

# MetaAlign: Coordinating Domain Alignment and Classification for Unsupervised Domain Adaptation

Guoqiang Wei<sup>1\*</sup>   Cuiling Lan<sup>2†</sup>   Wenjun Zeng<sup>2</sup>   Zhibo Chen<sup>1†</sup>

<sup>1</sup> University of Science and Technology of China   <sup>2</sup> Microsoft Research Asia

wgq7441@mail.ustc.edu.cn   {culan,wezeng}@microsoft.com   chenzhibo@ustc.edu.cn

## Abstract

For unsupervised domain adaptation (UDA), to alleviate the effect of domain shift, many approaches align the source and target domains in the feature space by adversarial learning or by explicitly aligning their statistics. However, the optimization objective of such domain alignment is generally not coordinated with that of the object classification task itself such that their descent directions for optimization may be inconsistent. This will reduce the effectiveness of domain alignment in improving the performance of UDA. In this paper, we aim to study and alleviate the optimization inconsistency problem between the domain alignment and classification tasks. We address this by proposing an effective meta-optimization based strategy dubbed MetaAlign, where we treat the domain alignment objective and the classification objective as the meta-train and meta-test tasks in a meta-learning scheme. MetaAlign encourages both tasks to be optimized in a coordinated way, which maximizes the inner product of the gradients of the two tasks during training. Experimental results demonstrate the effectiveness of our proposed method on top of various alignment-based baseline approaches, for tasks of object classification and object detection. MetaAlign helps achieve the state-of-the-art performance.

## 1. Introduction

With the advance of deep convolutional neural networks (CNN), computer vision tasks such as image classification and object detection have gained significant improvement [28, 48, 20]. In general, the trained models perform well on the testing dataset of which the distribution bears a resemblance to that of training dataset. However, in many practical scenarios, directly applying such trained models to a new domain usually suffers from significant performance degradation. There exist differences in data characteris-

tics/distributions between the training and testing domains, which are known as *domain shift* [60, 56]. This makes it hard to directly transfer knowledge learned from source to target. Annotation on the samples of the target domain can alleviate this problem but is expensive and time-consuming. Without the requirement of annotation on target samples, unsupervised domain adaption (UDA) attracts a lot of attention, which allows us to learn to adapt the model trained on source to target by exploiting the unlabeled target samples.

There has been a large spectrum of UDA methods developed in the literature. The major line among them attempts to align the distributions of source and target domains by learning domain-invariant representations, through directly minimizing the discrepancy between feature distributions of two domains [62, 36, 57, 56, 16, 71] or adversarially learning to enforce the feature representations to be indistinguishable by a domain discriminator [14, 15, 43, 37, 67, 5, 9, 41]. The former category of methods usually align domain distributions by employing explicit distribution similarity metrics, *e.g.*, momentum distance [44, 69], or the second-order correlation [56, 57, 45], between the source and target domains. The latter one borrows ideas from Generative Adversarial Networks [18] and use adversarial training to learn aligned feature representations. However, these alignment constraints/strategies are actually not designed specially for the object classification task. There is a lack of efficient coordination between the optimizations of these two tasks. During training, the optimization procedure of alignment may be inconsistent with that of the object classification task itself, which could hurt the learning of discriminative object features for classification and thus results in inferior object classification performance.

In this work, we aim to address this pervasive problem/challenge faced by alignment-based unsupervised domain adaptation methods, *i.e.*, the optimization inconsistency between the domain alignment task and the classification task itself. We propose a meta-learning [53, 58] based method dubbed MetaAlign to mitigate such inconsistency. Particularly, as illustrated in Fig. 1 (b), we treat the domain

\*This work was done when Guoqiang was an intern at MSRA.

†Corresponding Author.

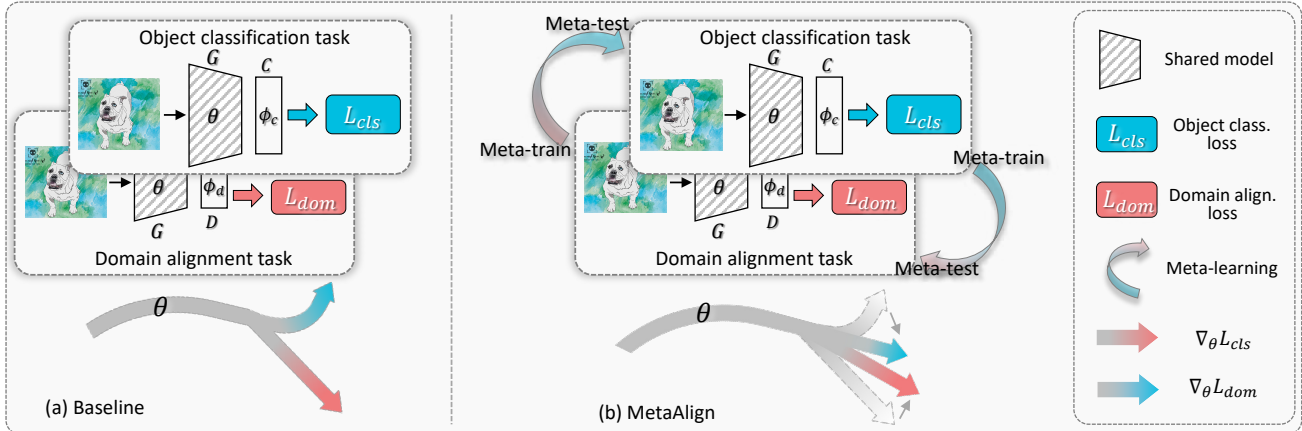


Figure 1: Illustration of our MetaAlign strategy which aims to encourage the optimization consistency between the domain alignment task and the object classification task for efficient UDA. (a) Previous approaches directly combine the optimization objective functions of the two tasks together (*i.e.*,  $\mathcal{L}_{dom} + \mathcal{L}_{cls}$ ), where the descent directions for optimizing the shared network parameters  $\theta$  from the two tasks may be inconsistent. (b) In contrast, we treat one of these two tasks as meta-train task and the other as meta-test task. We leverage this meta-optimization based strategy to enforce the consistency between their optimization gradients w.r.t.  $\theta$ . MetaAlign is generic and applicable to various domain alignment based UDAs.

alignment objective and the classification objective as two tasks in a meta-learning scheme, where we take one task as meta-train task to optimize the network and meanwhile we validate the optimization result on the other task (*i.e.*, meta-test task) for the same set of samples, during training. Meta-optimization across these two tasks encourages both to be optimized in a coordinated way. The theoretical analysis reveals that MetaAlign achieves this optimization coordination by maximizing the inner product of the gradients of the two tasks during training.

We summarize our contributions as follows:

- We pinpoint the problem/challenge in existing alignment-based UDA methods: the optimization inconsistency between the domain alignment task and the classification task itself. To address the problem, we propose a meta-optimization based strategy named MetaAlign to mitigate the inconsistency.
- The proposed MetaAlign strategy is generic and can be applied to *various domain alignment based UDA methods for object classification and detection* to enforce domain alignment while preserving the discrimination power of features for the recognition task.

