

Efficient and Versatile Robust Fine-Tuning of Zero-shot Models

Sungyeon Kim¹  Boseung Jeong¹  Donghyun Kim²  Suha Kwak^{1,3} 

¹ Department of Computer Science and Engineering, POSTECH, Korea

² Department of Artificial Intelligence, Korea University, Korea

³ Graduate School of Artificial Intelligence, POSTECH, Korea

¹{sungyeon.kim,boseung01,suha.kwak}@postech.ac.kr, ²d_kim@korea.ac.kr

<http://cvlab.postech.ac.kr/research/R-Adapter>

Abstract. Large-scale image-text pre-trained models enable zero-shot classification and provide consistent accuracy across various data distributions. Nonetheless, optimizing these models in downstream tasks typically requires fine-tuning, which reduces generalization to out-of-distribution (OOD) data and demands extensive computational resources. We introduce Robust Adapter (R-Adapter), a novel method for fine-tuning zero-shot models to downstream tasks while simultaneously addressing both these issues. Our method integrates lightweight modules into the pre-trained model and employs novel self-ensemble techniques to boost OOD robustness and reduce storage expenses substantially. Furthermore, we propose MPM-NCE loss designed for fine-tuning on vision-language downstream tasks. It ensures precise alignment of multiple image-text pairs and discriminative feature learning. By extending the benchmark for robust fine-tuning beyond classification to include diverse tasks such as cross-modal retrieval and open vocabulary segmentation, we demonstrate the broad applicability of R-Adapter. Our extensive experiments demonstrate that R-Adapter achieves state-of-the-art performance across a diverse set of tasks, tuning only 13% of the parameters of the CLIP encoders.

Keywords: robust fine-tuning; parameter-efficient fine-tuning; self-ensemble

1 Introduction

The emergence of large-scale models pre-trained jointly on image and text data [29, 39, 50] brings a paradigm shift in the field of computer vision. By aligning the embeddings of extensive image-text pairs, these models enable zero-shot inference and show a remarkable ability to generalize across diverse data distributions. Despite their impressive performance in a zero-shot context, they do not measure up to supervised learning models [51, 65], necessitating fine-tuning to unlock their full capabilities. While conventional full fine-tuning enhances task-specific performance, it introduces two major challenges: **1)** Full fine-tuning compromises the ability of the model to generalize to out-of-distribution (OOD) data, crucial for real-world applications where data variability is unpredictable. **2)** It

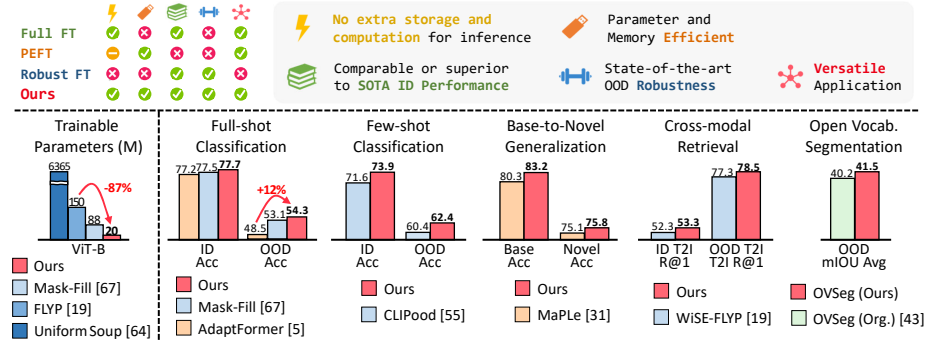


Fig. 1: We present Robust Adapter (R-Adapter), which combines the strengths of robust fine-tuning and parameter-efficient fine-tuning (PEFT). R-Adapter improves parameter and memory efficiency compared to existing robust fine-tuning (*e.g.*, Mask-fill [67], ModelSoup [64]) while being more robust compared to existing PEFT (*e.g.*, AdaptFormer [5], MaPLE [31]). Unlike most of existing robust fine-tuning, our method can apply to a wide range of tasks, and consistently outperforms current best methods on diverse tasks in both in-distribution (ID) and out-of-distribution (OOD).

demands substantial computational resources, memory, and storage, which is impractical given the growing size of large pre-trained models.

Recently, several fine-tuning approaches have been proposed to address these challenges. *Robust fine-tuning* [19, 35, 64, 65, 67] aims to fine-tune zero-shot models while preserving their robustness to OOD, and *Parameter-Efficient Fine-Tuning* (PEFT) [5, 26, 27, 31, 46, 49, 72] updates only a small set of parameters while keeping pre-trained parameters frozen. However, each approach addresses only one of the challenges while still falling short on the other. As shown in Fig. 1, existing robust fine-tuning methods still require tuning the entire model, making training expensive. Moreover, they have only targeted classification tasks, thus often training solely image encoder and excluding zero-shot inference capabilities from the model. On the other hand, PEFT significantly lags in performance compared to robust fine-tuning under distribution shifts. Their critical shortcomings highlight the need for new fine-tuning methods that *simultaneously address both challenges tackled by robust fine-tuning and PEFT separately*.

This paper presents **Robust Adapter (R-Adapter)**, a novel fine-tuning method for improving the robustness of PEFT while enhancing the efficiency of robust fine-tuning. Building upon the adapter-tuning approach [5, 46], where extra lightweight modules are added to a pre-trained model, R-Adapter incorporates novel self-ensemble strategies to enhance OOD robustness.

We take inspiration from the robustness gain observed when averaging multiple models in the weight-space [64, 65], yet implement this strategy within a single model via a unique way. This approach strikes a good balance between task-specific performance and robustness against distribution shifts, and at the same time significantly reduces storage costs. Specifically, R-Adapter achieves this through three self-ensemble techniques. It randomly drops the adapter mod-

ule, thereby dynamically generating and ensemble different subnetworks combining both the adapter and pre-trained layers in various configurations. Additionally, we accumulate adapter weights to form a temporal ensemble that captures all models derived throughout the learning process. Moreover, by re-scaling the weights of the adapter and integrating it into the pre-trained layer via re-parametrization, we enable a seamless linear interpolation between the weights of the pre-trained and fine-tuned models without two separate models.

Additionally, we propose the **Multi-Positive Margin NCE (MPM-NCE)** loss function designed for effective fine-tuning on vision-language downstream tasks. These tasks often involve intricate relations where multiple images can correspond to the same text, and vice versa. Unlike traditional contrastive loss, *i.e.*, InfoNCE [48, 57], which takes single positive pairs and therefore often leads to semantic mismatches in these relations, MPM-NCE accounts for multiple positive pairs and thus promotes more precise alignment across various image-text pairs. Moreover, MPM-NCE introduces an angular margin to penalize negative pairs, enabling the model to learn highly discriminative features critical for downstream tasks. Consequently, the proposed loss leads to significant improvement in task-specific performance, offering benefits in both ID and OOD contexts.

