Adapting to Domain Shift by Meta-Distillation from Mixture-of-Experts

Anonymous Author(s)
Affiliation
Address
email

Abstract

In this paper, we tackle the problem of domain shift. Most existing methods perform training on multiple source domains using a single model, and the same trained model is used on all unseen target domains. Such solutions are sub-optimal as each target domain exhibits its own speciality, which is not adapted. Furthermore, expecting the single-model training to learn extensive knowledge from the multiple source domains is counterintuitive. The model is more biased to learning only domain-invariant features and may result in negative knowledge transfer. In this work, we propose a novel framework for unsupervised test-time adaptation, which is formulated as a knowledge distillation process to address domain shift. Specifically, we incorporate with Mixture-of-Experts (MoE) as teachers, where each expert is separately trained on different source domains to maximize their speciality. Given a test-time target domain, a small set of unlabeled data is sampled to query the knowledge from MoE. As the source domains are correlated to the target domains, a transformer-based aggregator then combines the domain knowledge by examining the interconnection among them. The output is treated as a supervision signal to adapt a student prediction network toward the target domain. We further employ meta-learning to enforce the aggregator to distill positive knowledge and the student network to achieve fast adaptation. Extensive experiments demonstrate that the proposed method outperforms the state-of-the-art and validates the effectiveness of each proposed component.

1 Introduction

The emergence of deep models has achieved superior performance [26, 34]. Such unprecedented success is built on the strong assumption that the training and testing data are highly correlated (i.e., they are both sampled from the same data distribution). However, the assumption normally does not hold in real-world settings as the training data is infeasible to cover all the ever-changing deployment environments [33]. Reducing such distribution correlation is known as distribution shift, which significantly hampers the performance of deep models. Although human is more robust against the distribution shift, artificial learning-based systems suffer more from performance degradation.

One line of research aims to mitigate the distribution shift by exploiting some unlabeled data from a target domain is known as unsupervised domain adaptation (UDA) [20, 41, 22]. The unlabeled data is an estimation of the target distribution [72]. Therefore, UDA normally adapts to the target domain by transferring the source knowledge via a common feature space with less effect from domain discrepancy [66, 40]. However, UDA is less applicable for real-world scenarios as repetitive large-scale training is required for every target domain. In addition, collecting the data samples from a target domain in advance might be unavailable as the target could be unknown during training. Domain generalization (DG) [44, 24, 6] is an alternative line of research but more challenging as it

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

assumes the prior knowledge of the target domains is unknown. DG methods leverage multiple source domains for training and directly use the trained model on all unseen domains. As the domain-specific information for the target domains is not adapted, a generic model is sub-optimal [57, 13].

Test-time adaptation with DG [72] allows the model to exploit the unlabeled data during testing to overcome the limitation of using a flawed generic model for all unseen target domains. In ARM, meta-learning [21] is utilized for training the model as an initialization such that it can be adapted using the unlabeled data from unseen target domain before making the final inference. However, we observed that ARM only trains a single model, which is counterintuitive for the multi-source domain setting. There is a certain amount of correlation among the source domains while each of them also exhibits its own specific knowledge. When the number of source domains rises, data complexity dramatically increases, which impedes the exploration of the dataset thoroughly. Furthermore, real-world domains are not always balanced in data scales [33]. Therefore, the single-model training is more biased toward the domain-invariant features and dominant domains instead of the domain-specific features [11].

In this work, we propose to formulate the test-time adaptation as the process of knowledge distillation [28] from multiple source domains. Concretely, we propose to incorporate the concept of Mixture-of-Experts (MoE), which is a natural fit for the multi-source domain settings. The MoE models are treated as a teacher and separately trained on the corresponding domain to maximize their domain speciality. Given a new target domain, a few unlabeled data are collected to query the features from expert models. A transformer-based knowledge aggregator is proposed to examine the interconnection among queried knowledge and aggregate the correlated information toward the target domain. The output is then treated as a supervision signal to update a student prediction network to adapt to the target domain. The adapted student is then used for subsequent inference. We employ bi-level optimization as meta-learning to train the aggregator at the meta-level to improve generalization. The student network is also meta-trained to achieve fast adaptation via a few samples. Furthermore, we simulate the test-time out-of-distribution scenarios during training to align the training objective with the evaluation protocol.

The proposed method also provides additional advantages over ARM: 1) Our method provides a larger model capability to improve the generalization power; 2) Despite the higher computational cost, only the adapted student network is kept for inference, while the MoE models are discarded after adaptation. Therefore, our method is more flexible in designing the architectures for the teacher or student models. (e.g., designing compact models for the power-constrained environment); 3) Our method does not need to access the raw data of source domains but only needs their trained models. So, we can take advantage of private domains in a real-world setting where their data is inaccessible.

We name our method as **Meta-D**istillation of **MoE** (Meta-DMoE). Our contributions are as follows:

- We propose a novel unsupervised test-time adaptation framework that is tailored for multiple sources domain settings. Our framework employs the concept of MoE to allow each expert model to explore each source domain thoroughly. We formulate the adaptation process as knowledge distillation via aggregating the positive knowledge retrieved from MoE.
- The alignment between training and evaluation objectives via meta-learning improves the adaptation, hence the test-time generalization.
- We conduct extensive experiments to show the superiority of the proposed method among the state-of-the-arts and validate the effectiveness of each component of Meta-DMoE.
- We validate that our method is more flexible in real-world settings where computational power and data privacy are the concerns.

2 Related work

Domain shift Unsupervised Domain Adaptation (UDA) has been popular to address domain shift by transferring the knowledge from the labeled source domain to the unlabeled target domain [39, 35, 68]. It is achieved by developing domain-invariant via minimizing statistical discrepancy across domains [5, 47, 59]. Adversarial learning is also applied to develop indistinguishable feature space [22, 41, 46]. The first limitation of UDA is the assumption of the co-existence of source and target data, which is inapplicable when the target domain is unknown in advance. Furthermore, most of the algorithms focus on unrealistic single-source-single-target adaptation as source data normally come from multiple domains. Splitting the source data into various distinct domains and

exploring the unique characteristics of each domain and the dependencies among them strengthen the robustness [74, 64, 65]. Domain generalization (DG) is another line of research to alleviate the domain shift. DG aims to train a model on multiple source domains without accessing any prior information of the target domain, and expect it to perform well on unseen target domains. [24, 37, 43] learns the domain-invariant feature representation. [52, 63] exploits data augmentation strategies in data or feature space. However, deploying the same generic model to all unseen target domains fails to explore domain specialty and yields sub-optimal solutions. In contrast, our method further exploits the unlabeled target data and updates the trained model to each specific unseen target domain at test time.

