

# Distilling Large Vision-Language Model with Out-of-Distribution Generalizability

Xuanlin Li\* Yunhao Fang\* Minghua Liu Zhan Ling Zhuowen Tu Hao Su  
UC San Diego

## Abstract

*Large vision-language models have achieved outstanding performance, but their size and computational requirements make their deployment on resource-constrained devices and time-sensitive tasks impractical. Model distillation, the process of creating smaller, faster models that maintain the performance of larger models, is a promising direction towards the solution. This paper investigates the distillation of visual representations in large teacher vision-language models into lightweight student models using a small- or mid-scale dataset. Notably, this study focuses on open-vocabulary out-of-distribution (OOD) generalization, a challenging problem that has been overlooked in previous model distillation literature. We propose two principles from vision and language modality perspectives to enhance student’s OOD generalization: (1) by better imitating teacher’s visual representation space, and carefully promoting better coherence in vision-language alignment with the teacher; (2) by enriching the teacher’s language representations with informative and fine-grained semantic attributes to effectively distinguish between different labels. We propose several metrics and conduct extensive experiments to investigate their techniques. The results demonstrate significant improvements in zero-shot and few-shot student performance on open-vocabulary out-of-distribution classification, highlighting the effectiveness of our proposed approaches. Code released at [this link](#).*

## 1. Introduction

In recent years, there has been a significant growth and development of large vision language models (large VLMs) such as CLIP [39], GLIP [28], OFA [49], SimVLM [53], BEiT [50], Florence [60], and Flamingo [1], which have been pretrained on massive amounts of internet-scale data. These models have demonstrated enormous potential in a wide range of downstream applications, including classification/detection, image-to-text generation, and vision-language reasoning, especially for the out-of-distribution (OOD) samples and open-vocabulary settings. Despite their

potential, the large model sizes, high computational resource requirements, and inefficient inference speed of these models restrict their deployment on mobile and IoT edge devices [21, 44], as well as in scenarios like robotic control [4] that require rapid network inference. Thus, it would be ideal if we can obtain small and compact models that still possess strong generalizability towards diverse open-set concepts encountered in the real world.

We note that the internet-scale data has endowed foundation models with well-aligned vision and language representation spaces across diverse domains and datasets. Ideally, if such representation spaces can be perfectly distilled into smaller models, the resulting models can naturally possess similar open-set OOD generalization ability as their larger counterparts. However, it is usually not the case, as demonstrated in our later experiments. Therefore, we ask the following question: *How to effectively distill representation spaces of large vision-language teacher models to benefit OOD generalization of smaller student models?*

In this study, we investigate the principles and techniques for distilling visual representations from large vision-language models using small- to mid-scale datasets, with a specific focus on out-of-distribution (OOD) generalization. Although large-scale datasets with image-text pairs exist, as a pioneer work in this direction, we argue that small- to mid-scale datasets has many practical scenarios for vision application researchers (e.g., for robotics), because this is more flexible, allowing for faster research and development cycles with fewer resource dependencies. Additionally, as we shall see, extensive experiments yield a deeper understanding of the representation spaces of vision-language models, which can be beneficial for future tasks like distilling image-based foundation models for detection [22, 49], segmentation [24, 10], and other data modalities like 3D geometries [58, 36, 29]. To keep study focused and fundamental, we also choose the image classification task to explore many possible strategies.

Based on extensive experimental results, we propose to preserve the internal structure of the representation spaces. Specifically, we propose techniques to maintain the relationship between the visual and language representation spaces of the teacher models and a novel strategy to enhance the

\* Equal contributions. Corresponding authors: xul012@ucsd.edu, yuf026@ucsd.edu

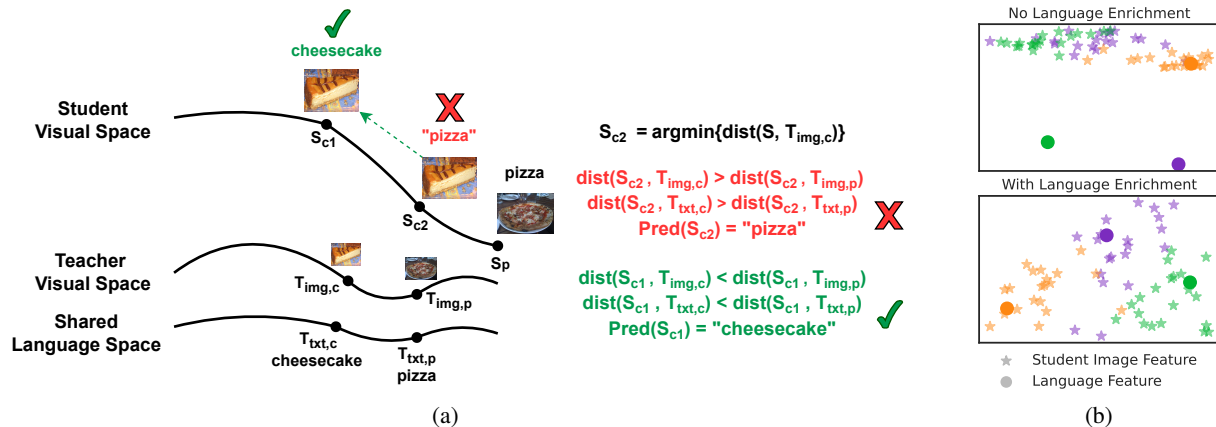


Figure 1: (a) Illustration of better teacher-student visual space alignments, motivated by our finding that achieving precise matching between teacher and student’s high-dimensional visual spaces is inherently challenging. While the student cheesecake image feature  $S_{c2}$  minimizes its distance to the teacher’s cheesecake image feature, it is even closer to the teacher’s pizza image feature. This discrepancy results in poor coherence in visual space and vision-language alignment with the teacher, causing classification mistakes. By replacing the “minimizing distance” requirement with the “relative distance” requirement, which encourages student visual features to be closer to their corresponding teacher visual features than other teacher visual features, we greatly enhance visual space and vision-language alignment coherence with the teacher, thereby improving OOD generalization. (b) UMAP embeddings of student visual features and text features before and after language representation enrichment with LLMs. Different colors denote different OOD classes on tiered-ImageNet. Language representation enrichment confers better clusterings of image features around their corresponding text features, enhancing OOD generalization.

text representation used during distillation. Across our study, we make the following contributions:

- 1) We motivate the ability for students to generalize towards out-of-distribution (OOD) concepts by designing several metrics that measure the visual representation space consistency and the vision-language alignment consistency between the student and the teacher vision-language model.
- 2) We find that by better imitating teacher’s visual representation space, and carefully promoting better coherence in vision-language alignment with the teacher, we substantially strengthen student’s OOD generalization ability.
- 3) We further improve student’s OOD generalization ability by enriching teacher’s language representations with more informative, finegrained, and meaningful semantic attributes to effectively distinguish between different labels.
- 4) We conduct a thorough experimental analysis to understand the efficacy and impact of our techniques on students’ OOD generalization ability.

## 2. Related Work

**Vision-Language Models (VLMs).** Recent years have witnessed tremendous progress on vision language models [39, 28, 49, 53, 50, 60, 1]. Various lines of approaches have been proposed to improve the downstream task performance of VLMs, such as prompt learning [65, 59, 61, 66], prompt engineering [38, 19, 30], and finetuning via vision-language fusion [43, 51, 25]. Different from these works that utilize existing VLMs on downstream tasks, our work

studies distilling large VLM’s feature space structures towards smaller student networks. Specifically, we investigate principles that effectively leverage the visual representation spaces and vision-language alignments in large VLMs to enhance the OOD generalization ability of students.

**Network Distillation on VLMs.** The teacher-student distillation framework [20] aims to distill teacher model knowledge into student models, thereby improving their downstream task performance. A classic setting considers teachers and students to be vision-only. In particular, distilling image representations from teacher visual backbones, such as CNNs [17] and Vision Transformers [12], have been well-studied [47, 34, 31, 6, 46, 37, 14, 56]. With the success of large-scale VLMs, recent works start to distill their powerful multimodal visual and language representations towards various downstream tasks, such as object detection [15, 32], visual-language reasoning [52, 8, 13], scene understanding [36], or adopt distillation during vision language model pretraining to further enhance their performance [26, 27, 11, 64, 45].

