
Unsupervised Domain Adaptation with Residual Transfer Networks

Mingsheng Long[†]
 Jianmin Wang[†]
 Michael I. Jordan[‡]

MINGSHENG@TSINGHUA.EDU.CN
 JIMWANG@TSINGHUA.EDU.CN
 JORDAN@BERKELEY.EDU

[†]School of Software, Tsinghua National Lab for Information Science and Technology, Tsinghua University, China

[‡]Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA, USA

Abstract

The recent success of deep neural networks relies on massive amounts of labeled data. For a target task where labeled data is unavailable, domain adaptation can transfer a learner from a different source domain. In this paper, we propose a new approach to domain adaptation in deep networks that can simultaneously learn adaptive classifiers and transferable features from labeled data in the source domain and unlabeled data in the target domain. We relax a shared-classifier assumption made by previous methods and assume that the source classifier and target classifier differ by a residual function. We enable classifier adaptation by plugging several layers into the deep network to explicitly learn the residual function with reference to the target classifier. We embed features of multiple layers into reproducing kernel Hilbert spaces (RKHSs) and match feature distributions for feature adaptation. The adaptation behaviors can be achieved in most feed-forward models by extending them with new residual layers and loss functions, which can be trained efficiently using standard back-propagation. Empirical evidence exhibits that the approach outperforms state of art methods on standard domain adaptation datasets.

1. Introduction

Deep neural networks have significantly improved the state of the art for a wide variety of machine-learning problems and applications. Unfortunately, these impressive gains in performance come only when massive amounts of labeled data are available for supervised training. Since manual labeling of sufficient training data for diverse application domains on-the-fly should be prohibitive, for problems short

of labeled data, there is strong incentive to establishing effective algorithms to reduce the labeling consumption, typically by leveraging off-the-shelf labeled data from a different but related source domain that is big enough for training large-scale deep networks. However, this learning paradigm suffers from the shift in data distributions across different domains, which poses a major obstacle in adapting predictive models for the target task. A typical example is to apply an object recognition model trained on manually tagged images to test images under substantial variations in pose, occlusion, or illumination (Torralba & Efros, 2011).

Learning a discriminative model in the presence of a shift between training and test distributions is known as domain adaptation, a special case of transfer learning (Pan & Yang, 2010). A rich line of approaches to domain adaptation have been proposed in the context of both shallow learning and deep learning, which bridge the source and target domains by learning domain-invariant feature representations without using target labels, such that the classifier learned from the source domain can also be applied to the target domain. In particular, recent studies have shown that deep neural networks can learn more transferable features for domain adaptation (Donahue et al., 2014; Yosinski et al., 2014), by **disentangling explanatory factors of variations underlying data samples, and grouping deep features hierarchically in accordance with their relatedness to invariant factors.** The latest advances have been achieved by embedding domain adaptation in the pipeline of deep feature learning to extract domain-invariant representations (Tzeng et al., 2014; Long et al., 2015; Ganin & Lempitsky, 2015; Tzeng et al., 2015).

The previous papers on deep domain adaptation worked under the assumption that the source classifier can be directly transferred to the target domain upon the learned domain-invariant feature representations. This assumption is rather strong as in practical applications, it is often infeasible to check whether the source classifier and target classifier can be shared or not. Hence, we focus in this paper on a more general (and safe) domain adaptation scenario in which the

source classifier and target classifier differ by a small perturbation function. The goal of this paper is to simultaneously learn adaptive classifiers and transferable features from labeled data in the source domain and unlabeled data in the target domain by embedding the adaptations of both classifiers and features in an end-to-end deep architecture.

This work is primarily motivated by He et al. (2015), the winner of the ImageNet ILSVRC 2015 challenge. We thus propose a novel Residual Transfer Network (RTN) approach to domain adaptation in deep networks which can simultaneously learn adaptive classifiers and transferable features, both are complementary to each other and can be combined. We relax the shared-classifier assumption made by previous methods and assume that the source and target classifiers differ by a small residual function. Thus we enable classifier adaptation by plugging several layers into the deep network to explicitly learn the residual function with reference to the target classifier. In this way, the source classifier and target classifier can be bridged tightly in the back-propagation process. The target classifier is made to better fit the target structures by exploiting the low-density separation criterion. We embed features of multiple layers into reproducing kernel Hilbert spaces (RKHSs) and match feature distributions for feature adaptation. We will show that the adaptation behaviors can be achieved in most feed-forward models by extending them with new residual layers and loss functions, which can be trained efficiently using standard back-propagation. Extensive empirical evidence suggests that the proposed RTN approach outperforms state of art methods on standard domain adaptation benchmarks.

The contributions of this paper are summarized as follows.

- (1) We propose a new residual transfer network for domain adaptation, where both classifiers and features are adapted.
- (2) We explore a deep residual learning framework for classifier adaptation, which does not require target labeled data. Our approach is generic since it can be used to add domain adaptation to almost all existing feed-forward architectures.

2. Related Work

This work is related to domain adaptation, a special case of transfer learning (Pan & Yang, 2010), which builds models that can bridge different domains or tasks, explicitly taking the cross-domain discrepancy into account. Transfer learning is to mitigate the burden of manual labeling for machine learning (Pan et al., 2011; Duan et al., 2012a; Zhang et al., 2013; Li et al., 2014; Wang & Schneider, 2014), computer vision (Saenko et al., 2010; Gopalan et al., 2011; Gong et al., 2012; Duan et al., 2012b; Hoffman et al., 2014) and natural language processing (Collobert et al., 2011). It is a consensus that the domain discrepancy in probability distributions of different domains should be formally reduced, either by shallow

learning or deep learning methods. Deep neural networks can learn abstract representations that disentangle different explanatory factors of variations behind data samples (Bengio et al., 2013) and manifest invariant factors underlying different populations that transfer well from original tasks to similar novel tasks (Yosinski et al., 2014). Thus deep neural networks have been explored for domain adaptation (Glorot et al., 2011; Oquab et al., 2013; Hoffman et al., 2014), multimodal and multi-task learning (Collobert et al., 2011; Ngiam et al., 2011; Gupta et al., 2015), where significant performance gains have been witnessed relative to prior shallow transfer learning methods.