We validate the effectiveness of MetaAlign on the image classification (unsupervised domain adaptation and domain generalization) and object detection (unsupervised domain adaptation) tasks. For image classification, we implement MetaAlign on top of various domain alignment based UDA baselines. Extensive experimental results demonstrate the effectiveness and applicability of MetaAlign and we achieve the state-of-the-art performance.

## 2. Related Work

**Unsupervised Domain Adaptation.** Unsupervised Domain Adaptation (UDA) aims to transfer the knowledge from a labeled source domain to an unlabeled target domain. Abundant UDA works focus on object classification or use it for their investigations. The mainstream approaches tend to address UDA by learning domain-invariant representation, to which our proposed method belongs. These approaches can be categorized into two categories.

One category explicitly reduces the domain discrepancy measured by some distribution discrepancy metrics. [62, 36, 39, 68] measure the domain similarity in terms of Maximum Mean Discrepancy (MMD) [3], while [56, 57, 44] introduce metrics based on second- or higher-order statistics. Another popular line learns domain-invariant representation using adversarial training. It has been widely studied [9, 5, 37, 61, 52, 6, 65, 35, 51, 40] since the seminal work DANN [14, 15]. In general, a domain discriminator is trained to distinguish the source domain from the target domain, meanwhile a feature extractor is trained to fool the discriminator to arrive at aligned features. SymNets[71] designs symmetric object classifiers which also play a role of domain discriminator. CDAN [37] conditions the adversarial model on the discriminative information conveyed in the classifier predictions. MCD [51] and STAR [40] build an adversarial framework to reduce the domain gap measured by the collision of two reduplicated object classifiers. GVB [9] balances adversarial training via constructing bridge layers on both the generator and discriminator.

These approaches all directly optimize domain alignment and classification tasks, while ignoring the optimization inconsistency between these two objectives. We pro-

pose a meta-optimization based strategy to mitigate the inconsistency for better UDA.

**UDA for Object Detection.** Learning domain adaptive deep object detector was first studied in DA-Faster [7], where they perform image-level and instance-level alignments via adversarial learning. SW-DA [50] reduces domain gap via aligning global-level and local-level features. Zhu *et al.* [73] propose to align domains at regions clustered by K-means. EPM [24] proposes the center-aware alignment based on the center map generated by an anchor-free detector [59]. Other variants exploit style-transfer [26], progressive alignment [25], and hierarchical alignment [74].

We also validate the effectiveness of our proposed MetaAlign on UDA for object detection task.

**Neural Network Meta Learning.** Meta-learning [53, 58] (a.k.a. learning to learn) has a long standing history. Recently, it has been widely applied to the optimization of deep neural networks [1, 34] and few-shot classification [27, 64]. Model-Agnostic Meta-Learning (MAML) [13] is proposed for few-shot learning and reinforcement learning, which aims to find good parameters initialization for fast adaptation to new tasks. [31, 46, 2] introduce meta-learning to Domain Generalization (DG) to synthesize the source-target domain shift during training. Li *et al.* [29] adopt MAML to provide better initialization condition for Multi-Source Domain Adaptation.

In this work, we pinpoint the underlying optimization inconsistency of the domain alignment objective and classification objective used for UDA. We are the first to mitigate it via a meta-optimization based strategy by treating the two objectives as meta-train and meta-test tasks respectively.

**Domain Generalization.** In contrast to UDA, Domain Generalization (DG) is applied in a more challenging scenario where target domain is inaccessible during training [42]. One category of methods for DG attempts to learn domain-invariant features [17, 33, 41], which borrows ideas from UDA. Li *et al.* [33] incorporate MMD as a constraint into the training of an adversarial autoencoder. Ghifary *et al.* [17] designs multi-task domain-specific decoders to help the training of the domain-invariant encoder. Matsuura *et al.* [41] use adversarial training to learn features invariant among predicted latent domains. Other categories exploit data augmentation [55, 66, 47], meta-learning [31, 2, 10], and auxiliary tasks [32].

Our MetaAlign can be applied to the first category of approaches for addressing the optimization inconsistency between domain alignment and classification.

### 3. Proposed MetaAlign for UDA

**Problem Formulation:** Unsupervised Domain Adaptation (UDA) aims to transfer the knowledge from labeled source domain to the unlabeled target domain. We mainly focus on object classification task but also investigate it for object

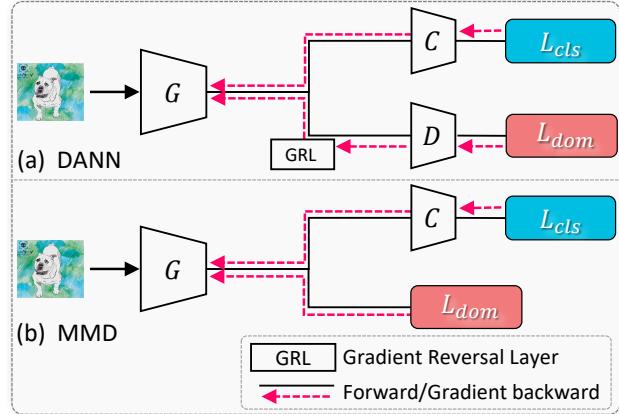


Figure 2: Two representative domain alignment based UDA frameworks. (a) DANN. (b) MMD.

detection. Without loss of generality, we take classification task as the instantiation to describe our approach.

For UDA classification, we denote the source domain as  $\mathcal{D}_S = \{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}_{i=1}^{N_s}$  with  $N_s$  labeled samples, where  $\mathbf{x}_i^s$  and  $\mathbf{y}_i^s$  denote the  $i^{th}$  sample and its class label respectively. The target domain is denoted as  $\mathcal{D}_T = \{\mathbf{x}_i^t\}_{i=1}^{N_t}$  with  $N_t$  unlabeled samples. Both domains share the same label space  $Y = \{1, 2, \dots, K\}$  with  $K$  object classes. UDA is expected to train model on  $\mathcal{D}_S$  and  $\mathcal{D}_T$ , to obtain high accuracy on the target test set.

The mainstream UDAs aim to align the source and target domains to alleviate the domain gap. Such alignments are, in general, not designed specially for classification task, *i.e.*, their optimization may not work harmoniously with that of object classification task. As a result, it may damage the discriminative power of features and thus impede attaining higher performance. To address this, as illustrated in Fig. 1, we introduce a meta-optimization based strategy MetaAlign to encourage the optimization consistency between domain alignment and object classification task itself.

To be self-included, we first describe several representative domain alignment based UDAs which we use as our baselines. Then, we introduce our MetaAlign to alleviate the above-mentioned optimization inconsistency problem.

#### 3.1. Recap of Alignment Based UDAs

Domain alignment based UDAs include adversarial training based [14, 15, 37] and explicit distribution similarity metric based [36, 38, 44, 56]. The core idea of the former category is to train a domain discriminator to distinguish source domain features from target domain features, meanwhile train the feature network to fool the discriminator to implicitly align domains. The latter one explicitly reduces domain discrepancy w.r.t. distribution discrepancy metrics, like Maximum Mean Discrepancy (MMD) [38, 36], moment distance [44], and second-order correlations [56].

