Domain2Vec: Domain Embedding for Unsupervised Domain Adaptation

Xingchao Peng^{*1}, Yichen Li^{*2}, and Kate Saenko^{1,3}

¹ Boston University, Boston, MA, USA
 ² Stanford University, Stanford, CA, USA
 ³ MIT-IBM Watson AI Lab, Boston, MA, USA
 {xpeng,saenko}@bu.edu, liyichen@stanford.edu

Abstract. Conventional unsupervised domain adaptation (UDA) studies the knowledge transfer between a limited number of domains. This neglects the more practical scenario where data are distributed in numerous different domains in the real world. A technique to measure domain similarity is critical for domain adaptation performance. To describe and learn relations between different domains, we propose a novel DOMAIN2VEC model to provide vectorial representations of visual domains based on joint learning of feature disentanglement and Gram matrix. To evaluate the effectiveness of our DOMAIN2VEC model, we create two large-scale cross-domain benchmarks. The first one is TINYDA, which contains 54 domains and about one million MNIST-style images. The second benchmark is DOMAINBANK , which is collected from 56 existing vision datasets. We demonstrate that our embedding is capable of predicting domain similarities that match our intuition about visual relations between different domains. Extensive experiments are conducted to demonstrate the power of our new datasets in benchmarking state-of-the-art multi-source domain adaptation methods, as well as the advantage of our proposed model. Data and code are available at https://github.com/VisionLearningGroup/Domain2Vec

Keywords: Unsupervised Domain Adaptation, Domain Vectorization

1 Introduction

Generalizing models learned on one visual domain to novel domains has been a major pursuit of machine learning in the quest for universal object recognition. The performance of the learned methods degrades significantly when tested on novel domains due to the presence of *domain shift* [1].

Recently, Unsupervised Domain Adaptation (UDA) methods have been proposed to mitigate domain gap. For example, several learning-based UDA models [2,3,4] incorporate Maximum Mean Discrepancy loss to minimize the domain discrepancy; other models propose different learning schema to align the marginal

^{*} These authors contributed equally.



Fig. 1. Our DOMAIN2VEC architecture achieve deep domain embedding by by joint learning of feature disentanglement and Gram matrix. We employ *domain disentanglement* (red lines) and *class disentanglement* (blue lines) to extract domain-specific features and category specific features, both trained adversarially. We further apply a mutual information minimizer to enhance the disentanglement.

feature distributions of the source and target domains, including aligning secondorder correlation [5,6], moment matching [7], GAN-based alignment [8,9,10], and adversarial domain confusion [11,12,13]. However, most of the current UDA methods consider domain adaptation between limited number of domains (usually one source domain and one target domain). In addition, the state-of-the-art UDA models mainly focus on aligning the feature distribution of the source domain with that of the target domain, and fail to consider the natural distance and relations between different domains. In the more practical scenarios where multiple domain exists and the relations between different domains are unclear, it is critical to evaluate the natural domain distances between source and target so to be able to select one or several domains from the source domain pool such that the target domain can achieve the best performance.

In this paper, we introduce the DOMAIN2VEC embedding to represent domains as elements of a vector space. Formally, given N distinct domains $\hat{\mathcal{D}} = \{\hat{\mathcal{D}}_1, \hat{\mathcal{D}}_2, ..., \hat{\mathcal{D}}_N\}^{\dagger}$ domains, the aim is the learn a domain to vector mapping $\Phi : \hat{\mathcal{D}} \to V$. We would like our DOMAIN2VEC to hold the following properties: (i) given two domains $\hat{\mathcal{D}}_i, \hat{\mathcal{D}}_j$, the accuracy of a model trained on $\hat{\mathcal{D}}_i$ and tested on $\hat{\mathcal{D}}_j$ should be negatively correlated to the domain distance in the vector space V, *i.e.* smaller domain distance leads to better cross-domain performance; (ii) the domain distance of two domains with building images $(\hat{\mathcal{D}}_i^{building}, \hat{\mathcal{D}}_j^{building})$ should be smaller than that of $(\hat{\mathcal{D}}_i^{building}, \hat{\mathcal{D}}_j^{car})$. Our domain embedding can be used to reason about the space of domains and solve many unsupervised domain adaptation problems. As a motivating example, we study the problem of selecting the best combination of source domains when a novel target domain emerges.

[†] In this literature, the calligraphic \mathcal{G}, \mathcal{D} denote Gram matrix and domains, and italic G, D denote feature generator and disentangler, respectively.

Computation of the DOMAIN2VEC embedding leverages a complementary term between the Gram matrix of deep representations and the disentangled *domain*specific feature. Gram Matrices are commonly used to build style representations that compute the correlations between different filter activations in a deep network [14]. Since activations of a deep network trained on a visual domain are a rich representation of the domain itself, we use Gram Matrix to capture the texture information of a domain and further obtain a stationary, multi-scale representation of the input domain. Specifically, given a domain defined by $\hat{\mathcal{D}} = \{x_j, y_j\}_{j=1}^{n_i}$ with n_i $(i \in [1, N])$ examples, we feed the data through a pre-train reference convolutional neural network which we call feature generator G, and compute the activations of the fully connected layer as the latent representation f_G , as shown in Figure 1. Inspired by the feature disentanglement idea [15], we introduce a disentangler D to disentangle f_G into domain-specific feature f_{ds} and category-specific feature f_{cs} . Finally, we compute the Gram matrix of the activations of the hidden convolutional layers in the feature extractor. Given a domain $\hat{\mathcal{D}} = \{x_j, y_j\}_{j=1}^{n_i}$, we average the domain-specific features of all the training examples in $\hat{\mathcal{D}}$ as the *prototype* of domain $\hat{\mathcal{D}}$. We utilize the concatenation of *prototype* and the diagonal entries of the average Gram matrix as the final embedding vector of domain $\hat{\mathcal{D}}$. We show this embedding encodes the intrinsic properties of the domains (Sec 4).

To evaluate our DOMAIN2VEC model, a large-scale benchmark with multiple domains is required. However, state-of-the-art cross-domain datasets contain only a limited number of domains. For example, the large-scale DomainNet [16] that contains six domains, and the Office-31 [17] benchmark that only has three domains. In this paper, we create two large-scale datasets to facilitate the research of multi-domain embedding. TINYDAdataset is by far the largest MNIST-style cross domain dataset. It contains 54 domains and about one million training examples. Following Ganin *et al* [12], the images are generated by blending different foreground shapes over patches randomly cropped from the background images. The second benchmark is DOMAINBANK, which contains 56 domains sampled from the existing popular computer vision datasets.

Based on TINYDA dataset, we validate our DOMAIN2VEC model's property on the negative correlation between the cross-domain performance and the domain distance computed by our model. Then, we show the effectiveness of our DOMAIN2VEC on multi-source domain adaptation. In addition, comprehensive experiments on DOMAINBANK benchmark with openset domain adaptation and partial domain adaptation schema demonstrate that our model achieves significant improvements over the state-of-the-art methods.

