

Architecture, Dataset and Model-Scale Agnostic Data-free Meta-Learning

Zixuan Hu¹ Li Shen^{2,*} Zhenyi Wang³ Tongliang Liu⁴ Chun Yuan^{1,*} Dacheng Tao²

¹Tsinghua Shenzhen International Graduate School, China; ²JD Explore Academy, China

³State University of New York at Buffalo, USA; ⁴The University of Sydney, Australia

huzixuan21@mails.tsinghua.edu.cn; mathshenli@gmail.com; zhenyiwa@buffalo.edu

tongliang.liu@sydney.edu.au; yuanc@sz.tsinghua.edu.cn; dacheng.tao@gmail.com

Abstract

The goal of data-free meta-learning is to learn useful prior knowledge from a collection of pre-trained models without accessing their training data. However, existing works only solve the problem in parameter space, which (i) ignore the fruitful data knowledge contained in the pre-trained models; (ii) can not scale to large-scale pre-trained models; (iii) can only meta-learn pre-trained models with the same network architecture. To address those issues, we propose a unified framework, dubbed **PURER**, which contains: (1) *e*Pisode *c*Urriculum *inve*Rsion (ECI) during data-free meta training; and (2) *inv*ErSION *calib*Ration following inner loop (ICFIL) during meta testing. During meta training, we propose ECI to perform pseudo episode training for learning to adapt fast to new unseen tasks. Specifically, we progressively synthesize a sequence of pseudo episodes by distilling the training data from each pre-trained model. The ECI adaptively increases the difficulty level of pseudo episodes according to the real-time feedback of the meta model. We formulate the optimization process of meta training with ECI as an adversarial form in an end-to-end manner. During meta testing, we further propose a simple plug-and-play supplement—ICFIL—only used during meta testing to narrow the gap between meta training and meta testing task distribution. Extensive experiments in various real-world scenarios show the superior performance of ours.

1. Introduction

Meta-learning [1, 28, 31] aims to learn useful prior knowledge (e.g., sensitive initialization) from a collection of similar tasks to facilitate the learning of new unseen tasks. Most meta-learning methods [3, 4, 7, 9, 15, 29, 30, 35, 36, 41, 42, 44] assume the access to the training and testing data of each task. However, this assumption is not always satisfied: many individuals and institutions only re-

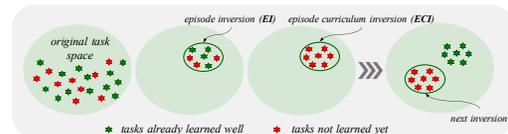


Figure 1. Episode Curriculum Inversion can improve the efficiency of pseudo episode training. At each episode, EI may repeatedly synthesize the tasks already learned well, while ECI only synthesizes harder tasks not learned yet.

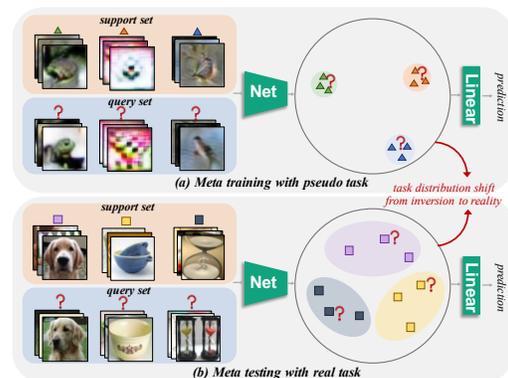


Figure 2. Task-distribution shift between meta training and testing. The pseudo data distilled from pre-trained models only contains partial semantic information learned by pre-trained models.

lease the pre-trained models instead of the data. This is due to data privacy, safety, or ethical issues in real-world scenarios, making the task-specific data difficult or impossible to acquire. For example, many pre-trained models with arbitrary architectures are released on GitHub without training data. However, when facing a new task, we need some prior knowledge learned from those pre-trained models so that the model can be adapted fast to the new task with few labeled examples. Thus, meta-learning from several pre-trained models without data becomes a critical problem, named *Data-free Meta-Learning (DFML)* [43].

Existing data-free meta-learning methods address the problem in the parameter space. Wang *et al.* [43] propose to meta-learn a black-box neural network by predicting the

*Corresponding authors: Li Shen and Chun Yuan

model parameters given the task embedding without data, which can be generalized to unseen tasks. The predicted model parameters with the average task embedding as input are served as the meta initialization for meta testing. However, this method has several drawbacks. First, they only merge the model in parameter space and ignore the underlying data knowledge that could be distilled from the pre-trained models. Second, their method can only be applied to small-scale pre-trained models since they use a neural network to predict the model parameters. Furthermore, their application scenarios are restricted to the case where all pre-trained models have the same architecture, limiting the real-world applicable scenarios.

In this work, we try to address all the above issues simultaneously in a unified framework, named **PURER** (see Fig. 3), which contains: (1) **eP**isode **cU**rriculum **inveR**sion (ECI) during data-free meta training; and (2) **invE**rsion **calibR**ation following inner loop (ICFIL) during meta testing, thus significantly expanding the application scenarios of DFML. During meta training, we propose ECI to perform pseudo episode training for learning to adapt fast to new unseen tasks. We progressively synthesize a sequence of pseudo episodes (tasks) by distilling the training data from each pre-trained model. ECI adaptively increases the difficulty level of pseudo episode according to the real-time feedback of the meta model. Specifically, we first introduce a small learnable dataset, named *dynamic dataset*. We initialize the dynamic dataset as Gaussian noise and progressively update it to better quality via one-step gradient descent for every iteration. For each episode, we construct a pseudo task by first sampling a subset of labels, and then sampling corresponding pseudo support data and query data. To improve the efficiency of pseudo episode training (see Fig. 1), we introduce the *curriculum mechanism* to synthesize episodes with an increasing level of difficulty. We steer the dynamic dataset towards appropriate difficulty so that only tasks not learned yet are considered at each iteration, which avoids repeatedly synthesizing the tasks already learned well. We design a *Gradient Switch* controlled by the real-time feedback from current meta model, to synthesize harder tasks only when the meta model has learned well on most tasks sampled from current dynamic dataset (see Fig. 3). Finally, we formulate the optimization process of meta training with ECI as an adversarial form in an end-to-end manner. We further propose a simple plug-and-play supplement—ICFIL—only used during meta testing to narrow the gap between meta training and meta testing task distribution (see Fig. 2). Overall, our proposed PURER can solve the DFML problem using the underlying data knowledge regardless of the dataset, scale and architecture of pre-trained models.