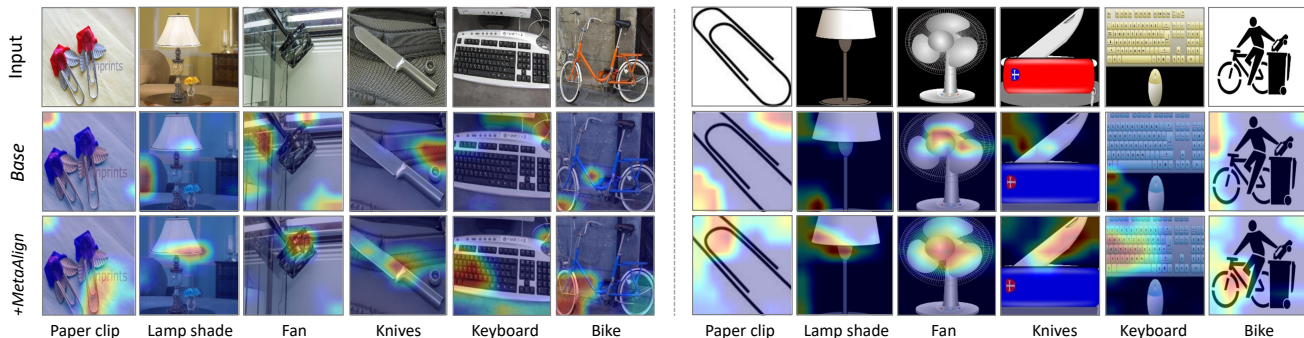


Figure 3: Visualization of the Grad-CAMs [54] w.r.t. the domain alignment task (domain classifier). The first row of left/right panels show the samples from source (Rw)/target (Cl) domains on Office-Home, while the second and third rows show the Grad-CAMs for Baseline (*Base*) and Baseline with MetaAlign (*+MetaAlign*), respectively.

**Adversarial Domain Adaptation.** To align the domain distributions via adversarial training, such methods try to optimize the object classification loss  $\mathcal{L}_{cls}$  and the domain alignment loss  $\mathcal{L}_{dom}$ , simultaneously. Typically, the network is equipped with a feature extractor/generator  $G$ , an object classifier  $C$ , and a domain discriminator/classifier  $D$ .

The specific object classification loss in the source domain  $\mathcal{D}_S$  is formulated as:

$$\mathcal{L}_{cls} = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathcal{L}_{ce}(C(G(\mathbf{x}_i^s)), \mathbf{y}_i^s). \quad (1)$$

where  $\mathcal{L}_{ce}$  is a typical cross-entropy loss.

The domain discriminator  $D$  is a two-class classification module, which aims to distinguish the target domain from source domain based on features from  $G$ . Adversarially,  $G$  aims to generate responses with aligned distributions for two domains to fool  $D$ . Particularly, the domain classification loss can be formulated as:

$$\begin{aligned} \mathcal{L}_{dom_{cls}} = & -\frac{1}{N_s} \sum_{i=1}^{N_s} \log(D(G(\mathbf{x}_i^s))) \\ & -\frac{1}{N_t} \sum_{j=1}^{N_t} \log(1 - D(G(\mathbf{x}_j^t))). \end{aligned} \quad (2)$$

We define the domain alignment loss as  $\mathcal{L}_{dom} = -\mathcal{L}_{dom_{cls}}$ . The more inseparable (larger  $\mathcal{L}_{dom_{cls}}$ ) two domains are, the smaller domain alignment loss  $\mathcal{L}_{dom}$  is. During training, we train  $D$  to maximize  $\mathcal{L}_{dom}$ , meanwhile  $\{G, C\}$  to minimize  $\mathcal{L}_{cls}$  and  $\mathcal{L}_{dom}$ :

$$\begin{aligned} & \max_D \mathcal{L}_{dom}, \\ & \min_{G,C} \mathcal{L}_{cls} + \mathcal{L}_{dom}, \end{aligned} \quad (3)$$

where we ignore the hyper-parameter  $\lambda$  (*i.e.*,  $\lambda \mathcal{L}_{dom}$ ) for balancing two losses, for simplicity. Actually, we keep  $\lambda$  same as that of the baselines [37, 9] in our experiments.

Fig. 2 (a) shows a seminal work *DANN* [14, 15], which constructs  $G$  as a CNN feature extractor. The extracted features from the two domains are fed to two task branches  $C$  and  $D$  simultaneously. Gradient Reversal Layer (GRL) [14], which flips the gradients to  $G$  from  $D$  during gradient back propagation, is used to simplify adversarial training. We also take *DANNPE* [9], an improved variant of *DANN*, as another strong baseline. It differs from *DANN* in two key aspects : 1) the input of  $D$  is the predicted classification probability; 2)  $D$  is prioritized on those easy-to-transfer samples by re-weighting with the entropy of object class prediction. Please see Supplementary for more details about *DANNPE*.

**Explicit Domain Alignment.** Without introducing additional domain discrimination modules, these methods directly reduce the distribution discrepancy between the features of source domain and target domain w.r.t. some discrepancy measurements/metrics. *MMD* [3] is a representative distribution discrepancy metric. It has been widely employed as the explicit domain alignment constraint for UDA [19, 36, 38]. Fig. 2 (b) illustrates one UDA framework with MMD constraint. Following [33], the domain alignment loss becomes:

$$\mathcal{L}_{dom} = \text{MMD}(\mathbf{F}^s, \mathbf{F}^t)^2. \quad (4)$$

where  $\mathbf{F}^s$  is the feature distribution of source domain. Please refer to Supplementary for more details.

### 3.2. Meta-learning to Align $\mathcal{L}_{cls}$ and $\mathcal{L}_{dom}$

Adversarial UDAs have brought significant performance improvement on multiple benchmarks. They promote the domain alignment, thus reduce the domain gap and enhance the transferability of the models to target domain. The optimization objective of domain alignment is to reduce discrepancy of features between source domain and target domain. However, without explicit coordination with the task of classification, the optimization direction of alignment

may be inconsistent with that of the classification task itself. Such inconsistency could hinder the optimization and lead to inferior performance.

The optimizations of domain alignment and classification can be considered as two tasks. In Fig. 3, we use *DANNPE* as our baseline to visualize the Grad-CAMs [54], which produce visual explanations for decisions and reflect the “important” region of the input for the predictions w.r.t. the domain alignment task. The second row (*Base*) denotes the Grad-CAMs obtained from the baseline, where the regions with higher responses indicate that their features are indistinguishable to  $D$  (considering the GRL has flipped the gradients). We can see that *Base* attains alignment usually on regions irrelevant to objects (e.g., backgrounds) or only on small partial regions of objects. The features of some foreground objects are still not aligned well, which would impede the transferability of the models on classification. It is well known that the foreground object regions are most discriminative for object classification [54, 72]. Not aligning the features of foreground objects well would damage the performance of classification.