As the real world is full of concepts unseen during model training, it is essential for student models to attain strong OOD generalization capabilities. However, prior works have rarely *carefully* studied how to design distillation approaches to enhance such ability. In this work, we propose principles and approaches to improve student’s OOD generalization ability, and we carefully analyze their efficacies through various metrics and extensive experiments.

### 3. Overview

**Problem Setup.** We distill a large vision-language teacher model  $T$  (e.g., CLIP ViT-L/14) to a small student image model  $S$  (e.g., ResNet18) by focusing on out-of-distribution (OOD) generalization for open-vocabulary object classification. We choose small- or mid-scale datasets to achieve the distillation so that the distillation process is flexible for fast research cycle and has less resource dependency.

The teacher consists of an image encoder  $T_{\text{img}}(\cdot)$  and a text encoder  $T_{\text{txt}}(\cdot)$ . During distillation, we keep the flexibility of the existing teacher text encoder for the open-set setting, and we let the student model  $S$  be vision-only, i.e.,  $S = S_{\text{img}}$ . Through this process, we hope that  $S$  not only achieves high prediction accuracy on seen labels  $\mathcal{Y}_{\text{id}}$ , but also attains strong generalization ability on out-of-distribution labels  $\mathcal{Y}_{\text{ood}}$ . In addition, we train students from scratch to avoid label contamination, allowing us to more carefully assess and understand their OOD generalization ability.

**Experiment Setup.** We are given a training (distillation) dataset  $\mathcal{X}_{\text{train}}$ , an in-distribution evaluation dataset  $\mathcal{X}_{\text{id}}$ , and an out-of-distribution evaluation dataset  $\mathcal{X}_{\text{ood}}$ . Each dataset consists of image-label pairs  $\{(\mathbf{x}, y)\}$ . For  $(\mathbf{x}, y) \in \mathcal{X}_{\text{train}} \cup \mathcal{X}_{\text{id}}$ , we have  $y \in \mathcal{Y}_{\text{id}}$ , and for  $(\mathbf{x}, y) \in \mathcal{X}_{\text{ood}}$ , we have  $y \in \mathcal{Y}_{\text{ood}}$ , where  $\mathcal{Y}_{\text{id}}, \mathcal{Y}_{\text{ood}}$  denote in-distribution and out-of-distribution label sets, and  $\mathcal{Y}_{\text{id}} \cap \mathcal{Y}_{\text{ood}} = \emptyset$ . We evaluate the student on  $\mathcal{X}_{\text{id}}$ , along with zero-shot and few-shot generalization on  $\mathcal{X}_{\text{ood}}$ . For few-shot learning, we finetune our student models (we will show later in our ablation study that finetuning achieves much higher performance on  $\mathcal{X}_{\text{ood}}$  than training-free retrieval). We also adopt balanced training batches where at most half of samples come from few-shot data and the rest from  $\mathcal{X}_{\text{train}}$ . This strategy leads to similar performance on  $\mathcal{X}_{\text{id}}$  before and after finetuning.

We adopt a diverse collection of recognition tasks using small to medium-scale datasets, including Caltech-Birds [55], StanfordCars [23], Flower102 [35], Food101 [3], SUN397 [57], and tiered-ImageNet [40, 9]. We split the dataset labels such that  $|\mathcal{Y}_{\text{id}}| = |\mathcal{Y}_{\text{ood}}|$ , except tiered-ImageNet, which comes with an existing split.

**Organization of Exposition.** Like other papers on understanding neural network behaviors [16, 2, 18, 48], we will propose techniques and invent metrics to help understand the effectiveness of each technique. Following the organization of a representative work of this kind [62], we present different techniques in a waterfall style and organize the presentation of each technique by grouping the approach, metric, and results together.

**Approach Overview.** Our approach aims to preserve the structure of the visual representation space and its relationship with the text representation space inherited from the teacher model to enhance the student’s OOD generalization ability. Specifically, (1) In Sec. 4, we show that by better imitating teacher’s visual representation space, and carefully

promoting better coherence in vision-language alignment with the teacher, students achieve substantially better OOD generalization. (2) In Sec. 5, we show that by enriching teacher’s language representations, there are more meaningful attributes to effectively distinguish between different labels, thereby further enhancing student’s OOD generalization ability. A summary of our main experimental findings is illustrated in Fig. 1.

## 4. Teacher-Student Visual Space and Vision-Language Alignments

### 4.1. Better Imitating Teacher’s Visual Representation Space

**Naive baseline without visual space alignment between teacher and student.** A straightforward approach to training a student model is to directly align its visual representation with teacher’s language representation through *only* the following contrastive loss:

$$\begin{aligned} \mathcal{L}_{\text{cls}}(\mathbf{x}, y) &= \sum_{y'} -\mathbf{1}_{y'=y} \log P_S(y'|\mathbf{x}) \\ P_S(y|\mathbf{x}) &= \frac{\exp(\cos(S(\mathbf{x}), T_{\text{txt}}(l(y)))/\tau)}{\sum_{y' \in \mathcal{Y}} \exp(\cos(S(\mathbf{x}), T_{\text{txt}}(l(y')))/\tau)} \end{aligned} \quad (1)$$

Here  $\tau$  is the temperature parameter.  $S(\mathbf{x})$  denotes the feature of  $\mathbf{x}$  extracted by a student image model  $S$ .  $l(y) = \text{prompt} + \text{description}(y)$  is a language generation function that maps the label  $y$  to its natural language representation, which consists of a prompt and a description of  $y$ . For this baseline, we use the suggested prompt “A photo of a” from the CLIP paper, and let  $\text{description}(y)$  to be simply the class label name (e.g., “lotus”). From here on, we assume that all teacher and student representations are normalized, i.e.,  $S(\mathbf{x}) \leftarrow \frac{S(\mathbf{x})}{\|S(\mathbf{x})\|_2}$ ,  $T_{\text{txt}}(l(y)) \leftarrow \frac{T_{\text{txt}}(l(y))}{\|T_{\text{txt}}(l(y))\|_2}$ . In this case,  $\cos(S(\mathbf{x}), T_{\text{txt}}(l(y))) = 1 - \frac{\|S(\mathbf{x}) - T_{\text{txt}}(l(y))\|_2^2}{2}$ .

Training student models solely using  $\mathcal{L}_{\text{cls}}$  does not enforce similarity between teacher and student visual representation space structures. Consequently, student visual space structures tend to overfit to training labels, hindering their OOD generalizability. This motivates us to introduce auxiliary losses to align teacher and student visual representation spaces, which we will describe next.

**Directly fitting teacher visual features.** We note that teacher’s visual representation space is well-aligned with language across diverse datasets, and such alignment demonstrates strong generalization across many domains. By imitating teacher’s visual space structure, we hope to enhance the ability for student’s visual space to generalize and extrapolate towards unseen concepts, thereby implicitly enhancing the generalizability of student’s vision-language alignment and improving its OOD generalization.