However, recent advances show that deep neural networks can learn abstract feature representations that can only reduce, but not remove, the cross-domain discrepancy (Glorot et al., 2011; Tzeng et al., 2014). Dataset shift has posed a bottleneck to the transferability of deep features, resulting in statistically unbounded risk for target tasks (Ben-David et al., 2007; Mansour et al., 2009; Ben-David et al., 2010). Some recent work addresses the aforementioned problem by deep domain adaptation, which bridges the two worlds of deep learning and domain adaptation (Tzeng et al., 2014; Long et al., 2015; Ganin & Lempitsky, 2015; Tzeng et al., 2015). They extend deep convolutional networks (CNNs) to domain adaptation either by adding one or multiple adaptation layers through which the mean embeddings of distributions are matched (Tzeng et al., 2014; Long et al., 2015), or by adding a fully connected subnetwork as a domain discriminator whilst the deep features are learned to confuse the domain discriminator in a domain-adversarial training paradigm (Ganin & Lempitsky, 2015; Tzeng et al., 2015). While performance was significantly improved, these state of the art methods may be restricted by the assumption that under the learned domain-invariant feature representations, the source classifier can be directly transferred to the target domain. In particular, this assumption may not hold when the source classifier and target classifier cannot be shared. As theoretically studied in (Ben-David et al., 2010), when the combined error of the ideal joint hypothesis is large, then there is no single classifier that performs well on both source and target domains, so we cannot find a good target classifier by directly transferring from the source domain.

This work is primarily motivated by He et al. (2015), the winner of the ImageNet ILSVRC 2015 challenge. They present a residual learning framework to ease the training of very deep networks (up to 152 layers), termed *residual nets*. The residual nets explicitly reformulate the layers as learning residual functions $\mathcal{F}(\mathbf{x})$ with reference to the layer inputs \mathbf{x} , instead of directly learning the unreferenced functions $\mathcal{H}(\mathbf{x}) = \mathcal{F}(\mathbf{x}) + \mathbf{x}$ as in traditional plain nets. The method focuses on standard deep learning in which training data and test data are drawn from identical dis-

tributions, hence it cannot be directly applied to domain adaptation. In this paper, we propose to bridge the source classifier $f_S(\mathbf{x})$ and target classifier $f_T(\mathbf{x})$ by the residual layers such that the cross-domain classifier mismatch can be explicitly modeled by the residual functions $\mathcal{F}(\mathbf{x})$ in an end-to-end deep learning architecture. Note that the idea of adapting source classifier to target domain by adding a perturbation function has been studied by (Yang et al., 2007; Duan et al., 2012b). However, these methods require target labeled data to learn the perturbation function, which cannot be applied to unsupervised domain adaptation, the focus of this study. Another crucial distinction is that their perturbation function is defined in the input space \mathbf{x} , while the input to our residual function is the target classifier $f_T(\mathbf{x})$, which can better capture the mismatch between the source and target classifiers. Finally, their source classifiers should be pre-learned in a separated step, and the prediction phase will be inefficient using nonlinear kernel functions.

3. Residual Transfer Networks

In unsupervised domain adaptation, we are provided with a *source* domain $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{n_s}$ of n_s labeled examples and a *target* domain $\mathcal{D}_t = \{\mathbf{x}_i^t\}_{i=1}^{n_t}$ of n_t unlabeled examples. The source domain and target domain are sampled from probability distributions p and q respectively, and note $p \neq q$. The goal of this paper is to craft a deep neural network that enables learning of transfer classifiers $y = f_s(\mathbf{x})$ and $y = f_t(\mathbf{x})$ to bridge the source-target discrepancy, such that the target risk $R_t(f_t) = \Pr_{(\mathbf{x}, y) \sim q} [f_t(\mathbf{x}) \neq y]$ is minimized by leveraging the source domain supervision.

The challenge of unsupervised domain adaptation arises in that the target domain has no labeled data, while the source classifier f_s trained on source domain \mathcal{D}_s cannot be directly applied to the target domain \mathcal{D}_t due to the distribution discrepancy $p(\mathbf{x}, y) \neq q(\mathbf{x}, y)$. The distribution discrepancy may give rise to mismatches in both features and classifiers, i.e. $p(\mathbf{x}) \neq q(\mathbf{x})$ and $f_s(\mathbf{x}) \neq f_t(\mathbf{x})$. Both mismatches should be fixed by the adaptation of features and classifiers to enable effective domain adaptation. Classifier adaptation is more difficult than feature adaptation because the target domain is fully unlabeled. Note that current state of the art deep feature adaptation methods (Long et al., 2015; Ganin & Lempitsky, 2015; Tzeng et al., 2015) generally assume shared classifier on their adaptive deep features. This paper assumes $f_s \neq f_t$ and presents an end-to-end deep learning architecture for joint adaptation of classifiers and features.

3.1. Convolutional Neural Network

We will extend from the breakthrough AlexNet architecture (Krizhevsky et al., 2012) that comprises five convolutional layers (*conv1–conv5*) and three fully connected layers

(*fc6–fc8*), where *conv1–fc7* are feature layers and *fc8* is the classifier layer. Specifically, each fully connected layer ℓ will learn a nonlinear mapping $\mathbf{h}_i^\ell = a^\ell(\mathbf{W}^\ell \mathbf{h}_i^{\ell-1} + \mathbf{b}^\ell)$, where \mathbf{h}_i^ℓ is the ℓ^{th} -layer hidden representation of example \mathbf{x}_i , \mathbf{W}^ℓ and \mathbf{b}^ℓ are the ℓ^{th} -layer weight and bias parameters, and a^ℓ is the ℓ^{th} -layer activation function, taken as rectifier units (ReLU) $a^\ell(\mathbf{x}) = \max(\mathbf{0}, \mathbf{x})$ for the feature layers or softmax units $a^\ell(\mathbf{x}) = e^{\mathbf{x}} / \sum_{j=1}^{|\mathbf{x}|} e^{x_j}$ for classifier layer. Denote by $\mathcal{F} = \{\mathbf{W}^\ell, \mathbf{b}^\ell\}_{\ell=1}^l$ the set of network parameters (l layers in total). The empirical error $E(\mathcal{D}_s; f_s)$ of source-classifier f_s on source data $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{n_s}$ is

$$\min_{f_s \in \mathcal{F}} E(\mathcal{D}_s; f_s) = \frac{1}{n_s} \sum_{i=1}^{n_s} L(f_s(\mathbf{x}_i^s), y_i^s), \quad (1)$$

where $f_s(\mathbf{x}_i^s)$ is the conditional probability that CNN assigns point \mathbf{x}_i^s to label y_i^s , $L(\cdot, \cdot)$ is the cross-entropy loss function $L(f_s(\mathbf{x}_i^s), y_i^s) = -\sum_{j=1}^c 1\{y_i^s = j\} \log f_j^s(\mathbf{x}_i^s)$, c is the number of classes, and $f_j^s(\mathbf{x}_i^s) = e^{h_{ij}^{s,l}} / \sum_{j'} e^{h_{ij'}^{s,l}}$ is the softmax function defined on the l^{th} -layer activation $\mathbf{h}_i^{s,l}$ that computes the probability of predicting point \mathbf{x}_i^s to class j . Based on the quantification study of feature transferability (Yosinski et al., 2014), **the convolutional layers *conv1–conv5* can learn generic features transferable across domains** (Yosinski et al., 2014). Hence, when adapting the pre-trained AlexNet model from the source domain to the target domain, we opt to fine-tune *conv1–conv5* such that the efficacy of feature co-adaptation can be preserved. **As we perform residual transfer and distribution matching only for fully connected layers and pooling layers, we will not elaborate the computational details of convolutional layers.**

