

Leveraging Normalization Layer in Adapters With Progressive Learning and Adaptive Distillation for Cross-Domain Few-Shot Learning

Yongjin Yang, Taehyeon Kim, Se-Young Yun

KAIST AI

{dyjdkd, potter32, yunseyoung}@kaist.ac.kr

Abstract

Cross-domain few-shot learning presents a formidable challenge, as models must be trained on base classes and then tested on novel classes from various domains with only a few samples at hand. While prior approaches have primarily focused on parameter-efficient methods of using adapters, they often overlook two critical issues: shifts in batch statistics and noisy sample statistics arising from domain discrepancy variations. In this paper, we introduce a novel generic framework that leverages normalization layer in adapters with Progressive Learning and Adaptive Distillation (ProLAD), marking two principal contributions. First, our methodology utilizes two separate adapters: one devoid of a normalization layer, which is more effective for similar domains, and another embedded with a normalization layer, designed to leverage the batch statistics of the target domain, thus proving effective for dissimilar domains. Second, to address the pitfalls of noisy statistics, we deploy two strategies: a progressive training of the two adapters and an adaptive distillation technique derived from features determined by the model solely with the adapter devoid of a normalization layer. Through this adaptive distillation, our approach functions as a modulator, controlling the primary adapter for adaptation, based on each domain. Evaluations on standard cross-domain few-shot learning benchmarks confirm that our technique outperforms existing state-of-the-art methodologies.

1 Introduction

Deep neural networks, while showing remarkable aptitude for visual recognition tasks when trained on large datasets, face a significant challenge when applied to real-world scenarios characterized by diverse domains. The challenges become even more pronounced when these networks encounter unseen data with scarce samples during operational deployment. This is where Cross-Domain Few-Shot Learning (CD-FSL) comes into play. CD-FSL aims to leverage knowledge from a source domain (e.g., ImageNet (Deng et al. 2009)), where abundant labeled data is available, to learn a predictive model for a target domain (e.g., Fungi (Schroeder and Cui 2018)) where labeled data is scarce. The core objective of this learning paradigm is to cultivate a model that can generalize effectively to new classes within the target domain using a limited set of examples.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

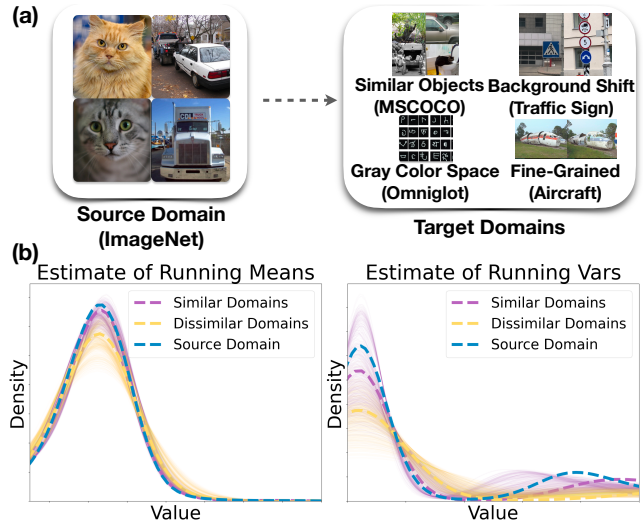


Figure 1: (a) Domain shift encountered during deployment in CD-FSL, (b) Kernel density estimation of batch statistics across different domains, where bold dotted lines represent the mean value of batch statistics for each domain and light solid lines represent the batch statistics of individual samples, showing high variations. Domain similarity is computed using Earth-Mover Distance (EMD) as outlined in (Cui et al. 2018; Oh et al. 2022). Detailed information is described in Appendix B.

The primary challenges encountered in CD-FSL stem from the diverse domain discrepancies (Li et al. 2020; Oh et al. 2022) and the scarcity of samples. As illustrated in Figure 1 (a), the model is required to adjust to new domains that may differ in aspects such as color space and background scenes, focal features essential for accurate classification, while using only a restricted number of samples. Adaptation from the source domain to the target domain is challenging since it easily poses a high risk of overfitting to a few samples. To address this challenge, a parameter-efficient approach that fine-tunes only linear adapters has gained attention in CD-FSL (Requeima et al. 2019; Bateni et al. 2020; Li, Liu, and Bilen 2022). This approach aims to balance the need for domain adaptation with the risk of overfitting by

limiting the number of parameters to be fine-tuned.

However, relying solely on linear adapters fails to address the distribution shift in latent features, which is induced by statistics in normalization layers biased towards the source set. Figure 1 (b) illustrates the batch statistics of the source (ImageNet) and target domains based on domain similarity in the Meta-Dataset (Triantafillou et al. 2019). It is observed that underlying batch statistics in the source and new domains can significantly differ in the case of dissimilar domains that do not share any similar classes with the source set. This bias in statistics towards the source set can render the additional learned parameters insufficient for domain adaptation (Li et al. 2016; Bilen and Vedaldi 2017; Du et al. 2020; Mirza et al. 2022). Therefore, the consideration of normalization layers in adapter-based approaches is necessary.

CD-FSL faces an additional challenge where attempting to incorporate statistics from target domain with a few samples may result in noisy batch statistics. This is evident in Figure 1 (b), as highlighted by the light solid lines, complicating the accurate representation of a domain statistics. This challenge becomes more pronounced for target domains similar to the source, where the statistics of the target domains substantially overlap with those of the source. As a result, integrating noisy data can disrupt the well-established statistics from the pre-trained batch normalization layers. Conversely, for dissimilar domains, even slightly noisy statistics better represent the statistics of target domain than the statistics of source. Given this scenario, there remains a demand for a method that can dynamically leverage statistics based on domain similarity considering the nature of diverse test domains of CD-FSL.