The main contributions of this paper are highlighted as follows: (i) we propose a novel learning paradigm of deep domain embedding and develop a DO-MAIN2VEC model to achieve the domain embedding; (ii) we collect two stateof-the-art benchmarks to facilitate research in multiple domain embedding and adaptation. (iii) we conduct extensive experiments on various domain adaptation settings to demonstrate the effectiveness of our proposed model.

2 Related Work

Vectorial Representation Learning Discovery of effective representations that capture salient semantics for a given task is a fundamental goal for perceptual learning. The individual dimensions in the vectorial embedding have no inherent meaning. Instead, it is the overall patterns of location and distance between vectors that machine learning takes advantage of. GloVe [18] models achieve global vectorial embbedings for word by training on the nonzero elements in a word-word co-occurrence matrix, rather than on the entire sparse matrix or on individual context windows in a large corpus. DECAF [19] investigates semi-supervised multi-task learning of deep convolutional representations, where representations are learned on a set of related problems but applied to new tasks which have too few training examples to learn a full deep representation. Modern state-of-the-art deep models [20,21,22,23,24] learn semantic representations with supervision and are applied to various vision and language processing tasks. Another work which is very related to our work is the TASK2VEC model [25] which leverage the Fisher Information Matrix as the vectorial representation of different tasks. However, the TASK2VEC model mainly consider the similarity between different tasks. In this work, we mainly focus on the same task and introduce a DOMAIN2VEC framework to achieve deep domain embedding for multiple domains. Specifically, DOMAIN2VEC is initially proposed in the work of Deshmukh et al [26]. However, their model is designed for domain generalization. Our model is developed independently for a different purpose.

Unsupervised Domain Adaptation Deep neural networks have achieved remarkable success on diverse vision tasks [22,27,28] but at the expense of tedious labor work on labeling data. Given a large-scale unlabeled dataset, it is expensive to annotate enough training data such that we can train a deep model that generalizes well to that dataset. Unsupervised Domain Adaptation [17,4,12,13,16,15,29] provides an alternative way by transferring knowledge from a different but related domain (source domain) to the domain of interest (target domain). Specifically, 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 discrepancybased UDA approaches [2,30,31,6], adversary-based approaches [32,11,33], and reconstruction-based approaches [34,8,9,35]. 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 [36,37,38], MSDA has been applied to many practical applications [39,40,16]. Specifically, Ben-David *et al* [36] introduce an $\mathcal{H}\Delta\mathcal{H}$ -divergence between the weighted combination of source domains and a target domain. Different from the previous work, we propose a DOMAIN2VEC model to evaluate the natural distances between different domains.

Deep Feature Disentanglement Deep neural networks are known to extract features where multiple hidden factors are highly entangled [41]. Learning disentangled representations can help to model the relevant factors of data variation

as well as evaluate the relations between different domains by extracting the domain-specific features. To this end, recent work [42,43,33,44] leverages generative adversarial networks (GANs) [45] or variational autoencoders (VAEs) [46] to learn the interpretable representations. Under the multi-domain setting, Liu et al. [33] propose a unified feature disentanglement framework to learn domain-invariant features from data across different domains. Odena et al. [44] introduce an auxiliary classifier GAN (AC-GAN) to achieve representation disentanglement under supervised setting. Recent work [47,15] propose to disentangle the features into a domain-invariant content space and a domain-specific attributes space, producing diverse outputs without paired training data. In this paper, we propose a cross-disentanglement schema to disentangle the deep features into domain-specific and category-specific features.

3 Domain2Vec

We define the domain vectorization task as follows: given N domains $\hat{\mathcal{D}} = \{\hat{\mathcal{D}}_1, \hat{\mathcal{D}}_2, ..., \hat{\mathcal{D}}_N\}$ domains, the aim is the learn a domain to vector mapping $\Phi : \hat{\mathcal{D}} \to V$, which is capable of predicting domain similarities that match our intuition about visual relations between different domains. Our DOMAIN2VEC includes two components: we first leverage feature disentanglement to generate the *domain-specific* features, and then we achieve deep domain embedding by the joint learning of Gram Matrix of the latent representations and the *domain-specific* features.

3.1 Feature Disentanglement

Given an image-label pair (x,y), a deep neural network is a family of function $p_{\theta}(y|x)$, trained to approximate the posterior p(y|x) by minimizing the cross entropy loss $H_{p_{\theta},\hat{p}}(y|x) = \mathbb{E}_{x,y}[-\log p_{\theta}(y|x)]$, where \hat{p} is the empirical distribution defined by the *i*-th domain $\hat{\mathcal{D}}_i = \{x_j, y_j\}_{j=1}^{n_i}$ with n_i training examples, $i \in [1, n]$. It is beneficial, especially in domain vectorization task, to think of the deep neural network as composed of two parts: a feature generator which computes the latent representations $f_{\theta} = \phi_{\theta}(x)$ of the input data, and a classifier which encodes the distribution p(y|x) given the representation f_{θ} .

The latent representations $f_{\theta} = \phi_{\theta}(x)$ are highly entangled with multiple hidden factors. We propose to disentangle the hidden representations to *domainspecific* and *category-specific* features. Figure 1 shows the proposed model. Given N domains, the feature extractor G maps the input data to a latent feature vector f_G , which contains both the domain-specific and category-specific factors. The disentangler D is trained to disentangle the feature f_G to *domain-specific* feature f_{ds} and *category-specific* feature f_{cs} with cross-entropy loss and adversarial training loss. The feature reconstructor R is responsible to recover f_G from (f_{ds}, f_{cs}) pair, aiming to keep the information integrity in the disentanglement process. To enhance the disentanglement, we follow Peng *et al* [15] to apply a mutual information minimizer between f_{ds} and f_{cs} . A category classifier C is

trained with class labels to predict the class distributions and a domain classifier DC is trained with domain labels to predict the domain distributions. In addition, the cross-adversarial training step removes domain information from f_{cs} and category information from f_{ds} . We next describe each component in detail.

Category Disentanglement Given an input image x, the feature generator G computes the latent representation f_G . Our category disentanglement is achieved by two-step adversarial training. First, we train the disentangler D and the k-way category classifier C to correctly predict the class labels, supervised by the cross-entropy loss:

$$\mathcal{L}_{ce}^{class} = -\sum_{i=1}^{N} \mathbb{E}_{(x,y_c)\sim\hat{\mathcal{D}}_i} \sum_{k=1}^{K} \mathbb{1}[k=y_c] log(C(f_{cs}))$$
(1)

where $f_{cs} = D(G(x))$ and y_c indicates the class label.

In the second step, we aim to remove the domain-specific information from f_{cs} . Using a well-trained domain classifier (achieved with Equation 3), we freeze the domain classifier DC and train the disentangler to generate f_{cs} , with the objective to fool the domain classifier. This can be achieved by minimizing the negative entropy of the predicted domain distribution:

$$\mathcal{L}_{ent}^{class} = -\sum_{i=1}^{N} \frac{1}{n_i} \sum_{j=1}^{n_i} \log DC(f_{cs})$$

$$\tag{2}$$

This adversarial training process corresponds to the blue dotted line in Figure 1. The above adversarial training process forces the generated *category-specific* feature f_{cs} contains no *domain-specific* information.