Our method is architecture, dataset and model-scale agnostic, thus substantially expanding the application scope

of DFML in real-world applications. We perform extensive experiments in various scenarios, including (i) **SS**: DFML with **S**ame dataset and **S**ame model architecture; (ii) **SH**: DFML with **S**ame dataset and **H**eterogeneous model architectures; (iii) **MH**: DFML with **M**ultiple datasets and **H**eterogeneous model architectures. For benchmarks of SS, SH and MH on CIFAR-FS and MiniImageNet, our method achieves significant performance gains in the range of 6.92% to 17.62%, 6.31% to 27.49% and 7.39% to 11.76%, respectively.

We summarize the main contributions as three-fold:

- We propose a new orthogonal perspective with respect to existing works to solve the data-free meta-learning problem by exploring the underlying data knowledge. Furthermore, our framework is architecture, dataset and model-scale agnostic, *i.e.*, it can be easily applied to various real-world scenarios.
- We propose a united framework, PURER, consisting of: (i) ECI to perform pseudo episode training with an increasing level of difficulty during meta training; (ii) ICFIL to narrow the gap between meta training and testing task distribution during meta testing.
- Our method achieves superior performance and outperforms the SOTA baselines by a large margin on various benchmarks of SS, SH and MH, which shows the effectiveness of our method.

2. Related Works

Meta-Learning. Meta-Learning [1, 28, 31] aims to learn general prior knowledge from a large collection of tasks such that the prior can be adapted fast to new unseen tasks. Most existing works [9–11, 13–16, 19–21, 23, 32, 37, 40, 42, 46–48, 50, 51, 53, 55] assume that the data associated with each task is available during meta training. Recently, data-free meta-learning [43] problem has attracted researchers’ attention. Wang *et al.* [43] propose to predict the meta initialization through a meta-trained black-box neural network. However, this method has several drawbacks. First, they focus on parameter space and ignore the underlying data knowledge that could be distilled from the pre-trained models. Second, their method can only be applied for small-scale pre-trained models since they use a neural network to predict the meta initialization. Furthermore, their application scenarios are only restricted to the case where all the pre-trained models have the same architecture, which reduces the applicable scenarios in real applications.

Model Inversion. Model inversion (MI) [12, 45, 56] aims to reconstruct training data from a pre-trained model. Existing works [5, 8, 12, 24, 26, 45, 54, 56, 57] synthesize images from pre-trained models via gradient descent. Those

works solely synthesize pseudo images without using external feedback, which causes the generated data to be suboptimal or unuseful for meta training. In contrast, our method uses the feedback from meta-learning model to adaptively adjust the synthesized images in an end-to-end way.

Curriculum Learning. Curriculum learning (CL) [2, 34, 38] aims to train a model in a meaningful order during loading training data, from easy to hard. Kumar *et al.* [18] propose self-paced learning (SPL) to provide automatic curriculum with a loss-based difficulty measure. Recently, Zhang *et al.* [52] propose a curriculum-based meta-learning method by forming harder tasks from each cluster, which merely applies to data-based meta-learning because it needs K-means clustering [25] before training. While in this work, we adaptively synthesize the curriculum on the fly under data-free setting to improve the efficiency of pseudo episode training.

3. Problem Setup

In this section, we first clarify the definition of Data-free Meta-Learning (DFML) problem. We then discuss the meta testing procedure for DFML problem.

3.1. Data-free Meta Learning Setup

We are given a collection of pre-trained models. Each pre-trained model is trained to solve a specific task without accessing their training data. The goal of DFML is to learn general prior knowledge (*e.g.*, sensitive initialization) which can be adapted fast to new unseen tasks. Here, we emphasize that the pre-trained models may have different architectures. Our method can work with arbitrary architectures of pre-trained models.

3.2. Meta Testing

During meta testing, several N -way K -shot tasks arrive together, which are called the target tasks. The classes appearing in the target tasks have never been seen during both pre-training and meta training. Each task contains a support set with N classes and K instances per class. The support set is used for adapting the meta initialization to the specific task. The query set is what the model actually needs to predict. Our goal is to adapt the model to the support set so that it can perform well on the query set. The final accuracy is measured by the average accuracy for these target tasks.

4. Methodology

In this section, we propose a unified framework PURER to solve DFML as illustrated in Fig. 3. In Section 4.2, we propose ECI to perform pseudo episode training by progressively synthesizing a sequence of pseudo episodes with an increasing level of difficulty during meta training. In Sec-

tion 4.3, we further propose ICFIL to eliminate the task-distribution shift issue during meta testing.

4.1. Preliminary: Episode Training

Our work aims to synthesize a sequence of pseudo episodes to perform pseudo episode training. Thus, we would like to introduce episode training briefly. In the case of MAML [9], each episode involves:

- *Outer Loop:* Update the *meta model* with the goal of improving the performance of *base model* on query set. The meta model works across episodes and learns good generalization from many episodes.
- *Inner Loop:* Perform *fast adaptation*. The *base model* takes the *meta model* as initialization and performs few steps of gradient descent over the support set. The *base model* works at the level of individual task.

Considering a per datum classification loss $l : \mathcal{X} \times \mathcal{Y} \times \Theta \rightarrow \mathbb{R}_+$, the *empirical loss* over a finite dataset \mathcal{B} is defined as $\mathcal{L}(\mathcal{B}; \theta) \triangleq \frac{1}{|\mathcal{B}|} \sum_{(x,y) \in \mathcal{B}} l(x, y; \theta)$. Thus, for a given task $\mathcal{T} = \{\mathcal{S}, \mathcal{Q}\}$, it becomes $\mathcal{L}(\mathcal{S}; \theta)$ and $\mathcal{L}(\mathcal{Q}; \theta)$. The bi-level optimization can be formulated as follows

$$\min_{\theta} \mathcal{L}_{outer}(\mathcal{T}; \theta) \triangleq \mathcal{L}(\mathcal{Q}; \theta^*), \quad (1a)$$

$$\begin{aligned} \text{s.t. } \theta^* &= \theta - \alpha_{inner} \nabla_{\theta} \mathcal{L}_{inner}(\mathcal{T}; \theta) \\ &\triangleq \theta - \alpha_{inner} \nabla_{\theta} \mathcal{L}(\mathcal{S}; \theta). \end{aligned} \quad (1b)$$

where α_{inner} is the step size in the inner loop.

