

DOMAIN GENERALIZATION USING LARGE PRE-TRAINED MODELS WITH MIXTURE-OF-ADAPTERS

Anonymous authors

Paper under double-blind review

ABSTRACT

Learning a robust vision model despite large distribution shift situations is an important task for model deployment in real-world settings. Especially, domain generalization (DG) algorithm aims to maintain the performance of a trained model on different distributions which were not seen during training. One of the most effective method has been leveraging the already learned rich knowledge of large pretrained models. However, naively tuning large models to DG tasks is practically infeasible due to memory limitations, extensive time requirements for training, and the risk of learned knowledge deterioration. Parameter-efficient fine-tuning (PEFT) methods have been used to reduce the high computational cost during training and efficiently adapt large models to downstream tasks. In this work for the first time we find that the use of adapters in PEFT methods not only reduce high computational cost during training but also serve as an effective regularizer for DG tasks. Surprisingly, a naive adapter implementation for large models achieve superior performance on common datasets. However, in situations of large distribution shifts, additional factors such as optimal amount of regularization due to the strength of distribution shifts should be considered for a sophisticated adapter implementation. As a result, we propose a mixture-of-expert based adapter fine-tuning method, dubbed as mixture-of-adapters (MoA). We employ multiple adapters that have varying capacities, and by using learnable routers, we allocate each token to a proper adapter. By using both PEFT and MoA methods we effectively alleviate the performance deterioration caused by distribution shifts and achieve state-of-the-art performance on diverse DG benchmarks.

1 INTRODUCTION

The goal of domain generalization (DG) is to well predict on domains that were unavailable during training (a.k.a unseen domains) (Gulrajani & Lopez-Paz, 2020). In DG settings, the model is trained on multiple source domains and evaluated on one unseen target domain (Zhou et al., 2021; Gulrajani & Lopez-Paz, 2020; Cha et al., 2021). Unlike domain adaptation, domain generalization is unable to access any information about the target domain. Therefore DG algorithms should fully exploit domain invariant features underlying in the source domain to well predict on the target domain (Seo et al., 2020; Gulrajani & Lopez-Paz, 2020; Arjovsky et al., 2019).

In recent times the usage of large pretrained models is gaining popularity in domain generalization fields (Cha et al., 2022; Mao et al., 2022; Lew et al., 2023; Li et al., 2023). Since large pretrained models already possess some extent of domain invariant knowledge (Cha et al., 2022), exploiting this knowledge into domain generalization has become a popular choice. There are few studies that have tried to train these models directly with empirical risk minimization (ERM) (Vapnik, 1998). Angarano et al. (2022) shows that ERM algorithm performs competitively well when accompanied with proper backbones like EfficientNet (Tan & Le, 2019), ViT (Dosovitskiy et al., 2020), DeiT (Touvron et al., 2021), LeViT (Graham et al., 2021) and ConViT (d’Ascoli et al., 2021). This inspired us to leverage large pretrained models on DG settings. However it is widely known that naively fine-tuning a large model is impractical, not only does it demand a large size VRAM caused by high peak memory in training phase, but it also requires a significant amount of training time due to numerous amount of parameters. Additionally, the pretrained feature extractor can be distorted by overfitting on source domains (Gao et al., 2021) and its contextual information can be harmed (Kumar et al., 2022; Mao et al., 2022; Wortsman et al., 2022), which overall results in a generalization

failure to out-of-distribution data. Fig. 1 compares the accuracy of various fine-tuning methods, including naive full fine-tuning, linear probing, and other partial fine-tuning methods, which also highlights their incapability to handle such distribution shifts.

This work represents a pioneering adoption of Parameter-Efficient Fine-Tuning (PEFT) methods in the context of domain generalization, expanding its usage beyond its traditional domains on transfer learning. PEFT methods (Houlsby et al., 2019; Karimi Mahabadi et al., 2021; He et al., 2022; Zaken et al., 2021; Hu et al., 2021; Ryu, 2023; Jia et al., 2022; Chen et al., 2022; Li & Liang, 2021; Sung et al., 2022a) mitigate the high cost of full fine-tuning of large pretrained models and aims to reach or exceed the performance of full fine-tuning or zero-shot performance by tuning only some parts of the model (Zaken et al., 2021) or by employing a small number of external learnable parameters (Hu et al., 2021; Ryu, 2023; Houlsby et al., 2019; Karimi Mahabadi et al., 2021; He et al., 2022; Jia et al., 2022). We compare various aspects of different trainable parameter settings (e.g., full fine-tuning, and adapter fine-tuning) and empirically verify that PEFT methods can act as an effective regularization during the training process and show that it can largely outperform a fully fine-tuned model and reach comparable performance with recent state-of-the-art methods (Cha et al., 2021; 2022; Arpit et al., 2022; Li et al., 2023) in DG settings. We also discover that the optimal strength of PEFT regularization differs, depending on the amount of a distribution shift.

As shown in Fig. 1, linear probing (Linear), bias tuning in attention layer (Bias (MSA)), bias tuning in both attention and MLP layer (Bias (MSA+MLP)) show comparable performance on PACS (Li et al., 2017) and VLCS (Fang et al., 2013) dataset. However on datasets with large distribution shifts, say TerraIncognita (Beery et al., 2018), Bias (MSA+MLP) shows significant improvement compared to the Bias (MSA) and Linear method. This shows that determining the proper amount of regularization by adjusting the trainable parameter is crucial to deal with different distribution shift. To handle this, we introduce mixture-of-expert based adapter architecture, called mixture-of-adapters (MoAs), that adequately manipulate the magnitude of regularization by employing adapters that have different capacity and routing each token to proper adapters. By doing so, we can further improve the performance on DG tasks by employing MoAs and learnable routers to handle the different intensity of distribution shift among various datasets and achieve state-of-the-art performance on DG benchmark.

2 RELATED WORK

Domain generalization. For the past decade many learning methods on how to learn domain invariant representations have been proposed in the domain generalization field. Empirical risk minimization (ERM) (Vapnik, 1998) which is one of the most simplest approaches, just minimizes the loss on each domains and trains the model. In DomainBed (Gulrajani & Lopez-Paz, 2020), ERM still remains effective within the benchmark’s restricted hyperparameter search space and model selection methods. DomainBed tested many DG methods such as IRM (Arjovsky et al., 2019), GroupDRO (Sagawa et al., 2019), MixUp (Xu et al., 2020), DANN (Ganin et al., 2016), and CORAL (Sun & Saenko, 2016) in a unified and contained experimental setup. SWAD (Cha et al., 2021) explored the relationship between flat loss surfaces and DG performance, achieving superior




Figure 1: Domain generalization benchmark results with varying trainable parameters in ViT-Base (Dosovitskiy et al., 2020) pretrained with CLIP of OpenAI (Radford et al., 2021). Y axis is accuracy. We use linear probing (denoted as Linear), bias tuning in attention layer (Bias (MSA)), bias tuning in both attention and MLP layer (Bias (MSA+MLP)), and full fine-tuning to show the accuracy change according to the trainable parameter change when using PEFT methods with large models. OH, TI, DN denotes OfficeHome (Venkateswara et al., 2017), TerraIncognita (Beery et al., 2018), and DomainNet (Peng et al., 2019), respectively.

results. Ensemble-of-Averages (EoA) (Arpit et al., 2022) perform model averaging in training time and ensemble them in test time. Current research is focused on leveraging knowledge from large pre-trained models, with MIRO (Cha et al., 2022) introducing an oracle model approximated by a large pretrained model to maximize the mutual information between the target model. Recently, methods that ensemble diverse models show remarkable performance in the DG benchmark. SIMPLE (Li et al., 2023) utilizes many different pretrained models from a ModelPool and extracts outputs from the frozen pretrained models, and trains a shallow dispatcher using these outputs.