This paper presents a novel framework which leverages normalization layers in adapters with Progressive Learning and Adaptive Distillation (**ProLAD**) that progressively trains adapters in two stages with adaptive distillation. Two types of adapters are utilized: Task Adapter with Normalization (**TAN**), which employs normalization layer that collects batch statistics, and Task Adapter (**TA**), which solely comprises a linear adapter. The **TAN** is optimized for dissimilar domains due to its normalization layer, while the **TA** is best suited for similar domains, ensuring that well-represented batch statistics remain undisturbed. During the adapter training process, our method implements a progressive learning approach, training **TA** followed by **TAN** sequentially, and integrate adaptive distillation. This adaptive distillation, informed by the estimation of domain similarity, controls the activation of each adapter. As a result, the model leans on **TA** for similar domains and primarily engages **TAN** when faced with dissimilar domains. We demonstrate that our **ProLAD** outperforms baseline methods on the popular benchmark Meta-Dataset, especially in more challenging settings where the feature extractor is trained on a single domain dataset and then tested on all the other domains.

2 Related Work

2.1 Cross-Domain Few-Shot Learning

The early algorithms for solving the cross-domain few-shot classification problem are based on meta-learning, which

can be categorized into three categories: model-based (Santoro et al. 2016; Munkhdalai and Yu 2017; Zhmoginov, Sandler, and Vladymyrov 2022), optimization-based (Finn, Abbeel, and Levine 2017; Antoniou, Edwards, and Storkey 2018), and metric-based (Vinyals et al. 2016; Snell, Swersky, and Zemel 2017). However, recent studies have demonstrated that fine-tuning methods outperform meta-learning methods on few-shot classification tasks (Chen et al. 2019; Dhillion et al. 2019; Chen et al. 2021; Tian et al. 2020; Chowdhury et al. 2021), especially for cross-domain few-shot learning (Guo et al. 2020), since meta-learning is more prone to overfitting on base classes (Dumoulin et al. 2021).

The common approach for cross-domain few-shot learning is to learn task-agnostic feature extractors using the base classes and then fine-tune with task-specific weights to handle the novel classes (Requeima et al. 2019; Dvornik, Schmid, and Mairal 2020; Tao et al. 2022; Bateni et al. 2020; Li, Liu, and Bilen 2021, 2022; Triantafillou et al. 2021; Liu et al. 2021). These task-specific weights are decided using an auxiliary system (Dvornik, Schmid, and Mairal 2020; Liu et al. 2020; Triantafillou et al. 2021; Requeima et al. 2019; Bateni et al. 2020; Liu et al. 2021), or through training from scratch regardless of the task (Tao et al. 2022; Li, Liu, and Bilen 2021, 2022). Recent methods have utilized adapters as task-specific weights to efficiently balance between domain adaptation and overfitting (Requeima et al. 2019; Bateni et al. 2020; Li, Liu, and Bilen 2022). These methods either adopt FiLM layers (Perez et al. 2018) as adapters (Requeima et al. 2019; Bateni et al. 2020) or residual adapters parameterized by matrix (Li, Liu, and Bilen 2022).

2.2 Knowledge Distillation for Few-shot Learning

Knowledge distillation is a popular technique in deep learning that transfers the knowledge from a large, pre-trained model (teacher) to a smaller, more efficient model (student) (Hinton, Vinyals, and Dean 2015). Knowledge distillation has also been applied to few-shot learning (Tian et al. 2020; Rajasegaran et al. 2020), primarily using self-distillation with identical architectures. Unlike previous approaches, we demonstrate that the knowledge distillation can be useful even during the fine-tuning stage with similar architectures, serving as a regularizer for few-shot training.

3 Method

This section presents a novel framework, termed Leveraging Normalization Layer in Adapters with Progressive Learning and Adaptive Distillation (**ProLAD**), designed for CD-FSL. Our approach incrementally trains two adapters, **TA** and **TAN** with adaptive distillation. An overview of our method is depicted in Figure 2, and the detailed pseudo-algorithm is provided in Appendix C.

3.1 Training Adapters with Progressive Learning and Distillation

Design of Adapters Before delving into each stage, we first provide an overview of the two adapters, **TAN** and **TA**, as depicted in Figure 3. First adapter, **TA**, parameterized by ϕ_w , employs 1×1 convolutional adapter embedded within

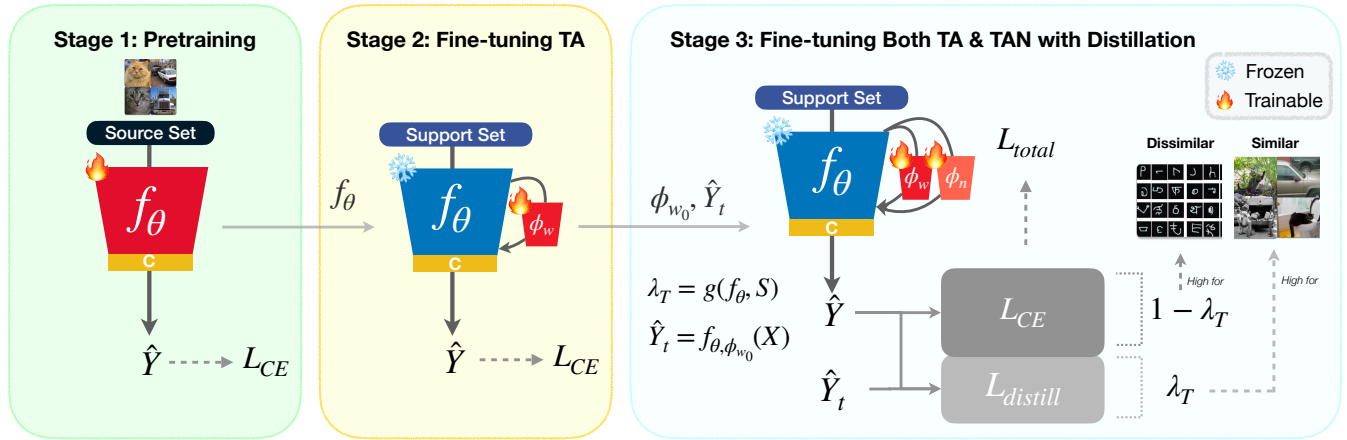


Figure 2: Overview of our method. In stage 1, the feature extractor is pretrained on the source dataset. In stage 2, an adapter **TA** is fine-tuned using the support set data, while the backbone remains frozen. In stage 3, with **TA** initialized from the outcomes of stage 2 and adaptive distillation applied using the logits from the stage 2 model for teacher prediction, both adapters are further fine-tuned on the support set. The classifier c is reinitialized at the beginning of each stage. The adaptive coefficient, λ_T , varies based on domain similarities.

each 3×3 convolutional layer without a normalization layer, following Li, Liu, and Bilen (2022). The role of **TA** is pivotal in adapting to similar domains, as it does not interfere with the well-represented batch statistics.