Our method enables zero-shot inference after fine-tuning, extending its applicability beyond image classification tasks to a wide range of applications. To show its versatility, we present a new evaluation benchmark for robust fine-tuning that includes five tasks: image classification tasks under three scenarios, cross-modal retrieval, and open-vocabulary segmentation. Extensive experiments demonstrate that our method achieves superior performance under distribution shift while using fewer parameters compared to existing robust fine-tuning and PEFT methods. The main contribution of this paper is four-fold:

- We introduce an efficient and versatile framework for robust fine-tuning that incorporates the strengths of both PEFT and robust fine-tuning. To the best of our knowledge, it is the first method to make the best of both worlds.
- We propose **R-Adapter** with self-ensemble techniques enabling weight-space ensemble using a single model with adapters. These techniques enhance robustness while reducing storage costs, as it does not need multiple models.
- We develop **MPM-NCE** loss tailored for fine-tuning, utilizing multiple positive pairs and introducing an angular margin. This loss ensures precise alignment of multiple image-text pairs and discriminative feature learning.
- For the first time, we extend the benchmark for robust fine-tuning beyond image classification to include tasks such as cross-modal retrieval and open vocabulary segmentation, allowing us to assess the broad applicability. As shown in Fig. 1, our method achieves state-of-the-art performance on diverse tasks while tuning only 13% of CLIP encoder parameters.

2 Related Work

Robust Fine-tuning. In the conventional practice of leveraging pre-trained models, linear probing or full fine-tuning are commonly used methods for fine-tuning pre-trained models. Kumar et al. [35] show that while fine-tuning achieves

higher accuracy on in-distribution (ID) data, it can distort pre-trained knowledge, reducing out-of-distribution (OOD) accuracy. To mitigate this, a two-step process involving linear probing followed by full fine-tuning has been suggested. Following this paradigm, ensembling-based robust fine-tuning approaches have been proposed in [64, 65]. WiSE-FT [65] ensembles weights of pre-trained fine-tuned models, improving accuracy on both ID and OOD data. FLYP [19] reuses the same contrastive formulation from pre-training for fine-tuning. Mask-Fill [67] promotes consistency between fine-tuned and pre-trained models on counterfactual samples. However, these require full fine-tuning or additional forward/backward passes, leading to high memory and computational demands. Given the substantial size of the foundation models, we aim to develop efficient and fast adaptation methods while improving ID and OOD accuracy. While earlier work primarily focuses on image classification tasks, we extend our investigation to a broader range of tasks, showing the versatility of our approach.

Parameter-Efficient Fine-Tuning. In the context of ever-growing model sizes, fine-tuning large-scale models for various downstream tasks presents a significant challenge, demanding substantial memory and computational resources. To solve this issue, PEFT has been proposed [5, 21, 26, 27, 30, 42, 49, 52]. These methods selectively update a limited portion of trainable parameters, while keeping pre-trained parameters frozen. The concept of low-rank adaptation [13, 27] is introduced to provide an approximation for the parameter update. Several methods only update additional learnable tokens during fine-tuning [37, 40, 56, 63, 71, 72] while freezing all the parameters. It is feasible to incorporate lightweight adapter modules [5, 26, 33, 46, 49] and only update these modules during fine-tuning. However, naïvely using additional learnable tokens and adapters could increase inference costs. RepAdapter [46] proposes a re-parameterization trick for adapters and achieves zero additional cost during inference. We propose R-Adapter which employs PEFT for efficient adaptation of large models to diverse downstream tasks and enhancing ID performance and OOD robustness.

Contrastive Learning. Contrastive loss has been explored in various fields including self-supervised learning [6, 22, 66], vision-language pre-training [29, 50], supervised learning [20, 32], metric learning [10, 34, 57], image captioning [11, 54], etc. Contrastive learning trains a model to differentiate between similar (positive) and dissimilar (negative) data sample pairs. Recently, contrastive learning on web-crawled image-caption data [50] has shown significant gains in zero-shot classification and domain robustness. FLYP [19] proposes a fine-tuned approach using the same contrastive learning formulation for image classification with class prompt templates, but this can cause class collision issues between the same positive classes. To address this, we introduce MPM-NCE which leverages multiple positive relations, considering the characteristics of downstream tasks.

3 Proposed Method

Our method is compatible with various zero-shot models [29, 39], but our research primarily centers on the most renowned model, CLIP [50]. In this section, we

first revisit the CLIP encoders [50] and their pre-training scheme (Sec. 3.1). Next, we define the problem setup (Sec. 3.2). Then our R-Adapter (Sec. 3.3) and MPM-NCE loss (Sec. 3.4) are introduced.

3.1 Preliminary

CLIP Encoders. CLIP consists of two encoders for extracting features from image and text, respectively. Each encoder is composed of a series of Transformer layers [61], each of which consists of Multi-Head Attention (MHA), Layer Normalization (LN), and Feed-Forward Network (FFN). Specifically, the l -th Transformer layer is formulated as follows:

$$\begin{aligned}\bar{X}_l &= \text{MHA}(\text{LN}(X_{l-1})) + X_{l-1}, \\ X_l &= \text{FFN}(\text{LN}(\bar{X}_l)) + \bar{X}_l.\end{aligned}\tag{1}$$

MHA involves k -head self-attention operations on queries, keys, and values, achieved via independent linear projections of the input; it is formulated by

$$\begin{aligned}\text{MHA}(X) &= [\text{Attn}^1(X), \dots, \text{Attn}^k(X)]W_O, \\ \text{Attn}^i(X) &= \text{softmax}((XW_Q^i)(XW_K^i)^\top / \sqrt{d_h})(XW_V^i),\end{aligned}\tag{2}$$

where $[\cdot, \cdot]$ denotes concatenation, and d_h is set to d/k . $W_Q^i \in \mathbb{R}^{d \times d_h}$, $W_K^i \in \mathbb{R}^{d \times d_h}$, $W_V^i \in \mathbb{R}^{d \times d_h}$ and $W_O \in \mathbb{R}^{d \times d}$ are linear projection matrices. FFN consists of two linear layers with a non-linear layer in between:

$$\text{FFN}(X) = \sigma(XW_1 + b_1)W_2 + b_2,\tag{3}$$

where $W_1 \in \mathbb{R}^{d \times 4d}$, $W_2 \in \mathbb{R}^{4d \times d}$, $b_1 \in \mathbb{R}^{4d}$, and $b_2 \in \mathbb{R}^d$ are the respective linear projection weights and biases; $\sigma(\cdot)$ denotes the GELU function.

Contrastive Learning. The CLIP encoders are trained to predict which text descriptions correspond to a given set of images and vice versa. This is achieved through contrastive learning using the InfoNCE loss [48], which forces image embeddings and their corresponding text embeddings to be close to each other and farther away from other text embeddings in a batch. Let $f(\cdot)$ and $g(\cdot)$ be the CLIP encoders for image and text, respectively. Given a batch with B image-text pairs $\mathcal{B} = \{(I_1, T_1), \dots, (I_B, T_B)\}$, the loss function is formulated by

$$\mathcal{L}(\mathcal{B}) = - \sum_{i=1}^B \left(\log \frac{e^{f_i \cdot g_i / \tau}}{\sum_{j=1}^B e^{f_i \cdot g_j / \tau}} + \log \frac{e^{f_i \cdot g_i / \tau}}{\sum_{j=1}^B e^{f_j \cdot g_i / \tau}} \right),\tag{4}$$

where $f_i = \frac{f(I_i)}{\|f(I_i)\|_2}$, $g_i = \frac{g(T_i)}{\|g(T_i)\|_2}$, τ denotes a learnable temperature parameter.