With the alignment objective and the classification objective separately assigned to the network, there is a lack of effective interaction between the domain alignment task and classification task. These two tasks may have different gradient descent directions of their optimizations, resulting in optimization inconsistency.

One natural question to ask arises: how to easily incorporate the optimization consistency constraint for domain alignment and classification?

We propose to promote the optimization consistency between these two tasks by exploring a meta-optimization strategy. We draw inspiration from Model-Agnostic Meta-Learning (MAML) [13], which generally separates the samples into meta-train splits and meta-test splits, and uses meta-learning to train the model to be learned quickly on meta-test samples given the knowledge in meta-train. Meta-Learning Domain Generalization (MLDG) [31] simulates training-test domain shift during training by synthesizing virtual test domain, with meta-optimization objective requiring that steps to improve training domain performance should also improve testing domain performance.

In our work, we leverage meta-optimization to coordinate the domain alignment task and classification task. Particularly, rather than splitting the samples into meta-train and meta-test as in [13, 31], we treat the domain alignment task and classification task as meta-train (or meta-test) and meta-test (or meta-train) for the same set of samples.

A UDA network is jointly optimized with classification objective and domain alignment objective. The learnable network parameters consists of shared parameters  $\theta$ , the parameters specific to domain alignment  $\phi_d$ , and the param-

---

### Algorithm 1 MetaAlign Optimization Algorithm

---

- 1: **Input:** Source and target data sets  $\mathcal{D}_S$  and  $\mathcal{D}_T$
  - 2: **Init:** parameters  $\Psi = \{\theta, \phi_c, \beta, \phi_d\}$ , learning rate  $\eta, \alpha$
  - 3: **for**  $t$  in iterations **do**
  - Meta-train:**
  - 4:   Compute domain alignment loss  $\mathcal{L}_{dom}$  ▷ Eq. (2)
  - 5:   Update  $\theta$  w.r.t.  $\mathcal{L}_{dom}$ :  

$$\theta_m^{t+1} \leftarrow \theta_m^t - \alpha \beta_m \nabla_{\theta_m^t} \mathcal{L}_{dom}(\theta^t, \phi_d^t)$$
  - Meta-test:**
  - 6:   Compute classification loss  $\mathcal{L}_{cls}(\theta^{t+1}, \phi_c^t)$  ▷ Eq. (1)
  - Meta optimization:**
  - 7:   Compute total loss  $\mathcal{L}_{total}$  ▷ Eq. (8)
  - 8:   Update model parameters:  

$$\Psi^{t+1} \leftarrow \Psi^t - \eta \nabla_{\Psi^t} \mathcal{L}_{total}$$
  - 9: **end for**
  - 10: **Output:**  $\theta, \phi_c$
- 

eters specific to classification  $\phi_c$ . The general optimization objective can be formulated as:

$$\min_{\theta, \phi_c} \max_{\phi_d} \mathcal{L}_{dom}(\theta, \phi_d) + \mathcal{L}_{cls}(\theta, \phi_c). \quad (5)$$

which does not handle the potential optimization inconsistency of the two tasks.

With the intuition that meta-test task (e.g., classification) will be used to evaluate the effect of the model optimization on meta-train task (e.g., domain alignment), the overall meta-optimization objective can be formulated as:

$$\min_{\theta, \phi_c} \max_{\phi_d} \mathcal{L}_{dom}(\theta, \phi_d) + \mathcal{L}_{cls}(\theta - \alpha \nabla_{\theta} \mathcal{L}_{dom}(\theta, \phi_d), \phi_c). \quad (6)$$

which aims to optimize both the loss of meta-train  $\mathcal{L}_{dom}$ , and that of meta-test  $\mathcal{L}_{cls}$  after updating  $\theta$  during meta-train by one gradient descent step:  $\theta' \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}_{dom}(\theta, \phi_d)$ , where  $\alpha$  denotes the meta-learning rate. To alleviate the computational complexity, similar to [10, 13], we omit higher-order ones during the back-propagation of gradients.

This meta-optimization enables the explicit interaction between the two tasks. Following [31], we analyse Eq. (6) by approximating the second term using its first-order Taylor expansion as:

$$\min_{\theta, \phi_c} \max_{\phi_d} \mathcal{L}_{dom}(\theta, \phi_d) + \mathcal{L}_{cls}(\theta, \phi_c) - \alpha \nabla_{\theta} \mathcal{L}_{cls}(\theta, \phi_c) \nabla_{\theta} \mathcal{L}_{dom}(\theta, \phi_d). \quad (7)$$

Compared with the general optimization objective Eq. (5), the additional last term in Eq. (7) maximizes the dot product of  $\nabla_{\theta} \mathcal{L}_{cls}$  and  $\nabla_{\theta} \mathcal{L}_{dom}$ , which encourages the consistency of the optimization directions (gradients) of two tasks. In this way, both domain alignment and object classification are optimized in a coordinated way. We refer to our method as *MetaAlign* which Aligns the domain alignment task and classification task with a Meta-optimization strategy.