3.2. Classifier Adaptation

Although the source classifier and target classifier are different, $f_s(\mathbf{x}) \neq f_t(\mathbf{x})$, they should be related to ensure the feasibility of domain adaptation. Hence it is reasonable to assume that $f_s(\mathbf{x})$ and $f_t(\mathbf{x})$ differ only by a small perturbation function Δf . Previous works (Yang et al., 2007; Duan et al., 2012b) assume that $f_t(\mathbf{x}) = f_s(\mathbf{x}) + \Delta f(\mathbf{x})$, where the perturbation $\Delta f(\mathbf{x})$ is a function of input feature \mathbf{x} . Unfortunately, these methods require target labeled data to learn the perturbation function, which cannot be applied to unsupervised domain adaptation, the focus of this study. We argue that the difficulty arises because no constraint is imposed to the perturbation function $\Delta f(\mathbf{x})$, therefore it is independent on the source labeled data and source classifier $f_s(\mathbf{x})$ so that it must be learned from target labeled data. How to bridge $f_s(\mathbf{x})$ and $f_t(\mathbf{x})$ in a tight way such that the perturbation Δf can be constrained and learned effectively is a critical challenge for unsupervised domain adaptation.

We are motivated by the deep residual learning framework presented by He et al. (2015) to win the ImageNet ILSVRC

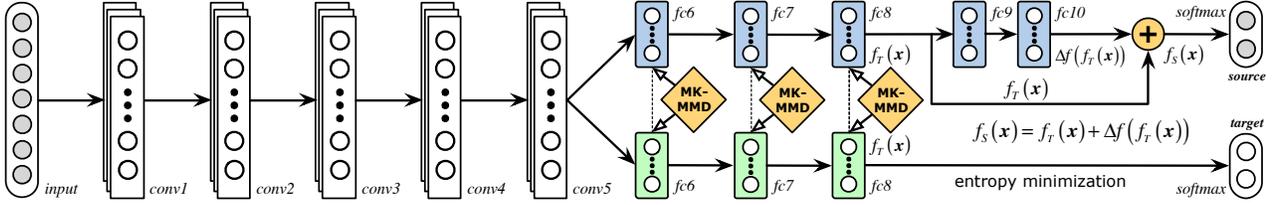


Figure 1. The *Residual Transfer Network* (RTN) architecture based on AlexNet for unsupervised domain adaptation. Since deep features eventually transition from general to specific along the network: (1) fully connected layers $fc6$ – $fc8$ are tailored to model task-specific structures, hence they are not safely transferable and should be adapted with MK-MMD minimization (Long et al., 2015); (2) supervised classifiers are not safely transferable, hence they are bridged by the residual layers $fc9$ – $fc10$ such that $f_S(\mathbf{x}) = f_T(\mathbf{x}) + \Delta f(f_T(\mathbf{x}))$.

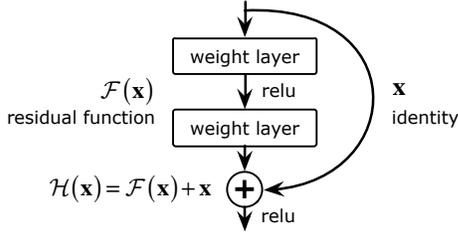


Figure 2. A building block for residual learning (He et al., 2015). Instead of using the residual block to model the feature mapping, we use it to bridge the source classifier $f_S(\mathbf{x})$ and target classifier $f_T(\mathbf{x})$ by $\mathbf{x} \triangleq f_T(\mathbf{x})$, $\mathcal{H}(\mathbf{x}) \triangleq f_S(\mathbf{x})$, and $\mathcal{F}(\mathbf{x}) \triangleq \Delta f(f_T(\mathbf{x}))$.

2015 challenge. Consider fitting $\mathcal{H}(\mathbf{x})$ as an original mapping by a few stacked layers (convolutional layers or fully connected layers) in Figure 2, where \mathbf{x} denotes the inputs to the first of these layers. If one hypothesizes that multiple nonlinear layers can asymptotically approximate complicated functions, then it is equivalent to hypothesize that they can asymptotically approximate the residual functions, i.e. $\mathcal{H}(\mathbf{x}) - \mathbf{x}$. Thus rather than expect stacked layers to approximate $\mathcal{H}(\mathbf{x})$, one explicitly lets these layers approximate a residual function $\mathcal{F}(\mathbf{x}) \triangleq \mathcal{H}(\mathbf{x}) - \mathbf{x}$, with the original function becoming $\mathcal{F}(\mathbf{x}) + \mathbf{x}$. The operation $\mathcal{F}(\mathbf{x}) + \mathbf{x}$ is performed by a shortcut connection and an element-wise addition, while the residual function is parameterized as $\mathcal{F}(\mathbf{x}; \{\mathbf{W}_\ell\})$ by residual layers within each building block. Although both forms are able to asymptotically approximate the desired functions, the ease of learning is different. **In reality, it is unlikely that identity mappings are optimal, but it should be easier to find the perturbations with reference to an identity mapping, than to learn the function as new.** The residual learning is the key contributor to the successful training of very deep networks (He et al., 2015).