3.2 Problem Setup

Our objective is to efficiently fine-tune a vision-language pre-trained model for various downstream tasks while preserving its inherent out-of-distribution

(OOD) generalization capability. While most existing robust-fine tuning methods are limited to classification tasks [65, 67], we *broaden the scope to robustly fine-tune the models for diverse downstream tasks* such as image classification, cross-modal retrieval, and open-vocabulary segmentation.

Given an image-text pre-trained model, the goal is its adaptation using an in-distribution (ID) training dataset $\mathcal{D}_{\mathcal{I}} = \{(I_i, T_i)\}_{i=1}^n$ for the target downstream task, where I denotes an image and T is a text description corresponding to the image. Concurrently, we aim to enhance the performance of the model on an OOD test dataset $\mathcal{D}_{\mathcal{O}} = \{(I_j, T_j)\}_{j=1}^m$. The ID and OOD datasets, $\mathcal{D}_{\mathcal{I}}$ and $\mathcal{D}_{\mathcal{O}}$, are sampled from different probability distributions, $p_{\mathcal{I}}(I, T)$ and $p_{\mathcal{O}}(I, T)$, respectively, exhibiting distribution shift when $p_{\mathcal{I}}(I, T) \neq p_{\mathcal{O}}(I, T)$. In classification tasks, T represents a text description of the target class which is constructed by sampling from a set of predefined templates (*e.g.*, “a photo of a {class}”) [19, 50]. For other vision-language tasks, T could be one of the captions associated with the image I [44, 68].

3.3 Robust Adapter (R-Adapter)

To achieve efficient and robust fine-tuning, we introduce R-Adapter. Our method is grounded in the PEFT framework, which freezes the pre-trained model while tuning a small number of additional learnable parameters. However, a naïve application of this framework in training can incur a significant bias towards in-distribution data (refer to Table 2). Drawing inspiration from observations that ensembles enhance generalizability across a wide range of distributions [28, 65], R-Adapter is designed with three novel self-ensembling strategies to enable robust fine-tuning without adding computational load during training and inference. In the following, we will introduce the design of R-adapter and then describe our three self-ensemble strategies.

Design of R-Adapter. R-Adapter builds upon the adapter-tuning framework where lightweight modules are added to a pre-trained model. Specifically, the adapter modules in R-Adapter adopt the simple version of the Houlsby Adapter [26] removing nonlinear layers and bias. The module is structured as a residual block composed of a weight matrix as follows:

$$h(X) = XW_{\text{adp}} + X, \quad (5)$$

where X means an output of a pre-trained block and $W_{\text{adp}} \in \mathbb{R}^{d \times d}$ is the weight matrix of our adapter. For full-shot learning, we maintain a full-rank structure for W_{adp} to preserve sufficient capacity. In the few-shot learning, we can adopt a bottleneck architecture by decomposing W_{adp} into a product of low-rank matrices BA , where $B \in \mathbb{R}^{d \times r}$, $A \in \mathbb{R}^{r \times d}$, and the rank $r \ll d$. This decomposition avoids over-parameterization and significantly reduces the number of parameters and computations. We deploy adapters per Transformer layer in *both image and text encoders*, positioned after MHA and FFN layers, as shown in Fig. 2.

Since our adapters lack nonlinearity in between, we can re-parameterize the adapter to remove extra computation overhead from adapter during inference

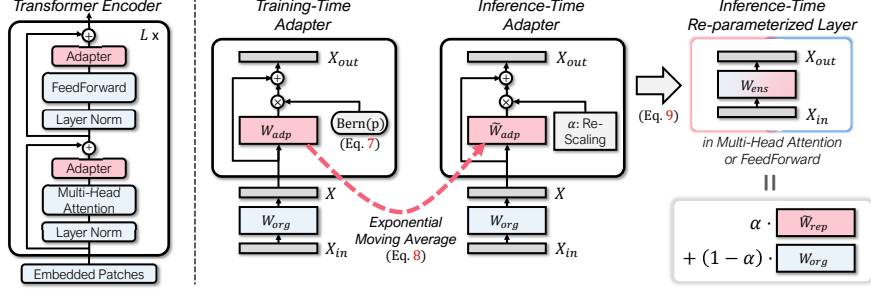


Fig. 2: An overview of R-Adapter. Each adapter is positioned after MHA and FFN layers. R-Adapter stochastically drops the adapters during training. Also, the weights of the adapters are accumulated using an exponential moving average during the training. At the evaluation, these weights are re-scaled by α and then re-parametrized to be integrated into their prior layers, resulting in a weight-space ensemble between the pre-trained layers and the re-parametrized layer without re-scaling.

by integrating it with the closest pre-trained layer [46]. The weights of the pre-trained layer preceding the adapter, denoted by W_{org} , is either W_O from MHA (Eq. 2) or W_2 in FFN (Eq. 3), and the corresponding bias b_{org} is b_2 in FFN (Eq. 3). Given the input to pre-trained layers X_{in} , the re-parametrization is then conducted by

$$\begin{aligned} h(X_{\text{in}}W_{\text{org}} + b_{\text{org}}) &= X_{\text{in}}W_{\text{org}}(W_{\text{adp}} + \mathbf{I}) + b_{\text{org}}W_{\text{adp}} + b_{\text{org}} \\ &= X_{\text{in}}W_{\text{rep}} + b_{\text{rep}}, \end{aligned} \quad (6)$$

where $\mathbf{I} \in \mathbb{R}^{d \times d}$ is the identity matrix, $W_{\text{rep}} = W_{\text{org}}(W_{\text{adp}} + \mathbf{I})$, and $b_{\text{rep}} = b_{\text{org}}(W_{\text{adp}} + \mathbf{I})$.

Dynamic Ensemble by Adapter Dropping. To enhance OOD robustness, R-Adapter employs a dynamic ensemble technique through adapter dropping. During only training, adapter modules are randomly deactivated as follows:

$$h(X) = \frac{\gamma}{1-p} \cdot XW_{\text{adp}} + X, \quad (7)$$

where γ is an independent variable drawn from Bernoulli($1-p$), and p is the drop probability of the adapter dropping. Unlike dropout [58] for feature sparsity or drop-path [36] for model depth reduction, our technique uniquely focuses on randomly disabling adapter layers while consistently supplying pre-trained features. Adapter dropping is not applied during inference, serving to create an ensemble of subnetworks that vary by the combination of both pre-trained and adapter layers. This strategy enables a dynamic ensemble of multiple models that retain both pre-trained knowledge and fine-tuned knowledge simultaneously and thus boost performance both on ID and OOD data. (see Table 2)

Temporal Ensemble by Accumulation. We advance the robustness of the model by incorporating a temporal ensemble strategy through the historical ac-

cumulation of adapter weights. The ensemble captures a broader understanding of the feature space by averaging the weights over multiple iterations during training [3, 28]. The weights of the accumulated adapter \tilde{W}_{adp} are updated via an exponential moving average:

$$\tilde{W}_{\text{adp}} \leftarrow m \cdot \tilde{W}_{\text{adp}} + (1 - m) \cdot W_{\text{adp}}, \quad (8)$$

where $m \in [0, 1]$ is the coefficient that controls the momentum update rate. This procedure is notably *memory-efficient* since only the parameters of adapters are momentum updated, not the parameters of the entire model. In inference time, we utilize the accumulated weights \tilde{W}_{adp} for Eq. 6, thereby produces re-parameterized weight \tilde{W}_{rep} and bias \tilde{b}_{rep} .