Domain Disentanglement To achieve deep domain embedding, disentangling *category-specific* features is not enough, as it fails to describe the relations between different domains. We introduce domain disentanglement to disentangle the domain-specific features from the latent representations. Previous adversarialalignment based UDA models [11,15] propose to leverage a domain classifier to classify the input feature as source or target. However, the proposed domain classifier is a binary classifier, which can not be applied to our case directly. Similar to category disentanglement, our domain disentanglement is achieved by two-step adversarial training. We first train the feature generator G and disentangler D to extract the domain-specific feature f_{ds} , supervised by domain labels and cross-entropy loss:

$$\mathcal{L}_{ce}^{domain} = -\mathbb{E}_{(x,y_d)\sim\hat{\mathcal{D}}} \sum_{k=1}^{N} \mathbb{1}[k = y_d] log(DC(f_{ds}))$$
(3)

where $f_{ds} = D(G(x))$ and y_d denotes the domain label.

In the second step, we aim to remove the category-specific information from f_{ds} . With a well-trained classifier C, we freeze the parameter weights of C and train the disentangler to generate f_{ds} to confuse the category classifier C. Similarly,

we can minimize the negative entropy of the predicted class distribution:

$$\mathcal{L}_{ent}^{domain} = -\sum_{i=1}^{N} \frac{1}{n_i} \sum_{j=1}^{n_i} \log C(f_{ds}) \tag{4}$$

This adversarial training process corresponds to the red dotted line in Figure 1. If a well-trained category classifier C is not able to predict the correct class labels, the category-specific information has been successfully removed from f_{ds} . **Feature Reconstruction** Previous literature [15] has shown that some information gets lost in the feature disentangle process, especially when the feature disentangler D is composed of several fully connected and RELU layers and it cannot guarantee the information integrity in the feature disentanglement process. We therefore introduce a feature reconstructor R to recover the original feature f_G with the disentangled *domain-specific* feature and *category-specific* feature. The feature reconstructor R has two input and will concatenate the (f_{ds}, f_{cs}) pair to a vector in the first layer. The feature vector is feed forward to several fully connected and RELU layers. Denoting the reconstructed feature as \hat{f}_G , we can train the feature reconstruction process with the following loss:

$$\mathcal{L}_{rec} = \|\hat{f}_G - f_G\|_F^2 + KL(q(z|f_G)||p(z))$$
(5)

where the first term aims at recovering the original features extracted by G, and the second term calculates *Kullback-Leibler divergence* which penalizes the deviation of latent features from the prior distribution $p(z_c)$ (as $z \sim \mathcal{N}(0, I)$).

Mutual Information Minimization The mutual information is a pivotal measure of the mutual dependence between two variables. To enhance the disentanglement, we minimize the mutual information between *category-specific* features and *domain-specific* features. Specifically, the mutual information is defined as:

$$I(f_{ds}; f_{cs}) = \int_{\mathcal{P} \times \mathcal{Q}} \log \frac{d\mathbb{P}_{\mathcal{P}\mathcal{Q}}}{d\mathbb{P}_{\mathcal{P}} \otimes \mathbb{P}_{\mathcal{Q}}} d\mathbb{P}_{\mathcal{P}\mathcal{Q}}$$
(6)

where $\mathbb{P}_{\mathcal{P}\mathcal{Q}}$ is the joint probability distribution of (f_{ds}, f_{cs}) , and $\mathbb{P}_{\mathcal{P}} = \int_{\mathcal{Q}} d\mathbb{P}_{\mathcal{P}\mathcal{Q}}$, $\mathbb{P}_{\mathcal{Q}} = \int_{\mathcal{Q}} d\mathbb{P}_{\mathcal{P}\mathcal{Q}}$ are the marginal probability of f_{ds} and f_{cs} , respectively. The conventional mutual information is only tractable for discrete variables, for a limited family of problems where the probability distributions are unknown [48]. To address this issue, we follow [15] to adopt the Mutual Information Neural Estimator (MINE) [48] to estimate the mutual information by leveraging a neural network $T_{\theta}: I(\mathcal{P}; \mathcal{Q}) = \sup_{\theta \in \Theta} \mathbb{E}_{\mathbb{P}_{\mathcal{P}\mathcal{Q}}}[T_{\theta}] - \log(\mathbb{E}_{\mathbb{P}_{\mathcal{P}} \otimes \mathbb{P}_{\mathcal{Q}}}[e^{T_{\theta}}])$. Practically, MINE can be calculated as $I(\mathcal{P}; \mathcal{Q}) = \int \int \mathbb{P}_{\mathcal{P}\mathcal{Q}}(p, q) T(p, q, \theta) - \log(\int \int \mathbb{P}_{\mathcal{P}}(p)\mathbb{P}_{\mathcal{Q}}(q)e^{T(p, q, \theta)})$. To avoid computing the integrals, we leverage Monte-Carlo integration to calculate the estimation:

$$I(\mathcal{P}, \mathcal{Q}) = \frac{1}{n} \sum_{i=1}^{n} T(p, q, \theta) - \log(\frac{1}{n} \sum_{i=1}^{n} e^{T(p, q', \theta)})$$
(7)

where (p,q) are sampled from the joint distribution, q' is sampled from the marginal distribution \mathbb{P}_{Q} , n is number of training examples, and $T(p,q,\theta)$ is the

neural network parameteralized by θ to estimate the mutual information between \mathcal{P} and \mathcal{Q} , we refer the reader to MINE [48] for more details.

3.2 Deep Domain Embedding

Our DOMAIN2VEC model learns a domain to vector mapping $\Phi : \hat{\mathcal{D}} \to V$ by jointly embeds the Gram matrix and *domain-specific* features. Specifically, given a domain $\hat{\mathcal{D}} = \{x_j, y_j\}_{j=1}^{n_i}$, we compute the disentangled features for all the training examples of $\hat{\mathcal{D}}$. The prototype of domain $\hat{\mathcal{D}}$ is defined as: $P_{\hat{\mathcal{D}}} = \frac{1}{n_i} \sum f_{ds}^j$, denoting the average of the *domain-specific* features of the examples in $\hat{\mathcal{D}}$. In addition, we compute the Gram matrix of the activations of the hidden convolutional layers in the feature extractor G. The Gram matrix build a style representation that computes the correlations between different filter responses. The feature correlations are given by the Gram matrix $\mathcal{G} \in \mathcal{R}^{n \times n}$, where \mathcal{G}_{ij} is the inner product between the vectorised feature map between i and j:

$$\mathcal{G}_{ij} = \sum_{k} F_{ik} F_{jk} \tag{8}$$

where F is the vectorised feature map of the hidden convolutional layers. Since the full Gram matrix is prohibitively large for deep features we make an approximation by only leveraging the entries in the subdiagonal, main diagonal, and superdiagonal of the Gram matrix \mathcal{G} . We utilize the concatenation of the *prototype* $P_{\hat{\mathcal{D}}}$ and the diagonals of the \mathcal{G} as the final embedding of domain $\hat{\mathcal{D}}$.