In contrast, **TAN**, parameterized by ϕ_n , is equipped with a Standard Normalization layer (SN), which is a batch normalization layer without affine layer that integrates the statistics of the target domain, combined with a group convolutional layer. The inclusion of this normalization layer in **TAN** is crucial for domain adaptation in dissimilar domains, as it gathers additional statistics from the target domain. Furthermore, we introduce an extra group convolution layer to accommodate larger receptive fields with fewer parameters, replacing the role of an affine layer. Detailed insights into these design choices are explored in Section 4.3.

The effects of these adapters are summarized in Table 1. More detailed results are presented in Section 4.3. As observed, **TAN** excels in dissimilar domains (+4.1%), highlighting the efficacy of the normalization layer in handling low domain similarity. This finding supports that incorporating the batch statistics of the target domain is essential for dissimilar domains. Conversely, **TA** performs better in similar domains (+2.8%), suggesting that the noisy nature of statistics can hinder the training process for high domain similarity. This observation paves the way for the introduction of an adaptive mechanism designed to amplify the activation of **TAN** for dissimilar domains and **TA** for similar ones, which will be discussed subsequently.

Step-by-Step Methods As discussed, training one of the two adapter choices is advantageous either for similar or dissimilar domains, but not both. This necessitates an adaptive method to function as a switch, allowing us to selectively employ primary adapters based on domain similarity. Motivated from this insight, our method comprises three distinct training stages with these adapters.

Adapter	Avg. Similar	Avg. Dissimilar
TA	68.0	81.7
TAN	65.2	85.8

Table 1: Summary of accuracy using each proposed adapter based on domain similarity.

In stage 1, we pretrain the feature extractor, parameterized by θ , using source data, as a transfer learning approach has been demonstrated to be effective for CD-FSL (Guo et al. 2020). Formally, let $D_b = \{(\mathbf{x}, y)\}$ be a source dataset comprising image and label pairs. We train the model on this source dataset using the cross-entropy loss, defined with dataset D and embedding function f , as follows:

$$L_{CE}(f, D) = - \sum_{\mathbf{x}, y \in D} y \log(f(\mathbf{x})) \quad (1)$$

For pretraining, we use $L_{CE}(f_\theta, D_b)$ as a loss function.

In stage 2, our primary objective is to train **TA** while keeping the feature extractor frozen and then save the features to be used as the teacher, denoted by \hat{Y}_t . This phase plays a crucial role in steering the subsequent training of **TAN**, as both the teacher features and the initialized parameters of **TA** are determined at this stage. In particular, let $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^{|S|}$ be a support set consisting of $|S|$ image-label pairs, which is divided into $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^{|S|}$ and $Y = \{y_i\}_{i=1}^{|S|}$. We train **TA** using $L_{CE}(f_{\theta, \phi_w}, S)$ as a loss function. Then, \hat{Y}_t is calculated with $\hat{Y}_t = f_{\theta, \phi_w}(X)$.

In stage 3, we proceed to train both **TAN** and **TA**, initializing **TA** with the values trained in the second stage, denoted as ϕ_{w_0} , to set a starting point for further training. Moreover,

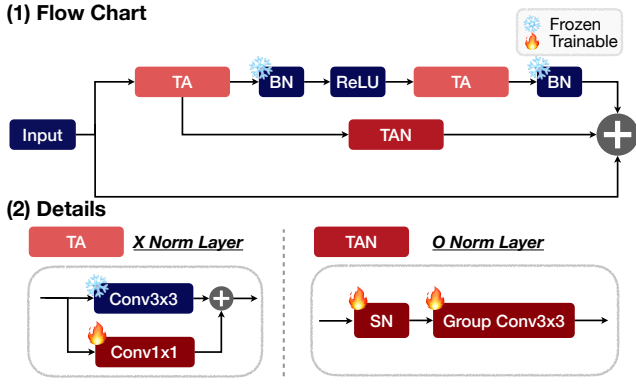


Figure 3: The entire structure of our adapters for each ResNet block. **TA** is 1×1 convolutional adapter parameterized by matrix. **TAN** consists of normalization layer (SN) without affine layer followed by group convolutional layer.

we apply distillation using the teacher \hat{Y}_t . This distillation is modulated based on domain similarities by an adaptive coefficient $\lambda_{\mathcal{T}}$. This coefficient is determined by the function $g(f_{\theta}, S)$, giving us an approximation of the domain shift magnitude. With this adaptive distillation, our method can select the required adapters depending on domain similarities. In the case of similar domains, our method prompts a more distillation and hence relies on **TA** for adaptation since the activation of **TAN** is reduced to follow the teacher \hat{Y}_t . On the contrary, for dissimilar domains, our method prioritizes the influence of the ground truth one-hot labels during training, applying minimal distillation and consequently allowing **TAN** to play a more dominant role in adaptation. As the classification depends on the cosine distance, distillation loss is computed using cosine similarity defined as:

$$\text{cos}(f_{\theta, \phi_n, \phi_w}, X, \hat{Y}_t) = \sum_{\mathbf{x} \in X, \hat{y}_t \in \hat{Y}_t} \frac{f_{\theta, \phi_n, \phi_w}(\mathbf{x})^T \hat{y}_t}{\|f_{\theta, \phi_n, \phi_w}(\mathbf{x})\|_2 \|\hat{y}_t\|_2} \quad (2)$$

where $L_{\text{distill}} = 1 - \text{cos}$. As a result, the final training loss for this stage is defined as:

$$L_{\text{total}} = (1 - \lambda_{\mathcal{T}}) \cdot L_{\text{CE}}(f_{\theta, \phi_w, \phi_n}, S) + \lambda_{\mathcal{T}} \cdot L_{\text{distill}}(f_{\theta, \phi_n, \phi_w}, X, \hat{Y}_t)$$

Subsequently, we explore potential methods to compute $\lambda_{\mathcal{T}}$.

