

NETWORK ARCHITECTURE SEARCH FOR DOMAIN ADAPTATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Deep networks have been used to learn transferable representations for domain adaptation. Existing deep domain adaptation methods systematically employ popular hand-crafted networks designed specifically for image-classification tasks, leading to sub-optimal domain adaptation performance. In this paper, we present Neural Architecture Search for Domain Adaptation (NASDA), a principle framework that leverages differentiable neural architecture search to derive the optimal network architecture for domain adaptation task. NASDA is designed with two novel training strategies: neural architecture search with multi-kernel Maximum Mean Discrepancy to derive the optimal architecture, and adversarial training between a feature generator and a batch of classifiers to consolidate the feature generator. We demonstrate experimentally that NASDA leads to state-of-the-art performance on several domain adaptation benchmarks.

1 INTRODUCTION

Supervised machine learning models (Φ) aim to minimize the empirical test error ($\epsilon(\Phi(\mathbf{x}), \mathbf{y})$) by optimizing Φ on training data (\mathbf{x}) and ground truth labels (\mathbf{y}), assuming that the training and testing data are sampled *i.i.d* from the same distribution. While in practical, the training and testing data are typically collected from related domains under different distributions, a phenomenon known as domain shift (or domain discrepancy) (Quionero-Candela et al., 2009). To avoid the cost of annotating each new test data, Unsupervised Domain Adaptation (UDA) tackles domain shift by transferring the knowledge learned from a rich-labeled source domain ($P(\mathbf{x}^s, \mathbf{y}^s)$) to the unlabeled target domain ($Q(\mathbf{x}^t)$). Recently unsupervised domain adaptation research has achieved significant progress with techniques like discrepancy alignment (Long et al., 2017; Tzeng et al., 2014; Ghifary et al., 2014; Peng & Saenko, 2018; Long et al., 2015; Sun & Saenko, 2016), adversarial alignment (Xu et al., 2019a; Liu & Tuzel, 2016; Tzeng et al., 2017; Liu et al., 2018a; Ganin & Lempitsky, 2015; Saito et al., 2018; Long et al., 2018), and reconstruction-based alignment (Yi et al., 2017; Zhu et al., 2017; Hoffman et al., 2018; Kim et al., 2017). While such models typically learn feature mapping from one domain ($\Phi(\mathbf{x}^s)$) to another ($\Phi(\mathbf{x}^t)$) or derive a joint representation across domains ($\Phi(\mathbf{x}^s) \otimes \Phi(\mathbf{x}^t)$), the developed models have limited capacities in deriving an optimal neural architecture specific for domain transfer.

To advance network designs, neural architecture search (NAS) automates the net architecture engineering process by reinforcement supervision (Zoph & Le, 2017) or through neuro-evolution (Real et al., 2019a). Conventional NAS models aim to derive neural architecture α along with the network parameters w , by solving a bilevel optimization problem (Anandalingam & Friesz, 1992): $\Phi_{\alpha, w} = \arg \min_{\alpha} \mathcal{L}_{val}(w^*(\alpha), \alpha)$ s.t. $w^*(\alpha) = \operatorname{argmin}_w \mathcal{L}_{train}(w, \alpha)$, where \mathcal{L}_{train} and \mathcal{L}_{val} indicate the training and validation loss, respectively. While recent works demonstrate competitive performance on tasks such as image classification (Zoph et al., 2018; Liu et al., 2018c;b; Real et al., 2019b) and object detection (Zoph & Le, 2017), designs of existing NAS algorithms typically assume that the training and testing domain are sampled from the same distribution, neglecting the scenario where two data domains or multiple feature distributions are of interest.

To efficiently devise a neural architecture across different data domains, we propose a novel learning task called Neural Architecture Search for Domain Adaptation (NASDA). The ultimate goal of NASDA is to minimize the validation loss of the target domain (\mathcal{L}_{val}^t). We postulate that a solution to NASDA should not only minimize validation loss of the source domain (\mathcal{L}_{val}^s), but should also

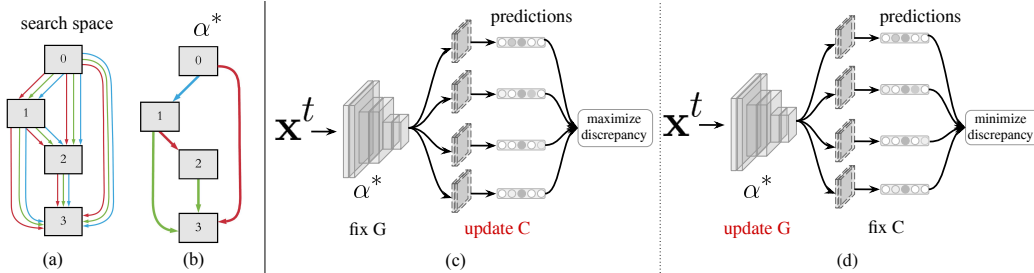


Figure 1: An overview of NASDA: (a) Continuous relaxation of the research space by placing a mixture of the candidate operations on each edge. (b) Inducing the final architecture by joint optimization of the neural architecture parameters α and network weights w , supervised by minimizing the validation loss on the source domain and reducing the domain discrepancy. (c)(d) Adversarial training of the derive feature generator G and classifiers C .

reduce the domain gap between the source and target. To this end, we propose a new NAS learning schema:

$$\Phi_{\alpha,w} = \operatorname{argmin}_{\alpha} \mathcal{L}_{val}^s(w^*(\alpha), \alpha) + \operatorname{disc}(\Phi^*(\mathbf{x}^s), \Phi^*(\mathbf{x}^t)) \quad (1)$$

$$\text{s.t. } w^*(\alpha) = \operatorname{argmin}_w \mathcal{L}_{train}^s(w, \alpha) \quad (2)$$

where $\Phi^* = \Phi_{\alpha,w^*(\alpha)}$, and $\operatorname{disc}(\Phi^*(\mathbf{x}^s), \Phi^*(\mathbf{x}^t))$ denotes the domain discrepancy between the source and target. Note that in unsupervised domain adaptation, \mathcal{L}_{train}^t and \mathcal{L}_{val}^t cannot be computed directly due to the lack of label in the target domain.