Parameter efficient fine-tuning. Leveraging large pretrained models for specific tasks involves fine-tuning, but fine-tuning all parameters is impractical. Recent approaches such as Parameter-efficient Fine-Tuning (PEFT) (Houlsby et al., 2019; He et al., 2021; Paul et al., 2022), focus on efficient fine-tuning by freezing most of the model parameters and optimizing only a few for a given task. Many successful PEFT approaches (Zaken et al., 2021; Hu et al., 2021; Ryu, 2023; Gao et al., 2020; He et al., 2021; Jia et al., 2022; Lester et al., 2021; Sung et al., 2022a;b; Zhang et al., 2022; 2023) adopt popular pretrained models to various downstream tasks. Among these methods, the use of adapters (Hu et al., 2021; Karimi Mahabadi et al., 2021; He et al., 2022) are adopted because of its high performance and efficient computation cost. Adapters are small modules trained on specific tasks which are inserted between network layers, where all the layers except for the adapters are frozen during training. LoRA (Hu et al., 2021), Compacter (Karimi Mahabadi et al., 2021) and KAdaptation (He et al., 2022) greatly reduced the number of trainable parameters by attaching a low-rank hypercomplex adapter layers in the transformer model and by decomposing the updated weight matrix into low rank matrices respectively.

Mixture-of-Experts. Mixture-of-Experts models (MoEs) are proposed to improve model performance by incorporating multiple subsets of parameters called ‘experts’ with routing algorithms conditioned by its input (Jacobs et al., 1991; Jordan & Jacobs, 1994; Eigen et al., 2013). Evolving from this, a type of model called sparse MoEs has become popular in both NLP (Shazeer et al., 2017; Lepikhin et al., 2020; Zoph et al., 2022; Fedus et al., 2022; Du et al., 2022) and vision (Ahmed et al., 2016; Gross et al., 2017; Yang et al., 2019; Wang et al., 2020; Riquelme et al., 2021) tasks lately. This is due to its capability to enhance model capacity while simultaneously reducing the substantial increase in the computational resources required for training. In the field of domain generalization, Li et al. (2022) incorporated MoE design into DG tasks, and proposed a Generalizable Mixture-of-Experts (GMOE) architecture to effectively handle distribution shifts.

3 PRELIMINARIES

3.1 DOMAIN GENERALIZATION

Let us denote the set of training domains as $\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^K$ where K implies the total number of training domains, and \mathcal{D}_i is the distribution over the input space. Also, define the set of target domains as $\mathcal{T} = \{\mathcal{T}_i\}_{i=1}^{K'}$ where K' denotes the total number of training domains. The training dataset is composed of $n_{\mathcal{D}_i}$ data points denoted as $(x_j^i, y_j^i)_{j=1}^{n_{\mathcal{D}_i}} \sim \mathcal{D}_i$ where x is the input and y is the target label for each training domain. In DG settings, the goal is to find the model parameter θ of a classifier f_θ which generalizes well on both \mathcal{D} and \mathcal{T} . To be more precise, for the ERM algorithm, we define

$$\mathcal{E}_{\mathcal{D}}(\theta) = \frac{1}{K} \sum_{i=1}^K \mathbb{E}_{(x^i, y^i) \sim \mathcal{D}_i} [l(f_\theta(x^i), y^i)] \quad (1)$$

where $l(\cdot, \cdot)$ denote a loss function and define $\mathcal{E}_{\mathcal{T}}$ in the similar manner. We minimize empirical risk $\hat{\mathcal{E}}_{\mathcal{D}}(\theta) = \frac{1}{K} \sum_{i=1}^K \frac{1}{n_{\mathcal{D}_i}} \sum_{j=1}^{n_{\mathcal{D}_i}} [l(f_\theta(x_j^i), y_j^i)]$ during training and expect the optimal parameter $\hat{\theta} = \arg \min_{\theta} \hat{\mathcal{E}}_{\mathcal{D}}(\theta)$ to be optimal for $\mathcal{E}_{\mathcal{T}}(\theta)$.

3.2 PARAMETER EFFICIENT ADAPTER

Parameter-Efficient Fine-Tuning (PEFT) efficiently handles enormous pretrained models to a downstream task and achieves significant performance unlike the direct fine-tuning methods where high cost of memory and computation are required. Previous methods of fine-tuning large models have been accomplished by partial fine-tuning (Touvron et al., 2022) methods. However, recent works




Figure 2: Flatness comparison of loss surfaces from models trained with full fine-tuning, LoRA, KAdaptation, and KAdaptation with Mixture-of-Adapter (denoted as KMoA) on PACS dataset Li et al. (2017). All visualizations are computed from test environment 0 (Art) domain.

(Houlsby et al., 2019; Hu et al., 2021; Karimi Mahabadi et al., 2021) reveal that employing a learnable layer in the frozen pretrained weights can surpass the performance of the partial fine-tuning method.

More specifically, Aghajanyan et al. (2020) verifies that dense neural network layers with full-rank matrices can be reduced to lower-rank subspaces. From this finding, Hu et al. (2021) constrains the updated weight $\Delta \mathbf{W} = \mathbf{B}\mathbf{A} \in \mathbb{R}^{k \times d}$ to have a low intrinsic rank which can be expressed as $\mathbf{h} = \mathbf{W}_0\mathbf{x} + \Delta \mathbf{W}\mathbf{x} = \mathbf{W}_0\mathbf{x} + \mathbf{B}\mathbf{A}\mathbf{x}$, where \mathbf{x} denotes input sequences of each module, $\mathbf{W}_0 \in \mathbb{R}^{k \times d}$ denotes the frozen weight matrix of a pretrained model, and $\mathbf{A} \in \mathbb{R}^{k \times r}$, and $\mathbf{B} \in \mathbb{R}^{r \times d}$ ($r < k, d$) are the trainable parameters during training.

In addition, KAdaptation (He et al., 2022) and Compacter (Karimi Mahabadi et al., 2021) both exploit Kronecker products to decompose the weight matrices into reduce trainable parameters. The Kronecker product between matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{p \times q}$ is denoted by $\mathbf{A} \otimes \mathbf{B} \in \mathbb{R}^{mp \times nq}$ and can be expressed as the following:

$$\mathbf{A} \otimes \mathbf{B} = \begin{pmatrix} a_{11}\mathbf{B} & \dots & a_{1n}\mathbf{B} \\ \vdots & \ddots & \vdots \\ a_{m1}\mathbf{B} & \dots & a_{mm}\mathbf{B} \end{pmatrix} \quad (2)$$

where a_{ij} indicates the element in the i -th row and the j -th column of \mathbf{A} . It decomposes the update weight matrix $\Delta \mathbf{W} = \sum_{i=1}^t \mathbf{A}_i \otimes \mathbf{B}_i \in \mathbb{R}^{k \times d}$ where t is a hyperparameter that decides the number of Kronecker products. Furthermore, \mathbf{W} can be expressed as the following: $\mathbf{W} = \sum_{i=1}^t \mathbf{A}_i \otimes \mathbf{B}_i = \sum_{i=1}^n \mathbf{A}_i \otimes (\mathbf{u}_i \mathbf{v}_i^\top)$ where slow-weights $\mathbf{A}_i \in \mathbb{R}^{t \times t}$ is shared across all layers, and fast-weights $\mathbf{B}_i \in \mathbb{R}^{\frac{k}{t} \times \frac{d}{t}}$ is decomposed into low-rank matrices $\mathbf{u}_i \in \mathbb{R}^{\frac{k}{t} \times r}$ and $\mathbf{v}_i \in \mathbb{R}^{r \times \frac{d}{t}}$ with $i \in \{1, \dots, t\}$. Compacter decomposes the weight matrix to the additional adapter layer whereas KAdaptation decomposes the update matrix $\Delta \mathbf{W}$ in the original layer.

4 ANALYSIS OF ADAPTERS FOR DOMAIN GENERALIZATION

In this section, we investigate how adapters affect loss landscape flatness by analyzing the maximum Hessian eigenvalue spectra and loss landscapes of various trained models. It is widely known that finding a flat minima during optimization is closely related to the generalization performance and robustness of trained models, as documented in prior studies (Izmailov et al., 2018; Keskar et al., 2016; Garipov et al., 2018; Foret et al., 2020; Cha et al., 2021; Park & Kim, 2022). Therefore, by comparing loss landscapes along with maximum Hessian eigenvalues, we can anticipate which model will generalize better to unseen domains. We show the results of our analysis in the following sections.