	CaltechBirds	StanfordCars	Flower102	Food101	SUN397	tiered-ImageNet	Average
CLIP ViT-L/14	70.0 / 70.5	79.3 / 78.3	74.4 / 84.1	90.5 / 91.2	72.8 / 74.4	71.1 / 76.3	76.4 / 79.1
CLIP RN50	57.1 / 56.4	53.2 / 56.3	59.3 / 65.3	76.5 / 78.3	65.0 / 66.3	55.7 / 62.0	61.1 / 64.1
Closed-Set Classification	48.1 / NA / 18.1	27.9 / NA / 10.1	77.1 / NA / 45.0	71.7 / NA / 30.3	57.8 / NA / 31.1	63.4 / NA / 31.2	57.7 / NA / 27.6
$\mathcal{L}_{\text{cls}}$	61.0 / 14.2 / 34.9	56.3 / 14.7 / 20.0	81.2 / 4.5 / 46.2	72.2 / 16.1 / 24.5	57.5 / 13.6 / 28.6	64.4 / 13.9 / 27.5	65.4 / 12.8 / 30.3
$\mathcal{L}_{\text{mse}}$	27.0 / 12.0 / 14.3	5.5 / 3.8 / 4.0	48.1 / 7.3 / 15.0	45.0 / 17.0 / 19.3	24.3 / 11.0 / 14.5	49.3 / 14.8 / 23.2	33.2 / 11.0 / 15.1
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{mse}}$	<b>63.7</b> / 17.4 / 36.2	62.2 / 18.8 / 35.1	82.6 / 6.3 / 46.0	72.3 / 19.0 / 35.5	57.1 / 15.3 / 29.4	66.2 / 14.9 / 28.5	67.4 / 15.3 / 35.1
$\mathcal{L}_{\text{im-cst}}$	42.1 / 21.3 / 29.1	33.2 / 13.7 / 20.0	54.8 / 13.3 / 27.3	70.0 / 34.9 / 36.8	45.2 / 22.8 / 27.2	46.3 / 22.8 / 30.8	48.6 / 21.5 / 28.5
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{im-cst}}$	60.9 / 20.4 / 37.6	59.6 / 18.3 / 31.2	82.4 / 12.7 / <b>52.5</b>	74.0 / 30.5 / 42.0	<b>62.5</b> / 18.8 / <b>35.2</b>	64.4 / 18.0 / 33.5	67.3 / 19.8 / 38.7
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{im-cst}} + \mathcal{L}_{\text{mse}}$	62.5 / 20.8 / <b>39.0</b>	59.6 / 19.0 / 33.1	82.6 / 12.0 / 48.7	<b>75.0</b> / 31.2 / 42.0	60.0 / 19.8 / <b>35.2</b>	67.0 / 19.4 / 34.6	67.8 / 20.3 / <b>38.8</b>
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{im-cst}} + \mathcal{L}_{\text{vlprox}} (+\mathcal{L}_{\text{mse}})^{\dagger}$	62.3 / 21.6 / <b>39.0</b>	<b>63.9</b> / 19.8 / 38.5	<b>82.7</b> / 14.6 / 52.0	74.3 / 32.0 / <b>43.2</b>	61.7 / 21.5 / 34.7	<b>67.5</b> / 20.5 / <b>35.3</b>	<b>68.7</b> / 21.7 / <b>40.5</b>
$\mathcal{L}_{\text{im-cst}} + \mathcal{L}_{\text{vlprox}} (+\mathcal{L}_{\text{mse}})$	45.3 / <b>21.9</b> / 30.4	46.5 / 17.8 / 26.9	66.9 / 13.5 / 35.4	71.4 / <b>35.2</b> / 40.0	52.0 / <b>23.1</b> / 28.8	57.5 / <b>23.0</b> / 33.2	56.6 / <b>22.4</b> / 32.5

Table 1: Comparison between student models trained without teacher-student visual representation space alignment ( $\mathcal{L}_{\text{cls}}$  only), with direct teacher visual feature fitting ( $+\mathcal{L}_{\text{mse}}$ ), with improved teacher-student visual space alignment ( $+\mathcal{L}_{\text{im-cst}}$ ), and with improved preservation of teacher’s vision-language alignment structure ( $+\mathcal{L}_{\text{vlprox}}$ ). We adopt ResNet18 as the student and CLIP ViT-L/14 as the teacher. The three numbers  $x_1/x_2/x_3$  in each entry denote the evaluation performance on  $\mathcal{X}_{\text{id}}$ , zero-shot performance on  $\mathcal{X}_{\text{ood}}$ , and 5-shot performance on  $\mathcal{X}_{\text{ood}}$ , respectively (“NA”=not applicable). As reference, we also report CLIP performance ( $x_1/x_2$  in each entry denote zero-shot results on  $\mathcal{X}_{\text{id}}$  and  $\mathcal{X}_{\text{ood}}$ ), along with a closed-set classification baseline that uses separate classifiers for  $\mathcal{Y}_{\text{id}}$  and  $\mathcal{Y}_{\text{ood}}$ . Note that CLIP was pretrained on LAION [42], where the training image-text pairs contain many concepts similar to those in  $\mathcal{Y}_{\text{id}}$  and  $\mathcal{Y}_{\text{ood}}$ , resulting in higher evaluation performance on  $\mathcal{X}_{\text{ood}}$ .

	Food101	SUN397
$\mathcal{L}_{\text{mse}}$	0.24 / 28.4°	0.36 / 34.9°
$\mathcal{L}_{\text{mse}}$ (RN50)	0.24 / 28.4°	0.35 / 34.4°
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{mse}}$	0.45 / 39.2°	0.71 / 49.8°
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{mse}} + \mathcal{L}_{\text{im-cst}}$	0.65 / 47.6°	0.82 / 53.8°
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{im-cst}}$	1.29 / 69.2°	1.28 / 68.9°

Table 2: Average MSE / degree difference between student visual features and teacher visual features for students trained with different strategies. We adopt ResNet18 as student and CLIP ViT-L/14 as teacher, except “ $\mathcal{L}_{\text{mse}}$  (RN50)”, where the student is ResNet50, and the teacher is CLIP ResNet50. We find that it is very challenging for students to *precisely* match teacher’s visual features through  $\mathcal{L}_{\text{mse}}$ . Moreover, in conjunction with Tab. 1, we find that a lower visual feature MSE is not necessary for better student OOD generalization.

A direct approach to achieve this is to align the teacher and student visual representations through the Mean Squared-Error (MSE) loss:

$$\mathcal{L}_{\text{mse}}(\mathbf{x}) = \|S(\mathbf{x}) - T_{\text{img}}(\mathbf{x})\|_2^2 \quad (2)$$

In Tab. 1, we show that adding  $\mathcal{L}_{\text{mse}}$  on top of  $\mathcal{L}_{\text{cls}}$  improves student OOD generalization. However, upon further examination, we find that *students face significant challenges in precisely reproducing teacher’s visual representations*, as evidenced by the substantial errors shown in Tab. 2. Such errors persist even when the student and teacher networks possess the same representation power (e.g., both ResNet50 networks). This phenomenon highlights that achieving precise matching between teacher and student’s high-dimensional visual feature spaces is inherently challenging, which can

<sup>†</sup>If  $\mathcal{L}_{\text{mse}}$  is under a parenthesis, then we report the better performance between adding and not adding  $\mathcal{L}_{\text{mse}}$ . However the impact of  $\mathcal{L}_{\text{mse}}$  is not significant.

be attributed to differences in weight initialization, training data, and the presence of local minima in the loss landscape. Moreover, we later find that when students struggle to precisely match teacher’s visual features, they also struggle to preserve teacher’s *local visual space structure and relative visual feature relationship* between different images, hindering their OOD generalization ability.

**Better imitating teacher visual representation space.** Since precisely matching teacher’s visual features is inherently challenging, we propose to augment the training objective with the following contrastive loss, which “softly” matches teacher’s visual features:

$$\mathcal{L}_{\text{im-cst}}(\mathbf{x}) = \frac{\exp(-\|S(\mathbf{x}) - T_{\text{img}}(\mathbf{x})\|_2^2/\tau)}{\sum_{\mathbf{x}'} \exp(-\|S(\mathbf{x}) - T_{\text{img}}(\mathbf{x}')\|_2^2/\tau)} \quad (3)$$

By combining  $\mathcal{L}_{\text{im-cst}}$  with  $\mathcal{L}_{\text{cls}}$ , we observe in Tab. 1 that the student exhibits significantly better zero-shot and few-shot OOD generalization ability across different datasets. Interestingly, such improvement is accompanied by a larger distance between student and teacher visual features, as shown in Tab. 2. This finding suggests that a lower MSE visual feature matching loss, which only captures the absolute distance to teacher’s visual features, does *not* necessarily imply better visual space consistency. This is because *MSE does not consider the coherence of relative visual feature relationships or local visual space structures between teacher and student*.

In the following paragraphs, we will develop several metrics to better assess the teacher-student visual space consistency. These metrics provide us with valuable insights into how  $\mathcal{L}_{\text{im-cst}}$  facilitates students to achieve closer visual space proximity to the teacher while yielding a deeper understanding of the teacher’s visual representation space.

