

Discriminative Sample-Guided and Parameter-Efficient Feature Space Adaptation for Cross-Domain Few-Shot Learning

Rashindrie Perera^{1,2} Saman Halgamuge¹

¹ University of Melbourne, Australia

² Peter MacCallum Cancer Centre, Australia

cdperera@student.unimelb.edu.au, saman@unimelb.edu.au

Abstract

In this paper, we look at cross-domain few-shot classification which presents the challenging task of learning new classes in previously unseen domains with few labelled examples. Existing methods, though somewhat effective, encounter several limitations, which we alleviate through two significant improvements. First, we introduce a lightweight parameter-efficient adaptation strategy to address overfitting associated with fine-tuning a large number of parameters on small datasets. This strategy employs a linear transformation of pre-trained features, significantly reducing the trainable parameter count. Second, we replace the traditional nearest centroid classifier with a discriminative sample-aware loss function, enhancing the model's sensitivity to the inter- and intra-class variances within the training set for improved clustering in feature space. Empirical evaluations on the Meta-Dataset benchmark showcase that our approach not only improves accuracy up to 7.7% and 5.3% on previously seen and unseen datasets, respectively, but also achieves the above performance while being at least $\sim 3\times$ more parameter-efficient than existing methods, establishing a new state-of-the-art in cross-domain few-shot learning. Our code is available at <https://github.com/rashindrie/DIPA>.

1. Introduction

Deep neural networks achieve remarkable performance when presented with abundant training data. However, collecting large datasets is not feasible in many applications due to bottlenecks in data (e.g., rare categories), or the cost of manual annotations. Inspired by this limitation, few-shot classification aims to learn a classifier to recognize new classes with only a limited number of samples per class [21, 22, 34]. In traditional few-shot settings, new classes arise within a previously seen domain but have no class overlap with previously seen classes. Therefore, early

works focused on learning to recognize new classes arising within a previously seen domain. However, in real-world scenarios, it is more likely that the new classes will arise from previously unseen domains. This challenging scenario of learning new classes in previously unseen domains is tackled in *cross-domain* few-shot learning problem, which is also the focus of this paper.

Existing methods typically address this challenge by first learning a set of *task-agnostic* (i.e., generalized) features from a large dataset. These features are then fine-tuned to a specific target task using a small training set (called *support set*), often comprising as few as five images per category [37, 55]. These two stages are called meta-training (or pre-training) and meta-testing, respectively. The performance of the fine-tuned model is subsequently evaluated on a separate set of test samples, the *query set*, where the objective is to accurately categorize each query sample into one of the classes represented in the support set. However, as illustrated in Table 1, existing methods frequently involve fine-tuning a considerable number of task-specific parameters during the meta-testing phase. This practice can lead to overfitting, especially when data is scarce [59]. Moreover, it has been shown that shallower neural network layers generally contain more general features that can be applied to new tasks without the need for explicit fine-tuning [60, 64]. Despite this, many existing studies [28, 37, 59] attempt to adapt the entire pre-trained model to the target task, which in a limited data regime, can increase the risk of overfitting.

On the other hand, existing works [37, 48, 53, 55] typically employ the nearest centroid classifier (NCC) [45, 52] for fine-tuning the task-specific parameters and for subsequent query classification. NCC assigns an image to the class with the closest centroid, where the centroid is the mean of the feature embeddings belonging to the class. It encourages each embedding to be closer to its respective class centroid than the centroids of the other classes. As a result, these classes form tight clusters in the feature embedding space (refer Fig. 1a). However, due to the lack of focus

Model	Backbone	# Total Params (T)	# Fine-tuned Params (F)	% Fine-tuned Params (F/T)*100
TSA [37], TSA+SSA [53]	ResNet-18	12.65 M	1.48 M	11.7%
PMF [28]	ViT-Sma11	21.72M	21.72M	100.0%
eTT [59]	ViT-Sma11	21.72M	1.95 M	9.0%
ATTNSCALE [2]	ViT-Sma11	21.72M	0.26 M	1.2%
Ours, $d_t = 12$	ViT-Sma11	21.72 M	0.10 M	0.5% ($\downarrow 0.7$)
Ours, $d_t = 9$	ViT-Sma11	21.72M	0.08 M	0.4% ($\downarrow 0.8$)
Ours, $d_t = 7$	ViT-Sma11	21.72M	0.06 M	0.3% ($\downarrow 0.9$)

Table 1. Comparison of the number of parameters fine-tuned during the meta-testing stage by recent state-of-the-art methods, including our method (last two rows). For our method, d_t refers to tuning depth, where $d_t = 12$ fine-tunes every layer of the model while $d_t = 7$ fine-tunes only the last 7 layers of the model.

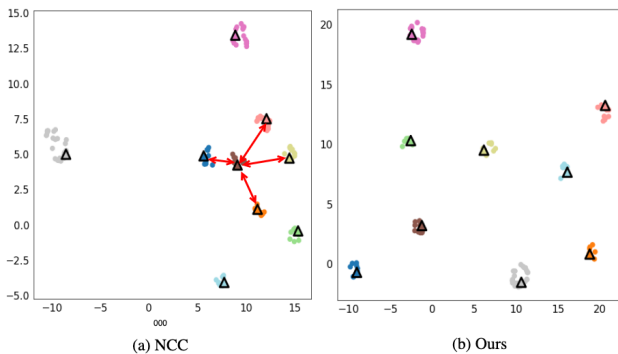


Figure 1. Visualization of feature space adaptation using the support set by employing (a) NCC and (b) Our approach for a few-shot classification task on ImageNet [51]. The Δ represents class centroids. The clusters formed by NCC are located close to each other, thereby, potentially generating confusing class centroids. In contrast, clusters formed in our approach are well separated.

on inter-class variance in NCC [45], there is a risk that the resulting clusters may not be adequately separated. This can lead to centroids being positioned too closely, as depicted in Fig. 1a, causing confusion during query classification.

Motivated by the above, as shown in Fig. 2, we propose to employ a set of lightweight task-specific parameters for adapting the task-agnostic feature backbone to an unseen domain. Unlike prior work [2, 28, 37, 59], we propose to employ a parameter-efficient linear transformation of features for the adaptation of a pre-trained model which significantly reduces the count of trainable parameters. Also, we propose to vary the tuning depth to fine-tune only the less transferable feature representations in the deeper layers of the network, thereby further reducing the number of trainable parameters, making it $3.3\times$ to $4.3\times$ more parameter-efficient than existing methods (refer Table 1 where $d_t = 9, 7$). Moreover, we propose to utilize the sample information present within the support set to identify challenging positive examples (i.e., the most dissimilar positive examples of a given class) or challenging nega-

tive examples (i.e., the most similar negative examples for a given class) for encouraging better intra and inter-class separation during fine-tuning to encourage better cluster formation in the feature space. Consequently, as shown in Fig. 1b, our approach forms well-separated clusters in the feature space, minimizing confusing centroids and thereby enhancing query classification accuracy. We advance few-shot learning with a straightforward and effective pipeline. We systematically evaluate our method on the standard cross-domain few-shot classification benchmark and demonstrate superior performance to all well-known methods.