4.2. Episode Curriculum Inversion (ECI)

We propose an ECI component to perform pseudo episode training for learning to adapt fast to new unseen tasks (see Fig. 3). We synthesize a sequence of pseudo episodes by distilling the training data from each pre-trained model. ECI adaptively increases the difficulty level of pseudo episodes according to the real-time feedback of the meta model.

Episode Inversion. We first propose the basic Episode Inversion (EI) to synthesize the pseudo task for each episode. At the beginning of meta training, we first introduce a small dynamic dataset \mathcal{D} initialized by Gaussian noise with mean $\mu = 0$ and standard deviation $\sigma = 1$. For each class, it only contains $K + M$ instances, consistent with N -way K -shot target task configuration with M instances each class in \mathcal{Q} . At each iteration, \mathcal{D} is dynamically updated to better quality by taking one-step gradient descent to minimize

$$\mathcal{L}_{inv}(\mathcal{D}) = \sum_{(\hat{x}, y) \in \mathcal{D}} l(\hat{x}, y; \psi) + \mathcal{R}_{prior}(\hat{x}) + \mathcal{R}_{feature}(\hat{x}). \quad (2)$$

ψ is the parameter of the corresponding pre-trained model. For each pair (\hat{x}, y) , \hat{x} is the pseudo image, and y is the manually assigned label to which we desire \hat{x} to belong.

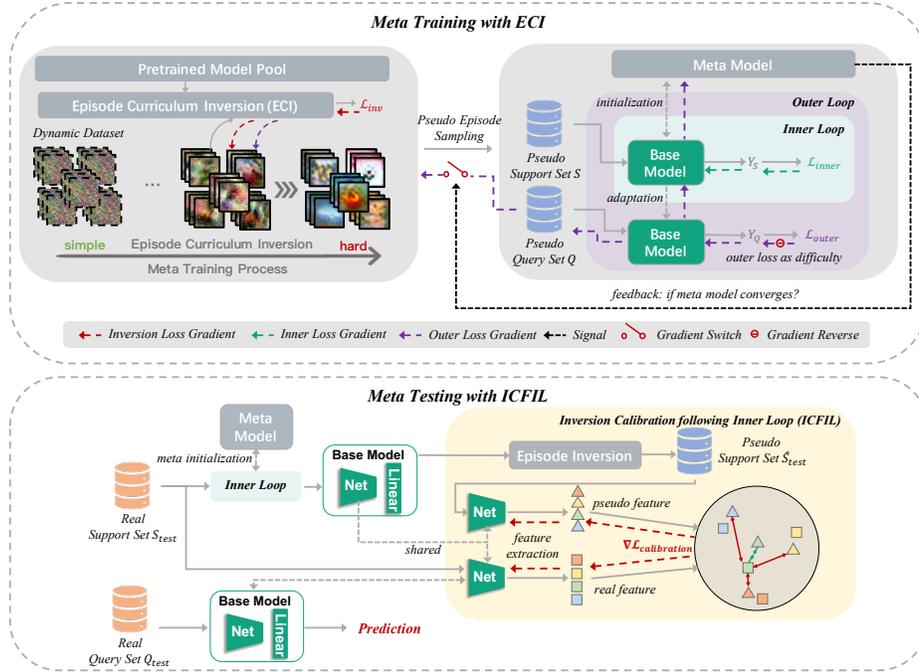


Figure 3. The overall pipeline of our proposed PURER consisting of ECI and ICFIL. For each episode during meta training, a pseudo episode is sampled from the dynamic dataset. The split pseudo support set and query set are used for the inner loop and outer loop of meta-learning. The real-time feedback of meta model controls the Gradient Switch. When the feedback is positive, the dynamic dataset is updated to synthesize harder tasks with larger outer loss for the next iteration by minimizing the reversed outer loss through gradient descent. During meta testing, the adapted base model after inner loop is calibrated via ICFIL. For brevity, we leave out the calibration for linear classifier head.

$l(\cdot)$ is the per datum classification loss (e.g., cross-entropy loss) mentioned in Sec. 4.1. $\mathcal{R}_{prior}(\cdot)$ is borrowed from DeepDream [27] to steer \hat{x} away from unreal images:

$$\mathcal{R}_{prior}(\hat{x}) = \alpha_{TV} \mathcal{R}_{TV}(\hat{x}) + \alpha_{l_2} \mathcal{R}_{l_2}(\hat{x}), \quad (3)$$

where \mathcal{R}_{TV} and \mathcal{R}_{l_2} are the total variance and l_2 norm of \hat{x} , with scaling factor α_{TV} and α_{l_2} . $\mathcal{R}_{feature}(\cdot)$ is a feature distribution regularization term used in DeepInversion [49] to minimize the distance between feature maps of pseudo images and original training images in each convolutional layer:

$$\mathcal{R}_{feature}(\hat{x}) = \sum_l \|\mu_l(\hat{x}) - \text{BN}_l(\text{running_mean})\| + \sum_l \|\sigma_l^2(\hat{x}) - \text{BN}_l(\text{running_variance})\|. \quad (4)$$

We feed a batch of pseudo images into the given pre-trained model to obtain feature maps in each convolutional layer. $\mu_l(\hat{x})$ and $\sigma_l^2(\hat{x})$ are the batch-wise mean and variance of those feature maps in the l^{th} convolutional layer. $\text{BN}_l(\text{running_mean})$ and $\text{BN}_l(\text{running_variance})$ are the running average mean and variance stored in the l^{th} Batch-Norm layer of pre-trained model, which is calculated during the pre-training period using original training images and can be obtained without access to original training images.

For each episode, we construct a pseudo task by first sampling a subset of labels, and then sampling the pseudo support data and query data from \mathcal{D} .