**Quantifying teacher-student visual space alignment.** To quantify how student models preserve teacher’s visual rep-

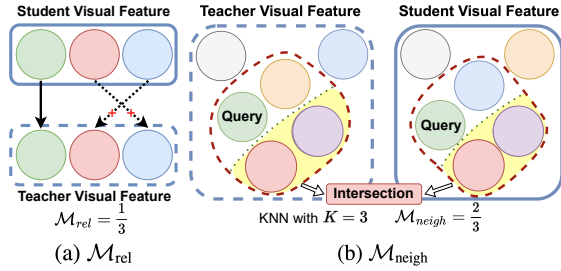


Figure 2: Illustrations of  $\mathcal{M}_{\text{rel}}$  and  $\mathcal{M}_{\text{neigh}}$  that measure the coherence of relative visual feature relationship and local visual space structure between the student and teacher. Each color corresponds to a single image. In (a), arrows indicate the nearest neighbors from student’s visual space to teacher’s visual space. Solid lines represent correct matching, while dashed lines indicate incorrect matching. In (b), the green circle represents a query image, and the red dashed regions represent the  $k$ -nearest neighbors (kNN) in both spaces. We measure the intersection between the two kNN sets.

representation space structures, we propose two metrics. For the first metric, we are motivated by the fact that if the relative relationships among teacher’s visual representations are preserved, then for an image  $\mathbf{x}$ ,  $S(\mathbf{x})$  tends to be the closest to  $T_{\text{img}}(\mathbf{x})$  rather than  $T_{\text{img}}(\mathbf{x}')$  of some other image  $\mathbf{x}' \neq \mathbf{x}$ . The latter scenario tends to break the relative feature relationship between  $\mathbf{x}$  and  $\mathbf{x}'$  that is originally present in the teacher, thereby causing the student to have distinct visual feature manifold structure from the teacher, as illustrated in Fig. 2(a). Furthermore, if  $S(\mathbf{x})$  is closer to  $T_{\text{img}}(\mathbf{x}')$ , then the language feature closest to  $S(\mathbf{x})$  often differs from the language feature closest to  $T_{\text{img}}(\mathbf{x})$ , and the latter is typically the ground truth language label. Thus, such discrepancy often results in erroneous label predictions. We can then formally define the first metric as follows:

$$\mathcal{M}_{\text{rel}}(\mathcal{X}) = \frac{\sum_{i=1}^{|\mathcal{X}|} \mathbf{1}[i = \text{argmin}_j \|T_{\text{img}}(\mathbf{x}_j) - S(\mathbf{x}_i)\|_2^2]}{|\mathcal{X}|} \quad (4)$$

The second metric measures the proximity between local neighborhoods of student visual features and teacher visual features. That is, for each image, how much the  $k$  images whose features have the closest proximity to it (excluding itself) overlap between the student and the teacher (see Fig. 2b for illustration). We can define this metric as follows:

$$\mathcal{M}_{\text{neigh}}(\mathcal{X}, k) = \frac{\sum_{i=1}^{|\mathcal{X}|} |\text{KNN}(S, \mathbf{x}_i, k) \cap \text{KNN}(T_{\text{img}}, \mathbf{x}_i, k)|}{k|\mathcal{X}|} \quad (5)$$

Here  $\text{KNN}(S, \mathbf{x}_i, k) = \{\text{argbottom}_k_j \|S(\mathbf{x}_j) - S(\mathbf{x}_i)\|_2\}$  outputs the  $k$  nearest-neighbor image-ids of  $S(\mathbf{x}_i)$ , and  $\text{KNN}(T_{\text{img}}, \mathbf{x}_i, k)$  is defined similarly. This metric is complementary to  $\mathcal{M}_{\text{rel}}$  as it only requires the set of neighbor

$\mathcal{M}_{\text{rel}} \uparrow$   $\mathcal{M}_{\text{neigh}} \uparrow$	$\mathcal{X}_{\text{train}}$   $\mathcal{X}_{\text{ood}}$		$k = 3$	$k = 5$	$k = 10$
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{mse}}$	0.030	0.004	0.13 / 0.06	0.18 / 0.07	0.27 / 0.08
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{mse}} + \mathcal{L}_{\text{im-cst}}$	0.305	0.022	0.20 / 0.10	0.25 / 0.11	0.34 / 0.13

Table 3: We evaluate  $\mathcal{M}_{\text{rel}}$  (middle 2 columns) and  $\mathcal{M}_{\text{neigh}}$  (right 3 columns) on different students to measure their coherence with teacher’s relative visual feature relationships and local visual feature structures (higher the better). For  $\mathcal{M}_{\text{neigh}}$ ,  $x_1/x_2$  in each entry denote  $\mathcal{M}_{\text{neigh}}(\mathcal{X}_{\text{train}})$  and  $\mathcal{M}_{\text{neigh}}(\mathcal{X}_{\text{ood}})$ , respectively. Metrics are evaluated on Flower102. More results in Appendix.

image ids to be identical between the student and the teacher, and does not enforce these neighbor images to have the same relative feature relationships.

### $\mathcal{L}_{\text{im-cst}}$ improves teacher-student visual space alignments through better relative and local visual space coherence.

We assess  $\mathcal{M}_{\text{rel}}$  and  $\mathcal{M}_{\text{neigh}}$  on different students and present the results in Tab. 3. We find that students trained solely with  $\mathcal{L}_{\text{mse}}$  and without  $\mathcal{L}_{\text{im-cst}}$  exhibit poor values for both  $\mathcal{M}_{\text{rel}}$  and  $\mathcal{M}_{\text{neigh}}$ . Notably,  $\mathcal{M}_{\text{rel}}(\mathcal{X}_{\text{train}})$  is only 0.03, indicating a large discrepancy between student and teacher visual spaces, even on the training set. On the other hand, after incorporating  $\mathcal{L}_{\text{im-cst}}$ , we observe significant improvements of  $\mathcal{M}_{\text{rel}}$  and  $\mathcal{M}_{\text{neigh}}$  on both  $\mathcal{X}_{\text{train}}$  and  $\mathcal{X}_{\text{ood}}$ . This demonstrates that the student exhibits much better consistency in relative and local visual structures with the teacher, both for seen and unseen concepts, thereby enhancing the generalization and extrapolation ability of the student’s visual space. A better-aligned visual space also implicitly enables better generalization and extrapolation in vision-language alignments (which we show later), contributing to improved OOD performance. Furthermore, our findings suggest a broader principle: when students face challenges in precisely matching the teacher’s visual space, the proximity of local visual feature structures and relative visual feature relationships have a greater impact on OOD generalization than the absolute distance to the teacher’s features.

## 4.2. Better Coherence with Teacher’s Vision-Language Alignment Structure

**Motivation.** In Section 4.1, we focused on improving the student’s OOD generalization ability by better aligning student-teacher visual spaces. Since teacher’s visual space is well-aligned with language across diverse concepts and domains, a better student coherence with teacher’s visual space *implicitly* leads to better vision-language (V-L) alignments. Naturally, an alternative perspective to improve student’s OOD generalization becomes to enhance its *explicit* V-L alignments and improve their coherence with the teacher’s, where we previously only used a simple contrastive V-L matching loss  $\mathcal{L}_{\text{cls}}$ . Another motivation to focus on *explicit* V-L alignments arises from our finding that they play an

essential role to ensure precise and accurate V-L alignments, especially when training on seen concepts or performing few-shot learning on novel concepts. *Relying solely on implicit V-L alignments is inadequate in these scenarios.* This is evident in Tab. 1, where *solely* utilizing the visual space alignment loss  $\mathcal{L}_{\text{im-cst}}$  yields better performance on 0-shot  $\mathcal{X}_{\text{ood}}$  (where classes are unseen) but worse performance on  $\mathcal{X}_{\text{id}}$  and 5-shot  $\mathcal{X}_{\text{ood}}$  (where classes are seen). On the other hand, by combining both implicit and explicit V-L alignment losses ( $\mathcal{L}_{\text{im-cst}} + \mathcal{L}_{\text{cls}}$ ), students excel in all of  $\mathcal{X}_{\text{id}}$ , 0-shot  $\mathcal{X}_{\text{ood}}$ , and 5-shot  $\mathcal{X}_{\text{ood}}$  scenarios. Therefore, by improving explicit V-L alignments, we not only hope to further enhance student’s 0-shot OOD generalization ability, but also improve their performance on familiar concepts and their ability to few-shot adapt to novel concepts.