Method	Ar→Cl	Ar→Pr	Ar→Rw	Cl→Ar	Cl→Pr	Cl→Rw	Pr→Ar	Pr→Cl	Pr→Rw	Rw→Ar	Rw→Cl	Rw→Pr	Avg
Source-Only [21]	34.9	50.0	58.0	37.4	41.9	46.2	38.5	31.2	60.4	53.9	41.2	59.9	46.1
MCD(CVPR'18)[51]	48.9	68.3	74.6	61.3	67.6	68.8	57.0	47.1	75.1	69.1	52.2	79.6	64.1
TAT(ICML'19)[35]	51.6	69.5	75.4	59.4	69.5	68.6	59.5	50.5	76.8	70.9	56.6	81.6	65.8
ALDA(AAAI'20)[5]	53.7	70.1	76.4	60.2	72.6	71.5	56.8	51.9	77.1	70.2	56.3	82.1	66.6
Sym(CVPR'19)[71]	47.7	72.9	78.5	64.2	71.3	74.2	63.6	47.6	79.4	73.8	50.8	82.6	67.2
TADA(AAAI'19)[67]	53.1	72.3	77.2	59.1	71.2	72.1	59.7	53.1	78.4	72.4	60.0	82.9	67.6
MDD(ICML'19)[70]	54.9	73.7	77.8	60.0	71.4	71.8	61.2	53.6	78.1	72.5	60.2	82.3	68.1
BNM(CVPR'20)[8]	56.2	73.7	79.0	63.1	73.6	74.0	62.4	54.8	80.7	72.4	58.9	83.5	69.4
MMD	49.1	67.0	74.7	54.5	62.9	65.7	55.3	45.7	74.5	68.1	52.5	78.6	62.3
+MetaAlign	49.4 <sub>†</sub>	67.2 <sub>†</sub>	75.5 <sub>†</sub>	58.6 <sub>†</sub>	64.7 <sub>†</sub>	67.2 <sub>†</sub>	55.5 <sub>†</sub>	46.1 <sub>†</sub>	74.8 <sub>†</sub>	69.0 <sub>†</sub>	52.1 <sub>↓</sub>	78.9 <sub>†</sub>	63.3 <sub>†</sub>
DANN(ICML'15)[15]†	45.8	63.4	71.9	53.6	61.9	62.6	49.1	39.7	73.0	64.6	47.8	77.8	59.2
+MetaAlign	48.6 <sub>†</sub>	69.5 <sub>†</sub>	76.0 <sub>†</sub>	58.1 <sub>†</sub>	65.7 <sub>†</sub>	68.3 <sub>†</sub>	54.9 <sub>†</sub>	44.4 <sub>†</sub>	75.3 <sub>†</sub>	68.5 <sub>†</sub>	50.8 <sub>†</sub>	80.1 <sub>†</sub>	63.3 <sub>†</sub>
CDAN(NeurIPS'18)[37]	50.7	70.6	76.0	57.6	70.0	70.0	57.4	50.9	77.3	70.9	56.7	81.6	65.8
+MetaAlign	55.2 <sub>†</sub>	70.5 <sub>↓</sub>	77.6 <sub>†</sub>	61.5 <sub>†</sub>	70.0 <sub>–</sub>	70.0 <sub>–</sub>	58.7 <sub>†</sub>	55.7 <sub>†</sub>	78.5 <sub>†</sub>	73.3 <sub>†</sub>	61.0 <sub>†</sub>	81.7 <sub>†</sub>	67.8 <sub>†</sub>
DANNPE	54.7	72.8	78.5	62.3	71.1	73.1	61.0	53.0	80.0	72.8	56.5	83.4	68.3
+MetaAlign	57.1 <sub>†</sub>	74.5 <sub>†</sub>	80.1 <sub>†</sub>	64.9 <sub>†</sub>	73.6 <sub>†</sub>	74.6 <sub>†</sub>	62.5 <sub>†</sub>	54.8 <sub>†</sub>	80.6 <sub>†</sub>	73.6 <sub>†</sub>	60.3 <sub>†</sub>	84.7 <sub>†</sub>	70.1 <sub>†</sub>
GVB(CVPR'20)[9]	57.0	74.7	79.8	64.6	74.1	74.6	65.2	55.1	81.0	<b>74.6</b>	59.7	84.3	70.4
+MetaAlign	<b>59.3<sub>†</sub></b>	<b>76.0<sub>†</sub></b>	<b>80.2<sub>†</sub></b>	<b>65.7<sub>†</sub></b>	<b>74.7<sub>†</sub></b>	<b>75.1<sub>†</sub></b>	<b>65.7<sub>†</sub></b>	<b>56.5<sub>†</sub></b>	<b>81.6<sub>†</sub></b>	74.1 <sub>↓</sub>	<b>61.1<sub>†</sub></b>	<b>85.2<sub>†</sub></b>	<b>71.3<sub>†</sub></b>

Table 1: Classification accuracy (%) of different UDAs on Office-Home with ResNet-50 as backbone. We re-implement all the adopted baselines for MetaAlign. † denotes our re-implemented result is different from the one reported in other papers.

In Eq. (7), we place the consistency constraint on all the parameters  $\theta$  (including different layers) shared by two tasks. Actually, different layers in a CNN learn features with different semantics. Intuitively, they should be treated differently when aligning their optimizations, where the optimization consistency for some layers may be more important than other layers. We propose to adaptively learn the importance levels for different groups of layers for better optimization. Particularly, we partition the layers into  $M$  groups (e.g., each convolutional block of ResNet as a group) and learn a scalar weight  $\beta_m$  for the  $m^{\text{th}}$  group. The optimization objective is thus formulated as:

$$\begin{aligned}
& \min_{\theta, \phi_c, \beta} \max_{\phi_d} \mathcal{L}_{dom}(\theta, \phi_d) \\
& + \mathcal{L}_{cls} \left( \left\{ \theta_m - \alpha \beta_m \nabla_{\theta_m} \mathcal{L}_{dom}(\theta, \phi_d) \right\}_{m=1}^M, \phi_c \right) \\
& + \mathcal{L}_{\beta}(\beta).
\end{aligned} \tag{8}$$

where  $\theta_m$  denotes the parameters of the  $m^{\text{th}} \in \{1, \dots, M\}$  group. To avoid trivial solution, we add  $L_1$  constraint on  $\beta = \{\beta_m\}_m$ :  $\mathcal{L}_{\beta} = \|\sum_{m=1}^M \beta_m - B\|_1$ , where  $B$  is a hyper-parameter. We denote the entire losses in Eq. (8) as  $\mathcal{L}_{total}$  for simplicity.

**Training.** In practice, we can iteratively choose one of the two tasks (domain alignment task and object classification task) as meta-train while the other as meta-test. We describe the training procedure in Alg. 1. For simplicity, we only show one case, i.e., the domain alignment task and object classification task are taken as meta-train and meta-test respectively. The other case is similar.

## 4. Experiments

In this section, we validate the effectiveness of our proposed MetaAlign on UDA for classification in Sec. 4.1 and for object detection in Sec. 4.2. Besides, we further show its effectiveness on DG for classification in Sec. 4.3. Due to the space constraint, we refer readers to Supplementary for more details for all experiments.

### 4.1. UDA for Classification

#### 4.1.1 Datasets and Settings

We conduct UDA experiments on two popular benchmarks **Office31** [49] and **Office-Home** [63]. 1) **Office31** is a standard benchmark for domain adaptive classification. It contains images of 31 categories, drawn from three domains: Amazon (A), Webcam (W), and DSLR (D). Following the typical setting [9, 71, 51], we evaluate the methods on one-source to one-target domain adaptation. 2) **Office-Home** is a more challenging recent dataset for UDA. It consists of images from 4 different domains: Art (Ar), Clip Art (Cl), Product (Pr), and Real-World (Rw). Each domain contains 65 object categories found typically in office and home environments. We evaluate our method in all the 12 one-source to one-target adaptation cases. All reported results are obtained from the average of multiple runs.

