

Correlated Adversarial Joint Discrepancy Adaptation Network

Youshan Zhang and Brian D. Davison

Computer Science and Engineering, Lehigh University, Bethlehem, PA, USA
{yoz217, bdd3}@lehigh.edu

Abstract—Domain adaptation aims to mitigate the domain shift problem when transferring knowledge from one domain into another similar but different domain. However, most existing works rely on extracting marginal features without considering class labels. Moreover, some methods name their model as so-called unsupervised domain adaptation while tuning the parameters using the target domain label. To address these issues, we propose a novel approach called correlated adversarial joint discrepancy adaptation network (*CAJNet*), which minimizes the joint discrepancy of two domains and achieves competitive performance with tuning parameters using the correlated label. By training the joint features, we can align the marginal and conditional distributions between the two domains. In addition, we introduce a probability-based top- \mathcal{K} correlated label (\mathcal{K} -label), which is a powerful indicator of the target domain and effective metric to tune parameters to aid predictions. Extensive experiments on benchmark datasets demonstrate significant improvements in classification accuracy over the state of the art.

Index Terms—Domain adaptation, Adversarial learning, Distribution alignment

I. INTRODUCTION

The availability of massive labeled training data is a prerequisite of machine learning models. Unfortunately, such a requirement cannot be met in many real scenarios. On the other hand, it is time-consuming and expensive to manually annotate data. Therefore, it is often necessary to transfer knowledge from an existing auxiliary labeled domain to a similar but different domain with limited or no labels. However, due to the phenomenon of data bias or domain shift [1], machine learning models do not generalize well from an existing domain to a novel unlabeled domain. Domain adaptation (DA) aims to leverage knowledge from an abundant labeled source domain to learn an effective predictor for the target domain with few or no labels, while mitigating the domain shift problem. In this paper, we focus on unsupervised domain adaptation (UDA), where the target domain has no labels.

Recently, deep neural network methods witness great success in UDA. Especially, adversarial learning shows its power in embedding in deep neural networks to learn feature representations to minimize the discrepancy between the source and target domains [2], [3]. Inspired by the generative adversarial network (GAN) [4], adversarial learning also contains a feature extractor and a domain discriminator. The domain discriminator aims to distinguish the source domain from the target domain, while the feature extractor aims to learn domain-invariant representations to fool the domain discriminator [2],

[3], [5]. The target domain risk is expected to be minimized via minimax optimization.

Although many methods achieve remarkable results in domain adaptation, they still suffer from two challenges: (1) the feature extractor often seeks for marginal features without considering the class label information; and (2) it is inappropriate to tune model hyperparameters using real target domain labels, as it violates the setting of UDA.

To address these challenges, we aggregate four different loss functions in one framework: classification loss, adversarial domain discrepancy loss, top- \mathcal{K} correlated loss, and domain alignment loss to reduce the joint discrepancy of two domains. Moreover, hyperparameters are properly updated using the proposed top- \mathcal{K} correlated label without the target domain label. Our contributions are three-fold:

- We propose a novel correlated adversarial joint discrepancy adaptation network (*CAJNet*) to adversarially minimize the joint domain discrepancy;
- We develop a probabilistic mechanism to compute the joint features for two domains to align both marginal and conditional distributions of source and target domains in a dynamic domain alignment setting;
- We introduce a top- \mathcal{K} correlated loss to help ensure predictions are locally consistent (with those of nearby examples) and use that loss to tune hyperparameters.

Experiments on three benchmark dataset (Office + Caltech-10, Office-31 and Office-Home) show that *CAJNet* achieves higher classification accuracy over state-of-the-art methods.

II. RELATED WORK

With the advent of GAN [4], adversarial learning models have been found to be an impactful mechanism for identifying invariant representations in domain adaptation. The Domain Adversarial Neural Network (DANN) considers a minimax loss to integrate a gradient reversal layer to promote the discrimination of source and target domains [5]. The Adversarial Discriminative Domain Adaptation (ADDA) method uses an inverted label GAN loss to split the source and target domain, and features can be learned separately [2]. The Joint Adaptation Network (JAN) [6] combined MMD with adversarial learning to align the joint distribution between two distributions of multiple domain-specific layers across domains. Domain-Symmetric Network (SymNet) [7] is a symmetrically designed source and target classifier based on an additional classifier. The proposed category level loss

can improve the domain level loss by learning the invariant features between two domains. Miyato et al. [8] incorporated virtual adversarial training (VAT) in semi-supervised contexts to smooth the output distributions as a regularization of deep networks. Later, Virtual Adversarial Domain Adaptation (VADA) improved adversarial feature adaptation using VAT. It generated adversarial examples against only the source classifier and adapted on the target domain [9]. Unlike VADA methods, Transferable Adversarial Training (TAT) adversarially generates transferable examples that fit the gap between source and target domain [3].

Through adversarial learning, the domain discrepancy can be largely reduced. Moreover, the feature extractor in deep networks could be large enough to align the feature level (marginal) distribution between domains. However, the features in typical adversarial learning still lack class label information (conditional distribution), *i.e.*, the joint distribution of category level and feature level have not been well addressed.