2. Related Work

Below we provide a more detailed discussion of the most related work.

Learning task-agnostic representations that are easily adaptable to new tasks is crucial for cross-domain generalization success [7, 37]. Many existing methods [3, 17, 37, 48, 52, 53, 56] rely on supervised learning during meta-training to learn task-agnostic representations, but this approach can lead to *supervision collapse* [16, 17], where models only focus on features important for the source dataset’s classes and thereby ignore semantically important features of unseen classes. To circumvent this, our approach employs self-supervised learning for pre-training a task-agnostic feature extractor, ensuring broader learning beyond the source dataset’s class labels [2, 28, 59]

During meta-testing, task-agnostic representations are adapted (or fine-tuned) to the target task using the support set. Simple CNAPS [3] and FLUTE [56] employ task-specific FiLM layers [47], which are connected serially to the backbone and involve affine transformations for feature extractor adaptation, with FiLM parameters estimated by a meta-trained auxiliary network. Unlike them, our method directly learns this adaptation on the support set, eliminating the need for auxiliary networks. TSA [37] adapts the full backbone of a pre-trained extractor using residual adapters, which are connected in parallel and involve matrix multiplications [49]. Similarly, eTT [59] and ATTNSCALE [2] utilize visual prompts or scaling matrices for task-specific adaptations. In contrast, our method simplifies this process by using lightweight, linear transformations for adapting features to the target task. URT [43] uses multiple domain-specific feature extractors with a task-specific fusion mechanism, which increases training costs. In contrast, we achieve adaptation with a single feature extractor and task-specific parameters. Additionally, rather than adapting the entire feature backbone like many existing works [2, 28, 37, 59], we focus on adapting the deeper network layers. Finally, we diverge from the common use of NCC loss for the fine-tuning process in prior studies [2, 28, 37, 52, 53, 55, 59] by employing a discriminative sample-aware loss function on the support set.

There are multiple existing works that do not fall under the typical task-agnostic and task-specific categorization discussed above. For instance, ProtoNet [52] utilizes class prototypes to classify query samples. CTX [17] extends ProtoNet by using attention mechanisms to create more task-aligned prototypes. SSA [53] improves upon TSA by augmenting the support set dataset to create more challenging training examples. PMF [28] fine-tunes the full feature backbone for the target domain with an added task sampling and learning rate selection strategy. Our approach is also benchmarked against these for a comprehensive performance comparison.

3. Methodology

A few-shot task typically contains support set $\mathcal{S} = \{x_i, y_i\}_{i=1}^{|\mathcal{S}|}$ containing $|\mathcal{S}|$ sample and label pairs and a query set $\mathcal{Q} = \{x_i, y_i\}_{i=1}^{|\mathcal{Q}|}$ containing $|\mathcal{Q}|$ sample and label pairs. The goal in few-shot classification is to learn a classifier using \mathcal{S} that accurately predicts the labels of \mathcal{Q} . Note that this paper focuses on the few-shot image classification problem, i.e. x and y denote an image and its label.

We follow existing work [20, 38, 59] and employ a two-step approach to solve this problem. In the first stage, we train a feature extractor f_θ to learn task-agnostic feature representations using a large dataset (called source dataset). In the second stage, we first adapt the task-agnostic features for a target task $(\mathcal{S}, \mathcal{Q})$ sampled from the target dataset, using task-specific weights learned from the support set \mathcal{S} . Subsequently, we use the adapted model to classify the query samples in \mathcal{Q} .

Neural architecture: As in [2, 28, 59], we employ a Vision Transformer (ViT) model as our backbone architecture.

3.1. Task-agnostic representation learning

Recent studies [2, 28, 59] propose utilizing the DINO self-supervised algorithm [9] to pre-train the task-agnostic feature representations. DINO focuses on learning from the ‘global-to-local’ relationship within image crops to derive insightful features. Our approach, however, draws on the emerging concept of masked image modelling (MIM) [14, 24, 25, 62] for pre-training the feature extractor f_θ . MIM involves masking certain patches of an image and prompting the model to reconstruct these masked portions utilizing the context of the unmasked patches. This method requires the model to infer not just the missing visual information but also the context in which it occurs, necessitating a deep understanding of the image content, thereby, facilitating the learning of semantically rich and generalized feature representations. As illustrated in our results, this generalization is key for learning task-agnostic feature representations that can easily adapt to new tasks.

3.2. Task-specific representation learning

The feature extractor f_θ trained using MIM is expected to provide a good starting point for more advanced image recognition tasks [62]. However, when dealing with new classes from unseen domains, these features often require fine-tuning to better handle the novel scenarios.

To adapt f_θ to the target task, we propose to attach a set of scaling and shifting offset parameters to the pre-trained ViT backbone to learn a linear transformation of features [41]. As discussed above, unlike existing works that employ matrix-based [2, 37] or prompt-based [59] calculations, we strictly apply these offsets as a linear transformation. Specifically, we only learn the amount of scale (γ) and shift (β) required for adapting the task-agnostic features to the target task, thereby, reducing the tunable parameter count to only 0.5% of the model parameters during meta-testing. Formally, in a ViT, for an input $x \in \mathbb{R}^{(P^2+1) \times e}$, where e is the embedding dimension, the output $y \in \mathbb{R}^{(P^2+1) \times e}$ (which is also the input for the next layer) is computed as,

$$y = \gamma \odot x + \beta \tag{1}$$

where $\gamma \in \mathbb{R}$ and $\beta \in \mathbb{R}$ are the amount of scale and shift applied on x , and \odot is the dot product. This approach draws intuition from the concept of feature distribution matching [41, 54]. Specifically, the aim is to tune the first-order (mean) and second-order (variance) statistics of feature distributions, enabling adjustment of pre-trained features to the target data distribution using only two parameters per layer.

We apply this linear transformation on the layer normalization, multi-layer perceptron, and multi-head self-attention layers of the ViT as illustrated in Fig. 2a. While the choice of tuned layers can be varied, we opt for our experiments to tune all three types of layers. We denote the final adapted feature extractor that includes the re-parameterized task-specific parameters as $f_{\hat{\theta}}$.