Inspired by the past works in NAS and unsupervised domain adaptation, we propose in this paper an instantiated NASDA model, which comprises of two training phases, as shown in Figure 1. The first is the neural architecture searching phase, aiming to derive an optimal neural architecture (α^*), following the learning schema of Equation 1,2. Inspired by Differentiable ARchiTecture Search (DARTS) (Liu et al., 2019a), we relax the search space to be continuous so that α can be optimized with respect to \mathcal{L}_{val}^s and $\operatorname{disc}(\Phi(\mathbf{x}^s), \Phi(\mathbf{x}^t))$ by gradient descent. Specifically, we enhance the feature transferability by embedding the hidden representations of the task-specific layers to a reproducing kernel Hilbert space where the mean embeddings can be explicitly matched by minimizing $\operatorname{disc}(\Phi(\mathbf{x}^s), \Phi(\mathbf{x}^t))$. We use multi-kernel Maximum Mean Discrepancy (MK-MMD) (Gretton et al., 2007) to evaluate the domain discrepancy.

The second training phase aims to learn a good feature generator with task-specific loss, based on the derived α^* from the first phase. To establish this goal, we use the derived deep neural network (Φ_{α^*}) as the feature generator (G) and devise an adversarial training process between G and a batch of classifiers C . The high-level intuition is to first diversify C in the training process, and train G to generate features such that the diversified C can have similar outputs. The training process is similar to Maximum Classifier Discrepancy framework (MCD) (Saito et al., 2018) except that we extend the dual-classifier in MCD to an ensembling of multiple classifiers. Experiments on standard UDA benchmarks demonstrate the effectiveness of our derived NASDA model in achieving significant improvements over state-of-the-art methods.

Our contributions of this paper are highlighted as follows:

- We formulate a novel dual-objective task of Neural Architecture Search for Domain Adaptation (NASDA), which optimize neural architecture for unsupervised domain adaptation, concerning both source performance objective and transfer learning objective.
- We propose an instantiated NASDA model that comprises two training stages, aiming to derive optimal architecture parameters α^* and feature extractor G , respectively. We are the first to show the effectiveness of MK-MMD in NAS process specified for domain adaptation.
- Extensive experiments on multiple cross-domain recognition tasks demonstrate that NASDA achieves significant improvements over traditional unsupervised domain adaptation models as well as state-of-the-art NAS-based methods.

2 RELATED WORK

Deep convolutional neural network has been dominating image recognition task. In recent years, many handcrafted architectures have been proposed, including VGG (Simonyan & Zisserman, 2014), ResNet (He et al., 2016), Inception (Szegedy et al., 2015), etc., all of which verifies the importance of human expertise in network design. Our work bridges domain adaptation and the emerging field of neural architecture search (NAS), a process of automating architecture engineering technique.

Neural Architecture Search Neural Architecture Search has become the mainstream approach to discover efficient and powerful network structures (Zoph & Le, 2017; Zoph et al., 2018). The automatically searched architectures have achieved highly competitive performance in tasks such as image classification (Liu et al., 2018c;b), object detection (Zoph et al., 2018), and semantic segmentation (Chen et al., 2018). Reinforce learning based NAS methods (Zoph & Le, 2017; Tan et al., 2019; Tan & Le, 2019) are usually computational intensive, thus hampering its usage with limited computational budget. To accelerate the search procedure, many techniques has been proposed and they mainly follow four directions: (1) estimating the actual performance with *lower fidelities*. Such lower fidelities include shorter training times (Zoph et al., 2018; Zela et al., 2018), training on a subset of the data (Klein et al., 2017), or on lower-resolution images. (2) estimating the performance based on the *learning curve extrapolation*. Domhan et al. (2015) propose to extrapolate initial learning curves and terminate those predicted to perform poorly. (3) *initializing* the novel architectures based on other well-trained architectures. Wei et al. (2016) introduce network morphisms to modify an architecture without changing the network objects, resulting in methods that only require a few GPU days (Elsken et al., 2017; Cai et al., 2018a; Jin et al., 2019; Cai et al., 2018b). (4) *one-shot* architecture search. One-shot NAS treats all architectures as different subgraphs of a supergraph and shares weights between architectures that have edges of this supergraph in common (Saxena & Verbeek, 2016; Liu et al., 2019b; Bender, 2018). DARTS (Liu et al., 2019a) places a mixture of candidate operations on each edge of the one-shot model and optimizes the weights of the candidate operations with a continuous relaxation of the search space. Inspired by DARTS (Liu et al., 2019a), our model employs differentiable architecture search to derive the optimal feature extractor for unsupervised domain adaptation.

Domain Adaptation Unsupervised domain adaptation (UDA) aims to transfer the knowledge learned from one or more labeled source domains to an unlabeled target domain. Various methods have been proposed, including discrepancy-based UDA approaches (Long et al., 2017; Tzeng et al., 2014; Ghifary et al., 2014; Peng & Saenko, 2018), adversary-based approaches (Liu & Tuzel, 2016; Tzeng et al., 2017; Liu et al., 2018a), and reconstruction-based approaches (Yi et al., 2017; Zhu et al., 2017; Hoffman et al., 2018; Kim et al., 2017). These models are typically designed to tackle single source to single target adaptation. Compared with single source adaptation, multi-source domain adaptation (MSDA) assumes that training data are collected from multiple sources. Originating from the theoretical analysis in (Ben-David et al., 2010; Mansour et al., 2009; Crammer et al., 2008), MSDA has been applied to many practical applications (Xu et al., 2018; Duan et al., 2012; Peng et al., 2019). Specifically, Ben-David et al. (2010) introduce an $\mathcal{H}\Delta\mathcal{H}$ -divergence between the weighted combination of source domains and a target domain. These models are developed using the existing hand-crafted network architecture. This property limits the capacity and versatility of domain adaptation as the backbones to extract the features are fixed. In contrast, we tackle the UDA from a different perspective, not yet considered in the UDA literature. We propose a novel dual-objective model of NASDA, which optimize neural architecture for unsupervised domain adaptation. We are the first to show the effectiveness of MK-MMD in NAS process which is designed specifically for domain adaptation.