III. METHOD

A. Problem.

For unsupervised domain adaptation, given a source domain $\mathcal{D}_S = \{\mathcal{X}_S^i, \mathcal{Y}_S^i\}_{i=1}^{N_S}$ of N_S labeled samples in C categories and a target domain $\mathcal{D}_T = \{\mathcal{X}_T^j\}_{j=1}^{N_T}$ of N_T samples without any labels (*i.e.*, \mathcal{Y}_T is unknown). The samples \mathcal{X}_S and \mathcal{X}_T obey the marginal distribution of P_S and P_T . The conditional distributions of the two domains are denoted as Q_S and Q_T . Due to the discrepancy between the two domains, the distributions are assumed to be different, *i.e.*, $P_S \neq P_T$ and $Q_S \neq Q_T$. Our ultimate goal is to learn a classifier \mathcal{F} under a feature extractor Φ , that ensures lower generalization error in the target domain.

B. Motivation.

Most existing models were trained only with the extracted marginal features without considering the conditional features. Also, hyperparameters in most deep learning-based models are tuned using target domain accuracy, thus incorrectly exploiting the target labels. A fairer alternative is to develop a metric for tuning to enable impartial assessment of expected performance in the target domain. Furthermore, most existing work aligns only one distribution (marginal, conditional, or joint) to reduce the domain discrepancy.

Differently, in this work, we propose first to compute marginal and conditional features and align marginal, conditional, and joint distributions to better align the two domains. In light of the value of using conditional features, we further develop a top- \mathcal{K} labeling approach for tuning parameters in the target domain.

C. Joint discrepancy in latent space.

Considering extracted features are from a pre-trained neural network that maps both domains into a latent space, the

marginal discrepancy $Dist^M$ between the source and target domain can be formalized as follows.

$$Dist^M(\mathcal{X}_S, \mathcal{X}_T) = \left\| \frac{1}{N_S} \sum_{i=1}^{N_S} \Phi(\mathcal{X}_S^i) - \frac{1}{N_T} \sum_{j=1}^{N_T} \Phi(\mathcal{X}_T^j) \right\|_2, \quad (1)$$

where $\|\cdot\|_2$ is the L2 norm, $\Phi(\cdot) \rightarrow \mathbb{R}^d$ is an activation function, which maps samples into a latent d -dimensions space, using any pre-trained neural network. Specifically, $\Phi(x_l) = \Phi(W_{(l)}x_{(l-1)} + b_{(l)})$, where $W_{(l)}, b_{(l)}, x_{(l)}$ are the model weight, bias and output of the l -th layer, and then $\Phi(\mathcal{X}_{S/T}) \in \mathbb{R}^{N_{S/T} \times d}$. For the feature extraction, l is the last fully connected layer ($d = 1,000$). In the latent space, the marginal distance $Dist^M$ is minimized via the pre-trained model, and extracted features are frequently used as the single input for traditional models. The conditional discrepancy $Dist^C$ of two domains should also be minimizing in the following equation.

$$Dist^C(\mathcal{Y}_S|\mathcal{X}_S, \mathcal{Y}_T|\mathcal{X}_T) = Dist^C \sum_{c=1}^C (\mathcal{Y}_{S^c}|\mathcal{X}_{S^c}, \mathcal{Y}_{T^c}|\mathcal{X}_{T^c}) \\ = \sum_{c=1}^C \left\| \frac{1}{N_S^c} \sum_{i=1}^{N_S^c} \Phi(\mathcal{X}_{S^c}^i) - \frac{1}{N_T^c} \sum_{j=1}^{N_T^c} \Phi(\mathcal{X}_{T^c}^j) \right\|_2, \quad (2)$$

where C is the number of categories, $\mathcal{Y}_{S^c}|\mathcal{X}_{S^c}$ (or $\mathcal{Y}_{T^c}|\mathcal{X}_{T^c}$) represents c^{th} category target (sub-domain) data in the source (or target) domain and N_S^c (or N_T^c) is the number of samples in the c^{th} category in the source or target domain. However, we cannot directly compute such a conditional distance due to the unknown \mathcal{Y}_T or $\Phi(\mathcal{X}_{T^c}^j)$. We hence need to reformulate the sub-domain data, so that it contains the C categories' label information. In the latent space, if we only consider the marginal discrepancy, supposing that $\mathcal{Y}_T^j = c$, the probability of marking a correct prediction of \mathcal{X}_T^j is

$$p_j^M = \sum_{i=1}^{N_S^c} p_{j|i}^M = \sum_{i=1}^{N_S^c} \frac{Dist^M(\mathcal{X}_S^i, \mathcal{X}_T^j)}{\sum_{i=1}^{N_S^c} Dist^M(\mathcal{X}_S^i, \mathcal{X}_T^j)} \\ = \frac{\sum_{i=1}^{N_S^c} \|\Phi(\mathcal{X}_S^i) - \Phi(\mathcal{X}_T^j)\|_2}{\sum_{i=1}^{N_S^c} \|\Phi(\mathcal{X}_S^i) - \Phi(\mathcal{X}_T^j)\|_2}. \quad (3)$$

Motivated thus, we model the target conditional features:

$$\Phi(\mathcal{X}_T)^C = \mathcal{Y}_{T^c}|\mathcal{X}_{T^c} \\ \approx \otimes_{j=1}^{N_T} \oplus_{c=1}^C \left\| \frac{1}{N_S^c} \sum_{i=1}^{N_S^c} \Phi(\mathcal{X}_S^i) - \Phi(\mathcal{X}_T^j) \right\|_2, \quad (4)$$

where \otimes and \oplus is the column and row matrix concatenation operation, respectively. One target sample has the C -dimensions and $\Phi(\mathcal{X}_T)^C \in \mathbb{R}^{N_T \times C}$. Similarly, the source conditional features can be denoted as $\Phi(\mathcal{X}_S)^C = \mathcal{Y}_{S^c}|\mathcal{X}_{S^c} \approx \otimes_{i=1}^{N_S} \oplus_{c=1}^C \left\| \frac{1}{N_S^c} \sum_{i=1}^{N_S^c} \Phi(\mathcal{X}_S^i) - \Phi(\mathcal{X}_S^{i'}) \right\|_2 \in \mathbb{R}^{N_S \times C}$ ($\mathcal{X}_S^{i'}$ is one element in \mathcal{X}_S). Therefore, we rewrite Eq. 2 as follows.