As elaborated above, the deep residual learning framework bridges the inputs and outputs of the residual layers by the shortcut connection (identity mapping) such that $\mathcal{H}(\mathbf{x}) = \mathcal{F}(\mathbf{x}) + \mathbf{x}$, which eases the learning of residual function $\mathcal{F}(\mathbf{x})$ (similar to the perturbation function across the source

and target classifiers). Based on this observation, we extend the CNN architecture (Figure 1) by plugging in the residual block (Figure 2). We formulate the residual block to bridge the source classifier $f_S(\mathbf{x})$ and target classifier $f_T(\mathbf{x})$ by $\mathbf{x} \triangleq f_T(\mathbf{x})$, $\mathcal{H}(\mathbf{x}) \triangleq f_S(\mathbf{x})$, and $\mathcal{F}(\mathbf{x}) \triangleq \Delta f(f_T(\mathbf{x}))$. Note that $f_S(\mathbf{x})$ is the outputs of the source-classifier layer $fc10$ and $f_T(\mathbf{x})$ is the outputs of the target-classifier layer $fc8$, both before softmax activation $\sigma(\cdot)$, hence $f_s(\mathbf{x}) \triangleq \sigma(f_S(\mathbf{x}))$, $f_t \triangleq \sigma(f_T(\mathbf{x}))$. We can bridge the source and target classifiers (before activation) by the residual block as

$$f_s(\mathbf{x}) = f_T(\mathbf{x}) + \Delta f(f_T(\mathbf{x})), \quad (2)$$

where we have used the unactivated classifiers f_S and f_T to ensure the final classifier f_s and f_t will output well-defined probabilities as the softmax classifier in deep neural networks. Both residual layers $fc9$ – $fc10$ are fully connected with $c \times c$ units, where c is the number of classes. We set the source classifier f_S as the outputs of the residual block to guarantee that it can be trained from the source-labeled data by the deep residual learning framework. In other words, if we set f_T as the outputs of the residual block, then we will be unable to learn it successfully as we do not have target labeled data and standard back-propagation will not work. The deep residual learning (He et al., 2015) ensures to output valid classifiers $|\Delta f(f_T(\mathbf{x}))| \ll |f_T(\mathbf{x})| \approx |f_S(\mathbf{x})|$. Note that, the residual learning framework will make the perturbation function $\Delta f(f_T(\mathbf{x}))$ dependent on both the target classifier $f_T(\mathbf{x})$ (the functional dependency) as well as the source classifier $f_S(\mathbf{x})$ (thanks to back-propagation).

Although we cast the classifier adaptation into the residual learning framework, while the residual learning framework guarantees that the target classifier f_t cannot deviate much from the source classifier f_s , we still cannot guarantee that f_t will fit the target-specific structures well. To address this problem, we further exploit the entropy minimization principle (Grandvalet & Bengio, 2004) for classifier adaptation, which favors the low-density separation between classes by minimizing the conditional-entropy $E(\mathcal{D}_t; f_t)$ of the class probability $p(y_i^t = j | \mathbf{x}_i^t; f_t) = f_j^t(\mathbf{x}_i^t)$ on target

data \mathcal{D}_t as

$$\min_{f_t \in \mathcal{F}} E(\mathcal{D}_t; f_t) = \frac{1}{n_t} \sum_{i=1}^{n_t} H(f_t(\mathbf{x}_i^t)), \quad (3)$$

where $f_t(\mathbf{x}_i^t)$ is the conditional probability that CNN assigns unlabeled-point \mathbf{x}_i^t to pseudo-label y_i^t , and $H(\cdot)$ is the conditional-entropy loss function defined as $H(f_t(\mathbf{x}_i^t)) = -\sum_{j=1}^c f_j^t(\mathbf{x}_i^t) \log f_j^t(\mathbf{x}_i^t)$, c is the number of classes, and $f_j^t(\mathbf{x}_i^t)$ is the probability of predicting point \mathbf{x}_i^t to class j . By minimizing the entropy penalty (3), the target classifier f_t is made directly accessible to target-unlabeled data and amend itself to pass through the target low-density regions.

3.3. Feature Adaptation

As classifier adaptation does not necessarily undo the mismatch in the feature distributions, we further perform feature adaptation to make domain adaptation more effective. The literature has revealed that the deep features learned by CNNs can disentangle the explanatory factors of variations underlying data distributions and facilitate knowledge transfer (Oquab et al., 2013; Bengio et al., 2013). **But deep features can only reduce, but not remove, the cross-domain distribution discrepancy** (Yosinski et al., 2014; Tzeng et al., 2014), which thus motivates the state of the art deep feature adaptation methods (Long et al., 2015; Ganin & Lempitsky, 2015; Tzeng et al., 2015). As feature adaptation has been addressed quite well, we adopt the deep adaptation network (DAN) (Long et al., 2015) to fine-tune CNN on labeled examples and require the feature distributions of the source and target to become similar under feature representations of fully connected layers $\mathcal{L} = \{fc6, fc7, fc8\}$, which is realized by minimizing a *multi-layer MK-MMD* penalty as

$$\min_{f_s, f_t \in \mathcal{F}} E(\mathcal{D}_s, \mathcal{D}_t; f_s, f_t, k) = \sum_{\ell \in \mathcal{L}} M_k^2(\mathcal{D}_s^\ell, \mathcal{D}_t^\ell), \quad (4)$$

where $\mathcal{D}_s^\ell \triangleq \{\mathbf{h}_i^{s,\ell}\}$ and $\mathcal{D}_t^\ell \triangleq \{\mathbf{h}_i^{t,\ell}\}$ are ℓ^{th} -layer hidden representations of the source and target points respectively, $M_k^2(\mathcal{D}_s^\ell, \mathcal{D}_t^\ell)$ is the MK-MMD between source and target evaluated using the ℓ^{th} -layer hidden representations, and \mathcal{L} is the indices of layers where MK-MMD penalty is active. Let \mathcal{H}_k be the reproducing kernel Hilbert space (RKHS) with kernel k . The multi-kernel maximum mean discrepancy (MK-MMD) between p and q is (Gretton et al., 2012b)

$$M_k^2(p, q) \triangleq \|\mathbb{E}_p[\phi(\mathbf{x}^s)] - \mathbb{E}_q[\phi(\mathbf{x}^t)]\|_{\mathcal{H}_k}^2, \quad (5)$$

where $\phi(\cdot)$ is the nonlinear feature mapping. An important property is $p = q$ iff $M_k^2(p, q) = 0$ (Gretton et al., 2012a). Characteristic kernel $k(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$ is defined as the convex combination of m PSD kernels $\{k_u\}$, $\mathcal{K} \triangleq$

$\{k = \sum_{u=1}^m \beta_u k_u : \sum_{u=1}^m \beta_u = 1, \beta_u \geq 0, \forall u\}$ where the kernel coefficients $\{\beta_u\}$ can be simply set to $1/m$ or can be learned by multiple kernel learning (Long et al., 2015).

3.4. Residual Transfer Network

Based on the aforementioned analysis, to enable effective unsupervised domain adaptation, we propose the *residual transfer network* (RTN) to be an integration of deep feature learning (1), classifier adaptation (2)–(3), and feature adaptation (4) in an end-to-end deep learning framework as

$$\begin{aligned} \min_{f_s = f_T + \Delta f(f_T)} & \frac{1}{n_s} \sum_{i=1}^{n_s} L(f_s(\mathbf{x}_i^s), y_i^s) \\ & + \frac{\gamma}{n_t} \sum_{i=1}^{n_t} H(f_t(\mathbf{x}_i^t)) \\ & + \lambda \sum_{\ell \in \mathcal{L}} M_k^2(\mathcal{D}_s^\ell, \mathcal{D}_t^\ell), \end{aligned} \quad (6)$$

where γ and λ are the tradeoff parameters for the entropy penalty (3) and multi-layer MK-MMD penalty (4) respectively. The proposed RTN model (6) is able to learn both adaptive classifiers and transferable features. As classifier adaptation proposed in this paper and feature adaptation studied in (Long et al., 2015; Ganin & Lempitsky, 2015) are tailored to adapt different layers of the deep networks, they are well complementing to each other to establish better performance and can be combined wherever necessary.