3.2 Domain Adaptive Coefficient for Distillation

An adaptive coefficient $\lambda_{\mathcal{T}}$ is essential for modulating the activation of **TAN**, thereby selectively utilizing additional target statistics based on domain similarity. Several approaches could be employed to design the g function to compute $\lambda_{\mathcal{T}}$, aiming to estimate domain similarity without additional parameters. In the following sections, we will explore two such variants: one based on similarities between features, and the other on performance metrics. An in-depth exploration of the impact of $\lambda_{\mathcal{T}}$ is detailed in Section 4.3.

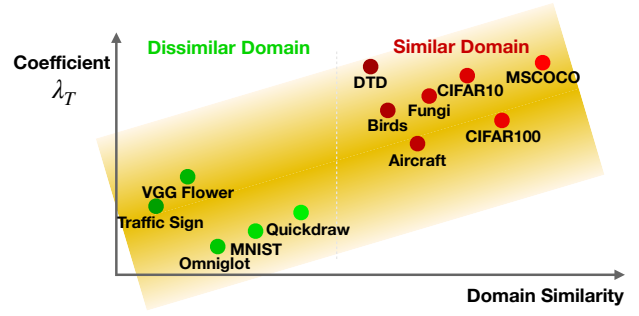


Figure 4: Correlation between our adaptive coefficient using similarity difference and domain similarity using EMD.

Difference between Inter-class and Feature Similarities

To estimate domain similarity, we utilize embedded feature similarities. Two metrics are defined in classification tasks: intra-class similarity, which measures the similarity of features within each class, and inter-class similarity, which evaluates the similarity between class prototypes. The distinction between these similarities helps determine how individual class features cluster relative to the entire feature distribution, as a larger difference in these similarities implies a more pronounced domain similarity. However, for classes represented by a single sample, computing the intra-class similarity is not feasible. There, our method uses feature similarities, which measure the similarity across all feature pairs, as an alternative. Hence, our method estimates domain similarity by comparing the distinctness of class prototypes to the similarity between pairs of individual samples.

To determine both types of similarities, we compute the cosine similarities separately. Let $\text{cos}_{ij}^{\text{feat}}$ denote the cosine similarity between the i -th and j -th features, and $\text{cos}_{ij}^{\text{inter-cl}}$ denote the similarity between the class prototypes of the i -th and j -th classes. Each cosine similarity is evaluated using Eq. (2). The adaptive distillation coefficient is derived by considering the difference between the computed similarities, as follows:

$$\lambda_{\mathcal{T}} = \exp(\beta \cdot (\frac{1}{N_s} \cdot \sum_{i < j} \text{cos}_{ij}^{\text{feat}} - \frac{1}{N_c} \cdot \sum_{i < j} \text{cos}_{ij}^{\text{inter-cl}} | -1)) \quad (3)$$

In this equation, β is a scaling factor, N_s denotes the number of feature pairs from support set and N_c represents the number of class pairs. We employ the exponential function due to its value range being constrained between zero and one. According to this equation, a larger difference between the two similarity measures results in a higher $\lambda_{\mathcal{T}}$, which is favorable for similar domains as intended. Coefficients calculated with Eq. (3) exhibit a high correlation with domain similarity, as illustrated in Figure 4, showcasing the efficacy of this metric as a reliable estimator of domain similarity. We refer to this technique as **ProLAD-sim**.

Support Set Loss & Accuracy Another metric to estimate domain similarity utilizes performance indicators from the pretrained feature extractor on the support set, such as loss

and accuracy. These metrics are especially relevant in understanding task distribution shift (Luo, Xu, and Xu 2022; Luo et al. 2023) or task difficulty (Oh et al. 2022), both correlating with domain shift. A closer domain similarity is indicated by a lower loss and higher accuracy. Specifically, let acc denote the accuracy and $L_{orig} = L_{CE}(f_{\theta}, S)$ be the loss on the support set S , where the loss is computed using Eq. (1). The adaptive coefficient is then given by:

$$\lambda_{\mathcal{T}} = \exp(-\beta \cdot L_{orig} \cdot (1 - acc)) \quad (4)$$

As for **ProLAD-sim**, we employ the exponential function with β as a scaling factor. This equation implies that higher losses and lower accuracies yield lower coefficients suited for dissimilar domains, while the inverse leads to higher coefficients optimal for similar domains. We term this approach **ProLAD-loss**.

4 Experiments

4.1 Experimental Setup

Datasets We utilize the Meta-Dataset (Triantafillou et al. 2019), which is a CD-FSL benchmark consisting of 13 different datasets. We evaluate our method using the standard setting as described in the original paper. For the Single-Domain Learning (SDL) setting, we first train the feature extractor on ImageNet and then test it on all 13 datasets in the Meta-Dataset. On the other hand, for the Multi-Domain Learning (MDL) setting, we train the feature extractor on 8 datasets and then evaluate its performance on all 13 datasets. Mean accuracy over the 600 episodes are reported with 95 confidence interval. For the analysis of domain similarities in the SDL setting, we assess domain similarities using the EMD distance (Rubner, Tomasi, and Guibas 1998). We classify Traffic Sign, VGG Flower, MNIST, Omniglot, and Quick Draw as the five most dissimilar domains, while the others are deemed similar. Details are in Appendix B.

