

# Online Meta-Learning for Multi-Source and Semi-Supervised Domain Adaptation

Da Li<sup>1</sup>[0000-0002-2101-2989] and Timothy Hospedales<sup>1,2</sup>[0000-0003-4867-7486]

<sup>1</sup> Samsung AI Center Cambridge, UK

<sup>2</sup> The University of Edinburgh, UK.

dali.academic@gmail.com; t.hospedales@ed.ac.uk

**Abstract.** Domain adaptation (DA) is the topical problem of adapting models from labelled source datasets so that they perform well on target datasets where only unlabelled or partially labelled data is available. Many methods have been proposed to address this problem through different ways to minimise the domain shift between source and target datasets. In this paper we take an orthogonal perspective and propose a framework to further enhance performance by meta-learning the initial conditions of existing DA algorithms. This is challenging compared to the more widely considered setting of few-shot meta-learning, due to the length of the computation graph involved. Therefore we propose an online shortest-path meta-learning framework that is both computationally tractable and practically effective for improving DA performance. We present variants for both multi-source unsupervised domain adaptation (MSDA), and semi-supervised domain adaptation (SSDA). Importantly, our approach is agnostic to the base adaptation algorithm, and can be applied to improve many techniques. Experimentally, we demonstrate improvements on classic (DANN) and recent (MCD and MME) techniques for MSDA and SSDA, and ultimately achieve state of the art results on several DA benchmarks including the largest scale DomainNet.

**Keywords:** meta-learning, domain adaptation

## 1 Introduction

Contemporary deep learning methods now provide excellent performance across a variety of computer vision tasks when ample annotated training data is available. However, this performance often degrades rapidly if models are applied to novel domains with very different data statistics from the training data, which is a problem known as domain shift. Meanwhile, data collection and annotation for every possible domain of application is expensive and sometimes impossible. This challenge has motivated extensive study in the area of domain adaptation (DA), which addresses training models that work well on a target domain using only unlabelled or partially labelled target data from that domain together with labelled data from a source domain.




Fig. 1: Left: Exact meta-learning of initial condition with inner loop training DA to convergence is intractable. Right: Online meta-learning alternates between meta-optimization and domain adaptation.

Several variants of the domain adaptation problem have been studied. Single-source domain adaptation (SDA) considers adaptation from a single source domain [5, 6], while multi-source domain adaptation (MSDA) considers improving cross-domain generalization by aggregating information across multiple sources [35, 47]. Unsupervised domain adaptation (UDA) learns solely from unlabelled data in the target domain [16, 43], while semi-supervised domain adaptation (SSDA) learns from a mixture of labelled and unlabelled target domain data [10, 9]. The main means of progress has been developing improved methods for aligning representations between source(s) and the target in order to improve generalization. These methods span distribution alignment, for example by maximum mean discrepancy (MMD) [27, 48], domain adversarial training [16, 43], and cycle consistent image transformation [26, 18].

In this paper we adopt a novel research perspective that is complementary to all these existing methods. Rather than proposing a new domain adaptation strategy, we study a meta-learning framework for improving these existing adaptation algorithms. Meta-learning (a.k.a. learning to learn) has a long history [44, 45], and has re-surfed recently, especially due to its efficacy in improving few-shot deep learning [50, 40, 12]. Meta-learning pipelines aim to improve learning by training some aspect of a learning algorithm such a comparison metric [50], model optimizer [40] or model initialisation [12], so as to improve outcomes according to some meta-objective such as few-shot learning efficacy [50, 40, 12] or learning speed [1]. In this paper we provide a first attempt to define a meta-learning framework for improving domain adaptive learning.

We take the perspective of meta-optimizing the *initial condition* [33, 12] of domain adaptive learning<sup>3</sup>. While there are several facets of algorithms that can be meta-learned such as hyper-parameters [14] and learning rates [24]; these are somewhat tied to the base learning algorithm (domain adaptive algorithm in our case). In contrast, our framework is *algorithm agnostic* in that it can be used to improve many existing gradient-based domain adaptation algorithms.

Furthermore we develop variants for both unsupervised multi-source adaptation, as well as semi-supervised single source adaptation, thus providing broad

<sup>3</sup> One may not think of domain adaptation as being sensitive to initial condition, but given the lack of target domain supervision to guide learning, different initialization can lead to a significant 10-15% difference in accuracy (see Supplementary material).

potential benefit to existing frameworks and application settings. In particular we demonstrate application of our framework to the classic domain adversarial neural network (DANN) [16] algorithm, as well as the recent maximum-classifier discrepancy (MCD) [43], and min-max entropy (MME) [41] algorithms.

Meta-learning can often be cleanly formalised as a bi-level optimization problem [14, 39]: where an outer loop optimizes the meta-parameter of interest (such as the initial condition in our case) with respect to some meta-loss (such as performance on a validation set); and the inner loop runs the learning algorithm conditioned on the chosen meta-parameter. This is tricky to apply directly in a domain adaptation scenario however, because: (i) The computation graph of the inner loop is typically long (unlike the popular few-shot meta-learning setting [12]), making meta-optimization intractable, and (ii) Especially in unsupervised domain adaptation, there is no labelled data in the target domain to define a supervised learning loss for the outer-loop meta-objective. We surmount these challenges by proposing a simple, fast and efficient meta-learning strategy based on online shortest path gradient descent [36], and defining meta-learning pipelines suited for meta-optimization of domain adaptation problems. Although motivated by initial condition learning, our online algorithm ultimately has the interpretation of intermittently performing meta-update(s) of the parameters in order to achieve the best outcome from the following DA updates (Fig 1).

Overall, our contributions are: (i) Introducing a meta-learning framework suitable for both multi-source and semi-supervised domain adaptation settings, (ii) We demonstrate the algorithm agnostic nature of our framework through its application to several base domain adaptation methods including MME [41], DANN [16] and MCD [43], (iii) Applying our meta-learner to these base adaptation methods, we achieve state of the art performance on several MSDA and SSDA benchmarks, including the largest-scale DA benchmark, DomainNet [38].

## 2 Related Work

**Single-Source Domain Adaptation** Single-source unsupervised domain adaptation (SDA) is a well established area [5, 48, 27, 16, 29, 6, 30, 43, 31, 19]. Theoretical results have bound the cross-domain generalization error in terms of domain divergence [4], and numerous algorithms have been proposed to reduce the divergence between source and target features. Representative approaches include minimising MMD distribution shift [48, 27] or Wasserstein distance [2, 51], adversarial training [16, 43, 20] or alignment by cycle-consistent image translation [26, 18]. Given the difficulty of SDA, studies have considered exploiting semi-supervised or multi-source adaptation where possible.

**Semi-Supervised Domain Adaptation** This setting assumes that besides the labelled source and unlabelled target domain data, there are a few labelled samples available in the target domain. Exploiting the few target labels allows better domain alignment compared to purely unsupervised approaches. Representative approaches are based on regularization [9], subspace learning [54], label smoothing [10] and entropy minimisation in the target domain [17]. The state

of the art method in this area, MME, extends the entropy minimisation idea to adversarial entropy minimisation in a deep network setting [41].

**Multi-Source Domain Adaptation** This setting assumes there are multiple labelled source domains for training. In deep learning, simply aggregating all source domains data together often already improves performance due to bigger datasets learning a stronger representation. Theoretical results based on  $\mathcal{H}$ -divergence [4] can still apply after aggregation, and existing SDA methods that attempt to reduce source-target divergence [5, 16, 43] can be used. Meanwhile, new generalization bounds for MSDA have been derived [55, 38], which motivate algorithms that align amongst source domains as well as between source and target. Nevertheless, practical deep network optimization is non-convex, and the degree of alignment achieved depends on the details of the optimization strategy. Therefore our paradigm of meta-learning the initial condition of optimization is compatible with, and complementary to, all this prior work.

**Meta-Learning for Neural Networks** Meta-Learning (learning to learn) [44, 46] has experienced a recent resurgence. This has largely been driven by its efficacy for few-shot deep learning via initial condition learning [12], optimizer learning [40] and embedding learning [50]. More generally it has been applied to improve optimization efficiency [1], reinforcement learning [37], gradient-based hyperparameter optimization [14] and neural architecture search [25]. We start from the perspective of MAML [12], in terms of focusing on learning initial conditions of neural network optimization. However besides the different application (domain adaptation vs few-shot learning), our resulting algorithm is very different as we end up performing meta-optimization online *while* solving the target task rather than in advance of solving it [12, 39]. A few recent studies also attempted online meta-learning [53, 13], but these are designed specifically to backprop through RL [53] or few-shot supervised [13] learning. Meta-learning with domain adaptation in the inner loop has not been studied before now.