Varying the tuning depth. As discussed above, recent works [2, 28, 37, 59] commonly fine-tune the full backbone for the novel task. Instead, our method varies the extent of task-specific adaption of a pre-trained model for the novel task. This approach enables customized adaptation to new tasks by selectively fine-tuning pre-trained representations across the model’s layers to meet the specific requirements of each task. Specifically, we vary the depth d_t of task-specific parameters h_ψ attached to the pre-trained model (i.e., the depth of layers that are fine-tuned on the target task) where $0 < d_t < L$, and $h_\psi = \{\psi_j\}$, where $j = (L - d_t + 1), \dots, 12\}$ and $\psi_j = \{(\gamma_m, \beta_m)\}$, where $m = 1, \dots, 6\}$. This approach, which strategically reduces the number of tunable parameters during the meta-testing phase as detailed in Table 1, is designed to further mitigate the risk of overfitting on novel tasks.




Figure 2. Illustration of our framework during meta-testing. **(a)** A set of task-specific parameters $(\gamma_m, \beta_m) \in h_\psi$ are attached to a ViT backbone f_θ , up to a pre-defined tuning depth d_t , where $m = 1, \dots, 6$. **(b)** h_ψ is fine-tuned on the support set using a set of learnable anchors A_ϕ . **(c)** Query images are classified by assigning them to the nearest class centre using the fine-tuned model $f_{\hat{\theta}}$. L : Number of layers in the ViT, d_f : feature fusion depth, z_l : [cls] output from l^{th} ViT layer, Z : fused feature embedding, Norm: Layer Normalization, MLP: Multi-Layer Perceptron, and MHA: Multi-head Attention.

3.2.1 Discriminative sample-guided feature adaptation

We aim to provide discriminative sample-guided supervisory signals to the model, considering the relations between the samples in the support set. For this, we propose to utilize a discriminative sample-aware loss function. Similar to NCC, we assign a learnable class representative (prototype) for each class in the support set. However, unlike NCC, we treat these prototypes as *anchors* (denoted by A_ϕ) to which the samples can be attracted or repelled to learn a task-adapted embedding space. Moreover, unlike NCC which only associates the samples of a given class to its class prototype, we propose to utilize the anchors to associate with all the samples in the support set. This can be denoted as,

$$l_{A_\phi}(X) = \frac{1}{|A|} \sum_{a \in A} \left\{ \log(1 + \sum_{x \in X_a^+} e^{\alpha(\delta - s(x, a))}) + \log(1 + \sum_{x \in X_a^-} e^{\alpha(s(x, a) + \delta)}) \right\} \quad (2)$$

where for a set of embedding vectors X and an anchor $a \in A$, the loss encourages the cosine similarity $s(x, a)$ between a feature vector $x \in X$ and a to be larger (i.e., greater than a user-defined margin δ), if a corresponds to the anchor of the class x belongs to (denoted as $x \in X_a^+$), or smaller (i.e., less than $-\delta$) if x belongs to a different class (denoted as $x \in X_a^-$). $\alpha > 0$ is a user-defined scaling factor. $l_{A_\phi}(X)$ is known as the proxy anchor loss [32].

Specific to our work is the manner in which the loss modifies the feature space by considering the sample information present within the support set, as illustrated in Fig. 3.

Here, for a given anchor $a \in A$, it tries to pull a and its most challenging positive example (for instance, a_r and r_3 in Fig. 3a) together. Similarly, it tries to push a and its most challenging negative example (for instance, a_r and b_3 in Fig. 3b) apart. The gradient is larger (thicker lines) when the positive feature vector is far from a , and when the negative feature vector is close to a . In this manner, Eq. 2 considers the *relative difficulty (hardness)* of each sample based on the inter and intra-class variation present in the feature embedding space to determine the relative strength of pull and push force to be applied for each sample. As demonstrated in our results, such variance-guided gradients provide better supervisory signals for fine-tuning in comparison to the conventional NCC.




Figure 3. Loss gradient calculation during training. The example illustrates a scenario with three unique classes red, green and blue, denoted by r, g, b . The anchor of each class is coloured in black and denoted as a_{class} . **(a)** The gradients for positive samples of class r are computed based on the relative hardness of all positive samples, so as to pull harder positives more strongly (thicker black lines). **(b)** The gradient calculation for negative samples for class r considers the distribution of all negative samples and tries to push harder negatives more strongly (thicker black lines).

3.3. Feature fusion

Features from shallower layers of the network have been observed to be more transferrable to unseen domains than those from deeper layers [9, 60]. Yet, existing works [2, 2, 6, 28, 37, 59] only use the features from the last layer of the model to represent an image, i.e. feature depth $d_f = 1$. In contrast to them, we extract the $[cls]$ embedding from the last $d_f > 1$ layers of the ViT and concatenate them to form the final feature embedding for an input image as,

$$Z = \text{concat}(z_L, z_{L-1}, \dots, z_{L-d_f+1}) \quad (3)$$

where z_L is the $[cls]$ embedding output for an image from the L^{th} layer, $d_f > 1$ is a hyper-parameter, $z \in \mathbb{R}^{\text{dim}}$ and $Z \in \mathbb{R}^{h \times \text{dim}}$. Although this straightforward fusion strategy requires no additional training as some existing feature fusion strategies [40, 57, 61], we observe significant performance gains when we increase the fusion depth.

3.4. Query classification

Following the feature space adaptation, we combine the adapted feature extractor $f_{\hat{\theta}}$ with a NCC classifier to classify query samples. In NCC, a prototype c_n is calculated for each class n in the support set as the mean of all feature representations belonging to that class, defined as:

$$c_n = \frac{1}{|S_n|} \sum_{(x,y) \in S_n} f_{\hat{\theta}}(x), \quad n = 1, \dots, N \quad (4)$$

Here, $(x, y) \in S_n$ represents pairs of feature embeddings and their corresponding labels in the support class S_n . To classify a sample x_q , we employ cosine similarity as the distance metric $d(f_{\hat{\theta}}(x_q), c_n)$ to assign x_q to the class with the nearest centroid c_n [37, 52].

4. Experiments

4.1. Experimental Settings

Dataset: We use the standard cross-domain benchmark dataset, Meta-Dataset [55]. It contains images from 13 diverse datasets. We follow the standard protocol [55] for 1) multi-domain learning (MDL), where the train sets of eight datasets from Meta-Dataset are used for pre-training and 2) single-domain learning (SDL), where the train set of only ImageNet is used for pre-training. For fair comparisons, we also follow [28] for single-domain learning with extra data (SDL-E) setting, where the entire ImageNet dataset is used for pre-training. Refer to Supplementary for more details.