#### 4.1.2 Ablation Study

**Effectiveness of MetaAlign on Various Baselines.** Our proposed MetaAlign is generic which could be applied to

Method	A $\rightarrow$ W	A $\rightarrow$ D	W $\rightarrow$ A	W $\rightarrow$ D	D $\rightarrow$ A	D $\rightarrow$ W	Avg
Source-Only [21]	68.4 $\pm$ 0.2	68.9 $\pm$ 0.2	60.7 $\pm$ 0.3	99.3 $\pm$ 0.1	62.5 $\pm$ 0.3	96.7 $\pm$ 0.1	76.1
CDAN (NeurIPS'18) [37]	94.1 $\pm$ 0.1	92.9 $\pm$ 0.2	69.3 $\pm$ 0.3	<b>100.0</b> $\pm$ 0	71.0 $\pm$ 0.3	98.6 $\pm$ 0.1	87.7
TAT (ICML'19) [35]	92.5 $\pm$ 0.3	93.2 $\pm$ 0.2	72.1 $\pm$ 0.3	<b>100.0</b> $\pm$ 0	73.1 $\pm$ 0.3	<b>99.3</b> $\pm$ 0.1	88.4
TADA (AAAI'19) [67]	94.3 $\pm$ 0.3	91.6 $\pm$ 0.3	73.0 $\pm$ 0.3	99.8 $\pm$ 0.2	72.9 $\pm$ 0.2	98.7 $\pm$ 0.1	88.4
Sym (CVPR'19) [71]	90.8 $\pm$ 0.1	93.9 $\pm$ 0.5	72.5 $\pm$ 0.5	<b>100.0</b> $\pm$ 0	74.6 $\pm$ 0.6	98.8 $\pm$ 0.3	88.4
BNM (CVPR'20) [8]	92.8	92.9	73.8	<b>100.0</b>	73.5	98.8	88.6
ALDA (AAAI'20) [5]	<b>95.6</b> $\pm$ 0.5	94.0 $\pm$ 0.4	72.5 $\pm$ 0.2	<b>100.0</b> $\pm$ 0	72.2 $\pm$ 0.4	97.7 $\pm$ 0.1	88.7
MDD (ICML'19) [70]	94.5 $\pm$ 0.3	93.5 $\pm$ 0.2	72.2 $\pm$ 0.1	<b>100.0</b> $\pm$ 0	74.6 $\pm$ 0.3	98.4 $\pm$ 0.1	88.9
DANNPE	92.5 $\pm$ 0.5	89.9 $\pm$ 0.3	71.1 $\pm$ 0.3	99.9 $\pm$ 0.1	70.7 $\pm$ 0.4	98.5 $\pm$ 0.3	87.0
+MetaAlign	93.9 $\pm$ 0.4 $\uparrow$	91.6 $\pm$ 0.3 $\uparrow$	<b>74.1</b> $\pm$ 0.2 $\uparrow$	<b>100.0</b> $\pm$ 0 $\uparrow$	73.7 $\pm$ 0.2 $\uparrow$	98.7 $\pm$ 0.2 $\uparrow$	88.7 $\uparrow$
GVB (CVPR'20) [9] $\uparrow$	92.0 $\pm$ 0.3	91.4 $\pm$ 0.5	73.4 $\pm$ 0.1	<b>100.0</b> $\pm$ 0	74.9 $\pm$ 0.5	98.7 $\pm$ 0	88.3
+MetaAlign	93.0 $\pm$ 0.5 $\uparrow$	<b>94.5</b> $\pm$ 0.3 $\uparrow$	73.6 $\pm$ 0 $\uparrow$	<b>100.0</b> $\pm$ 0 $\_$	<b>75.0</b> $\pm$ 0.3 $\uparrow$	98.6 $\pm$ 0 $\_$	<b>89.2</b> $\uparrow$

Table 2: Classification accuracy (mean  $\pm$  std %) of different UDAs on Office31 with ResNet-50 as backbone. We re-implement all the adopted baselines.  $\uparrow$  denotes the result is different from the reported one in the original paper.

Source-Only	DANNPE	+MetaAlign w/o $\beta$	+MetaAlign
46.1	68.3	69.7	70.1

Table 3: Ablation study on Office-Home with DANNPE as the baseline.




Figure 4: t-SNE visualization of features learned by Source-Only (left), DANNPE (middle) and DANNPE+MetaAlign (right).

alleviate the optimization inconsistency of most existing domain alignment based UDAs. We use various alignment-based UDAs as our baselines to validate the effectiveness of MeanAlign. Specifically, we adopt five baselines. 1) *MMD*, 2) *DANN*, and 3) *DANNPE* have been described in Sec. 3.1 in detail. 4) *CDAN* [37] aligns domain at class-level. 5) *GVB* [9] is a most recent state-of-the-art method with enhanced  $C$  and  $D$ . Note that *MMD* explicitly reduces the discrepancy of domains, while the others belong to adversarial learning based approaches.

Table 1 shows the comparisons on Office-Home. Our MetaAlign consistently improves the accuracy of all the five baselines, *i.e.*, 1.0%, 3.9%, 2.0%, 1.8%, 0.9% on average for *MMD*, *DANN*, *CDAN*, *DANNPE*, *GVB*, respectively, regardless of the design differences on  $\mathcal{L}_{dom}$ . With the help of MetaAlign, domain alignment and classification are optimized in a coordinated way, resulting in more efficient optimization.

**Effectiveness of Re-weighting with  $\beta$  in MetaAlign.** In our design,  $\beta_m, m = 1, \dots, M$  in Eq. (8) are learned to

allocate the levels of importance for different layer groups on the consistency constraint. We validate its effectiveness on top of the baseline *DANNPE* on Office-Home in Table 3. Our MetaAlign improves over *DANNPE* by 1.8% in accuracy. Without  $\beta$ , the gain decreases to 1.4%.

**Which Task as Meta-Train?** We could treat any one of the two tasks as meta-train and the other as meta-test. We could also iteratively exchange their roles during training. Experimental results show their results are very close ( $< 0.3\%$  accuracy). The explanation lies in that these settings have the same optimization objective as Eq. (7).

### 4.1.3 Comparisons with State-of-the-Arts

To compare with previous state-of-the-art UDAs, we incorporate our MetaAlign optimization strategy into the recent strong UDA method *GVB* [9], termed as *GVB+MetaAlign*. Table 1 and Table 2 show the comparisons with the state-of-the-art approaches on Office-Home and Office31, respectively. *GVB+MetaAlign* outperforms *GVB* and achieves the best performance on both datasets.

### 4.1.4 Feature Visualization

As analysed in Sec. 3.2, we expect MetaAlign to enforce the domain alignment task and object classification task to be optimized in a coordinated way. To validate it, we visualize the Grad-CAMs [54] w.r.t. the domain alignment task in Fig. 3. With MetaAlign, the domain alignment task focuses on the regions more related to foreground objects compared with Baseline. These regions play the most important role for object classification task [72, 54], which is also validated in the Supplementary. Aligning these features helps improve object classification accuracy indeed.

We also visualize the learned features by t-SNE [50], on task Pr  $\rightarrow$  Cl in Fig. 4. It is shown that Source-Only works well only in source domain but poorly in target domain

Source → Target	AGG	MMD-AAE[33] (CVPR'18)	CrossGrad[55] (ICLR'18)	MetaReg[2] (NeurIPS'18)	JiGen[4] (CVPR'19)	MLDG[31] (AAAI'18)	MASF[10] (NeurIPS'19)	Epi-FCR[32] (ICCV'19)	MMLD[41] (AAAI'20)	DANNPE	+MetaAlign
A.C.S → P	94.4	96.0	94.0	94.3	96.0	94.3	95.0	93.9	<b>96.1</b>	95.2	95.5
C.P.S → A	77.6	75.2	78.7	79.5	79.4	79.5	80.3	<b>82.1</b>	81.3	76.5	78.5
A.P.S → C	73.9	72.7	73.3	75.4	77.3	75.3	77.2	77.0	77.2	77.2	<b>77.8</b>
A.C.P → S	70.3	64.2	65.1	72.2	71.4	71.5	71.7	73.0	72.3	74.1	<b>75.7</b>
Avg.	79.1	77.0	77.8	80.4	80.5	80.7	81.0	81.5	81.8	80.7	<b>81.9</b>

Table 4: Accuracy (%) of different domain generalization methods on PACS with ResNet-18 as backbone. Best in bold.

