

# Leveraging Vision-Language Models for Improving Domain Generalization in Image Classification

Sravanti Addepalli\* Ashish Ramayee Asokan\* Lakshay Sharma R. Venkatesh Babu  
Vision and AI Lab, Indian Institute of Science, Bangalore

## Abstract

Vision-Language Models (VLMs) such as CLIP are trained on large amounts of image-text pairs, resulting in remarkable generalization across several data distributions. However, in several cases, their expensive training and data collection/curation costs do not justify the end application. This motivates a vendor-client paradigm, where a vendor trains a large-scale VLM and grants only input-output access to clients on a pay-per-query basis in a black-box setting. The client aims to minimize inference cost by distilling the VLM to a student model using the limited available task-specific data, and further deploying this student model in the downstream application. While naive distillation largely improves the In-Domain (ID) accuracy of the student, it fails to transfer the superior out-of-distribution (OOD) generalization of the VLM teacher using the limited available labeled images. To mitigate this, we propose Vision-Language to Vision - Align, Distill, Predict (VL2V-ADiP), which first aligns the vision and language modalities of the teacher model with the vision modality of a pre-trained student model, and further distills the aligned VLM representations to the student. This maximally retains the pre-trained features of the student, while also incorporating the rich representations of the VLM image encoder and the superior generalization of the text embeddings. The proposed approach achieves state-of-the-art results on the standard Domain Generalization benchmarks in a black-box teacher setting as well as a white-box setting where the weights of the VLM are accessible. Project page: <http://val.cds.iisc.ac.in/VL2V-ADiP/>

## 1. Introduction

While the initial success of Deep Learning was predominantly driven by training specialized models for each task or dataset [26, 29], recent research on foundation models [24, 34, 47, 65] eliminates the need for this by training generic models jointly over several modalities using large-

\*Equal Contribution.

Correspondence to Sravanti Addepalli <sravantia@iisc.ac.in>, Ashish Ramayee Asokan <ashish.ramayee@gmail.com>

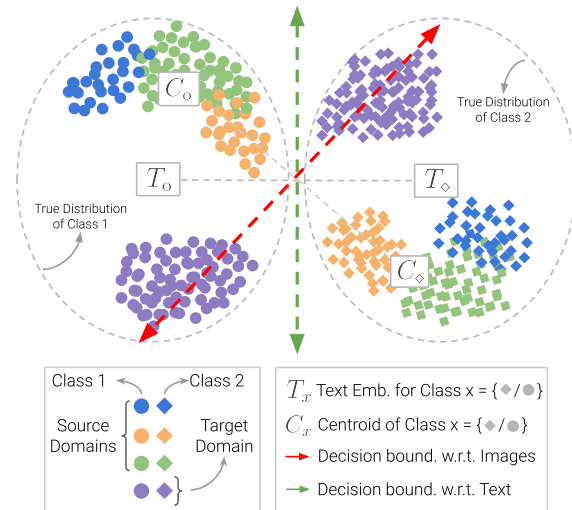


Figure 1. **Schematic diagram showing class and domain distributions in the shared text/ image embedding space of a VLM:** VLMs learn highly specialized image representations that are not domain invariant. Thus, a linear classifier (red decision boundary) that is trained over the vision encoder using limited training data cannot generalize well to the target domain (shown in purple). On the other hand, generic text embeddings such as “A photo of a class” represent the core concept of a class by virtue of their training method and vast training data. Thus, they generalize effectively across domains, and a zero-shot classifier (green decision boundary) aligns better with the true distribution of classes.

scale data. The use of both image and language modalities in large-scale Vision-Language Models (VLMs) enables their use in several applications including zero-shot classification, where the embedding of an image is compared with text embeddings for “a photo of a {class}” corresponding to every class during inference and the class with highest similarity is predicted. For example, the LiT (private) model [65] has been trained on 4 billion image-text pairs scraped from the web, and it achieves 85.2% zero-shot accuracy on ImageNet. VLMs also demonstrate extraordinary performance across several distributions, owing to the vast diversity of distributions seen during their training [47].

The remarkable OOD generalization of VLMs makes them suitable for use in several applications, either directly

or after finetuning on the downstream target dataset. Although several open-sourced models such as CLIP [47] are easily accessible, they may not be suitable for use in specialized applications. For example, the training data may have to be curated to reduce biases in the model or to ensure data privacy. It may also be important to ensure that the web-crawled data is clean for adoption in critical applications such as autonomous driving, to avoid data-poisoning attacks [4, 7]. Further, general-purpose models cannot be used in specialized applications such as medical diagnosis, where a VLM trained on data containing expert comments can be invaluable. This motivates the need for training highly specialized VLMs, which can be expensive due to the high costs associated with data collection/ curation and training.

A given downstream application may not always justify the cost associated with training such large-scale models. However, since these VLMs are generic and applicable to several downstream applications, one vendor could train such a model, and make it available to several clients in a black-box setting on a pay-per-query basis. In this case, the client can minimize their inference time costs by distilling the VLM to a student model using the limited available data specific to the downstream task, and deploying the distilled model rather than the VLM. Since the key motivation of using VLMs is their generalization to several domains, reliable OOD performance of the distilled model is crucial.

In this work, we consider the problem of distilling a multimodal Vision-Language model to a unimodal Vision model using limited task-specific downstream data, targeted toward better OOD performance with respect to the training domains used for distillation. While standard distillation of VLMs would require training data comprising of images and their associated text, in this case, the training data consists of only labeled images, similar to a classification task. Moreover, since the superior generalization abilities of VLMs are a result of the large-scale data ( $\sim 400$  million in the case of CLIP) used for their training, it is challenging to transfer these generalization abilities using a limited dataset comprising around 10,000 (0.01 million) images [16, 30, 60]. Lastly, since the VLM is considered to be a black box, the image encoder’s weights cannot be used as initialization to the vision model for further fine-tuning. However, since the downstream task-specific data is limited and is insufficient to train models from scratch, we assume that the student is initialized with weights from a model pre-trained on a publicly available dataset such as ImageNet-1K [26], as is common in a domain-generalization setting [19].