$$Dist^C(\mathcal{Y}_S|\mathcal{X}_S, \mathcal{Y}_T|\mathcal{X}_T) \approx Dist^C(\Phi(\mathcal{X}_S)^C, \Phi(\mathcal{X}_T)^C) \\ = \left\| \frac{1}{N_S} \sum_{i=1}^{N_S} \Phi(\mathcal{X}_S^i)^C - \frac{1}{N_T} \sum_{j=1}^{N_T} \Phi(\mathcal{X}_T^j)^C \right\|_2 \quad (5)$$

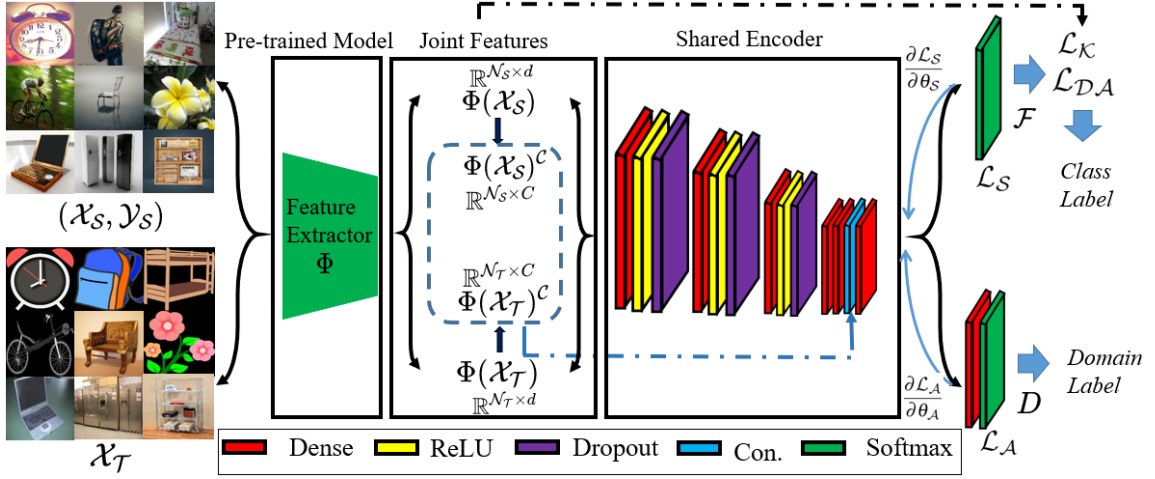


Fig. 1: The architecture of the *CAJNet* model. We first extract marginal features $\Phi(\mathcal{X}_{S/T}) \in \mathbb{R}^{N_{S/T} \times d}$ for both source and target domains via Φ using a pre-trained model, and then compute the conditional features $\Phi(\mathcal{X}_{S/T})^C \in \mathbb{R}^{N_{S/T} \times C}$ (blue dashed rectangle), which will be concatenated in the shared encoder (blue dashed line). The shared encoder is jointly trained with marginal and conditional features. Also, the parameters in the shared encoder are updated by the backward gradients from the class label classifier and the domain label predictor. In addition, the joint features are included in aligning marginal and conditional distributions in \mathcal{L}_{DA} (\mathcal{L}_S : classification loss, \mathcal{L}_A : adversarial domain loss, \mathcal{L}_K : top- \mathcal{K} correlated loss, and \mathcal{L}_{DA} : domain alignment loss, Con.: concatenate, \mathcal{F} : source classifier and D : discriminator, / means or).

If we only consider the conditional discrepancy, the probability of making correct prediction of \mathcal{X}_T^j is denoted as:

$$p_j^c = \sum_{i=1}^{N_S^c} p_{j|i}^c = \sum_{i=1}^{N_S^c} \frac{Dist^c(\mathcal{Y}_S^i | \mathcal{X}_S^i, \mathcal{Y}_T^j | \mathcal{X}_T^j)}{\sum_{i=1}^{N_S} Dist^c(\mathcal{Y}_S^i | \mathcal{X}_S^i, \mathcal{Y}_T^j | \mathcal{X}_T^j)} \quad (6)$$

$$= \frac{\sum_{i=1}^{N_S^c} \|\Phi(\mathcal{X}_S^i)^c - \Phi(\mathcal{X}_T^j)^c\|_2}{\sum_{i=1}^{N_S} \|\Phi(\mathcal{X}_S^i)^c - \Phi(\mathcal{X}_T^j)^c\|_2}$$

Substituting Eq. 1 and Eq. 5, we define the **joint discrepancy** in the latent space in Eq. 7.

$$Dist(\mathcal{D}_S, \mathcal{D}_T) = Dist^M(\mathcal{X}_S, \mathcal{X}_T) + Dist^c(\Phi(\mathcal{X}_S)^c, \Phi(\mathcal{X}_T)^c) \quad (7)$$

The **joint features** of the source and target domain are defined as $\mathcal{J}(\mathcal{X}_S) = \Phi(\mathcal{X}_S) \odot \Phi(\mathcal{X}_S)^c$ and $\mathcal{J}(\mathcal{X}_T) = \Phi(\mathcal{X}_T) \odot \Phi(\mathcal{X}_T)^c$, where \odot is the feature conjunct function.