4.1 LOSS LANDSCAPES

As shown in (Park & Kim, 2022; Garipov et al., 2018; Li et al., 2018) randomly perturbing the trained weights with random direction vectors, we can obtain variations of the loss value, which can be used to predict the models generalization ability. We evaluate on PACS (Li et al., 2017) dataset, one of the dataset in the DG benchmark (Gulrajani & Lopez-Paz, 2020). As shown in Fig. 2,

Table 1: Quantitative evaluation on domain generalization datasets with different PEFT methods. Following DomainBed (Gulrajani & Lopez-Paz, 2020), we evaluate our algorithm on PACS (Li et al., 2017), VLCS (Fang et al., 2013), OfficeHome (Venkateswara et al., 2017), TerraIncognita (Beery et al., 2018) and DomainNet (Peng et al., 2019). For fair comparison, we use training-domain-validation model selection strategy. Reported performance of each algorithms are from their papers. We indicate the pretraining method and dataset as the form of $\text{PRETRAIN}_{\text{DATASET}}$, and IG3B denotes Instagram-3B dataset from Zellers et al. (2018). Results with our method are colored in blue.

Algorithm	Architecture	Pretraining	PACS	VLCS	OfficeHome	TerraIn.	DomainNet	Avg.	#Param.	Trainable #Param.
MIRO	RegNetY-16GF	SWAG _{IG3B}	97.4 ± 0.2	79.9 ± 0.6	80.4 ± 0.2	58.9 ± 1.3	53.8 ± 0.1	74.1	167.2M	83.6M
SMA	RegNetY-16GF	SWAG _{IG3B}	95.5 ± 0.0	80.7 ± 0.1	82.0 ± 0.0	59.7 ± 0.0	60.0 ± 0.0	75.6	83.6M	83.6M
MIRO+SWAD	RegNetY-16GF	SWAG _{IG3B}	96.8 ± 0.2	81.7 ± 0.1	83.3 ± 0.1	64.3 ± 0.3	60.7 ± 0.0	77.3	167.2M	83.6M
<i>Methods with Parameter-Efficient Fine-Tuning (PEFT)</i>										
ERM (Baseline)	ViT-B/16	CLIP _{LAION2B}	85.8 ± 2.1	78.5 ± 0.9	78.1 ± 0.8	41.0 ± 1.6	52.2 ± 0.1	67.1	85.8M	85.8M
ERM _{LoRA, r=2}	ViT-B/16	CLIP _{LAION2B}	96.4 ± 0.6	82.6 ± 0.6	86.7 ± 0.3	46.1 ± 1.7	61.5 ± 0.1	74.7	85.9M	0.11M
ERM _{KAdaptation}	ViT-B/16	CLIP _{LAION2B}	97.5 ± 0.1	83.0 ± 0.1	90.3 ± 0.1	51.9 ± 0.5	62.7 ± 0.0	<u>77.1</u>	85.9M	0.14M

compared to the case of full fine-tuning (Fig. 2a) where we trained all of the model parameters, model parameters tuned with LoRA (Hu et al., 2021; Ryu, 2023) (Fig. 2b) and KAdaptation (He et al., 2022) (Fig. 2c) displays a more flatter loss surface. Additionally, KAdaptation shows a flatter loss surface than LoRA. We show further visualizations of the loss landscape on additional test environments on PACS in Appendix A.1.

4.2 MAXIMUM HESSIAN EIGENVALUE SPECTRA

Drawing loss landscapes from the optimal point to random directions sometimes fail to fully represent the shape of loss surfaces due to the high dimensionality of the loss surface (Garipov et al., 2018; Li et al., 2018). Therefore, following Park & Kim (2022) we calculate the top-5 Hessian eigenvalues and show their spectra.

Analogous to loss landscape results in the previous section, we observe the same phenomena with the use of adapters, say LoRA and KAdaptation. As depicted in Fig. 3, the top-5 Hessian eigenvalues are further zero-concentrated than the fully fine-tuned model. These discrepancy implies that employing adapters with large pre-trained models results in a flatter loss surface around the optimal point. This is because the Hessian matrix serves as an indicator of the curvature of the loss surface. Since the flatness of loss surfaces has a substantial impact on model performance in DG tasks Cha et al. (2021), it can be concluded that adapter fine-tuning with large pre-trained models effectively facilitate the performance in domain generalization. We show visualizations of max Hessian eigenvalue spectra on additional test environments on PACS dataset in Appendix A.2.

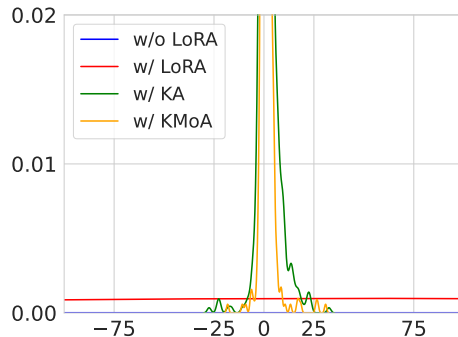


Figure 3: Hessian max eigenvalue spectra of test environment 0 (Art) in PACS dataset. KA and KMoA denotes KAdaptation and Mixture-of-Adapter with KAdaptation method, respectively. In the order of w/o LoRA, w/LoRA, w/KA, and w/KMoA, the hessian eigenvalue spectrum concentrates to zero and eigenvalues are suppressed accordingly.

4.3 PARAMETER-EFFICIENT ADAPTER FOR DOMAIN GENERALIZATION

Previously, we demonstrated that models trained with parameter-efficient fine-tuning methods have the potential to reach a more generalizable optimization point. Nevertheless, it’s challenging to come to a conclusive decision that a flatter loss surface directly translates to better generalization performance. Consequently, in the following section, we conduct a comprehensive evaluation of various fine-tuning methods and make comparisons between them. Finally, we discuss about the most effective practical fine-tuning approach and explore strategies to further enhance performance.

In Tab. 1 we show results from various paper on five datasets, namely PACS (Li et al., 2017), VLCS (Fang et al., 2013), OfficeHome (Venkateswara et al., 2017), TerraIncognita Beery et al. (2018) and DomainNet (Peng et al., 2019), and from the observations above, we evaluate ERM with various PEFT adapter methods in standard DG benchmark form Gulrajani & Lopez-Paz (2020). Some works like Cha et al. (2022) report the results of full fine-tuning a ViT-B/16 model with an ERM algorithm. However just by training a small adapter layer or an attention layer, we observe a greater performance improvement. Compared to the methods that use additional regularization or ensembling, our parameter efficient training approach with naive ERM algorithm reaches comparable accuracy and even state-of-the-art for the OfficeHome dataset. Additionally KAdaptation method achieves the best average accuracy of all the methods. These results validate our observation in Sec. 4.1 and 4.2, demonstrating that a proper selection of an adapter method can boost the performance adequately.

5 MIXTURE-OF-ADAPTERS (MOA)

In this section we show the effectiveness of our proposed method, mixture-of-adapter, which can tune large models with low cost and achieve great performance on DG tasks. As revealed in Li et al. (2022), mixture-of-expert model architectures with the use of cosine routing is more effective on DG tasks. Experts also handle some parts of the visual attributes, therefore reducing incorrect token allocation caused by intra-domain similarities. We bring it to our adapter-based method, which maintains the computational efficiency and achieves better results in domain generalization.

As demonstrated in Sec. 1, we utilize adapters to handle various amounts of distribution shift among the domains. Specifically, adapters, each with distinct capacities at every MoA layer, are employed to select the top-k outputs which are averaged and then integrated with the original layer’s outputs. We manipulate the capacities of each adapter by adjusting its inner rank.