While knowledge distillation [20] using the labeled task-specific images improves the In-Domain accuracy on the downstream task, the improvement in OOD accuracy is marginal due to the limited size of the dataset. To address this, we first analyze the robustness of the text and image embeddings from the VLM and highlight the impor-

tance of text embeddings for better OOD generalization (see Fig.1). Further, we propose *VL2V-ADiP: Vision-Language to Vision - Align, Distill, Predict*, to firstly align the features of the pre-trained student model with the text and image modalities of the VLM teacher, and further fine-tune the student backbone in a distillation framework. This improves the robustness of the student, while also aligning the student embeddings to the VLM’s text embeddings corresponding to the respective class, thereby making the latter suitable for use as a classifier in the student without the need for further training. We summarize our contributions below:

- We investigate the robustness of image and text embeddings of VLMs, and highlight the importance of the text embeddings for better OOD generalization.
- To demonstrate the superior generalization of text embeddings, we propose VL2V-SD - a self-distillation approach, that improves the OOD generalization of the VLM’s pre-trained vision encoder using supervision from its own text encoder in a white-box VLM setting.
- We further propose VL2V-ADiP - a black-box distillation approach that effectively combines the features of a pre-trained vision model with the text and vision encoders of a VLM teacher for better OOD generalization.
- We demonstrate state-of-the-art results on standard Domain Generalization benchmark datasets in both white-box and black-box settings of the VLM teacher.

## 2. Related Works

**Vision-Language Models (VLMs):** VLMs are trained on large-scale datasets of image-text pairs crawled from the web [10, 53, 58]. Contrastive pre-training methods [24, 34, 47, 65] maximize the similarity between the embeddings of matching image-text pairs. More recent efforts focus on scaling up image-text pre-training to larger datasets [34], better supervision with unimodal and multimodal data [56], data-efficient pre-training [36], and fine-tuning VLMs for downstream tasks [18, 28, 63]. Another line of work focuses on leveraging CLIP [47] to improve the OOD generalization in image classification [40, 55].

**Domain Generalization (DG):** Prior works learn domain invariant representations by using augmentations [41, 42, 50], feature alignment across domains [17, 33, 35, 57], and disentangling domain-specific and domain-invariant features [11, 38, 45]. Gulrajani *et.al* [19] demonstrate that a well-tuned ERM model is comparable to several past DG algorithms [1, 5, 21, 27, 31, 32, 52, 62, 66]. SWAD [8] is a generic strategy that performs weight-averaging across different model snapshots in the optimal basin during training. MIRO [9] proposes Mutual Information (MI) regularization to maximally retain the pre-trained model’s representations. SAGM [61] aligns gradient directions between the  $\epsilon$ -perturbed loss and the empirical risk. We demonstrate performance gains when compared to the existing DG methods

Table 1. **CLIP on DG datasets:** Performance (%) of CLIP ViT-B/16 model on using different embeddings in the classifier head for computing similarity w.r.t. the test image during inference (T.E.: Text Embedding, I.E.: Image Embedding).

Embedding used for computing similarity	OH	TI	VLCS	PACS	Avg.
<b>E1:</b> Text embedding for "A photo of a {class}"	82.36	34.19	82.08	96.10	73.68
<b>E2:</b> Avg. text embedding for "A {domain} of a {class}" across all train domains	83.70	35.55	82.28	96.21	74.44
<b>E3:</b> Avg. image embedding of each class (Source)	71.37	33.99	48.21	79.03	58.15
<b>E4:</b> Avg. image embedding of each class (Target)	78.21	38.69	69.31	93.08	69.82
<b>E5:</b> Avg. image embedding of 10 images per class closest to test image (Source)	76.42	39.33	76.42	92.15	71.08
<b>E6:</b> Avg. image embedding of 10 images per class closest to test image (Target)	84.86	85.38	87.88	98.32	89.11

in both white-box and black-box settings of the VLM. More recently, several methods perform diverse training of multiple models, followed by weight-averaging to improve OOD generalization [2, 23, 48, 49]. However, they incur the enormous expense of training several diverse models. We thus restrict our comparison to DART [23], where the number of models trained is lower than the others. Another direction has been the generation of better prompts for improved generalization of VLMs [39, 54, 67]. These methods can be integrated with the proposed approach.

**Knowledge Distillation (KD):** Knowledge distillation aims to transfer knowledge from a powerful teacher model to a compact student model. Output-distillation [20] uses the teacher’s softmax outputs as guidance for the student. Feature-distillation methods [12–14, 37, 43, 51] use the teacher’s intermediate representations as guidance since they contain richer information. The proposed method operates in the black box distillation setup that allows only input-output access to the teacher network and restricts access to intermediate features. We thus compare the proposed black-box approach only with methods that do not access intermediate features [13, 20, 22].

### 3. Notations

We consider the problem of Knowledge Distillation (KD) from VLMs (teacher) to vision models (student) for improved OOD generalization. We denote the VLM’s text embedding for the input, “A photo of a {class}” for class  $c$  as  $\mathbf{T}_c$ . The image embedding corresponding to the image  $x$  is denoted as  $\mathbf{I}_x$ . The student model’s features corresponding to an input  $x$  are denoted as  $\mathbf{F}_x$ . These features when projected to the same dimension as the VLM embeddings ( $\mathbf{I}_x$  and  $\mathbf{T}_c$ ), are denoted as  $\mathbf{PF}_x$ . The cosine similarity between two vectors  $\mathbf{a}$  and  $\mathbf{b}$  is denoted as  $\cos(\mathbf{a}, \mathbf{b})$ . While we use CLIP [47] as the teacher for most of our analysis and experiments, we show the compatibility of the proposed approach with other VLMs [34, 36, 56, 64] in Suppl. Sec-3.3.

## 4. Robustness of CLIP embeddings

In this section, we present an overview of the training and zero-shot evaluation of the CLIP model, followed by a study on the robustness of its image and text embeddings.