Evaluation: We sample 600 N -way- K -shot tasks from the test set of each dataset and report the average accuracy and 95% confidence score. Here, N denotes the number of classes and K denotes the number of examples per class present in the task where the convention is to sample the N ,

K and the number of query images uniformly at random (refer [55] Appendix.3 for more details). We evaluate the Meta-Dataset on MDL, SDL and SDL-E settings.

Architecture: We limit our experiments to the standard ViT-small architecture [18]. Recent ViT-based works [28, 59], including ours, typically resize input images to 224×224 pixels, deviating from earlier approaches that use 84×84 pixels [37, 56]. We acknowledge that higher resolution may impact classification performance due to richer visual information. Therefore, we include backbone architecture details in our results for transparent method comparison.

Pre-training: We follow the strategy proposed by [62] and mostly stick to the hyperparameter settings reported in their work. A batch size of 128, 800 epochs and 4 Nvidia A100 GPUs with 80GB each is used for pre-training. We refer to Supplementary for more details.

Fine-tuning: We experimentally determine the values for the hyperparameters, where $\delta = 0.1$, $\alpha = 32$, and $d_f = 4$ give optimal performance across domains. γ and β parameters were initialized with constant values of one and zero, i.e., without any randomness, ensuring the reproducibility of our results. We use two separate NAdam optimizers [19] with the learning rate of 0.005 for h_{ψ} and 5 for A_{ϕ} . We experimentally find that this learning rate combination works well across most tasks. Fine-tuning is performed across 80 iterations on a single Nvidia A100 GPU with 80GB. We refer to Supplementary Section 10.3 for more details on hyperparameters. Finally, we report results using $d_t = 7$ for domains seen during pre-training (in-domain) and $d_t = 9$ for previously unseen domains (out-of-domain) as they provided the optimal average results across domains. Section 6 provides more discussion on this choice.

5. Main Results

5.1. Comparison to state-of-the-art methods

From here onwards, we refer to our method as *DIPA*, **DI**iscriminative-sample-guided and **P**arameter-efficient feature **A**daptation. We evaluate DIPA with the feature extractors pre-trained under the MDL, SDL and SDL-E settings on Meta-Dataset and compare it to existing state-of-the-art methods in Table 2. To facilitate fair comparisons, given the varying meta (or pre)-training strategies and backbone architectures in prior research, we include these details alongside the results. The table is organized into two sections: one for in-domain (seen) and another for out-of-domain (unseen) dataset accuracies, including their overall averages. Additionally, we include a column for the average accuracy across commonly unseen domains, as some prior works do not report results for all the unseen domains considered in our work.

Multi-domain feature extractor. In Table 2 (rows A0-A6), we benchmark against state-of-the-art methods that

ID	Method	SS PT	Sup. MT	Back-bone	ImageNet Omniglot Aircraft Birds Textures QuickDraw Fungi VGGFlower										TrafficSign MSCOCO MNIST CIFAR10 CIFAR100					Avg. Seen	Avg. Unseen (com.)	Avg. Unseen (all)	Avg. All															
					In-domain					Out-of-domain					Out-of-domain																							
PT/MT=MDL																																						
A0	Simple CNAPS [3]		✓	RN18	58.4 ± 1.1	91.6 ± 0.6	82.0 ± 0.7	74.8 ± 0.9	68.8 ± 0.9	76.5 ± 0.8	46.6 ± 1.0	90.5 ± 0.5	57.2 ± 1.0	48.9 ± 1.1	94.6 ± 0.4	74.9 ± 0.7	61.3 ± 1.1	73.7	53.1	67.4	71.2																	
A1	URT [43]		✓	RN18	56.8 ± 1.1	94.2 ± 0.4	85.8 ± 0.5	76.2 ± 0.8	71.6 ± 0.7	82.4 ± 0.6	64.0 ± 1.0	87.9 ± 0.6	48.2 ± 1.1	51.5 ± 1.1	90.6 ± 0.5	67.0 ± 0.8	57.3 ± 1.0	77.4	49.9	62.9	71.8																	
A2	FLUTE [56]		✓	RN18	58.6 ± 1.0	92.0 ± 0.6	82.8 ± 0.7	75.3 ± 0.8	71.2 ± 0.8	77.3 ± 0.7	48.5 ± 1.0	90.5 ± 0.5	63.0 ± 1.0	52.8 ± 1.1	96.2 ± 0.3	75.4 ± 0.8	62.0 ± 1.0	74.5	57.9	69.9	72.7																	
A3	TSA [37]		✓	RN18	59.5 ± 1.0	94.9 ± 0.4	89.9 ± 0.4	81.1 ± 0.8	77.5 ± 0.7	81.7 ± 0.6	66.3 ± 0.8	92.2 ± 0.5	82.8 ± 1.0	57.6 ± 1.0	96.7 ± 0.4	82.9 ± 0.7	70.4 ± 0.9	80.4	70.2	78.1	79.5																	
A4	TSA + SSA [53]		✓	RN18	58.9 ± 1.1	95.6 ± 0.4	90.0 ± 0.5	82.2 ± 0.7	77.6 ± 0.7	82.7 ± 0.7	66.6 ± 0.8	93.0 ± 0.5	84.9 ± 1.1	58.1 ± 1.0	98.5 ± 0.4	82.9 ± 0.7	70.8 ± 0.9	80.8	71.5	79.0	80.1																	
A5	eTT [59]		✓	ViT-s	67.4 ± 1.0	78.1 ± 1.2	79.9 ± 1.1	85.9 ± 0.9	87.6 ± 0.6	71.3 ± 0.9	61.8 ± 1.1	96.6 ± 0.5	85.1 ± 0.9	62.3 ± 1.0			79.6	73.7	73.7	78.3																		
A6	DIPA		✓	ViT-s	70.9 ± 1.0	84.7 ± 1.1	86.3 ± 1.0	90.8 ± 0.8	88.6 ± 0.5	75.3 ± 0.8	66.6 ± 1.1	97.9 ± 0.3	91.3 ± 1.0	64.8 ± 1.0	96.9 ± 0.5	87.4 ± 0.6	81.2 ± 0.8	82.6	78.1	84.3	83.3																	
																	(+1.8)	(+4.4)	(+5.3)	(+3.2)																		
PT=SDL-E, MT=MDL																																						
B0	PMF [28]		✓	ViT-s	74.6	91.8	88.3	91.0	86.6	79.2	74.2	94.1	88.9	62.6			85.0	75.8	75.8	83.1																		
B1	ATTNSCALE [2]		✓	ViT-s	80.9	78.8	86.7	85.8	74.4	59.0	95.9	91.4	91.4	61.0			80.2	76.2	76.2	79.3																		
B2	DIPA		✓	ViT-s	77.3 ± 0.7	83.9 ± 1.1	86.0 ± 1.1	91.1 ± 0.7	88.8 ± 0.5	75.9 ± 0.8	62.4 ± 1.1	97.7 ± 0.3	91.7 ± 0.8	66.5 ± 0.9	97.2 ± 0.5	92.2 ± 0.5	84.5 ± 0.7	82.9	79.1	86.4	84.2																	
																	(+2.9)	(+10.2)	(+1.1)																			
PT/MT=SDL																																						
					In-domain					Out-of-domain																												
C0	ProtoNet [16]		✓	RN34	53.7 ± 1.1	68.5 ± 1.3	58.0 ± 1.0	74.1 ± 0.9	68.8 ± 0.8	53.3 ± 1.1	40.7 ± 1.1	87.0 ± 0.7	58.1 ± 1.1	41.7 ± 1.1			53.7	61.1	61.1	60.4																		
C1	CTX [16]		✓	RN34	62.8 ± 1.1	82.2 ± 1.0	79.5 ± 0.9	80.6 ± 0.9	75.6 ± 0.6	72.7 ± 0.8	51.6 ± 1.1	95.3 ± 0.4	82.7 ± 0.8	59.9 ± 1.0			62.8	75.6	75.6	74.3																		
C2	TSA [37]		✓	RN34	63.7 ± 1.0	82.6 ± 1.1	80.1 ± 1.0	83.4 ± 0.8	79.6 ± 0.7	71.0 ± 0.8	51.4 ± 1.2	94.0 ± 0.5	81.7 ± 0.9	61.7 ± 0.9	94.6 ± 0.5	86.0 ± 0.6	78.3 ± 0.8	63.7	76.2	78.7	77.5																	
C3	DIPA		✓	ViT-s	71.4 ± 0.9	84.3 ± 1.2	86.7 ± 1.0	88.2 ± 0.9	87.1 ± 0.6	74.6 ± 0.8	61.4 ± 1.2	97.4 ± 0.4	88.9 ± 1.0	65.2 ± 1.0	97.1 ± 0.5	88.5 ± 0.6	81.5 ± 0.8	71.4	81.5	83.4	82.5																	
																	(+7.7)	(+5.3)	(+4.7)	(+5.0)																		
PT=SDL-E, MT=SDL																																						
D0	PMF [28]		✓	ViT-s	74.7	80.7	76.8	85.0	86.6	71.3	54.8	94.6	88.3	62.6			74.7	77.9	77.9	77.5																		
D1	DIPA		✓	ViT-s	77.3 ± 0.7	84.1 ± 1.2	87.1 ± 1.0	90.5 ± 0.7	87.3 ± 0.6	75.4 ± 0.8	60.9 ± 1.1	97.5 ± 0.4	91.7 ± 0.8	66.5 ± 0.9	97.2 ± 0.5	92.2 ± 0.5	84.5 ± 0.7	77.3	82.3	84.6	84.0																	
																	(+2.6)	(+4.4)	(+6.7)	(+6.5)																		