In terms of learning with multiple domains a few studies [21, 3, 11] have considered meta-learning for multi-source *domain generalization*, which evaluates the ability of models to generalise directly without adaptation. In practice these methods use supervised learning in their inner optimization. No meta-learning method has been proposed for the domain adaptation problem addressed here.

### 3 Methodology

#### 3.1 Background

**Unsupervised Domain Adaptation** Domain adaptation techniques aim to reduce the domain shift between source domain(s)  $\mathcal{D}_S$  and target domain  $\mathcal{D}_T$ , in order that a model trained on labels from  $\mathcal{D}_S$  performs well when deployed on  $\mathcal{D}_T$ . Commonly such algorithms train a model  $\theta$  with a loss  $\mathcal{L}_{\text{uda}}$  that breaks down into a term for supervised learning on the source domain  $\mathcal{L}_{\text{sup}}$  and an adaptation loss  $\mathcal{L}_a$  that attempts to align the target and source data

$$\mathcal{L}_{\text{uda}}(\theta, \mathcal{D}_S, \overline{\mathcal{D}}_T) = \mathcal{L}_{\text{sup}}(\theta, \mathcal{D}_S) + \lambda \mathcal{L}_a(\theta, \mathcal{D}_S, \overline{\mathcal{D}}_T). \quad (1)$$

**Algorithm 1** Meta-Update Initial Condition

---

```

Function UpdateIC( $\Theta_0, \mathcal{D}_{\text{tr}}, \mathcal{D}_{\text{val}}, \mathcal{L}_{\text{inner}}, \mathcal{L}_{\text{outer}}$ ) :
for  $j = 1, 2, \dots, J$  do                                     // Inner-level optimization
     $\Theta_j = \Theta_{j-1} - \alpha \nabla_{\Theta_{j-1}} \mathcal{L}_{\text{inner}}(\Theta_{j-1}, (\mathcal{D}_{\text{tr}})_j)$ 
end for
 $\Theta_0 = \Theta_0 - \alpha \nabla_{\Theta_0} \mathcal{L}_{\text{outer}}(\Theta_J, \mathcal{D}_{\text{val}})$                                      // Outer-level step
Output:  $\Theta_0$ 

```

---

We use notation  $\mathcal{D}_S$  and  $\overline{\mathcal{D}}_T$  to indicate that the source and target domains contain labelled and unlabelled data respectively. Many existing domain adaptation algorithms [48, 16, 43, 38] fit into this template, differing in their definition of the domain alignment loss  $\mathcal{L}_a$ . In the case of multi-source adaptation [38],  $\mathcal{D}_S$  may contain several source domains  $\mathcal{D}_S = \{D_1, \dots, D_N\}$  and the first supervised learning term  $\mathcal{L}_{\text{sup}}$  sums the performance on all of these.

**Semi-Supervised Domain Adaptation** In the SSDA setting [41], we assume a sparse set of labelled target data  $\mathcal{D}_T$  is provided along with a large set of unlabelled target data  $\overline{\mathcal{D}}_T$ . The goal is to learn a model that fits both the source and few-shot target labels  $\mathcal{L}_{\text{sup}}$ , while also aligning the unlabelled target data to the source with an adaptation loss  $\mathcal{L}_a$ .

$$\begin{aligned} \mathcal{L}_{\text{ssda}}(\Theta, \mathcal{D}_S, \overline{\mathcal{D}}_T, \mathcal{D}_T) = & \mathcal{L}_{\text{sup}}(\Theta, \mathcal{D}_S) + \mathcal{L}_{\text{sup}}(\Theta, \mathcal{D}_T) \\ & + \lambda \mathcal{L}_a(\Theta, \mathcal{D}_S, \overline{\mathcal{D}}_T) \end{aligned} \quad (2)$$

Several existing algorithms [28, 41, 16] fit this template and hence can potentially be optimized by our framework.

**Meta Learning Model Initialisation** The problem of meta-learning the initial condition of an optimization can be seen as a bi-level optimization problem [14, 39]. In this view there is a standard task-specific (inner) algorithm of interest whose initial condition we wish to optimize, and an outer-level meta-algorithm that optimizes that initial condition. This setup can be described as

$$\Theta = \underbrace{\operatorname{argmin}_{\Theta} \mathcal{L}_{\text{outer}}}_{\text{Outer-level}} \left( \overbrace{\mathcal{L}_{\text{inner}}(\Theta, \mathcal{D}_{\text{tr}})}^{\text{Inner-level}}, \mathcal{D}_{\text{val}} \right) \quad (3)$$

where  $\mathcal{L}_{\text{inner}}(\Theta, \mathcal{D}_{\text{tr}})$  denotes the standard loss of the base task-specific algorithm on its training set.  $\mathcal{L}_{\text{outer}}(\Theta^*, \mathcal{D}_{\text{val}})$  denotes the validation set loss *after*  $\mathcal{L}_{\text{inner}}$  has been optimized, ( $\Theta^* = \operatorname{argmin} \mathcal{L}_{\text{inner}}$ ), when starting from the initial condition set by the outer optimization. The overall goal in Eq. 3 above is thus to set the initial condition of base algorithm  $\mathcal{L}_{\text{inner}}$  such that it achieves minimum loss on the validation set. When both losses are differentiable we can in principle solve Eq. 3 by taking gradient steps on  $\mathcal{L}_{\text{outer}}$  as shown in Algorithm 1. However, such exact meta-learning requires backpropagating through the path of the inner optimization, which is costly and inaccurate for a long computation graph.

**Algorithm 2** Meta-Update Initial Condition: SPG

---

**Function** UpdateIC( $\theta_0, \mathcal{D}_{\text{tr}}, \mathcal{D}_{\text{val}}, \mathcal{L}_{\text{inner}}, \mathcal{L}_{\text{outer}}$ ) :  
 $\tilde{\theta}_0 = \text{copy}(\theta_0)$   
**for**  $j = 1, 2, \dots, J$  **do** // Inner-level optimization  
 $\tilde{\theta}_j = \tilde{\theta}_{j-1} - \alpha \nabla_{\tilde{\theta}_{j-1}} \mathcal{L}_{\text{inner}}(\tilde{\theta}_{j-1}, (\mathcal{D}_{\text{tr}})_j)$   
**end for**  
 $\nabla_{\tilde{\theta}_0}^{\text{short}} = \theta_0 - \tilde{\theta}_J$  // Outer-level step  
 $\theta_0 = \theta_0 - \alpha \nabla_{\theta_0} \mathcal{L}_{\text{outer}}(\theta_0 - \nabla_{\tilde{\theta}_0}^{\text{short}}, \mathcal{D}_{\text{val}})$   
**Output:**  $\theta_0$

---

**3.2 Meta-Learning for Domain Adaptation**

**Overview** For meta domain adaptation, we would like to instantiate the initial condition learning idea summarised earlier in Eq. 3 in order to initialize popular domain adaptation algorithms such as [16, 43, 41] that can be represented as problems in the form of Eqs. 1 and 2, so as to to maximise the resulting performance in the target domain upon deployment. To this end we will introduce in the following section appropriate definitions of the inner and outer tasks, as well as a tractable optimization strategy.

**Multi-Source Domain Adaptation** Suppose we have an adequate algorithm to optimize for initial conditions as required in Eq. 3. How could we apply this idea to multi-source unsupervised domain adaptation setting, given that there is no target domain training data to take the role of  $\mathcal{D}_{\text{val}}$  in providing the metric for outer loop optimization of the initial condition? Our idea is that in the *multi-source* domain adaptation setting, we can split available source domains into disjoint meta-training and meta-testing domains  $\mathcal{D}_S = \mathcal{D}_S^{\text{mtr}} \cup \mathcal{D}_S^{\text{mte}}$ , where we actually have labels for both. Now we can let  $\mathcal{L}_{\text{inner}}$  be an unsupervised domain method [16, 43]  $\mathcal{L}_{\text{inner}} := \mathcal{L}_{\text{uda}}$ , and ask it to adapt from meta-train to the unlabelled meta-test domain. In the outer loop, we can then use the labels of the meta-test domain as a validation set to evaluate the adaptation performance via a supervised loss  $\mathcal{L}_{\text{outer}} := \mathcal{L}_{\text{sup}}$ , such as cross-entropy. Thus we aim to find an initial condition for our base domain adaptation method  $\mathcal{L}_{\text{uda}}$  that enables it to adapt effectively between source domains

$$\theta_0 = \underset{\theta_0}{\operatorname{argmin}} \sum_{\mathcal{D}_S^{\text{mtr}}, \mathcal{D}_S^{\text{mte}} \sim \mathcal{D}_S} \mathcal{L}_{\text{sup}}(\mathcal{L}_{\text{uda}}(\mathcal{D}_S^{\text{mtr}}, \overline{\mathcal{D}}_S^{\text{mte}}; \theta_0), \mathcal{D}_S^{\text{mte}}) \quad (4)$$

where we use  $\mathcal{L}(\cdot; \theta_0)$  to denote optimizing a loss from starting point  $\theta_0$ . This initial condition could be optimized by taking gradient descent steps on the outer supervised loss using UpdateIC from Algorithm 1. The resulting  $\theta_0$  is suited to adapting between all source domains hence should also be good for adapting to the target domain. Thus we would finally instantiate the same UDA algorithm using the learned initial condition, but this time between the full set of source domains, and the true unlabelled target domain  $\overline{\mathcal{D}}_T$ .