Test-time adaptation (TTA) TTA constructs supervision signals from unlabeled data to update the generic model before inference. Sun *et al.* [57] use rotation prediction to update the model during inference. Chi *et al.* [13] and Li *et al.* [38] reconstruct the input images to achieve internal-learning to better restore the blurry images and estimate the human pose. ARM [72] incorporates test-time adaptation with DG which meta-learns an model that is capable of adapting to unseen target domains before making inference. Instead of adapting to every data sample, our method only updates once for each target domain using a fixed number of examples.

Meta-learning Meta-learning aims to learn the learning process by episodic learning. The existing meta-learning methods can be categorised as model-based [51, 48, 8], metric-based [54], and optimization-based [21]. In our method, we incorporate meta-learning with knowledge distillation by jointly learning a student model initialization and a knowledge aggregator for fast adaptation.

Mixture-of-Experts The goal of MoE [31] is to decompose the whole training set into many subsets, which are independently learned by different models. It has been successfully applied in image recognition models to improve the accuracy [1]. MoE is also popular in scaling up the architectures. As each expert is independently trained, sparse selection methods are developed to select a subset of the MoE during inference to increase the network capacity [36, 19, 25]. In contrast, our method utilizes all the experts to extract and combine the knowledge for positive knowledge transfer.

3 Preliminaries

In this section, we describe the problem setting and discuss the adaptive model. We mainly follow the test-time unsupervised adaptation as in [72]. Specifically, we define a set of N source domains $\mathcal{D}_\mathcal{S} = \{{\mathcal{D}_\mathcal{S}}^i\}_{i=1}^N$ and M target domains $\mathcal{D}_\mathcal{T} = \{{\mathcal{D}_\mathcal{T}}^j\}_{j=1}^M$. The physical definition of a domain varies and depends on the applications or data collection methods. It could be a specific dataset, user or location. Let $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ denote the input and corresponding label, respectively. Each of the source domains contains the data in the form of input-output pairs: ${\mathcal{D}_\mathcal{S}}^i = \{(x_\mathcal{S}^z, y_\mathcal{S}^z)\}_{z=1}^{Z_i}$. In contrast, each of the target domains contains only the unlabeled data: ${\mathcal{D}_\mathcal{T}}^j = \{(x_\mathcal{T}^k)\}_{k=1}^{K_j}$. For well-designed datasets, (e.g. [27, 16]), all the source or target domains have the same number of data samples. Such condition is not ubiquitous for real-world scenarios (i.e. $Z_{i_1} \neq Z_{i_2}$ if $i_1 \neq i_2$ and $K_{j_1} \neq K_{j_2}$ if $j_1 \neq j_2$) where data imbalance always exists [33]. It further challenges the generalization with a broader range of real-world distribution shifts instead of finite synthetic ones. Generic domain shift tasks focus on the out-of-distribution setting where the source and target domains are non-overlapping (i.e. $\mathcal{D}_\mathcal{S} \cap \mathcal{D}_\mathcal{T} = \varnothing$), but the label spaces of both domains are the same (i.e. $\mathcal{Y}_\mathcal{S} = \mathcal{Y}_\mathcal{T}$).

Conventional DG methods perform training on $\mathcal{D}_\mathcal{S}$ and make minimal assumption on the testing scenarios [56, 3, 29]. Therefore, the same generic model is directly applied to all target domains $\mathcal{D}_\mathcal{T}$, which leads to non-optimal solutions [57]. In fact, for each ${\mathcal{D}_\mathcal{T}}^j$, some unlabeled data are readily available which provide certain prior knowledge for that target distribution. Adaptive Risk Minimization (ARM) [72] assumes that a batch of unlabeled input data $\mathbf{x}$ approximates the input distribution p_x which provides useful information about $p_{y|x}$. Based on the assumption, an unsupervised test-time adaptation [48, 23] is proposed. The fundamental concept is to adapt the model to the specific domain using $\mathbf{x}$. Overall, ARM aims to minimize the following objective $\mathcal{L}(\cdot,\cdot)$ over all training domains:

$$\sum_{{\mathcal{D}_\mathcal{S}}^j \in \mathcal{D}_\mathcal{S}} \sum_{(\mathbf{x},\mathbf{y}) \in {\mathcal{D}_\mathcal{S}}^j} \mathcal{L}(\mathbf{y}, f(\mathbf{x};\theta')), \text{where } \theta' = h(\mathbf{x}, \theta; \phi). \tag{1}$$

$\mathbf{y}$ is the labels for $\mathbf{x}$. $f(\mathbf{x};\theta)$ denotes the prediction model parameterized by θ. $h(\cdot;\phi)$ is an adaptation function parameterized by ϕ. It receives the original θ of f and the unlabeled data $\mathbf{x}$ to adapt θ to θ'.

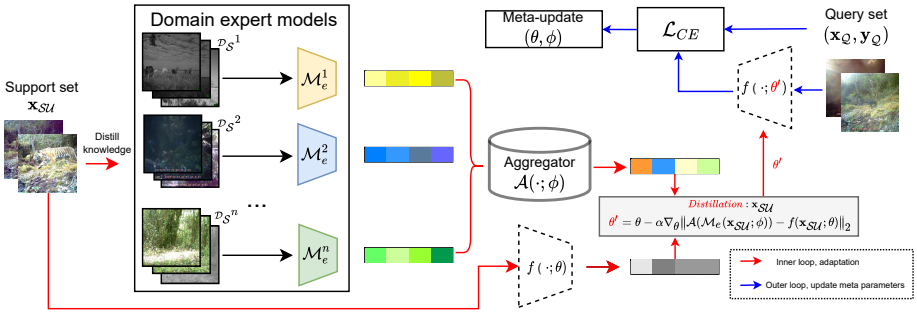

Figure 1: Overview of the training of Meta-DMoE. We first sample disjoint support set $\mathbf{x}_{\mathcal{SU}}$ and query set $(\mathbf{x}_{\mathcal{Q}}, \mathbf{y}_{\mathcal{Q}})$ from a training domain. $\mathbf{x}_{\mathcal{SU}}$ is sent to the expert models $\mathcal{M}$ to query their domain-specific knowledge. An aggregator $\mathcal{A}(\cdot;\phi)$ then combines the information and generates a supervision signal to update the $f(\cdot;\theta)$ via knowledge distillation. The updated $f(\cdot;\theta')$ is evaluated using the labeled query set to update the meta-parameters.