3 NEURAL ARCHITECTURE SEARCH FOR DOMAIN ADAPTATION

In unsupervised domain adaptation, we are given a source domain $\mathcal{D}_s = \{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}_{i=1}^{n_s}$ of n_s labeled examples and a target domain $\mathcal{D}_t = \{\mathbf{x}_j^t\}_{j=1}^{n_t}$ of n_t unlabeled examples. The source domain and target domain are sampled from joint distributions $P(\mathbf{x}^s, \mathbf{y}^s)$ and $Q(\mathbf{x}^t, \mathbf{y}^t)$, respectively. The goal of this paper is to leverage NAS to derive a deep network $G: \mathbf{x} \mapsto \mathbf{y}$, which is optimal for reducing the shifts in data distributions across domains, such that the target risk $\epsilon_t(G) = \mathbb{E}_{(\mathbf{x}^t, \mathbf{y}^t) \sim Q} [G(\mathbf{x}^t) \neq \mathbf{y}^t]$ is minimized. We will start by introducing some preliminary background in Section 3.1. We then describe how to incorporate the MK-MMD into the neural architecture searching framework in

Section 3.2. Finally, we introduce the adversarial training between our derived deep network and a batch of classifiers in Section 3.3. An overview of our model can be seen in Algorithm 1.

3.1 PRELIMINARY: DARTS

In this work, we leverage DARTS (Liu et al., 2019a) as our baseline framework. Our goal is to search for a robust cell and apply it to a network that is optimal to achieve domain alignment between \mathcal{D}_s and \mathcal{D}_t . Following Zoph et al. (2018), we search for a computation cell as the building block of the final architecture. The final convolutional network for domain adaptation can be stacked from the learned cell. A cell is defined as a directed acyclic graph (DAG) of L nodes, $\{x^i\}_{i=1}^N$, where each node $x^{(i)}$ is a latent representation and each directed edge $e^{(i,j)}$ is associated with some operation $o^{(i,j)}$ that transforms $x^{(i)}$. DARTS (Liu et al., 2019a) assumes that cells contain two input nodes and a single output node. To make the search space continuous, DARTS relaxes the categorical choice of a particular operation to a softmax over all possible operations and is thus formulated as:

$$\bar{o}^{(i,j)}(x) = \sum_{o \in \mathcal{O}} \frac{\exp(\alpha_o^{(i,j)})}{\sum_{o' \in \mathcal{O}} \exp(\alpha_{o'}^{(i,j)})} o(x) \quad (3)$$

where \mathcal{O} denotes the set of candidate operations and $i < j$ so that *skip-connect* can be applied. An intermediate node can be represented as $x_j = \sum_{i < j} o^{(i,j)}(x_i)$. The task of architecture search then reduces to learning a set of continuous variables $\alpha = \{\alpha^{(i,j)}\}$. At the end of search, a discrete architecture can be obtained by replacing each mixed operation $\bar{o}^{(i,j)}$ with the most likely operation, i.e., $o^{*(i,j)} = \operatorname{argmax}_{o \in \mathcal{O}} \alpha_o^{(i,j)}$ and $\alpha^* = \{o^{*(i,j)}\}$.

3.2 SEARCHING NEURAL ARCHITECTURE

Denote by \mathcal{L}_{train} and \mathcal{L}_{val} the training loss and validation loss, respectively. Conventional neural architecture search models aim to derive $\Phi_{\alpha,w}$ by solving a bilevel optimization problem (Anandalingam & Friesz, 1992): $\Phi_{\alpha,w} = \operatorname{argmin}_{\alpha} \mathcal{L}_{val}(w^*(\alpha), \alpha)$ s.t. $w^*(\alpha) = \operatorname{argmin}_w \mathcal{L}_{train}(w, \alpha)$. While recent work (Zoph et al., 2018; Liu et al., 2018c) have show promising performance on tasks such as image classification and object detection, the existing models assume that the training data and testing data are sampled from the same distributions. Our goal is to jointly learn the architecture α and the weights w within all the mixed operations (e.g. weights of the convolution filters) so that the derived model Φ_{w^*,α^*} can transfer knowledge from \mathcal{D}_s to \mathcal{D}_t with some simple domain adaption guidance. Initialized by Equation 1, we leverage multi-kernel Maximum Mean Discrepancy (Gretton et al., 2007) to evaluate $\operatorname{disc}(\Phi^*(\mathbf{x}^s), \Phi^*(\mathbf{x}^t))$.

MK-MMD Denote by \mathcal{H}_k be the Reproducing Kernel Hilbert Space (RKHS) endowed with a characteristic kernel k . The *mean embedding* of distribution p in \mathcal{H}_k is a unique element $\mu_k(P)$ such that $\mathbf{E}_{\mathbf{x} \sim P} f(\mathbf{x}) = \langle f(\mathbf{x}), \mu_k(P) \rangle_{\mathcal{H}_k}$ for all $f \in \mathcal{H}_k$. The MK-MMD $d_k(P, Q)$ between probability distributions P and Q is defined as the RKHS distance between the mean embeddings of P and Q . The squared formulation of MK-MMD is defined as

$$d_k^2(P, Q) \triangleq \|\mathbf{E}_P[\Phi_{\alpha}(\mathbf{x}^s)] - \mathbf{E}_Q[\Phi_{\alpha}(\mathbf{x}^t)]\|_{\mathcal{H}_k}^2. \quad (4)$$