$$\theta = \underset{\theta}{\operatorname{argmin}} \mathcal{L}_{\text{uda}}(\mathcal{D}_S, \overline{\mathcal{D}}_T; \theta_0) \quad (5)$$

**An Online Solution** While conceptually simple, the problem with the direct approach above is that it requires completing domain-adaptive training multiple times in the inner optimization. Instead we propose to perform *online* meta-learning [53, 56] by alternating between steps of meta-optimization of Eq. 4 and steps on the final unsupervised domain adaptation problem in Eq. 5. That is, we iterate

$$\begin{aligned}\Theta &= \text{UpdateIC}(\Theta, (\mathcal{D}_S^{\text{mtr}}) \cup (\overline{\mathcal{D}}_S^{\text{mte}}), (\mathcal{D}_S^{\text{mte}}), \mathcal{L}_{\text{uda}}, \mathcal{L}_{\text{sup}}) \\ \Theta &= \Theta - \alpha \nabla_{\Theta} \mathcal{L}_{\text{uda}}(\Theta, (\mathcal{D}_S), (\overline{\mathcal{D}}_T))\end{aligned}\quad (6)$$

where  $(\mathcal{D})$  denotes minibatch sampling from the corresponding dataset, and we call  $\text{UpdateIC}(\cdot)$  with a small number of inner-loop optimizations such as  $J = 1$ .

Our method, summarised in Figure 1 and Algorithm 3, performs meta-learning online, by simultaneously solving the meta-objective and the target task. It translates to tuning the initial condition between taking optimization steps on the target DA task. This avoids the intractability and instability of back-propagating through the long computational graph in the exact approach that meta-optimizes  $\Theta_0$  to completion before doing DA. Online meta-learning is also potentially advantageous in practice due to improving optimization throughout training rather than only at the start – c.f. the vanilla exact method, where the impact of the initial condition on the final outcome is very indirect.

**Semi-Supervised Domain Adaptation** We next consider how to adapt the ideas introduced previously to the semi-supervised domain adaptation setting. In the MSDA setting above, we divided source domains into meta-train and meta-test, used unlabeled data from meta test to drive adaptation, and then used meta-test labels to validate the adaptation performance. In SSDA we do not have multiple source domains with which to use such a meta-train/meta-test split strategy, but we do have a small amount of labeled data in the target domain that we can use to validate adaptation performance and drive initial condition optimization. By analogy to Eq. 4, we can aim to find the initial condition for the unsupervised component  $\mathcal{L}_{\text{uda}}$  of an SSDA method. But now we can use the few labelled examples  $\mathcal{D}_T$  to validate the adaptation in the outer loop.

$$\Theta_0 = \underset{\Theta_0}{\text{argmin}} \sum \mathcal{L}_{\text{sup}}(\mathcal{L}_{\text{uda}}(\mathcal{D}_S, \overline{\mathcal{D}}_T; \Theta_0), \mathcal{D}_T) \quad (7)$$

The learned initial condition can then be used to instantiate the final semi-supervised domain adaptive training.

$$\Theta = \underset{\Theta}{\text{argmin}} \mathcal{L}_{\text{ssda}}(\Theta, \mathcal{D}_S, \overline{\mathcal{D}}_T, \mathcal{D}_T; \Theta_0) \quad (8)$$

**An Online Solution** The exact meta SSDA approach above suffers from the same limitations as exact MetaMSDA. So we again apply online meta-learning by iterating between meta-optimization of Eq. 7 and the final supervised domain adaptation problem of Eq. 9.

$$\begin{aligned}\Theta &= \text{UpdateIC}(\Theta, (\mathcal{D}_S) \cup (\overline{\mathcal{D}}_T), (\mathcal{D}_T), \mathcal{L}_{\text{uda}}, \mathcal{L}_{\text{sup}}) \\ \Theta &= \Theta - \alpha \nabla_{\Theta} \mathcal{L}_{\text{ssda}}(\Theta, \mathcal{D}_S, \overline{\mathcal{D}}_T, \mathcal{D}_T; \Theta)\end{aligned}\quad (9)$$

**Algorithm 3** Online Meta Learning: Multi-Source DA

---

**Input:** N source domains  $\mathcal{D}_S = [D_1, D_2, \dots, D_N]$  and unlabelled target domain  $\overline{\mathcal{D}}_T$ .  
**Initialise:** Model parameters  $\theta$ , learning rate  $\alpha$ , task loss  $\mathcal{L}_{\text{sup}}$ , UDA method  $\mathcal{L}_{\text{uda}}$ .  
**for**  $i = 1, 2, \dots, I$  **do**  
     $\theta = \text{UpdateIC}(\theta, \mathcal{D}_S^{\text{mtr}} \cup \overline{\mathcal{D}}_S^{\text{mte}}, \mathcal{D}_S^{\text{mte}}, \mathcal{L}_{\text{uda}}, \mathcal{L}_{\text{sup}})$   
    **for**  $s = 1, 2, \dots, S$  **do**  
         $\theta = \theta - \alpha \nabla_{\theta} \mathcal{L}_{\text{uda}}(\theta, (\mathcal{D}_S)_i, (\overline{\mathcal{D}}_T)_i)$       // Domain Adaptation Training  
    **end for**  
**end for**  
**Output:**  $\theta$

---

The final procedure is summarized in Algorithm 4.

### 3.3 Shortest Path optimization

**Meta-Learning Model Initialisation** As described so far, our meta-learning approach to domain adaptation relies on the ability meta-optimize initial conditions using gradient descent steps as described in Algorithm 1. Such steps evaluate a meta-gradient that depends on the parameter  $\theta^*$  output by the base domain adaptation algorithm

$$\theta_0 = \theta_0 - \alpha \overbrace{\nabla_{\theta} \mathcal{L}_{\text{sup}}(\theta^*, \mathcal{D}_{\text{val}})}^{\text{Meta Gradient}} \quad (10)$$

Evaluating the meta-gradient directly is impractical because: (i) The inner loop that runs the base domain adaptation algorithm may take multiple gradient descent iterations  $j = 1 \dots J$ . This will trigger a large chain of higher-order gradients  $\nabla_{\theta_0} \mathcal{L}_{\text{inner}}(\cdot), \dots, \nabla_{\theta_{J-1}} \mathcal{L}_{\text{inner}}(\cdot)$ . (ii) More fundamentally, several state of the art domain adaptation algorithms [43, 41] use multiple optimization steps when making updates on  $\mathcal{L}_{\text{inner}}$ . For example, to adversarially train the deep feature extractor and classifier modules of the model in  $\theta$ . Taking gradient steps on  $\mathcal{L}_{\text{outer}}(\theta^*)$  thus triggers higher-order gradients, even if one only takes a single step  $J = 1$  of domain adaptation optimization.