Eliminating Sparsity The *domain-specific* feature and the Gram matrix are high sparsity data, which hampers the effectiveness of our DOMAIN2VEC model. To address this issue, we leverage dimensionality reduction technique to decrease the dimensionality. Empirically, we start by using PCA to reduce the dimensionality of the data to a specific length. Then we leverage Stochastic Neighbor Embedding [49] to reduce the dimensionality to our desired one.

Optimization Our model is trained in an end-to-end fashion. We train the feature extractor G, category and domain disentanglement component D, MINE and the reconstructor R iteratively with Stochastic Gradient Descent [50] or Adam [51] optimizer. The overall optimization objective is:

$$\mathcal{L} = w_1 \mathcal{L}^{class} + w_2 \mathcal{L}^{domain} + w_3 \mathcal{L}_{rec} + w_4 \mathcal{I}(f_{ds}, f_{cs}) \tag{9}$$

where w_1, w_2, w_3, w_4 are the hyper-parameters, $\mathcal{L}^{class} = \mathcal{L}^{class}_{ce} + \alpha \mathcal{L}^{class}_{ent}, \mathcal{L}^{domain} = \mathcal{L}^{domain}_{ce} + \alpha \mathcal{L}^{domain}_{ent}$ denote the category disentanglement loss and domain disentanglement loss.

4 Experiments

We test DOMAIN2VEC on two large-scale datasets we created. Our experiments aim to test both qualitative properties of the domain embedding and its performance on multi-source domain adaptation, openset domain adaptation and partial domain adaptation. In the main paper, we only report major results; more implementation details are provided in the Appendix.

$\mathrm{FG}/\mathrm{Mode}$	Black BG	White BG	Color	Grayscale	Original
MNIST	50419	50419	50419	50 - 19	50419
USPS	65473	65473	65478	65373	65473
FashionMNIST	all 🐴 🕌 🕌 👗				i i i fi

Table 1. Illustration of TINYDA dataset. More images are in supplementary material.

4.1 Dataset

To evaluate the domain-to-vector mapping ability of our DOMAIN2VEC model, a large-scale dataset with multiple domains is desired. Unfortunately, existing UDA benchmarks [17,52,16,53] only contain limited number of domains. These datasets provide limit benchmarking ability for our DOMAIN2VEC model. To address this issue, we collect two datasets for multiple domain embedding and adaptation, *i.e.*, TINYDA and DOMAINBANK.

TinyDA We create our by far the largest MNIST-style cross domain dataset to data, TINYDA. This dataset contains 54 domains and about one million MNISTstyle training examples. We generate our TINYDA dataset by blending different foreground shapes over patches randomly extracted from background images. This operation is formally defined for two images I^1 , I^2 as $I_{ijk}^{\text{out}} = \|I_{ijk}^1 - I_{ijk}^2\|$, where i, j are the coordinates of a pixel and k is the channel index. The foreground shapes are from MNIST [54], USPS [55], EMNIST [56], KMNIST [57], QMNIST [58]. and FashionMNIST [59]. Specifically, the MNIST, USPS, QMNIST contains digit images; EMNIST dataset includes images of MNIST-style English characters; KMNIST dataset is composed of images of Japanese characters; FashionMNIST dataset contains MNIST-style images about fashion. The background images are randomly cropped from CIFAR10 [60] or BSDS500 [61] dataset. We perform three different post-processes to our rendered images: (1) replace the background with black patch, (2) replace the background with white patch, (3) convert the images to grayscale. The three post-processes, together with the original foreground images and the generated color images, form a dataset with five different modes, *i.e.* White Background (WB), Black Background (BB), GrayScale image (GS), Color (Cr) image, Original image(Or).

DomainBank[‡] To evaluate our DOMAIN2VEC model on state-of-the-art computer vision datasets, we collect a large-scale benchmark, named DOMAINBANK. The images of DOMAINBANK dataset are sampled from 56 existing popular computer vision datasets such as COCO [62], CALTECH-256 [63], PASCAL [64], VisDA [53], DomainNet [16], *etc.* We choose the dataset with different image modalities, illuminations, camera perspectives *etc.* to increase the diversity of the domains. In total, we collect 339,772 images with image-level and domain-level annotations. Different from TINYDA, the categories of different domains in Do-MAINBANK are not identical. This property makes DOMAINBANK a good testbed for Openset Domain Adaptation [65,66] and Partial Domain Adaptation [67].

[‡] In this dataset, the *domain* is defined by datasets. The data from different genres or times typically have different underlying distributions.

_					_																						
L		KMNIST						EMNIST											Fash	ionM	NIST						
L			BSDS CIFAR			1 F		BSDS		CIFAR					BSDS				CIFAR								
L			WB BB Or Cr GS	WB BB Cr GS				WB	BB	Or	Cr	GS	WB	BB	Cr	GS			WB	BB	Or	Cr	GS	WB	BB	Cr	GS
Γ		WB	89.8 13.3 12.4 16.8 16.4	88.0 12.8 14.9 14.6	Γ	Τ	WB	86.6	2.9	2.8	8.1	8.6	83.2	5.1	6.9	7.5	П	WE	83.5	16.9	29.9	27.0	25.6	80.7	16.7	27.3 2	24.9
	s	BB	12.5 94.1 94.3 32.9 30.4	11.5 92.6 23.3 22.2		S	BB	3.6	87.3	88.0	23.4	18.1	4.2	82.8	14.9	13.4		BB	23.6	84.5	85.4	38.1	36.6	21.1	81.7	28.92	28.9
	ß	Or	8.4 56.9 95.4 35.2 32.9	9.3 62.6 24.7 23.2		ß	Or	12.0	31.1	91.3	33.4	32.2	11.1	33.6	21.1	21.2	56) Or	15.1	53.6	87.0	33.0	33.2	14.8	52.2	23.82	25.1
E	B	\mathbf{Cr}	73.4 68.6 89.8 84.2 69.1	66.2 66.5 70.9 56.5	E	; lei	Cr	59.1	47.0	85.8	80.0	60.8	47.9	42.0	60.0	42.7	Ξŕ	Cr	75.6	68.6	85.2	81.6	74.4	69.9	54.7	75.67	71.3
E		GS	72.7 64.0 87.9 67.4 74.1	68.7 66.7 55.1 59.4		3	GS	59.4	46.7	82.5	56.1	65.9	52.2	46.8	41.2	44.6	2	GS	72.3	66.3	83.5	71.5	77.6	70.2	61.9	69.57	73.2
S	2	WB	83.8 17.0 16.2 18.6 18.9	81.2 15.1 18.8 18.0			WB	87.8	13.9	4.5	15.3	16.7	86.1	12.2	13.0	13.6	1:BC	WE	82.9	18.1	27.2	28.5	28.6	81.8	17.0	29.62	29.3
Г	AR	BB	13.1 90.0 91.2 26.0 24.1	11.8 88.8 18.8 17.9	ſ	Ĩ	BB	2.1	85.4	87.1	18.1	17.1	1.9	82.7	12.0	12.5	as	BB	21.1	84.8	86.2	29.1	28.4	18.1	82.3	22.1 2	23.3
L	E	\mathbf{Cr}	66.5 65.8 85.3 81.4 68.8	61.6 65.7 76.1 65.7		旧	Cr	58.2	48.9	83.5	76.1	59.6	48.4	44.7	67.8	55.0		Cr	75.1	67.9	85.1	82.2	75.6	72.4	62.4	78.67	16.6
		GS	64.5 60.5 85.8 58.0 70.7	60.8 63.4 56.7 66.8		Γ	GS	46.6	46.5	81.1	48.1	63.2	43.8	48.8	45.3	57.4		GS	67.9	61.8	82.2	65.2	77.0	66.3	58.0	68.7 7	76.3