As training deep CNNs requires a large amount of labeled data that is prohibitive for many domain adaptation applications, we start with the AlexNet model pre-trained on ImageNet 2012 data and fine-tune it as (Long et al., 2015). The training of RTN mainly follows standard back-propagation, including the residual layers for classifier adaptation (He et al., 2015). But the optimization of MK-MMD penalty (4) requires carefully-designed algorithm as detailed in (Long et al., 2015) for linear-time training with back-propagation.

Discussion: The proposed idea of residual transfer is quite general and can be readily extended in several scenarios. In heterogeneous domain adaptation where the feature representations are different across domains (Li et al., 2014), the residual transfer can be applied to the feature layers (e.g. $fc6$ – $fc7$) of the deep network to bridge the source and target representations. (2) In multi-task learning (Chu et al., 2015) that solves multiple potentially correlated tasks together to improve performance of all these tasks by sharing statistic strength, the residual transfer can also be extended to bridge the gap between specific task and the shared task. We will leave these possible extensions to the future work.

4. Experiments

We evaluate the proposed residual transfer network against state of the art transfer learning and deep learning methods on unsupervised domain adaptation benchmarks. Datasets, codes, and configurations will be made available publicly.

4.1. Setup

Office-31 (Saenko et al., 2010) is a standard benchmark for domain adaptation, comprising 4,652 images distributed in 31 classes collected from three distinct domains: *Amazon* (**A**), which contains images downloaded from amazon.com, *Webcam* (**W**) and *DSLR* (**D**), which contain images taken by web camera and digital SLR camera in an office environment with different photographic settings, respectively. We evaluate all methods across three transfer tasks $\mathbf{A} \rightarrow \mathbf{W}$, $\mathbf{D} \rightarrow \mathbf{W}$ and $\mathbf{W} \rightarrow \mathbf{D}$, which are commonly adopted by previous deep transfer learning methods (Donahue et al., 2014; Tzeng et al., 2014; Ganin & Lempitsky, 2015), and across another three transfer tasks $\mathbf{A} \rightarrow \mathbf{D}$, $\mathbf{D} \rightarrow \mathbf{A}$ and $\mathbf{W} \rightarrow \mathbf{A}$ as presented in (Long et al., 2015; Tzeng et al., 2015).

Office-Caltech (Gong et al., 2012) is built by selecting the 10 common categories shared by the *Office-31* and *Caltech-256* (**C**) (Griffin et al., 2007) datasets, and is widely used by previous methods (Long et al., 2013; Sun et al., 2016). We consider all domain combinations of this dataset, and build 12 transfer tasks: $\mathbf{A} \rightarrow \mathbf{W}$, $\mathbf{D} \rightarrow \mathbf{W}$, $\mathbf{W} \rightarrow \mathbf{D}$, $\mathbf{A} \rightarrow \mathbf{D}$, $\mathbf{D} \rightarrow \mathbf{A}$, $\mathbf{W} \rightarrow \mathbf{A}$, $\mathbf{A} \rightarrow \mathbf{C}$, $\mathbf{W} \rightarrow \mathbf{C}$, $\mathbf{D} \rightarrow \mathbf{C}$, $\mathbf{C} \rightarrow \mathbf{A}$, $\mathbf{C} \rightarrow \mathbf{W}$, and $\mathbf{C} \rightarrow \mathbf{D}$. While *Office-31* has more categories and is more difficult for domain adaptation algorithms, *Office-Caltech* provides more transfer tasks to enable an unbiased look at the dataset bias issue (Torralba & Efros, 2011). We adopt DeCAF₇ (Donahue et al., 2014) features for shallow transfer methods and original images for deep transfer methods.

We compare with state of the art transfer and deep learning methods: Transfer Component Analysis (**TCA**) (Pan et al., 2011), Geodesic Flow Kernel (**GFK**) (Gong et al., 2012), Deep Domain Confusion (**DDC**) (Tzeng et al., 2014), Deep Adaptation Network (**DAN**) (Long et al., 2015), and Reverse Gradient (**RevGrad**) (Ganin & Lempitsky, 2015). TCA is a conventional transfer learning method based on MMD-regularized Kernel PCA. GFK is a widely-adopted manifold learning method that interpolates across an infinite number of intermediate subspaces to bridge the source and target. DDC is the first method that maximizes domain invariance by adding to AlexNet an adaptation layer using linear-kernel MMD. DAN learns more transferable features by embedding deep features of multiple task-specific layers to reproducing kernel Hilbert spaces (RKHSs) and matching different distributions optimally by multi-kernel MMD. RevGrad boosts domain adaptation by making source and

target indistinguishable for a discriminative domain classifier through adversarial training (Goodfellow et al., 2014). To go deeper with the efficacy of *residual transfer* with *entropy minimization*, we evaluate several variants of RTN: (1) **RTN-res**, by removing the residual layers from RTN; (2) **RTN-ent**, by removing the entropy penalty from RTN; (3) **RTN-mmd**, by removing the MMD penalty from RTN. Note that RTN-mmd is no longer built upon DAN, and this experiment will suggest that residual transfer of classifiers and MMD adaptation of features are well complementary.

We follow standard evaluation protocols for unsupervised domain adaptation (Saenko et al., 2010; Long et al., 2015). For the *Office-31* dataset, we adopt the sampling protocol (Saenko et al., 2010) that randomly samples the source domain with 20 labeled examples per category for Amazon (**A**) and 8 labeled examples per category for Webcam (**W**) and DSLR (**D**). For the *Office-Caltech* dataset, we adopt the full-sampling protocol (Long et al., 2015) that uses all labeled source examples and all unlabeled target examples. All unlabeled examples in the target domain are used for learning transfer classifiers. We compare the average classification accuracy of each method on five random experiments, and report the standard error of the classification accuracies by different experiments of the same transfer task.