The goal of ARM is to learn both (θ, ϕ). To mimic the test-time adaptation (i.e., adapt before prediction), it follows the episodic learning as in meta-learning [21]. Specifically, each episode processes a domain by performing unsupervised adaptation using $\mathbf{x}$ and $h(\cdot;\phi)$ in the inner loop to obtain $f(\cdot;\theta')$. The outer loop evaluates the adapted $f(\cdot;\theta')$ using the true label to perform meta-update. ARM is a general framework that can be incorporated with existing meta-learning approaches with different forms of adaptation module $h(\cdot;\cdot)$ [21, 23].

However, several shortcomings are observed with respect to the generalization. The episodic learning processes one domain at a time, which has clear boundaries among the domains. The overall setting is equivalent to the multi-source domain setting, which is proven to be more effective than learning from a single domain [43, 73] as most of the domains are correlated to each other [2]. However, it is counterintuitive to learn all the domain knowledge in one single model as each domain has specialized semantics or low-level features [53]. Therefore, the single-model method in ARM is sub-optimal due to: 1) some domains may contain competitive information, which leads to negative knowledge transfer [55]. It may tend to learn the ambiguous feature representations instead of capturing all the domain-specific information [67]; 2) not all the domains are equally important [64], and the learning might be biased as the data in each domain is imbalanced in real-world [33].

4 Proposed approach

In this section, we explicitly formulate the test-time adaptation as a knowledge transfer process to distill the knowledge from MoE. The proposed method is learned via meta-learning to mimic the test-time out-of-distribution scenarios and ensure positive knowledge transfer.

4.1 Meta-distillation from mixture-of-experts

Overview Fig. 1 shows the method overview. We wish to explicitly transfer valid knowledge from various domains to elevate the generalization of unseen domains. Concretely, we define MoE as $\mathcal{M} = \{\mathcal{M}^i\}_{i=1}^N$ to represent the domain-specific models. Each $\mathcal{M}^i$ is separately trained using standard supervised learning on source domain ${\mathcal{D}_{\mathcal{S}}}^i$ to learn its discriminative features. We propose the test-time adaptation as the unsupervised knowledge distillation [28] to learn the knowledge from MoE. Therefore, we treat $\mathcal{M}$ as the teacher and aim to distill its knowledge to a student prediction network $f(\cdot;\theta)$ to achieve adaptation. To do so, we sample a batch of unlabeled $\mathbf{x}$ from a target domain, and pass it to $\mathcal{M}$ to query their domain-specific knowledge $\{\mathcal{M}^i(\mathbf{x})\}_{i=1}^N$. That knowledge is then forwarded to a knowledge aggregator $\mathcal{A}(\cdot;\phi)$. The aggregator is learned to explore the interconnection among domain knowledge and yield knowledge composition towards that target domain. The output of $\mathcal{A}(\cdot;\phi)$ is treated as the supervision signal to update $f(\mathbf{x};\theta)$. Once the adapted θ' is obtained, $f(\cdot;\theta')$ is used to test the rest of the data in that domain. The overall framework follows the effective few-shot learning where $\mathbf{x}$ is treated as unlabeled support set [62, 54, 21].

Training Meta-DMoE Properly training the (θ, ϕ) is critical to improve the generalization on unseen domains. First, $\mathcal{A}(\cdot, \phi)$ performs as a mechanism that explores and mixes the input knowledge. Thus, it is not supposed to be biased to any training data. Second, the conventional distillation process requires large numbers of data samples and learning iterations [28, 2]. The repetitive large-scale training is inapplicable in real-world applications. To mitigate the aforementioned challenges, we follow the meta-learning paradigm [21]. Such bi-level optimization enforces the $\mathcal{A}(\cdot, \phi)$ to learn beyond any specific knowledge [71] and allows the student prediction network $f(\cdot; \theta)$ to achieve fast adaptation. Specifically, We first split the data samples in each source domain ${\mathcal{D}_\mathcal{S}}^i$ into disjoint support and query sets. The unlabeled support set $(\mathbf{x}_{\mathcal{SU}})$ is used to perform adaptation via knowledge distillation, while the labeled query set $(\mathbf{x}_\mathcal{Q}, \mathbf{y}_\mathcal{Q})$ is used to evaluate the adapted parameters to explicitly test the generalization on unseen data.

Algorithm 1 Training for Meta-DMoE

Require: $\{{\mathcal{D}_\mathcal{S}}^i\}_{i=1}^N$: data of source domains; α, β: learning rates; B: meta batch size

1: **// Pretrain domain-specific MoE models**
2: **for** i=1,...,N **do**
3: Train the domain-specific model $\mathcal{M}^i$ using ${\mathcal{D}_\mathcal{S}}^i$.
4: **end for**
5: **// Meta-train aggregator** $\mathcal{A}(\cdot; \phi)$ **and student model** $f(\cdot, \theta_e; \theta_c)$
6: **Initialize**: ϕ, θ_e, θ_c
7: **while** *not converged* **do**
8: Sample a batch of B source domains $\{{\mathcal{D}_\mathcal{S}}^b\}^B$, reset batch loss $\mathcal{L}_\mathcal{B} = 0$
9: **for** each ${\mathcal{D}_\mathcal{S}}^b$ **do**
10: Sample support and query set: $(\mathbf{x}_{\mathcal{SU}}), (\mathbf{x}_\mathcal{Q}, \mathbf{y}_\mathcal{Q}) \sim {\mathcal{D}_\mathcal{S}}^b$
11: $\mathcal{M}'_e(\mathbf{x}_{\mathcal{SU}}; \phi) = \{\mathcal{M}^i_e(\mathbf{x}_{\mathcal{SU}}; \phi)\}_{i=1}^N$, mask $\mathcal{M}^i_e(\mathbf{x}_{\mathcal{SU}}; \phi)$ with $\mathbf{0}$ if $b = i$
12: Perform adaptation via knowledge distillation from MoE:
13: $\theta'_e = \theta_e - \alpha \nabla_{\theta_e} \left\| \mathcal{A}(\mathcal{M}'_e(\mathbf{x}_{\mathcal{SU}}; \phi)) - f(\mathbf{x}_{\mathcal{SU}}; \theta_e) \right\|_2$
14: Evaluate the adapted θ'_e using query set and accumulate the loss:
15: $\mathcal{L}_\mathcal{B} = \mathcal{L}_\mathcal{B} + \mathcal{L}_{CE}(\mathbf{y}_\mathcal{Q}, f(\mathbf{x}_\mathcal{Q}; \theta'_e, \theta_c))$
16: **end for**
17: Update ϕ, θ_e, θ_c for the current meta batch:
18: $(\phi, \theta_e, \theta_c) \leftarrow (\phi, \theta_e, \theta_c) - \beta \nabla_{(\phi, \theta_e, \theta_c)} \mathcal{L}_\mathcal{B}$
19: **end while**