**Explicitly enhancing teacher-student vision-language alignment coherency.** We note that while  $\mathcal{L}_{\text{cls}}$  performs explicit V-L alignments, it has the limitation of indiscriminately pushing an image away from all non ground-truth language features, therefore disregarding teacher’s relative alignment relationship between the same image and different language features. Furthermore, we find that even though preserving teacher’s relative V-L alignment structure is desirable, it may not always be perfect due to potential misalignments between teacher image features and their corresponding language labels. These misalignments can introduce inconsistent noise during distillation, ultimately harming student performance.

Motivated by these observations, we propose to augment our training objective with  $\mathcal{L}_{\text{vlprox}}$ , which effectively and carefully preserves the teacher’s vision-language alignment structure while accounting for potential misalignments:

$$\begin{aligned} \mathcal{L}_{\text{vlprox}}(\mathbf{x}, k) &= I(\mathbf{x}) \cdot \mathcal{D}_{\text{KL}}(P_{T, \text{topk}}(\cdot|\mathbf{x}) || P_{S, \text{topk}}(\cdot|\mathbf{x})) \\ I(\mathbf{x}) &= \mathbf{1}[\text{argmax}_y P_T(y|\mathbf{x}) = \text{label}(\mathbf{x})] \\ P_{\cdot, \text{topk}}(y|\mathbf{x}) &= \frac{\mathbf{1}_{y \in Y_{\text{topk}}} P_{\cdot}(y|\mathbf{x})}{\sum_{y \in Y_{\text{topk}}} P_{\cdot}(y|\mathbf{x})}; Y_{\text{topk}} = \text{argtopk}_y P_T(y|\mathbf{x}) \end{aligned} \quad (6)$$

Here  $P_T$  and  $P_S$  denote teacher and student label probabilities;  $I(\cdot)$  filters out images misaligned with language labels; and  $k$  controls the number of most-similar language features for each image. In our implementations, we find a larger  $k$  beneficial for OOD generalization, and we choose  $k = 256$ .

We demonstrate the effectiveness of  $\mathcal{L}_{\text{vlprox}}$  in Tab. 1. We find that by combining  $\mathcal{L}_{\text{vlprox}}$  with  $\mathcal{L}_{\text{cls}}$  and  $\mathcal{L}_{\text{im-cst}}$ , we further improve student’s ability to generalize towards OOD concepts. Interestingly, we also observe that while  $\mathcal{L}_{\text{cls}}$  and  $\mathcal{L}_{\text{vlprox}}$  both explicitly perform V-L alignments, adding them together yields significantly better student performance on  $\mathcal{X}_{\text{id}}$  and 5-shot  $\mathcal{X}_{\text{ood}}$  than solely keeping  $\mathcal{L}_{\text{vlprox}}$ . This observation is distinct from those in the traditional model distillation literature [20], where distilling teacher logits alone from vision-only models typically produces good student performance.

$\mathcal{M}_{\text{vlalign}} \downarrow$	$k = 2$	$k = 3$	$k = 5$
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{mse}}$	0.20 / 0.50	0.68 / 1.45	2.67 / 4.73
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{mse}} + \mathcal{L}_{\text{im-cst}}$	0.18 / 0.43	0.62 / 1.3	2.52 / 4.24
$\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{mse}} + \mathcal{L}_{\text{im-cst}} + \mathcal{L}_{\text{vlprox}}$	0.17 / 0.39	0.59 / 1.17	2.17 / 4.20

Table 4: We evaluate  $\mathcal{M}_{\text{vlalign}}$  (lower the better) on different students to measure their proximity with the teacher vision-language alignment structure.  $x_1/x_2$  in each entry denote  $\mathcal{M}_{\text{vlalign}}(\mathcal{X}_{\text{train}})$  and  $\mathcal{M}_{\text{vlalign}}(\mathcal{X}_{\text{ood}})$ , respectively. Metrics are evaluated on Flower102. More results are in Appendix.

**Quantifying teacher-student vision-language alignment coherency.** To better understand the vision-language alignment proximity between teacher and student, we propose a metric  $\mathcal{M}_{\text{vlalign}}$  to quantify such proximity:

$$\begin{aligned} \mathcal{M}_{\text{vlalign}}(\mathcal{X}, k) &= \frac{\sum_{i=1}^{|\mathcal{X}|} \#\text{reverse\_pairs}(\text{arrS}(i, k))}{|\mathcal{X}|} \\ \text{arrS}(i, k) &= [||S(\mathbf{x}_i) - T_{\text{txt}}(l(y_j))||_2]_{j \in \mathcal{I}(i, k)} \\ \mathcal{I}(i, k) &= \text{argtopk}([-||T_{\text{im}}(\mathbf{x}_i) - T_{\text{txt}}(l(y_j))||_2]_{j=1}^{|\mathcal{Y}|}) \end{aligned} \quad (7)$$

Here  $\#\text{reverse\_pairs}(\cdot)$  measures the number of reverse pairs  $\{(i, j) : a_i > a_j\}$  in an array. Overall,  $\mathcal{M}_{\text{vlalign}}$  measures the extent to which the distance ordering of the  $k$  nearest language features to the teacher image feature differs from the distance ordering of the same  $k$  language features with respect to the student image feature. In other words, a lower value of  $\mathcal{M}_{\text{vlalign}}$  indicates a higher degree of proximity to the teacher’s vision-language alignment, as teacher’s relative language orderings with respect to each image becomes better preserved.

We assess  $\mathcal{M}_{\text{vlalign}}$  on different students and present the results in Tab. 4. We find that  $\mathcal{L}_{\text{im-cst}}$  not only fosters strong student-teacher visual space alignments, but also *implicitly* facilitates effective V-L alignments along this process. This observation illustrates that a better relative and local visual space coherence with the teacher (indicated by better  $\mathcal{M}_{\text{rel}}$  and  $\mathcal{M}_{\text{neigh}}$ ) can lead to enhanced alignment coherence in the V-L space. Additionally, the inclusion of  $\mathcal{L}_{\text{vlprox}}$  further strengthens teacher-student V-L alignments. Such improved alignments are then effectively transferred to unseen concepts, demonstrating that the student has acquired better V-L alignment structures that generalize well to OOD scenarios.

## 5. Language Representation Enrichment

In Sec. 4, we focused on improving the imitation of teacher’s visual space and promoting better coherence with teacher’s vision-language alignment. Throughout this process, we kept the language representations fixed. However, the quality of language representations also plays a pivotal role in student learning and inference. Ideally, language representations should be capture precise, finegrained, and

	CaltechBirds	StanfordCars	Flower102	Food101	SUN397	tiered-ImageNet	Average
Tab. 1 best (no lang enrichment)	62.3 / 21.6 / 39.0	63.9 / 19.8 / 38.5	82.7 / 14.6 / 52.0	74.3 / 32.0 / 43.2	<b>61.7</b> / 21.5 / 34.7	67.5 / 20.5 / 35.3	68.7 / 21.7 / 40.5
Semantic Details	62.0 / <b>23.2</b> / 40.4	63.6 / 20.0 / 37.5	82.4 / 17.6 / 52.7	74.8 / 33.9 / 43.7	60.8 / 23.3 / 36.8	<b>69.8</b> / 23.5 / <b>36.2</b>	68.9 / 23.6 / 41.2
Auxiliary Captions	<b>62.5</b> / 21.4 / <b>41.0</b>	<b>65.5</b> / 19.0 / 38.1	81.9 / 14.3 / 52.5	75.4 / 33.3 / <b>44.0</b>	61.6 / 22.1 / 36.9	68.7 / 21.0 / 34.4	69.3 / 21.9 / 41.2
Semantics + Caption	62.0 / 22.7 / 39.8	64.9 / <b>20.4</b> / <b>39.7</b>	<b>83.7</b> / <b>18.2</b> / <b>53.4</b>	<b>75.6</b> / <b>35.7</b> / 42.9	61.0 / <b>24.0</b> / <b>37.5</b>	68.9 / <b>23.6</b> / 35.8	<b>69.4</b> / <b>24.1</b> / <b>41.5</b>

Table 5: Comparison between different language representation enrichment strategies. The three numbers  $x_1/x_2/x_3$  in each entry denote the evaluation performance on  $\mathcal{X}_{id}$ , zero-shot performance on  $\mathcal{X}_{ood}$ , and 5-shot performance on  $\mathcal{X}_{ood}$ , respectively.