In this paper, we consider the case of combining Gaussian kernels with injective functions f_{Φ} , where $k(x, x') = \exp(-\|f_{\Phi}(x) - f_{\Phi}(x')\|^2)$. Inspired by Long et al. (2015), the characteristic kernel associated with the feature map Φ , $k(\mathbf{x}^s, \mathbf{x}^t) = \langle \Phi(\mathbf{x}^s), \Phi(\mathbf{x}^t) \rangle$, is defined as the convex combination of n positive semidefinite kernels $\{k_u\}$,

$$\mathcal{K} \triangleq \left\{ k = \sum_{u=1}^n \beta_u k_u : \sum_{u=1}^n \beta_u = 1, \beta_u \geq 0, \forall u \right\}, \quad (5)$$

where the constraints on $\{\beta_u\}$ are imposed to guarantee that the k is characteristic. In practice we use finite samples from distributions to estimate MMD distance. Given $\mathbf{X}^s = \{\mathbf{x}_1^s, \dots, \mathbf{x}_m^s\} \sim P$ and $\mathbf{X}^t = \{\mathbf{x}_1^t, \dots, \mathbf{x}_m^t\} \sim Q$, one estimator of $d_k^2(P, Q)$ is

$$\hat{d}_k^2(P, Q) = \frac{1}{\binom{m}{2}} \sum_{i \neq i'} k(\mathbf{x}_i^s, \mathbf{x}_{i'}^s) - \frac{2}{\binom{m}{2}} \sum_{i \neq j} k(\mathbf{x}_i^s, \mathbf{x}_j^t) + \frac{1}{\binom{m}{2}} \sum_{j \neq j'} k(\mathbf{x}_j^t, \mathbf{x}_{j'}^t). \quad (6)$$

Algorithm 1 Neural Architecture Search for Domain Adaptation**Phase I: Searching Neural Architecture**

- 1: Create a mixed operation $o^{(i,j)}$ parametrized by $\alpha^{(i,j)}$ for each edge (i, j)
- 2: **while** not converged **do**
- 3: Update architecture α by $\frac{\partial}{\partial \alpha} \mathcal{L}_{val}^s \left(w - \xi \frac{\partial}{\partial w} \mathcal{L}_{train}^s(w, \alpha), \alpha \right) + \lambda \frac{\partial}{\partial \alpha} \left(\hat{d}_k^2(\Phi(\mathbf{x}^s), \Phi(\mathbf{x}^t)) \right)$
- 4: Update weights w by descending $\frac{\partial}{\partial w} \mathcal{L}_{train}^s(w, \alpha)$
- 5: **end while**
- 6: Derive the final architecture based on the learned α^* .

Phase II: Adversarial Training for Domain Adaptation

- 1: Stack feature generator G based on α^* , initialize classifiers C
- 2: **while** not converged **do**
- 3: Step one: Train G and C with $\mathcal{L}^s(\mathbf{x}^s, \mathbf{y}^s) = -\mathbb{E}_{(\mathbf{x}^s, \mathbf{y}^s) \sim \mathcal{D}^s} \sum_{k=1}^K \mathbf{1}_{[k=\mathbf{y}^s]} \log p(\mathbf{y}^s | \mathbf{x}^s)$
- 4: Step two: Fix G , train C with loss: $\mathcal{L}^s(\mathbf{x}^s, \mathbf{y}^s) - \mathcal{L}_{adv}(\mathbf{x}^t)$ (Eq. 13)
- 5: Step three: Fix C , train G with loss: $\mathcal{L}_{adv}(\mathbf{x}^t)$
- 6: **end while**

The merit of multi-kernel MMD lies in its differentiability such that it can be easily incorporated into the deep network. However, the computation of the $\hat{d}_k^2(P, Q)$ incurs a complexity of $O(m^2)$, which is undesirable in the differentiable architecture search framework. In this paper, we use the unbiased estimation of MK-MMD (Gretton et al., 2012) which can be computed with linear complexity.

NAS for Domain Adaptation Denote by \mathcal{L}_{train}^s and \mathcal{L}_{val}^s the training loss and validation loss on the source domain, respectively. Both losses are affected by the architecture α as well as by the weights w in the network. The goal for NASDA is to find α^* that minimizes the validation loss $\mathcal{L}_{val}^t(w^*, \alpha^*)$ on the target domain, where the weights w^* associated with the architecture are obtained by minimizing the training loss $w^* = \operatorname{argmin}_w \mathcal{L}_{train}^s(w, \alpha^*)$. Due to the lack of labels in the target domain, it is prohibitive to compute \mathcal{L}_{val}^t directly, hampering the assumption of previous gradient-based NAS algorithms (Liu et al., 2019a; Chen et al., 2019). Instead, we derive α^* by minimizing the validation loss $\mathcal{L}_{val}^s(w^*, \alpha^*)$ on the source domain plus the domain discrepancy, $\operatorname{disc}(\Phi(\mathbf{x}^s), \Phi(\mathbf{x}^t))$, as shown in Equation 1.

Inspired by the gradient-based hyperparameter optimization (Franceschi et al., 2018; Pedregosa, 2016; Maclaurin et al., 2015), we set the architecture parameters α as a special type of hyperparameter. This implies a bilevel optimization problem (Anandalingam & Friesz, 1992) with α as the upper-level variable and w as the lower-level variable. In practice, we utilize the MK-MMD to evaluate the domain discrepancy. The optimization can be summarized as follows:

$$\Phi_{\alpha, w} = \operatorname{argmin}_{\alpha} \left(\mathcal{L}_{val}^s(w^*(\alpha), \alpha) + \lambda \hat{d}_k^2(\Phi(\mathbf{x}^s), \Phi(\mathbf{x}^t)) \right) \quad (7)$$

$$\text{s.t. } w^*(\alpha) = \operatorname{argmin}_w \mathcal{L}_{train}^s(w, \alpha) \quad (8)$$

where λ is the trade-off hyperparameter between the source validation loss and the MK-MMD loss.