The student prediction network $f(\cdot; \theta)$ can be decoupled as a feature extractor θ_e and classifier θ_c. Unsupervised knowledge distillation can be achieved via the softened output [28] or intermediate features [70] from $\mathcal{M}$. The former one allows the whole student network $\theta = (\theta_e, \theta_c)$ to be adaptive, while the latter one allows partial or complete θ_e to adapt to $\mathbf{x}$, depending on the features utilized. We follow [45] to only adapt θ_e in the inner loop while keeping the θ_c fixed. Thus, the adaptation process is achieved by distilling the knowledge via the aggregated features:

$$DIST(\mathbf{x}_{\mathcal{SU}}, \mathcal{M}_e, \phi, \theta_e) = \theta'_e = \theta_e - \alpha \nabla_{\theta_e} \left\| \mathcal{A}(\mathcal{M}_e(\mathbf{x}_{\mathcal{SU}}); \phi) - f(\mathbf{x}_{\mathcal{SU}}; \theta_e) \right\|_2, \tag{2}$$

where α denotes the adaptation learning rate, $\mathcal{M}_e$ is the feature extractor of MoE models which extracts the features before the classifier, and $\|\cdot\|_2$ measures the L_2 distance. The goal is to obtain an updated θ'_e such that the extracted features of $f(\mathbf{x}_{\mathcal{SU}}; \theta'_e)$ is closer to the aggregated features. The overall learning objective of Meta-DMoE is to minimize the following expected loss:

$$\arg\min_{\theta_e, \theta_c, \phi} \sum_{{\mathcal{D}_\mathcal{S}}^j \in \mathcal{D}_\mathcal{S}} \sum_{\substack{(\mathbf{x}_{\mathcal{SU}}) \in {\mathcal{D}_\mathcal{S}}^j \\ (\mathbf{x}_\mathcal{Q}, \mathbf{y}_\mathcal{Q}) \in {\mathcal{D}_\mathcal{S}}^j}} \mathcal{L}_{CE}(\mathbf{y}_\mathcal{Q}, f(\mathbf{x}_\mathcal{Q}; \theta'_e, \theta_c)), \text{where } \theta'_e = DIST(\mathbf{x}_{\mathcal{SU}}, \mathcal{M}_e, \phi, \theta_e), \tag{3}$$

where $\mathcal{L}_{CE}$ is the cross-entropy loss. Alg. 1 demonstrates our full training procedure. To smooth the meta gradient and stabilize the training, we process a batch of episodes before each meta-update.

Since the training domains overlap for the MoE and meta-training, we simulate the test-time out-of distribution by excluding the corresponding expert model in each episode. To do so, we multiple the features by $\mathbf{0}$ to mask them out. $\mathcal{M}'_e$ in L11 of Alg. 1 denotes such operation. Therefore, the adaptation is enforced to use the knowledge that is aggregated from other domains.

4.2 Fully learned explicit knowledge aggregator

Aggregating the knowledge from distinct domains requires exploring the relation among them to ensure the relevant knowledge transfer. Prior works design more hand-engineered solutions to combine the knowledge or choose data samples that are closer to the target domain for knowledge transfer [2, 74]. A superior alternative is to replace the hand-designed pipelines with the fully learned solutions [15, 9]. Thus we follow the same trend and allow the aggregator $\mathcal{A}(\cdot; \phi)$ to be fully meta-learned without many manual designs except defining its architecture.

We observe that the self-attention mechanism is quite suitable where interaction among domain knowledge can be computed. Therefore, we use a transformer encoder as the aggregator [18, 61]. The

encoder consists of multi-head self-attention and multi-layer perceptron blocks with layernorm [4] applied before each block, and residual connection applied after each block. We refer the readers to the supplementary material for the detailed architecture and computation. We concatenate the output features from the MoE models in the domain dimension as $Concat[\mathcal{M}_e^1(\mathbf{x}), \mathcal{M}_e^2(\mathbf{x}), ..., \mathcal{M}_e^N(\mathbf{x})] \in \mathbb{R}^{N \times d}$, where d is the feature dimension. The aggregator $\mathcal{A}(\cdot; \phi)$ processes the input tensor to obtain the aggregated feature $\mathbf{F} \in \mathbb{R}^d$, which is used as a supervision signal for test-time adaptation.

4.3 More constrained real-world settings

In this section, we investigate two critical settings for real-world applications but have drawn less attention from the prior works: limitation on computational resources and data privacy.

Constraint on computational cost In real-world deployment environments, the computational power might be highly constrained (e.g., smartphones). It requires fast inference and compact models. However, the reduction in learning capabilities greatly hinders the generalization as some methods utilize only a single model regardless of the data complexity. On the other hand, when the number of domain data scales up, methods relying on adaptation on every data sample [72] will experience inefficiency. In contrast, our method only needs to perform adaptation once for every unseen domain. Only the final $f(\cdot; \theta')$ is used for inference. To investigate the impact on generalization caused by reducing the model size, we replace the models (only $f(\cdot; \theta)$ for us) with MobileNet V2 [50].