For all methods, we either follow the procedures for model selection explained in their respective papers, or conduct cross-validation on labeled source data if parameter selection strategies are not explained. For MMD-based methods (TCA, DDC, DAN, and RTN), we use Gaussian kernel with bandwidth b set to median pairwise squared distances on training data, i.e. median heuristic (Gretton et al., 2012b). We follow (Long et al., 2015) and use multi-kernel MMD for DAN and RTN, and a family of m Gaussian kernels by varying bandwidth $b_u \in [2^{-8}b, 2^8b]$ with a multiplicative step-size of $2^{1/2}$. We conduct cross-validation on labeled source data to select parameters of RTN (Long et al., 2015). We implement all deep methods based on the **Caffe** (Jia et al., 2014) deep-learning framework, and fine-tune from Caffe-trained models of AlexNet (Krizhevsky et al., 2012) pre-trained on ImageNet (Russakovsky et al., 2014). For RTN, We fine-tune convolutional layers *conv1-conv5* and fully connected layers *fc6-fc7*, and train classifier layer *fc8* and residual layers *fc9-fc10*, all through standard back-propagation. Since the classifier and residual layers are trained from scratch, we set their learning rate to be 10 times that of the lower layers. We use mini-batch SGD with momentum of 0.9 and the learning rate annealing strategy implemented in Caffe, and cross-validate base learning rate between 10^{-5} and 10^{-2} by a multiplicative step-size $10^{1/2}$. To suppress noisy predictions (due to random initialization) that bias the entropy penalty at early stages of the training procedure, we set parameter $\gamma = 0$ within 1000 mini-batch iterations and set it to the cross-

Table 1. Accuracy on *Office-31* dataset under sampling protocol (Saenko et al., 2010) for unsupervised adaptation.

Method	A \rightarrow W	D \rightarrow W	W \rightarrow D	A \rightarrow D	D \rightarrow A	W \rightarrow A	Avg
TCA (Pan et al., 2011)	59.0 \pm 0.4	90.2 \pm 0.2	88.2 \pm 0.4	57.8 \pm 0.4	51.6 \pm 0.5	47.9 \pm 0.5	65.8
GFK (Gong et al., 2012)	58.4 \pm 0.6	93.6 \pm 0.4	91.0 \pm 0.5	58.6 \pm 0.6	52.4 \pm 0.6	46.1 \pm 0.5	66.7
AlexNet (Krizhevsky et al., 2012)	60.4 \pm 0.5	94.0 \pm 0.3	92.2 \pm 0.3	58.5 \pm 0.4	46.0 \pm 0.6	49.0 \pm 0.5	66.7
DDC (Tzeng et al., 2014)	62.7 \pm 0.5	94.1 \pm 0.3	92.9 \pm 0.3	60.0 \pm 0.4	47.3 \pm 0.5	48.8 \pm 0.6	67.6
RevGrad (Ganin & Lempitsky, 2015)	67.3 \pm 0.5	93.7 \pm 0.3	94.0 \pm 0.3	-	-	-	-
DAN (Long et al., 2015)	66.0 \pm 0.5	94.3 \pm 0.3	95.2 \pm 0.3	63.2 \pm 0.4	50.0 \pm 0.5	51.1 \pm 0.5	70.0
RTN-res (no residual)	68.9 \pm 0.5	95.3 \pm 0.3	95.5 \pm 0.2	65.7 \pm 0.4	52.1 \pm 0.6	52.7 \pm 0.6	71.7
RTN-ent (no entropy)	64.5 \pm 0.4	94.2 \pm 0.3	94.2 \pm 0.3	63.0 \pm 0.5	53.8 \pm 0.5	51.7 \pm 0.4	70.2
RTN-mmd (no MMD)	65.6 \pm 0.4	96.1 \pm 0.3	96.0\pm0.3	59.3 \pm 0.4	52.1 \pm 0.4	51.0 \pm 0.3	70.1
RTN (all)	70.2\pm0.6	96.6\pm0.5	95.5 \pm 0.4	66.3\pm0.3	54.9\pm0.6	53.1\pm0.4	72.8

 Table 2. Accuracy on *Office-Caltech* dataset under full protocol (Long et al., 2015) for unsupervised adaptation.

Method	A \rightarrow W	D \rightarrow W	W \rightarrow D	A \rightarrow D	D \rightarrow A	W \rightarrow A	A \rightarrow C	W \rightarrow C	D \rightarrow C	C \rightarrow A	C \rightarrow W	C \rightarrow D	Avg
TCA (Pan et al., 2011)	84.4	96.9	99.4	82.8	90.4	85.6	81.2	75.5	79.6	92.1	88.1	87.9	87.0
GFK (Gong et al., 2012)	89.5	97.0	98.1	86.0	89.8	88.5	76.2	77.1	77.9	90.7	78.0	77.1	85.5
AlexNet (Krizhevsky et al., 2012)	83.1	97.6	100.0	88.3	89.0	83.1	84.0	77.9	81.0	91.3	83.2	89.1	87.3
DDC (Tzeng et al., 2014)	86.1	98.2	100.0	89.0	89.5	84.9	85.0	78.0	81.1	91.9	85.4	88.8	88.2
DAN (Long et al., 2015)	93.8	99.0	100.0	92.4	92.0	92.1	85.1	84.3	82.4	92.0	90.6	90.5	91.2
RTN-res (no residual)	96.1	99.0	100.0	92.8	94.3	93.4	88.0	87.3	82.4	93.5	96.3	91.4	92.9
RTN-ent (no entropy)	94.6	99.0	100.0	93.2	92.8	92.9	85.9	85.1	83.2	92.8	91.4	91.3	92.0
RTN-mmd (no MMD)	92.2	98.9	100.0	90.3	92.6	88.8	86.3	83.7	82.3	93.1	94.3	90.5	91.1
RTN (all)	97.0	98.8	100.0	94.6	95.5	93.1	88.5	88.4	84.3	94.4	96.6	92.9	93.7

validated value afterwards.

4.2. Results and Discussion

The classification accuracy results of unsupervised domain adaptation on the six transfer tasks generated from *Office-31* are shown in Table 1, and the results on the twelve transfer tasks generated from *Office-Caltech* are shown in Table 2. Note that the results of RevGrad under the sampling protocol (Saenko et al., 2010) are directly reported from (Sun et al., 2016), as the original paper only provided results under the full-sampling protocol (Ganin & Lempitsky, 2015). The RTN model based on AlexNet (Figure 1) outperforms all comparison methods on most transfer tasks. In particular, RTN substantially improves the classification accuracy on hard transfer tasks, e.g. $A \rightarrow W$ and $A \rightarrow D$, where the source and target are substantially different, and achieves comparable classification accuracy on easy transfer tasks, $D \rightarrow W$ and $W \rightarrow D$, where source and target are similar (Saenko et al., 2010). The encouraging results suggest that RTN is able to learn more adaptive classifiers as well as more transferable features for effective domain adaptation.