Table 2. Comparison of state-of-the-art methods on Meta-Dataset using MDL, SDL and SDL-E settings where PT: Pre-training, MT: Meta-Training, RN: ResNet, ViT-s: ViT-small, com.: common, Avg.: Average, SS PT: indicates self-supervised pre-training and Sup. MT: indicates supervised meta-training. Mean accuracy and 95% confidence interval are reported, where available.

meta (or pre)-train solely under the MDL setting. Notably, DIPA outperforms eTT [59], which employs self-supervised pre-training, across all common seen and unseen domains (10 out of 10) and is up to $32.5\times$ more parameter-efficient (refer Table 1, $d_t = 7$ and 9). When compared to methods using supervised meta-training for their feature extractor f_θ (such as Simple CNAPS [3], URT [43], FLUTE [56], TSA [37], and SSA [53]), DIPA shows superior performance in most domains, particularly outperforming them in 4 out of 5 unseen domains. Improving performance in unseen domains is a significant challenge due to the vast difference between seen and unseen domains and the limited availability of labelled samples for new tasks. DIPA addresses this by employing lightweight linear transformations for feature adaptation, together with a discriminative sample-guided loss function. While TSA+SSA [53] also shows competitive results using MixStyle-like augmentation strategies [63], DIPA achieves even better performance in most unseen domains without such augmentations¹ while being up to $24.7\times$ more parameter-efficient (see Table 1). Specifically, DIPA significantly surpasses TSA+SSA [53] in Traffic Sign (+6.4), MS-COCO (+6.7), CIFAR-10 (+4.5), and CIFAR-100 (+10.4).

Multi-domain feature extractor with additional pre-training. In Table 2 (rows B0-B2), we compare DIPA with PMF [28] and ATTNSCALE [2], which evaluate their models under the MDL setting and use SDL-E for self-supervised pre-training. Note that, unlike DIPA and ATTNSCALE, PMF also employs MDL for meta-training. DIPA outperforms ATTNSCALE, achieving a 2.9% higher performance in unseen domains and 2.7% in seen domains, and is up to $4.3\times$ more parameter-efficient (see Table

¹Incorporating similar augmentation strategies as in [53] into DIPA could lead to further improvements, a prospect we reserve for future exploration.

1). Despite PMF incorporating an additional meta-training stage and task-specific learning rate selection [28], DIPA exceeds PMF’s unseen domain performance by 3.3%. This enhanced performance is achieved with significant computational efficiency, as DIPA requires only two stages: pre-training and fine-tuning, thus eliminating the need for additional meta-training or learning rate selection. Furthermore, DIPA’s fine-tuning process is remarkably efficient, using only up to 0.08% of the model parameters, a stark contrast to PMF’s 100% parameter utilization (see Table 1).

Single-domain feature extractor. We also evaluate our method using a feature extractor that is trained solely on the ImageNet domain [51], referred to as the single-domain learning (SDL) setting. The SDL setting poses greater challenges compared to multi-domain scenarios, given that the model is exposed to only one training domain but evaluated across multiple, including the test split of ImageNet and additional diverse domains. In Table 2 (rows C0-C3), we benchmark our method against existing methods (ProtoNet [52], CTX [17], and TSA [37]) with published results under this setting. Our method surpasses these methods across all 13 domains. It is important to acknowledge that our model utilizes self-supervised pre-training together with a backbone architecture that is substantially different from those compared here, which may affect the fairness of this comparison. Nonetheless, our method still demonstrates significant performance gains, with improvements of 7.7%, 4.7%, and 5% for seen, unseen (all), and all domains, respectively.

Single-domain feature extractor with additional pre-training. In Table 2 (rows D0-D1), we present a comparison of DIPA with the latest state-of-the-art method, PMF [28]. Here, PMF employs SDL-E for pre-training and an additional SDL stage for meta-training, a step not used in DIPA. Despite the absence of this meta-training stage, DIPA demonstrates superior performance across all 13 domains.