Methods	bike	bird	car	cat	dog	person	mAP
Source Only	68.8	46.8	37.2	32.7	21.3	60.7	44.6
BDC-Faster[50](CVPR'19)	68.6	48.3	47.2	26.5	21.7	60.5	45.5
WST+BSR[26](ICCV'19)	75.6	45.8	<b>49.3</b>	34.1	30.3	64.1	49.9
MAF[22](ICCV'19)	73.4	55.7	46.4	36.8	28.9	60.8	50.3
DT-UDA[25](CVPR'18)	82.8	47.0	40.2	34.6	35.3	62.5	50.4
ATF[23](ECCV'20)	78.8	<b>59.9</b>	47.9	<b>41.0</b>	34.8	66.9	54.9
W-DA[50](CVPR'19)	66.4	53.7	43.8	37.9	31.9	65.3	49.8
W-DA+MetaAlign	74.9 <sub>↑</sub>	54.0 <sub>↑</sub>	43.7 <sub>↓</sub>	38.1 <sub>↑</sub>	35.2 <sub>↑</sub>	66.9 <sub>↑</sub>	52.1 <sub>↑</sub>
SW-DA[50](CVPR'19)	76.1	52.7	49.1	36.3	40.2	66.3	53.5
SW-DA+MetaAlign	<b>83.7<sub>↑</sub></b>	53.2 <sub>↑</sub>	48.7 <sub>↓</sub>	38.7 <sub>↑</sub>	<b>42.0<sub>↑</sub></b>	<b>67.2<sub>↑</sub></b>	<b>55.6<sub>↑</sub></b>

Table 5: Performance of UDAs for object detection from Pascal VOC to Watercolor2k in terms of mAP. Best in bold.

without domain alignment. *DANNPE* aligns domains well via adversarial learning. Further, employing MetaAlign arrives at much better alignment results, where the clusters are more compact and less data points scatter at the boundaries between clusters. The visualization result further validates the effectiveness of MetaAlign for domain alignment based UDA.

## 4.2. Experiments on UDA for Object Detection

Adversarial learning has also been exploited in UDA for object detection [7, 50]. It is natural to align foreground objects across domains in this task, where the aforementioned optimization inconsistency issue still exists. Our MetaAlign is generic and is expected to work well for this task. We take a Faster RCNN[48] based SW-DA [50] as the baseline, and conduct UDA experiments from Pascal VOC [11, 12] to Watercolor2k [25]. All reported results are obtained from the average of multiple runs. Please refer to Supplementary for details about datasets, experimental settings, and the introduction of competitors.

As shown in Table 5, our MetaAlign strategy improves the mAP of the two baselines *W-DA* (*SW-DA* without local alignment) and *SW-DA* by 2.3 (4.6%) and 1.9 (3.5%), respectively. The latter one achieves state-of-the-art performance, compared with recent methods. Note that the results on ‘bird’ are unstable due to the insufficient data (the number of bird bounding boxes are about 2.6% of all bounding boxes in the dataset). Some qualitative comparisons (see Supplementary) demonstrate MetaAlign improves the object classification accuracy of the predicted bounding boxes, thanks to the coordination between domain alignment and

object classification. These validate that our MetaAlign strategy is compatible with UDAs for different vision tasks.

## 4.3. Experiments on Domain Generalization

Learning domain-invariant features is also widely explored in DG. Therefore we further apply our MetaAlign on DG to validate its generalizability. *DANNPE* is originally designed for UDA. However, its goal, to learn domain-invariant features, fits well with DG. Therefore we repurpose it as a baseline for DG to incorporate our MetaAlign. We perform experiments on PACS [30]. All reported results are obtained from the average of multiple runs. Please refer to Supplementary for more details about the dataset, settings, and the description of competitors.

As shown in Table 4, *DANNPE* also works for DG with about 1.6% improvement over *AGG*. With our proposed MetaAlign strategy, the accuracy is further improved by 1.2%. We reckon that MetaAlign enforces the domain-invariant features and classification discriminative features to be learned in concert. *DANNPE+MetaAlign* outperforms previous state-of-the-art methods on the average accuracy, especially in the most challenging scenario where the target domain is *Sketch*. *Sketch* has the largest domain gap with other domains, therefore, more powerful domain-invariant features are required.

## 5. Conclusion

In this paper, we pinpoint the optimization inconsistency problem between the domain alignment task and the classification task itself in alignment-based UDAs. To mitigate it, we propose a meta-optimization based strategy named MetaAlign, which treats one of these two tasks as meta-train and the other as meta-test. The analysis of the optimization objective of MetaAlign reveals that the two tasks will be optimized in a coordinated way. The experimental results validate that MetaAlign is applicable to various alignment-based UDAs for classification and detection.

## 6. Acknowledgments

This work was supported in part by NSFC under Grant U1908209, 61632001 and the National Key Research and Development Program of China 2018AAA0101400.