Weight-space Ensemble by Re-scaling. Finally, we introduce a strategy that establishes a weight-space ensemble between the pre-trained and fine-tuned layers through re-scaling with re-parameterization. The conventional weight-space ensemble (WiSE-FT) [65] linearly interpolates between the weights of the original pre-trained parameters and the fine-tuned parameters, thus requiring storing both separate models. In contrast, we evolve this concept by employing the re-parameterized weights \tilde{W}_{rep} as the weights of a fine-tuned layer. We streamline the weight-space ensemble within a single model to be implemented simply by re-scaling the weights of the adapter and re-parameterizing them at inference. This process is expressed as follows:

$$\begin{aligned} \underbrace{\alpha \tilde{W}_{\text{rep}} + (1 - \alpha) W_{\text{org}}}_{\text{Weight-space Ensemble}} &= \underbrace{\alpha W_{\text{org}} \tilde{W}_{\text{adp}} + \alpha W_{\text{org}} + (1 - \alpha) W_{\text{org}}}_{\text{Re-scaling}} \\ &= \underbrace{W_{\text{org}} (\alpha \tilde{W}_{\text{adp}} + \mathbf{I})}_{\text{Re-parameterization}} = W_{\text{ens}}, \end{aligned} \quad (9)$$

where W_{ens} denotes the ensembled weights, and α is a re-scaling coefficient. The coefficient α serves as an interpolation factor, adjusting the balance between the original pre-trained weights W_{org} and the adjusted weights of the fine-tuned layer. This technique not only improves accuracy under distribution shifts but also maintains high performance on the ID data. Crucially, unlike WiSE-FT, our method does not require maintaining two separate full models in storage, thus facilitating weight-space ensemble *more storage-efficiently*.

3.4 MPM-NCE Loss for Downstream Task

To enhance learning for downstream tasks, it is crucial to use loss functions that align closely with the characteristics of the tasks. Vision-language tasks often involve multiple correspondences between modalities. For instance, in classification tasks, using different text templates for the same class can result in multiple text descriptions matching a single image, and naturally the reverse is true as well. This situation also occurs in cross-modal retrieval tasks with images and captions. When adapting zero-shot models to new tasks, a common approach is to use the InfoNCE loss used for pre-training. However, this loss is not ideal for

tasks where multiple positive samples exist, as it considers a single positive pair. Moreover, InfoNCE learns the ordering between positive and negative samples, which may not lead to sufficiently discriminative features for downstream tasks.

To address these limitations, we propose MPM-NCE Loss, designed to accommodate the multi-positive nature of these tasks while enhancing the discriminative power of the learned embeddings. This loss function has two pivotal improvements. First, we use soft labels that assign equal probability to multiple positive pairs. The formulation of the soft label is given as follows:

$$\tilde{y}_{ij} = \frac{(1 - \epsilon) \cdot y_{ij}}{|P(i)|} + \frac{\epsilon \cdot (1 - y_{ij})}{B - |P(i)|} \in [0, 1], \quad (10)$$

where $y_{ij} \in \{0, 1\}$ indicates the positive relation between samples i and j , $P(i)$ is the set of positive samples of sample i including itself and ϵ is a label smoothing noise [59]. This soft label ensures the correct alignment of multiple image-text pairs in downstream tasks. Additionally, the soft labels can include ϵ , reducing overfitting risks by introducing a minor perturbation to the labels.

The second improvement is the addition of a margin δ applied to negative pairs. This margin enhances the discrimination of learned features by ensuring that negative pairs are not only distinct but separated by a certain threshold. Incorporating these improvements, our MPM-NCE is formulated as follows:

$$\mathcal{L}(\mathcal{B}) = - \sum_{i,j=1}^B \left(\tilde{y}_{ij} \log \frac{e^{(f_i \cdot g_j + \delta_{ij})/\tau}}{\sum_{k=1}^B e^{(f_i \cdot g_k + \delta_{ik})/\tau}} + \tilde{y}_{ji} \log \frac{e^{(f_j \cdot g_i + \delta_{ji})/\tau}}{\sum_{k=1}^B e^{(f_k \cdot g_i + \delta_{ki})/\tau}} \right), \quad (11)$$

where the temperature τ is set to a constant value of 0.01, and δ_{ij} is 0 for positive relations and δ for the rest. Consequently, MPM-NCE loss encourages the model to correctly align multiple image-text pairs and learn discriminative features, leading to notable improvements in performance under ID and OOD.

4 Experiments

We first demonstrate the robustness of R-Adapter against natural distribution shifts for image classification and its efficiency (Sec. 4.3). We then analyze the effectiveness of proposed components in R-Adapter and MPM-NCE loss, including ensemble techniques and loss, compared to existing approaches and also conduct an ablation study on hyperparameters. Furthermore, we validate the versatility of R-Adapter by extending it to broader tasks such as few-shot classification (Sec. 4.4), cross-modal retrieval (Sec. 4.5), open-vocabulary segmentation (Sec. 4.6), and base-to-novel generalization (in Appendix).

4.1 Datasets

Image Classification. We use ImageNet (IN) [12] as the ID dataset for fine-tuning; we evaluate the robustness of the models on five standard OOD datasets

with different distribution shifts, following prior work [19, 35, 50, 65, 67]: ImageNetV2 (IN-V2) [53], ImageNet-R (IN-R) [24], ImageNet-Sketch (IN-Sketch) [62], ObjectNet [1], and ImageNet-A (IN-A) [25]. Note that these datasets except ObjectNet are also used in a few-shot setting following previous work [46, 55, 71, 72].

Cross-Modal Retrieval. We utilize two standard benchmarks for image-text cross-modal retrieval, COCO [44] as ID and Flickr30K [68] as OOD. For these two datasets, each image is associated with the corresponding five captions.

Open-Vocabulary Segmentation. Following our baseline method [43], we train a CLIP model on the COCO Captions dataset [7] and test it on several OOD benchmarks: ADE20K [70] (A-150 and A-847 category versions), Pascal Context [47] (PC-59 and PC-459 category versions), and Pascal VOC [16].

4.2 Implementation Details

Network Architectures. We adopt the pre-trained CLIP models from OpenAI [50] with four different sizes of image encoder, ViT-B/32, ViT-B/16, ViT-L/14, and ViT-L/14@336px [15].

Network Optimization. Our model is trained using AdamW without weight decay for 10 epochs, except for open vocabulary segmentation which is trained for 5 epochs following previous work [43]. The initial learning rate is set to $5e-4$, using a cosine scheduling with 500 warm-up steps. We closely follow the settings in [19] for full-shot classification, [55] for few-shot classification, and [43] for open-vocabulary segmentation. More details are in the appendix.