 Table 2. Experimental results on TINYDA. The column-wise domains are source domains, the row-wise domains are the target domain.



Fig. 2. Deep domain embedding results of our DOMAIN2VEC model on TINYDA dataset: (a) t-SNE plot of the embedding result (color indicates different domain); (b) Domain knowledge graph. The size and color of the circles indicate the number of training examples and the degree of that domain, respectively. The width of the edge shows the domain distance between two domains. (c) The final deep domain embedding of our DOMAIN2VEC model. (Best viewed in color. Zoom in to see details.)

4.2 Experiments on TinyDA

Domain Embedding Results We apply our devised DOMAIN2VEC model to TINYDA dataset to achieve deep domain embedding. The results are shown in Figure 2. Specifically, the domain knowledge graph shows the relations between different domains. The nodes in the graph show the deep domain embedding. For each domain, we connect it with five closest neighboring domains with a edge weighted by their domain distance. The size and the color of the nodes are correlated with the number of training images in that domain and the degree of that domain, respectively. To validate that the domain distance computed with our DOMAIN2VEC is negatively correlated with the cross-domain performance, we conduct extensive experiments to calculate the cross-domain results on TINYDA dataset, as shown in Table 2. We split the cross-domain results in three sub-tables for Japanese characters (KMNIST), English characters (EMNIST) and fashion items (FashionMNIST), respectively. In each sub-table, the column-wise domains are selected as the source domain and the row-wise domains are selected as the target domain.

Standards	Models	$MNIST \rightarrow USPS$	$MNIST \rightarrow QMNIST$	$USPS \rightarrow MNIST$	$USPS \rightarrow QMNIST$	$QMNIST \rightarrow MNIST$	$QMNIST \rightarrow USPS$	Avg
	Source Only	17.7 ± 0.21	83.4±0.55	16.4 ± 0.32	16.3 ± 0.25	83.1 ± 0.32	20.2 ± 0.31	39.5±0.32
	DAN [4]	21.4 ± 0.27	87.1 ± 0.64	19.7±0.37	19.9±0.34	85.7±0.34	21.8 ± 0.37	42.6±0.39
	RTN [69]	18.0 ± 0.28	85.0 ± 0.58	18.8 ± 0.37	20.0 ± 0.26	84.2 ± 0.42	21.3 ± 0.34	41.2±0.38
Single	JAN [2]	21.7 ± 0.27	87.6 ± 0.64	19.4 ± 0.42	18.0±0.29	87.2 ± 0.36	25.1 ± 0.33	43.2±0.39
Best	DANN [12]	21.2 ± 0.25	86.1 ± 0.55	20.1 ± 0.31	19.4±0.24	86.6 ± 0.38	24.0 ± 0.34	42.9±0.34
	ADDA [11]	20.3 ± 0.31	88.1 ± 0.63	18.3 ± 0.46	21.4 ± 0.38	88.5±0.39	25.9 ± 0.43	43.8±0.43
	SE [29]	13.6 ± 0.42	78.1 ± 0.87	10.7 ± 0.62	11.8 ± 0.50	80.1 ± 0.64	17.0 ± 0.55	35.2±0.60
	MCD [13]	23.8 ± 0.33	89.0 ± 0.61	22.3 ± 0.36	19.6±0.26	86.7 ± 0.36	22.6 ± 0.41	44.0±0.39
	Source Only	20.2 ± 0.23	85.7±0.59	19.2 ± 0.42	20.5±0.37	85.1±0.25	19.2 ± 0.40	41.6±0.38
	DAN [4]	19.8 ± 0.30	85.4 ± 0.64	22.4 ± 0.43	21.9±0.49	88.0±0.33	19.2 ± 0.48	42.8±0.45
	RTN [69]	22.9 ± 0.27	88.2 ± 0.72	19.9 ± 0.54	23.2±0.49	88.1 ± 0.29	20.6 ± 0.53	43.8±0.47
Source	JAN [2]	21.8 ± 0.29	88.1 ± 0.59	22.2 ± 0.50	23.9±0.45	89.5±0.36	22.3 ± 0.46	44.6±0.44
Combine	DANN [12]	22.3 ± 0.31	87.1 ± 0.65	22.1 ± 0.47	21.0 ± 0.46	84.7±0.35	19.3 ± 0.43	42.8±0.45
	ADDA [11]	25.2 ± 0.24	87.9 ± 0.61	20.5 ± 0.46	22.0 ± 0.36	88.1 ± 0.25	20.7 ± 0.49	44.1 ±0.40
	SE [29]	19.4 ± 0.28	82.8 ± 0.68	19.3 ± 0.45	19.3±0.45	84.3 ± 0.34	18.9 ± 0.48	40.7±0.45
	MCD [13]	23.20 ± 0.3	91.2 ± 0.68	21.6 ± 0.46	25.8±0.37	86.9±0.33	23.0 ± 0.42	45.3±0.43
	M ³ SDA [16]	25.5 ± 0.26	91.6±0.63	22.2 ± 0.43	25.8±0.43	90.7±0.30	24.8 ± 0.41	46.8±0.41
Multi-	DCTN [39]	25.5 ± 0.28	93.10±0.7	22.9 ± 0.41	29.5±0.47	91.2 ± 0.29	26.5 ± 0.48	48.1±0.44
Source	Domain2Vec- α	27.8 ± 0.27	94.3±0.64	24.3 ± 0.52	27.1±0.39	89.2 ± 0.26	28.1±0.41	48.5±0.42
	Domain2Vec- <i>β</i>	28.2 ± 0.31	94.5±0.63	27.6 ± 0.41	29.3±0.39	91.5±0.26	27.2 ± 0.42	49.7±0.40
Table	3. MSI	DA resul	ts on the	TinyDA	dataset.	Our model	DOMAIN2	$VEC - \alpha$

and DOMAIN2VEC- β achieves 48.5% and 49.7% accuracy, outperforming baselines.)