Curriculum Mechanism. Directly synthesizing the pseudo images without any feedback may repeatedly generate not optimal or not useful data for meta training (see Fig. 1). For example, we may repeatedly synthesize the tasks already learned well, thus increasing the unnecessary inversion cost. To find an effective episode inversion path, we propose to perform EI with curriculum mechanism, i.e. ECI. We aim to steer the dynamic dataset \mathcal{D} towards appropriate difficulty so that only tasks not learned yet are considered at each iteration, which avoids repeatedly synthesizing the tasks already learned well. We characterize the difficulty of \mathcal{D} for current meta model (parameterized by θ) as the expected outer loss over candidate tasks, i.e. $\mathbb{E}_{\mathcal{T} \in \mathcal{D}} [\mathcal{L}_{outer}(\mathcal{T}; \theta)]$. Specifically, a large outer loss presents a hard task; a task which has been learned well by current meta model, however, leads to a small outer loss. Thus, as illustrated in Fig. 3, we design a loss-based criteria to judge whether the meta model has learned well on tasks sampled from current \mathcal{D} . Specifically, at each iteration, for a batch of episodes $\{\mathcal{T}_i\}$, if the sum of the outer loss (training accuracy) over $\{\mathcal{T}_i\}$ has not decreased (increased) for 6 consecutive iterations, the meta model will send a pos-

Algorithm 1: Data-free Meta Training with ECI

```
1 REQUIRE Meta-model parameters  $\theta$ ; dynamic dataset
    $\mathcal{D}$ ; step size  $\alpha_{inner}$ ,  $\alpha_{outer}$  and  $\beta$ ; scaling factor  $\lambda$ .
2 Initialize  $\mathcal{D}$  and  $\theta$ 
3 while not done do
   // Episode Curriculum Inversion
4   Evaluate  $\mathcal{L}_{inv}(\mathcal{D})$  w.r.t. Eq. (2)
5   if feedback  $\Omega$  is Positive then
6     Randomly sample a task  $\mathcal{T}$  from  $\mathcal{D}$ 
7     Evaluate  $\mathcal{L}_{outer}(\mathcal{T}; \theta)$  w.r.t. Eq. (1)
8     Update  $\mathcal{D} \leftarrow$ 
        $\mathcal{D} - \beta * \nabla_{\mathcal{D}} (\mathcal{L}_{inv}(\mathcal{D}) - \lambda * \mathcal{L}_{outer}(\mathcal{T}; \theta))$ 
9   else
10    Update  $\mathcal{D} \leftarrow \mathcal{D} - \beta * \nabla_{\mathcal{D}} (\mathcal{L}_{inv}(\mathcal{D}))$ 
   // Meta Training
11  Randomly generate a batch of tasks  $\mathcal{T}_i$  from  $\mathcal{D}$ 
12  for all  $\mathcal{T}_i$  do
13    Evaluate  $\mathcal{L}_{outer}(\mathcal{T}_i; \theta)$  w.r.t. Eq. (1)
14  Update  $\theta \leftarrow \theta - \alpha_{outer} * \nabla_{\theta} \sum_{\mathcal{T}_i} \mathcal{L}_{outer}(\mathcal{T}_i; \theta)$ 
   // Update Feedback
15  Check feedback  $\Omega$ 
```

itive feedback, indicating that harder tasks are needed for the next iteration; else negative. To synthesize harder tasks only when the feedback is positive, we design the *Gradient Switch* controlled by the real-time feedback. Specifically, if the meta model has converged, the positive feedback will turn on the Gradient Switch so that the reversed gradient of outer loss can flow backwards to \mathcal{D} (see Fig. 3). In this way, we will update \mathcal{D} to maximize the outer loss through gradient descent, which means we can obtain harder tasks with a larger outer loss for the next iteration. The Gradient Switch controlled by feedback Ω works as an indicator function, i.e.,

$$\mathbb{I}(\Omega) = \begin{cases} 1, & \text{if } \Omega \text{ is positive;} \\ 0, & \text{if } \Omega \text{ is negative.} \end{cases} \quad (5)$$

In this way, ECI finds an effective episode inversion path by ignoring the tasks already learned well, thus improving the efficiency of pseudo episode training.

Adversarial Optimization. To adaptively synthesize episodes from easy to hard along with the meta training process, we perform pseudo episode training by adversarially updating the meta model (parameterized by θ) and dynamic dataset \mathcal{D} in an end-to-end manner. We search the optimal meta model parameters θ by minimizing the expected outer loss over candidate tasks from \mathcal{D} , i.e., $\mathbb{E}_{\mathcal{T} \in \mathcal{D}} [\mathcal{L}_{outer}(\mathcal{T}; \theta)]$, while we update the dynamic dataset \mathcal{D} to synthesize harder tasks with higher quality by minimizing reversed outer loss $-\mathbb{E}_{\mathcal{T} \in \mathcal{D}} [\mathcal{L}_{outer}(\mathcal{T}; \theta)]$ (when feedback Ω is positive) and inversion loss $\mathcal{L}_{inv}(\mathcal{D})$. Thus, we formulate the overall opti-

Algorithm 2: Meta Testing with ICFIL

```
1 REQUIRE The meta testing task  $\mathcal{T}_{test} = \{\mathcal{S}_{test}, \mathcal{Q}_{test}\}$ ;
   meta initialization  $\theta$ .
   // Inner Loop
2 Obtain  $\{\varphi_{test}^*, \mathbf{W}_{test}^*\}$  via fast adaptation over  $\mathcal{S}_{test}$ 
3 Obtain pseudo support set  $\hat{\mathcal{S}}_{test}$  w.r.t. Eq. (2)
   // Calibration
4 Calibrate  $\varphi_{test}^*$  by minimizing Eq. (7)
5 Train a new linear classifier head
   // Prediction
6 Make predictions on the query set  $\mathcal{Q}_{test}$ 
```

mization process as an adversarial form:

$$\min_{\theta} \max_{\mathcal{D}} \mathbb{E}_{\mathcal{T} \in \mathcal{D}} [-\mathcal{L}_{inv}(\mathcal{D}) + \mathbb{I}(\Omega) * \mathcal{L}_{outer}(\mathcal{T}; \theta)]. \quad (6)$$

To the end, we summarize the detailed procedure of ECI for data-free meta training in Alg. 1.

4.3. Inversion Calibration following Inner Loop

Pseudo images synthesized via ECI can not cover all semantic information of real images, thus leading to a gap between meta training and testing task distribution (see Fig. 2). In other words, the task-distribution shift issue in the data-free setting is much more significant than that in the setting of data-based meta-learning. To address this issue, we propose a simple plug-and-play supplement—ICFIL—only used during meta testing, to further narrow the gap between meta training and testing task distribution.