Substituting Eq. 3 and Eq. 6, the probability of making correct prediction of \mathcal{X}_T^j considering both marginal and conditional discrepancies is given by:

$$p_j^{M,c} = \sum_{i=1}^{N_S^c} (p_{j|i}^M \circ p_{j|i}^c)$$

$$= \frac{\sum_{i=1}^{N_S^c} (\|\Phi(\mathcal{X}_S^i) - \Phi(\mathcal{X}_T^j)\|_2 + \|\Phi(\mathcal{X}_S^i)^c - \Phi(\mathcal{X}_T^j)^c\|_2)}{\sum_{i=1}^{N_S} (\|\Phi(\mathcal{X}_S^i) - \Phi(\mathcal{X}_T^j)\|_2 + \|\Phi(\mathcal{X}_S^i)^c - \Phi(\mathcal{X}_T^j)^c\|_2)} \quad (8)$$

where \circ is the probability conjunct function.

Since we want to minimize the joint discrepancy of the two domains in Eq. 7, it is equivalent to maximize $\sum_{j=1}^{N_T} p_j^{M,c}$ in the target domain. To realize this, our proposed *CAJNet* optimizes a four-component loss function using the joint features, as shown in Fig. 1. In summary, the discrepancies of both the

marginal and conditional distributions between the source and target can be jointly reduced by minimizing $Dist(\mathcal{D}_S, \mathcal{D}_T)$. Specifically, the discrepancy between P_S and P_T can be reduced by minimizing $Dist^M$, whereas the mismatches of Q_S and Q_T can be decreased by minimizing $Dist^c$.

\mathcal{Y}_{pred}	$\mathcal{Y}_K(\mathcal{K} = 3)$	$\mathcal{Y}_{pred}[\mathcal{Y}_K]$	$M(\mathcal{Y}_{pred}[\mathcal{Y}_K])$
\mathcal{X}^1	0	2 3 4	0
\mathcal{X}^2	1	1 3 4	0
\mathcal{X}^3	0	1 2 4	0
\mathcal{X}^4	0	5 1 3	0
\mathcal{X}^5	1	4 1 3	0

Fig. 2: An example of top- \mathcal{K} (here, 3) correlated labels in updating predicted labels. Given five examples (\mathcal{X}^1 to \mathcal{X}^5), the prediction is \mathcal{Y}_{pred} . The predictions of the five examples contain two classes, while the true labels are all zeros. There are two incorrectly predicted results (\mathcal{Y}_{pred}^2 and \mathcal{Y}_{pred}^5). The top-3 \mathcal{Y}_K shows the top 3 instances that should be in the same class (e.g., \mathcal{X}^1 should have the same class as \mathcal{X}^2 , \mathcal{X}^3 and \mathcal{X}^4). The $M(\mathcal{Y}_{pred}[\mathcal{Y}_K])$ changes the predicted labels and (in this case) outputs the same label as the truth label.

D. Top \mathcal{K} correlated label (\mathcal{K} -label).

Although there are no labels available in the target domain, model parameters still need to be tuned based on target domain information for pursuing higher transfer accuracy. We hence calculate the top- \mathcal{K} correlated labels for the target domain. Eq. 8 considers the probability of cross domains, while in the same domain, we can also denote a similar probability $p_{j|n}$ to indicate whether \mathcal{X}_T^j would pick \mathcal{X}_T^n ($n \neq j$) as its neighbor.

For nearby data points, $p_{j|n}$ is relatively high, whereas for widely separated data points, $p_{j|n}$ will be almost infinitesimal.

The probabilistic similarity between \mathcal{X}_T^j and \mathcal{X}_T^n is defined:

$$p_{j|n}^{\mathcal{M},c} = p_{j|n}^{\mathcal{M}} \circ p_{j|n}^c$$

$$= \frac{\|\Phi(\mathcal{X}_T^n) - \Phi(\mathcal{X}_T^j)\|_2 + \|\Phi(\mathcal{X}_T^n)^c - \Phi(\mathcal{X}_T^j)^c\|_2}{\sum_{n=1}^{\mathcal{N}_T} (\|\Phi(\mathcal{X}_T^n) - \Phi(\mathcal{X}_T^j)\|_2 + \|\Phi(\mathcal{X}_T^n)^c - \Phi(\mathcal{X}_T^j)^c\|_2)}.$$

This probability matrix in the target domain is $P_{T_m} = \otimes_{j=1}^{\mathcal{N}_T} \oplus_{n=1}^{\mathcal{N}_T} p_{j|n}^{\mathcal{M},c}$ ($p_{j|j}^{\mathcal{M},c} = 0$). Hence, $P_{T_m} \in \mathbb{R}^{\mathcal{N}_T \times \mathcal{N}_T}$ with its elements $0 \leq P_{jn} \leq 1$, also has the constraint of $\sum_{n=1}^{\mathcal{N}_T} P_{jn} = 1$, $\forall n \in \{1, \dots, \mathcal{N}_T\}$. The \mathcal{K} -label measures the index of the most similar sample in the same domain, which is reformulated as $\mathcal{Y}_{T_k} = \otimes_{j=1}^{\mathcal{N}_T} \uparrow (P_{T_m}^j) \in \mathbb{R}^{\mathcal{N}_T \times \mathcal{K}}$, $\mathcal{Y}_{S_k} = \otimes_{i=1}^{\mathcal{N}_S} \uparrow (P_{S_m}^i) \in \mathbb{R}^{\mathcal{N}_S \times \mathcal{K}}$, where \uparrow is an operation that first descending sort the probability matrix $P_{T_m}^j/P_{S_m}^i$ and then return the first \mathcal{K} index of associated probability. The defined top \mathcal{K} -label contains the k nearest neighborhood, that has the top k highest probability for selecting one sample in the same category. Therefore, the top- \mathcal{K} correlated loss in the target domain is defined in Eq. 9.