Specifically, it shows a 2.6% improvement in seen domains, 4.4% in commonly unseen domains, and 6.5% across all domains. This performance is achieved without requiring additional computational steps or the extensive fine-tuning of the entire backbone, as is necessary for PMF (see details in Table 1).

Overall, our results demonstrate that DIPA significantly outperforms state-of-the-art methods in cross-domain scenarios within both seen and unseen domains while being more parameter efficient.

6. Additional Analysis

In this section, we conduct ablation studies to evaluate each step in our method and validate their impact on improving the few-shot classification performance.

6.1. Impact of varying the fine-tuning and query classification approaches

First, we compare our fine-tuning strategy l_{A_ϕ} against the conventional mean-embedding-based NCC (NCC_{mean}) in Table 3 (rows A0 and A2). Here, employing l_{A_ϕ} for fine-tuning significantly improves classification performance, yielding up to 2.8% improvement in seen domains and 1.6% in unseen domains compared to using NCC_{mean} . Moreover, in terms of cluster formation, on the ImageNet dataset, l_{A_ϕ} consistently outperformed NCC_{mean} in terms of inter-cluster distance, intra-cluster distance, and silhouette index by 1%, 6% and 1%, respectively.

Next, we compare our query classification strategy - NCC_{mean} - with another alternative. As mentioned above, during fine-tuning, DIPA learns a set of anchors A_ϕ that provide strong supervision to form well-separated and compact clusters (refer Fig. 1b). This provides an opportunity to use the fine-tuned anchors in lieu of the mean class centroids used in NCC_{mean} for query classification. We denote this variant as NCC_{A_ϕ} . However, while anchors A_ϕ provide strong supervisory signals for cluster formations, we observe that the anchors are placed with a small offset from the mean centroid in some cases (refer Supplementary Fig. 8), suggesting that the mean class centroid which encapsulates the summary of the compact cluster might be a more representative class descriptor during inference. This hypothesis is also supported by our results, where combining l_{A_ϕ} with NCC_{mean} yielded an optimal performance increase of 3.9% and 5.3% on seen and unseen domains compared to $l_{A_\phi} + NCC_{A_\phi}$. Therefore, we employ A_ϕ anchors for cluster formation during fine-tuning and thereafter, use the well-defined clusters to obtain mean class centroids for query classification during inference.

6.2. Impact of varying the number of tuned layers

We study the impact of varying the number of tuned layers in f_θ . Fig. 4 illustrates the average accuracy variation

for each dataset in Meta-Dataset as we vary the number of adapted layers d_t , where $0 \leq d_t \leq L$ and $L = 12$ for ViT-small. From the results, we observe the following: 1) A universal value of d_t that provides optimal results for all domains is not available, 2) A general trend is present where domains that are relatively less challenging such as VGG Flower, CIFAR-10 etc., report higher accuracies for smaller values of d_t . In contrast, the relatively more challenging domains, such as Quickdraw, and Traffic Sign, report higher accuracies for larger values of d_t , 3) Most domains report higher accuracies with $d_t < 12$. i.e., our approach does not need task-specific parameters attached to each layer of the model (as in TSA [17], and eTT [59]) nor fine-tuning the entire backbone (as in PMF [28]). This facilitates an even further reduction in the count of trainable parameters. For example, as shown in Table 1, by reducing d_t from 12 to 7, we achieve a reduction in the parameter count by 0.2%.




Figure 4. Variation of accuracies for each dataset in Meta-Dataset as d_t varies in the MDL setting. Average results are reported while more detailed results can be found in Supplementary Table 8. The dotted lines represent the relatively more challenging datasets. For each dataset, the value of d_t that reports the highest accuracy is annotated with a dot.




Figure 5. Variation of average seen, unseen and total average accuracies across all domains in Meta-Dataset as the number of tuned layers d_t varies in the MDL setting. Average results are reported and more detailed results can be found in Supplementary Table 8. The value of d_t that reports the highest average for each line is annotated with a dot.

Fig. 4 demonstrates that customizing the depth d_t , on

	Fine-tune stage	Inference stage	ImageNet	Omniglot	Aircraft	Birds	Textures	QuickDraw	Fungi	VGGFlower	TrafficSign	MSCOCO	MNIST	CIFAR10	CIFAR100	Avg. Seen	Avg. Unseen	Avg. All
A0	NCC_{mean}	NCC_{mean}	68.0 ± 1.0	79.0 ± 1.4	81.8 ± 1.1	88.2 ± 0.9	87.3 ± 0.6	73.2 ± 0.8	63.4 ± 1.0	97.2 ± 0.4	89.0 ± 1.0	63.3 ± 1.0	95.2 ± 0.6	87.2 ± 0.7	78.8 ± 0.9	79.8	82.7	80.9
A1	l_{A_ϕ}	NCC_{A_ϕ}	65.5 ± 1.0	83.4 ± 1.1	81.8 ± 1.0	88.6 ± 0.8	87.0 ± 0.6	68.4 ± 0.9	57.0 ± 1.2	97.5 ± 0.3	86.7 ± 1.1	52.4 ± 1.2	96.8 ± 0.5	82.8 ± 0.8	76.3 ± 0.9	78.7	79.0	78.8
A2	l_{A_ϕ}	NCC_{mean}	70.9 ± 1.0	84.7 ± 1.1	86.3 ± 1.0	90.8 ± 0.8	88.6 ± 0.5	75.3 ± 0.8	66.6 ± 1.1	97.9 ± 0.3	91.3 ± 1.0	64.8 ± 1.0	96.9 ± 0.5	87.4 ± 0.6	81.2 ± 0.8	82.6	84.3	83.3

Table 3. Comparing the impact of using the traditional NCC (NCC_{mean}) vs. l_{A_ϕ} loss for fine-tuning and NCC_{mean} vs anchor-based NCC (NCC_{A_ϕ}) for inference. The results are reported for the MDL setting where the first eight datasets are seen during training and the last five datasets are unseen and used for testing only. Mean accuracy and 95% confidence interval are reported.

a per-domain basis, yields the best performance. This customization, however, adds extra user-defined parameters to our model. To simplify, we streamline our approach to use just two d_t values: $d_t = 7$ for seen domains and $d_t = 9$ for unseen domains, based on peak average performance as illustrated in Fig. 5. Moreover, opting for a higher d_t value for unseen domains aligns with the intuition that a more extensive adaptation, via tuning additional layers, is typically beneficial when dealing with unseen domains.