From the experimental results shown in Table 2, we make several interesting observations. (i) For each sub-table, the performances of training and testing on the same domain (gray background) are better than cross-domain performance, except a few outliers (pink background, mainly between MNIST, USPS, and QMNIST). (ii) The cross-domain performance is negatively correlated with the domain distance (illustrated in Figure 2(b)). We leverage Pearson correlation coefficient (PCC) [68] to quantitatively demonstrate the negative correlation. The PCC can be computed as $\rho_{xy} = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2} \sqrt{\sum_i (y_i - \bar{y})^2}}$. We set the cross-domain performance and the domain distance as two variables. The PCC that we compute for our case is -0.774, which demonstrates that our DOMAIN2VEC successfully encodes the natural domain distance.

Multi-Source Domain Adaptation On TinyDA Our TINYDA dataset contains 54 domains. In our experiments, we consider the MSDA between digit datasets, *i.e.* MNIST, USPS, and QMNIST dataset, resulting in six MSDA settings. We choose the "grayscale" (GS) with CIFAR10 background as the target domain. For the source domains, we remove the two "grayscale" domains and leverage the remaining seven domains as the source domain.

State-of-the-art multi-source domain adaptation algorithms tackle MSDA task by adversarial alignment [39] or matching the moments of the domains [16]. However, these models neglect the effect of domain distance. We incorporate our DOMAIN2VEC model to the previous work [39,16], and devise two models, DOMAIN2VEC- α and DOMAIN2VEC- β . For DOMAIN2VEC- α , we follow the implementation of Peng *et al* [16], and separately align each source domain with the target by matching the moments of their feature distributions. For DOMAIN2VEC- β , we follow the framework and hyper-parameters of Xu *et al* [39], and separately align each source domain with the target using adversarial learning. For both frameworks, we weight each source-target alignment pairs by the distance computed using our DOMAIN2VEC model. Following [39], we compare MSDA results with two other evaluation standards: (i) *single best*, reporting the single best-preforming source transfer result on the test set, and (ii) *source com*-

Target	VisDA	Ytb BBox	PASCAL	COCO	Average				
Source Only	$53.4{\pm}0.4$	$67.2 {\pm} 0.4$	74.8 ± 0.4	$80.4{\pm}0.3$	68.9				
Openset SVM [70]	$53.9{\pm}0.5$	$68.6{\pm}0.4$	$77.7{\pm}0.4$	$82.1{\pm}0.4$	70.6				
AutoDIAL	$54.2{\pm}0.5$	$68.1{\pm}0.5$	$75.9{\pm}0.4$	$83.4{\pm}0.4$	70.4				
AODA [71]	$56.4{\pm}0.5$	$69.7{\pm}0.4$	$76.7{\pm}0.4$	$82.3{\pm}0.4$	71.3				
Domain2Vec	$56.6{\pm}0.4$	$70.6{\pm}0.4$	$81.3{\pm}0.4$	$86.8{\pm}0.4$	73.8				
Table 4. Openset domain adaption on the DOMAINBANK dataset.									

bine, combining the source domains to a single domain and performing traditional single-source single target adaptation. The high-level motivations of these two evaluation schema are: the first metric evaluates whether MSDA can boost the best single source UDA results; the second standard measures whether MSDA can outperform the trivial baseline which combines the multiple source domains as a single domain.

For both single best and source combine experiment setting, we compare with following methods: Deep Alignment Network (**DAN**) [4], Joint Adaptation Network (**JAN**) [2], Domain Adversarial Neural Network (**DANN**) [12], Residual Transfer Network (**RTN**) [69], Adversarial Deep Domain Adaptation (**ADDA**) [11], Maximum Classifier Discrepancy (**MCD**) [13], and Self-Ensembling (**SE**) [29]. For multi-source domain adaptation, we compare to Deep Cocktail Network (**DCTN**) [39] and Moment Matching for Multi-source Domain Adaptation (**M³SDA**) [16].

The experimental results are shown in Table 3. The DOMAIN2VEC- α and DOMAIN2VEC- β achieve an **48.5%** and **49.7%** average accuracy, outperforming other baselines. The results demonstrate that our models outperform the *single best* UDA results, the *combine source* results, and can boost the multi-source baselines. We argue that the performance improvement is due to the good domain embedding of our DOMAIN2VEC model.

4.3 Experiments on DomainBank

Domain Embedding Results Similar to the experiments for TINYDA dataset, we apply our devised DOMAIN2VEC model to DOMAINBANK dataset. The results are shown in Figure 3. Since our DOMAINBANK dataset is collected from multiple existing computer vision dataset, the categories of different domains in DOMAINBANK are not identical. It is not feasible to compute the cross-domain performance directly like TINYDA. However, we can still make the following interesting observations: (i) Domains with similar contents tend to form a cluster. For example, the domain containing buildings $(\hat{D}^{building})$ are close to each other in terms of the domain distance. The domains containing faces share the same property. (ii) The domains which contains artistic images are scattered in the exterior side of the embedding and are distinct from the domains which contains images in the wild. For example, the "cartoon", "syn", "quickdraw", "sketch", "logo" domains are distributed in the exterior side of the embedding space.



Fig. 3. Domain embedding results of our DOMAIN2VEC model on DOMAINBANK dataset.

Openset Domain Adaptation on DomainBank Open-set domain adaptation (ODA) considers classification when the target domain contains categories unknown (unlabelled) in the source domain. Our DOMAINBANK dataset provides a good testbed for openset domain as the categories of different domains are not identical. Since DOMAINBANK contains 56 domains, it is infeasible to explore all the (source, target) domain combinations. Instead, we demo our model on the following four transfer setting: DomainNet [16] \rightarrow VisDA [53], DomainNet \rightarrow Youtube BBox [72], DomainNet \rightarrow PASCAL [64], DomainNet \rightarrow COCO. Since Domain-Net [16] contains the largest number of domains, it is best fit as a source domain for our openset adaptation setting. The experimental results are shown in Table 4. Our model achieves **73.8%** accuracy, outperforming the compared baselines.

Partial Domain Adaptation on DomainBank In partial domain adaptation, the source domain label space is a superset of the target domain label space. Specifically, our model utilizes the idea of PADA [67], which trains a partial adversarial alignment network to tackle the partial domain adaptation task. We compute the domain distance between the sub-domains in the source training data (DomainNet) and apply the domain distance as weight in the partial adversarial alignment process. We consider the following four partial domain adaptation setting: DomainNet [16] \rightarrow VisDA [53], DomainNet \rightarrow Youtube BBox [72], DomainNet \rightarrow PASCAL [64], DomainNet \rightarrow COCO. The experimental results are shown in Table 5. Our model achieves 65.5% accuracy, outperforming the compared baselines. The results demonstrate that our model can boost the performance in partial domain adaptation setting.