**Shortest Path optimization** To obtain the meta gradient in Eq. 10 efficiently, we use shortest-path gradient (SPG) [36]. Before optimising the inner loop, we copy parameters  $\theta_0$  as  $\tilde{\theta}_0$  and use  $\tilde{\theta}_0$  in the inner-level algorithm. Then, after finishing the inner loop we get the shortest-path gradient between  $\tilde{\theta}_J$  and  $\theta_0$ .

$$\nabla_{\theta_0}^{\text{short}} = \theta_0 - \tilde{\theta}_J \quad (11)$$

Each meta-gradient step (Eq. 10) is then approximated as

$$\theta_0 = \theta_0 - \alpha \nabla_{\theta_0} \mathcal{L}_{\text{sup}}(\theta_0 - \nabla_{\theta_0}^{\text{short}}, \mathcal{D}_{\text{val}}) \quad (12)$$

**Summary** We now have an efficient implementation of UpdateIC for updating



**Algorithm 4** Online Meta Learning: Semi-Supervised DA

---

**Input:**  $N$  source domains  $\mathcal{D}_S = [D_1, D_2, \dots, D_N]$ , labelled and unlabelled target domain data  $\mathcal{D}_T$  and  $\bar{\mathcal{D}}_T$ .  
**Initialise:** Model params  $\Theta$ , learning rate  $\alpha$ , task loss  $\mathcal{L}_{\text{sup}}$ , UDA method  $\mathcal{L}_{\text{uda}}$ .  
**for**  $i = 1, 2, \dots, I$  **do**  
     $\Theta = \text{UpdateIC}(\Theta, \mathcal{D}_S \cup \bar{\mathcal{D}}_T, \mathcal{D}_T, \mathcal{L}_{\text{uda}}, \mathcal{L}_{\text{sup}})$   
    **for**  $s = 1, 2, \dots, S$  **do**  
         $\Theta = \Theta - \alpha \nabla_{\Theta} (\mathcal{L}_{\text{sup}}(\Theta, (\mathcal{D}_T)_i) + \mathcal{L}_{\text{uda}}(\Theta, (\mathcal{D}_S)_i, (\bar{\mathcal{D}}_T)_i))$   
    // Domain Adaptation Training  
    **end for**  
**end for**  
**Output:**  $\Theta$

---

initial conditions as summarised in Algorithm 2. This shortest path approximation has the advantage of allowing efficient initial condition updates both for multiple iterations of inner loop optimization  $J > 1$ , as well as for inner loop domain adaptation algorithms that use multiple steps [43, 41]. We use this implementation for the MSDA and SSDA methods in Algorithms 3 and 4.

## 4 Experiments

**Datasets** We evaluate our method on several multi-source domain adaptation benchmarks including PACS [22], Office-Home [49] and DomainNet [38]; as well as on the semi-supervised setting of Office-Home and DomainNet.

**Base Domain Adaptation Algorithms and Ablation** Our Meta-DA framework is designed to complement existing base domain adaptation algorithms. We evaluate it in conjunction with Domain Adversarial Neural Networks (DANN, [16]) – as a representative classic approach to deep domain adaptation; as well as Maximum Classifier Discrepancy (MCD, [43]) and MinMax Entropy (MME, [41]) – as examples of state of the art multi-source and semi-supervised domain adaptation algorithms respectively. Our goal is to evaluate whether our Meta-DA framework can improve these base learners. We note that the MCD algorithm has two variants: (1) A multi-step variant that alternates between updating the classifiers and several steps of updating the feature extractor and (2) A one-step variant that uses a gradient reversal [16] layer so that classifier and feature extractor can be updated in a single gradient step. We evaluate both of these. **Sequential Meta-Learning:** As an ablation, we also consider an alternative fast meta-learning approach that performs all meta-updates at the start of learning, before doing DA; rather than performing meta-updates online with DA as in our proposed Meta-DA algorithms.

### 4.1 Multi-Source Domain Adaptation

**PACS: Dataset** PACS [22] was initially proposed for domain generalization and had been subsequently been re-purposed [34, 7] for multi-source domain

Method	C,P,S $\mapsto$ A	A,P,S $\mapsto$ C	A,C,S $\mapsto$ P	A,C,P $\mapsto$ S	Ave.
Source-only	77.85	74.86	95.73	67.74	79.05
DIAL [34]	87.30	85.50	97.00	66.80	84.15
DDiscovery [34]	<b>87.70</b>	<b>86.90</b>	97.00	69.60	85.30
JiGen [7]	84.88	81.07	<b>97.96</b>	<b>79.05</b>	85.74
DANN [16]	84.77	83.83	96.29	69.61	83.62
Meta-DANN (Ours)	87.30	84.90	96.89	73.22	85.58
MCD (n=4) [43]	86.32	84.51	<b>97.31</b>	71.01	84.79
Meta-MCD (n=4) (Ours)	<b>87.40</b>	<b>86.18</b>	97.13	<b>78.26</b>	<b>87.24</b>
MCD (os) [43]	85.99	82.89	97.24	74.49	85.15
Meta-MCD (os) (Ours)	86.67	84.94	96.23	77.70	<b>86.39</b>

Table 1: Multi-Source DA results on PACS. Bold: Best. Red: Second Best.

Method	C,P,R $\mapsto$ A	A,P,R $\mapsto$ C	A,C,R $\mapsto$ P	A,C,P $\mapsto$ R	Ave.
Source-only	67.04	56.04	80.74	82.86	71.67
DSBN [8]	-	-	-	83.00	-
M <sup>3</sup> SDA- $\beta$ [38]	67.20	58.58	79.05	81.18	71.50
DANN [16]	68.23	58.90	79.70	83.08	72.48
Meta-DANN (Ours)	<b>70.62</b>	59.13	80.24	82.79	73.20
MCD [43]	69.84	59.84	80.92	82.67	73.32
Meta-MCD (Ours)	70.21	<b>60.50</b>	<b>81.17</b>	<b>83.43</b>	<b>73.83</b>

Table 2: Multi-Source Domain adaptation on Office-Home.

adaptation. This dataset has four diverse domains: (A)rt painting, (C)artoon, (P)hoto and (S)ketch with seven object categories ‘dog’, ‘elephant’, ‘giraffe’, ‘guitar’, ‘house’, ‘horse’ and ‘person’ with 9991 images in total. **Setting** We follow the setting in [7] and perform leave-one-domain out evaluation, setting each domain as the adaptation target in turn. As per [7], we use the ImageNet pre-trained ResNet-18 as our feature extractor for fair comparison. We train with M-SGD (batch size=32, learning rate= $2 \times 10^{-3}$ , momentum=0.9, weight decay= $10^{-4}$ ). All the models are trained for 5k iterations before testing. **Results** From the results in Table 1, we can see that: (i) Several recent methods with published results on PACS achieve similar performance, with JiGen [7] performing best. We additionally evaluate DANN and MCD including one-step MCD (os) and multi-step MCD (n=4) variants, and find that one-step MCD performs similarly to JiGen. (ii) Applying our Meta-DA framework to DANN and MCD boosts all three base domain adaptation methods by 1.96%, 2.5% and 1.2% respectively. (iii) In particular, our Meta-MCD surpasses the previous state of the art performance set by JiGen. Together these results show the broad applicability and absolute efficacy of our method. Based on these results we focus on the better performing single-step MCD in the following evaluations.

**Office-Home: Dataset and Settings** Office-Home was initially proposed for the single-source domain adaptation, containing  $\approx 15,500$  images from four domains ‘artistic’, ‘clip art’, ‘product’ and ‘real-world’ with 65 different categories. We follow the setting in [8] and use ImageNet pretrained ResNet-50 as our backbone. We train all models with M-SGD (batch size=32, learning rate= $10^{-3}$ , mo-