Hyperparameters. The drop probability p is set to 0.2. The momentum update rate m in Eq. 8 is set to 0.999. The margin δ in Eq. 11 is 0.05. For classification tasks, following the WiSE-FT [65], we use the re-scaling coefficient α in Eq. 9 of 0.5. For cross-modal retrieval and open vocabulary segmentation tasks, we set α to its optimal values of 0.8 and 0.4, respectively. We set the smoothing coefficient ϵ in Eq. 10 to 0.05 for classification, and 0 for other tasks.

4.3 ImageNet Classification Under Distribution Shifts

Main Results. We compare our method with zero-shot, conventional fine-tuning approach, and previous robust fine-tuning methods, WiSE-FT [65], LP-FT [35], Model Soup [64], FLYP [19], and Mask-Fill [67] on the in-distribution (ID) dataset and five out-of-distribution (OOD) datasets. We report the performance of WiSE-FT with the default mixing coefficient of 0.5. We take the *uniform soup* as the default method of [64]. The results and the number of trainable parameters are summarized in Table 1. Specifically, our method improves the previous state of the art by a significant margin as 1.2%p and 1.5%p in terms of OOD avg. with CLIP ViT-B/32 and CLIP ViT-B/16, respectively; even though our method only requires much less tunable parameters (20.5M) than the others (>80M). Moreover, our method scales efficiently to the CLIP ViT-L/14@336px model, showing a notable 2%p improvement in OOD performance over WiSE-FT. While the Uniform Soup achieves superior results on IN

Table 1: Top-1 accuracy of models with different robust fine-tuning on ImageNet (ID) and OOD datasets. “OOD avg” is the average accuracy across the five OOD datasets. Entries in **green** indicate fewer parameters than full fine-tuning, and **red** use more.

Methods	Trainable Params (M)	ID IN	Out-Of-Distribution (OOD)					
			OOD avg	IN-V2	IN-R	IN-Sketch	ObjectNet	IN-A
<i>CLIP ViT-B/32</i>								
Zero-Shot [50]	\times	63.4	48.7	55.9	69.3	42.3	44.5	31.4
Fine-Tuning (FT)	88.4	75.9	44.2	64.7	57.0	39.8	39.5	20.0
WiSE-FT [65]	88.4	76.6	52.4	66.6	<u>70.2</u>	47.1	46.3	31.9
Uniform Soup [64]	6364.8	80.0	51.6	68.6	66.6	<u>47.7</u>	46.1	29.2
Mask-Fill [67]	88.4	77.5	<u>53.1</u>	67.1	69.7	46.9	<u>48.0</u>	<u>33.8</u>
Ours	20.5	<u>77.7</u>	54.3	<u>67.7</u>	70.8	47.8	49.7	35.6
<i>CLIP ViT-B/16</i>								
Zero-Shot [50]	\times	68.3	58.4	61.9	77.6	48.3	54.0	50.1
Fine-Tuning (FT)	86.7	80.7	52.8	70.4	64.0	45.1	49.1	35.2
LP-FT [35]	86.7	81.7	60.3	72.1	73.5	50.3	58.2	47.6
WiSE-FT [65]	86.7	81.7	63.0	72.8	<u>78.7</u>	53.9	57.3	52.2
FLYP [19]	149.6	<u>82.6</u>	60.2	73.0	71.4	48.1	58.7	49.6
WiSE-FLYP [19]	149.6	82.9	63.1	<u>73.5</u>	76.0	53.0	60.8	52.3
Mask-Fill [67]	86.7	82.4	<u>63.3</u>	73.4	78.1	53.4	57.9	<u>53.5</u>
Ours	20.5	82.0	64.8	73.6	79.1	53.9	<u>59.7</u>	57.5
<i>CLIP ViT-L/14@336px</i>								
WiSE-FT [65]	305.1	86.8	<u>76.9</u>	<u>79.5</u>	<u>89.4</u>	64.7	<u>71.1</u>	<u>79.9</u>
Ours	64.5	86.8	78.9	79.6	89.9	<u>64.1</u>	73.3	82.4

Table 2: Ablation study on key components of our method and comparison with the other adapter-tuning methods using full-rank structure. The experiments are performed on the ImageNet classification with ViT-B/32. The last row (**E10**) corresponds to our default configuration. DO: Dropout in Adapters. DP: Drop-path in pre-trained layers. AD: Adapter Dropping. AC: Accumulation. RS: Re-scaling. LS: Label Smoothing.

Exp No.	Adapter Design (w/ Full-Rank)	Regularization					Loss			Accuracy	
		DO	DP	AD	AC	RS	InfoNCE	MPM-NCE	LS	ID	OOD avg
B1	AdaptFormer [5]	✓	X	X	X	X	✓	X	X	77.2	48.5
B2	RepAdapter [46]	✓	✓	X	X	X	✓	X	X	77.2	48.3
E0	R-Adapter (Ours)	X	X	X	X	X	✓	X	X	77.5 (↑ 0.0)	47.7 (↑ 0.0)
E1		✓	X	X	X	X	✓	X	X	77.6 (↑ 0.1)	48.7 (↑ 1.1)
E2		X	✓	X	X	X	✓	X	X	77.4 (↓ 0.1)	47.9 (↑ 0.2)
E3		X	X	✓	X	X	✓	X	X	77.8 (↑ 0.3)	49.6 (↑ 1.9)
E4		X	X	X	✓	X	✓	X	X	77.4 (↓ 0.1)	47.8 (↑ 0.1)
E5		X	X	X	X	✓	✓	X	X	76.5 (↓ 1.0)	53.5 (↑ 5.8)
E6		X	X	✓	✓	X	✓	X	X	77.9 (↑ 0.4)	49.9 (↑ 2.2)
E7		X	X	✓	✓	✓	✓	X	X	76.6 (↓ 0.9)	53.7 (↑ 6.0)
E8		X	X	✓	✓	✓	✓	✓	✓	76.9 (↓ 0.6)	54.0 (↑ 6.3)
E9		X	X	✓	✓	✓	X	✓	X	77.5 (↑ 0.0)	53.9 (↑ 6.2)
E10		X	X	✓	✓	✓	X	✓	✓	77.7 (↑ 0.2)	54.3 (↑ 6.6)

and IN-V2, it involves a complex ensemble of fine-tuned models, leading to increased computational and resource demands. In contrast, our method offers a cost-efficient approach to enhancing robustness, as evidenced by the pronounced gains observed in the most distribution-shifted datasets, IN-A and IN-R.

Effectiveness of Key Components. In our ablation study, we evaluate the impact of key components and compare our method with AdaptFormer and

Table 3: Ablation study on hyperparameters on the ImageNet classification task with ViT-B/32. The last column shows the average accuracy across the five OOD datasets. gray corresponds to our default setting. “w/ SP” indicates the considering single positive without soft labels as InfoNCE, but employing a margin δ of 0.05.