6.3. Impact of varying the feature fusion depth

We compare the performance variation for few-shot classification by varying the feature fusion depth d_f and report the average results in Table 4. First, we note that even without feature fusion (when $d_f = 1$), our method outperforms existing methods by a significant margin. Next, we observe that as we increase d_f in DIPA, the average accuracies generally show an increasing trend up to $d_f = 4, 6$, validating our choice of fusing features from multiple layers. However, the performance shows a gradual declining trend afterwards. One possible reason is that combining a higher degree of shallower layers (containing more generic patterns) will cause the model to lose focus on important domain-specific features in the deeper layers. This is also reflected in the results where $d_f = 8, 12$ report even lower accuracies than $d_f = 1, 2$.

Method	d_f	dim	Avg. Seen	Avg. Unseen	Avg. all
TSA + SSA [53]	1	512	80.8	79.0	80.1
eTT [59]	1	384	79.6	73.7	78.3
DIPA	1	384	82.3	83.7	82.9
	2	768	82.4	83.7	82.9
	4	1536	82.6	84.3	83.3
	6	2304	82.6	83.9	83.1
	8	3072	82.0	83.5	82.6
	12	4608	81.5	83.2	82.1

Table 4. Variation of average (Avg.) accuracies for seen, unseen and all domains as the feature fusion depth d_f and its corresponding feature representation’s dimension (dim) under the MDL setting. Mean results are reported and more detailed results can be found in Supplementary Table 9.

6.4. Impact of Pre-training

In Table 5, we evaluate the impact of using MIM for pre-training, as opposed to the more commonly used DINO pre-training in existing methods [2, 28, 59]. For our experiments, we leveraged the SDL-E setting, utilizing the readily available DINO checkpoint [9]. First, we compare the

DIPA model with DINO pre-training to the DINO-based PMF [28] model, which represents the state-of-the-art in the SDL-E setting. As shown in the first two rows of Table 5, DIPA improves performance by 5.8% compared to PMF, demonstrating its effectiveness even without our preferred self-supervised pre-training task. Next, we substitute the DINO-based feature extractor with a MIM-based one (refer Table 5 row 3). This change leads to a further increase in accuracy up to 6.5% over PMF, validating the effectiveness of our selected pre-training algorithm.

Method	Pre-Train	Avg. Seen	Avg. Unseen (com.)	Avg. Unseen (all)	Avg. All
PMF [28]	DINO	74.7	77.9	77.9	77.5
DIPA	DINO	75.9 (+1.2)	82.0 (+4.1)	83.9 (+6.0)	83.3 (+5.8)
DIPA	MIM	77.3 (+2.6)	82.3 (+4.4)	84.6 (+6.7)	84.0 (+6.5)

Table 5. The impact of varying the pre-training algorithms in SDL-E setting. Average (Avg.) results are reported and more detailed results can be found in Supplementary Table 10.

7. Conclusions and Limitations

In this study, we explore the efficient adaptation of neural networks for few-shot classification. Our approach leverages extremely lightweight linear transformations, optimized by a discriminative sample-aware loss function, to learn new classes and domains with a limited number of labelled samples. This method achieves state-of-the-art performance on the challenging Meta-Dataset benchmark while ensuring parameter efficiency.

Our method is not without limitations. The current approach applies a fixed linear transformation to every layer of the pre-trained model. Future improvements could enable these transformations to be defined flexibly, layer by layer, to suit the specific requirements of the target task. Moreover, instead of restricting tuning depth to only two values for seen and unseen datasets, which may lead to suboptimal outcomes, future research could explore defining optimal tuning depths customized for each dataset and task.

8. Acknowledgements

R.P. acknowledges the Melbourne Graduate Research scholarship and GCI Top-up scholarship. S.H. acknowledges ARC grant DP210101135. The authors thank Hai-Hang Wu, Tamasha Malepathirana, Nisal Ranasinghe and Maneesha Perera for proofreading.

References

- [1] Mohamed Afham, Salman Khan, Muhammad Haris Khan, Muzammal Naseer, and Fahad Shabbaz Khan. Rich semantics improve few-shot learning. In *32nd British Machine Vision Conference 2021, BMVC 2021, Online, November 22-25, 2021*, page 152. BMVA Press, 2021. [5](#)
- [2] Samyadeep Basu, Daniela Massiceti, Shell Xu Hu, and Soheil Feizi. Strong baselines for parameter efficient few-shot fine-tuning. *CoRR*, abs/2304.01917, 2023. [2](#), [3](#), [5](#), [6](#), [8](#)
- [3] Peyman Bateni, Jarred Barber, Raghav Goyal, Vaden Masrani, Jan-Willem van de Meent, Leonid Sigal, and Frank Wood. Beyond simple meta-learning: Multi-purpose models for multi-domain, active and continual few-shot learning. *CoRR*, abs/2201.05151, 2022. [2](#), [6](#)
- [4] Peyman Bateni, Jarred Barber, Jan-Willem Van de Meent, and Frank Wood. Enhancing few-shot image classification with unlabelled examples. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2796–2805, 2022. [5](#)
- [5] Luca Bertinetto, João F. Henriques, Philip H. S. Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. [1](#)
- [6] Hakan Bilen and Andrea Vedaldi. Universal representations: The missing link between faces, text, planktons, and cat breeds. *CoRR*, abs/1701.07275, 2017. [5](#)
- [7] Hakan Bilen and Andrea Vedaldi. Universal representations: The missing link between faces, text, planktons, and cat breeds. *arXiv preprint arXiv:1701.07275*, 2017. [2](#)
- [8] Schroeder Brigit and Cui Yin. Fgvcx fungi classification challenge, 2018. [1](#)
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. [3](#), [5](#), [8](#), [1](#)
- [10] Da Chen, Yuefeng Chen, Yuhong Li, Feng Mao, Yuan He, and Hui Xue. Self-supervised learning for few-shot image classification. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1745–1749. IEEE, 2021. [5](#)
- [11] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. [5](#)
- [12] Yinbo Chen, Zhuang Liu, Huijuan Xu, Trevor Darrell, and Xiaolong Wang. Meta-baseline: Exploring simple meta-learning for few-shot learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9062–9071, 2021. [5](#)
- [13] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014. [1](#)
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [3](#)
- [15] Guneet Singh Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [5](#)
- [16] Carl Doersch, Ankush Gupta, and Andrew Zisserman. CrossTransformers: Spatially-aware few-shot transfer. *Advances in Neural Information Processing Systems*, 2020-Decem(NeurIPS), 2020. [2](#), [6](#)
- [17] Carl Doersch, Ankush Gupta, and Andrew Zisserman. Crosstransformers: spatially-aware few-shot transfer. *Advances in Neural Information Processing Systems*, 33: 21981–21993, 2020. [2](#), [3](#), [6](#), [7](#)
- [18] 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. [5](#)
- [19] Timothy Dozat. Incorporating nesterov momentum into adam. 2016. [5](#)
- [20] Nikita Dvornik, Cordelia Schmid, and Julien Mairal. Selecting Relevant Features from a Multi-domain Representation for Few-shot Classification. *Eccv 2020*, 12355 LNCS:769–786, 2020. [3](#)
- [21] Li Fei-Fei, Robert Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(4):594–611, 2006. [1](#)
- [22] Michael Fink. Object classification from a single example utilizing class relevance metrics. In *Advances in Neural Information Processing Systems 17 [Neural Information Processing Systems, NIPS 2004, December 13-18, 2004, Vancouver, British Columbia, Canada]*, pages 449–456, 2004. [1](#)
- [23] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. Boosting few-shot visual learning with self-supervision. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8059–8068, 2019. [5](#)
- [24] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. [3](#)
- [25] Markus Hiller, Rongkai Ma, Mehrtash Harandi, and Tom Drummond. Rethinking Generalization in Few-Shot Classification. (NeurIPS), 2022. [3](#), [5](#)
- [26] Sebastian Houben, Johannes Stallkamp, Jan Salmen, Marc Schlipsing, and Christian Igel. Detection of traffic signs in real-world images: The german traffic sign detection benchmark. In *The 2013 international joint conference on neural networks (IJCNN)*, pages 1–8. Ieee, 2013. [1](#)
- [27] Shell Xu Hu, Pablo Garcia Moreno, Yang Xiao, Xi Shen, Guillaume Obozinski, Neil D. Lawrence, and Andreas C.