### 4.1. CLIP training and zero-shot prediction

CLIP is a VLM that consists of an image encoder and a text encoder trained jointly on 400 million web-scraped image-text pairs [47]. The outputs of CLIP are the text and image embeddings for the respective text and image inputs. CLIP is trained using a contrastive loss which enforces similarity between the representations of positive image-text pairs and disparity between the representations of negative image-text pairs in the training minibatch. The training objective of CLIP is shown in Eq.1, where the embeddings of the  $i^{th}$  image-text pair are represented as  $\mathbf{I}_i$  and  $\mathbf{T}_i$  respectively.:

$$\mathcal{L}_{\text{CLIP}}(x_i) = -\log \frac{\exp(\cos(\mathbf{I}_i, \mathbf{T}_i))}{\sum_{c=1}^C \exp(\cos(\mathbf{I}_i, \mathbf{T}_c))} \quad (1)$$

The use of both text and image modalities during training enables the effective use of text embeddings as a zero-shot classifier, where inference can be performed directly without access to training data. For this, the text embeddings  $\mathbf{T}_c$  corresponding to the captions “A photo of a {class}” for each of the  $C$  classes are compared with the image embedding  $\mathbf{I}_i$ , and the class with maximum similarity is predicted as shown in Eq.2. Thus, the text embeddings  $\mathbf{T}_c$  can be viewed as the weights of a classifier for a CLIP backbone.

$$\hat{y}_i = \operatorname{argmax}_c \cos(\mathbf{I}_i, \mathbf{T}_c) \quad (2)$$

### 4.2. Characteristics of image and text embeddings

CLIP [47] shows remarkable zero-shot performance as shown in Table-1 (E1), and outperforms even an ImageNet pre-trained model that is explicitly fine-tuned using the downstream dataset (Table-3) on several DG datasets. To understand its superior zero-shot performance, we discuss below the robustness of its image and text embeddings.

**Characteristics of the Image Encoder:** In image classification tasks, the representations learned using standard ERM training are expected to exhibit invariances to several factors of variation due to: (a) the classification objective that enforces similar representations for different variations within a given class, which are different from the representations for other classes and (b) augmentations such as color jitter that enforce additional invariance. However, CLIP is trained using detailed captions for each image, such as “A brown cat sitting on a sofa”, “A black cat standing on two limbs”, and “A black dog with long ears and a lot of fur”. This allows CLIP to learn rich specialized representations for each attribute, with higher intra-class variance when compared to standard ERM training.

**Characteristics of the Text Encoder:** Robustness to distribution shifts can be achieved either by incorporating the required domain-invariance in the feature extractor, or by domain-invariant feature selection at the classification head. As discussed, CLIP’s image encoder produces input-specific representations and is not domain invariant. The superior generalization of CLIP on different datasets and domains is thus a result of the robust text embeddings. Although descriptive captions such as “A brown cat sitting on a sofa” are expected to have unique embeddings based on pose, location, etc., generic captions such as “A photo of a {class}” represent the core concept of the class.

**Empirical observations:** In Table-1, we present results to demonstrate the robustness of CLIP text embeddings on the multi-source Domain Generalization benchmarks, where each dataset consists of images from  $n$  domains. The results presented are an average of  $n$  cases, where each case corresponds to the training on a given set of  $n - 1$  domains and testing on the remaining unseen domain referred to as the target domain. Firstly, CLIP zero-shot results (E1) show the robustness of text embeddings across different distributions. The robustness of the text embeddings can be improved by explicitly enforcing domain invariance for better generalization to unseen domains (E2). However, the same level of robustness is not seen when we use an average of all image embeddings across the same domains for a given class, as shown in E3. This is due to the fact that while the text embeddings for “A {domain} of a {class}” are obtained by training over a large dataset (400 million images), the image embeddings are an average of  $\sim 0.01$  million images on the downstream dataset. This is illustrated in Fig.1, where the red decision boundary trained using image embeddings misclassifies images from the unseen target domain, while the green classifier trained using the generic text embeddings is closer to an ideal decision boundary. However, an average image embedding on the target domain improves results significantly (E4). The accuracy improves further when an average embedding corresponding to a few (10) source domain images closest (in terms of cosine similarity of image embeddings) to the target image is used for each class during prediction (E5). A similar experiment using images from the target domain (E6), achieves considerable improvements over the zero-shot accuracy in E1. We can consider E6 as an In-(target)-Domain setting where a labeled hold-out validation set is used during prediction for obtaining the average embeddings. Therefore, it is possible to outperform the generalized representations of the text embeddings only when labeled images from the target domain are available. *Thus, in a DG setting where the target domain is inaccessible, the generic text embeddings provide the best robustness across distribution shifts.* We therefore use the text embeddings to maximally transfer their robustness to the student model.

## 5. Proposed Approach: VL2V-ADiP

In this section, we first present the knowledge distillation framework adapted to a VLM-to-vision distillation setting. We further propose a self-distillation approach VL2V-SD, to enhance the robustness of the VLM image encoder, which motivates the proposed approach VL2V-ADiP.

### 5.1. Distillation from VLMs to Vision models

The standard framework of Knowledge Distillation (KD) [20] applies in the unimodal setting, for instance, where the teacher and student are both vision models. To adapt this to a VLM-to-vision distillation setting, a classification model (which we refer to as a VLM-classifier) is first constructed using the VLM. The feature extractor of the VLM-classifier comes from the image encoder of the VLM, and the linear classifier head obtains its weights from the text embeddings  $\mathbf{T}_c$  corresponding to the captions “a photo of a {class}” for all  $C$  classes. For every image-label pair  $(x_i, y_i)$ , the VLM-classifier first computes the image embedding  $\mathbf{I}_{x_i}$ , and further computes its similarity w.r.t. each of the text embeddings  $\mathbf{T}_c$ . The similarity scores  $\cos(\mathbf{I}_{x_i}, \mathbf{T}_c)$  are considered to be analogous to the logit layer, over which the softmax function is applied. The final distillation loss of the student that is minimized during training is shown below in Eq.4 for a minibatch of size  $n$ . Here,  $f_S(\cdot)$  and  $f_T(\cdot)$  represent the softmax representations of the student and teacher respectively, CE represents the Cross-Entropy loss, and KL represents the Kullback-Leibler divergence. Temperature scaling is applied to the logits before computing softmax [20]. We tune over the range - {0.1, 1, 10} on the in-domain validation set, and report the baseline best results in Table-3.

$$f_T(x_i) = \text{softmax}(\cos(\mathbf{I}_{x_i}, \mathbf{T}_c)) \quad (3)$$

$$\mathcal{L}_{\text{KD}} = \frac{1}{n} \sum_{i=1}^n \{ \text{CE}(f_S(x_i), y_i) + \lambda \cdot \text{KL}(f_T(x_i) || f_S(x_i)) \} \quad (4)$$

### 5.2. Self-Distillation from Text to Image encoders

As discussed in Section-4.2, the image encoder of CLIP is trained to represent all factors of variations in an image that are described in the corresponding caption used during its training. Therefore, it does not exhibit the invariances to factors such as color, texture, lighting, and pose, that are required for an image classification task. On the other hand, text prompts allow much better control over invariances in text-embeddings. For example, a prompt such as “A brown cat sitting on a sofa” uniquely represents the color, location, and pose of the cat in addition to the object itself, whereas a prompt “A photo of a cat” is invariant to such variations.