(a) Rank of Adapter				(b) Loss Variations			(c) p in Eq. 7			(d) m in Eq. 8		
Rank	#Params	ID	OOD	Loss	ID	OOD	p	ID	OOD	m	ID	OOD
4	0.25M	72.5	51.7	$\delta = 0$	77.1	54.0	0	77.6	53.3	0	77.8	54.0
8	0.49M	73.4	52.4	$\delta = 0.02$	77.5	54.3	0.1	77.9	54.0	0.9	77.8	54.3
16	0.98M	74.5	52.5	$\delta = 0.05$	77.7	54.3	0.2	77.7	54.3	0.99	77.8	54.3
128	7.84M	76.7	53.7	$\delta = 0.1$	77.8	53.8	0.3	77.6	54.4	0.999	77.7	54.3
Full	20.45M	77.7	54.3	w/ SP	77.2	47.0	0.5	77.1	54.2	0.9999	77.0	54.3

RepAdapter, both trained with the FLYP scheme, as shown in Table 2. Despite using regularization techniques like dropout (DO) and drop-path (DP), these methods perform poorly in out-of-distribution (OOD) settings, revealing the limitations of naively combining PEFT with robust fine-tuning. Our base R-Adapter model (E0) also falls short in OOD accuracy. However, using Adapter Dropping (AD) improves OOD accuracy by 1.9% and in-distribution (ID) accuracy by 0.3% (E1, E2, and E3). Accumulation (AC) and Re-scaling (RS) are crucial for OOD robustness (E4 and E5), with RS boosting OOD performance by 5.8% despite a slight reduction in ID performance. Combining our regularization techniques mitigates this reduction and further enhances OOD accuracy (E6 and E7). MPM-NCE outperforms InfoNCE in both ID and OOD settings by 0.9% and 0.2%, respectively (E7 and E9). While label smoothing (LS) with InfoNCE can reduce ID performance due to semantic misalignments, MPM-NCE with LS improves both ID and OOD performance by maintaining accurate alignment and providing additional regularization (E10). Our default model, the R-Adapter trained with MPM-NCE loss, significantly advances ID performance and OOD robustness over existing adapter techniques (B1, B2, and **E10**).

Effect of Hyperparameters. We investigate the effects of the rank of the adapter module r and various hyperparameters, including the drop probability p in Eq. 7, the momentum update rate m in Eq. 8 and the margin of our loss δ in Eq. 11. Table 3a reveals that increasing the rank of the adapter enhances performance, due to improved model capacity. This result aligns with findings in [4] that more parameters yield better results in data-rich environments. Table 3 shows gradual performance gains with a margin δ up to 0.05, but using a margin with a single positive reduces OOD performance. As shown in Table 3c and 3d, each hyperparameter brings performance improvement compared to when it is set to 0, regardless of specific values. Fig. 3 shows the impact of varying the re-scaling coefficient α in Eq. 9. Compared to WiSE-FT [65], our method shows less sensitivity to changes in α , maintaining superior performance across various settings.




Fig. 3: Performance of our method varying re-scaling coefficient α against WiSE-FT.

Table 4: Top-1 accuracy for adapting CLIP to 16-shot ImageNet classification on ID and OOD datasets. OOD avg is the average accuracy across the four OOD datasets. “ r -Rank” denotes our models with adapters employing low-rank decomposition while “Full-Rank” is no decomposition. All methods adopt CLIP ViT-B/16 as the backbone.

Methods	Trainable Params (M)	ID IN	Out-Of-Distribution (OOD)				
			OOD avg	IN-V2	IN-R	IN-Sketch	IN-A
Zero-Shot [50]	\times	68.3	58.4	61.9	77.6	48.3	50.1
CoOp [72]	> 0.01	71.5	59.3	64.2	75.2	48.0	49.7
CoCoOp [71]	0.03	71.0	59.9	64.1	76.2	48.8	50.6
RepAdapter-T [46]	0.27	71.9	60.4	64.8	76.5	49.3	51.1
CLIPood [55]	86.70	71.6	60.4	64.9	77.2	49.3	50.4
Ours (1-Rank)	0.06	71.7	<u>61.6</u>	65.3	<u>78.6</u>	<u>50.3</u>	52.3
Ours (4-Rank)	0.25	72.0	<u>61.6</u>	65.1	<u>78.6</u>	50.0	52.6
Ours (8-Rank)	0.49	<u>72.4</u>	<u>61.6</u>	<u>65.7</u>	<u>78.6</u>	49.8	<u>52.4</u>
Ours (Full-Rank)	20.45	73.9	62.4	67.0	79.1	51.2	52.3

Table 5: Cross-modal retrieval performance on the COCO (5K test set) and Flickr30K datasets in Recall at K (R@K). B and L denote the use of 12-layer and 24-layer transformer encoders, respectively. FLYP $_L$ training has failed due to memory constraints.

Methods	Training Dataset	COCO						Flickr30K					
		Text-to-Img			Img-to-Text			Text-to-Img			Img-to-Text		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
Unicoder-VL $_B$ [38]	Same as Test	46.7	76.0	85.3	62.3	87.1	92.8	71.5	90.9	94.9	86.2	96.3	99.0
Uniter $_L$ [8]	Same as Test	52.9	79.9	88.0	65.7	88.6	93.8	75.6	94.1	96.8	87.3	98.0	99.2
VILLA $_L$ [17]	Same as Test	-	-	-	-	-	-	76.3	94.2	96.8	87.9	97.5	98.8
Oscar $_L$ [41]	Same as Test	<u>57.5</u>	<u>82.8</u>	89.8	<u>73.5</u>	<u>92.2</u>	<u>96.0</u>	-	-	-	-	-	-
ERNIE-ViL $_L$ [69]	Same as Test	-	-	-	-	-	-	76.7	93.6	96.4	88.7	98.0	99.2
CLIP $_B$ [50]	\times	33.1	58.4	69.0	52.5	76.7	84.7	62.1	85.7	91.9	82.2	96.6	99.0
FLYP $_B$ [19]	COCO	51.7	77.6	86.0	69.7	88.7	93.9	76.3	94.2	96.8	89.0	98.2	99.5
WiSE-FLYP $_B$ [19]	COCO	52.3	77.7	85.8	70.3	89.3	94.0	77.3	94.6	97.2	91.0	98.6	99.3
Ours $_B$	COCO	53.5	79.0	87.0	71.6	90.2	94.4	<u>78.4</u>	<u>95.0</u>	<u>97.5</u>	<u>91.9</u>	<u>98.7</u>	99.6
Ours $_L$	COCO	58.1	58.1	<u>89.0</u>	75.8	92.9	96.2	83.4	96.9	98.6	95.9	99.4	99.6

4.4 Few-Shot ImageNet Classification

We investigate the robustness of our model when training images are limited, focusing on 16-shot few-shot classification on both ID and OOD datasets. We compare our model with the existing PEFT methods [46, 71, 72] and robust fine-tuning techniques [55]. As shown in Table 4, full-rank R-adapter outperforms the state of the art [55] on all datasets, despite requiring four times fewer trainable parameters. Furthermore, our model with a rank-1 adapter surpasses CoOp and CoCoOp by 2.3% and 1.7% in average OOD top-1 accuracy, with a similar number of tunable parameters. This demonstrates that our method maintains strong generalization on OOD datasets even with extremely minimal parameters.