4.4 Ablation Study

Our model is composed of multiple component. To demonstrate the effectiveness of each component, we perform the ablation study analysis. Table 6 shows the ablation results on TINYDA dataset. We observe that the performance drops in most of the experiments when Mutual information minimization and Gram matrix information are **not** applied. The experimental results demonstrate

Target	VisDA	Ytb BBox	PASCAL	COCO	Average				
Source Only	$34.5{\pm}0.5$	$74.3 {\pm} 0.4$	68.2 ± 0.3	$76.4{\pm}0.2$	63.3				
AdaBN	$35.1{\pm}0.5$	$75.6{\pm}0.5$	$68.2{\pm}0.4$	$78.1{\pm}0.4$	64.2				
AutoDIAL [73]	$35.2{\pm}0.6$	$74.0{\pm}0.4$	$68.5{\pm}0.4$	$77.6{\pm}0.4$	63.8				
PADA [67]	$34.2{\pm}0.6$	$76.8{\pm}0.4$	$69.7{\pm}0.3$	$77.7{\pm}0.4$	64.6				
Domain2Vec	$36.6{\pm}0.5$	$76.8{\pm}0.4$	$70.0{\pm}0.3$	$78.8{\pm}0.4$	65.5				
Table 5. Partial domain adaption on the DOMAINBANK dataset.									

target	target			$T \rightarrow QMN$	$USPS \rightarrow M$.	NIST	USPS	$\rightarrow QMNIST$	Avg		
D2V	D2V		31 94	94.5 ± 0.63		27.6 ± 0.41		29.3 ± 0.39		44.9	
D2V w/o . C	Tram	$28.5 \pm 0.$	29 92	$.4{\pm}0.5$	56	25.5 ± 0	0.29	27.	7 ± 0.26	43.5	
D2V w/o . Mutual		$26.7 \pm 0.$	27 94	94.1 ± 0.49		27.9 ± 0.35		$27.4 {\pm} 0.41$		44.0	
target	VisDA	Ytb BBox	PASCAL	coco	Avg	S VisDA	Ytb	BBox	PASCAL	coco	Avg
D2V	56.6	70.6	81.3	86.8	73.8	36.6	76	6.8	70.0	78.8	65.5
D2V w/o. Gram	54.5	68.4	80.5	85.4	72.2	2 34.5	77	7.1	65.4	77.9	63.7

D2V w/o. Mutual 55.2 69.3 81.4 85.7 72.9 35.4 73.5 67.8 77.5 63.5 **Table 6.** The ablation study results show that the Mutual information minimizing and Gram matrix information is essential for our model. The above table shows ablation experiments performed on the TINYDA dataset. The table below shows ablation experiments on DOMAINBANK dataset (openset DA on the left, partial DA on the right).

the effectiveness of the mutual information minimization and Gram matrix information.

5 Conclusion

In this paper, we have proposed a novel learning paradigm to explore the natural relations between different domains. We introduced the deep domain embedding task and proposed DOMAIN2VEC to achieve domain-to-vector mapping with joint learning of Gram Matrix of the latent representations and feature disentanglement. We have collected and evaluated two state-of-the-art domain adaptation datasets, TINYDA and DOMAINBANK. These two datasets are challenging due to the presence of notable domain gaps and a large number of domains. Extensive experiments has been conducted, both qualitatively and quantitatively, on the two benchmarks we collected to demonstrate the effectiveness of our proposed model. We also show that our model can facilitate multi-source domain adaptation, openset domain adaptation and partial domain adaptation. We hope the learning schema we proposed and the benchmarks we collected will be beneficial for the future domain adaptation research.

Acknowledgements

We thank the anonymous reviewers for their comments and suggestions. This work was partially supported by NSF and Honda Research Institute.

15

References

- Gretton, A., Smola, A.J., Huang, J., Schmittfull, M., Borgwardt, K.M.: Covariate Shift by Kernel Mean Matching. In: Dataset Shift in Machine Learning. MIT Press (2009)
- Long, M., Zhu, H., Wang, J., Jordan, M.I.: 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. (2017) 2208–2217
- Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. arXiv preprint arXiv:1412.3474 (2014)
- Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In Bach, F., Blei, D., eds.: Proceedings of the 32nd International Conference on Machine Learning. Volume 37 of Proceedings of Machine Learning Research., Lille, France, PMLR (07–09 Jul 2015) 97–105
- Sun, B., Feng, J., Saenko, K.: Return of frustratingly easy domain adaptation. In: AAAI. Volume 6. (2016) 8
- Peng, X., Saenko, K.: 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. (2018) 1982–1991
- Zellinger, W., Grubinger, T., Lughofer, E., Natschläger, T., Saminger-Platz, S.: Central moment discrepancy (CMD) for domain-invariant representation learning. CoRR abs/1702.08811 (2017)
- Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Computer Vision (ICCV), 2017 IEEE International Conference on. (2017)
- Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: CyCADA: Cycle-consistent adversarial domain adaptation. In Dy, J., Krause, A., eds.: Proceedings of the 35th International Conference on Machine Learning. Volume 80 of Proceedings of Machine Learning Research., Stockholmsmässan, Stockholm Sweden, PMLR (10–15 Jul 2018) 1989–1998
- Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: Advances in Neural Information Processing Systems. (2017) 700–708
- Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: Computer Vision and Pattern Recognition (CVPR). Volume 1. (2017) 4
- Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In Bach, F., Blei, D., eds.: Proceedings of the 32nd International Conference on Machine Learning. Volume 37 of Proceedings of Machine Learning Research., Lille, France, PMLR (07–09 Jul 2015) 1180–1189
- Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (June 2018)
- Gatys, L.A., Ecker, A.S., Bethge, M.: A neural algorithm of artistic style. arXiv preprint arXiv:1508.06576 (2015)
- Peng, X., Zijun, H., Sun, X., Saenkp, K.: Domain agnostic learning with disentangled representations. arXiv preprint arXiv:1904.12347 (2019)
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: Proceedings of the IEEE International Conference on Computer Vision. (2019) 1406–1415

- 16 Peng et al.
- 17. Saenko, K., Kulis, B., Fritz, M., Darrell, T.: Adapting visual category models to new domains. In: European conference on computer vision, Springer (2010) 213–226
- Pennington, J., Socher, R., Manning, C.: Glove: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). (2014) 1532–1543
- Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., Darrell, T.: Decaf: A deep convolutional activation feature for generic visual recognition. In: International conference on machine learning. (2014) 647–655
- Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems. (2012) 1097–1105
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR abs/1409.1556 (2014)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2016) 770–778
- Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on, IEEE (2017) 5987–5995
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2017) 4700–4708
- Achille, A., Lam, M., Tewari, R., Ravichandran, A., Maji, S., Fowlkes, C.C., Soatto, S., Perona, P.: Task2vec: Task embedding for meta-learning. In: Proceedings of the IEEE International Conference on Computer Vision. (2019) 6430–6439
- Deshmukh, A.A., Bansal, A., Rastogi, A.: Domain2vec: Deep domain generalization. arXiv preprint arXiv:1807.02919 (2018)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: Advances in Neural Information Processing Systems (NIPS). (2015)
- He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Computer Vision (ICCV), 2017 IEEE International Conference on, IEEE (2017) 2980–2988
- French, G., Mackiewicz, M., Fisher, M.: Self-ensembling for visual domain adaptation. In: International Conference on Learning Representations. (2018)
- Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. arXiv preprint arXiv:1412.3474 (2014)
- Ghifary, M., Kleijn, W.B., Zhang, M.: Domain adaptive neural networks for object recognition. In: Pacific Rim international conference on artificial intelligence, Springer (2014) 898–904
- Liu, M.Y., Tuzel, O.: Coupled generative adversarial networks. In: Advances in neural information processing systems. (2016) 469–477
- Liu, A.H., Liu, Y., Yeh, Y., Wang, Y.F.: A unified feature disentangler for multidomain image translation and manipulation. CoRR abs/1809.01361 (2018)
- Yi, Z., Zhang, H.R., Tan, P., Gong, M.: Dualgan: Unsupervised dual learning for image-to-image translation. In: ICCV. (2017) 2868–2876
- 35. Kim, T., Cha, M., Kim, H., Lee, J.K., Kim, J.: Learning to discover cross-domain relations with generative adversarial networks. In Precup, D., Teh, Y.W., eds.: Proceedings of the 34th International Conference on Machine Learning. Volume 70 of Proceedings of Machine Learning Research., International Convention Centre, Sydney, Australia, PMLR (06–11 Aug 2017) 1857–1865