Data privacy regulation Large-scale training data is normally collected from various venues. However, some venues may have privacy regulations enforced. Their data might not be accessible but the models that are trained using the private data are available. To simulate such an environment, we split the training source domains into two splits: private domains (${\mathcal{D}_\mathcal{S}}^{pri}$) and public domains (${\mathcal{D}_\mathcal{S}}^{pub}$). We use ${\mathcal{D}_\mathcal{S}}^{pri}$ to train MoE models and ${\mathcal{D}_\mathcal{S}}^{pub}$ for the subsequent meta-training. Since ARM and other methods only utilize the data as input, we train them on ${\mathcal{D}_\mathcal{S}}^{pub}$. We conduct the experiments with data privacy on top of the power-constrained scenario and show the superiority of the proposed method in the experiment section. For details of the settings, please refer the supplementary material.

5 Experiements

5.1 Datasets and implementation details

Datasets and evaluation metrics. In this work, we mainly evaluate our method on the real-world domain shift scenarios. Drastic variation in deployment conditions normally exists in nature, such as a change in illumination, background, and time. It shows a huge domain gap between deployment environments and imposes challenges to the algorithm's robustness. Thus, we test our methods on the large-scale distribution shift benchmark WILDS [33], which reflects a diverse range of real-world distribution shifts. Following [72], we mainly perform experiments on five image testbeds, iWildCam [10], Camelyon17 [7],RxRx1 [58] and FMoW [14] and PovertyMap [69]. In each benchmark dataset, a domain represents a distribution over data that is similar in some way, such as images collected from the same camera trap or satellite images taken in the same locations. We follow the same evaluation metrics as in [33] to compute: accuracy, Macro F1, worst-case (WC) accuracy, Pearson correlation (r), and its worst-case counterpart.

Network architecture. We follow WILDS [33] to use ResNet18 & 50 [26] or DenseNet101 [30] for the expert models $\{\mathcal{M}^i\}_{i=1}^N$ and student network $f(\cdot, ; \theta)$. Also, we use a single-layer transformer encoder block[61] as the knowledge aggregator $\mathcal{A}(\cdot; \phi)$. To investigate the resource-constrained and privacy-sensitive scenarios, we utilize MobileNet V2 [50] with a width multiplier of 0.25.

Pre-training domain-specific models. WILDS benchmark is highly imbalanced in data size, and some classes have empty input data set; we found that using every single domain to train an expert is unstable, and sometimes it cannot converge. Inspired by [42], we propose to cluster the training domains into N super domains and use each super-domain to train the expert models. Specifically, we set $N = \{10, 5, 3, 4, 3\}$ for iWildCam, Camelyon17, RxRx1, FMoW and Poverty Map, respectively. We use ImageNet [17] pre-trained model as the initialization and separately train the models using Adam optimizer [32] with a learning rate of $1e^{-4}$ and a decay of 0.96 per epoch.

Meta-training and testing. We first pre-train the aggregator and student network [12]. After that, the model is further trained using Alg. 1 for 15 epochs with a fixed learning rate of $3e^{-4}$ for α and

Table 1: Comparison with the state-of-the-arts on the WILDS image testbeds and out-of-distribution setting. Metric means and standard deviations are reported across replicates. Our proposed method performs well across all problems and achieves the best results on 4 out of 5 datasets.

	iWildCam		Camelyon17	RxRx1	FMoW		PovertyMap	
Method	Acc	Macro F1	Acc	Acc	WC Acc	Avg Acc	WC Pearson r	Pearson r
ERM	71.6 (2.5)	31.0 (1.3)	70.3 (6.4)	29.9 (0.4)	32.3 (1.25)	**53.0 (0.55)**	0.45 (0.06)	0.78 (0.04)
CORAL	73.3 (4.3)	32.8 (0.1)	59.5 (7.7)	28.4 (0.3)	31.7 (1.24)	50.5 (0.36)	0.44 (0.06)	0.78 (0.05)
Group DRO	72.7 (2.1)	23.9 (2.0)	68.4 (7.3)	23.0 (0.3)	30.8 (0.81)	52.1 (0.5)	0.39 (0.06)	0.75 (0.07)
IRM	59.8 (3.7)	15.1 (4.9)	64.2 (8.1)	8.2 (1.1)	30.0 (1.37)	50.8 (0.13)	0.43 (0.07)	0.77 (0.05)
ARM-CML	70.5 (0.6)	28.6 (0.1)	84.2 (1.4)	17.3 (1.8)	27.2 (0.38)	45.7 (0.28)	0.37 (0.08)	0.75 (0.04)
ARM-BN	70.3 (2.4)	23.7 (2.7)	87.2 (0.9)	**31.2 (0.1)**	24.6 (0.04)	42.0 (0.21)	0.49 (0.21)	**0.84 (0.05)**
ARM-LL	71.4 (0.6)	27.4 (0.8)	84.2 (2.6)	24.3 (0.3)	22.1 (0.46)	42.7 (0.71)	0.41 (0.04)	0.76 (0.04)
Ours (w/o mask)	**74.1 (0.4)**	**35.1 (0.9)**	**90.8 (1.3)**	29.6 (0.5)	**36.8 (1.01)**	50.6 (0.20)	**0.52 (0.04)**	0.80 (0.03)
Ours	**77.2 (0.3)**	**34.0 (0.6)**	**91.4 (1.5)**	29.8 (0.4)	**35.4 (0.58)**	52.5 (0.18)	**0.51 (0.04)**	0.80 (0.03)

$3e^{-5}$ for β. During meta-testing, we use Line 13 of Alg. 1 to adapt before making a prediction for every testing domain. Specifically, we set the number of examples for adaptation at test time = $\{24, 64, 75, 64, 64\}$ for iWildCam, Camelyon17, RxRx1, FMoW, and Poverty Map, respectively. For both meta-training and testing, we perform one gradient update for adaptation on the unseen target domain. We refer the readers to the supplementary materials for more detailed information.

5.2 Main results

In this section, we compare the proposed method with prior approaches showing on WILDS leader-board [33], including non-adaptive methods: CORAL [56], ERM [33], IRM [3], Group DRO [49] and adaptive methods used in ARM [72] (CML, BN and LL). We directly copy the available results from the WILDS leader-board or their corresponding paper. As for the missing ones, we conduct experiments using their provided source code with default hyperparameters. Table 1 reports the comparison with the state-of-the-art. Our proposed method performs well across all datasets and increases both worst-case and average accuracy compared to other methods. Our proposed method achieves the best performance on 4 out of 5 benchmark datasets. ARM [72] applies meta-learning approach to learn how to adapt to unseen domains with unlabeled data. However, their method is greatly bounded by using a single model to exploit knowledge from multiple source domains. Instead, our proposed method is more fitted to multi-source domain settings and meta-trains an aggregator that properly mixtures the knowledge from multiple domain-specific experts. As a result, our method outperforms ARM-CML, BN and LL by 9.5%, 9.8%, 8.1% for iWildCam, 8.5%, 4.8%, 8.5% for Camelyon17 and 14.8%, 25.0%, 22.9% for FMoW in terms of average accuracy. Furthermore, we also evaluate our method without masking the in-distribution domain in MoE models (Ours w/o mask) during meta training (Line 10-11 of Alg. 1), where the sampled domain is overlapped with MoE. It violates the generalization to unseen target domains during testing. As most of the performance dropped, it reflects the importance of aligning the training and evaluation objectives.