$$\mathcal{L}_{\mathcal{K}}(\mathcal{Y}_{pred}, \mathcal{Y}_{\mathcal{K}}) = \frac{1}{\mathcal{N}_{S/T}} \sum_{i/j=1}^{\mathcal{N}_{S/T}} |\mathcal{Y}_{pred}^{i/j} - M(\mathcal{Y}_{pred}[\mathcal{Y}_{\mathcal{K}}^{i/j}])| \quad (9)$$

where $|\cdot|$ is the absolute value operation, $\mathcal{Y}_{pred}^{i/j}$ is the prediction of either source domain or target domain from source classifier in Eq. 10, and $\mathcal{Y}_{\mathcal{K}}^{i/j}$ is the correlated label with size $\mathcal{N}_{S/T} \times \mathcal{K}$; it shows the top- \mathcal{K} index, which is highly related to the instance that should be in the same class. $\mathcal{Y}_{pred}[\mathcal{Y}_{\mathcal{K}}]$ is the updated matrix for the predicted labels and $M(\cdot)$ selects the most frequent labels in the updated matrix as shown in Fig. 2. Therefore, the loss measures how different the predicted label is to its nearest neighbors, which are assumed to be from the same domain.

E. Source classifier.

The task in the source domain is trained using the typical cross-entropy loss:

$$\mathcal{L}_S(\mathcal{F}(\mathcal{J}(\mathcal{X}_S)), \mathcal{Y}_S) = -\frac{1}{\mathcal{N}_S} \sum_{i=1}^{\mathcal{N}_S} \sum_{c=1}^C \mathcal{Y}_{S_c}^i \log(\mathcal{F}(\mathcal{J}(\mathcal{X}_{S_c}^i))), \quad (10)$$

where $\mathcal{Y}_{S_c}^i \in [0, 1]^C$ is the binary indicator of each class c in true label, and $\mathcal{F}(\mathcal{J}(\mathcal{X}_{S_c}^i))$ is the predicted probability of class c using classifier \mathcal{F} .

F. Adversarial domain loss.

Given the feature representation output of shared encoder, we can learn a discriminator D as shown in Fig. 1, which can distinguish the two domains using following loss function:

$$\mathcal{L}_A(\mathcal{J}(\mathcal{X}_S), \mathcal{J}(\mathcal{X}_T)) = -\frac{1}{\mathcal{N}_S} \sum_{i=1}^{\mathcal{N}_S} \log(1 - D(\mathcal{J}(\mathcal{X}_S^i)))$$

$$- \frac{1}{\mathcal{N}_T} \sum_{j=1}^{\mathcal{N}_T} \log(D(\mathcal{J}(\mathcal{X}_T^j))) \quad (11)$$

G. Shared Encoder.

The shared encoder begins with three repeated blocks and each block has a dense layer, a ‘‘ReLU’’ activation layer, and a dropout layer. The numbers of units of the dense layer are 512, 128, and 64, respectively. The rate of the Dropout layer is 0.5. It ends with a dense layer (the number of units is the number of classes in each dataset). The shared encoder is jointly optimized by both the source classifier and the domain labels. Note that the input of shared encoder is the joint features: the marginal features $\Phi(\mathcal{X}_S)$ are first fed into the repeated blocks and then the conditional features $\Phi(\mathcal{X}_S^c)$ are concatenated.

Let $\mathcal{F}_{\mathcal{E}}(\cdot, \theta_{\mathcal{E}})$ be the output of shared encoder with parameters of $\theta_{\mathcal{E}}$. In addition, let $\mathcal{F}_S(\cdot, \theta_S)$ be output of class label classifier with parameters of θ_S and $\mathcal{F}_A(\cdot, \theta_A)$ be output of domain label predictor with parameters of θ_A . Therefore, the shared encoder is optimized by both class label classifier and domain label predictor. The parameters are updated as:

$$\theta_S \leftarrow \theta_S - \epsilon \frac{\partial \mathcal{L}_S}{\partial \theta_S}, \quad \theta_A \leftarrow \epsilon \tau \theta_A \frac{\partial \mathcal{L}_A}{\partial \theta_A}, \quad (12)$$

$$\theta_{\mathcal{E}} \leftarrow \theta_{\mathcal{E}} - \epsilon \left(\frac{\partial \mathcal{L}_S}{\partial \theta_{\mathcal{E}}} - \tau \theta_A \frac{\partial \mathcal{L}_A}{\partial \theta_{\mathcal{E}}} \right),$$

where ϵ is learning rate and τ is the adaptation factor.