Method	inf,pnt,qdr, rel,skt $\mapsto$ clp	pnt,qdr, rel,skt $\mapsto$ inf	clp,inf,qdr, rel,skt $\mapsto$ pnt	clp,inf,qdr, rel,skt $\mapsto$ qdr	clp,inf,qdr, qdr,skt $\mapsto$ rel	clp,inf,qdr, rel,qdr,rel $\mapsto$ skt	Ave.
Source-only	47.6 $\pm$ 0.52	13.0 $\pm$ 0.41	38.1 $\pm$ 0.45	13.3 $\pm$ 0.39	51.9 $\pm$ 0.85	33.7 $\pm$ 0.54	32.9 $\pm$ 0.54
Various Backbones [38]							
DAN [27]	45.4 $\pm$ 0.49	12.8 $\pm$ 0.86	36.2 $\pm$ 0.58	15.3 $\pm$ 0.37	48.6 $\pm$ 0.72	34.0 $\pm$ 0.54	32.1 $\pm$ 0.59
RTN [29]	44.2 $\pm$ 0.57	12.6 $\pm$ 0.73	35.3 $\pm$ 0.59	14.6 $\pm$ 0.76	48.4 $\pm$ 0.67	31.7 $\pm$ 0.73	31.1 $\pm$ 0.68
JAN [30]	40.9 $\pm$ 0.43	11.1 $\pm$ 0.61	35.4 $\pm$ 0.50	12.1 $\pm$ 0.67	45.8 $\pm$ 0.59	32.3 $\pm$ 0.63	29.6 $\pm$ 0.57
DANN [16]	45.5 $\pm$ 0.59	13.1 $\pm$ 0.72	37.0 $\pm$ 0.69	13.2 $\pm$ 0.77	48.9 $\pm$ 0.65	31.8 $\pm$ 0.62	32.6 $\pm$ 0.68
ADDA [47]	47.5 $\pm$ 0.76	11.4 $\pm$ 0.67	36.7 $\pm$ 0.53	14.7 $\pm$ 0.50	49.1 $\pm$ 0.82	33.5 $\pm$ 0.49	32.2 $\pm$ 0.63
SE [15]	24.7 $\pm$ 0.32	3.9 $\pm$ 0.47	12.7 $\pm$ 0.35	7.1 $\pm$ 0.46	22.8 $\pm$ 0.51	9.1 $\pm$ 0.49	16.1 $\pm$ 0.43
MCD [43]	54.3 $\pm$ 0.64	22.1 $\pm$ 0.70	45.7 $\pm$ 0.63	7.6 $\pm$ 0.49	58.4 $\pm$ 0.65	43.5 $\pm$ 0.57	38.5 $\pm$ 0.61
DCNTN [52]	48.6 $\pm$ 0.73	23.5 $\pm$ 0.59	48.8 $\pm$ 0.63	7.2 $\pm$ 0.46	53.5 $\pm$ 0.56	47.3 $\pm$ 0.47	38.2 $\pm$ 0.57
M <sup>3</sup> SDA- $\beta$ [38]	58.6 $\pm$ 0.53	26.0 $\pm$ 0.89	52.3 $\pm$ 0.55	6.3 $\pm$ 0.58	62.7 $\pm$ 0.51	49.5 $\pm$ 0.76	42.6 $\pm$ 0.64
ResNet-18							
Source-only	56.58 $\pm$ 0.16	18.97 $\pm$ 0.10	45.95 $\pm$ 0.16	11.52 $\pm$ 0.15	60.79 $\pm$ 0.17	43.70 $\pm$ 0.03	39.58 $\pm$ 0.09
DANN [16]	56.34 $\pm$ 0.12	18.66 $\pm$ 0.09	47.09 $\pm$ 0.08	12.27 $\pm$ 0.12	61.34 $\pm$ 0.07	45.26 $\pm$ 0.34	40.16 $\pm$ 0.12
Meta-DANN (Ours)	57.26 $\pm$ 0.17	<b>19.24<math>\pm</math>0.09</b>	47.29 $\pm$ 0.16	13.38 $\pm$ 0.15	61.21 $\pm$ 0.13	45.53 $\pm$ 0.17	40.65 $\pm$ 0.04
MCD [43]	57.64 $\pm$ 0.28	18.71 $\pm$ 0.10	<b>47.82<math>\pm</math>0.09</b>	12.64 $\pm$ 0.16	<b>61.69<math>\pm</math>0.10</b>	45.61 $\pm$ 0.01	40.69 $\pm$ 0.05
Meta-MCD (Ours)	<b>58.37<math>\pm</math>0.21</b>	19.09 $\pm$ 0.08	47.63 $\pm$ 0.12	<b>13.70<math>\pm</math>0.14</b>	61.30 $\pm$ 0.18	<b>45.90<math>\pm</math>0.18</b>	<b>41.00<math>\pm</math>0.05</b>
ResNet-34							
Source-only	61.50 $\pm$ 0.06	21.10 $\pm$ 0.07	49.13 $\pm$ 0.06	13.03 $\pm$ 0.18	64.14 $\pm$ 0.10	48.19 $\pm$ 0.12	42.85 $\pm$ 0.05
DANN [16]	60.95 $\pm$ 0.05	20.91 $\pm$ 0.11	50.35 $\pm$ 0.08	14.53 $\pm$ 0.06	64.73 $\pm$ 0.02	49.88 $\pm$ 0.27	43.56 $\pm$ 0.04
Meta-DANN (Ours)	61.39 $\pm$ 0.03	<b>21.53<math>\pm</math>0.14</b>	50.49 $\pm$ 0.29	15.31 $\pm$ 0.28	64.33 $\pm$ 0.09	49.87 $\pm$ 0.25	43.82 $\pm$ 0.07
MCD [43]	62.21 $\pm$ 0.12	20.49 $\pm$ 0.08	<b>50.87<math>\pm</math>0.10</b>	14.66 $\pm$ 0.30	<b>64.78<math>\pm</math>0.06</b>	50.10 $\pm$ 0.11	43.85 $\pm$ 0.05
Meta-MCD (Ours)	<b>62.81<math>\pm</math>0.22</b>	21.37 $\pm$ 0.07	50.53 $\pm$ 0.08	<b>15.47<math>\pm</math>0.22</b>	64.58 $\pm$ 0.16	<b>50.40<math>\pm</math>0.12</b>	<b>44.19<math>\pm</math>0.07</b>

Table 3: Multi-Source Domain adaptation on DomainNet dataset.

Method	R $\rightarrow$ C	P $\rightarrow$ C	P $\rightarrow$ A	A $\rightarrow$ C	C $\rightarrow$ A	Ave.
S+T	44.6	44.4	36.1	38.8	37.5	40.3
AlexNet						
DANN [16]	47.2	44.4	36.1	39.8	38.6	41.2
ADR [42]	45.0	38.9	36.3	40.0	37.3	39.5
CDAN [28]	41.8	35.8	32.0	34.5	27.9	34.4
ENT [17]	44.9	41.2	34.6	37.8	31.8	38.1
MME [41]	<b>51.2</b>	47.2	<b>40.7</b>	43.8	<b>44.7</b>	45.5
Meta-MME (Ours)	50.3	<b>48.3</b>	40.3	<b>44.5</b>	44.5	<b>45.6</b>
ResNet-34						
S+T	57.4	54.5	59.9	56.2	57.6	57.1
ENT [17]	62.8	61.8	65.4	62.1	65.8	63.6
MME [41]	64.9	63.8	65.0	63.0	66.6	64.7
Meta-MME (Ours)	<b>65.2</b>	<b>64.5</b>	<b>66.7</b>	<b>63.3</b>	<b>67.5</b>	<b>65.4</b>

Table 4: Semi-supervised domain adaptation: Office-Home.

mentum=0.9 and weight decay= $10^{-4}$ ) for  $3k$  iterations. **Results** From Table 2, we see that MCD achieves the best performance among the baselines. Applying our meta-learning framework improves both baselines by a small amount, and Meta-MCD achieves state-of-the-art performance on this benchmark.

**DomainNet: Dataset** DomainNet is a recently benchmark [38] for multi-source domain adaptation in object recognition. It is the largest-scale DA benchmark so far, with  $\approx 0.6$ m images across six domains and 345 categories. **Settings** We follow the official train/test split protocol [38]<sup>4</sup>. Various feature extraction backbones were used in the original paper [38], making it hard to compare results. We use ImageNet pre-trained ResNet-18 and ResNet-34 for our own implementations to facilitate direct comparison. We use M-SGD to train all the competitors (batch size=32, learning rate=0.001, momentum=0.9, weight de-

<sup>4</sup> Other settings such as optimizer, iterations and data augmentation are not clearly stated in [38], making it hard to replicate their results.

Method	R $\rightarrow$ C	R $\rightarrow$ P	P $\rightarrow$ C	C $\rightarrow$ S	S $\rightarrow$ P	R $\rightarrow$ S	P $\rightarrow$ R	Ave
S+T	47.1	45.0	44.9	36.4	38.4	33.3	58.7	43.4
DANN [16]	46.1	43.8	41.0	36.5	38.9	33.4	57.3	42.4
ADR [42]	46.2	44.4	43.6	36.4	38.9	32.4	57.3	42.7
CDAN [28]	46.8	45.0	42.3	29.5	33.7	31.3	58.7	41.0
ENT [17]	45.5	42.6	40.4	31.1	29.6	29.6	60.0	39.8
MME [41]	55.6	49.0	51.7	39.4	43.0	37.9	<b>60.7</b>	48.2
Meta-MME (Ours)	<b>56.4</b>	<b>50.2</b>	<b>51.9</b>	<b>39.6</b>	<b>43.7</b>	<b>38.7</b>	<b>60.7</b>	<b>48.8</b>
S+T	60.0	62.2	59.4	55.0	59.5	50.1	73.9	60.0
DANN [16]	59.8	62.8	59.6	55.4	59.9	54.9	72.2	60.7
ADR [42]	60.7	61.9	60.7	54.4	59.9	51.1	74.2	60.4
CDAN [28]	69.0	67.3	68.4	57.8	65.3	59.0	78.5	66.5
ENT [17]	71.0	69.2	71.1	60.0	62.1	61.1	78.6	67.6
MME [41]	72.2	69.7	71.7	61.8	66.8	61.9	78.5	68.9
Meta-MME (Ours)	<b>73.5</b>	<b>70.3</b>	<b>72.8</b>	<b>62.8</b>	<b>68.0</b>	<b>63.8</b>	<b>79.2</b>	<b>70.1</b>

Table 5: Semi-supervised DA on DomainNet.