Visualization of adapted feature To evaluate the capability of adaptation via learning discriminative representations on unseen target domains, we compare the t-SNE [60] feature visualization using the same test domain sampled from iWildCam and Camelyon17 datasets. ERM utilizes single model and standard supervised training without adaptation, therefore, we set it as the baseline. Figure 2 shows the comparison, where each color denotes a class and each point represents a data sample. It is clear that our method obtains better clustered and more discriminative features.

5.3 Results under constrained real-world settings

Constraint on computational cost Computational power is always limited in real-world deployment scenarios, such as edge devices. Efficiency and adaptation ability should be both considered. Thus, we replace our student model and the models in other methods with MobileNet V2. As reported in Table 2, our proposed method still outperforms prior methods. Since the MoE model is only used for knowledge transfer, our method is more flexible in designing the student architecture for different scenarios. We also report multiply-Accumulate operations (MACS) for inference and time complexity on adaptation. As ARM needs to make adaptation before inference on every example, its adaptation cost scales linearly with the number of examples. Our proposed method performs better in accuracy and requires much less computational cost for adaptation, as reported in Table 3.

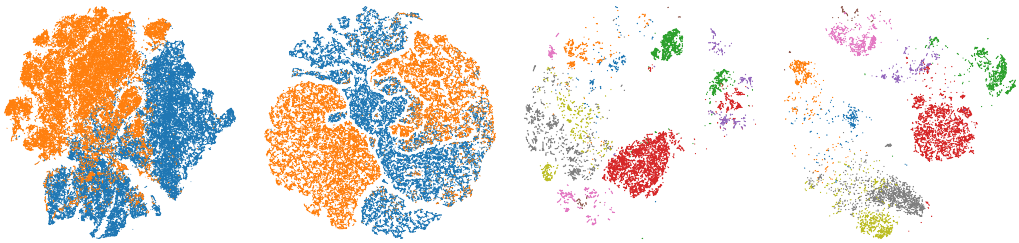

(a) Camelyon17-ERM (b) Camelyon17-Ours (c) iWildCam-ERM (d) iWildCam-Ours

Figure 2: t-SNE visualization of adapted features at test-time. We directly utilize features adapted to the same unseen target domains from ERM and our proposed method in Camelyon17 and WildCam datasets, respectively. Our resulting features show more discriminative decision boundaries.

Table 2: Comparison of WILDS testbeds using MobileNet V2. Reducing the model size hampers the learning capability. Our method shows a better trade-off as the knowledge is distilled from MoE.

	iWildCam		Camelyon17	RxRx1	FMoW		PovertyMap	
Method	Acc	Macro F1	Acc	Acc	WC Acc	Avg Acc	WC Pearson r	Pearson r
ERM	56.7 (0.7)	17.5 (1.2)	69.0 (8.8)	14.3 (0.2)	15.7 (0.68)	**40.0 (0.11)**	0.39 (0.05)	0.77 (0.04)
CORAL	**61.5 (1.7)**	17.6 (0.1)	75.9 (6.9)	12.6 (0.1)	22.7 (0.76)	31.0 (0.32)	**0.44 (0.06)**	**0.79 (0.04)**
ARM-CML	58.2 (0.8)	15.8 (0.6)	74.9 (4.6)	14.0 (1.4)	21.1 (0.33)	30.0 (0.13)	0.41 (0.05)	0.76 (0.03)
ARM-BN	54.8 (0.6)	13.8 (0.2)	85.6 (1.6)	14.9 (0.1)	17.9 (1.82)	29.0 (0.69)	0.42 (0.05)	0.76 (0.03)
ARM-LL	57.5 (0.5)	12.6 (0.8)	84.8 (1.7)	15.0 (0.2)	17.1 (0.22)	30.3 (0.54)	0.39 (0.07)	0.76 (0.02)
Ours	59.5 (0.7)	**19.7 (0.5)**	**87.1 (2.3)**	**15.1 (0.4)**	**26.9 (0.67)**	37.9 (0.31)	**0.44 (0.04)**	0.77 (0.03)

Constraint on privacy regulation On top of computational limitations, privacy-regulated scenarios are common in the real world. It introduces new challenges as the raw data is inaccessible. Our method does not need to access the raw data but the trained models, which greatly mitigates such regulation. Thus, as shown in Table 4, our method does not suffer from much performance degradation compared to other methods that require access to the private raw data.

Table 3: Adaptation efficiency evaluated on iWildCam using MobileNet V2. Our method not only outperforms prior methods but also keeps constant time complexity in test-time adaptation.

Method	Acc / Macro-F1	MACS	Complexity
ERM	56.7 / 17.5	7.18×10^7	N/A
ARM-CML	58.2 / 15.8	7.73×10^7	O(n)
ARM-LL	57.5 / 12.6	7.18×10^7	O(n)
Ours	59.5 / 19.7	7.18×10^7	O(1)

Table 4: Results on privacy-related regulation setting evaluated on iWildCam and FMoW using MobileNet V2. Without privacy considered in the design, prior methods can only exploit public data and thus achieve far worse performance.

	iWildCam		FMoW	
Method	Acc	Macro-F1	WC Acc	Acc
ERM	51.2	11.2	22.5	**35.4**
CORAL	50.2	11.1	18.1	25.4
ARM-CML	42.7	7.5	16.8	24.1
ARM-BN	46.9	8.7	14.2	22.2
ARM-LL	46.8	9.3	13.7	22.6
Ours	**54.7**	**14.2**	**24.4**	33.8

5.4 Ablation studies

In this section, we conduct ablation studies on iWildCam to analyze various components of the proposed method. We also seek to answer the two key questions: 1) Does the number of experts affect the capability of capturing knowledge from multi-source domains? 2) Is meta-learning superior than standard supervised learning under the knowledge distillation framework?