Motivated by this, we propose *Vision-Language to Vision Self Distillation (VL2V-SD)*, where the invariances of generic text embeddings are distilled to the image encoder



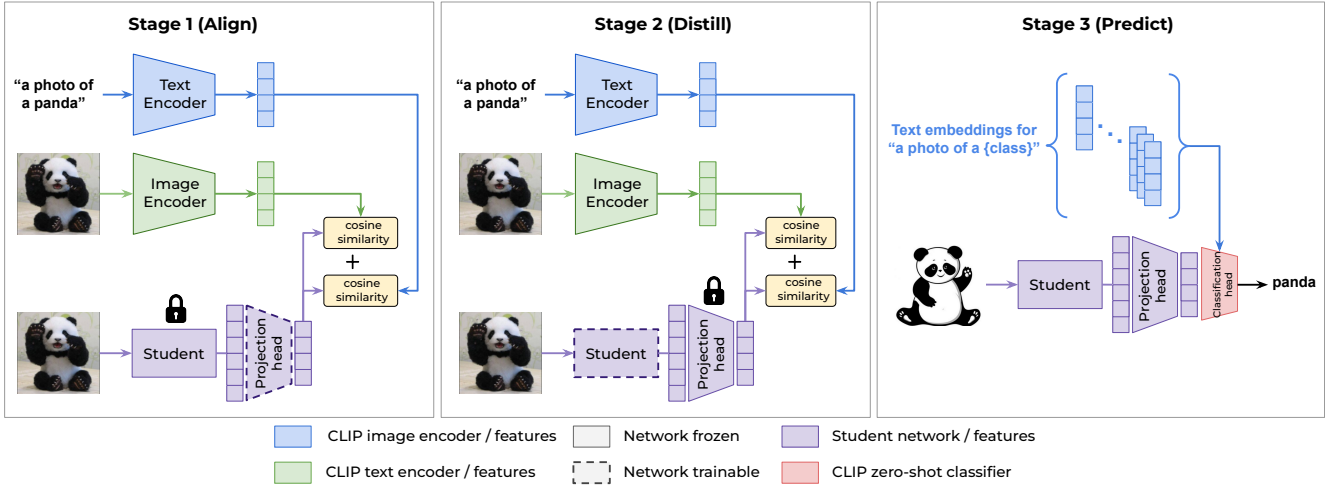


Figure 2. Overview of the proposed approach VL2V-ADiP, consisting of (a) Align, (b) Distill and (c) Predict Stages for Black-Box Distillation from Vision-Language to Vision (VL2V) models.

using the downstream dataset. We minimize the loss shown in Eq.5 for training, where  $\mathbf{T}_{y_i}$  represents the embedding for “A photo of a {class}” corresponding to the ground truth class  $y_i$ , and  $\mathbf{I}_{x_i}^t, \mathbf{I}_{x_i}^s$  represent the image embeddings of the teacher and student respectively for an image  $x_i$ .

$$\mathcal{L}_{SD} = -\frac{1}{2n} \sum_{i=1}^n \{ \cos(\mathbf{I}_{x_i}^s, \mathbf{T}_{y_i}) + \cos(\mathbf{I}_{x_i}^s, \mathbf{I}_{x_i}^t) \} \quad (5)$$

The image encoders of both teacher and student are the same (VLM’s image encoder) initially, and the student’s image encoder is updated during training. The text encoder is frozen during training. While the first term enforces the image encoder to learn representations that match the corresponding text embeddings more accurately on the given dataset, the second term ensures that there is no representation collapse, and retains the rich features learned by the image encoder. The classifier is composed of text embeddings corresponding to the prompts “a photo of a {class}” for all classes, as was the case in the VLM-classifier discussed in Section-5.1. We present the results of VL2V-SD in Table-2, where we note significant improvements over both CLIP zero-shot and the SOTA DG methods - SWAD [8], MIRO [9], DART [23], and SAGM [61] on the ViT-B/16 architecture with CLIP initialization. We demonstrate significant gains over the CLIP fine-tuning methods LP-FT [28], FLYP [18], CLIPood [55], and WiSE-FT [63]. We also achieve significant gains over a distillation method RISE [22], despite using the standard text prompt while RISE uses an expanded set of prompts. For WiSE-FT, we present the results with the optimal  $\alpha$  that achieves the best performance by using the search space  $\alpha = [0.1 - 0.9]$ . Similarly, for RISE, we present the results with the optimal hyperparameters  $\lambda_1, \lambda_2, \lambda_3$  based on their search space. We incorporate SWAD [8] in all our runs for the baselines as well

Table 2. **White-Box setting (CLIP initialization):** Performance (%) of the proposed approach VL2V-SD, compared to the existing methods. ViT-B/16 architecture is used. (S) denotes SWAD [8].

Method	OH	TI	VLCS	PACS	DN	Avg-ID	Avg-OOD
Zero-shot [47]	82.40	34.10	82.30	96.50	57.70	-	70.60
SWAD [8]	81.01	42.92	79.13	91.35	57.92	89.05	70.47
MIRO (S) [9]	84.80	<b>59.30</b>	82.30	96.44	60.47	<b>91.00</b>	76.66
DART (S) [23]	80.93	51.24	80.38	93.43	59.32	89.25	73.06
SAGM (S) [61]	83.40	58.64	82.05	94.31	59.05	89.74	75.49
LP-FT (S) [28]	81.17	47.26	80.88	92.92	57.04	88.97	71.85
FLYP (S) [18]	82.76	33.25	66.64	78.53	57.41	78.94	63.72
CLIPood (S) [55]	83.31	46.28	77.19	93.16	57.78	69.90	71.55
WiSE-FT (S) [63]	86.32	54.50	82.88	<b>97.29</b>	58.01	88.35	75.80
RISE (S) [22]	78.39	49.61	80.62	93.25	55.37	87.91	71.45
<b>VL2V-SD (Ours)</b>	<b>87.38</b>	58.54	<b>83.25</b>	96.68	<b>62.79</b>	89.99	<b>77.73</b>

as the proposed approach (denoted by “(S)”) since it is a generic technique that improves the generalization of any given method. However, we do not incorporate SWAD with WiSE-FT since it already performs weight averaging of the pre-trained and fine-tuned models. We present the results without SWAD in Suppl. Sec-3.2.