The mixture-of-experts (MoE) layer utilizing the cosine router $G(\mathbf{x})$ with embedding \mathbf{E} and adapter A_{r_i} with inner rank r_i can be denoted as:

$$\begin{aligned} f_{\text{MoE}}(\mathbf{x}) &= \sum_{i=1}^N G(\mathbf{x})_i A_{r_i}(\mathbf{x}) \\ &= \sum_{i=1}^N \text{TOP}_k \left(\text{Softmax} \left(\frac{\mathbf{E}^T \mathbf{W}_x}{\tau \|\mathbf{W}_x\| \|\mathbf{E}\|} \right) \right)_i A_{r_i}(\mathbf{x}) \end{aligned} \quad (3)$$

where $A_{r_i}(\mathbf{x})$ is the output of an adapter with a different inner rank r_i and τ is a learnable temperature term of a softmax layer. In detail, we incorporate adapters in the large pretrained model by attaching them to attention sub modules. Following the implementation of traditional mixture-of-expert (MoE) methods, routers are adopted to dispatch tokens to the appropriate expert adapters. An adapter can be of any form such as LoRA (Hu et al., 2021), one of the most popular method in PEFT, Compacter (Karimi Mahabadi et al., 2021), and KAdaptation (He et al., 2022). In the case of KAdaptation, we choose the inner rank r_i as the additional decomposition of \mathbf{B}_i . Visualizations of our architecture is shown in Fig. 4




Figure 4: Architecture of proposed Mixture-of-Adapters (MoA). \mathbf{W}_0 , \mathbf{x}_{in} , and \mathbf{x}_{out} denotes original pretrained weight, input, and output tokens in multi-head self-attention (MHSA). Adapter in Fig. 4a can be any adapter-based PEFT methods like LoRA, Compacter, KAdaptation, etc, and Router in Fig. 4b can be linear or cosine router that commonly used in Mixture-of-Expert methods.

Table 2: Quantitative evaluation on domain generalization datasets with Mixture-of-Adapter (MoA) methods. The numbers marked with † are reported results from Cha et al. (2022). OpenAI denotes the private dataset used to train CLIP (Radford et al., 2021), and ModelPool-A and B denotes the set of pretrained models used in Li et al. (2023). Results with ensemble methods are colored in gray, and results with our method are colored in blue.

Algorithm	Architecture	Pretraining	PACS	VLCS	OfficeHome	TerraIn.	DomainNet	Avg.	#Param.	Trainable #Param.
ERM†	ViT-B/16	CLIP _{OpenAI}	83.4 ± 0.5	75.9 ± 1.3	66.4 ± 0.5	35.3 ± 0.8	44.4 ± 0.6	61.1	85.8M	85.8M
MIRO†	ViT-B/16	CLIP _{OpenAI}	95.6 ± 0.8	82.2 ± 0.3	82.5 ± 0.1	54.3 ± 0.4	54.0 ± 0.3	73.7	172M	85.8M
ERM†	RegNetY-16GF	SWAG _{IG3B}	89.6 ± 0.4	78.6 ± 0.3	71.9 ± 0.6	51.4 ± 1.8	48.5 ± 0.6	68.0	83.6M	83.6M
MIRO†	RegNetY-16GF	SWAG _{IG3B}	97.4 ± 0.2	79.9 ± 0.6	80.4 ± 0.2	58.9 ± 1.3	53.8 ± 0.1	74.1	167.2M	83.6M
ERM+SWAD†	RegNetY-16GF	SWAG _{IG3B}	94.7 ± 0.2	79.7 ± 0.2	80.0 ± 0.1	57.9 ± 0.7	53.6 ± 0.6	73.2	83.6M	83.6M
MIRO+SWAD†	RegNetY-16GF	SWAG _{IG3B}	96.8 ± 0.2	81.7 ± 0.1	83.3 ± 0.1	64.3 ± 0.3	60.7 ± 0.0	77.3	167.2M	83.6M
SMA	RegNetY-16GF	SWAG _{IG3B}	95.5 ± 0.0	80.7 ± 0.1	82.0 ± 0.0	59.7 ± 0.0	60.0 ± 0.0	75.6	83.6M	83.6M
<i>Methods with Parameter-Efficient Fine-Tuning (PEFT)</i>										
ERM (Baseline)	ViT-B/16	CLIP _{LAION2B}	85.8 ± 2.1	78.5 ± 0.9	78.1 ± 0.8	41.0 ± 1.6	52.2 ± 0.1	67.1	85.8M	85.8M
ERM _{Compacter}	ViT-B/16	CLIP _{LAION2B}	94.1 ± 0.4	81.0 ± 0.5	83.0 ± 0.1	35.9 ± 0.7	56.2 ± 1.2	70.0	85.9M	0.10M
ERM _{Attention}	ViT-B/16	CLIP _{LAION2B}	93.8 ± 0.6	82.0 ± 0.3	85.9 ± 0.4	51.4 ± 0.8	57.2 ± 0.1	74.1	85.8M	28.4M
ERM _{LoRA, r=2}	ViT-B/16	CLIP _{LAION2B}	96.4 ± 0.6	82.6 ± 0.6	86.7 ± 0.3	46.1 ± 1.7	61.5 ± 0.1	74.7	85.9M	0.11M
ERM _{KAdaptation}	ViT-B/16	CLIP _{LAION2B}	97.5 ± 0.1	83.0 ± 0.1	90.3 ± 0.1	51.9 ± 0.5	<u>62.7</u> ± 0.0	77.1	85.9M	0.14M
<i>Methods with Mixture-of-Adapter</i>										
ERM _{LoRA-MoA}	ViT-B/16	CLIP _{LAION2B}	96.9 ± 0.3	82.8 ± 0.7	89.5 ± 0.2	49.2 ± 2.4	62.2 ± 0.0	75.9	87.2M	1.5M
ERM _{KAdaptation-MoA}	ViT-B/16	CLIP _{LAION2B}	97.4 ± 0.2	<u>83.1</u> ± 0.3	<u>90.6</u> ± 0.0	52.8 ± 1.4	<u>62.7</u> ± 0.1	77.3	87.3M	1.5M
<i>Methods Using Ensemble</i>										
EoA	RegNetY-16GF	SWAG _{IG3B}	95.8	81.1	83.9	61.1	60.9	76.6	> 500M	-
SIMPLE	ModelPool-A	ModelPool-A	88.6 ± 0.4	79.9 ± 0.5	84.6 ± 0.5	57.6 ± 0.8	49.2 ± 1.1	72.0	> 1,000M	0.9M
SIMPLE†	ModelPool-B	ModelPool-B	99.0 ± 0.1	82.7 ± 0.4	87.7 ± 0.4	<u>59.0</u> ± 0.6	61.9 ± 0.5	78.1	> 1,000M	0.9M
ERM _{KAdaptation-MoA-Ensemble}	ViT-B/16	CLIP _{LAION2B}	97.6	83.4	90.9	54.3	63.1	77.9	261.3M	-

6 EXPERIMENTS

Experimental details. We use the standard benchmark DomainBed (Gulrajani & Lopez-Paz, 2020) for training and evaluating the performance on the domain generalization task. Following this, we use fixed hyperparameters within the same backbone model for all the experiments. We train five types of models which are full fine-tuning, attention tuning, LoRA (Hu et al., 2021), KAdaptation (He et al., 2022), and Compacter (Karimi Mahabadi et al., 2021). We employ CLIP (Radford et al., 2021) trained ViT (Dosovitskiy et al., 2020) as our initialization model because CLIP carries strong zero-shot ability which can be used to obtain favorable generalization performance. In detail we use OpenCLIP (Ilharco et al., 2021), an open-source re-implementation of CLIP model with LAION-2B dataset (Schuhmann et al., 2022), for all our experiments. We describe additional implementation details in Appendix B.1

Evaluation protocols and datasets. For a fair comparison, we employ DomainBed evaluation protocols (Cha et al., 2021; Gulrajani & Lopez-Paz, 2020). The following five benchmark datasets are used: PACS (Li et al., 2017), VLCS (Fang et al., 2013), OfficeHome (Venkateswara et al., 2017), TerraIncognita (Beery et al., 2018), and DomainNet (Peng et al., 2019). Using a *leave-one-out cross-validation*, all performance scores are evaluated by averaging all the cases that use a single domain as the target domain and the others as the source domains. Experiment is repeated three times and 20% percent of source domain data is left out for validation purposes. Lastly model selection (training-domain validation) and hyperparameter search procedure is referenced from DomainBed (Gulrajani & Lopez-Paz, 2020). We perform three runs with different random seeds for each setting and report their mean and standard deviation to show the training randomness. In ablation studies, we keep all the random seeds fixed and conduct the experiment.