	StanfordCars	tiered-ImageNet
No language enrichment	<b>63.9</b> / 19.8 / <b>38.5</b>	67.5 / 20.5 / 35.3
Prompt in Tab. 5	63.3 / <b>20.0</b> / 37.5	<b>69.8</b> / 23.5 / 36.2
More Succinct	63.0 / 18.9 / 37.6	68.8 / 23.1 / 36.8
More Detailed	62.9 / 19.0 / 35.1	69.3 / <b>24.2</b> / <b>37.7</b>
More Distinct	<b>63.9</b> / 19.7 / 37.1	69.2 / 23.3 / 37.0

Table 6: Results on leveraging different prompts to control semantic details of label descriptions generated by ChatGPT.

meaningful semantic attributes, such that the student can effectively distinguish between different labels. We therefore ask the following question: can we leverage better and richer teacher language representations to further enhance student’s OOD generalization ability? We propose the following candidate strategies:

**Enriching semantic details of label descriptions by prompting LLMs.** Previously, when we generate language representations  $l(y) = \text{prompt} + \text{description}(y)$  for student training, we adopted a simple strategy. In particular, for the description of a label  $y$ , we merely used its label name, e.g., “lotus”. However, these simplistic descriptions overlook many finegrained properties of semantic categories, such as the shape, color, and texture of flowers, along with the description of their petals, leaves, and stems. In addition, we hope to automatically and efficiently generate enriched language descriptions for a wide range of labels, ensuring scalability for an arbitrary number of labels. Motivated by the recent progress on instruction-finetuned large language models (LLMs) [33, 7, 41, 54], which have demonstrated impressive sequence generation abilities given user prompts, we find these models well-suited for our goal. Therefore, we propose to use ChatGPT [5, 33] to generate category descriptions. We prompt ChatGPT with the following instruction: “Use a single sentence to describe the appearance and shape of {cls}. Only describe the shape and appearance”. This allows ChatGPT to generate informative, finegrained, and meaningful descriptions for target classes (e.g., “large, round, flat leaves; tall, slender stems; delicate petals in shades of pink, white, or yellow”), while keeping sequence lengths within CLIP text encoder’s limit. We then set  $\text{description}(y)$  by concatenating “a photo of {cls}” with ChatGPT-generated class descriptions. We still keep the same vision-language alignment losses ( $\mathcal{L}_{cls}$  and  $\mathcal{L}_{vprox}$ ) as before.

**Augmenting text through auxiliary captions.** Currently, during student training, there is only one language description per category, i.e.,  $|\{l(y) : (\mathbf{x}, y) \in \mathcal{X}_{train}\}| = |\mathcal{Y}_{id}|$ .

On the other hand, the number of training images significantly exceeds the number of labels, i.e.,  $|\mathcal{X}_{train}| \gg |\mathcal{Y}_{id}|$ . We therefore wish to generate language descriptions for each individual image, such that we can substantially enrich the number of language features during student training, which potentially benefits student performance. To achieve this, we propose using OFA [49] to generate captions for each image, resulting in a new dataset  $\{(\mathbf{x}, \text{cap}(\mathbf{x}), y) : (\mathbf{x}, y) \in \mathcal{X}_{train}\}$  augmented with captions. During student training, besides using the same vision-language alignment losses  $\mathcal{L}_{cls}$  and  $\mathcal{L}_{vprox}$  as before, we also adopt the following auxiliary loss:

$$\mathcal{L}_{cap}(\mathbf{x}) = \frac{\exp(\cos(S(\mathbf{x}), T_{txt}(\text{cap}(\mathbf{x}))) / \tau)}{\sum_{(\mathbf{x}', y') : y' \neq y} \exp(\cos(S(\mathbf{x}'), T_{txt}(\text{cap}(\mathbf{x}')) / \tau)} \quad (8)$$

The loss pushes  $\mathbf{x}$  and its corresponding caption  $\text{cap}(\mathbf{x})$  together while pulling away from captions belonging to different categories. Our preliminary experiments show that distinguishing captions belonging to the same category could degrade student performance as they are usually similar. *Note that we only incorporate captions for the auxiliary loss during student training. For student inference and label predictions, we continue to use the same  $l(y)$  as before.*

## 5.1. Analyzing the Efficacy of Different Strategies

We adopt the aforementioned language-enriching strategies for student learning, and we present the results in Tab. 5. We find that combining LLM-enriched label descriptions with auxiliary captions yields the best OOD generalization. However, upon analyzing their individual effectiveness, we find that *LLM-enriched label descriptions provide significantly better zero-shot OOD benefit than auxiliary captions, and solely relying on auxiliary captions only marginally improves zero-shot OOD generalizability.* Upon further analysis, we find that many generated captions only broadly describe objects and are much less informative than ChatGPT-generated descriptions for distinguishing fine-grained categories. For instance, in the StanfordCars dataset, a generated caption for an “Acura Integra Type R 2001” image is “a white car is parked in a field”, and solely relying on the white color provides little information to distinguish different car categories. Consequently, captions have limited impact on enhancing the generalizability of student’s vision-language alignment structures.

Given the significant benefits of LLM-enriched label descriptions, we are particularly interested in exploring differ-

ent prompts to *control how ChatGPT generates semantic details* and their influence on OOD generalizability. We design the following prompts: **More Succinct**: “Use a single sentence to broadly describe the appearance and shape of {cls}. Don’t give too much details. Only describe the shape and appearance.” **More Detailed**: “Use a single sentence and short, simple, descriptive phrases to describe the detailed appearance and detailed shape of {cls}.” **More Distinct**: “Use a single sentence to describe the unique, distinctive appearance and shape of {cls}. Only describe the unique, distinctive shape and appearance.”

We compare these prompts in Tab. 6. Interestingly, we observe that generating more detailed semantic descriptions on labels does not always perform better. We conjecture that this is because (1) LLM-generated details are not grounded in specific images, causing some attributes to be invisible and confusing the students; (2) the teacher CLIP is trained on LAION [42], where most language descriptions do not contain many finegrained appearance details, so CLIP’s text embeddings are not very sensitive to some of these details. Additionally, we find that explicitly prompting ChatGPT to generate more concise text descriptions could be still helpful. Upon further analysis, we find that the resulting generations remain highly descriptive, albeit with slightly fewer details (e.g., when describing a “trumbone”, the more concise description becomes “a brass instrument with a long cylindrical tube curved into an elongated S shape with a flared bell at the end”, whereas under our original prompt, additional details like “a sliding U-shaped section called the slide” are included).

## 5.2. Comparing Language Representation Spaces Before and After Enrichment

To gain further insights into the efficacy of our language enrichment strategies, in this section, we analyze the changes in language space structures before and after adopting such strategies. On the out-of-distribution evaluation dataset of Flower102, we obtain text features  $\{T_{\text{txt}}(l_{\text{new}}(y))\}_{y=1}^{|\mathcal{Y}_{\text{ood}}|}$  and  $\{T_{\text{txt}}(l_{\text{old}}(y))\}_{y=1}^{|\mathcal{Y}_{\text{ood}}|}$ , where  $l_{\text{new}}$  denotes the language generation function with LLM enrichment, and  $l_{\text{old}}$  uses simple label names as label descriptions. We also obtain the average caption features for each class, namely  $\{\text{mean}[\{T_{\text{txt}}(\text{cap}(\mathbf{x})) : \text{label}(\mathbf{x}) = y\}]\}_{y=1}^{|\mathcal{Y}_{\text{ood}}|}$ . We then mean-center these 3 sets of text features, perform Singular-Value Decomposition, and plot the resulting eigenvalues in Fig. 3. We observe that language features of  $l_{\text{new}}$  contain many more large eigenvalues than  $l_{\text{old}}$ , demonstrating that the text space generated by  $l_{\text{new}}$  confers more independent and meaningful attributes to effectively distinguish between different classes. On the other hand, the language space of auxiliary captions only contains a small number of meaningful attributes, indicating that auxiliary captions are not very helpful for distinguishing between different labels. Additionally, in con-

junction with Tab. 5, we find that the advantages of richer attributes are concealed when evaluating students on the in-distribution classes, since students tend to overfit the existing attributes of these labels during training. However, for OOD scenarios, the presence of more finegrained attributes allows OOD text features to be more precisely aligned with image features, especially when no vision-language alignment training is done on these OOD samples (see Fig. 1b for an illustration). As a result, students exhibit better OOD generalization ability.