4.5 Cross-Modal Retrieval

We evaluate our model on COCO [44] and Flickr30K [68] for cross-modal retrieval, where the model is only fine-tuned on the COCO dataset. Since most

Table 6: Comparison of mIoU results between the OVSeg fine-tuned with our method and existing open-vocabulary segmentation models. Note that OVSeg (Org.) is trained in two stages, starting with full CLIP model fine-tuning followed by mask prompt tuning, whereas OVSeg (Ours) involves single-stage adapter training.

Methods	Backbone	A-847	PC-459	A-150	A-59	PAS-20
ZegFormer [14]	R-50 [23]	-	-	16.4	-	80.7
OpenSeg [18]	R-101 [23]	4.0	6.5	15.3	36.9	60.0
LSeg+ [18]	Eff-B7 [60]	3.8	7.8	18.0	46.5	-
OpenSeg [18]	Eff-B7 [60]	6.3	9.0	21.1	42.1	-
OVSeg (Org.) [43]	Swin-B [45]	9.0	12.4	29.6	55.7	94.5
OVSeg (Ours)	Swin-B [45]	10.3 (↑ 1.3)	12.8 (↑ 0.4)	29.5 (↓ 0.1)	58.4 (↑ 2.7)	96.4 (↑ 1.9)

previous methods for robust fine-tuning are limited to the classification task only, we compare our method with FLYP [19] from our re-implementation. We further compare ours with supervised specialists [8, 17, 38, 41, 69]. As shown in Table 5, our method outperforms FLYP and its weight-ensemble (WiSE-FLYP) in terms of all evaluation metrics both on COCO and Flickr30K. Moreover, our method using CLIP ViT-L/14 surpasses the supervised specialists that have a similar size and are trained on both datasets, respectively. Note that although we do not utilize Flickr30K in training, it outperforms supervised methods.

4.6 Open-Vocabulary Segmentation

Our method can enhance open-vocabulary segmentation performance when used for fine-tuning the CLIP model within the OVSeg framework [43]. We use full-rank adapters in the CLIP image and text encoders of OVSeg, fine-tuning them while keeping the pre-trained encoders frozen. Following the OVSeg setup, we employ MaskFormer [9] with Swin-B [45] as a mask proposal network, trained on the COCO-Stuff dataset [2]. The masked image classification model using CLIP ViT-L/14 is trained on masked images from COCO Captions [7] with our method and evaluated on five unseen datasets, as shown in Table 6. Compared to the original OVSeg model, OVSeg model fine-tuned with our method shows significant performance improvements, with mIoU increases of 1.3%, 0.4%, 2.6%, and 1.9% on A-847, PC-459, PC-59, and PAS-20, respectively. These results confirm that our method enhances generalization for unseen classes, showing it to be a promising approach for the open-vocabulary segmentation task.

5 Conclusion

We have introduced a novel approach for fine-tuning image-text models, emphasizing parameter efficiency and robustness to out-of-distribution data. By incorporating R-Adapter with self-ensembling techniques and MPM-NCE loss function, our method surpasses existing methods in robustness and efficiency. Moreover, its adaptability is confirmed by its successful application to diverse tasks. We believe that our method will greatly facilitate making the fine-tuning of zero-shot models much more broadly and easily accessible.

Acknowledgements

This work was supported by NRF grants (NRF-2021R1A2C3012728–30%, NRF-2018R1A5A1060031–30%, RS-2024-00341514–25%) and IITP grants (RS-2019-II191906–10%, Artificial Intelligence Graduate School Program - POSTECH, RS-2019-II190079–5%, Artificial Intelligence Graduate School Program - Korea University) funded by Ministry of Science and ICT, Korea.

References

1. Barbu, A., Mayo, D., Alverio, J., Luo, W., Wang, C., Gutfreund, D., Tenenbaum, J., Katz, B.: Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *Advances in neural information processing systems* **32** (2019) 10
2. Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018) 14
3. Cha, J., Chun, S., Lee, K., Cho, H.C., Park, S., Lee, Y., Park, S.: Swad: Domain generalization by seeking flat minima. *Proc. Neural Information Processing Systems (NeurIPS)* (2021) 8
4. Chen, G., Liu, F., Meng, Z., Liang, S.: Revisiting parameter-efficient tuning: Are we really there yet? *arXiv preprint arXiv:2202.07962* (2022) 12
5. Chen, S., Ge, C., Tong, Z., Wang, J., Song, Y., Wang, J., Luo, P.: Adaptformer: Adapting vision transformers for scalable visual recognition. In: *Proc. Neural Information Processing Systems (NeurIPS)* (2022) 2, 4, 11
6. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297* (2020) 4
7. Chen, X., Fang, H., Lin, T.Y., Vedantam, R., Gupta, S., Dollár, P., Zitnick, C.L.: Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325* (2015) 10, 14
8. Chen, Y.C., Li, L., Yu, L., El Kholy, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Universal image-text representation learning. In: *Proc. European Conference on Computer Vision (ECCV)* (2020) 13, 14
9. Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems* (2021) 14
10. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2005) 4
11. Dai, B., Lin, D.: Contrastive learning for image captioning. In: *Advances in Neural Information Processing Systems*. pp. 898–907 (2017) 4
12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2009) 9
13. Dettmers, T., Pagnoni, A., Holtzman, A., Zettlemoyer, L.: Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314* (2023) 4
14. Ding, J., Xue, N., Xia, G.S., Dai, D.: Decoupling zero-shot semantic segmentation. In: *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2022) 14

15. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: Proc. International Conference on Learning Representations (ICLR) (2021) 10
16. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision (IJCV)* (2010) 10
17. Gan, Z., Chen, Y.C., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-scale adversarial training for vision-and-language representation learning. *Advances in Neural Information Processing Systems* (2020) 13, 14
18. Ghiasi, G., Gu, X., Cui, Y., Lin, T.Y.: Scaling open-vocabulary image segmentation with image-level labels. In: Proc. European Conference on Computer Vision (ECCV) (2022) 14
19. Goyal, S., Kumar, A., Garg, S., Kolter, Z., Raghunathan, A.: Finetune like you pretrain: Improved finetuning of zero-shot vision models. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023) 2, 4, 6, 10, 11, 13, 14
20. Graf, F., Hofer, C., Niethammer, M., Kwitt, R.: Dissecting supervised contrastive learning. In: *International Conference on Machine Learning*. pp. 3821–3830. PMLR (2021) 4
21. He, J., Zhou, C., Ma, X., Berg-Kirkpatrick, T., Neubig, G.: Towards a unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366* (2021) 4
22. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020) 4
23. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016) 14
24. Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.: The many faces of robustness: A critical analysis of out-of-distribution generalization. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2021) 10
25. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 10
26. Hounsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S.: Parameter-efficient transfer learning for nlp. In: Proc. International Conference on Machine Learning (ICML). PMLR (2019) 2, 4, 6
27. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021) 2, 4
28. Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D., Wilson, A.G.: Averaging weights leads to wider optima and better generalization (2018) 6, 8
29. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: Proc. International Conference on Machine Learning (ICML) (2021) 1, 4
30. Jia, M., Tang, L., Chen, B.C., Cardie, C., Belongie, S., Hariharan, B., Lim, S.N.: Visual prompt tuning. In: Proc. European Conference on Computer Vision (ECCV). Springer (2022) 4