6.1 RESULTS FOR MIXTURE-OF-ADAPTER

As described in Sec. 5, we implement Mixture-of-Adapter (MoA) with LoRA (Hu et al., 2021) and KAdaptation, and report the benchmark result in Table 2. In addition to the performance gains by applying PEFT, results after employing MoA are increased consistently at all datasets. In the case of LoRA-MoA, average performance is increased by 1.2pp compared to the results when using only LoRA. Also, KAdaptation-MoA method achieves state-of-the-art results on VLCS, OfficeHome, DomainNet datasets, exceeding the result of the original KAdaptation by 0.2pp in average, validating the analyses about loss surface in Fig. 2. Additionally, following Arpit et al. (2022), ensembling the three weights obtained from different seeds, we further enhance the performance by 0.6pp in average, achieving best results in VLCS, OfficeHome, and DomainNet.

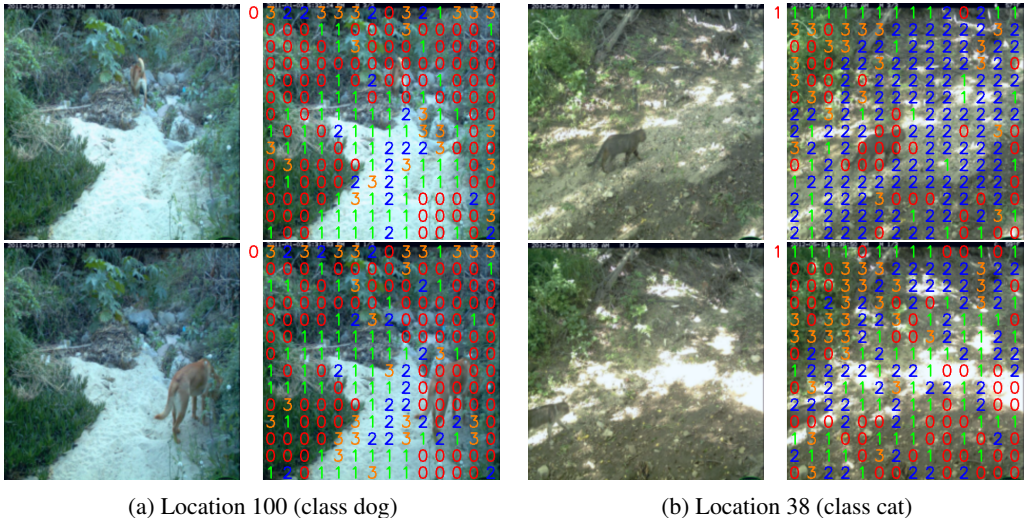


Figure 5: Visualizations of routed indices of each patch in TerraIncognita (Beery et al., 2018) dataset. Left column is original image, and in right column we indicate where each patch is routed. Upper and lower images were taken from same location but different time, therefore they have same background but different object (dog, cat) shape and location.

Analysis on the router and experts To understand why using multiple adapters with a router improves performance, we conduct experiments on what the router sees and what each expert learns. To investigate the role of an adapter, we monitor the routing path of each token and identify the expert to which it is directed. This visualization is performed on TerraIncognita dataset. In Fig. 5, the image is divided into patches and each patch is assigned a number, where the number corresponds to the routed expert. We also reveal that the router tends to cluster tokens in areas with semantic information (e.g., object foreground or object outlines). As an example, consider the dog depicted in Fig. 5a and the cat in Fig. 5b, which exhibit varying positions between the first and second rows. Images sharing the same location have consistent backgrounds, resulting in shared expert routing for background tokens. However, objects and their positions can differ, causing object-related tokens to be routed to different experts. In conclusion, the importance of routing numerous tokens to their respective adapters greatly enhances the model’s ability to capture semantic information in images, thereby enabling the model to effectively navigate challenging distribution shift scenarios in domain generalization tasks. We show more visualizations of routed patches in Appendix A.3.

6.2 ABLATION STUDY

Effect of fine-tuning in pretrained models We observe that using a fine-tuned model with a smaller dataset shows degraded performance. Specifically, LAION-2B (Schuhmann et al., 2022) pretrained CLIP-ViT model almost consistently outperforms the LAION-2B pretrained, ImageNet Deng et al. (2009) fine-tuned CLIP-ViT model across all the adapter methods in terms of accuracy except for TerraIncognita dataset, as illustrated in Fig. 6. This can be attributed to the superior generalization capabilities of larger models and the adapter’s capacity to preserve the knowledge from the pretrained model. These findings align with previous literature, such as Kumar et al. (2022), which suggests that fine-tuning entire large models can degrade the learned representations.

Ablation study on routing strategy in Mixture-of-Adapter We conduct an ablation study on three components proposed in Li et al. (2022), which are cosine router, auxiliary loss \mathcal{L}_{aux} , and the layer configuration on the location of the MoE layer. Li et al. (2022) demonstrates that employing a cosine router in their Generalizable Mixture-of-Experts (GMoE) architecture yields better performance when contrasted with employing a linear router. Our approach significantly deviates from theirs, as each of our expert adapters exhibit distinct capacities, and our MoA is attached to an attention layer. Thus we test both the linear and cosine routers. The auxiliary loss function (\mathcal{L}_{aux}) balances the amount of token allocation to each experts, and the layer configuration called ‘Every 2’ or ‘Last 2’ in their work denotes the attachment of MoE layer with every two layers or only on

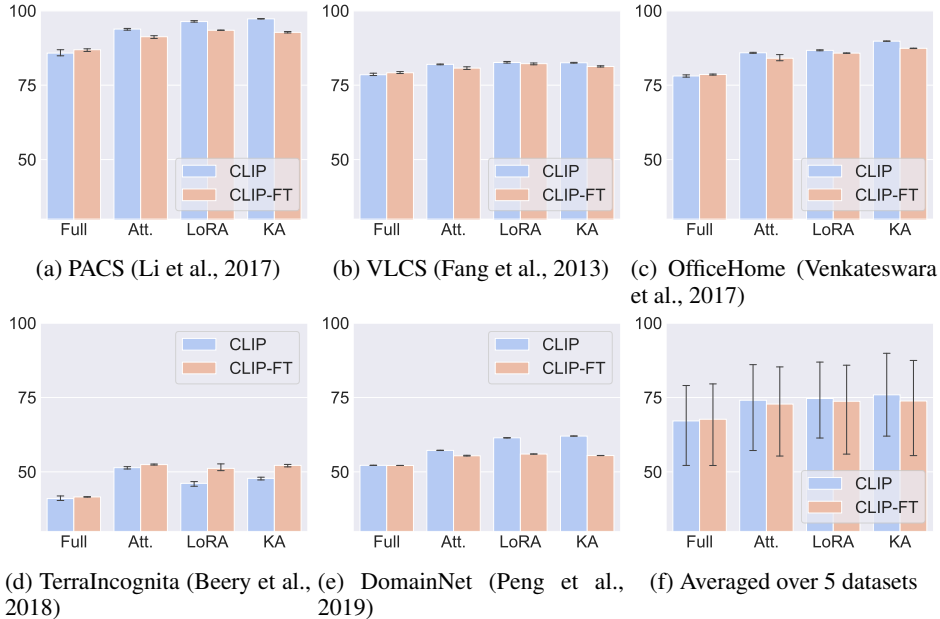


Figure 6: Performance comparison between original (not fine-tuned) CLIP and ImageNet fine-tuned CLIP from timm (Wightman, 2019) (denoted as CLIP-FT) on different fine-tuning strategies. Full, Att., LoRA, and KA denotes full fine-tuning, attention-only tuning (Touvron et al., 2022), LoRA, and KAdaptation, respectively.

Table 3: Performance comparison on different mixture-of-adapter settings. We perform experiments with dropping each components independently from our best setting which is highlighted in blue. ‘w/o \mathcal{L}_{aux} ’ denotes the results without auxiliary loss that used in (Li et al., 2022), ‘Cosine \rightarrow Linear’ denotes the result when we changed the Cosine router to a Linear router, and ‘Every \rightarrow Last’ denotes the result when we changed the adapter attached layer from Every, which attaches adapters in every two layer, to Last, which attaches adapters in last two layers.