Implementation Details We use the standard ResNet18 (He et al. 2016) model as the feature extractor, which is a commonly used architecture for CD-FSL (Bateni et al. 2020; Dvornik, Schmid, and Mairal 2020; Liu et al. 2020; Li, Liu, and Bilen 2021). For the SDL setting, we adopt the training procedure proposed in Dvornik, Schmid, and Mairal (2020); Li, Liu, and Bilen (2022). In the MDL setting, we use Universal Representation Learning (URL) (Li, Liu, and Bilen 2021) to train on multiple domains to train across multiple domains, thereby avoiding negative transfer.

For fine-tuning adapters, we utilize the Adadelta optimizer (Zeiler 2012) with a learning rate of 0.5. We set the scaling coefficient β in distillation to the value of 1.5 for both variants of **ProLAD**. We use a large momentum size of 0.8 for normalization layer because the feature values can change significantly between epochs. Additionally, we employ a group size of 8 for the group convolutional network. More details can be found in the Appendix A.

Baselines For the baselines, we compare our methods to state-of-the-art methods, including standard fine-tuning and Protonet-based methods (Triantafillou et al. 2019). We also consider adapters with meta-net such as Simple CNAPS

(Bateni et al. 2020), URT (Liu et al. 2020), FLUTE (Triantafillou et al. 2021), tri-M (Liu et al. 2021), Transductive CNAPS (Bateni et al. 2022), selective method like SUR (Dvornik, Schmid, and Mairal 2020), hyperparameter-optimization-based method of BOHB (Saikia, Brox, and Schmid 2020), and simple linear adapter-based methods like TSA (Li, Liu, and Bilen 2022) and URL (Li, Liu, and Bilen 2021). For TSA (Li, Liu, and Bilen 2022) in the SDL setting, we have reproduced the results ourselves, incorporating a data shuffling modification.

4.2 Main Results

Single-Domain Learning (SDL) Table 2 presents the results of our proposed method on the Meta-Dataset with the SDL setting. A feature extractor is trained on the ImageNet dataset and evaluation is carried out on all datasets, including ImageNet, with different classes. Both of our methods surpass other techniques in terms of average accuracy. Specifically, for the majority of domains, our methods excel: **ProLAD-loss** performs better in 8 out of 13 domains, and **ProLAD-sim** in 11 out of 13. Significant improvements are observed in domains distinctly dissimilar to ImageNet (+4.9%), such as Omniglot (+7.9%) and Traffic Sign (+7.7%), attributed to the high activation of **TAN**. For domains closely related to the base domain, like CIFAR-10 (+0.5%), CIFAR-100 (-0.6%), and MSCOCO (-0.3%), our method demonstrates performance comparable to TSA (Li, Liu, and Bilen 2022) (+0.8%), which trains **TA** from scratch, showing that our method primarily utilizes **TA** for similar domains.

Additionally, for most domains, **ProLAD-sim** that employs Eq. (3) for coefficient demonstrates superior performance over **ProLAD-loss**, which utilizes Eq. (4) for its coefficient. This performance indicates a more favorable scenario for computing $\lambda_{\mathcal{T}}$: a stronger correlation with domain similarity over other considerations. While loss and accuracy on the support set correlate with domain similarity, they more closely align with the notion of task-level shift (Luo, Xu, and Xu 2022; Luo et al. 2023) and task difficulty (Oh et al. 2022). However, variations in similarities directly correlate with domain similarity, as illustrated in Figure 4. Given the evident superiority of **ProLAD-sim** over **ProLAD-loss**, our subsequent experiments will primarily focus on **ProLAD-sim**.

Multi-Domain Learning (MDL) Table 3 presents the overall results of our proposed method on the Meta-Dataset under the MDL setting. More detailed results for each dataset, along with their respective confidence intervals, is presented in Appendix D. The MDL setting is considered less challenging than the SDL setting because the feature extractor is pre-trained on multiple domains, specifically eight domains, which increases the likelihood that some of the domain properties of the test set have already been learned. Nevertheless, both **ProLAD-loss** and **ProLAD-sim** still outperform other methods in both seen and unseen domains, although not to the same degree as in the SDL setting, given the limited presence of domains significantly distinctive from the source. Overall, the results demonstrate the ef-