31. Khattak, M.U., Rasheed, H., Maaz, M., Khan, S., Khan, F.S.: Maple: Multi-modal prompt learning. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023) [2](#)
32. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: Proc. Neural Information Processing Systems (NeurIPS) (2020) [4](#)
33. Kim, S., Kim, D., Kwak, S.: Universal metric learning with parameter-efficient transfer learning. arXiv preprint arXiv:2309.08944 (2023) [4](#)
34. Kim, S., Kim, D., Cho, M., Kwak, S.: Embedding transfer with label relaxation for improved metric learning. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
35. Kumar, A., Raghunathan, A., Jones, R., Ma, T., Liang, P.: Fine-tuning can distort pretrained features and underperform out-of-distribution. arXiv preprint arXiv:2202.10054 (2022) [2](#), [3](#), [10](#), [11](#)
36. Larsson, G., Maire, M., Shakhnarovich, G.: Fractalnet: Ultra-deep neural networks without residuals. In: International Conference on Learning Representations (2016) [7](#)
37. Lester, B., Al-Rfou, R., Constant, N.: The power of scale for parameter-efficient prompt tuning. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (2021) [4](#)
38. Li, G., Duan, N., Fang, Y., Gong, M., Jiang, D.: Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In: Proc. AAAI Conference on Artificial Intelligence (AAAI) (2020) [13](#), [14](#)
39. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: Proc. International Conference on Machine Learning (ICML). PMLR (2022) [1](#), [4](#)
40. Li, X.L., Liang, P.: Prefix-tuning: Optimizing continuous prompts for generation. arXiv preprint arXiv:2101.00190 (2021) [4](#)
41. Li, X., Yin, X., Li, C., Zhang, P., Hu, X., Zhang, L., Wang, L., Hu, H., Dong, L., Wei, F., et al.: Oscar: Object-semantics aligned pre-training for vision-language tasks. In: Proc. European Conference on Computer Vision (ECCV) (2020) [13](#), [14](#)
42. Lian, D., Zhou, D., Feng, J., Wang, X.: Scaling & shifting your features: A new baseline for efficient model tuning. Advances in Neural Information Processing Systems **35**, 109–123 (2022) [4](#)
43. Liang, F., Wu, B., Dai, X., Li, K., Zhao, Y., Zhang, H., Zhang, P., Vajda, P., Marculescu, D.: Open-vocabulary semantic segmentation with mask-adapted clip. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023) [10](#), [14](#)
44. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: Proc. European Conference on Computer Vision (ECCV) (2014) [6](#), [10](#), [13](#)
45. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proc. IEEE International Conference on Computer Vision (ICCV) (2021) [14](#)
46. Luo, G., Huang, M., Zhou, Y., Sun, X., Jiang, G., Wang, Z., Ji, R.: Towards efficient visual adaption via structural re-parameterization. arXiv preprint arXiv:2302.08106 (2023) [2](#), [4](#), [7](#), [10](#), [11](#), [13](#)
47. Mottaghi, R., Chen, X., Liu, X., Cho, N.G., Lee, S.W., Fidler, S., Urtasun, R., Yuille, A.: The role of context for object detection and semantic segmentation in the wild. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014) [10](#)

48. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018) [3](#), [5](#)
49. Pfeiffer, J., Kamath, A., Rücklé, A., Cho, K., Gurevych, I.: Adapterfusion: Non-destructive task composition for transfer learning. arXiv preprint arXiv:2005.00247 (2020) [2](#), [4](#)
50. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: Proc. International Conference on Machine Learning (ICML) (2021) [1](#), [4](#), [5](#), [6](#), [10](#), [11](#), [13](#)
51. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: Proc. International Conference on Machine Learning (ICML) (2021) [1](#)
52. Rebuffi, S.A., Bilen, H., Vedaldi, A.: Learning multiple visual domains with residual adapters. In: Proc. Neural Information Processing Systems (NeurIPS) (2017) [4](#)
53. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do imagenet classifiers generalize to imagenet? In: Proc. International Conference on Machine Learning (ICML). PMLR (2019) [10](#)
54. Sarto, S., Barraco, M., Cornia, M., Baraldi, L., Cucchiara, R.: Positive-augmented contrastive learning for image and video captioning evaluation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6914–6924 (2023) [4](#)
55. Shu, Y., Guo, X., Wu, J., Wang, X., Wang, J., Long, M.: Clipood: Generalizing clip to out-of-distributions. In: Proc. International Conference on Machine Learning (ICML) (2023) [10](#), [13](#)
56. Smith, J.S., Karlinsky, L., Gutta, V., Cascante-Bonilla, P., Kim, D., Arbelles, A., Panda, R., Feris, R., Kira, Z.: Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. arXiv preprint arXiv:2211.13218 (2022) [4](#)
57. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: Proc. Neural Information Processing Systems (NeurIPS) (2016) [3](#), [4](#)
58. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning Research (JMLR) **15**, 1929–1958 (2014) [7](#)
59. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) [9](#)
60. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: Proc. International Conference on Machine Learning (ICML) (2019) [14](#)
61. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Proc. Neural Information Processing Systems (NeurIPS) (2017) [5](#)
62. Wang, H., Ge, S., Lipton, Z., Xing, E.P.: Learning robust global representations by penalizing local predictive power. Advances in Neural Information Processing Systems **32** (2019) [10](#)
63. Wang, Z., Zhang, Z., Lee, C.Y., Zhang, H., Sun, R., Ren, X., Su, G., Perot, V., Dy, J., Pfister, T.: Learning to prompt for continual learning. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022) [4](#)

64. Wortsman, M., Ilharco, G., Gadre, S.Y., Roelofs, R., Gontijo-Lopes, R., Morcos, A.S., Namkoong, H., Farhadi, A., Carmon, Y., Kornblith, S., et al.: Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In: Proc. International Conference on Machine Learning (ICML) (2022) [2](#), [4](#), [10](#), [11](#)
65. Wortsman, M., Ilharco, G., Kim, J.W., Li, M., Kornblith, S., Roelofs, R., Lopes, R.G., Hajishirzi, H., Farhadi, A., Namkoong, H., et al.: Robust fine-tuning of zero-shot models. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022) [1](#), [2](#), [4](#), [6](#), [8](#), [10](#), [11](#), [12](#)
66. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [4](#)
67. Xiao, Y., Tang, Z., Wei, P., Liu, C., Lin, L.: Masked images are counterfactual samples for robust fine-tuning. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023) [2](#), [4](#), [6](#), [10](#), [11](#)
68. Young, P., Lai, A., Hodosh, M., Hockenmaier, J.: From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. Transactions of the Association for Computational Linguistics (2014) [6](#), [10](#), [13](#)
69. Yu, F., Tang, J., Yin, W., Sun, Y., Tian, H., Wu, H., Wang, H.: Ernie-vil: Knowledge enhanced vision-language representations through scene graphs. In: Proc. AAAI Conference on Artificial Intelligence (AAAI) (2021) [13](#), [14](#)
70. Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. International Journal of Computer Vision (IJCV) (2019) [10](#)
71. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Conditional prompt learning for vision-language models. In: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022) [4](#), [10](#), [13](#)
72. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. International Journal of Computer Vision (IJCV) (2022) [2](#), [4](#), [10](#), [13](#)