Adapter	Changed component	PACS	VLCS	OfficeHome	TerraIn.	DomainNet	Avg.	#Param.	Trainable #Param.
	Original	97.5	82.8	90.6	53.1	62.6	77.3	87.3M	1.5M
KAdaptaion-MoA	w/o \mathcal{L}_{aux}	97.3	82.6	90.5	54.0	62.6	77.4	87.3M	1.5M
	Cosine \rightarrow Linear	97.5	82.3	90.3	51.5	62.6	76.9	86.1M	0.33M
	Every \rightarrow Last	97.2	82.2	90.2	47.5	62.4	75.9	86.3M	0.51M

the last two layers. Therefore we also conduct experiments on these settings and report the whole results in Table 3. In our setting, employing a cosine router showed a consistent performance increase for all domains. Also attaching an adapter with every two layers brought better performance than using MoA only on the last two layers. While there was a minimal difference between using and not using \mathcal{L}_{aux} (within the error margin), using the auxiliary loss yielded more better results for the four datasets (PACS, VLCS, OfficeHome, DomainNet). Thus we decided to incorporate the auxiliary loss \mathcal{L}_{aux} when conducting our main experiment.

7 CONCLUSION

We have shown that using only parameter-efficient fine-tuning can outperform or be competitive with previous state-of-the-art domain generalization algorithms. Additionally, we propose integrating extra adapters with learnable routers to handle various distribution shift situations. This allows us to achieve state-of-the-art results without the need for ensemble methods, and when we do employ ensembling, our results are on par with those achieved by ensemble approaches. We observe the powerful efficacy of PEFT and large models about domain generalization during this work, and expect that this research can motivate the way of research directions of domain generalization towards robustly fine-tuning the large pretrained models.

REFERENCES

- Armen Aghajanyan, Luke Zettlemoyer, and Sonal Gupta. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. *arXiv preprint arXiv:2012.13255*, 2020.
- Karim Ahmed, Mohammad Haris Baig, and Lorenzo Torresani. Network of experts for large-scale image categorization. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, pp. 516–532. Springer, 2016.
- Simone Angarano, Mauro Martini, Francesco Salvetti, Vittorio Mazzia, and Marcello Chiaberge. Back-to-bones: Rediscovering the role of backbones in domain generalization. *arXiv preprint arXiv:2209.01121*, 2022.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Devansh Arpit, Huan Wang, Yingbo Zhou, and Caiming Xiong. Ensemble of averages: Improving model selection and boosting performance in domain generalization. *Advances in Neural Information Processing Systems*, 35:8265–8277, 2022.
- Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 456–473, 2018.
- Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems*, 34:22405–22418, 2021.
- Junbum Cha, Kyungjae Lee, Sungrae Park, and Sanghyuk Chun. Domain generalization by mutual-information regularization with pre-trained models. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIII*, pp. 440–457. Springer, 2022.
- Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. *arXiv preprint arXiv:2205.08534*, 2022.
- 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*, pp. 248–255. Ieee, 2009.
- 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. *arXiv preprint arXiv:2010.11929*, 2020.
- Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, pp. 5547–5569. PMLR, 2022.
- Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In *International Conference on Machine Learning*, pp. 2286–2296. PMLR, 2021.
- David Eigen, Marc’Aurelio Ranzato, and Ilya Sutskever. Learning factored representations in a deep mixture of experts. *arXiv preprint arXiv:1312.4314*, 2013.
- Chen Fang, Ye Xu, and Daniel N Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1657–1664, 2013.
- William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 23(1): 5232–5270, 2022.

- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*, 2020.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021.
- Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*, 2020.
- Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31, 2018.
- Benjamin Graham, Alaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet’s clothing for faster inference. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 12259–12269, 2021.
- Sam Gross, Marc’Aurelio Ranzato, and Arthur Szlam. Hard mixtures of experts for large scale weakly supervised vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6865–6873, 2017.
- Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2020.
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366*, 2021.
- Xuehai He, Chunyuan Li, Pengchuan Zhang, Jianwei Yang, and Xin Eric Wang. Parameter-efficient fine-tuning for vision transformers. *arXiv preprint arXiv:2203.16329*, 2022.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pp. 2790–2799. PMLR, 2019.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021.
- Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, July 2021. URL <https://doi.org/10.5281/zenodo.5143773>.
- Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, pp. 709–727. Springer, 2022.
- Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994.

- Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers. *Advances in Neural Information Processing Systems*, 34:1022–1035, 2021.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Ananya Kumar, Aditi Raghunathan, Robbie Matthew Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International Conference on Learning Representations*, 2022.
- Dmitry Lepikhin, Hyoungho Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*, 2020.
- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021.
- Byunggyu Lew, Donghyun Son, and Buru Chang. Gradient estimation for unseen domain risk minimization with pre-trained models. *arXiv preprint arXiv:2302.01497*, 2023.
- Bo Li, Yifei Shen, Jingkan Yang, Yezhen Wang, Jiawei Ren, Tong Che, Jun Zhang, and Ziwei Liu. Sparse mixture-of-experts are domain generalizable learners. *arXiv preprint arXiv:2206.04046*, 2022.
- Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pp. 5542–5550, 2017.
- Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems*, 31, 2018.
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- Ziyue Li, Kan Ren, Xinyang Jiang, Yifei Shen, Haipeng Zhang, and Dongsheng Li. Simple: Specialized model-sample matching for domain generalization. In *The Eleventh International Conference on Learning Representations*, 2023.
- Xiaofeng Mao, Yuefeng Chen, Xiaojun Jia, Rong Zhang, Hui Xue, and Zhao Li. Context-aware robust fine-tuning. *arXiv preprint arXiv:2211.16175*, 2022.
- Namuk Park and Songkuk Kim. How do vision transformers work? In *International Conference on Learning Representations*, 2022.
- Lysandre Debut Younes Belkada Sayak Paul, Sourab Mangrulkar, and Sylvain Gugger. Peft: State-of-the-art parameter-efficient fine-tuning methods, 2022.
- Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1406–1415, 2019.
- 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*, pp. 8748–8763. PMLR, 2021.
- Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems*, 34:8583–8595, 2021.

- Simo Ryu. lora. <https://github.com/cloneofsimon/lora>, 2023.
- Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*, 2019.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*, 2022.
- Seonguk Seo, Yumin Suh, Dongwan Kim, Geeho Kim, Jongwoo Han, and Bohyung Han. Learning to optimize domain specific normalization for domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pp. 68–83. Springer, 2020.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part III 14*, pp. 443–450. Springer, 2016.
- Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. Lst: Ladder side-tuning for parameter and memory efficient transfer learning. *Advances in Neural Information Processing Systems*, 35:12991–13005, 2022a.
- Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. Vi-adapter: Parameter-efficient transfer learning for vision-and-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5227–5237, 2022b.
- Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pp. 6105–6114. PMLR, 2019.
- 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*, pp. 10347–10357. PMLR, 2021.
- Hugo Touvron, Matthieu Cord, Alaaeldin El-Nouby, Jakob Verbeek, and Hervé Jégou. Three things everyone should know about vision transformers. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV*, pp. 497–515. Springer, 2022.
- Vladimir Vapnik. *Statistical learning theory*. Wiley. New York, 1998.
- Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5018–5027, 2017.
- Xin Wang, Fisher Yu, Lisa Dunlap, Yi-An Ma, Ruth Wang, Azalia Mirhoseini, Trevor Darrell, and Joseph E Gonzalez. Deep mixture of experts via shallow embedding. In *Uncertainty in artificial intelligence*, pp. 552–562. PMLR, 2020.
- Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019.
- Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7959–7971, 2022.

- Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 6502–6509, 2020.
- Brandon Yang, Gabriel Bender, Quoc V Le, and Jiquan Ngiam. Condconv: Conditionally parameterized convolutions for efficient inference. *Advances in neural information processing systems*, 32, 2019.
- Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*, 2021.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. Swag: A large-scale adversarial dataset for grounded commonsense inference. *arXiv preprint arXiv:1808.05326*, 2018.
- Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8552–8562, 2022.
- Renrui Zhang, Xiangfei Hu, Bohao Li, Siyuan Huang, Hanqiu Deng, Hongsheng Li, Yu Qiao, and Peng Gao. Prompt, generate, then cache: Cascade of foundation models makes strong few-shot learners. *arXiv preprint arXiv:2303.02151*, 2023.
- Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*, 2021.
- Barret Zoph, Irwan Bello, Sameer Kumar, Nan Du, Yanping Huang, Jeff Dean, Noam Shazeer, and William Fedus. St-moe: Designing stable and transferable sparse expert models. *arXiv preprint arXiv:2202.08906*, 2022.

Appendix

In this appendix, we provide a more detailed explanations of the implementation details, results and visualizations of additional analyses, and discuss limitations of this work.