From the results, we can make insightful observations. (1) Standard deep-learning methods (AlexNet) perform comparably with traditional *shallow* transfer-learning methods with deep DeCAF₇ features as input (TCA and GFK). The only difference between these two sets of methods is that AlexNet can take the advantage of supervised fine-tuning on the source-labeled data, while TCA and GFK can take

benefits of their domain adaptation procedures. This result confirms the current practice that supervised fine-tuning is important for transferring source classifier to target domain (Oquab et al., 2013), and sustains the recent discovery that deep neural networks learn abstract feature representation, which can only reduce, but not remove, the cross-domain discrepancy (Yosinski et al., 2014). This reveals that the two worlds of deep learning and domain adaptation are not compatible with each other in the two-step pipeline, which motivates carefully-designed deep adaptation architectures to unify them. (2) Deep-transfer learning methods that reduce the domain discrepancy by domain-adaptive deep networks (DDC, DAN and RevGrad) substantially outperform standard deep learning methods (AlexNet) and traditional shallow transfer-learning methods with deep features as the input (TCA and GFK). This confirms that incorporating domain-adaptation modules to deep networks can improve domain adaptation performance. By adapting source-target distributions in multiple task-specific layers using optimal multi-kernel two-sample matching, DAN performs the best in general among the prior deep-transfer learning methods. (3) The proposed residual transfer network (RTN) method performs the best and sets a new state of the art record on these datasets. Different from all the previous deep-transfer learning methods that only adapt the feature layers of deep neural networks to learn more transferable features, RTN further adapts the classifier layers to bridge the source and target classifiers in an end-to-end residual learning framework, which can correct the classifier mismatch effectively.

To go deeper into the modules of RTN, we show the results

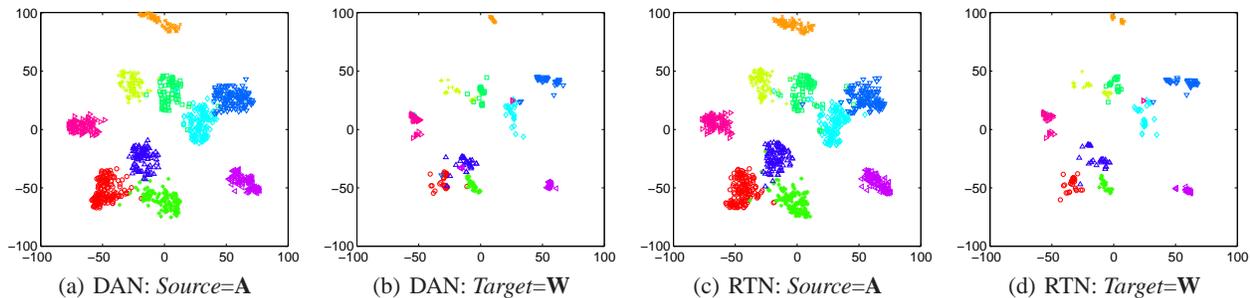


Figure 3. Visualization of classifier predictions: t-SNE of DAN on source (a) and target (b); t-SNE of RTN on source (c) and target (d).

of RTN variants: RTN-res (no residual transfer), RTN-ent (no entropy penalty), and RTN-mmd (no MMD penalty). The results in Tables 1 and 2 well validate our motivation. (1) RTN-res achieves much better results than RevGrad and DAN, but substantially underperforms RTN. This highlights the importance of the *residual transfer* of classifier layers for learning more adaptive classifiers. This is critical as in practical applications, there is no guarantee that the source classifier and target classifier can be safely shared. (2) RTN-ent also achieves much better results than DAN and RevGrad, but substantially underperforms RTN. This highlights the importance of *entropy minimization* for low-density separation, which exploits the cluster structure of target-unlabeled data such that the target-classifier can be better adapted to the target data. Otherwise, the residual function may tend to learn a useless zero mapping such that the source and target classifiers are nearly identical (He et al., 2015). (3) RTN-mmd performs comparably with the best baseline DAN, but substantially underperforms RTN. This evidence suggests the residual transfer of classifiers devised in this paper is as effective as the MMD adaptation of features (Long et al., 2015). Since these two methods are tailored to adapt different layers of the deep networks, they are well complementing each other to establish better performance as witnessed by the best performance of RTN.

4.3. Empirical Analysis

Visualization of Predictions: We illustrate the adaptivity of classifiers by visualizing in Figures 3(a)–3(d) the t-SNE embeddings (Donahue et al., 2014) of the predictions made by the classifiers of DAN and RTN for transfer task $\mathbf{A} \rightarrow \mathbf{W}$, respectively. We can make the following observations. (1) The predictions made by DAN in Figure 3(a)–3(b) show that the target categories are not well discriminated by the source classifier, which implies that target data is not compatible with the source classifier. Hence the source and target classifiers should not be assumed to be identical, which, whatsoever, has been a common assumption made by all prior deep domain adaptation methods (Tzeng et al., 2014;

Long et al., 2015; Ganin & Lempitsky, 2015; Tzeng et al., 2015). (2) The predictions made by RTN in Figures 3(c)–3(d) show that the target categories are discriminated much better by the target classifier, which suggests that residual transfer of classifiers proposed in this paper is a reasonable extension to the previous deep domain adaptation methods. RTN learns more adaptive classifiers as well as more transferable features to enable effective deep domain adaptation.

Analysis of Layer Responses: We illustrate in Figure 4(a) the average magnitudes and standard deviations of the layer responses (He et al., 2015), which are the outputs of $f_T(\mathbf{x})$ ($fc8$ layer), $\Delta f(f_T(\mathbf{x}))$ ($fc10$ layer), and $f_S(\mathbf{x})$ (sum operator) before other nonlinearity (ReLU/addition/softmax), respectively. For the residual transfer network, this analysis exposes the response strength of the residual functions. The results reveal that the residual functions $\Delta f(f_T(\mathbf{x}))$ have generally much smaller responses than the shortcut functions $f_T(\mathbf{x})$. These results support our original motivation that the residual functions might be generally closer to zero than the non-residual functions, since they characterize the small gap between the source classifier and target classifier. The small residual function can be learned more effectively through the residual learning framework (He et al., 2015).

Parameter Sensitivity: Besides the MMD penalty parameter λ as DAN (Long et al., 2015), the RTN model involves another entropy penalty parameter γ . We perform sensitivity analysis for it on transfer tasks $\mathbf{A} \rightarrow \mathbf{W}$ (31 classes) and $\mathbf{C} \rightarrow \mathbf{W}$ (10 classes) by varying the parameter of interest in $\{0.01, 0.04, 0.07, 0.1, 0.4, 0.7, 1.0\}$. The results are shown in Figures 4(b), with the best results of the baseline shown as dashed lines. We observe that the accuracy of RTN first increases and then decreases as γ varies and demonstrates a desirable bell-shaped curve. This justifies our motivation of jointly learning deep features and adapting classifiers using residual layers and entropy penalty in deep nets, as a good trade-off between them can promote transfer performance.