Approximate Architecture Search Equation 7,8 imply that directly optimizing the architecture gradient is prohibitive due to the expensive inner optimization. Inspired by DARTS (Liu et al., 2019a), we approximate $w^*(\alpha)$ by adapting w using only a single training step, without solving the optimization in Equation 8 by training until convergence. This idea has been adopted and proven to be effective in meta-learning for model transfer (Finn et al., 2017), gradient-based hyperparameter tuning (Luketina et al., 2016) and unrolled generative adversarial networks. We therefore propose a simple approximation scheme as follows:

$$\begin{aligned} & \frac{\partial}{\partial \alpha} \left(\mathcal{L}_{val}^s(w^*(\alpha), \alpha) + \lambda \hat{d}_k^2(\Phi(\mathbf{x}^s), \Phi(\mathbf{x}^t)) \right) \\ & \approx \frac{\partial}{\partial \alpha} \mathcal{L}_{val}^s \left(w - \xi \frac{\partial}{\partial w} \mathcal{L}_{train}^s(w, \alpha), \alpha \right) + \lambda \frac{\partial}{\partial \alpha} \left(\hat{d}_k^2(\Phi(\mathbf{x}^s), \Phi(\mathbf{x}^t)) \right) \end{aligned} \quad (9)$$

where $w - \xi \frac{\partial}{\partial w} \mathcal{L}_{train}^s(w, \alpha)$ denotes weight for one-step forward model and ξ is the learning rate for a step of inner optimization. Note Equation 9 reduces to $\nabla_{\alpha} \mathcal{L}_{val}^s(w, \alpha)$ if w is already a local optimum for the inner optimization and thus $\nabla_w \mathcal{L}_{train}^s(w, \alpha) = 0$.

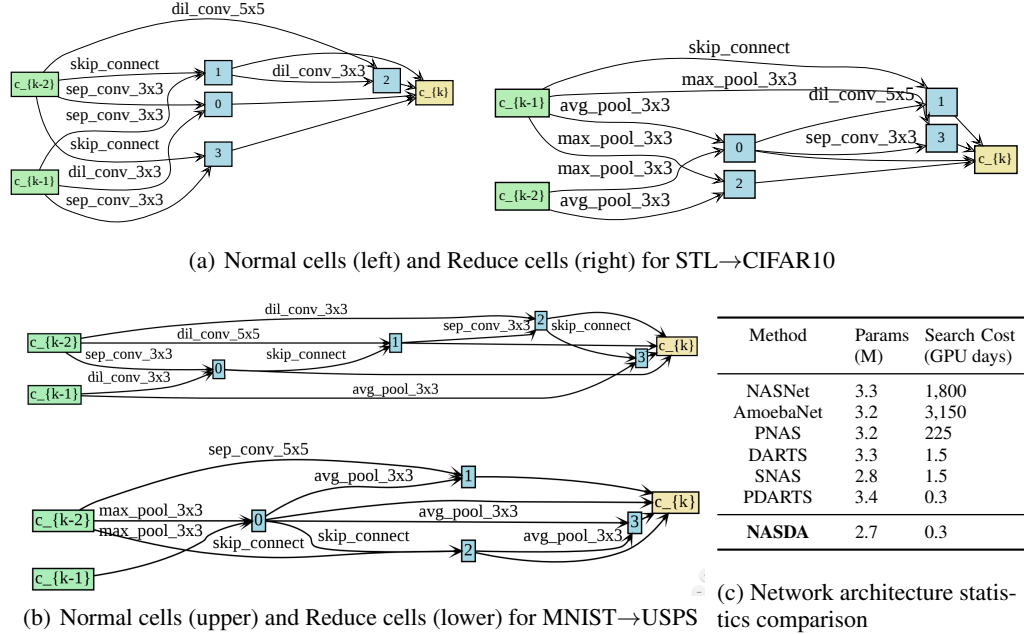


Figure 2: (a) Neural architecture for STL→CIFAR10 task. (b) Neural architecture for MNIST→USPS results. (c) Comparison between our NASDA model and state-of-the-art NAS models.

The second term of Equation 9 can be computed directly with some forward and backward passes. For the first term, applying chain rule to the approximate architecture gradient yields

$$\frac{\partial}{\partial \alpha} \mathcal{L}_{val}^s(w', \alpha) - \xi \left(\frac{\partial^2}{\partial \alpha \partial w} \mathcal{L}_{train}^s(w, \alpha) \frac{\partial}{\partial w'} \mathcal{L}_{val}^s(w', \alpha) \right) \quad (10)$$

where $w' = w - \xi \frac{\partial}{\partial w} \mathcal{L}_{train}^s(w, \alpha)$. The expression above contains an expensive matrix-vector product in its second term. We leverage the central difference approximation to reduce the computation complexity. Specifically, let η be a small scalar and $w^\pm = w \pm \eta \frac{\partial}{\partial w'} \mathcal{L}_{val}^s(w', \alpha)$. Then:

$$\frac{\partial^2}{\partial \alpha \partial w} \mathcal{L}_{train}^s(w, \alpha) \frac{\partial}{\partial w'} \mathcal{L}_{val}^s(w', \alpha) \approx \frac{\frac{\partial}{\partial \alpha} \mathcal{L}_{train}^s(w^+, \alpha) - \frac{\partial}{\partial \alpha} \mathcal{L}_{train}^s(w^-, \alpha)}{2\eta} \quad (11)$$

Evaluating the central difference only requires two forward passes for the weights and two backward passes for α , reducing the complexity from quadratic to linear.

3.3 ADVERSARIAL TRAINING FOR DOMAIN ADAPTATION

By neural architecture searching from Section 3.2, we have derived the optimal cell structure (α^*) for domain adaptation. We then stack the cells to derive our feature generator G . In this section, we describe how do we consolidate G by an adversarial training of G and the classifiers C . Assume C includes N independent classifiers $\{C^{(i)}\}_{i=1}^N$ and denote $p_i(\mathbf{y}|\mathbf{x})$ as the K -way probabilistic outputs of $C^{(i)}$, where K is the category number.

The high-level intuition is to consolidate the feature generator G such that it can make the diversified C generate similar outputs. To this end, our training process include three steps: (1) train G and C on \mathcal{D}_s to obtain task-specific features, (2) fix G and train C to make $\{C^{(i)}\}_{i=1}^N$ have diversified output, (3) fix C and train G to minimize the output discrepancy between C . Related techniques have been used in Saito et al. (2018); Kumar et al. (2018).