Furthermore, we analyze the average cosine similarity between all pairs of text features within each of the three sets we introduced earlier. We present the results in Tab. 7. We observe that LLM-enriched semantic details allow text features to naturally separate further apart from each other, making it easier to distinguish between different classes. In contrast, auxiliary captions show limited effectiveness in achieving such separation between classes.




Figure 3: Top 10 eigenvalues of text features.

	Cos. Sim.
w/o enrichment	0.5156
ChatGPT	0.4462
Caption	0.5572

Table 7: Average cosine similarity between all pairs of text features.

## 6. Ablations and Applications

In this section, we introduce ablation studies to enhance our main findings. We also illustrate an example to adopt our previous findings towards novel tasks and domains.

**Different designs on visual space and vision-language alignments.** We ablate on our designs of visual space and vision-language alignments and present the results in Tab. 8. We first explore various designs of  $\mathcal{L}_{\text{vprox}}$ , which aims to improve the *vision-language* alignment coherency with the teacher. We find that enhancing coherency with a greater number of most-similar language features by increasing  $k$  improves student OOD generalization. Additionally, it is also beneficial to filter out teacher’s image features that are misaligned with their corresponding language labels. Interestingly, we observe a different pattern when considering the teacher-student *visual space* alignments. In this case, we did not find filtering out misaligned image features helpful, and imitating the entirety of teacher’s visual space yields better student OOD generalization.

**Different student visual backbones.** In our prior experiments, we adopted ResNet as our student visual backbone. In this section, we further investigate whether our findings are consistent across different student network architectures. We adopt ViT-B/32 [12] as our student, while keeping CLIP ViT-



	Flower102	tiered-ImageNet
Main Paper Tab. 5 best	82.4 / 17.6 / 52.7	69.8 / 23.5 / 36.2
No filtering in $\mathcal{L}_{\text{vprox}}$	83.3 / 16.8 / 51.5	69.6 / 23.2 / 35.3
$k = 10$ for $\mathcal{L}_{\text{vprox}}$	83.4 / 17.0 / 51.4	68.4 / 22.6 / 34.6
$k = 1$ for $\mathcal{L}_{\text{vprox}}$	81.4 / 15.9 / 51.4	66.7 / 21.0 / 35.4
With filtering in $\mathcal{L}_{\text{im-cst}}$	83.6 / 16.6 / 51.0	68.5 / 23.1 / 33.4

Table 8: Ablation studies on (1) different designs of  $\mathcal{L}_{\text{vprox}}$ , and (2) whether to imitate the entirety of the teacher’s visual space in  $\mathcal{L}_{\text{im-cst}}$ .

$\mathcal{L}_{\text{cls}}$	$\mathcal{L}_{\text{im-cst}}$	$\mathcal{L}_{\text{vprox}}$	Semantics	Flower102	CaltechBirds	SUN397
✓				77.9 / 5.3 / 31.3	19.7 / 6.2 / 10.6	39.6 / 8.6 / 13.5
✓	✓			78.1 / 11.0 / 46.0	21.5 / 7.7 / 11.1	44.2 / 13.7 / 18.3
✓	✓		✓	77.8 / 11.9 / 46.5	<b>22.2 / 8.8</b> / 12.7	44.3 / 14.7 / 19.5
✓	✓	✓	✓	<b>78.9 / 13.1 / 47.7</b>	<b>22.2</b> / 8.5 / <b>13.1</b>	<b>45.0 / 15.4 / 21.0</b>

Table 9: Distilling CLIP ViT-L/14 into a student ViT-B/32 network. In each entry,  $x_1/x_2/x_3$  denote  $\mathcal{X}_{\text{id}}$  / zero-shot  $\mathcal{X}_{\text{ood}}$  / 5-shot  $\mathcal{X}_{\text{ood}}$  performance.

	CaltechBirds	SUN397	tiered-ImageNet
5-shot Finetune	<b>40.36</b>	<b>36.81</b>	<b>36.17</b>
5-shot Retrieval	24.39	24.40	24.27

Table 10: Comparison between finetuning student visual backbone vs. training-free retrieval on  $\mathcal{X}_{\text{ood}}$ . Finetuning student backbones significantly outperforms training-free retrieval.

L/14 as our VLM teacher. Results are shown in Tab. 9. We observe that even though we apply strong data augmentations (RandAugment) when training the student ViT-B/32, it still suffers from severe overfitting (the training accuracy on  $\mathcal{X}_{\text{train}}$  is  $> 90\%$ , which is significantly higher than the accuracy on  $\mathcal{X}_{\text{id}}$  in Tab. 9). The in-distribution and OOD generalization performance of the ViT-B/32 student are also worse than the ResNet student. Nevertheless, we still observe that better teacher-student visual space alignment, improved coherence in vision-language alignment, and language representation enrichment all enhance student’s OOD generalization ability, aligning with our previous findings.

**Few-shot learning strategies on  $\mathcal{X}_{\text{ood}}$ .** Previously, we finetune the student visual networks during few shot learning on  $\mathcal{X}_{\text{ood}}$ . Alternatively, we could adopt a finetune-free retrieval-based strategy like [63], which has been shown effective on large VLMs. We compare these two strategies in Tab. 10. Further implementation details are presented in Appendix. Interestingly, we observe that the retrieval-based strategy significantly underperforms finetuning student visual backbones. We hypothesize that this phenomenon arises because our student networks are trained on small to medium-scale datasets, and have been exposed to far fewer concepts than their teacher VLMs. Consequently, student’s image feature structure and vision-language alignment structure are less generalizable on OOD concepts, making them less suitable for retrieval-based approaches.

	$\mathcal{X}_{\text{id}}$	$\mathcal{Y}_{\text{id}}$ on $\mathcal{X}_{\text{ood}}$	$\mathcal{Y}_{\text{ood}}$ on $\mathcal{X}_{\text{ood}}$
Closed-Set	96.4 / 96.5	NA / 86.2	NA / 87.7
$\mathcal{L}_{\text{cls}}$	96.9 / 97.2	79.3 / 85.3	71.7 / 87.3
+ $\mathcal{L}_{\text{im-cst}}$	<b>99.2 / 99.2</b>	84.0 / 91.9	76.3 / 88.3
+ Semantic Enrich	98.2 / 99.0	<b>84.3 / 92.0</b>	<b>83.0 / 89.6</b>

(a) Overall accuracy over all YCB objects

	$\mathcal{X}_{\text{id}}$	$\mathcal{Y}_{\text{id}}$ on $\mathcal{X}_{\text{ood}}$	$\mathcal{Y}_{\text{ood}}$ on $\mathcal{X}_{\text{ood}}$
Closed-Set	91.6 / 91.9	NA / 57.8	NA / 35.9
$\mathcal{L}_{\text{cls}}$	94.0 / 94.4	46.5 / 54.7	23.3 / 32.8
+ $\mathcal{L}_{\text{im-cst}}$	<b>98.1 / 98.0</b>	55.3 / <b>70.8</b>	<b>23.7 / 47.3</b>
+ Semantic Enrich	97.2 / 97.4	<b>55.6 / 70.8</b>	11.7 / <b>50.7</b>

(b) F1-measure over objects that exist in observations.