A ADDITIONAL ANALYSIS

In this section, we show additional analysis about loss landscapes, max Hessian eigenvalue spectra, and analyses about routed tokens.

A.1 COMPARISONS OF LOSS LANDSCAPE VISUALIZATIONS

We show loss landscapes for all test environments in PACS dataset (Li et al., 2017) in Fig. 7. Similar with the visualizations in main paper, the other test environments have a tendency that fully fine-tuned models show most sharp loss landscape. But trained models with LoRA and KAdaptation shows much more flatter loss landscapes, especially KAdaptation have most flat loss landscape.

A.2 COMPARISONS OF MAX HESSIAN EIGENVALUE SPECTRA

We show the max Hessian eigenvalue spectra evaluated on the all domains from PACS dataset (Li et al., 2017) in Fig. 8. Every domain shows consistent shape of max Hessian eigenvalue spectra that they shows more zero-concentrated shape in the order of w/o LoRA, w/LoRA, w/ KA and w/KMoA.

A.3 VISUALIZATIONS OF ROUTED PATCHES IN PACS AND TERRAINCOGNITA DATASET.

We additionally show the visualizations of routed patch indices in Fig. 9 on PACS dataset (Li et al., 2017), and Fig. 10, 11 on TerraIncognita dataset (Beery et al., 2018). All images are visualizations from the last adapter-attached transformer layer, layer 10. Similar with the findings from main paper, we can observe that same indices are clustered at the regions where having semantic meanings.

B ADDITIONAL EXPERIMENTAL DETAILS

B.1 IMPLEMENTATION DETAILS

Pretrained vision transformer models are obtained from the timm (Wightman, 2019) library. Adam (Kingma & Ba, 2014) optimizer is used for model optimization along with a learning rate of $5e - 5$. A batch size of 32 per domain is used for the ViT-Base model. We run 15,000 iterations on DomainNet and 5,000 for others, and evaluate at every 500 iteration steps for DomainNet, 200 steps for others. For the adapters, we arbitrarily choose the inner rank so that all adapter methods have approximately the same number of trainable parameters ($\sim 0.1M$). The LoRA (Hu et al., 2021) and KAdaptation (He et al., 2022) weights are attached to the QKV projection matrix and projection weight matrix in the self-attention layer. Also for Compacter (Karimi Mahabadi et al., 2021), we attach the adapter after the MLP layer in every blocks. Every adapter except LoRA are implemented using their official repository, and for LoRA, we implement it using unofficial implementation (Ryu, 2023). For pretrained models, we use models from timm (Wightman, 2019) library for all our experiments, and we perform all experiment on one server with 8 NVIDIA RTX3090 GPUs. We'll release training code for all our experiments and trained weights in public.

B.2 DETAILS ABOUT HYPERPARAMETERS

We report all the hyperparameters that are used in our experiments and analysis in Table 4.

C LIMITATIONS

Our method heavily relies on the performance of large pretrained models, hence using a better pretrained model can lead to improved performance. But, such models are limited and require a sub-

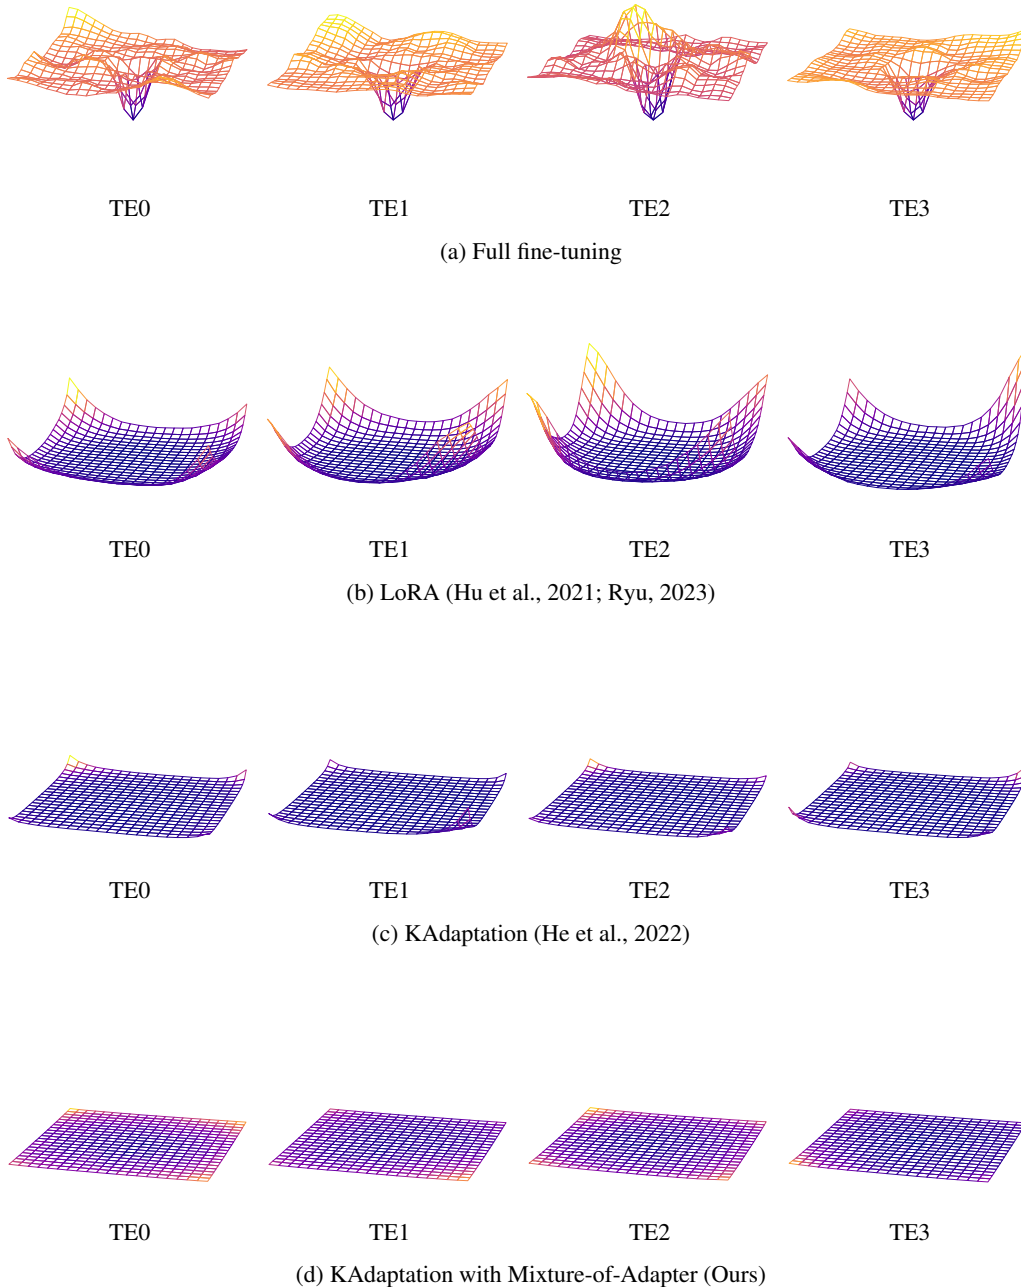


Figure 7: Flatness comparison of loss surfaces trained with full fine-tuning, LoRA, KAdaptation, and KAdaptation with mixture-of-expert on PACS dataset (Li et al., 2017).

stantial amount of time and cost for training. These weakness also exist in methods like MIRO (Cha et al., 2022) or SIMPLE (Li et al., 2023), and the availability of high-performance open-source models like OpenCLIP (Ilharco et al., 2021) can alleviate these drawbacks. Our approach may not yield significantly better performance on datasets that are more challenging than TerraIncognita (Beery et al., 2018) since it has fewer trainable parameters compared to fully fine-tuned DG algorithms. However, our method offers flexibility in adjusting trainable parameters through the control of inner




Figure 8: Comparison of max Hessian eigenvalue spectra trained with full fine-tuning, LoRA, and KAdaptation on PACS dataset (Li et al., 2017). In all test environments we can observe that our proposed KMoA shows the most zero-concentrated eigenvalue spectrum, and other two adapter method, KA and LoRA also shows smaller max Hessian eigenvalue distribution compared to the fully fine-tuned case (w/o LoRA).