	Finetune	Protonet	Proto-MAML	BOHB	Simple CNAPS	FLUTE	TSA	ProLAD-loss	ProLAD-sim
ImageNet	45.8 ± 1.1	50.5 ± 1.1	49.5 ± 1.1	51.9 ± 1.1	54.8 ± 1.2	46.9 ± 1.1	57.1 ± 1.2	57.8 ± 1.2	57.2 ± 1.1
Omniglot	60.9 ± 1.6	60.0 ± 1.4	63.4 ± 1.3	67.6 ± 1.2	62.0 ± 1.3	61.6 ± 1.4	76.2 ± 1.2	82.7 ± 1.1	84.1 ± 1.2
Aircraft	68.7 ± 1.3	53.1 ± 1.0	56.0 ± 1.0	54.1 ± 0.9	49.2 ± 0.9	48.5 ± 1.0	71.9 ± 1.0	74.8 ± 1.2	76.1 ± 1.2
Birds	57.3 ± 1.3	68.8 ± 1.0	68.7 ± 1.0	70.7 ± 0.9	66.5 ± 1.0	47.9 ± 1.0	74.4 ± 0.9	74.1 ± 1.0	75.5 ± 1.0
Textures	69.0 ± 0.9	66.6 ± 0.8	66.5 ± 0.8	68.3 ± 0.8	71.6 ± 0.7	63.8 ± 0.8	76.9 ± 0.7	77.1 ± 0.8	77.7 ± 0.8
Quick Draw	42.6 ± 1.2	49.0 ± 1.1	51.5 ± 1.0	50.3 ± 1.0	56.6 ± 1.0	57.5 ± 1.0	66.4 ± 0.9	69.6 ± 1.0	70.6 ± 1.0
Fungi	38.2 ± 1.0	39.7 ± 1.1	40.0 ± 1.1	41.4 ± 1.1	37.5 ± 1.2	31.8 ± 1.0	46.7 ± 1.1	46.1 ± 1.1	46.8 ± 1.2
VGG Flower	85.5 ± 0.7	85.3 ± 0.8	87.2 ± 0.7	87.3 ± 0.6	82.1 ± 0.9	80.1 ± 0.9	91.2 ± 0.5	92.2 ± 0.6	92.9 ± 0.6
Traffic Sign	66.8 ± 1.3	47.1 ± 1.1	48.8 ± 1.1	51.8 ± 1.0	63.1 ± 1.1	46.5 ± 1.1	81.9 ± 1.0	89.6 ± 0.9	89.4 ± 0.9
MSCOCO	34.9 ± 1.0	41.0 ± 1.1	43.7 ± 1.1	48.0 ± 1.0	41.4 ± 1.0	45.8 ± 1.0	55.7 ± 1.0	54.6 ± 1.1	55.4 ± 1.1
MNIST	-	-	-	-	-	80.8 ± 0.8	92.9 ± 0.6	95.5 ± 0.5	95.8 ± 0.5
CIFAR-10	-	-	-	-	-	65.4 ± 0.8	79.4 ± 0.7	79.2 ± 0.8	79.7 ± 0.8
CIFAR-100	-	-	-	-	-	52.7 ± 1.1	70.9 ± 1.0	70.1 ± 0.9	70.3 ± 1.0
Avg. seen	45.8	50.5	49.5	51.9	54.8	46.9	57.1	57.8	57.2
Avg. unseen	58.2	56.7	58.4	60.0	59.3	53.2	73.7	75.4	76.2
Avg. similar	53.6	53.8	55.0	56.5	53.2	51.0	68.0	68.0	68.8
Avg. dissimilar	64.0	60.4	62.7	64.3	66.0	65.3	81.7	85.9	86.6
Avg. all	57.0	56.1	57.5	59.2	58.9	52.6	72.4	74.0	74.7

Table 2: Comparison to state-of-the-art methods on Meta-Dataset with the single-domain setting where the feature extractor is trained only on ImageNet and then test on all datasets. Mean accuracy, 95 confidence interval are reported.

	Simple CNAPS	Transductive CNAPS	SUR	URT	FLUTE	tri-M	URL	TSA	ProLAD-loss	ProLAD-sim
Avg. seen	74.6	75.1	75.2	76.7	76.2	74.5	80.0	80.2	80.6	80.9
Avg. unseen	65.8	66.5	63.1	62.2	69.9	72.9	69.3	77.2	78.2	78.6
Avg. all	71.2	71.8	70.5	71.1	73.8	73.9	75.9	79.0	79.7	80.0

Table 3: Comparison to state-of-the-art methods on Meta-Dataset with the multi-domain setting where the feature extractor is trained on 8 datasets and then test on all datasets. Mean accuracy is reported.

fectiveness of our method in the easier MDL setting as well.

4.3 Ablation Study

Performance of Each Adapter Table 4 shows the performance of each adapter in the SDL setting. **TA** excels in domains closely related to the base domains (+2.6%), whereas **TAN** performs better on tasks with low domain similarity (+4.1%). These findings corroborate our hypothesis: while incorporating batch statistics from the target domain aids performance in dissimilar domains, it might be less beneficial for similar domains due to the noisy statistics from few-shot scenarios. Furthermore, training both adapters from scratch yields performance on par with **TAN**, suggesting that the activation strength of **TAN** surpasses that of **TA**. Thus, in the absence of distillation, activation of **TAN** predominantly serves the adaptation, highlighting that distillation can guide the selection of primary adapters during fine-tuning. Our approach exhibits performance akin to **TA** in similar domains and mirrors **TAN** in dissimilar domains, implying that our method effectively increases the activation of optimal adapter based on the domain similarity.

Progressive Learning and Adaptive Distillation Table 5 illustrates the impact of progressive learning and adaptive

distillation in stage 3. For this analysis, we exclusively employ **ProLAD-sim**, where the coefficient is determined using Eq. (3). Collectively, both progressive learning and adaptive distillation emerge as crucial for regularization during the concurrent training of both adapters; this is evidenced by the peak performance exhibited by **ProLAD**. Notably, adaptive distillation considerably enhances performance, indicating the pivotal role it plays in dynamically controlling the activation of the two adapters.

Efficacy of Adaptive Coefficient Figure 5 illustrates the performance based on a fixed coefficient $\lambda_{\mathcal{T}}$ in conjunction with the average coefficient of **ProLAD-sim** and **ProLAD-loss**. For similar domains, the results suggest that while a higher coefficient is generally favorable, an excessively large coefficient leads to decreased performance. This decline stems from the potential risks of relying solely on distillation, given the difference architectures of student and teacher. For dissimilar domains, a lower coefficient is typically preferred as **TAN** performs better in these scenarios. The coefficient of **ProLAD-sim** closely matches the optimal fixed coefficient and outperforms all fixed coefficients. This is because, within the same domain, it can modulate the distillation based on target classes. In contrast,

Adapter	Image-net	Omni-glot	Airc-craft	Birds	Tex-tures	Quick Draw	Fun-gi	VGG Flower	Traffic Sign	MS -COCO	MN -IST	CIFAR -10	CIFAR -100	Avg. Similar	Avg. Dis-similar	Avg. all
TA	57.1	76.2	71.9	74.4	76.9	66.4	46.7	91.2	81.9	55.7	92.9	79.4	70.9	68.0	81.7	72.4
TAN	55.2	84.1	73.2	71.6	75.7	69.0	42.5	91.2	89.8	50.9	94.9	76.2	66.6	65.2	85.8	72.4
TAN + TA	55.4	83.4	72.1	71.7	74.3	69.0	42.4	91.6	89.7	49.6	95.7	75.1	65.9	64.4	85.9	72.0
ProLAD-sim	57.2	84.1	76.1	75.5	77.7	70.6	46.8	92.9	89.4	55.4	95.8	79.7	70.3	68.8	86.6	74.7

Table 4: The ablation study on the adapters in the SDL setting. The term ‘Adapter’ refers to the combination of adapters trained from scratch.