First, we train both G and C to classify the source samples correctly with cross-entropy loss. This step is crucial as it enables G and C to extract the task-specific features. The training objective is $\min_{G,C} \mathcal{L}^s(\mathbf{x}^s, \mathbf{y}^s)$ and the loss function is defined as follows:

$$\mathcal{L}^s(\mathbf{x}^s, \mathbf{y}^s) = -\mathbb{E}_{(\mathbf{x}^s, \mathbf{y}^s) \sim \mathcal{D}^s} \sum_{k=1}^K \mathbf{1}_{[k=\mathbf{y}^s]} \log p(\mathbf{y}^s | \mathbf{x}^s) \quad (12)$$

In the second step, we are aiming to diversify C . To establish this goal, we fix G and train C to increase the discrepancy of C 's output. To avoid mode collapse (e.g. $C^{(1)}$ outputs all zeros and $C^{(2)}$ output all ones), we add $\mathcal{L}^s(\mathbf{x}^s, \mathbf{y}^s)$ as a regularizer in the training process. The high-level intuition is that we do not expect C to forget the information learned in the first step in the training process. The training objective is $\min_C \mathcal{L}^s(\mathbf{x}^s, \mathbf{y}^s) - \mathcal{L}_{\text{adv}}(\mathbf{x}^t)$, where the adversarial loss is defined as:

$$\mathcal{L}_{\text{adv}}(\mathbf{x}^t) = \mathbb{E}_{\mathbf{x}^t \sim \mathcal{D}^t} \sum_{i=1}^{N-1} \sum_{j=i+1}^N \|(p_i(\mathbf{y}|\mathbf{x}^t) - p_j(\mathbf{y}|\mathbf{x}^t))\|_1 \quad (13)$$

In the last step, we are trying to consolidate the feature generator G by training G to extract generalizable representations such that the discrepancy of C 's output is minimized. To achieve this goal, we fix the diversified classifiers C and train G with the adversarial loss (defined in Equation 13). The training objective is $\min_G \mathcal{L}_{\text{adv}}(\mathbf{x}^t)$.

4 EXPERIMENTS

We compare the proposed NASDA model with many state-of-the-art UDA baselines on multiple benchmarks. In the main paper, we only report major results; more details are provided in the supplementary material. All of our experiments are implemented in the PyTorch platform.

In the architecture search phase, we use $\lambda=1$ for all the searching experiments. We leverage the ReLU-Conv-BN order for convolutional operations, and each separable convolution is always applied twice. Our search space \mathcal{O} includes the following operations: 3×3 and 5×5 separable convolutions, 3×3 and 5×5 dilated separable convolutions, 3×3 max pooling, identity, and *zero*. Our convolutional cell consists of $N = 7$ nodes. Cells located at the $\frac{1}{3}$ and $\frac{2}{3}$ of the total depth of the network are reduction cells. The architecture encoding therefore is $(\alpha_{\text{normal}}, \alpha_{\text{reduce}})$, where α_{normal} is shared by all the normal cells and α_{reduce} is shared by all the reduction cells.

4.1 SETUP

Digits We investigate three digits datasets: **MNIST**, **USPS**, and Street View House Numbers (**SVHN**). We adopt the evaluation protocol of CyCADA (Hoffman et al., 2018) with three transfer tasks: USPS to MNIST (**U** \rightarrow **M**), MNIST to USPS (**M** \rightarrow **U**), and SVHN to MNIST (**S** \rightarrow **M**). We train our model using the training sets: MNIST (60,000), USPS (7,291), standard SVHN train (73,257).

STL \rightarrow **CIFAR10** Both CIFAR10 (Krizhevsky et al., 2009) and STL (Coates et al., 2011) are both 10-class image datasets. These two datasets contain nine overlapping classes. We remove the ‘frog’ class in CIFAR10 and the ‘monkey’ class in STL datasets as they have no equivalent in the other dataset, resulting in a 9-class problem. The STL images were down-scaled to 32×32 resolution to match that of CIFAR10.

SYN SIGNS \rightarrow **GTSRB** We evaluated the adaptation from synthetic traffic sign dataset called SYN SIGNS (Moiseev et al., 2013) to real-world sign dataset called GTSRB (Stallkamp et al., 2011). These datasets contain 43 classes.

We compare our NASDA model with state-of-the-art DA methods: Deep Adaptation Network (**DAN**) (Long et al., 2015), Domain Adversarial Neural Network (**DANN**) (Ganin & Lempitsky, 2015), Domain Separation Network (**DSN**) (Bousmalis et al., 2016), Coupled Generative Adversarial Networks (**CoGAN**) (Liu & Tuzel, 2016), Maximum Classifier Discrepancy (**MCD**) (Saito et al., 2018), Generate to Adapt (**G2A**) (Sankaranarayanan et al., 2018), Stochastic Neighborhood Embedding (**d-SNE**) (Xu et al., 2019b), Associative Domain Adaptation (**ASSOC**) (Haeusser et al., 2017).

Method	STL \rightarrow CIFAR10
DANN	56.9
MCD	69.2
DWT	71.2
SE	74.2
G2A	72.8
VADA	73.5
DIRT-T	<u>75.3</u>
NASNet+Phase II	67.3
AmoebaNet+Phase II	67.0
DARTS+Phase II	68.8
PDARTS+Phase II	66.0
NASDA	76.8

Table 1: Accuracy (%) STL \rightarrow CIFAR-10

Method	M \rightarrow U	U \rightarrow M	S \rightarrow M	Avg(digits)	Method	SYN SIGNS \rightarrow GTSRB
DAN	81.1	-	71.1	76.1	DAN	91.1
DANN	85.1	73.0	71.1	76.4	DANN	88.7
DSN	91.3*	-	82.7	-	DSN	93.1
CoGAN	91.2*	89.1	-	-	CORAL	86.9
MCD	94.2	94.1	96.2	94.8	MCD	94.4
G2A	95.0	90.8	92.4	92.7	ASSOC	82.8
SBADA-GAN	95.3	97.6	76.1	89.7	SRDA-RAN	93.6
d-SNE	99.0	98.7	96.5	<u>98.1</u>	DADRL	94.6
NASDA	98.0	98.7	98.6	98.4	NASDA	96.7

Table 2: Accuracy (%) on Digits and Traffic Signs for unsupervised domain adaptation.