Number of domain-specific experts

We investigate the impact of exploiting multiple experts to store domain-specific knowledge separately. Specifically, we keep the total number of data for experts pretraining fixed and report the results using various number of expert models. The experiments in Table 5 validate the benefits of using more domain-specific experts.

Training scheme To verify the effectiveness of meta-learning, we investigate three training schemes: random initialization, pre-train, and meta-train. To pre-train the aggregator, we add a classifier layer to its aggregated output and follow the standard supervised training scheme. For fair comparisons, we use the same testing scheme, including the number of updates and images for adaptation. Table 6

reports the results of different training scheme combinations. We observe that the randomly initialized student model struggles to learn with only a few-shot data. And the pre-trained aggregator brings weaker adaptation guidance to the student network as the aggregator is not learned to distill. In contrast, our bi-level optimization-based training scheme enforces the aggregator to choose more correlated knowledge from multiple experts to improve the adaptation of the student model. Therefore, the meta-learned aggregator is more optimal (row 1 vs. row 2). Furthermore, our meta-distillation training process simulates the adaptation in testing scenarios, which aligns with the training objective and evaluation protocol. Hence, for both meta-trained aggregator and student models, it gains additional improvement (row 3 vs. row 4).

Aggregator and distillation methods Table 7 reports the affects of various aggregators including two hand designed operators: Max and Average pooling, and two MLP based methods: Weighted sum (MLP-WS) and Projector (MLP-P) (details are provided in the supplement). We found that the fully learned transformer-based aggregator is crucial for mixing domain-specific features. Another important design choice in our proposed framework in the form of knowledge: distilling the teacher model's logits, intermediate features, or both. We show evaluation results of those three forms of knowledge in Table 8

Table 5: Results on the number of domain-specific experts. More experts increase the learning capacity to better explore each source domain, thus, improving generalization.

# of experts	**2**	**5**	**7**	**10**
Accuracy	70.4	74.1	76.4	77.2
Macro-F1	30.6	32.3	33.7	34.0

Table 6: Evaluation of training schemes. Using both meta-learned aggregator and student model improves generalization as they are learned towards test-time adaptation.

Train Scheme		Metrics	
Aggregator	Student	Acc	Macro-F1
Pretrain	Random	6.2	0.1
Meta	Random	32.7	0.5
Pretrain	Meta	74.8	32.9
Meta	Meta	77.2	34.0

Table 7: Comparison between different aggregator methods. The transformer explores interconnection, which gives the best result.

	Max	Ave.	MLP-WS	MLP-P	Trans.(ours)
Acc.	69.2	69.7	70.7	73.7	77.2
M-F1	29.2	25.0	32.8	32.7	34.0

Table 8: Comparison between different distillation methods. Distilling only the feature extractor yields the best generalization.

	Logits	Logits + Feat.	Feat. (Ours)
Accuracy	72.1	73.1	77.2
Marco-F1	26.4	26.9	34.0

6 Discussion

We present Meta-DMoE, a framework for adaptation towards domain shift using unlabeled examples at test time. We formulate the adaptation as a knowledge distillation process and devise a meta-learning algorithm to guide the student network to fast adapt to unseen target domains via transferring the aggregated knowledge from multiple sources domain-specific models. We demonstrate that Meta-DMoE is the state-of-the-art on four benchmarks. And it is competitive under two constrained real-world settings with a limited computational budget and domain data privacy regulation.

Limitations. As discussed in Section 5.4, Meta-DMoE can improve the capacity to capture complex knowledge from multi-source domains by increasing the number of experts. However, to compute the aggregated knowledge from domain-specific experts, every expert model needs to have one feed-forward pass. As a result, computational cost of adaptation scales linearly with the number of experts. Furthermore, to add or remove any domain-specific expert, both the aggregator and the student network require to be retrained. Thus, enabling a sparse-gated Meta-DMoE to encourage efficiency and scalability could be a valuable future direction, where a gating module determines a sparse combination of domain-specific experts to be used for each target domain.

Social impact. Tackling domain shift problems can have positive social impacts as it helps to elevate the model accuracy in real-world scenarios, e.g. applications in healthcare, and self-driving cars. In healthcare, domain shift occurs when a trained model is applied to patients in different hospitals. In this case, model performance might dramatically decrease, which leads to severe consequences. Tackling domain shifts helps to ensure that models can work well on new data, which can ultimately lead to better patient care. We believe our work is a small step toward the goal of adapting to domain shift.

References

[1] Karim Ahmed, Mohammad Haris Baig, and Lorenzo Torresani. Network of experts for large-scale image categorization. In *European Conference on Computer Vision*, 2016.

[2] Sk Miraj Ahmed, Dripta S Raychaudhuri, Sujoy Paul, Samet Oymak, and Amit K Roy-Chowdhury. Unsupervised multi-source domain adaptation without access to source data. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[3] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.

[4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

[5] Mahsa Baktashmotlagh, Mehrtash Harandi, and Mathieu Salzmann. Distribution-matching embedding for visual domain adaptation. *Journal of Machine Learning Research*, 17:Article–number, 2016.

[6] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. In *Advances in Neural Information Processing Systems*, 2018.

[7] Peter Bandi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermsen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging*, 2018.

[8] Peyman Bateni, Raghav Goyal, Vaden Masrani, Frank Wood, and Leonid Sigal. Improved few-shot visual classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.

[9] Shawn Beaulieu, Lapo Frati, Thomas Miconi, Joel Lehman, Kenneth O Stanley, Jeff Clune, and Nick Cheney. Learning to continually learn. In *European Conference on Artificial Intelligence*, 2020.

[10] Sara Beery, Elijah Cole, and Arvi Gjoka. The iwildcam 2020 competition dataset. *arXiv preprint arXiv:2004.10340*, 2020.

[11] Prithvijit Chattopadhyay, Yogesh Balaji, and Judy Hoffman. Learning to balance specificity and invariance for in and out of domain generalization. In *European Conference on Computer Vision*, 2020.

[12] Yinbo Chen, Zhuang Liu, Huijuan Xu, Trevor Darrell, and Xiaolong Wang. Meta-baseline: exploring simple meta-learning for few-shot learning. In *IEEE International Conference on Computer Vision*, 2021.