Method	Image-net	Omni-glot	Airc-craft	Birds	Tex-tures	Quick Draw	Fungi	VGG Flower	Traffic Sign	MS -COCO	MNIST	CIFAR10	CIFAR100	Avg. all
S	55.4	83.4	72.1	71.7	74.3	69.0	42.4	91.6	89.7	49.6	95.7	75.1	65.9	72.0
P	56.2	83.3	74.3	71.2	73.8	69.7	43.4	91.4	90.7	50.7	95.5	73.9	66.8	72.4
D + A	56.6	82.7	74.5	74.1	76.8	69.5	45.5	92.3	89.7	54.7	95.9	79.3	69.3	73.9
P + D + A (ours)	57.2	84.1	76.1	75.5	77.7	70.6	46.8	92.9	89.4	55.4	95.8	79.7	70.3	74.7

Table 5: Results from the ablation study examining progressive learning and adaptive distillation on the Meta-Dataset within the SDL setting. ‘S’ denotes training both adapters from scratch, ‘P’ represents progressive learning, ‘D’ signifies distillation, while ‘A’ corresponds to adaptive coefficient calculated with Eq. (3).




Figure 5: Performance by coefficient λ_T based on domain similarities. The green mark denotes the maximum value point with a fixed coefficient. The red and purple marks represent the performance using Eq. (3) and Eq. (4), respectively, where the value of x is the average coefficient across all samples. Our estimated coefficient value closely aligns with the ideal coefficient of a fixed value.

ProLAD-loss calculates λ_T similarly to the optimal value for similar domains but differs for dissimilar domains. This difference arises because even though these dissimilar domains are distinct from the source, they are relatively easy tasks, resulting in higher accuracy. Comprehensive results for each dataset can be found in Appendix D.

Components of TAN Adapter Finally, we present an ablation study on the individual components of **TAN** when combined with **ProLAD-sim**, as provided in Table 6. Full results are given in Appendix D. Our findings indicate that normalization is pivotal for adaptation; removing SN leads to a decline in performance. Furthermore, our analysis suggests that the FiLM layer is not adequate for feature modulation since its combination with SN does not enhance performance. While the group convolution underperforms

#	Architecture	#params	Avg. seen	Avg. unseen	Avg. all
(1)	FiLM	2.38%	57.2	74.4	73.0
(2)	Conv	8.57%	57.2	74.8	73.4
(3)	GroupConv	9.35%	56.0	73.9	72.6
(4)	SN + Conv	8.57%	56.0	75.6	74.1
(5)	SN + FiLM (BN)	2.38%	57.1	74.5	73.1
(6)	SN + GroupConv (Ours)	9.35%	57.2	76.2	74.7

Table 6: Ablation study on the **TAN** components. ‘Conv’ denotes a 1×1 convolutional layer, ‘FiLM’ refers to a FiLM layer. ‘#params’ indicates the number of parameters added for adapters relative to the total parameters in ResNet18.

on its own, its combination with SN delivers the best performance, underscoring the significance of large receptive fields in feature modification when the features are normalized. In summary, our proposed design surpasses other combinations with a similar number of additional parameters.

5 Conclusion

In this paper, we present a method of leveraging normalization layer in adapters with progressive learning and adaptive distillation. Our method employs two distinct adapters: one equipped with a normalization layer that leverages statistics from the target domain, and another without it. These adapters are trained using progressive learning and adaptive distillation. Our approach adeptly addresses the challenges posed by diverse test domains in CD-FSL by selecting the primary adapter based on domain similarity. Experimental results on the standard benchmark, Meta-Dataset, highlight our method’s outstanding performance.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by Korea government (MSIT) [No. 2021-0-00907, Development of Adaptive and Lightweight Edge-Collaborative Analysis Technology for Enabling Proactively Immediate Response and Rapid Learning, 90%] and [No. 2019-0-00075, Artificial Intelligence Graduate School Program (KAIST), 10%]. We thank Jimin Lee for helping us create figures of our motivation and methods.