### 5.3. VL2V - Align, Distill, Predict (VL2V-ADiP)

While VL2V-SD is very effective in improving the performance of the VLM in a white-box setting, it does not address the problem of *black-box distillation* that we consider in this work, since it assumes access to the weights of the VLM model. In this section, we discuss how this method can be adapted to the black-box setting and present the proposed approach *VL2V-ADiP*. Since the amount of downstream task-specific data is assumed to be limited, the student is initialized with the best available pre-trained model such as ImageNet [26].

In the self-distillation case seen in Section-5.2, the goal was to induce domain invariance from the text encoder to the image encoder, while ensuring that the rich features

learned by the image encoder are not lost. Whereas, in the black-box setting, the goal is to distill the rich features learned by the image-encoder of the VLM and the domain invariance from the text embeddings. While the two cases can have a similar loss formulation as shown in Eq.5, the method VL2V-SD cannot directly be applied here since the feature dimension of the student model may be different from the dimension of VLM’s image and text embeddings. Even if the dimensions had matched, the features of the ImageNet pre-trained backbone and VLM’s embeddings would be misaligned. Directly enforcing a similarity loss in such a case would lead to forgetting of the pre-trained features. To address this, we propose *VL2V-ADiP*, *Vision-Language to Vision - Align, Distill, Predict* as shown in Fig.2 (Ref: Alg. 1 in Suppl.). We propose to first *Align* the features of the pre-trained student model with VLM embeddings using a linear projection head on the student backbone in Stage-1, as shown in Fig.2(a). For this, the student backbone is frozen and only the projection head is trained using the following loss, where  $\mathbf{PF}_{x_i}^s$  represents the projected features of the student model for the input  $x_i$ :

$$\mathcal{L}_{ADiP} = -\frac{1}{2n} \sum_{i=1}^n \{ \cos(\mathbf{PF}_{x_i}^s, \mathbf{T}_{y_i}) + \cos(\mathbf{PF}_{x_i}^s, \mathbf{I}_{x_i}^t) \} \quad (6)$$

Next, the aligned student features are refined using the VLM’s image and text embeddings in the *Distill* step as shown in Fig.2(b) (Stage-2). For this step, the linear projection head is frozen, and the feature extractor is trained using the same loss as Stage-1, as shown in Eq.6. While we give equal weight to both loss terms in Eq.6, we present the impact of varying this in Suppl. Sec-4. Finally, in Stage-3 (Predict), the VLM’s embeddings corresponding to “a photo of a {class}” for all  $C$  classes are used as weights of the classification head, and the class with the highest similarity to the given image embedding is predicted, as shown in Eq.2. In addition to learning the rich features from the VLM teacher’s image-encoder and invariances from its text encoder, enforcing similarity between the features of the student backbone and the text embeddings of the corresponding classes also makes the text embeddings more suitable for use as a classifier in Stage-3 (Predict). Thus, although Stages 1 and 2 explicitly train only the student backbone and projection head, they implicitly impact the effectiveness of the classification head as well. This facilitates the use of these embeddings directly as a classifier without the need for fine-tuning further on the downstream dataset. As shown in Table-5, further fine-tuning degrades the performance, since only In-Domain (ID) data is available for fine-tuning, which would make the classifier forget the domain invariance that it inherently possesses. However, fine-tuning with limited In-Domain data improves the ID accuracy as shown in the table.

## 6. Experiments and Results

### 6.1. Evaluation Details

Towards evaluating the OOD generalization of image classification models, we consider the five popular Domain Generalization (DG) datasets - OfficeHome (OH) [60], PACS [30], VLCS [16], Terra-Incognita (TI) [3] and DomainNet (DN) [44]. Each of the datasets consists of  $d$  domains, where  $d = 4$  for the first four datasets and  $d = 6$  for DomainNet. We present more details on the datasets in Suppl. Sec-2. We follow the DomainBed [19] framework for training and evaluation, where each domain is split into training and validation sets, and training is performed on  $d - 1$  domains while evaluation is done on the  $d^{th}$  unseen domain. Thus, for every dataset,  $d$  models are trained, by leaving one domain out for evaluation each time. The average accuracy across these  $d$  runs on the unseen test domain (OOD accuracy) and the average In-domain (ID) validation split accuracy are reported. The unseen test domain is neither used for training nor for hyperparameter tuning [19].

### 6.2. Training Details

The number of training iterations is set to 5000 for OfficeHome, PACS, VLCS, and Terra-Incognita, and to 15000 for DomainNet, as is standard [8, 9]. We use the Adam optimizer with a constant learning rate of  $5 \times 10^{-5}$  for all our experiments except CLIPood (S) in Table-2, for which we use the AdamW optimizer as specified in [55]. These settings are fixed for all datasets, and we do not introduce any additional hyperparameters in the proposed algorithm. Our primary evaluations consider a ViT-B/16 architecture [15] for both the VLM (teacher) Image-encoder and the student model to enable a fair comparison with existing methods, that use the same architecture. Since one of the key motivations of distillation is to train a low-capacity model, we present results with lower-capacity student architectures as well. We use a CLIP teacher model, which uses a Transformer architecture [59] with modifications described by Radford et al. [46] for the text encoder. The student model and our primary baselines for VL2V-ADiP use an ImageNet pre-trained initialization [26], as is standard in DG [19]. We obtain the results of baseline methods on ViTs in Tables 2 and 3 by integrating the official code of the respective method with the MIRO [9] code base to ensure a fair comparison. Additionally, we show results with other teacher VLMs in Suppl. Sec-3.3.

### 6.3. Comparison with the SOTA

We present the ID and OOD results of the proposed approach VL2V-ADiP when compared to the DG and KD baselines in Table-3. Gulrajani and Lopez-Paz [19] show that ERM training is a very strong baseline in DG and outperforms several older methods that were proposed. Thus,

Table 3. **SOTA comparison with ImageNet initialization:** Performance (%) of the proposed approach VL2V-ADiP, compared to existing KD and DG methods with SWAD (S). ViT-B/16 with ImageNet-1K initialization is used as the student.