Domain2Vec: Domain Embedding for Unsupervised Domain Adaptation

17

- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. Machine learning 79(1-2) (2010) 151–175
- Mansour, Y., Mohri, M., Rostamizadeh, A., R, A.: Domain adaptation with multiple sources. In Koller, D., Schuurmans, D., Bengio, Y., Bottou, L., eds.: Advances in Neural Information Processing Systems 21. Curran Associates, Inc. (2009) 1041–1048
- Crammer, K., Kearns, M., Wortman, J.: Learning from multiple sources. Journal of Machine Learning Research 9(Aug) (2008) 1757–1774
- 39. Xu, R., Chen, Z., Zuo, W., Yan, J., Lin, L.: Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 3964–3973
- Duan, L., Xu, D., Chang, S.F.: 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, IEEE (2012) 1338–1345
- Zhuang, F., Cheng, X., Luo, P., Pan, S.J., He, Q.: Supervised representation learning: Transfer learning with deep autoencoders. In: IJCAI. (2015) 4119–4125
- Mathieu, M.F., Zhao, J.J., Zhao, J., Ramesh, A., Sprechmann, P., LeCun, Y.: Disentangling factors of variation in deep representation using adversarial training. In: Advances in Neural Information Processing Systems. (2016) 5040–5048
- Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I., Frey, B.: Adversarial autoencoders. ICLR workshop (2016)
- 44. Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier GANs. In Precup, D., Teh, Y.W., eds.: Proceedings of the 34th International Conference on Machine Learning. Volume 70 of Proceedings of Machine Learning Research., International Convention Centre, Sydney, Australia, PMLR (06–11 Aug 2017) 2642–2651
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. (2014) 2672–2680
- 46. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
- Lee, H.Y., Tseng, H.Y., Huang, J.B., Singh, M., Yang, M.H.: Diverse image-toimage translation via disentangled representations. In Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y., eds.: Computer Vision – ECCV 2018, Cham, Springer International Publishing (2018) 36–52
- Belghazi, M.I., Baratin, A., Rajeshwar, S., Ozair, S., Bengio, Y., Courville, A., Hjelm, D.: Mutual information neural estimation. In Dy, J., Krause, A., eds.: Proceedings of the 35th International Conference on Machine Learning. Volume 80 of Proceedings of Machine Learning Research., Stockholmsmässan, Stockholm Sweden, PMLR (10–15 Jul 2018) 531–540
- Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research 9(Nov) (2008) 2579–2605
- Kiefer, J., Wolfowitz, J., et al.: Stochastic estimation of the maximum of a regression function. The Annals of Mathematical Statistics 23(3) (1952) 462–466
- 51. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: (IEEE) Conference on Computer Vision and Pattern Recognition (CVPR). (2017)

- 18 Peng et al.
- 53. Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., Saenko, K.: Visda: The visual domain adaptation challenge. arXiv preprint arXiv:1710.06924 (2017)
- LeCun, Y., Boser, B., Denker, J., Henderson, D., Howard, R., Hubbard, W., Jackel, L.: Backpropagation applied to handwritten zip code recognition. Neural Computation (1989)
- 55. Hull, J.J.: A database for handwritten text recognition research. IEEE Transactions on pattern analysis and machine intelligence **16**(5) (1994) 550–554
- Cohen, G., Afshar, S., Tapson, J., van Schaik, A.: Emnist: an extension of mnist to handwritten letters. arXiv preprint arXiv:1702.05373 (2017)
- Clanuwat, T., Bober-Irizar, M., Kitamoto, A., Lamb, A., Yamamoto, K., Ha, D.: Deep learning for classical japanese literature. arXiv preprint arXiv:1812.01718 (2018)
- Yadav, C., Bottou, L.: Cold case: The lost mnist digits. arXiv preprint arXiv:1905.10498 (2019)
- Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv:1708.07747 (2017)
- Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. Technical report, Citeseer (2009)
- Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. IEEE Trans. Pattern Anal. Mach. Intell. 33(5) (May 2011) 898–916
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: European conference on computer vision, Springer (2014) 740–755
- 63. Griffin, G., Holub, A., Perona, P.: Caltech-256 object category dataset. (2007)
- Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. International Journal of Computer Vision 88(2) (June 2010) 303–338
- 65. Busto, P.P., Gall, J.: Open set domain adaptation. In: The IEEE International Conference on Computer Vision (ICCV). Volume 1. (2017)
- Busto, P.P., Iqbal, A., Gall, J.: Open set domain adaptation for image and action recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence (2018)
- Cao, Z., Ma, L., Long, M., Wang, J.: Partial adversarial domain adaptation. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 135–150
- Benesty, J., Chen, J., Huang, Y., Cohen, I.: Pearson correlation coefficient. In: Noise reduction in speech processing. Springer (2009) 1–4
- Long, M., Zhu, H., Wang, J., Jordan, M.I.: Unsupervised domain adaptation with residual transfer networks. In: Advances in Neural Information Processing Systems. (2016) 136–144
- Jain, L.P., Scheirer, W.J., Boult, T.E.: Multi-class open set recognition using probability of inclusion. In: European Conference on Computer Vision, Springer (2014) 393–409
- Saito, K., Yamamoto, S., Ushiku, Y., Harada, T.: Open set domain adaptation by backpropagation. CoRR abs/1804.10427 (2018)
- 72. Real, E., Shlens, J., Mazzocchi, S., Pan, X., Vanhoucke, V.: Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2017) 5296–5305

Domain2Vec: Domain Embedding for Unsupervised Domain Adaptation

 Cariucci, F.M., Porzi, L., Caputo, B., Ricci, E., Bulo, S.R.: Autodial: Automatic domain alignment layers. In: 2017 IEEE International Conference on Computer Vision (ICCV), IEEE (2017) 5077–5085