For an N -way target task $\mathcal{T}_{test} = \{\mathcal{S}_{test}, \mathcal{Q}_{test}\}$ during meta testing, we are given the meta initialization θ according to Alg. 1. The base model initialized by θ comprises a backbone $\phi : \mathbb{R}^{N_{in}} \rightarrow \mathbb{R}^{N_f}$ (parameterized by φ) and a linear classifier head parameterized by the weight matrix $\mathbf{W} \in \mathbb{R}^{N_f \times N}$, where N_{in} , N_f are the dimension of the input and embedding, respectively. We first perform fast adaptation over \mathcal{S}_{test} . The adapted parameters are $\theta_{test}^* = \{\varphi_{test}^*, \mathbf{W}_{test}^*\}$. Note that this is precisely identical to the standard meta testing procedure for regular meta-learning. The ICFIL does not assume access to query set \mathcal{Q}_{test} or any additional data and it can be viewed as a simple plug-and-play supplement for the standard meta testing period. We use ICFIL to calibrate the adapted base model before it predicts on the query set \mathcal{Q}_{test} . As illustrated in Fig. 3, we first synthesize a batch of pseudo images $\hat{\mathcal{S}}_{test}$ from the adapted base model according to Eq. (2). Inspired by the supervised contrastive learning [17], for a given image \mathbf{x} in \mathcal{S}_{test} , we define all pseudo images $\hat{\mathbf{x}}$ with the same label as positive samples $\hat{\mathbf{x}}^+$ while the others as negative samples $\hat{\mathbf{x}}^-$. The core idea behind ICFIL is to force the base model to focus on the overlapping information in both pseudo and real images, and ignore the others. The calibra-

tion loss is formulated as:

$$\mathcal{L}_{calibration}(\phi) = - \sum_{\mathbf{x} \in \mathcal{S}_{test}} \sum_{\hat{\mathbf{x}}^+} \log \frac{\exp[\phi(\mathbf{x})^T \phi(\hat{\mathbf{x}}^+)/\tau]}{\sum_{\hat{\mathbf{x}}} \exp[\phi(\mathbf{x})^T \phi(\hat{\mathbf{x}})/\tau]}.$$
(7)

where temperature parameter τ is widely used to help discriminate positive and negative samples. We then freeze the backbone and train a new linear classifier head by minimizing a standard classification loss (e.g., cross-entropy loss) over \mathcal{S}_{test} . The whole procedure for meta testing with IC-FIL is provided in Alg. 2.

5. Experiments

We conduct extensive experiments in various real-world scenarios to demonstrate the effectiveness of our proposed PURER framework. In Sec. 5.1, we present the basic experiment setup. In Secs. 5.2 to 5.4, we provide experiments about DFML in various scenarios, including SS, SH and MH. In Sec. 5.5, we provide more analysis of PURER.

5.1. Experiments Setup

Baselines. We compare PURER with four typical baselines: **(i) Random.** Randomly initialize the base model before the adaptation of each target task. **(ii) Average.** Average all pre-trained models to initialize the base model before the adaptation of each target task. **(iii) OTA [33].** Calculate the weighted average of all pre-trained models via optimal transport as the initialization of the base model before the adaptation of each target task. **(iv) DRO [43].** Meta-learn a neural network to predict the model parameters given the task embedding. The predicted model parameters with the average task embedding as input are served as the meta initialization. This is the first work to solve DFML by directly predicting the meta initialization in parameter space.

Datasets for meta testing. CIFAR-FS [3] and MiniImageNet [37] are commonly used in meta-learning, which consist of 100 classes with 600 images per class, respectively. We adopt a standard dataset split as [43]: 64 classes for meta training, 16 classes for meta validation and 20 classes for meta testing. All splits are non-overlapping. Note that we have no access to the meta training data in DFML setting.

Evaluation metric. We evaluate the performance by the average accuracy and standard deviation over 600 unseen target tasks sampled from meta testing dataset.

Scenarios. (i) **SS.** All pre-trained models are trained on Same dataset with Same architecture. (ii) **SH.** All pre-trained models are trained on Same dataset but with Heterogeneous architectures. (iii) **MH.** All pre-trained models are trained on Multiple datasets with Heterogeneous architectures. Existing DFML work, i.e., baseline DRO [42] and baselines Average and OTA merely apply to SS because they require pre-trained models have

the same architecture. We argue that SH and MH are more widely applicable scenarios in real-world applications and our PURER is a unified framework for addressing SS, SH and MH simultaneously.

5.2. Experiments of DFML in SS

Overview. We first perform experiments in SS scenario. We construct a collection of N -way tasks from the meta-training dataset and pre-train the model via standard supervised learning for each task. The collection of pre-trained models is used as the given resource models for DFML.

Implementation details. We take Conv4 as the architecture of all pre-trained models and the meta model, the same as regular meta-learning works [6, 9, 22, 39]. We take one-step gradient descent to perform fast adaptation for both meta training and testing. More implementation details are provided in Appendix A.

Table 1. Compare to baselines in SS scenario.

SS	Method	1-shot	5-shot
CIFAR-FS 5-way	Random	21.65 ± 0.45	21.59 ± 0.45
	Average	28.12 ± 0.62	32.15 ± 0.64
	OTA	29.10 ± 0.65	34.33 ± 0.67
	DRO	23.92 ± 0.49	24.34 ± 0.49
	Ours	38.66 ± 0.78	51.95 ± 0.79
MiniImageNet 5-way	Random	22.45 ± 0.41	23.48 ± 0.45
	Average	22.87 ± 0.39	26.13 ± 0.43
	OTA	24.22 ± 0.53	27.22 ± 0.59
	DRO	23.96 ± 0.42	25.81 ± 0.41
	Ours	31.14 ± 0.63	40.86 ± 0.64

Results. Tab. 1 shows the results for 5-way classification on CIFAR-FS and MiniImageNet under the scenario of SS. For CIFAR-FS, our method outperforms the best baseline by 9.56% and 17.62% for 1-shot and 5-shot learning, respectively. For MiniImageNet, our method outperforms the best baseline by 6.92% and 13.64% for 1-shot and 5-shot learning, respectively. The results show that random initialization is insufficient to perform fast adaptation. Simply averaging and OTA perform worse because: i) they fuse models layer-wise, but each pre-trained model are trained to solve different tasks, thus lacking the precise correspondence in parameter space; ii) they lack meta-learning objectives so that the fused model can not generalize well to new unseen tasks. DRO also performs not better due to the small number of pre-trained models. DRO trains a neural network to predict the meta initialization, thus requiring a relatively large number of pre-trained models as training resources. Our proposed method performs best because we leverage the fruitful underlying data knowledge contained in each pre-trained model instead of focusing on parameter space. We perform pseudo episode training with meta-learning ob-

jectives, thus learning to adapt fast to new unseen tasks with few labeled data.

5.3. Experiments of DFML in SH

Overview. To verify the broad applicability of our proposed method, we perform experiments under a more realistic and challenging scenario, **SH**. For each task, we pre-train the model with a randomly selected architecture.