4.2 EMPIRICAL RESULTS

Neural Architecture Search Results We show the neural architecture search results in Figure 2. We can observe that our model contains more “avg_pool” and “5x5_conv” layer than other NAS model. This will make our model more generic as the average pooling method smooths out the image and hence the model is not congested with sharp features and domain-specific features. We also show that our NASDA model contains less parameters and takes less time to converge compared with state-of-the-art NAS architectures. Another interesting finding is that our NASDA contains more sequential connections in both Normal and Reduce cells when trained on MNIST \rightarrow USPS.

Unsupervised Domain Adaptation Results The UDA results for Digits and SYN SIGNS \rightarrow GTSRB are reported in Table 2, with results of baselines directly reported from the original papers if the protocol is the same (numbers with * indicates training on partial data). The NASDA model achieves a **98.4%** average accuracy for Digits dataset, outperforming other baselines. For SYN SIGNS \rightarrow GTSRB task, our model gets comparable results with state-of-the-art baselines. The results demonstrate the effectiveness of our NASDA model on small images.

The UDA results on the STL \rightarrow CIFAR10 recognition task are reported in Table 1. Our model achieves a performance of **76.8%**, outperforming all the baselines. To compare our search neural architecture with previous NAS models, we replace the neural architecture in G with other NAS models. Other training settings in the second phase are identical to our model. As such, we derive **NASNet+Phase II** (Zoph et al., 2018), **AmoebaNet+Phase II** (Shah et al., 2018), **DARTS+Phase II** (Liu et al., 2019a), and **PDARTS+Phase II** (Chen et al., 2019) models. The results in Table 1 demonstrate that our model outperform other NAS based model by a large margin, which shows the effectiveness of our model in unsupervised domain adaptation. Specifically, we set **PDARTS+Phase II** as an ablation study to demonstrate the effectiveness of our task-specific design in learning domain-adaption aware features.

Analysis To dive deeper into the training process of our NASDA model, we plot the T-SNE embedding of the weights of C in USPS \rightarrow MNIST in Figure 3. This is achieved by recording the weights of all the classifiers at each epoch. The black dot indicates epoch zero, which is the common starting point. The color from light to dark corresponds to the epoch number from small to large. The T-SNE plots clearly show that the classifiers are diverged from each other, demonstrating the effectiveness of the second step of our NASDA training described in Section 3.2.




Figure 3: T-SNE for 4 classifiers

5 CONCLUSION

In this paper, we first formulate a novel dual-objective task of Neural Architecture Search for Domain Adaptation (NASDA) to invigorate the design of transfer-aware network architectures. Towards tackling the NASDA task, we have proposed a novel learning framework that leverages MK-MMD to guide the neural architecture search process. Instead of aligning the features from existing handcrafted backbones, our model directly searches for the optimal neural architecture specific for domain adaptation. Furthermore, we have introduced the ways to consolidate the feature generator, which is stacked from the searched architecture, in order to boost the UDA performance. Extensive empirical evaluations on UDA benchmarks have demonstrated the efficacy of the proposed model against several state-of-the-art domain adaptation algorithms.

REFERENCES

- G Anandalingam and Terry L Friesz. Hierarchical optimization: An introduction. *Annals of Operations Research*, 34(1):1–11, 1992.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- Gabriel Bender. Understanding and simplifying one-shot architecture search. *International Conference on Machine Learning*, 2018.
- Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *Advances in neural information processing systems*, pp. 343–351, 2016.
- Han Cai, Tianyao Chen, Weinan Zhang, Yong Yu, and Jun Wang. Efficient architecture search by network transformation. In *Thirty-Second AAAI conference on artificial intelligence*, 2018a.
- Han Cai, Jiacheng Yang, Weinan Zhang, Song Han, and Yong Yu. Path-level network transformation for efficient architecture search. *International Conference on Machine Learning*, 2018b.
- Liang-Chieh Chen, Maxwell Collins, Yukun Zhu, George Papandreou, Barret Zoph, Florian Schroff, Hartwig Adam, and Jon Shlens. Searching for efficient multi-scale architectures for dense image prediction. In *Advances in neural information processing systems*, pp. 8699–8710, 2018.
- Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1294–1303, 2019.
- Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 215–223, 2011.
- Koby Crammer, Michael Kearns, and Jennifer Wortman. Learning from multiple sources. *Journal of Machine Learning Research*, 9(Aug):1757–1774, 2008.
- Tobias Domhan, Jost Tobias Springenberg, and Frank Hutter. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- Lixin Duan, Dong Xu, and Shih-Fu Chang. Exploiting web images for event recognition in consumer videos: A multiple source domain adaptation approach. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pp. 1338–1345. IEEE, 2012.
- Thomas Elsken, Jan-Hendrik Metzen, and Frank Hutter. Simple and efficient architecture search for convolutional neural networks. *NIPS Workshop on Meta-Learning*, 2017.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1126–1135, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/finn17a.html>.
- Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil. Bilevel programming for hyperparameter optimization and meta-learning. *ICML*, 2018.
- Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1180–1189, Lille, France, 07–09 Jul 2015. PMLR. URL <http://proceedings.mlr.press/v37/ganin15.html>.
- Muhammad Ghifary, W Bastiaan Kleijn, and Mengjie Zhang. Domain adaptive neural networks for object recognition. In *Pacific Rim international conference on artificial intelligence*, pp. 898–904. Springer, 2014.