Method	OH	TI	VLCS	PACS	DN	Avg-ID	Avg-OOD
ERM-LP (S)	71.48	31.35	77.52	67.02	36.65	73.99	56.81
ERM Full Fine-tuning (S)	83.22	50.05	80.33	90.28	56.10	89.31	72.00
LP-FT (S) [28]	81.55	51.61	80.17	91.20	56.03	<b>90.03</b>	72.11
SimKD (S) [13]	66.76	81.01	<b>83.92</b>	28.24	49.42	68.24	61.87
KD (S) [20]	82.73	48.40	80.48	91.46	56.11	89.20	71.84
MIRO (S) [9]	80.09	50.29	81.10	89.50	55.75	88.71	71.35
DART (S) [23]	83.75	49.68	77.29	90.55	58.05	88.54	71.86
SAGM (S) [61]	82.22	53.24	79.60	90.02	55.66	89.22	72.15
Text2Concept (S) [40]	70.24	26.46	64.77	79.03	23.26	53.15	52.82
RISE (S) [22]	83.48	52.55	83.70	93.54	56.58	88.91	73.97
<b>VL2V-ADiP (Ours)</b>	<b>85.74</b>	<b>55.43</b>	81.90	<b>94.94</b>	<b>59.38</b>	89.02	<b>75.48</b>

Table 4. **Distillation to lower capacity student models:** Performance (%) of the proposed approach VL2V-ADiP (denoted as Ours), when compared to the KD baseline [20] and ERM (S) or SWAD [8] with different student architectures. The teacher architecture is ViT-B/16. (S) denotes SWAD.

Student	Method	OH	TI	VLCS	PACS	DN	Avg-ID	Avg-OOD
ViT-B/16 (86M)	ERM (S)	83.22	50.05	80.33	90.28	56.10	<b>89.31</b>	72.00
	KD (S)	82.73	48.40	80.48	91.46	56.17	89.20	71.85
	Ours (S)	<b>85.74</b>	<b>55.43</b>	<b>81.90</b>	<b>94.94</b>	<b>59.38</b>	89.02	<b>75.48</b>
ViT-S/16 (22M)	ERM (S)	78.58	49.40	78.72	85.80	52.09	85.66	68.92
	KD (S)	78.14	50.11	79.14	85.97	52.05	<b>87.43</b>	69.08
	Ours (S)	<b>81.22</b>	<b>52.47</b>	<b>81.44</b>	<b>89.32</b>	<b>54.23</b>	86.59	<b>71.73</b>
DeiT-S/16 (22M)	ERM (S)	74.95	47.85	79.37	<b>89.22</b>	49.18	86.58	68.12
	KD (S)	74.65	48.11	78.86	88.14	49.10	<b>86.66</b>	67.77
	Ours (S)	<b>77.63</b>	<b>48.72</b>	<b>81.89</b>	88.97	<b>50.37</b>	85.28	<b>69.52</b>
ResNet-50 (26M)	ERM (S)	70.85	49.47	<b>79.50</b>	<b>88.05</b>	46.43	84.88	66.86
	KD (S)	70.67	51.22	78.63	87.23	46.31	<b>85.04</b>	66.81
	Ours (S)	<b>74.42</b>	<b>53.46</b>	79.23	86.72	<b>47.74</b>	84.73	<b>68.31</b>

ERM training of only the linear layer (ERM-LP) and ERM full fine-tuning (ERM-FFT) are two important baselines we consider. Further, we present the results of important SOTA DG methods - SWAD [8], MIRO [9], DART [23], and SAGM [61]. Additionally, we also show results with Text2Concept [40] and LP-FT [28] - simple fine-tuning strategies for improved OOD generalization on downstream tasks. While there are several methods that outperform ERM [6, 25, 41, 57], we do not present all of them here since the above-listed recent methods outperform them. Since we incorporate SWAD in the proposed approach, we present all the baseline results by integrating SWAD with each of them and denote them with a suffix “(S)” in the tables. Additionally, we also present three distillation results - (a) knowledge distillation from the VLM teacher to the vision student model as discussed in Section-5.1 and (b) distillation from the VLM teacher using SimKD [13], and (c) distillation from the VLM teacher using RISE [22]. Since the feature distillation methods listed in Sec. 2 cannot be extended to the black box distillation setting, we do not compare with them. The proposed approach VL2V-ADiP achieves  $\sim 1.5\%$  improvement on average OOD accuracy across all datasets with respect to the best baselines, with comparable ID accuracy. We note that Text2Concept scales

poorly on smaller datasets compared to the results shown in the paper [40]. SimKD [13] is a method intended for unimodal distillation and thus fails for the multimodal setup. Furthermore, we outperform the method RISE [22] by a large margin despite the latter using expanded text prompts. This shows that the proposed distillation method VL2V-ADiP effectively distills from VLMs with lesser prompt engineering and text supervision compared to RISE.

#### 6.4. Distillation to lower capacity student models

We present results of distilling from a CLIP ViT-B/16 teacher to different student architectures and capacities (ViT-S/16, DeiT-S, ResNet-50) in Table-4. We observe substantial gains over the Knowledge Distillation and ERM (S) or SWAD baseline across all student architectures. We note that the distillation performance is best when the architecture (ViT) of the teacher matches that of the student. For example, although ViT-S/16, DeiT-S/16, and ResNet-50 have a similar number of parameters, the performance of the ViT-S/16 student is the best, showing the architecture plays a role in the nature of representations learned, thereby allowing better transfer when they match.

#### 6.5. Ablation Study

We present an ablation study of the proposed method in Table-5. Firstly, we study the impact of combining stages-1 (Align) and 2 (Distill), i.e., training the linear projection layer and the backbone jointly in A1. This results in a 2.2% drop in ID accuracy and a 6.3% drop in OOD accuracy, highlighting that the pre-trained ImageNet features are retained better when they are aligned before distillation. In A2, the linear projection head is also trained in Stage-2. The results of this case are similar to the proposed approach, showing that even when the linear head is changed, it does not get updated much after Stage-1 (Align). We next analyze the importance of distilling from both image and text encoders of CLIP by using only one of them individually in A3 and A4. In both cases, we note a significant drop in the OOD accuracy highlighting the importance of distilling from both encoders. Distilling only from the Image encoder has a larger drop due to the lack of alignment between the representations and the classifier head which is composed of CLIP text embeddings corresponding to each of the classes.