H. Dynamic distribution alignment.

We can get the prediction of the target domain from the class label classifier. However, we can further improve the predicted accuracy by employing a dynamic distribution alignment, which can dynamically balance the marginal and conditional distribution and update the predicted labels in the target domain. Manifold Embedded Distribution Alignment (MEDA), proposed by Wang et al. [10], aligns learned features from manifold learning. However, Zhang et al. showed that there are defects in the MEDA model, which cannot estimate the geodesic of sub-source and sub-target domains [11]. We modified the domain alignment loss as follows:

$$\mathcal{L}_{DA}(\mathcal{D}_S, \mathcal{D}_T) = \arg \min \mathcal{L}_S(\mathcal{F}(\mathcal{J}(\mathcal{X}_S)), \mathcal{Y}_S) + \eta \|\mathcal{F}\|_{\mathcal{K}}^2$$

$$+ \lambda \overline{D}_{\mathcal{F}}(\mathcal{D}_S, \mathcal{D}_T) + \rho R_{\mathcal{F}}(\mathcal{D}_S, \mathcal{D}_T)$$

where \mathcal{F} is the classifier from the shared encoder, \mathcal{L}_G is the sum of squares loss; $\|\mathcal{F}\|^2$ is the squared norm of \mathcal{F} ; and the first two terms minimize the structure risk of shared encoder. $\overline{D}_{\mathcal{F}}(\cdot, \cdot)$ represents the dynamic distribution alignment; $R_{\mathcal{F}}(\cdot, \cdot)$ is a Laplacian regularization; η, λ , and ρ are regularization parameters. Specifically, $\overline{D}_{\mathcal{F}}(\mathcal{D}_S, \mathcal{D}_T) = (1 - \mu) D_{\mathcal{F}}(P_S, P_T) + \mu \sum_{c=1}^C D_{\mathcal{F}}^c(Q_S, Q_T)$, where μ is an adaptive factor to balance the marginal distribution (P_S, P_T), and conditional distribution (Q_S, Q_T) [10].

I. CAJNet model.

The framework diagram of our proposed CAJNet model is depicted in Fig. 1. Our model minimizes the following objective function:

$$\mathcal{L}(\mathcal{X}_S, \mathcal{Y}_S, \mathcal{X}_T, \mathcal{Y}_{T_k}) = \arg \min (\mathcal{L}_S(\mathcal{F}(\mathcal{J}(\mathcal{X}_S)), \mathcal{Y}_S)$$

$$+ \mathcal{L}_A(\mathcal{J}(\mathcal{X}_S), \mathcal{J}(\mathcal{X}_T)) + \mathcal{L}_{\mathcal{K}}(\mathcal{Y}_{pred}, \mathcal{Y}_{T_k}) + \mathcal{L}_{DA}(\mathcal{D}_S, \mathcal{D}_T)) \quad (13)$$

where \mathcal{F} is the source classifier in Eq. 10; \mathcal{L}_S is the cross-entropy loss; \mathcal{Y}_{pred} is the predicted label and $\mathcal{Y}_{\mathcal{T}_K}$ is the top- \mathcal{K} correlated label, which shows the \mathcal{K} most highly related samples. $\mathcal{J}(\mathcal{X}_S)$ and $\mathcal{J}(\mathcal{X}_T)$ are the joint features. The \mathcal{L}_A , \mathcal{L}_K , and \mathcal{L}_{DA} represent the adversarial domain loss, top- \mathcal{K} correlated loss, and the domain alignment loss.

In summary, to effectively align both marginal and conditional distributions, we integrate the aforementioned four different loss functions in our proposed *CAJNet* model. The correlated adversarial learning mechanism further improves the domain invariant and leads to better results. The performance of different loss functions is shown in Sec. V.

IV. EXPERIMENTS

A. Datasets.

Office + Caltech-10 [1] consists of Office 10 and Caltech 10 datasets with 2,533 images in 10 classes of four domains: Amazon (A), Webcam (W), DSLR (D) and Caltech (C). $C \rightarrow A$ represents learning knowledge from domain C which is applied to domain A. **Office-31** [12] consists of 4,110 images in 31 classes from three domains: Amazon (A), Webcam (W), and DSLR (D). **Office-Home** [13] contains 15,588 images from 65 categories. It has four domains: Art (Ar), Clipart (Cl), Product (Pr) and Real-World (Rw).

B. Implementation details.

We implement our approach using PyTorch with an Nvidia GeForce 1080 Ti GPU and extract features for the three datasets from a fine-tuned ResNet50 network [14]. The 1,000 features are then extracted from the last fully connected layer [15]. Parameters in domain distribution alignment are $\eta = 0.1$, $\lambda = 10$, and $\rho = 10$, which are fixed based on previous research [10], and $\tau = 0.31$ is from [16]. Learning rate ($\epsilon = 0.001$), batch size (32), and number of epochs (1000) are determined by the performance of source domain. $\mathcal{K} = 3$ is tuned by minimizing \mathcal{L}_K . We compare our results with 11 state of the art (including traditional methods and deep networks).

C. Results.

The performance on Office + Caltech-10, Office-Home and Office-31 are shown in Tabs. I-III. For a fair comparison, we highlight in bold those methods that are re-implemented using our extracted features, and other methods are directly reported from their original papers. Our *CAJNet* model outperforms all state-of-the-art methods in terms of average accuracy (especially in the Office-Home dataset). It is compelling that our *CAJNet* model substantially enhances the classification accuracy on difficult adaptation tasks (e.g., $D \rightarrow A$ task in the Office-31 dataset and the challenging Office-Home dataset, which has a larger number of categories and different domains are visually dissimilar). To demonstrate the contribution of the proposed joint features, we also report the performance of marginal features alone (*CAJNet* - \mathcal{C}), conditional features alone (*CAJNet* - \mathcal{M}), and the joint features (*CAJNet*) in Tab. III. We can find that the marginal features provide slightly higher accuracy than the proposed conditional features. This

is because marginal features have more information than conditional features ($d \gg C$). However, the conditional features are still important in improving performance since *CAJNet* with joint features always has the highest accuracy, especially in the difficult Office-Home dataset. These experiments demonstrate the ability of the *CAJNet* model to align the marginal and conditional distributions of two domains.