Implementation details. For each task, we pre-train the model with an architecture randomly selected from Conv4, ResNet-10 and ResNet-18. Compared to Conv4, ResNet-10 and ResNet-18 are larger-scale neural networks. We take Conv4 as the meta model architecture. More implementation details are provided in Appendix A.

Results. Tab. 2 shows the results for 5-way classification on CIFAR-FS and MiniImageNet under the scenario of SH. For CIFAR-FS, our proposed method outperforms the baseline by 17.50% and 27.49% for 1-shot and 5-shot learning, respectively. For MiniImageNet, our method outperforms the baseline by 6.31% and 11.71% for 1-shot and 5-shot learning, respectively. Ours can apply to this real-world scenario and perform the best because we leverage the underlying data knowledge regardless of the scale and architecture of pre-trained models.

Table 2. Compare to baselines in SH scenario. Pre-trained models are trained with heterogeneous architectures (Conv4, ResNet-10 and ResNet-18).

SH	Method	1-shot	5-shot
CIFAR-FS 5-way	Random	21.65 ± 0.45	21.59 ± 0.45
	Ours	39.15 ± 0.70	49.08 ± 0.74
MiniImageNet 5-way	Random	22.45 ± 0.41	23.48 ± 0.45
	Ours	28.76 ± 0.60	35.19 ± 0.64

5.4. Experiments of DFML in MH

Overview. We further perform experiments in **MH** scenario, which provides a more challenging test. We construct a collection of tasks from different meta training datasets. For each task, we pre-train the model with a randomly selected architecture. In this way, we obtain a collection of pre-trained models trained on various datasets with different architectures.

Implementation details. During meta training, we construct each task from the meta training dataset randomly selected from CIFAR-FS and MiniImageNet. We pre-train the model with an architecture randomly selected from Conv4, ResNet-10 and ResNet-18. We take Conv4 as the meta model architecture. During meta testing, we report the average accuracy over 600 tasks sampled from both CIFAR-FS and MiniImageNet meta testing datasets. More implementation details are provided in Appendix A.

Results. Tab. 3 shows the results for 5-way classification under the scenario of MH. Our proposed method outperforms the baseline by 7.39% and 11.76% for 5-way 1-shot and 5-shot learning, respectively. Ours can apply to this more challenging scenario because our method is dataset-agnostic and architecture-agnostic. In other word, PURER works well on pre-trained models trained on various datasets with different architectures.

Table 3. Compare to baselines in MH scenario. Pre-trained models are from multiple datasets (CIFAR-100 and MiniImageNet) with heterogeneous architectures (Conv4, ResNet-10 and ResNet-18).

MH	Method	1-shot	5-shot
5-way	Random	21.11 ± 0.41	21.34 ± 0.40
	Ours	28.50 ± 0.63	33.10 ± 0.69

5.5. Ablation Study

Effectiveness of each component in PUERE. Tab. 4 shows the results of ablation studies in SS scenario. We first introduce a vanilla baseline (EI) performing pseudo episode training without curriculum mechanism and meta testing without ICFIL. EI still outperforms the best baselines in Tab. 1 by 3.8% and 6.55% for 1-shot and 5-shot learning, respectively, which verifies the effectiveness of the basic idea of DFML leveraging the underlying data knowledge. To further evaluate the effectiveness of curriculum mechanism, we append curriculum to EI, *i.e.*, ECI. We can observe that with curriculum mechanism, the performance can be improved by 1.09% and 2.93% for 1-shot and 5-shot learning, respectively. Similarly, we adopt ICFIL during meta testing for EI, *i.e.*, EI + ICFIL, and the performance can be improved by 1.84% and 2.89% for 1-shot and 5-shot learning, respectively. By introducing both curriculum mechanism for episode inversion and ICFIL for meta testing, we achieve the best performance with a boosting gain of 3.12% and 7.09% for 1-shot and 5-shot learning, respectively, which demonstrates the effectiveness of the joint schema.

Table 4. Ablation studies on MiniImageNet in SS scenario.

SS	Component		Accuracy	
	Curriculum	ICFIL	5-way 1-shot	5-way 5-shot
EI			28.02 ± 0.58	33.77 ± 0.63
	✓		29.11 ± 0.59	36.70 ± 0.64
		✓	29.86 ± 0.66	36.66 ± 0.66
Ours	✓	✓	31.14 ± 0.63	40.86 ± 0.64

Effect of the number of training class for each task. To evaluate the performance difference for the different number of training class for each task, we perform the experiments for 10-way classification problem in SS scenario.

Table 5. Effect of the number of training class for each task in SS.

SS	Method	1-shot	5-shot
CIFAR-FS 10-way	Random	10.26 ± 0.22	10.32 ± 0.22
	Average	12.23 ± 0.30	14.86 ± 0.30
	OTA	13.40 ± 0.32	15.17 ± 0.36
	DRO	11.29 ± 0.25	11.56 ± 0.26
	Ours	22.88 ± 0.41	36.19 ± 0.45
MiniImageNet 10-way	Random	11.20 ± 0.23	11.34 ± 0.24
	Average	12.83 ± 0.27	13.94 ± 0.27
	OTA	11.92 ± 0.25	13.25 ± 0.29
	DRO	11.78 ± 0.20	13.08 ± 0.24
	Ours	18.16 ± 0.37	26.92 ± 0.37

Table 6. Compare to MAML in SS scenario. †: use real meta-training dataset of equal size to the dynamic dataset used in DFML. Grey: data-based meta-learning method.

SS	Method	1-shot	5-shot
CIFAR-FS 5-way	MAML†	34.18 ± 0.75	41.20 ± 0.71
	Ours	38.66 ± 0.78	51.95 ± 0.79
CIFAR-FS 10-way	MAML†	19.05 ± 0.38	21.04 ± 0.35
	Ours	22.88 ± 0.41	36.19 ± 0.45

Tab. 5 shows the results. For CIFAR-FS, our method outperforms the best baseline by 9.48% and 21.02% for 1-shot and 5-shot learning, respectively. For MiniImageNet, our method outperforms the best baseline by 5.33% and 12.98% for 1-shot and 5-shot learning, respectively. Compared to 5-way classification, the accuracy of 10-way classification is lower because it is more challenging. Ours outperforms all baselines in both 5-way and 10-way classification problems. **Comparison to MAML.** To evaluate the effectiveness of DFML, we compare our method with regular data-based meta-learning work MAML [9] under the scenario of SS. The results are shown in Tab. 6. We introduce MAML†, which uses the real meta training dataset of equal size to the dynamic dataset used in DFML. For 5-way classification on CIFAR-FS, our method outperforms MAML† by 4.48% and 10.75% for 1-shot and 5-shot learning, respectively. For 10-way classification on CIFAR-FS, our method outperforms MAML† by 3.83% and 15.15% for 1-shot and 5-shot learning, respectively. MAML† performs worse because MAML† uses a fixed meta training dataset, which only provides a limited number of tasks in a random order. In contrast, PURER can provide a more diverse collection of tasks with an increasing level of difficulty because the dynamic dataset can be adaptively updated according to the real-time feedback of meta-learning model.