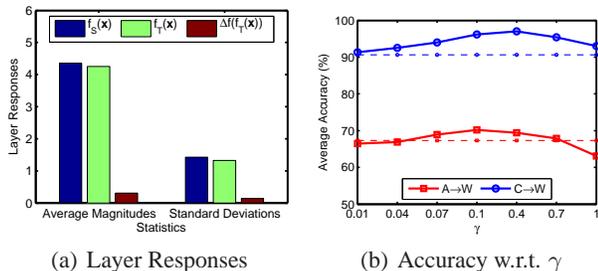


Figure 4. Empirical analysis: (a) analysis of layer responses for RTN; (b) sensitivity of γ (dashed lines show best baseline results).

5. Conclusion

This paper presented a novel approach to unsupervised domain adaptation of deep networks, which enables simultaneous learning of adaptive classifiers and transferable features. Similar to many prior domain adaptation techniques, the feature adaptation is achieved through matching the distributions of features across domains. However, unlike all previous techniques, the proposed approach also supports classifier adaptation, which is accomplished through a new residual transfer module to bridge the source classifier and target classifier. This new ingredient makes the approach a good complement to existing techniques. The approach can be trained by standard back-propagation, which is scalable and can be implemented by any deep learning package. We will constitute further evaluations on larger-scale tasks and semi-supervised domain adaptation settings as future work.

References

- Ben-David, S., Blitzer, J., Crammer, K., and Pereira, F. Analysis of representations for domain adaptation. In *NIPS*, 2007.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. A theory of learning from different domains. *MLJ*, 79(1-2):151–175, 2010.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *TPAMI*, 35(8):1798–1828, 2013.
- Chu, X., Ouyang, W., Yang, W., and Wang, X. Multi-task recurrent neural network for immediacy prediction. In *ICCV*, pp. 3352–3360, 2015.
- Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., and Kuksa, P. Natural language processing (almost) from scratch. *JMLR*, 12:2493–2537, 2011.
- Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., and Darrell, T. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, 2014.
- Duan, L., Tsang, I. W., and Xu, D. Domain transfer multiple kernel learning. *TPAMI*, 34(3):465–479, 2012a.
- Duan, L., Xu, D., Tsang, I. W., and Luo, J. Visual event recognition in videos by learning from web data. *TPAMI*, 34(9):1667–1680, 2012b.
- Ganin, Y. and Lempitsky, V. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015.
- Glorot, X., Bordes, A., and Bengio, Y. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*, 2011.
- Gong, B., Shi, Y., Sha, F., and Grauman, K. Geodesic flow kernel for unsupervised domain adaptation. In *CVPR*, 2012.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *NIPS*, pp. 2672–2680, 2014.
- Gopalan, R., Li, R., and Chellappa, R. Domain adaptation for object recognition: An unsupervised approach. In *ICCV*. IEEE, 2011.
- Grandvalet, Y. and Bengio, Y. Semi-supervised learning by entropy minimization. In *NIPS*, pp. 529–536, 2004.
- Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., and Smola, A. A kernel two-sample test. *JMLR*, 13:723–773, March 2012a.
- Gretton, A., Sriperumbudur, B., Sejdinovic, D., Strathmann, H., Balakrishnan, S., Pontil, M., and Fukumizu, K. Optimal kernel choice for large-scale two-sample tests. In *NIPS*, 2012b.
- Griffin, G., Holub, A., and Perona, P. Caltech-256 object category dataset. Technical report, California Institute of Technology, 2007.
- Gupta, S., Hoffman, J., and Malik, J. Cross modal distillation for supervision transfer. *arXiv preprint arXiv:1507.00448*, 2015.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- Hoffman, J., Guadarrama, S., Tzeng, E., Hu, R., Donahue, J., Girshick, R., Darrell, T., and Saenko, K. LSDA: Large scale detection through adaptation. In *NIPS*, 2014.

- Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., and Darrell, T. Caffe: Convolutional architecture for fast feature embedding. In *MM*, 2014.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- Li, W., Duan, L., Xu, D., and Tsang, I. W. Learning with augmented features for supervised and semi-supervised heterogeneous domain adaptation. *TPAMI*, 36(6):1134–1148, 2014.
- Long, M., Wang, J., Ding, G., Sun, J., and Yu, P. S. Transfer feature learning with joint distribution adaptation. In *ICCV*, 2013.
- Long, M., Cao, Y., Wang, J., and Jordan, M. I. Learning transferable features with deep adaptation networks. In *ICML*, 2015.
- Mansour, Y., Mohri, M., and Rostamizadeh, A. Domain adaptation: Learning bounds and algorithms. In *COLT*, 2009.
- Ngiam, J., Khosla, A., Kim, M., Nam, J., Lee, H., and Ng, A. Y. Multimodal deep learning. In *ICML*, 2011.
- Oquab, M., Bottou, L., Laptev, I., and Sivic, J. Learning and transferring mid-level image representations using convolutional neural networks. In *CVPR*, June 2013.
- Pan, S. J. and Yang, Q. A survey on transfer learning. *TKDE*, 22(10):1345–1359, 2010.
- Pan, S. J., Tsang, I. W., Kwok, J. T., and Yang, Q. Domain adaptation via transfer component analysis. *TNNLS*, 22(2):199–210, 2011.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. 2014.
- Saenko, K., Kulis, B., Fritz, M., and Darrell, T. Adapting visual category models to new domains. In *ECCV*, 2010.
- Sun, B., Feng, J., and Saenko, K. Return of frustratingly easy domain adaptation. In *AAAI*, 2016.
- Torralba, A. and Efros, A. A. Unbiased look at dataset bias. In *CVPR*, 2011.
- Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., and Darrell, T. Deep domain confusion: Maximizing for domain invariance. 2014.
- Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., and Darrell, T. Simultaneous deep transfer across domains and tasks. In *ICCV*, 2015.
- Wang, X. and Schneider, J. Flexible transfer learning under support and model shift. In *NIPS*, 2014.
- Yang, J., Yan, R., and Hauptmann, A. G. Cross-domain video concept detection using adaptive svms. In *MM*, pp. 188–197. ACM, 2007.
- Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. How transferable are features in deep neural networks? In *NIPS*, 2014.
- Zhang, K., Schölkopf, B., Muandet, K., and Wang, Z. Domain adaptation under target and conditional shift. In *ICML*, 2013.