- Arthur Gretton, Karsten M Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex J Smola. A kernel method for the two-sample-problem. In *Advances in neural information processing systems*, pp. 513–520, 2007.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773, 2012.
- P. Haeusser, T. Frerix, A. Mordvintsev, and D. Cremers. Associative domain adaptation. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2784–2792, 2017.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1989–1998, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL <http://proceedings.mlr.press/v80/hoffman18a.html>.
- Haifeng Jin, Qingquan Song, and Xia Hu. Auto-keras: An efficient neural architecture search system. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1946–1956, 2019.
- Taeksoo Kim, Moon-su Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1857–1865, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/kim17a.html>.
- Aaron Klein, Stefan Falkner, Simon Bartels, Philipp Hennig, and Frank Hutter. Fast Bayesian Optimization of Machine Learning Hyperparameters on Large Datasets. In Aarti Singh and Jerry Zhu (eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 528–536, Fort Lauderdale, FL, USA, 20–22 Apr 2017. PMLR. URL <http://proceedings.mlr.press/v54/klein17a.html>.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Abhishek Kumar, Prasanna Sattigeri, Kahini Wadhawan, Leonid Karlinsky, Rogerio Feris, Bill Freeman, and Gregory Wornell. Co-regularized alignment for unsupervised domain adaptation. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 31*, pp. 9345–9356. Curran Associates, Inc., 2018. URL <http://papers.nips.cc/paper/8146-co-regularized-alignment-for-unsupervised-domain-adaptation.pdf>.
- Alexander H. Liu, Yen-Cheng Liu, Yu-Ying Yeh, and Yu-Chiang Frank Wang. A unified feature disentangler for multi-domain image translation and manipulation. *CoRR*, abs/1809.01361, 2018a. URL <http://arxiv.org/abs/1809.01361>.
- Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 19–34, 2018b.
- Hanxiao Liu, Karen Simonyan, Oriol Vinyals, Chrisantha Fernando, and Koray Kavukcuoglu. Hierarchical representations for efficient architecture search. *ICLR*, 2018c.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *International Conference on Learning Representations*, 2019a.

- Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *International Conference on Learning Representations*, 2019b.
- Ming-Yu Liu and Oncel Tuzel. Coupled generative adversarial networks. In *Advances in neural information processing systems*, pp. 469–477, 2016.
- Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 97–105, Lille, France, 07–09 Jul 2015. PMLR. URL <http://proceedings.mlr.press/v37/long15.html>.
- Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. Deep transfer learning with joint adaptation networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pp. 2208–2217, 2017. URL <http://proceedings.mlr.press/v70/long17a.html>.
- Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pp. 1640–1650, 2018.
- Jelena Luketina, Mathias Berglund, Klaus Greff, and Tapani Raiko. Scalable gradient-based tuning of continuous regularization hyperparameters. In *International conference on machine learning*, pp. 2952–2960, 2016.
- Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning*, pp. 2113–2122, 2015.
- Yishay Mansour, Mehryar Mohri, Afshin Rostamizadeh, and A R. Domain adaptation with multiple sources. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou (eds.), *Advances in Neural Information Processing Systems 21*, pp. 1041–1048. Curran Associates, Inc., 2009.
- Boris Moiseev, Artem Konev, Alexander Chigorin, and Anton Konushin. Evaluation of traffic sign recognition methods trained on synthetically generated data. In Jacques Blanc-Talon, Andrzej Kasinski, Wilfried Philips, Dan Popescu, and Paul Scheunders (eds.), *Advanced Concepts for Intelligent Vision Systems*, pp. 576–583, Cham, 2013. Springer International Publishing. ISBN 978-3-319-02895-8.
- Fabian Pedregosa. Hyperparameter optimization with approximate gradient. 2016.
- Xingchao Peng and Kate Saenko. Synthetic to real adaptation with generative correlation alignment networks. In *2018 IEEE Winter Conference on Applications of Computer Vision, WACV 2018, Lake Tahoe, NV, USA, March 12-15, 2018*, pp. 1982–1991, 2018. doi: 10.1109/WACV.2018.00219. URL <https://doi.org/10.1109/WACV.2018.00219>.
- Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1406–1415, 2019.
- Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence. *Dataset Shift in Machine Learning*. The MIT Press, 2009. ISBN 0262170051, 9780262170055.
- Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. Aging evolution for image classifier architecture search. In *AAAI 2019*, 2019a.
- Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *Proceedings of the aai conference on artificial intelligence*, volume 33, pp. 4780–4789, 2019b.
- Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

- Swami Sankaranarayanan, Yogesh Balaji, Carlos D Castillo, and Rama Chellappa. Generate to adapt: Aligning domains using generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8503–8512, 2018.
- Shreyas Saxena and Jakob Verbeek. Convolutional neural fabrics. In *Advances in Neural Information Processing Systems*, pp. 4053–4061, 2016.
- Syed Asif Raza Shah, Wenji Wu, Qiming Lu, Liang Zhang, Sajith Sasidharan, Phil DeMar, Chin Guok, John Macauley, Eric Pouyoul, Jin Kim, et al. Amoebanet: An sdn-enabled network service for big data science. *Journal of Network and Computer Applications*, 119:70–82, 2018.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel. The German Traffic Sign Recognition Benchmark: A multi-class classification competition. In *IEEE International Joint Conference on Neural Networks*, pp. 1453–1460, 2011.
- Baochen Sun and Kate Saenko. Deep CORAL: correlation alignment for deep domain adaptation. *CoRR*, abs/1607.01719, 2016. URL <http://arxiv.org/abs/1607.01719>.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9, 2015.
- Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 6105–6114, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v97/tan19a.html>.
- Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2820–2828, 2019.
- Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.
- Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Computer Vision and Pattern Recognition (CVPR)*, volume 1, pp. 4, 2017.
- Tao Wei, Changhu Wang, Yong Rui, and Chang Wen Chen. Network morphism. In *International Conference on Machine Learning*, pp. 564–572, 2016.
- Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. *arXiv preprint arXiv:1912.01805*, 2019a.
- Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3964–3973, 2018.
- Xiang Xu, Xiong Zhou, Ragav Venkatesan, Gurumurthy Swaminathan, and Orchid Majumder. d-sne: Domain adaptation using stochastic neighborhood embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2497–2506, 2019b.
- Zili Yi, Hao (Richard) Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *ICCV*, pp. 2868–2876, 2017.
- Arber Zela, Aaron Klein, Stefan Falkner, and Frank Hutter. Towards automated deep learning: Efficient joint neural architecture and hyperparameter search. *ICML 2018 Workshop on AutoML (AutoML 2018)*, 2018.
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.

Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *ICLR*, 2017.

Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8697–8710, 2018.