In Stage-3 (Predict) of the proposed approach, the CLIP text embeddings corresponding to the classes are directly used as the classifier weights. We explore the impact of training this classifier further using cross-entropy loss over the softmax representations shown in Eq.3, in experiments A5-A8. In A5 and A6, the classifier head is first initialized using the text embeddings of CLIP, and in A7 and A8, a random initialization is used for the classification head. In A5 and A7, only the classifier head is trained, whereas in A6 and A8, the full network is trained. All these experi-

Table 5. **Ablation Study of the proposed VL2V-ADiP:** Performance (%) on ViT-B/16 model with ImageNet initialization.

Method - Changes done w.r.t. VL2V-ADiP (Ours)	OH	TI	VLCS	PACS	Avg-ID	Avg-OOD
<b>VL2V-ADiP (Ours)</b>	85.74	55.43	81.90	94.94	92.74	79.50
A1: Combining "Align" and "Distill" stages	74.55	52.99	80.14	85.29	90.55	73.24
A2: Without freezing projection head in Stage-2	86.13	56.68	81.78	93.75	93.00	79.59
A3: Distilling only from Text encoder in Stages-1 and 2	83.18	47.02	79.83	90.86	91.87	75.22
A4: Distilling only from Image encoder in Stages-1 and 2	78.97	28.97	82.24	89.97	74.09	70.04
A5: Finetuning CLIP classifier-head in Stage-3 (CE loss) - CLIP init classifier	84.47	49.28	81.30	93.54	93.06	77.15
A6: Finetuning full network in Stage-3 (CE loss) - CLIP init classifier	83.88	49.83	80.13	92.25	93.37	76.52
A7: Finetuning classifier-head in Stage-3 (CE loss) - random init classifier	84.63	49.55	81.29	93.57	93.02	77.26
A8: Finetuning full network in Stage-3 (CE loss) - random init classifier	83.16	50.03	79.78	92.15	93.14	76.28




Figure 3. OOD accuracy (%) of the proposed approach when compared to KD and ERM baselines for select classes/ domains in OfficeHome and Terra-Incognita datasets, where the (a) ERM-IN (linear) (b) CLIP Zero-shot (c) ERM-CLIP (linear) performance is poor.

ments result in a drop in the OOD accuracy, indicating that training on the downstream dataset using the cross-entropy loss can destroy the domain invariances learned from CLIP. Since the impact of training the full network is higher, the drop is higher in this case when compared to finetuning only the classification head. In all cases, the In-Domain accuracy improves when compared to the proposed approach, with higher gains in the case when the full network is trained, confirming that training without the supervision of the CLIP teacher indeed overfits to the training domains.

We also compare the ERM and KD baselines with the proposed approach in some of the extreme cases in Fig.3, where the CLIP Image encoder is significantly better than the ERM model (a), and vice-versa (c). For this, we select the classes/domains with the highest difference (40-50%) between the ERM-CLIP (Linear) and ERM-ImageNet (IN) (Linear) baselines from the OfficeHome dataset. We also present the case where the CLIP zero-shot accuracy is poor in OfficeHome (22%) and Terra-Incognita datasets (0%) in Fig.3 (b). The KD baseline relies a lot on the ImageNet pre-trained backbone, hence it is poor when ERM-IN is poor, and is close to our approach when ERM-IN is good. The ERM fine-tuning baseline is best when the ImageNet features are much better than CLIP Image encoder’s features, but it is poor in the other case, as expected. Lastly, even when CLIP zero-shot accuracy is poor (close to 0%), the proposed approach achieves reasonable accuracy, since Stages-1 and 2 align the text embeddings better to the feature extractor, making them more suitable as a classifier.

## 7. Conclusion

In this work, we aim to leverage the superior generalization of large-scale Vision-Language Models towards improving the OOD generalization of Vision models. We consider a practical scenario where a client gets only black-box access to the model on a pay-per-query basis from the vendor, motivating the need for distilling the model first, and further using the student model during inference. Towards this, we first highlight the unique aspects of the image and text encoders of VLMs and further propose a self-distillation (SD) approach - VL2V-SD, to distill the superior generalization of the VLM’s text encoder to its image encoder, while retaining the rich representations of the latter. We further adapt this to a black-box distillation setting, by firstly projecting the representations of the pre-trained feature-extractor (student) to a space that is aligned with embeddings of the VLM teacher, and further finetuning the backbone using the same. The use of both text and image embeddings induces the rich representations of the VLM and its superior generalization to the student model, while also making the text embeddings more suitable for use in the classification head. Both proposed approaches (VL2V-SD and VL2V-ADiP) achieve substantial gains over prior methods on the popular Domain Generalization datasets.

**Acknowledgements.** This work is supported by the Kotak IISc AI-ML Center and a research grant CRG/2021/005925 from SERB, DST, Govt. of India. Sravanti Addepalli is supported by the Google PhD Fellowship and CII-SERB Prime Minister’s Fellowship for Doctoral Research.