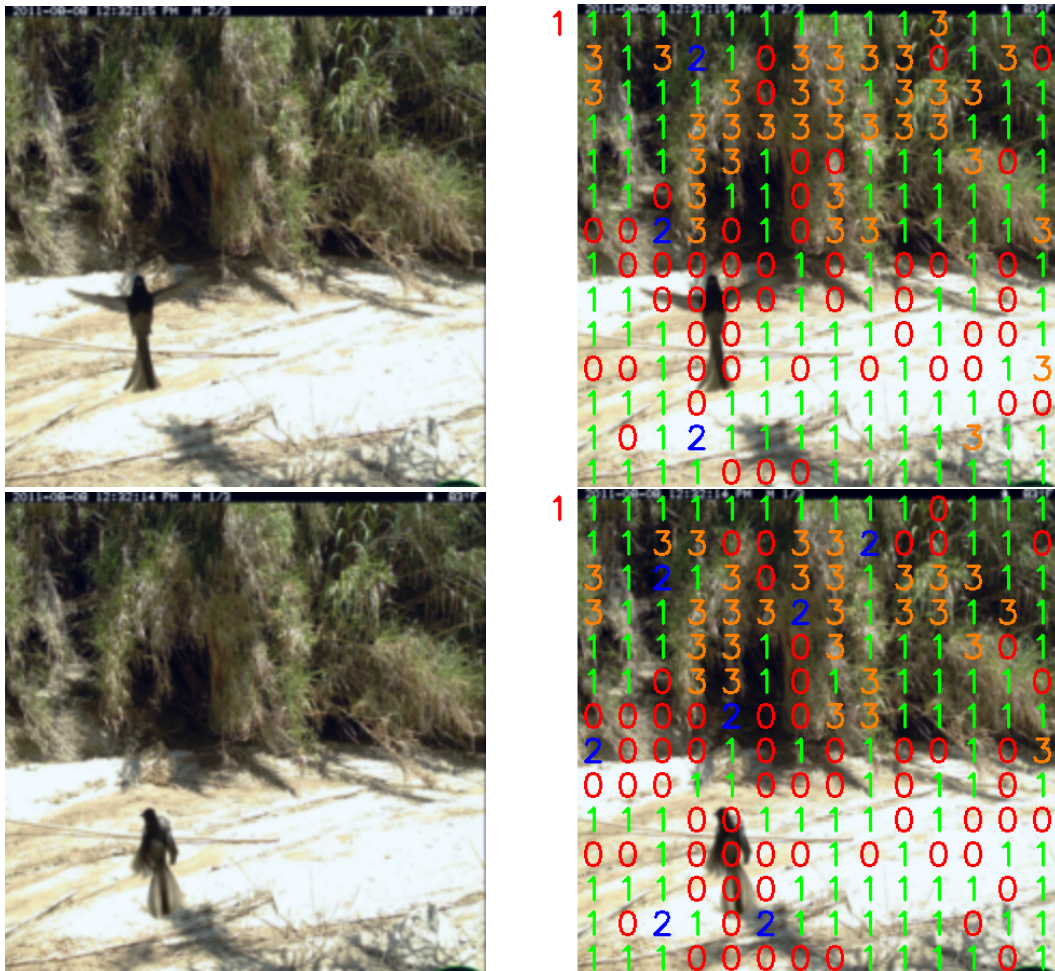
Table 4: List of hyperparameters used in experiments on domain generalization benchmarks.

Hyperparameter	LoRA	KAdaptation	LoRA-MoA	KMoA
# of Experts		N/A	4	4
Router		N/A	Cosine	
Router Top-k		N/A	Top-1	
Rank of adapter (r_i)	2	2	[1, 2, 4, 8]	
# of Kroneker products (t)	N/A	64	N/A	64
Batch size	160 (DomainNet), 96 (Otherwise)			
Learning rate	$5e - 5$			
Optimizer	Adam			

rank r_i and allows for obtaining the optimal rank through hyperparameter search, addressing this limitation effectively.

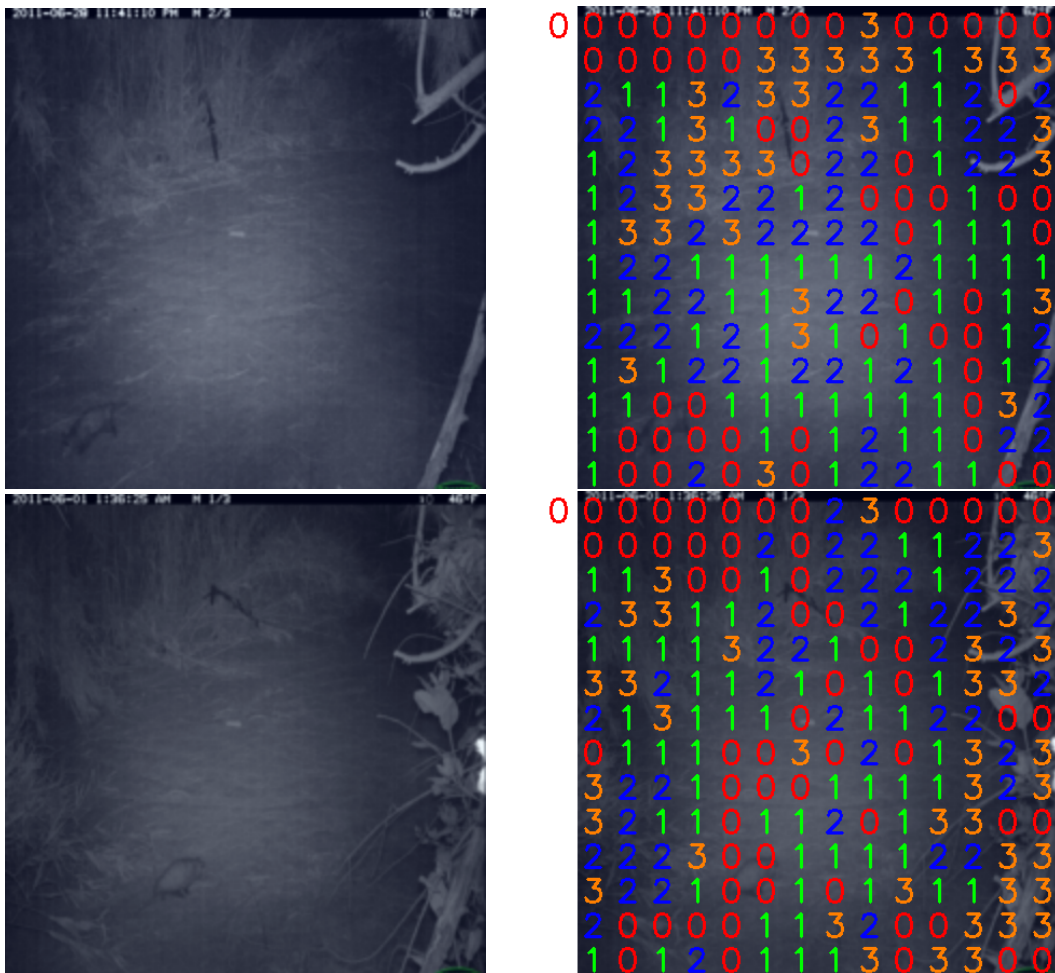


Figure 9: Visualizations of routed indices of each patch. We show a total of seven classes in PACS dataset (Li et al., 2017), with one class per row in the order of ‘Dog’, ‘Elephant’, ‘Giraffe’, ‘Guitar’, ‘Horse’, ‘House’, ‘Person’. Also, in each column, the same domains are located in the order of ‘Art Painting’, ‘Cartoon’, ‘Photo’, and ‘Sketch’.



Location 43

Figure 10: Visualizations of routed indices of each patch in TerraIncognita (Beery et al., 2018) dataset. Left column is original image, and in right column we indicate where each patch is routed. Upper and lower images were took from same location but different time, therefore they have same background but different object (bird) shape and location.



Location 46

Figure 11: Visualizations of routed indices of each patch in TerraIncognita (Beery et al., 2018) dataset. Left column is original image, and in right column we indicate where each patch is routed. Upper and lower images were taken from the same location but at different times, therefore they share the same background but feature different object (possum) shapes and positions.