- Damianou. Empirical bayes transductive meta-learning with synthetic gradients. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. 5
- [28] Shell Xu Hu, Da Li, Jan Stuhmer, Minyoung Kim, and Timothy M. Hospedales. Pushing the Limits of Simple Pipelines for Few-Shot Learning: External Data and Fine-Tuning Make a Difference. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2022-June:9058–9067, 2022. 1, 2, 3, 5, 6, 7, 8, 4
- [29] Yuqing Hu, Vincent Gripon, and Stéphane Pateux. Leveraging the feature distribution in transfer-based few-shot learning. In *International Conference on Artificial Neural Networks*, pages 487–499. Springer, 2021. 5
- [30] Kai Huang, Jie Geng, Wen Jiang, Xinyang Deng, and Zhe Xu. Pseudo-loss confidence metric for semi-supervised few-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8671–8680, 2021. 5
- [31] Jonas Jongejan, Henry Rowley, Takashi Kawashima, Jongmin Kim, and Nick Fox-Gieg. The quick, draw!-ai experiment. *Mount View, CA, accessed Feb*, 17(2018):4, 2016. 1
- [32] Sungyeon Kim, Dongwon Kim, Minsu Cho, and Suha Kwak. Proxy anchor loss for deep metric learning. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3235–3244, 2020. 4
- [33] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 1
- [34] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015. 1
- [35] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 1
- [36] Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10657–10665, 2019. 5
- [37] Wei-Hong Li, Xialei Liu, and Hakan Bilen. Cross-domain few-shot learning with task-specific adapters. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 7151–7160. IEEE, 2022. 1, 2, 3, 5, 6
- [38] Wei-Hong Li, Xialei Liu, and Hakan Bilen. Universal representation learning from multiple domains for few-shot classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9526–9535, 2021. 3
- [39] Xinzhe Li, Qianru Sun, Yaoyao Liu, Qin Zhou, Shibao Zheng, Tat-Seng Chua, and Bernt Schiele. Learning to self-train for semi-supervised few-shot classification. *Advances in neural information processing systems*, 32, 2019. 5
- [40] Xiaoxu Li, Xiaochen Yang, Zhanyu Ma, and Jing-Hao Xue. Deep metric learning for few-shot image classification: A review of recent developments. *Pattern Recognition*, page 109381, 2023. 5
- [41] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & Shifting Your Features: A New Baseline for Efficient Model Tuning. (NeurIPS):1–20, 2022. 3
- [42] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 1
- [43] Lu Liu, William L. Hamilton, Guodong Long, Jing Jiang, and Hugo Larochelle. A universal representation transformer layer for few-shot image classification. 2021. 2, 6
- [44] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 1
- [45] Thomas Mensink, Jakob Verbeek, Florent Perronnin, and Gabriela Csurka. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE transactions on pattern analysis and machine intelligence*, 35(11):2624–2637, 2013. 1, 2
- [46] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, pages 722–729. IEEE, 2008. 1
- [47] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual reasoning with a general conditioning layer. *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pages 3942–3951, 2018. 2
- [48] Xiaorong Qin, Xinhang Song, and Shuqiang Jiang. Bi-level meta-learning for few-shot domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15900–15910, 2023. 1, 2
- [49] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Efficient parametrization of multi-domain deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8119–8127, 2018. 2
- [50] Pau Rodríguez, Issam Laradji, Alexandre Drouin, and Alexandre Lacoste. Embedding propagation: Smoother manifold for few-shot classification. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16*, pages 121–138. Springer, 2020. 5
- [51] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015. 2, 6, 1
- [52] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017. 1, 2, 3, 5, 6
- [53] Manogna Sreenivas and Soma Biswas. Similar class style augmentation for efficient cross-domain few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4589–4597, 2023. 1, 2, 3, 6, 8

- [54] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI conference on artificial intelligence*, 2016. 3
- [55] Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, and Hugo Larochelle. Meta-dataset: A dataset of datasets for learning to learn from few examples. 2020. 1, 2, 5
- [56] Eleni Triantafillou, Hugo Larochelle, Richard S. Zemel, and Vincent Dumoulin. Learning a universal template for few-shot dataset generalization. 139:10424–10433, 2021. 2, 5, 6
- [57] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, number Nips, pages 3637–3645, 2016. 5, 1
- [58] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 1
- [59] Chengming Xu, Siqian Yang, Yabiao Wang, Zhanxiong Wang, Yanwei Fu, and Xiangyang Xue. Exploring Efficient Few-shot Adaptation for Vision Transformers. (62076067), 2023. 1, 2, 3, 5, 6, 7, 8
- [60] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? *Advances in Neural Information Processing Systems*, 4(January):3320–3328, 2014. 1, 5
- [61] Chi Zhang, Yujun Cai, Guosheng Lin, and Chunhua Shen. Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12203–12213, 2020. 5
- [62] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *International Conference on Learning Representations (ICLR)*, 2022. 3, 5, 1
- [63] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*, 2021. 6
- [64] Yixiong Zou, Shanghang Zhang, Jianpeng Yu, Yonghong Tian, and José MF Moura. Revisiting mid-level patterns for cross-domain few-shot recognition. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 741–749, 2021. 1