Contributions from curriculum mechanism. Fig. 4 shows the testing accuracy with and without curriculum mechanism, respectively. After 4000 iterations, we adopt curriculum mechanism to adaptively synthesize harder tasks. Without curriculum mechanism, the performance converges to a lower peak and even begins to decline, which can be explained as overfitting to the tasks of single-level difficulty. In contrast, with curriculum mechanism, the meta model

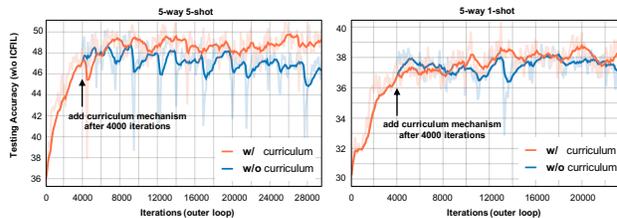


Figure 4. Effect of curriculum mechanism on testing performance on CIFAR-FS in SS scenario. Solid curve: smoothed performance curve. Transparent curve: original performance curve.

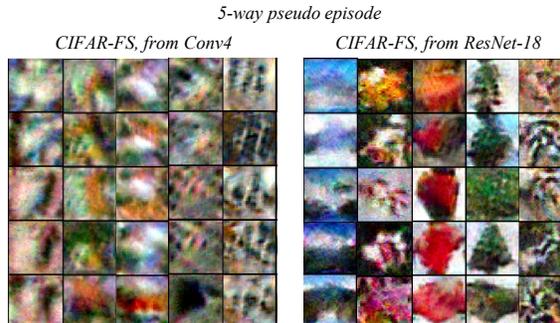


Figure 5. (left) pseudo images synthesized from Conv4. (right) pseudo images synthesized from ResNet-18. Each column corresponds to one class.

can keep learning from harder tasks and finally achieves a peak with higher testing performance.

Visualizations of ECI. Fig. 5 shows visualizations of 5-way pseudo episodes synthesized by ECI from models pre-trained on CIFAR-FS with architectures of Conv4 and ResNet-18, respectively. We can observe that compared with Conv4, a deeper pre-trained model (ResNet-18) can synthesize images of better quality.

Hyperparameter sensitivity. Results in Appendix B show that the performance of our method is stable with the changes of λ value.

6. Conclusion

We propose a new orthogonal perspective with respect to existing works to solve DFML problem for exploring underlying data knowledge. Our framework PURER is architecture, dataset and model-scale agnostic, thus working well in various real-world scenarios. It contains two components: (i) ECI to perform pseudo episode training with an increasing order of difficulty during meta training; (ii) ICFIL to settle the task-distribution shift issue during meta testing. Extensive experiments show the effectiveness of PURER.

Acknowledgement. This work is supported by the National Key R&D Program of China (2022YFB4701400/4701402), the SZSTC Grant (JCYJ20190809172201639, WDZC20200820200655001), Shenzhen Key Laboratory (ZDSYS20210623092001004).

References

- [1] Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gescei. On the optimization of a synaptic learning rule. In *Optimality in Biological and Artificial Networks?*, pages 281–303. Routledge, 2013. 1, 2
- [2] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009. 3
- [3] Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. *arXiv preprint arXiv:1805.08136*, 2018. 1, 6
- [4] John Bronskill, Daniela Massiceti, Massimiliano Patacchiola, Katja Hofmann, Sebastian Nowozin, and Richard Turner. Memory efficient meta-learning with large images. *Advances in Neural Information Processing Systems*, 34:24327–24339, 2021. 1
- [5] Akshay Chawla, Hongxu Yin, Pavlo Molchanov, and Jose Alvarez. Data-free knowledge distillation for object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3289–3298, 2021. 2
- [6] Jiaxin Chen, Li-Ming Zhan, Xiao-Ming Wu, and Fu-lai Chung. Variational metric scaling for metric-based meta-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 3478–3485, 2020. 6, 11
- [7] Qi Chen, Changjian Shui, and Mario Marchand. Generalization bounds for meta-learning: An information-theoretic analysis. *Advances in Neural Information Processing Systems*, 34:25878–25890, 2021. 1
- [8] Xiang Deng and Zhongfei Zhang. Graph-free knowledge distillation for graph neural networks. *arXiv preprint arXiv:2105.07519*, 2021. 2
- [9] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017. 1, 2, 3, 6, 8, 11
- [10] Chelsea Finn, Aravind Rajeswaran, Sham Kakade, and Sergey Levine. Online meta-learning. In *International Conference on Machine Learning*, pages 1920–1930. PMLR, 2019. 2
- [11] Chelsea Finn, Kelvin Xu, and Sergey Levine. Probabilistic model-agnostic meta-learning. *Advances in neural information processing systems*, 31, 2018. 2
- [12] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pages 1322–1333, 2015. 2
- [13] James Harrison, Apoorva Sharma, Chelsea Finn, and Marco Pavone. Continuous meta-learning without tasks. *Advances in neural information processing systems*, 33:17571–17581, 2020. 2
- [14] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Cross attention network for few-shot classification. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [15] Kyle Hsu, Sergey Levine, and Chelsea Finn. Un-supervised learning via meta-learning. *arXiv preprint arXiv:1810.02334*, 2018. 1, 2
- [16] Mikhail Khodak, Maria-Florina F Balcan, and Ameet S Talwalkar. Adaptive gradient-based meta-learning methods. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [17] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020. 5
- [18] M Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. *Advances in neural information processing systems*, 23, 2010. 3
- [19] Aoxue Li, Weiran Huang, Xu Lan, Jiashi Feng, Zhenguo Li, and Liwei Wang. Boosting few-shot learning with adaptive margin loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12576–12584, 2020. 2
- [20] Aoxue Li, Tiange Luo, Tao Xiang, Weiran Huang, and Liwei Wang. Few-shot learning with global class representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9715–9724, 2019. 2
- [21] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 2
- [22] Chenghao Liu, Zhihao Wang, Doyen Sahoo, Yuan Fang, Kun Zhang, and Steven CH Hoi. Adaptive task sampling for meta-learning. In *European Conference on Computer Vision*, pages 752–769. Springer, 2020. 6, 11
- [23] Lu Liu, Tianyi Zhou, Guodong Long, Jing Jiang, and Chengqi Zhang. Learning to propagate for graph meta-learning. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [24] Zechun Liu, Zhiqiang Shen, Yun Long, Eric Xing, Kwang-Ting Cheng, and Chas Leichner. Data-free neural architecture search via recursive label calibration. *arXiv preprint arXiv:2112.02086*, 2021. 2
- [25] Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982. 3
- [26] Raphael Gontijo Lopes, Stefano Fenu, and Thad Starner. Data-free knowledge distillation for deep neural networks. *arXiv preprint arXiv:1710.07535*, 2017. 2
- [27] Alexander Mordvintsev, Christopher Olah, and Mike Tyka. Inceptionism: Going deeper into neural networks, 2015. 4
- [28] Devang K Naik and Richard J Mammone. Meta-neural networks that learn by learning. In *[Proceedings 1992] IJCNN International Joint Conference on Neural Networks*, volume 1, pages 437–442. IEEE, 1992. 1, 2
- [29] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018. 1
- [30] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. *Advances in neural information processing systems*, 32, 2019. 1