Setting	Method	C,P,S $\rightarrow$ A	A,P,S $\rightarrow$ C	A,C,S $\rightarrow$ P	A,C,P $\rightarrow$ S	Ave.
DG	MetaReg [3]	83.7	77.2	95.5	70.3	81.7
	Epi-FCR [23]	82.1	77.0	93.9	73.0	81.5
	MASF [11]	80.3	77.2	95.0	71.7	81.0
DA	MCD [43]	86.3	84.5	97.3	71.0	84.8
	Meta-MCD	87.4	86.2	97.1	78.3	87.2

Table 6: Comparison between DG and DA methods on PACS.

cay=0.0001) for 10k iterations<sup>5</sup>. We re-train the model three times to generate standard deviations. **Results** From the results in Table 3, we can see that: (i) The top group of results from [38] show that the dataset is a much more challenging domain adaptation benchmark than previous ones. Most existing domain adaptation methods (typically tuned on small-scale benchmarks) fail to improve over the source-only baseline according to the results in [38]. (ii) The middle group of ResNet-18 results show that our MCD experiment achieves comparable results to those in [38]. (iii) Our Meta-MCD and Meta-DANN methods provide a small but consistent improvement over the corresponding MCD and DANN baselines for both ResNet-18 and ResNet-34 backbones. While the improvement margins are relatively small, this is a significant outcome as the results show that the base DA methods already struggle to make a large improvement over the source-only baseline when using ResNet-18/34; and also the multi-run standard deviation is small compared to the margins. (iv) Overall our Meta-MCD achieves state-of-the-art performance on the benchmark by a small margin.

## 4.2 Semi-Supervised Domain Adaptation

**Office-Home: Setting** We follow the setting in [41]. We focus on 3-shot learning in the target domain (three annotated examples only per category), and focus on the five most difficult source-target domain pairs. We use the ImageNet pretrained AlexNet and ResNet-34 as backbone models. We train all the

<sup>5</sup> We tried training with up to 50k, and found it did not lead to clear improvement. So, we train all models for 10k iterations to minimise cost.




Fig. 2: Vanilla DA vs seq. and online meta (T:MSDA, B:SSDA).




Fig. 3: Performance across weight space slices defined by a common initial condition  $\Theta_0$  and MCD and Meta-MCD solutions ( $\Theta_{\text{Meta-MCD}}$  and  $\Theta_{\text{MCD}}$  respectively). MSDA PACS benchmark with Sketch target.

models with M-SGD, with batch size 24 for labelled source and target domains and 48 for the unlabelled target as in [41], learning rate is  $10^{-2}$  and  $10^{-3}$  for the fully-connected and the rest trainable layers. We also use horizontal-flipping and random-cropping data augmentation for training images. **Results** From the results in Table 4, we can see that our Meta-MME does not impact performance on AlexNet. However, for a modern ResNet-34 architecture, Meta-MME provides a visible  $\sim 0.8\%$  accuracy gain over the MME baseline, which results in the state-of-the-art performance of SSDA on this benchmark.

**DomainNet: Settings** We evaluate DomainNet for 1-1 few-shot domain adaptation as in [41]. We evaluate both AlexNet and modern ResNet-34 backbones, and apply our meta-learning method on MME. As per [41], we train our models using M-SGD where the initial learning rate is 0.01 for the fully-connected layers and 0.001 for the rest of trainable layers. During the training we use the annealing strategy in [16] to decay the learning rate, and use the same batch size as [41].

**Results** From the results in Table 5, we can see our Meta-MME improves on the accuracy of the base MME algorithm in all pairwise transfer choices, and also for both backbones. These results show the consistent effectiveness of our method, as well as improving state-of-the-art for DomainNet SSDA.

### 4.3 Further Analysis

**Discussion** Our final online algorithm can be understood as performing DA with periodic meta-updates that adjust parameters to optimize the impact of the following DA steps. From the perspective of any given DA step, the role of the preceding meta-update is to tune its initial condition.

**Non-Meta vs Sequential vs Online Meta** This work is the first to propose meta-learning to improve domain adaptation, and in particular to contribute an efficient and effective online meta-learning algorithm for initial condition training. Exact meta learning is intractable to compare. However, this section we compare our online meta update with the alternative sequential approximation, and non-meta alternatives for both MSDA and SSDA using A,C,P $\rightarrow$ S and R $\rightarrow$ C as examples. For fair comparison, we control the number of meta-updates (UpdateIC) and base DA updates available to both sequential and online

meta-learning methods to the same amount. Figure 2 shows that: (1) Sequential meta-learning method already improves the performance on the target domain comparing to vanilla domain adaptation, which confirms the potential for improvement by refining model initialization. (2) The sequential strategy has a slight advantage early in DA training, which makes sense, as all meta-updates occur in advance. But overall our online method that interleaves meta-updates and DA updates leads to higher test accuracy.

**Computational Cost** Our Meta-DA imposes only a small computational overhead over the base DA algorithm. For example, comparing Meta-MCD and MCD on ResNet-34 DomainNet, the time per iteration is 2.96s vs 2.49s respectively<sup>6</sup>.

**Weight-Space Illustration** To investigate our method’s mechanism, we train MCD and Meta-MCD from a common initial condition on MSDA PACS when ‘Sketch’ is the target domain. We use the initial condition  $\Theta_0$  and two different solutions ( $\Theta_{\text{Meta-MCD}}$  and  $\Theta_{\text{MCD}}$ ) to define a plane in weight-space and colour it according to the performance at each point. We can see from Figure 3(a) that Meta-MCD finds a solution with greater test accuracy. Figures 3(b) and (c) break down the training loss components. We can see that, in this slice, both methods managed to minimize MCD’s adaptation (classifier discrepancy) loss  $\mathcal{L}_a$  adequately, but MCD failed to minimize the supervised loss as well as Meta-MCD ( $\Theta_{\text{Meta-MCD}}$  is closer to the minima than  $\Theta_{\text{MCD}}$ ). Note that both methods were trained to convergence in generating these solutions. This suggests that Meta-MCD’s meta-optimization step using meta-train/meta-test splits materially benefits the optimization dynamics of the downstream MSDA task.

**Model Agnostic** We emphasize that, although we focused on DANN, MCD and MME, our MetaDA framework can apply to any base DA algorithm. Supplementary C shows some results for JiGen and M<sup>3</sup>SDA algorithms.

**Comparison between DA and DG methods** As a highly related topical problem to domain adaptation, domain generalization assumes no access to the target domain data during the training. DA and DG methods are rarely directly compared. Now we compare our Meta-MCD and MCD with some state of the art DG methods on PACS as shown in Table 6. From the results, we can see that generally DA methods outperform the DG methods with a noticeable margin, which is expected as DA methods ‘see’ the target domain data at training.

## 5 Conclusion

We proposed a meta-learning pipeline to improve domain adaptation by initial condition optimization. Our online shortest-path solution is efficient and effective, and provides a consistent boost to several domain adaptation algorithms, improving state of the art in both multi-source and semi-supervised settings. Our approach is agnostic to the base adaptation method, and can potentially be used to improve many DA algorithms that fit a very general template. In future we aim to meta-learn other DA hyper-parameters beyond initial conditions.

<sup>6</sup> Using GeForce RTX 2080 GPU. Xeon Gold 6130 CPU @ 2.10GHz.