## References

- [1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 2
- [2] Devansh Arpit, Huan Wang, Yingbo Zhou, and Caiming Xiong. Ensemble of averages: Improving model selection and boosting performance in domain generalization. *NeurIPS*, 2022. 3
- [3] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *ECCV*. Springer, 2018. 6
- [4] Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines. In *ICML*. PMLR, 2012. 2
- [5] Gilles Blanchard, Aniket Anand Deshmukh, Ürun Dogan, Gyemin Lee, and Clayton Scott. Domain generalization by marginal transfer learning. *JMLR*, 2021. 2
- [6] Manh-Ha Bui, Toan Tran, Anh Tran, and Dinh Phung. Exploiting domain-specific features to enhance domain generalization. *NeurIPS*, 2021. 7
- [7] Nicholas Carlini, Matthew Jagielski, Christopher A Choquette-Choo, Daniel Paleka, Will Pearce, Hyrum Anderson, Andreas Terzis, Kurt Thomas, and Florian Tramèr. Poisoning web-scale training datasets is practical. *arXiv preprint arXiv:2302.10149*, 2023. 2
- [8] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. *NeurIPS*, 2021. 2, 5, 6, 7
- [9] Junbum Cha, Kyungjae Lee, Sungrae Park, and Sanghyuk Chun. Domain generalization by mutual-information regularization with pre-trained models. *ECCV*, 2022. 2, 5, 6, 7
- [10] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12M: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *CVPR*, 2021. 2
- [11] Prithvijit Chattopadhyay, Yogesh Balaji, and Judy Hoffman. Learning to balance specificity and invariance for in and out of domain generalization. In *ECCV*. Springer, 2020. 2
- [12] Defang Chen, Jian-Ping Mei, Yuan Zhang, Can Wang, Zhe Wang, Yan Feng, and Chun Chen. Cross-layer distillation with semantic calibration. In *AAAI*, 2021. 3
- [13] Defang Chen, Jian-Ping Mei, Hailin Zhang, Can Wang, Yan Feng, and Chun Chen. Knowledge distillation with the reused teacher classifier. In *CVPR*, 2022. 3, 7
- [14] Pengguang Chen, Shu Liu, Hengshuang Zhao, and Jiaya Jia. Distilling knowledge via knowledge review. In *CVPR*, 2021. 3
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 6
- [16] Chen Fang, Ye Xu, and Daniel N. Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *ICCV*, 2013. 2, 6
- [17] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 2016. 2
- [18] Sachin Goyal, Ananya Kumar, Sankalp Garg, Zico Kolter, and Aditi Raghunathan. Finetune like you pretrain: Improved finetuning of zero-shot vision models. In *CVPR*, 2023. 2, 5
- [19] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *ICLR*, 2021. 2, 6
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2, 3, 4, 7
- [21] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *ECCV*. Springer, 2020. 2
- [22] Zeyi Huang, Andy Zhou, Zijian Ling, Mu Cai, Haohan Wang, and Yong Jae Lee. A sentence speaks a thousand images: Domain generalization through distilling clip with language guidance. In *ICCV*, 2023. 3, 5, 7
- [23] Samyak Jain, Sravanti Addepalli, Pawan Kumar Sahu, Priyam Dey, and R Venkatesh Babu. Dart: Diversify-aggregate-repeat training improves generalization of neural networks. In *CVPR*, 2023. 3, 5, 7
- [24] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*. PMLR, 2021. 1, 2
- [25] Daehee Kim, Youngjun Yoo, Seunghyun Park, Jinkyu Kim, and Jaekoo Lee. Selfreg: Self-supervised contrastive regularization for domain generalization. In *ICCV*, 2021. 7
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012. 1, 2, 5, 6
- [27] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghui Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). In *ICML*. PMLR, 2021. 2
- [28] Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *ICLR*, 2022. 2, 5, 7
- [29] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 2015. 1
- [30] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Deeper, broader and artier domain generalization. In *ICCV*, 2017. 2, 6
- [31] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018. 2
- [32] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C. Kot. Domain generalization with adversarial feature learning. In *CVPR*, 2018. 2
- [33] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *CVPR*, 2018. 2

- [34] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*. PMLR, 2022. 1, 2, 3
- [35] Ya Li, Mingming Gong, Xinmei Tian, Tongliang Liu, and Dacheng Tao. Domain generalization via conditional invariant representations. In *AAAI*, 2018. 2
- [36] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In *ICLR*, 2022. 2, 3
- [37] Dongyang Liu, Meina Kan, Shiguang Shan, and Xilin Chen. Function-consistent feature distillation. In *ICLR*, 2023. 3
- [38] Fangrui Lv, Jian Liang, Shuang Li, Bin Zang, Chi Harold Liu, Ziteng Wang, and Di Liu. Causality inspired representation learning for domain generalization. In *CVPR*, 2022. 2
- [39] Sachit Menon and Carl Vondrick. Visual classification via description from large language models. In *ICLR*, 2023. 3
- [40] Mazda Moayeri, Keivan Rezaei, Maziar Sanjabi, and Soheil Feizi. Text-to-concept (and back) via cross-model alignment. *PMLR*, 2023. 2, 7
- [41] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias. In *CVPR*, 2021. 2, 7
- [42] Oren Nuriel, Sagie Benaim, and Lior Wolf. Permuted adain: Reducing the bias towards global statistics in image classification. In *CVPR*, 2021. 2
- [43] Peyman Passban, Yimeng Wu, Mehdi Rezagholizadeh, and Qun Liu. Alp-kd: Attention-based layer projection for knowledge distillation. In *AAAI*, 2021. 3
- [44] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, 2019. 6
- [45] Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. Efficient domain generalization via common-specific low-rank decomposition. In *ICML*. PMLR, 2020. 2
- [46] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019. 6
- [47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*. PMLR, 2021. 1, 2, 3, 5
- [48] Alexandre Rame, Matthieu Kirchmeyer, Thibaud Rahier, Alain Rakotomamonjy, Patrick Gallinari, and Matthieu Cord. Diverse weight averaging for out-of-distribution generalization. *NeurIPS*, 2022. 3
- [49] Alexandre Rame, Kartik Ahuja, Jianyu Zhang, Matthieu Cord, Léon Bottou, and David Lopez-Paz. Model ratatouille: Recycling diverse models for out-of-distribution generalization. In *ICML*. PMLR, 2023. 3
- [50] Alexander Robey, George J Pappas, and Hamed Hassani. Model-based domain generalization. *NeurIPS*, 2021. 2
- [51] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015. 3
- [52] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*, 2019. 2
- [53] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 2
- [54] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time prompt tuning for zero-shot generalization in vision-language models. In *NeurIPS*, 2022. 3
- [55] Yang Shu, Xingzhuo Guo, Jialong Wu, Ximei Wang, Jianmin Wang, and Mingsheng Long. Clipood: Generalizing clip to out-of-distributions. *PMLR*, 2023. 2, 5, 6
- [56] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *CVPR*, 2022. 2, 3
- [57] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*. Springer, 2016. 2, 7
- [58] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 2016. 2
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 6
- [60] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 2017. 2, 6
- [61] Pengfei Wang, Zhaoxiang Zhang, Zhen Lei, and Lei Zhang. Sharpness-aware gradient matching for domain generalization. In *CVPR*, 2023. 2, 5, 7
- [62] Yufei Wang, Haoliang Li, and Alex C Kot. Heterogeneous domain generalization via domain mixup. In *ICASSP*. IEEE, 2020. 2
- [63] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *CVPR*, 2022. 2, 5
- [64] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. In *ICLR*, 2022. 3
- [65] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *CVPR*, 2022. 1, 2

- [66] Marvin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: Learning to adapt to domain shift. *NeurIPS*, 2021. [2](#)
- [67] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 2022. [3](#)