[13] Zhixiang Chi, Yang Wang, Yuanhao Yu, and Jin Tang. Test-time fast adaptation for dynamic scene deblurring via meta-auxiliary learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[14] Gordon Christie, Neil Fendley, James Wilson, and Ryan Mukherjee. Functional map of the world. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[15] Jeff Clune. Ai-gas: Ai-generating algorithms, an alternate paradigm for producing general artificial intelligence. *arXiv preprint arXiv:1905.10985*, 2019.

[16] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist to handwritten letters. In *International Joint Conference on Neural Networks*, 2017.

[17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.

[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. In *International Conference on Learning Representations*, 2021.

[19] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961*, 2021.

[20] Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *IEEE International Conference on Computer Vision*, 2013.

[21] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Lachine Learning*, 2017.

[22] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.

[23] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *International Conference on Machine Learning*, 2018.

[24] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *IEEE International Conference on Computer Vision*, 2015.

[25] Sam Gross, Marc'Aurelio Ranzato, and Arthur Szlam. Hard mixtures of experts for large scale weakly supervised vision. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

[27] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.

[28] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. In *Advances in Neural Information Processing Systems*, 2015.

[29] Weihua Hu, Gang Niu, Issei Sato, and Masashi Sugiyama. Does distributionally robust supervised learning give robust classifiers? In *International Conference on Machine Learning*, 2018.

[30] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[31] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3(1):79–87, 1991.

[32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2014.

[33] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, 2021.

[34] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012.

[35] Vinod K Kurmi, Venkatesh K Subramanian, and Vinay P Namboodiri. Domain impression: A source data free domain adaptation method. In *IEEE Winter Conference on Applications of Computer Vision*, 2021.

[36] Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. In *International Conference on Learning Representations*, 2020.

[37] Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. In *European Conference on Computer Vision*, 2018.

[38] Yizhuo Li, Miao Hao, Zonglin Di, Nitesh Bharadwaj Gundavarapu, and Xiaolong Wang. Test-time personalization with a transformer for human pose estimation. *Advances in Neural Information Processing Systems*, 2021.

[39] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning*, 2020.

[40] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International Conference on Machine Learning*, 2015.

[41] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, 2018.

[42] Daniela Massiceti, Luisa Zintgraf, John Bronskill, Lida Theodorou, Matthew Tobias Harris, Edward Cutrell, Cecily Morrison, Katja Hofmann, and Simone Stumpf. Orbit: A real-world few-shot dataset for teachable object recognition. In *IEEE International Conference on Computer Vision*, 2021.

[43] Toshihiko Matsuura and Tatsuya Harada. Domain generalization using a mixture of multiple latent domains. In *AAAI Conference on Artificial Intelligence*, 2020.

[44] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, 2013.

[45] Jaehoon Oh, Hyungjun Yoo, ChangHwan Kim, and Se-Young Yun. Boil: Towards representation change for few-shot learning. In *International Conference on Learning Representations*, 2021.

[46] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. In *AAAI Conference on Artificial Intelligence*, 2018.

[47] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *IEEE International Conference on Computer Vision*, 2019.

[48] James Requeima, Jonathan Gordon, John Bronskill, Sebastian Nowozin, and Richard E Turner. Fast and flexible multi-task classification using conditional neural adaptive processes. *Advances in Neural Information Processing Systems*, 2019.

[49] Shiori Sagawa*, Pang Wei Koh*, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2020.

[50] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[51] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International Conference on Machine Learning*, 2016.

[52] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. In *International Conference on Learning Representations*, 2018.

[53] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations*, 2017.

[54] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, 2017.

[55] Trevor Standley, Amir Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Which tasks should be learned together in multi-task learning? In *International Conference on Machine Learning*, 2020.

[56] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European Conference on Computer Vision*, 2016.

[57] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International Conference on Machine Learning*, 2020.

[58] J. Taylor, B. Earnshaw, B. Mabey, M. Victors, and J. Yosinski. Rxrx1: An image set for cellular morphological variation across many experimental batches. In *International Conference on Learning Representations*, 2019.

[59] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.

[60] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(11), 2008.

[61] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.

[62] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, 2016.

[63] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In *Advances in Neural Information Processing Systems*, 2018.

[64] Junfeng Wen, Russell Greiner, and Dale Schuurmans. Domain aggregation networks for multi-source domain adaptation. In *International Conference on Machine Learning*, 2020.

[65] Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[66] Ruijia Xu, Guanbin Li, Jihan Yang, and Liang Lin. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *IEEE International Conference on Computer Vision*, 2019.

[67] Luyu Yang, Yogesh Balaji, Ser-Nam Lim, and Abhinav Shrivastava. Curriculum manager for source selection in multi-source domain adaptation. In *European Conference on Computer Vision*, 2020.

[68] Shiqi Yang, Yaxing Wang, Joost van de Weijer, Luis Herranz, and Shangling Jui. Generalized source-free domain adaptation. In *IEEE International Conference on Computer Vision*, 2021.

[69] Christopher Yeh, Anthony Perez, Anne Driscoll, George Azzari, Zhongyi Tang, David Lobell, Stefano Ermon, and Marshall Burke. Using publicly available satellite imagery and deep learning to understand economic well-being in africa. *Nature Communications*, 2020.

[70] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[71] Chi Zhang, Nan Song, Guosheng Lin, Yun Zheng, Pan Pan, and Yinghui Xu. Few-shot incremental learning with continually evolved classifiers. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[72] Marvin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: Learning to adapt to domain shift. In *Advances in Neural Information Processing Systems*, 2021.

[73] Han Zhao, Shanghang Zhang, Guanhang Wu, José MF Moura, Joao P Costeira, and Geoffrey J Gordon. Adversarial multiple source domain adaptation. In *Advances in Neural Information Processing Systems*, 2018.

[74] Sicheng Zhao, Guangzhi Wang, Shanghang Zhang, Yang Gu, Yaxian Li, Zhichao Song, Pengfei Xu, Runbo Hu, Hua Chai, and Kurt Keutzer. Multi-source distilling domain adaptation. In *AAAI Conference on Artificial Intelligence*, 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Section 1.
   (b) Did you describe the limitations of your work? [Yes] See Section 6
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See supplemental material
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See supplemental material

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See 5.1

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 5.2

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See supplemental material

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes] See [33]

(b) Did you mention the license of the assets? [Yes] See supplemental material

(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]