## References

1. Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M.W., Pfau, D., Schaul, T., Shillingford, B., De Freitas, N.: Learning to learn by gradient descent by gradient descent. In: *NeurIPS* (2016)
2. Balaji, Y., Chellappa, R., Feizi, S.: Normalized wasserstein distance for mixture distributions with applications in adversarial learning and domain adaptation. In: *ICCV* (2019)
3. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: *NeurIPS* (2018)
4. Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. *Machine learning* (2010)
5. Ben-David, S., Blitzer, J., Crammer, K., Pereira, F.: Analysis of representations for domain adaptation. In: *NeurIPS* (2006)
6. Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., Erhan, D.: Domain separation networks. In: *NeurIPS* (2016)
7. Carlucci, F.M., D’Innocente, A., Bucci, S., Caputo, B., Tommasi, T.: Domain generalization by solving jigsaw puzzles. In: *CVPR* (2019)
8. Chang, W.G., You, T., Seo, S., Kwak, S., Han, B.: Domain-specific batch normalization for unsupervised domain adaptation. In: *CVPR* (2019)
9. Daumé, H.: Frustratingly easy domain adaptation. In: *ACL* (2007)
10. Donahue, J., Hoffman, J., Rodner, E., Saenko, K., Darrell, T.: Semi-supervised domain adaptation with instance constraints. In: *CVPR* (2013)
11. Dou, Q., Castro, D.C., Kamnitsas, K., Glocker, B.: Domain generalization via model-agnostic learning of semantic features. In: *NeurIPS* (2019)
12. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: *ICML* (2017)
13. Finn, C., Rajeswaran, A., Kakade, S.M., Levine, S.: Online meta-learning. In: *ICML* (2019)
14. Franceschi, L., Frasconi, P., Salzo, S., Grazzi, R., Pontil, M.: Bilevel programming for hyperparameter optimization and meta-learning. In: *ICML* (2018)
15. French, G., Mackiewicz, M., Fisher, M.: Self-ensembling for visual domain adaptation. In: *ICLR* (2018)
16. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *JMLR* (2016)
17. Grandvalet, Y., Bengio, Y.: Semi-supervised learning by entropy minimization. In: *NeurIPS* (2005)
18. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A.A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: *ICML* (2018)
19. Kim, M., Sahu, P., Gholami, B., Pavlovic, V.: Unsupervised visual domain adaptation: A deep max-margin gaussian process approach. In: *CVPR* (2019)
20. Lee, C.Y., Batra, T., Baig, M.H., Ulbricht, D.: Sliced wasserstein discrepancy for unsupervised domain adaptation. In: *CVPR* (2019)
21. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.: Learning to generalize: Meta-learning for domain generalization. In: *AAAI* (2018)
22. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: *ICCV* (2017)

23. Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
24. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-sgd: Learning to learn quickly for few-shot learning. In: arXiv (2017)
25. Liu, H., Simonyan, K., Yang, Y.: Darts: Differentiable architecture search. In: ICLR (2019)
26. Liu, M.Y., Tuzel, O.: Coupled generative adversarial networks. In: NeurIPS (2016)
27. Long, M., Cao, Y., Wang, J., Jordan, M.I.: Learning transferable features with deep adaptation networks. ICML (2015)
28. Long, M., Cao, Z., Wang, J., Jordan, M.I.: Conditional adversarial domain adaptation. In: NeurIPS (2018)
29. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Unsupervised domain adaptation with residual transfer networks. In: NeurIPS (2016)
30. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Deep transfer learning with joint adaptation networks. In: ICML (2017)
31. Luo, Y., Zheng, L., Guan, T., Yu, J., Yang, Y.: Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In: CVPR (2019)
32. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov), 2579–2605 (2008)
33. Maclaurin, D., Duvenaud, D., Adams, R.: Gradient-based hyperparameter optimization through reversible learning. In: ICML (2015)
34. Mancini, M., Porzi, L., Rota Bulò, S., Caputo, B., Ricci, E.: Boosting domain adaptation by discovering latent domains. In: CVPR (2018)
35. Mansour, Y., Mohri, M., Rostamizadeh, A.: Domain adaptation with multiple sources. In: NeurIPS (2009)
36. Nichol, A., Achiam, J., Schulman, J.: On first-order meta-learning algorithms. arXiv (2018)
37. Parisotto, E., Ghosh, S., Yalamanchi, S.B., Chinnabireddy, V., Wu, Y., Salakhutdinov, R.: Concurrent meta reinforcement learning. arXiv (2019)
38. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: CVPR (2019)
39. Rajeswaran, A., Finn, C., Kakade, S., Levine, S.: Meta-learning with implicit gradients. In: NeurIPS (2019)
40. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: ICLR (2016)
41. Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised domain adaptation via minimax entropy. In: ICCV (2019)
42. Saito, K., Ushiku, Y., Harada, T., Saenko, K.: Adversarial dropout regularization. In: ICLR (2018)
43. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: CVPR (2018)
44. Schmidhuber, J.: Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation* (1992)
45. Schmidhuber, J., Zhao, J., Wiering, M.: Shifting inductive bias with success-story algorithm, adaptive levin search, and incremental self-improvement. *Machine Learning* (1997)
46. Thrun, S., Pratt, L. (eds.): *Learning to Learn*. Kluwer Academic Publishers (1998)
47. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: CVPR (2017)



48. Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. arXiv (2014)
49. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: (IEEE) Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
50. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: NeurIPS (2016)
51. Xu, P., Gurram, P., Whipps, G., Chellappa, R.: Wasserstein distance based domain adaptation for object detection. arXiv (2019)
52. Xu, R., Chen, Z., Zuo, W., Yan, J., Lin, L.: Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In: CVPR (2018)
53. Xu, Z., van Hasselt, H.P., Silver, D.: Meta-gradient reinforcement learning. In: NeurIPS (2018)
54. Yao, T., Pan, Y., Ngo, C.W., Li, H., Mei, T.: Semi-supervised domain adaptation with subspace learning for visual recognition. In: CVPR (2015)
55. Zhao, H., Zhang, S., Wu, G., Moura, J.M., Costeira, J.P., Gordon, G.J.: Adversarial multiple source domain adaptation. In: NeurIPS (2018)
56. Zheng, Z., Oh, J., Singh, S.: On learning intrinsic rewards for policy gradient methods. In: NeurIPS (2018)

## A Short-Path Gradient Descent

Optimizing Eq. 3 naively by Algorithm 1 would be costly and ineffective. It is costly because in the case of domain adaptation (unlike for example, few-shot learning [12], the inner loop requires many iterations). So back-propagating through the whole optimization path to update the initial  $\Theta$  in the outer loop will produce multiple high-order gradients. For example, if the inner loop applies  $j$  iterations, we will have

$$\begin{aligned}\Theta^{(1)} &= \Theta - \alpha \nabla_{\Theta^{(0)}} \mathcal{L}_{\text{uda}}(\cdot) \\ &\dots \\ \Theta^{(j)} &= \Theta^{(j-1)} - \alpha \nabla_{\Theta^{(j-1)}} \mathcal{L}_{\text{uda}}(\cdot)\end{aligned}\tag{13}$$

then the outer loop will update the initial condition as

$$\Theta^* = \Theta - \alpha \overbrace{\nabla_{\Theta} \mathcal{L}_{\text{sup}}(\Theta^{(j)}, \mathcal{D}_{\text{val}})}^{\text{Meta Gradient}}\tag{14}$$

where higher-order gradient will be required for all items  $\nabla_{\Theta^{(0)}} \mathcal{L}_{\text{uda}}(\cdot), \dots, \nabla_{\Theta^{(j-1)}} \mathcal{L}_{\text{uda}}(\cdot)$  in the update of Eq. 14.

One intuitive way of eliminating higher-order gradients for computing Eq. 14 is making  $\nabla_{\Theta^{(0)}} \mathcal{L}_{\text{uda}}(\cdot), \dots, \nabla_{\Theta^{(j-1)}} \mathcal{L}_{\text{uda}}(\cdot)$  constant during the optimization. Then, Eq. 14 is equivalent to

$$\Theta^* = \Theta - \alpha \overbrace{\nabla_{\Theta^{(j)}} \mathcal{L}_{\text{sup}}(\Theta^{(j)}, \mathcal{D}_{\text{val}})}^{\text{First-order Meta Gradient}}\tag{15}$$

However, in order to compute Eq. 15, one still needs to store the optimization path of Eq. 13 in memory and back-propagate through it to optimize  $\Theta$ , which requires high computational load. Therefore, we propose a practical solution an iterative meta-learning algorithm to iteratively optimize the model parameters during training.

