference on computer vision*, 15190–15200.
- Khoee, A. G.; Yu, Y.; and Feldt, R. 2024. Domain generalization through meta-learning: a survey. *Artificial Intelligence Review*, 57(10): 285.
- Lai, Z.; Vesdapunt, N.; Zhou, N.; Wu, J.; Huynh, C. P.; Li, X.; Fu, K. K.; and Chuah, C.-N. 2023. Padclip: Pseudo-labeling with adaptive debiasing in clip for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16155–16165.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. 2018a. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2017. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, 5542–5550.
- Li, J.; Jing, M.; Lu, K.; Zhu, L.; and Shen, H. T. 2019. Locality preserving joint transfer for domain adaptation. *IEEE Transactions on Image Processing*, 28(12): 6103–6115.
- Li, X.; Li, J.; Du, Z.; Zhu, L.; and Shen, H. T. 2025a. Unified Modality Separation: A Vision-Language Framework for Unsupervised Domain Adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(11): 10604–10618.

- Li, X.; Li, J.; Li, F.; Zhu, L.; Yang, Y.; and Shen, H. T. 2025b. Generalizing Vision-Language Models to Novel Domains: A Comprehensive Survey. *arXiv preprint arXiv:2506.18504*.
- Li, X.; Li, Y.; Du, Z.; Li, F.; Lu, K.; and Li, J. 2024a. Split to Merge: Unifying Separated Modalities for Unsupervised Domain Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23364–23374.
- Li, X.; Zhang, D.; Du, Z.; Zhu, L.; Chen, Z.; and Li, J. 2025c. PatAug: Augmentation of Augmentation for Test-Time Adaptation. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 5080–5089.
- Li, Y.; Gong, M.; Tian, X.; Liu, T.; and Tao, D. 2018b. Domain generalization via conditional invariant representations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Li, Y.; Tian, X.; Gong, M.; Liu, Y.; Liu, T.; Zhang, K.; and Tao, D. 2018c. Deep domain generalization via conditional invariant adversarial networks. In *Proceedings of the European conference on computer vision (ECCV)*, 624–639.
- Li, Z.; Li, X.; Fu, X.; Zhang, X.; Wang, W.; Chen, S.; and Yang, J. 2024b. Promptkd: Unsupervised prompt distillation for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26617–26626.
- Liang, J.; Sheng, L.; Wang, Z.; He, R.; and Tan, T. 2024. Realistic Unsupervised CLIP Fine-tuning with Universal Entropy Optimization. In *Forty-first International Conference on Machine Learning*.
- Loshchilov, I. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Peng, X.; Bai, Q.; Xia, X.; Huang, Z.; Saenko, K.; and Wang, B. 2019. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1406–1415.
- 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.
- Recht, B.; Roelofs, R.; Schmidt, L.; and Shankar, V. 2019. Do imagenet classifiers generalize to imagenet? In *International conference on machine learning*, 5389–5400. PMLR.
- Shu, Y.; Guo, X.; Wu, J.; Wang, X.; Wang, J.; and Long, M. 2023. Clipood: Generalizing clip to out-of-distributions. In *International Conference on Machine Learning*, 31716–31731. PMLR.
- Tanwisuth, K.; Zhang, S.; Zheng, H.; He, P.; and Zhou, M. 2023. POUF: Prompt-oriented unsupervised fine-tuning for large pre-trained models. In *International Conference on Machine Learning*, 33816–33832. PMLR.
- Torralba, A.; and Efros, A. A. 2011. Unbiased look at dataset bias. In *CVPR 2011*, 1521–1528. IEEE.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5018–5027.
- Wang, H.; Ge, S.; Lipton, Z.; and Xing, E. P. 2019. Learning robust global representations by penalizing local predictive power. *Advances in Neural Information Processing Systems*, 32.
- Wortsman, M.; Ilharco, G.; Kim, J. W.; Li, M.; Kornblith, S.; Roelofs, R.; Lopes, R. G.; Hajishirzi, H.; Farhadi, A.; Namkoong, H.; et al. 2022. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7959–7971.
- Yang, Y.; Ko, J.; and Yun, S.-Y. 2023. Towards Difficulty-Agnostic Efficient Transfer Learning for Vision-Language Models. *arXiv preprint arXiv:2311.15569*.
- Zhang, J.; Wu, S.; Gao, L.; Shen, H. T.; and Song, J. 2024. Dept: Decoupled prompt tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12924–12933.
- Zhao, Y.; Zhong, Z.; Zhao, N.; Sebe, N.; and Lee, G. H. 2024. Style-hallucinated dual consistency learning: A unified framework for visual domain generalization. *International Journal of Computer Vision*, 132(3): 837–853.
- Zhong, T.; Chi, Z.; Gu, L.; Wang, Y.; Yu, Y.; and Tang, J. 2022. Meta-dmoe: Adapting to domain shift by meta-distillation from mixture-of-experts. *Advances in Neural Information Processing Systems*, 35: 22243–22257.
- Zhou, K.; Liu, Z.; Qiao, Y.; Xiang, T.; and Loy, C. C. 2022a. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 4396–4415.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022b. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16816–16825.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022c. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9): 2337–2348.