V. DISCUSSION

There are two compelling advantages of the *CAJNet* model. First, we compute the joint features that minimize the joint discrepancy of two domains and align the joint distribution simultaneously. Secondly, we adversarially learn the domain invariant features, while the proposed top- \mathcal{K} label further updated the predicted labels via the favor of its neighbors. However, we observe that our model is compromised in some tasks ($W \rightarrow C$ in Office + Caltech-10 and $D \rightarrow W$ in Office-31 dataset) and so we cannot guarantee that our model always beats all other methods. We also conduct top- \mathcal{K} label analysis and perform an ablation study to determine the impact of different loss functions using Office-31 dataset. Other datasets show the same effect and hence are ignored.

Top- \mathcal{K} label analysis: \mathcal{K} is the only hyperparameter in *CAJNet*, controlling how many neighbors will be used for updated labels. Fig. 3 shows the impact of \mathcal{K} on \mathcal{L}_K loss using three tasks in the Office-Home dataset. We find that \mathcal{L}_K achieves the smallest value when $\mathcal{K} = 3$. \mathcal{L}_K increases when $\mathcal{K} > 4$ since the further away, the more likely noise labels will be included. Other tasks show the same effect. Therefore, the top- \mathcal{K} label is effective in updating predictions if correlated labels are not in the same class.

Ablation study: To better demonstrate the effects of different loss functions on classification accuracy, we present an ablation study in Tab. IV (\mathcal{A} : adversarial domain loss, \mathcal{K} : top- \mathcal{K} correlated loss, and \mathcal{DA} : distribution alignment loss). Notice that, \mathcal{L}_S is required during the training, and cannot be eliminated. “*CAJNet* - $\mathcal{A}/\mathcal{K}/\mathcal{DA}$ ” is implemented without adversarial domain loss, top- \mathcal{K} correlated loss, and distribution alignment loss. It is a simple model, which only reduces the source risk without minimizing the domain discrepancy. With the increasing of the number of loss functions, the robustness of our model keeps improving. Therefore, the proposed top- \mathcal{K} approach is effective in improving performance, and we can conclude that all of the loss functions are helpful and important in minimizing the target domain risk.




Fig. 3: Effect of different \mathcal{K} values on \mathcal{L}_K loss.

TABLE I: Accuracy (%) on Office + Caltech-10 dataset

Task	C→A	C→W	C→D	A→C	A→W	A→D	W→C	W→A	W→D	D→C	D→A	D→W	Ave.
JDA [17]	95.3	96.3	96.8	93.9	95.9	95.5	93.5	95.7	100	93.3	95.5	96.9	95.7
CORAL [18]	95.6	96.3	98.1	95.2	89.8	94.3	93.9	95.7	100	94.0	96.2	98.6	95.6
MEDA [10]	96.0	99.3	98.1	94.2	99.0	100	94.6	96.5	100	94.1	96.1	99.3	97.3
RTN [19]	93.7	96.9	94.2	88.1	95.2	95.5	86.6	92.5	100	84.6	93.8	99.2	93.4
MDDA [20]	93.6	95.2	93.4	89.1	95.7	96.6	86.5	94.8	100	84.7	94.7	99.4	93.6
SSD [21]	96.2	96.2	96.2	94.9	98.3	98.3	98.1	98.1	98.7	96.2	96.2	99.3	97.2
CAJNet	96.2	100	100	94.6	100	100	94.2	96.7	100	94.6	96.7	100	97.8

TABLE II: Accuracy (%) on Office-Home dataset

Task	Ar→Cl	Ar→Pr	Ar→Rw	Cl→Ar	Cl→Pr	Cl→Rw	Pr→Ar	Pr→Cl	Pr→Rw	Rw→Ar	Rw→Cl	Rw→Pr	Ave.
JDA [17]	47.4	72.8	76.1	60.7	68.6	70.5	66.0	49.1	76.4	69.6	52.5	79.7	65.8
CORAL [18]	48.0	78.7	80.9	65.7	74.7	75.5	68.4	49.8	80.7	73.0	50.1	82.4	69.0
MEDA [10]	48.5	74.5	78.8	64.8	76.1	75.2	67.4	49.1	79.7	72.2	51.7	81.5	68.3
DANN [16]	45.6	59.3	70.1	47.0	58.5	60.9	46.1	43.7	68.5	63.2	51.8	76.8	57.6
JAN [6]	45.9	61.2	68.9	50.4	59.7	61.0	45.8	43.4	70.3	63.9	52.4	76.8	58.3
TADA [22]	53.1	72.3	77.2	59.1	71.2	72.1	59.7	53.1	78.4	72.4	60.0	82.9	67.6
SymNets [7]	47.7	72.9	78.5	64.2	71.3	74.2	64.2	48.8	79.5	74.5	52.6	82.7	67.6
CAJNet	58.0	82.9	83.9	74.5	84.2	83.7	74.8	57.9	85.0	76.6	60.1	87.4	75.8

TABLE III: Accuracy (%) on Office-31 dataset

Task	A→W	A→D	W→A	W→D	D→A	D→W	Ave.
JDA [17]	79.1	79.7	72.9	97.4	71.0	94.2	82.4
CORAL [18]	88.9	87.6	74.7	99.2	73.0	96.7	86.7
MEDA [10]	90.8	91.4	74.6	97.2	75.4	96.0	87.6
JAN [6]	85.4	84.7	70.0	99.8	68.6	97.4	84.3
TADA [22]	94.3	91.6	73.0	99.8	72.9	98.7	88.4
SymNets [7]	90.8	93.9	72.5	100	74.6	98.8	88.4
CAN [23]	94.5	95.0	77.0	99.8	78.0	99.1	90.6
<i>CAJNet</i> - \mathcal{M}	93.7	93.0	78.5	98.5	80.1	97.9	90.3
<i>CAJNet</i> - \mathcal{C}	95.0	95.3	79.7	98.8	79.9	98.0	91.1
CAJNet	96.0	95.6	79.9	99.2	80.5	98.4	91.6