## References

- [1] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. In *NeurIPS*, pages 3981–3989, 2016. [3](#)
- [2] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. In *NeurIPS*, pages 998–1008, 2018. [3](#), [8](#)
- [3] Karsten M Borgwardt, Arthur Gretton, Malte J Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alex J Smola. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):e49–e57, 2006. [2](#), [4](#)
- [4] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *CVPR*, pages 2229–2238, 2019. [8](#)
- [5] Minghao Chen, Shuai Zhao, Haifeng Liu, and Deng Cai. Adversarial-learned loss for domain adaptation. In *AAAI*, pages 3521–3528, 2020. [1](#), [2](#), [6](#), [7](#)
- [6] Qingchao Chen, Yang Liu, Zhaowen Wang, Ian Wassell, and Kevin Chetty. Re-weighted adversarial adaptation network for unsupervised domain adaptation. In *CVPR*, pages 7976–7985, 2018. [2](#)
- [7] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *CVPR*, pages 3339–3348, 2018. [3](#), [8](#)
- [8] Shuhao Cui, Shuhui Wang, Junbao Zhuo, Liang Li, Qingming Huang, and Qi Tian. Towards discriminability and diversity: Batch nuclear-norm maximization under label insufficient situations. In *CVPR*, pages 3941–3950, 2020. [6](#), [7](#)
- [9] Shuhao Cui, Shuhui Wang, Junbao Zhuo, Chi Su, Qingming Huang, and Qi Tian. Gradually vanishing bridge for adversarial domain adaptation. In *CVPR*, pages 12455–12464, 2020. [1](#), [2](#), [4](#), [6](#), [7](#)
- [10] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. In *NeurIPS*, pages 6450–6461, 2019. [3](#), [5](#), [8](#)
- [11] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>. [8](#)
- [12] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. [8](#)
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, page 1126–1135, 2017. [3](#), [5](#)
- [14] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015. [1](#), [2](#), [3](#), [4](#)
- [15] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1):2096–2030, 2016. [1](#), [2](#), [3](#), [4](#), [6](#)
- [16] Bo Geng, Dacheng Tao, and Chao Xu. Daml: Domain adaptation metric learning. *IEEE Transactions on Image Processing*, 20(10):2980–2989, 2011. [1](#)
- [17] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *ICCV*, pages 2551–2559, 2015. [3](#)
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014. [1](#)
- [19] Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex J Smola. A kernel method for the two-sample-problem. In *NeurIPS*, pages 513–520, 2007. [4](#)
- [20] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *ICCV*, pages 2980–2988, 2017. [1](#)
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [6](#), [7](#)
- [22] Zhenwei He and Lei Zhang. Multi-adversarial faster-rcnn for unrestricted object detection. In *ICCV*, pages 6668–6677, 2019. [8](#)
- [23] Zhenwei He and Lei Zhang. Domain adaptive object detection via asymmetric tri-way faster-rcnn. *ECCV*, 2020. [8](#)
- [24] Cheng-Chun Hsu, Yi-Hsuan Tsai, Yen-Yu Lin, and Ming-Hsuan Yang. Every pixel matters: Center-aware feature alignment for domain adaptive object detector. In *ECCV*, pages 733–748. [3](#)
- [25] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *CVPR*, pages 5001–5009, 2018. [3](#), [8](#)
- [26] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Chang-ick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In *ICCV*, pages 6092–6101, 2019. [3](#), [8](#)
- [27] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, 2015. [3](#)
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, page 1097–1105, 2012. [1](#)
- [29] Da Li and Timothy Hospedales. Online meta-learning for multi-source and semi-supervised domain adaptation. In *ECCV*, pages 382–403, 2020. [3](#)
- [30] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *ICCV*, pages 5542–5550, 2017. [8](#)
- [31] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. *AAAI*, 2018. [3](#), [5](#), [8](#)

- [32] Da Li, Jianshu Zhang, Yongxin Yang, Cong Liu, Yi-Zhe Song, and Timothy M Hospedales. Episodic training for domain generalization. In *ICCV*, pages 1446–1455, 2019. 3, 8
- [33] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *CVPR*, pages 5400–5409, 2018. 3, 4, 8
- [34] Ke Li and Jitendra Malik. Learning to optimize neural nets. *arXiv preprint arXiv:1703.00441*, 2017. 3
- [35] Hong Liu, Mingsheng Long, Jianmin Wang, and Michael Jordan. Transferable adversarial training: A general approach to adapting deep classifiers. In *ICML*, pages 4013–4022, 2019. 2, 6, 7
- [36] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, pages 97–105, 2015. 1, 2, 3, 4
- [37] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, pages 1645–1655, 2018. 1, 2, 3, 4, 6, 7
- [38] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. In *NeurIPS*, pages 136–144, 2016. 3, 4
- [39] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, pages 2208–2217, 2017. 2
- [40] Zhihe Lu, Yongxin Yang, Xiatian Zhu, Cong Liu, Yi-Zhe Song, and Tao Xiang. Stochastic classifiers for unsupervised domain adaptation. In *CVPR*, pages 9111–9120, 2020. 2
- [41] Toshihiko Matsuura and Tatsuya Harada. Domain generalization using a mixture of multiple latent domains. In *AAAI*, pages 11749–11756, 2020. 1, 3, 8
- [42] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *ICML*, pages 10–18, 2013. 3
- [43] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. In *AAAI*, 2018. 1
- [44] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, pages 1406–1415, 2019. 1, 2, 3
- [45] Xingchao Peng and Kate Saenko. Synthetic to real adaptation with generative correlation alignment networks. In *WACV*, pages 1982–1991, 2018. 1
- [46] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *CVPR*, pages 12556–12565, 2020. 3
- [47] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *CVPR*, pages 12556–12565, 2020. 3
- [48] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, pages 91–99, 2015. 1, 8
- [49] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, pages 213–226, 2010. 6
- [50] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *CVPR*, pages 6956–6965, 2019. 3, 7, 8
- [51] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, pages 3723–3732, 2018. 2, 6
- [52] Swami Sankaranarayanan, Yogesh Balaji, Carlos D Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *CVPR*, pages 8503–8512, 2018. 2
- [53] Jürgen Schmidhuber. On learning how to learn learning strategies. 1995. 1, 3
- [54] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pages 618–626, 2017. 4, 5, 7
- [55] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. In *ICLR*, 2018. 3, 8
- [56] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *AAAI*, 2016. 1, 2, 3
- [57] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, pages 443–450, 2016. 1, 2
- [58] Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012. 1, 3
- [59] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, pages 9627–9636, 2019. 3
- [60] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR*, pages 1521–1528, 2011. 1
- [61] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pages 7167–7176, 2017. 2
- [62] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014. 1, 2
- [63] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, pages 5018–5027, 2017. 6
- [64] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, pages 3630–3638, 2016. 3
- [65] Riccardo Volpi, Pietro Morerio, Silvio Savarese, and Vittorio Murino. Adversarial feature augmentation for unsupervised domain adaptation. In *CVPR*, pages 5495–5504, 2018. 2
- [66] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In *NeurIPS*, pages 5334–5344, 2018. 3
- [67] Ximei Wang, Liang Li, Weirui Ye, Mingsheng Long, and Jianmin Wang. Transferable attention for domain adaptation. In *AAAI*, volume 33, pages 5345–5352, 2019. 1, 6, 7

- [68] Hongliang Yan, Yukang Ding, Peihua Li, Qilong Wang, Yong Xu, and Wangmeng Zuo. Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *CVPR*, pages 2272–2281, 2017. [2](#)
- [69] Werner Zellinger, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. Central moment discrepancy (cmd) for domain-invariant representation learning. In *ICLR*, 2017. [1](#)
- [70] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *ICML*, pages 7404–7413, 2019. [6](#), [7](#)
- [71] Yabin Zhang, Hui Tang, Kui Jia, and Mingkui Tan. Domain-symmetric networks for adversarial domain adaptation. In *CVPR*, pages 5031–5040, 2019. [1](#), [2](#), [6](#), [7](#)
- [72] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016. [5](#), [7](#)
- [73] Xinge Zhu, Jiangmiao Pang, Ceyuan Yang, Jianping Shi, and Dahua Lin. Adapting object detectors via selective cross-domain alignment. In *CVPR*, pages 687–696, 2019. [3](#)
- [74] Chenfan Zhuang, Xintong Han, Weilin Huang, and Matthew Scott. ifan: Image-instance full alignment networks for adaptive object detection. In *AAAI*, volume 34, pages 13122–13129, 2020. [3](#)