References

- Antoniou, A.; Edwards, H.; and Storkey, A. 2018. How to train your MAML. *arXiv preprint arXiv:1810.09502*.
- Batani, P.; Barber, J.; Van de Meent, J.-W.; and Wood, F. 2022. Enhancing few-shot image classification with unlabelled examples. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2796–2805.
- Batani, P.; Goyal, R.; Masrani, V.; Wood, F.; and Sigal, L. 2020. Improved few-shot visual classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14493–14502.
- Bilen, H.; and Vedaldi, A. 2017. Universal representations: The missing link between faces, text, planktons, and cat breeds. *arXiv preprint arXiv:1701.07275*.
- Chen, W.-Y.; Liu, Y.-C.; Kira, Z.; Wang, Y.-C. F.; and Huang, J.-B. 2019. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232*.
- Chen, Y.; Liu, Z.; Xu, H.; Darrell, T.; and Wang, X. 2021. Meta-baseline: Exploring simple meta-learning for few-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9062–9071.
- Chowdhury, A.; Jiang, M.; Chaudhuri, S.; and Jermaine, C. 2021. Few-shot image classification: Just use a library of pre-trained feature extractors and a simple classifier. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9445–9454.
- Cui, Y.; Song, Y.; Sun, C.; Howard, A.; and Belongie, S. 2018. Large scale fine-grained categorization and domain-specific transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4109–4118.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dhillon, G. S.; Chaudhari, P.; Ravichandran, A.; and Soatto, S. 2019. A baseline for few-shot image classification. *arXiv preprint arXiv:1909.02729*.
- Du, Y.; Zhen, X.; Shao, L.; and Snoek, C. G. 2020. Metanorm: Learning to normalize few-shot batches across domains. In *International Conference on Learning Representations*.
- Dumoulin, V.; Houlisby, N.; Evci, U.; Zhai, X.; Goroshin, R.; Gelly, S.; and Larochelle, H. 2021. Comparing transfer and meta learning approaches on a unified few-shot classification benchmark. *arXiv preprint arXiv:2104.02638*.
- Dvornik, N.; Schmid, C.; and Mairal, J. 2020. Selecting relevant features from a multi-domain representation for few-shot classification. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, 769–786. Springer.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Guo, Y.; Codella, N. C.; Karlinsky, L.; Codella, J. V.; Smith, J. R.; Saenko, K.; Rosing, T.; and Feris, R. 2020. A broader study of cross-domain few-shot learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVII 16*, 124–141. Springer.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Li, H.; Chaudhari, P.; Yang, H.; Lam, M.; Ravichandran, A.; Bhotika, R.; and Soatto, S. 2020. Rethinking the hyperparameters for fine-tuning. *arXiv preprint arXiv:2002.11770*.
- Li, W.-H.; Liu, X.; and Bilen, H. 2021. Universal representation learning from multiple domains for few-shot classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9526–9535.
- Li, W.-H.; Liu, X.; and Bilen, H. 2022. Cross-domain few-shot learning with task-specific adapters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7161–7170.
- Li, Y.; Wang, N.; Shi, J.; Liu, J.; and Hou, X. 2016. Revisiting batch normalization for practical domain adaptation. *arXiv preprint arXiv:1603.04779*.
- Liu, L.; Hamilton, W.; Long, G.; Jiang, J.; and Larochelle, H. 2020. A universal representation transformer layer for few-shot image classification. *arXiv preprint arXiv:2006.11702*.
- Liu, Y.; Lee, J.; Zhu, L.; Chen, L.; Shi, H.; and Yang, Y. 2021. A multi-mode modulator for multi-domain few-shot classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8453–8462.
- Luo, X.; Wu, H.; Zhang, J.; Gao, L.; Xu, J.; and Song, J. 2023. A Closer Look at Few-shot Classification Again. *arXiv preprint arXiv:2301.12246*.
- Luo, X.; Xu, J.; and Xu, Z. 2022. Channel importance matters in few-shot image classification. In *International conference on machine learning*, 14542–14559. PMLR.
- Mirza, M. J.; Micorek, J.; Possegger, H.; and Bischof, H. 2022. The norm must go on: Dynamic unsupervised domain adaptation by normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14765–14775.

- Munkhdalai, T.; and Yu, H. 2017. Meta networks. In *International conference on machine learning*, 2554–2563. PMLR.
- Oh, J.; Kim, S.; Ho, N.; Kim, J.-H.; Song, H.; and Yun, S.-Y. 2022. Understanding Cross-Domain Few-Shot Learning Based on Domain Similarity and Few-Shot Difficulty. In *Advances in Neural Information Processing Systems*.
- Perez, E.; Strub, F.; De Vries, H.; Dumoulin, V.; and Courville, A. 2018. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Rajasegaran, J.; Khan, S.; Hayat, M.; Khan, F. S.; and Shah, M. 2020. Self-supervised knowledge distillation for few-shot learning. *arXiv preprint arXiv:2006.09785*.
- Requeima, J.; Gordon, J.; Bronskill, J.; Nowozin, S.; and Turner, R. E. 2019. Fast and flexible multi-task classification using conditional neural adaptive processes. *Advances in Neural Information Processing Systems*, 32.
- Rubner, Y.; Tomasi, C.; and Guibas, L. J. 1998. A metric for distributions with applications to image databases. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*, 59–66. IEEE.
- Saikia, T.; Brox, T.; and Schmid, C. 2020. Optimized generic feature learning for few-shot classification across domains. *arXiv preprint arXiv:2001.07926*.
- Santoro, A.; Bartunov, S.; Botvinick, M.; Wierstra, D.; and Lillicrap, T. 2016. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, 1842–1850. PMLR.
- Schroeder, B.; and Cui, Y. 2018. Fgvx fungi classification challenge 2018. Available online: github.com/visipedia/fgvx_fungi_comp (accessed on 14 July 2021).
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30.
- Tao, R.; Zhang, H.; Zheng, Y.; and Savvides, M. 2022. Powering finetuning in few-shot learning: Domain-agnostic bias reduction with selected sampling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 8467–8475.
- Tian, Y.; Wang, Y.; Krishnan, D.; Tenenbaum, J. B.; and Isola, P. 2020. Rethinking few-shot image classification: a good embedding is all you need? In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, 266–282. Springer.
- Triantafillou, E.; Larochelle, H.; Zemel, R.; and Dumoulin, V. 2021. Learning a universal template for few-shot dataset generalization. In *International Conference on Machine Learning*, 10424–10433. PMLR.
- Triantafillou, E.; Zhu, T.; Dumoulin, V.; Lamblin, P.; Evci, U.; Xu, K.; Goroshin, R.; Gelada, C.; Swersky, K.; Manzagol, P.-A.; et al. 2019. Meta-dataset: A dataset of datasets for learning to learn from few examples. *arXiv preprint arXiv:1903.03096*.
- Vinyals, O.; Blundell, C.; Lillicrap, T.; Wierstra, D.; et al. 2016. Matching networks for one shot learning. *Advances in neural information processing systems*, 29.
- Zeiler, M. D. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Zhmoginov, A.; Sandler, M.; and Vladymyrov, M. 2022. Hypertransformer: Model generation for supervised and semi-supervised few-shot learning. In *International Conference on Machine Learning*, 27075–27098. PMLR.