- [31] Jürgen Schmidhuber. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. PhD thesis, Technische Universität München, 1987. 1, 2
- [32] Christian Simon, Piotr Koniusz, and Mehrtaash Harandi. Meta-learning for multi-label few-shot classification. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3951–3960, 2022. 2
- [33] Sidak Pal Singh and Martin Jaggi. Model fusion via optimal transport. *Advances in Neural Information Processing Systems*, 33:22045–22055, 2020. 6
- [34] Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. Curriculum learning: A survey. *International Journal of Computer Vision*, pages 1–40, 2022. 3
- [35] Yue Sun, Adhyyan Narang, Ibrahim Gulluk, Samet Oymak, and Maryam Fazel. Towards sample-efficient overparameterized meta-learning. *Advances in Neural Information Processing Systems*, 34:28156–28168, 2021. 1
- [36] Joaquin Vanschoren. *Meta-learning*. In *Automated machine learning*, pages 35–61. Springer, Cham, 2019. 1
- [37] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016. 2, 6
- [38] Xin Wang, Yudong Chen, and Wenwu Zhu. A survey on curriculum learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):4555–4576, 2022. 3
- [39] Yan Wang, Wei-Lun Chao, Kilian Q Weinberger, and Laurens van der Maaten. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning. *arXiv preprint arXiv:1911.04623*, 2019. 6, 11
- [40] Zhenyi Wang, Tieshang Duan, Le Fang, Qiuling Suo, and Mingchen Gao. Meta learning on a sequence of imbalanced domains with difficulty awareness. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8947–8957, 2021. 2
- [41] Zhenyi Wang, Li Shen, Tieshang Duan, Donglin Zhan, Le Fang, and Mingchen Gao. Learning to learn and remember super long multi-domain task sequence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7982–7992, 2022. 1
- [42] Zhenyi Wang, Li Shen, Le Fang, Qiuling Suo, Donglin Zhan, Tieshang Duan, and Mingchen Gao. Meta-learning with less forgetting on large-scale non-stationary task distributions. In *European Conference on Computer Vision*, pages 221–238. Springer, 2022. 1, 2, 6
- [43] Zhenyi Wang, Xiaoyang Wang, Li Shen, Qiuling Suo, Kaiqiang Song, Dong Yu, Yan Shen, and Mingchen Gao. Meta-learning without data via wasserstein distributionally-robust model fusion. In *The 38th Conference on Uncertainty in Artificial Intelligence*, 2022. 1, 2, 6
- [44] Zhenyi Wang, Yang Zhao, Ping Yu, Ruiyi Zhang, and Changyou Chen. Bayesian meta sampling for fast uncertainty adaptation. In *International Conference on Learning Representations*, 2019. 1
- [45] Xi Wu, Matthew Fredrikson, Somesh Jha, and Jeffrey F Naughton. A methodology for formalizing model-inversion attacks. In *2016 IEEE 29th Computer Security Foundations Symposium (CSF)*, pages 355–370. IEEE, 2016. 2
- [46] Shuo Yang, Lu Liu, and Min Xu. Free lunch for few-shot learning: Distribution calibration. *arXiv preprint arXiv:2101.06395*, 2021. 2
- [47] Huaxiu Yao, Yu Wang, Ying Wei, Peilin Zhao, Mehrdad Mahdavi, Defu Lian, and Chelsea Finn. Meta-learning with an adaptive task scheduler. *Advances in Neural Information Processing Systems*, 34:7497–7509, 2021. 2
- [48] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8808–8817, 2020. 2
- [49] Hongxu Yin, Pavlo Molchanov, Jose M Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via deep-inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8715–8724, 2020. 4
- [50] Mingzhang Yin, George Tucker, Mingyuan Zhou, Sergey Levine, and Chelsea Finn. Meta-learning without memorization. *arXiv preprint arXiv:1912.03820*, 2019. 2
- [51] Jaesik Yoon, Taesup Kim, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn. Bayesian model-agnostic meta-learning. *Advances in neural information processing systems*, 31, 2018. 2
- [52] Ji Zhang, Jingkuan Song, Yazhou Yao, and Lianli Gao. Curriculum-based meta-learning. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1838–1846, 2021. 3
- [53] Jian Zhang, Chenglong Zhao, Bingbing Ni, Minghao Xu, and Xiaokang Yang. Variational few-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1685–1694, 2019. 2
- [54] Lin Zhang, Li Shen, Liang Ding, Dacheng Tao, and Ling-Yu Duan. Fine-tuning global model via data-free knowledge distillation for non-iid federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10174–10183, 2022. 2
- [55] Ruixiang Zhang, Tong Che, Zoubin Ghahramani, Yoshua Bengio, and Yangqiu Song. Metagan: An adversarial approach to few-shot learning. *Advances in neural information processing systems*, 31, 2018. 2
- [56] Yuheng Zhang, Ruoxi Jia, Hengzhi Pei, Wenxiao Wang, Bo Li, and Dawn Song. The secret revealer: Generative model-inversion attacks against deep neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 253–261, 2020. 2
- [57] Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous federated learning. In *International Conference on Machine Learning*, pages 12878–12889. PMLR, 2021. 2