Table 11: Results on grasp feasibility prediction in the Pick-Clutter task. In each entry,  $x_1/x_2$  denote zero-shot and few-shot evaluation results, respectively. ( $\mathcal{L}_{\text{vprox}}$  is not available for this experiment since multiple objects exist in an observation)

**Applications.** We additionally demonstrate in Fig. 4 and Tab. 11 that we can adopt our previous findings for improving student’s OOD generalization ability towards novel tasks and domains like robotics (more details in Appendix).



Figure 4: Illustration of the PickClutter task. Given the name of a target object, the robot agent needs to decide whether the corresponding object is graspable based on the current visual observation.

## 7. Conclusion

In this work, we studied distillation of large teacher vision-language models into lightweight student models by focusing on open-vocabulary out-of-distribution (OOD) generalization for object classification using small to medium-scale datasets. We investigated strengthening students’ OOD generalization ability from two key perspectives: first, by better imitating teacher’s visual representation space and carefully promoting better teacher-student vision-language alignment coherence; and second, by enhancing the teacher’s language representations with informative and meaningful semantic attributes to effectively differentiate between different labels. We analyzed the efficacy and impact of our techniques by introducing metrics and conducting a comprehensive experimental analysis. Along this process, we significantly improve student’s zero-shot and few-shot generalization performance on open-vocabulary OOD classification tasks.

**Acknowledgements.** We sincerely thank Jiayuan Gu from UC San Diego for their valuable discussions and feedback. This work is in part supported by Qualcomm AI and AI Institute for Learning-Enabled Optimization at Scale (TI-LOS).

## References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022. 1, 2
- [2] Devansh Arpit, Stanislaw Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 233–242. PMLR, 06–11 Aug 2017. 3
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 - mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014. 3
- [4] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022. 1
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 7
- [6] Pengguang Chen, Shu Liu, Hengshuang Zhao, and Jiaya Jia. Distilling knowledge via knowledge review. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5008–5017, 2021. 2
- [7] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022. 7
- [8] Wenliang Dai, Lu Hou, Lifeng Shang, Xin Jiang, Qun Liu, and Pascale Fung. Enabling multimodal generation on clip via vision-language knowledge distillation. *arXiv preprint arXiv:2203.06386*, 2022. 2
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 3
- [10] Zheng Ding, Jieke Wang, and Zhuowen Tu. Open-vocabulary panoptic segmentation with maskclip. *arXiv preprint arXiv:2208.08984*, 2022. 1
- [11] Xiaoyi Dong, Yinglin Zheng, Jianmin Bao, Ting Zhang, Dongdong Chen, Hao Yang, Ming Zeng, Weiming Zhang, Lu Yuan, Dong Chen, et al. Maskclip: Masked self-distillation advances contrastive language-image pretraining. *arXiv preprint arXiv:2208.12262*, 2022. 2
- [12] 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. *International Conference on Learning Representations*, 2021. 2, 8
- [13] Zhiyuan Fang, Jianfeng Wang, Xiaowei Hu, Lijuan Wang, Yezhou Yang, and Zicheng Liu. Compressing visual-linguistic model via knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1428–1438, 2021. 2
- [14] Zhiyuan Fang, Jianfeng Wang, Lijuan Wang, Lei Zhang, Yezhou Yang, and Zicheng Liu. Seed: Self-supervised distillation for visual representation. *International Conference on Learning Representations*, 2021. 2
- [15] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021. 2
- [16] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR, 06–11 Aug 2017. 3
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [18] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 3
- [19] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2
- [20] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *ArXiv*, abs/1503.02531, 2015. 2, 6
- [21] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 1
- [22] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790, 2021. 1
- [23] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *2013 IEEE International Conference on Computer Vision Workshops*, pages 554–561, 2013. 3
- [24] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *arXiv preprint arXiv:2201.03546*, 2022. 1
- [25] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 2

- [26] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2
- [27] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. 2
- [28] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10965–10975, 2022. 1, 2
- [29] Minghua Liu, Ruoxi Shi, Kaiming Kuang, Yin hao Zhu, Xuanlin Li, Shizhong Han, Hong Cai, Fatih Porikli, and Hao Su. Openshape: Scaling up 3d shape representation towards open-world understanding. *arXiv preprint arXiv:2305.10764*, 2023. 1
- [30] Vivian Liu and Lydia B Chilton. Design guidelines for prompt engineering text-to-image generative models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–23, 2022. 2
- [31] Yufan Liu, Jiajiong Cao, Bing Li, Chunfeng Yuan, Weiming Hu, Yangxi Li, and Yunqiang Duan. Knowledge distillation via instance relationship graph. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7096–7104, 2019. 2
- [32] Zongyang Ma, Guan Luo, Jin Gao, Liang Li, Yuxin Chen, Shaoru Wang, Congxuan Zhang, and Weiming Hu. Open-vocabulary one-stage detection with hierarchical visual-language knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14074–14083, 2022. 2
- [33] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*, 2022. 7
- [34] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3967–3976, 2019. 2
- [35] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3498–3505, 2012. 3
- [36] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. Open-scene: 3d scene understanding with open vocabularies. *arXiv preprint arXiv:2211.15654*, 2022. 1, 2
- [37] Zhiliang Peng, Li Dong, Hangbo Bao, Qixiang Ye, and Furu Wei. A unified view of masked image modeling. *arXiv preprint arXiv:2210.10615*, 2022. 2
- [38] Sarah Pratt, Rosanne Liu, and Ali Farhadi. What does a platypus look like? generating customized prompts for zero-shot image classification. *arXiv preprint arXiv:2209.03320*, 2022. 2
- [39] 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 2
- [40] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. *arXiv preprint arXiv:1803.00676*, 2018. 3
- [41] Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. Multitask prompted training enables zero-shot task generalization. *arXiv preprint arXiv:2110.08207*, 2021. 7
- [42] 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. 4, 8
- [43] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. How much can clip benefit vision-and-language tasks? *arXiv preprint arXiv:2107.06383*, 2021. 2
- [44] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 1
- [45] Changyao Tian, Wenhai Wang, Xizhou Zhu, Xiaogang Wang, Jifeng Dai, and Yu Qiao. Vl-ltr: Learning class-wise visual-linguistic representation for long-tailed visual recognition. *European Conference on Computer Vision*, 2022. 2
- [46] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*, 2019. 2
- [47] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR, 2021. 2
- [48] Feng Wang and Huaping Liu. Understanding the behaviour of contrastive loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2495–2504, 2021. 3
- [49] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. *International Conference on Machine Learning*, 2022. 1, 2, 7
- [50] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022. 1, 2

- [51] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022. [2](#)
- [52] Zhecan Wang, Noel Codella, Yen-Chun Chen, Luwei Zhou, Jianwei Yang, Xiyang Dai, Bin Xiao, Haoxuan You, Shih-Fu Chang, and Lu Yuan. Clip-td: Clip targeted distillation for vision-language tasks. *arXiv preprint arXiv:2201.05729*, 2022. [2](#)
- [53] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. *International Conference on Learning Representations*, 2022. [1](#), [2](#)
- [54] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021. [7](#)
- [55] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. Technical Report CNS-TR-201, Caltech, 2010. [3](#)
- [56] Kan Wu, Jinnian Zhang, Houwen Peng, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Tinyvit: Fast pretraining distillation for small vision transformers. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXI*, pages 68–85. Springer, 2022. [2](#)
- [57] Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3485–3492, 2010. [3](#)
- [58] Le Xue, Mingfei Gao, Chen Xing, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Nieves, and Silvio Savarese. Ulip: Learning unified representation of language, image and point cloud for 3d understanding. *arXiv preprint arXiv:2212.05171*, 2022. [1](#)
- [59] Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. Cpt: Colorful prompt tuning for pre-trained vision-language models. *arXiv preprint arXiv:2109.11797*, 2021. [2](#)
- [60] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. [1](#), [2](#)
- [61] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Unified vision and language prompt learning. *arXiv preprint arXiv:2210.07225*, 2022. [2](#)
- [62] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017. [3](#)
- [63] Renrui Zhang, Rongyao Fang, Wei Zhang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free clip-adapter for better vision-language modeling. *arXiv preprint arXiv:2111.03930*, 2021. [9](#)
- [64] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16793–16803, 2022. [2](#)
- [65] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [2](#)
- [66] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [2](#)