**Shortest Path Optimization** To obtain the meta gradient in Eq. 15 in a more efficient way, we propose a more scalable and efficient meta-learning method using shortest-path gradient (S-P.G.) [36]. Before the optimization of Eq. 13, we copy the parameters  $\Theta$  as  $\tilde{\Theta}^{(0)}$  and use  $\tilde{\Theta}^{(0)}$  in the inner-level algorithm.

$$\tilde{\Theta}^{(j)} = \begin{cases} \tilde{\Theta}^{(0)} - \alpha \nabla_{\Theta^{(0)}} \mathcal{L}_{\text{uda}}(\tilde{\Theta}^{(0)}, \mathcal{D}_{\text{tr}}), \\ \dots \\ \tilde{\Theta}^{(j-1)} - \alpha \nabla_{\Theta^{(0)}} \mathcal{L}_{\text{uda}}(\tilde{\Theta}^{(j-1)}, \mathcal{D}_{\text{tr}}) \end{cases}\tag{16}$$

then, after finishing the optimization in Eq. 16, we can get the shortest-path gradient between two items  $\tilde{\Theta}_i^{(j)}$  and  $\Theta_i$ .

$$\nabla_{\Theta}^{\text{short}} = \Theta - \tilde{\Theta}^{(j)}\tag{17}$$

Different from Eq. 15, we use this shortest-path gradient  $\nabla_{\Theta}^{\text{short}}$  and initial parameter  $\Theta$  to compute  $\mathcal{L}_{\text{sup}}(\cdot)$  as,

$$\mathcal{L}_{\text{sup}}(\Theta_i - \nabla_{\Theta_i}^{\text{short}}, \mathcal{D}_{\text{val}}) \quad (18)$$

Then, one-step meta update of Eq. 18 will be,

$$\begin{aligned} \Theta_i^* &= \Theta_i - \alpha \nabla_{\Theta_i} \mathcal{L}_{\text{sup}}(\Theta_i - \nabla_{\Theta_i}^{\text{short}}, \mathcal{D}_{\text{val}}) \\ &= \Theta_i - \alpha \nabla_{\Theta_i - \nabla_{\Theta_i}^{\text{short}}} \mathcal{L}_{\text{sup}}(\Theta_i - \nabla_{\Theta_i}^{\text{short}}, \mathcal{D}_{\text{val}}) \\ &= \Theta_i - \alpha \nabla_{\tilde{\Theta}_i^{(j)}} \mathcal{L}_{\text{sup}}(\tilde{\Theta}_i^{(j)}, \mathcal{D}_{\text{val}}) \end{aligned} \quad (19)$$

**Effectiveness:** We can see that one update of Eq. 19 corresponds to that of Eq. 15, which proves that using shortest-path optimization has the equivalent effectiveness to the first-order meta optimization. **Scalability/Efficiency:** The computation memory of the first-order meta-learning increases linearly with the inner-loop update steps, which is constrained by the total GPU memory. However, for the shortest-path optimization, storing the optimization graph is no longer necessary, which makes it scalable and efficient. We also experimentally evaluate that one step shortest-path optimization is 7x faster than one-step first-order meta optimization in our setting. The overall algorithm flow is shown in Algorithm 2.

## B Additional Illustrative Schematics

To better explain the contrast between our online meta-learning domain adaptation approach with the sequential meta-learning approach, we add a schematic illustration in Figure 4. The main difference between sequential and online meta-learning approaches is how do we distribute the meta and DA updates. Sequential meta-learning approach performs meta updates and DA updates sequentially. And online meta-learning conducts the alternative meta and DA updates throughout the whole training procedure.

## C Additional Experiments

**Visualization of the Learned Features** We visualize the learned features of MCD and Meta-MCD on PACS when sketch is the target domain as shown in Figure 5. We can see that both MCD and Meta-MCD can learn discriminative features. However, the features learned by Meta-MCD is more separable than vanilla MCD. This explains why our Meta-MCD performs better than the vanilla MCD method.

**Effect of varying  $S$**  Our online meta-learning method has iteration hyperparameters  $S$  and  $J$ . We fix  $J = 1$  throughout, and analyze the effect of varying  $S$  here using the DomainNet MSDA experiment with ResNet-18. The result in Table 7 shows that MetaDA is rather insensitive to this hyperparameter.

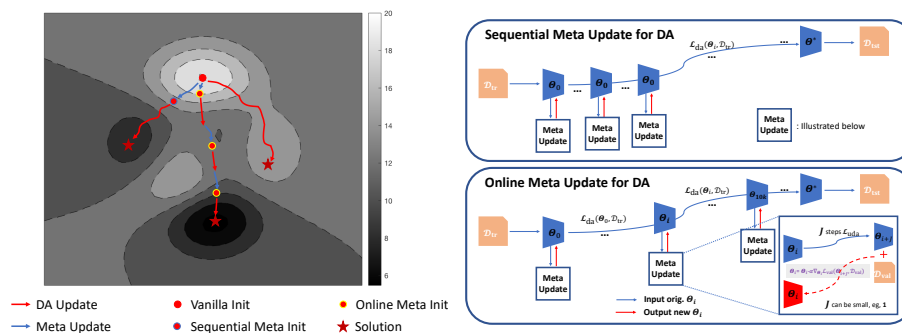


Fig. 4: Illustrative schematics of sequential and online meta domain adaptation. Left: Optimization paths of different approaches on domain adaptation loss (shading). (Solid line) Vanilla gradient descent on a DA objective from a fixed start point. (Multi-segment line) Online meta-learning iterates meta and gradient descent updates. (Two segment line) Sequential meta-learning provides an alternative approximation: update initial condition, then perform gradient descent. Right: (Top) Sequential meta-learning performs meta updates and DA updates sequentially. (Bottom) Online meta-learning alternates between meta-optimization and domain adaptation.

Method	Meta-MCD ( $S=3$ )	Meta-MCD ( $S=5$ )	Meta-MCD ( $S=10$ )
DomainNet (ave.)	41.02	40.98	40.93

Table 7: MetaDA is insensitive to the update ratio hyperparameter  $S$  – Results for MSDA ResNet-18 performance on DomainNet.

**Varying the Number of Source Domains in MSDA** For multi-source DA, the performance of both Meta-DA and the baselines is expected to drop with fewer sources (same for SSDA if fewer labeled target domain points). To disentangle the impact of the number of sources for baseline vs Meta-DA we compare MSDA by Meta-MCD on PACS with 2 vs 3 sources. The results for Meta-MCD vs vanilla MCD are 82.30% vs 80.07% (two source, gap 2.23%) and 87.24% vs 84.79% (three source, gap 2.45%). Meta-DA margin is similar with reduction of domains. Most difference is accounted for by the impact on the base DA algorithm.

**Other base DA methods** Besides the base DA methods evaluated in the main paper (DANN, MCD and MME), our method is applicable to any base domain adaptation method. We use the published code of JiGen<sup>7</sup> and M<sup>3</sup>SDA<sup>8</sup>, and further apply our Meta-DA on the existing code. The results are shown in Table 8 and 9. From the results, we can see that our Meta-JiGen and Meta-M<sup>3</sup>SDA- $\beta$  improves over the base methods by 3.42% and 1.2% accuracy respectively,

<sup>7</sup> <https://github.com/fmcarlucci/JigenDG>

<sup>8</sup> <https://github.com/VisionLearningGroup/VisionLearningGroup.github.io>




Fig. 5: t-SNE [32] visualization of learned MCD (left) and Meta-MCD (right) features on PACS (sketch as target domain). Different colors indicate different categories.

Method	C,P,S→A	A,P,S→C	A,C,S→P	A,C,P→S	Ave.
JiGen [7]	84.88	81.07	97.96	79.05	85.74
JiGen*	81.54	85.88	97.25	68.21	83.22
Meta-JiGen	85.21	86.13	97.31	77.91	86.64 (+3.42)

Table 8: Test accuracy on PACS. \* our run.

which confirms our Meta-DA’s generality. The reason we excluded these from the main results is that: (i) Re-running JiGen’s published code on our compute environment failed to replicate their published numbers. (ii)  $M^3SDA$  as a base algorithm is very slow to run comprehensive experiments on. Nevertheless, these results provide further evidence that Meta-DA can be a useful module going forward to plug in and improve future new base DA methods as well as those evaluated here.

**Initialization Dependence of Domain Adaptation** One may not think of domain adaptation as being sensitive to initial condition, but given the lack of target domain supervision to guide learning, different initialization can lead to a significant difference in accuracy. To illustrate this we re-ran MCD-based DA on PACS with sketch target using different initializations. From the results in Tab 10, we can see that both different classic initialization heuristics, and simple perturbation of a given initial condition with noise can lead to significant differences in final performance. This confirms that studying methods for tuning initialization provide a valid research direction for advancing DA performance.

Method	mt,up,sv,sv	mm,up,sv,sv	mt,mm,sv,sv	mt,mm,up,sv	mt,mm,up,sv	Ave.
	→mm	→mt	→up	→sv	→sy	
$M^3SDA-\beta$ [38]	72.82	98.43	96.14	81.32	89.58	87.65
Meta- $M^3SDA-\beta$	71.73	98.79	97.80	84.81	91.12	88.85 (+1.2)

Table 9: Test accuracy on Digit-Five.

Classifier Init	Kaiming $\mathcal{U}$	Xavier $\mathcal{U}$	Kaiming $\mathcal{N}$	Xavier $\mathcal{N}$
	74.49	73.02	64.27	73.66
Feat. Extr. Init	No perturb + $\epsilon \in \mathcal{N}(0, 0.01)$ + $\epsilon \in \mathcal{N}(0, 0.02)$ + $\epsilon \in \mathcal{N}(0, 0.03)$			
	74.49	71.85	59.99	52.18

Table 10: Test accuracy of MCD on PACS (sketch) with different initialization.