TABLE IV: Ablation study on Office-31 dataset

Task	A→W	A→D	W→A	W→D	D→A	D→W	Ave.
<i>CAJNet</i> - $\mathcal{A}/\mathcal{K}/\mathcal{D}\mathcal{A}$	91.8	88.4	78.2	98.2	77.5	97.4	88.6
<i>CAJNet</i> - $\mathcal{K}/\mathcal{D}\mathcal{A}$	92.6	89.9	78.5	98.2	78.9	97.6	89.3
<i>CAJNet</i> - \mathcal{A}/\mathcal{K}	93.3	94.4	79.2	98.2	79.5	97.4	90.3
<i>CAJNet</i> - $\mathcal{A}/\mathcal{D}\mathcal{A}$	93.8	94.6	79.2	98.5	79.5	97.6	90.5
<i>CAJNet</i> - \mathcal{K}	93.9	94.6	79.4	98.7	79.8	97.9	90.7
<i>CAJNet</i> - $\mathcal{D}\mathcal{A}$	95.6	94.6	79.0	99.2	78.9	98.0	90.9
<i>CAJNet</i> - \mathcal{A}	95.9	95.3	79.3	99.2	79.8	98.0	91.3
CAJNet	96.0	95.6	79.9	99.2	80.5	98.4	91.6

VI. CONCLUSION

We have proposed a correlated adversarial joint discrepancy adaptation network method (*CAJNet*) to overcome prior limitations in tuning parameters and aligning joint distributions of two domains by minimizing a four component loss function. The defined top- \mathcal{K} labels can help to further correct predicted labels of the two domains. Explicit domain-invariant features are learned through such a cross-domain training scheme. Experiments on three benchmark datasets show the robustness of our proposed *CAJNet* model.

REFERENCES

- [1] B. Gong, Y. Shi, F. Sha, and K. Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *CVPR*, pages 2066–2073, 2012.
- [2] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pages 7167–7176, 2017.
- [3] H. Liu, M. Long, J. Wang, and M. Jordan. Transferable adversarial training: A general approach to adapting deep classifiers. In *ICML*, pages 4013–4022, 2019.
- [4] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014.
- [5] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17(1):2096–2030, 2016.
- [6] M. Long, H. Zhu, J. Wang, and M. I. Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, volume 70, pages 2208–2217. JMLR.org, 2017.
- [7] Y. Zhang, H. Tang, K. Jia, and M. Tan. Domain-symmetric networks for adversarial domain adaptation. In *CVPR*, pages 5031–5040, 2019.
- [8] T. Miyato, S. Maeda, M. Koyama, and S. Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *TPAMI*, 41(8):1979–1993, 2018.
- [9] R. Shu, H. H. Bui, H. Narui, and S. Ermon. A DIRT-T approach to unsupervised domain adaptation. In *ICLR*, 2018.
- [10] J. Wang, W. Feng, Y. Chen, H. Yu, M. Huang, and P. S. Yu. Visual domain adaptation with manifold embedded distribution alignment. In *ACM MM*, pages 402–410, 2018.
- [11] Y. Zhang, S. Xie, and B. D. Davison. Transductive learning via improved geodesic sampling. In *BMVC*, page 122, 2019.
- [12] K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting visual category models to new domains. In *ECCV*, pages 213–226. Springer, 2010.
- [13] H. Venkateswara, J. Eusebio, S. Chakraborty, and S. Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, pages 5018–5027, 2017.
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [15] Y. Zhang and B. D. Davison. Impact of imagenet model selection on domain adaptation. In *WACVW*, pages 173–182, 2020.
- [16] M. Ghifary, W. B. Kleijn, and M. Zhang. Domain adaptive neural networks for object recognition. In *PRICAI*, pages 898–904. Springer, 2014.
- [17] M. Long, J. Wang, G. Ding, J. Sun, and P. S. Yu. Transfer feature learning with joint distribution adaptation. In *ICCV*, pages 2200–2207, 2013.
- [18] B. Sun, J. Feng, and K. Saenko. Correlation alignment for unsupervised domain adaptation. In *Domain Adaptation in Computer Vision Applications*, pages 153–171. Springer, 2017.
- [19] M. Long, H. Zhu, J. Wang, and M. I. Jordan. Unsupervised domain adaptation with residual transfer networks. In *NIPS*, pages 136–144, 2016.
- [20] M. M. Rahman, C. Fookes, M. Baktashmotlagh, and S. Sridharan. On minimum discrepancy estimation for deep domain adaptation. In *Domain Adaptation for Visual Understanding*, pages 81–94. Springer, 2020.
- [21] Y. Zhang and B. D. Davison. Domain adaptation for object recognition using subspace sampling demons. *Multimedia Tools and Applications*, pages 1–20, 2020.
- [22] X. Wang, L. Li, W. Ye, M. Long, and J. Wang. Transferable attention for domain adaptation. In *AAAI*, volume 33, pages 5345–5352, 2019.
- [23] G. Kang, L. Jiang, Y. Yang, and A. G. Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *CVPR*, pages 4893–4902, 2019.