

Domain Generalization for Object Recognition with Multi-task Autoencoders

Muhammad Ghifary W. Bastiaan Kleijn Mengjie Zhang David Balduzzi
Victoria University of Wellington

{muhammad.ghifary,bastiaan.kleijn,mengjie.zhang}@ecs.vuw.ac.nz, david.balduzzi@vuw.ac.nz

Abstract

The problem of domain generalization is to take knowledge acquired from a number of related domains where training data is available, and to then successfully apply it to previously unseen domains. We propose a new feature learning algorithm, Multi-Task Autoencoder (MTAE), that provides good generalization performance for cross-domain object recognition.

Our algorithm extends the standard denoising autoencoder framework by substituting artificially induced corruption with naturally occurring inter-domain variability in the appearance of objects. Instead of reconstructing images from noisy versions, MTAE learns to transform the original image into analogs in multiple related domains. It thereby learns features that are robust to variations across domains. The learnt features are then used as inputs to a classifier.

We evaluated the performance of the algorithm on benchmark image recognition datasets, where the task is to learn features from multiple datasets and to then predict the image label from unseen datasets. We found that (denoising) MTAE outperforms alternative autoencoder-based models as well as the current state-of-the-art algorithms for domain generalization.

1. Introduction

Recent years have seen dramatic advances in object recognition by deep learning algorithms [23, 11, 32]. Much of the increased performance derives from applying large networks to massive labeled datasets such as PASCAL VOC [14] and ImageNet [22]. Unfortunately, *dataset bias* – which can include factors such as backgrounds, camera viewpoints and illumination – often causes algorithms to generalize poorly *across* datasets [35] and significantly limits their usefulness in practical applications. Developing algorithms that are *invariant* to dataset bias is therefore a compelling problem.

Problem definition. In object recognition, the “visual world” can be considered as decomposing into *views* (e.g. perspectives or lighting conditions) corresponding to do-

main. For example, frontal-views and 45° rotated-views correspond to two different domains. Alternatively, we can associate views or domains with standard image datasets such as PASCAL VOC2007 [14], and Office [31].

The problem of learning from multiple *source* domains and testing on unseen *target* domains is referred to as *domain generalization* [6, 26]. A *domain* is a probability distribution \mathbb{P}_k from which samples $\{x_i, y_i\}_{i=1}^{N_k}$ are drawn. Source domains provide training samples, whereas distinct target domains are used for testing. In the standard supervised learning framework, it is assumed that the source and target domains coincide. Dataset bias becomes a significant problem when training and test domains differ: applying a classifier trained on one dataset to images sampled from another typically results in poor performance [35, 18]. The goal of this paper is to learn features that improve generalization performance across domains.

Contribution. The challenge is to build a system that recognizes objects in previously *unseen* datasets, given one or multiple training datasets. We introduce Multi-task Autoencoder (MTAE), a feature learning algorithm that uses a multi-task strategy [8, 34] to learn unbiased object features, where the *task* is the data reconstruction.

Autoencoders were introduced to address the problem of ‘backpropagation without a teacher’ by using *inputs as labels* – and learning to reconstruct them with minimal distortion [28, 5]. Denoising autoencoders in particular are a powerful basic circuit for unsupervised representation learning [36]. Intuitively, corrupting inputs forces autoencoders to learn representations that are robust to noise.

This paper proposes a broader view: that autoencoders are *generic circuits for learning invariant features*. The main contribution is a new training strategy based on naturally occurring transformations such as: rotations in viewing angle, dilations in apparent object size, and shifts in lighting conditions. The resulting Multi-Task Autoencoder learns features that are robust to real-world image variability, and therefore generalize well across domains. Extensive experiments show that MTAE with a denoising criterion outperforms the prior state-of-the-art in domain generalization over various cross-dataset recognition tasks.

2. Related work

Domain generalization has recently attracted attention in classification tasks, including automatic gating of flow cytometry data [6, 26] and object recognition [16, 21, 38]. Khosla *et al.* [21] proposed a multi-task max-margin classifier, which we refer to as Undo-Bias, that explicitly encodes dataset-specific biases in feature space. These biases are used to push the dataset-specific weights to be similar to the global weights. Fang *et al.* [16] developed Unbiased Metric Learning (UML) based on learning to rank framework. Validated on weakly-labeled web images, UML produces a less biased distance metric that provides good object recognition performance. and validated on weakly-labeled web images. More recently, Xu *et al.* [38] extended an exemplar-SVM to domain generalization by adding a nuclear norm-based regularizer that captures the likelihoods of all positive samples. The proposed model is denoted by LRE-SVM.

Other works in object recognition exist that address a similar problem, in the sense of having unknown targets, where the unseen dataset contains noisy images that are not in the training set [17, 33]. However, these were designed to be noise-specific and may suffer from dataset bias when observing objects with different types of noise.

A closely related task to domain generalization is domain adaptation, where unlabeled samples from the target dataset are available during training. Many domain adaptation algorithms have been proposed for object recognition (see, *e.g.*, [2, 31]). Domain adaptation algorithms are not readily applicable to domain generalization, since no information is available about the target domain.

Our proposed algorithm is based on the feature learning approach. Feature learning has been of a great interest in the machine learning community since the emergence of deep learning (see [4] and references therein). Some feature learning methods have been successfully applied to domain adaptation or transfer learning applications [9, 13]. To our best knowledge, there is no prior work along these lines on the more difficult problem of domain generalization, *i.e.*, to create useful representations without observing the target domain.

3. The Proposed Approach

Our goal is to learn features that provide good domain generalization. To do so, we extend the autoencoder [7] into a model that jointly learns multiple data-reconstruction tasks taken from related domains. Our strategy is motivated by prior work demonstrating that learning from multiple related tasks can improve performance on a novel, yet related, task – relative to methods trained on a single-task [1, 3, 8, 34].

3.1. Autoencoders

Autoencoders (AE) have become established as a pre-training model for deep learning [5]. The autoencoder training consists of two stages: 1) *encoding* and 2) *decoding*. Given an unlabeled input $\mathbf{x} \in \mathbb{R}^{d_x}$, a single hidden layer autoencoder $f_{\Theta}(\mathbf{x}) : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$ can be formulated as

$$\begin{aligned} \mathbf{h} &= \sigma_{\text{enc}}(\mathbf{W}^{\top} \mathbf{x}) \\ \hat{\mathbf{x}} &= \sigma_{\text{dec}}(\mathbf{V}^{\top} \mathbf{h}) = f_{\Theta}(\mathbf{x}), \end{aligned} \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{d_x \times d_y}$, $\mathbf{V} \in \mathbb{R}^{d_y \times d_x}$ are *input-to-hidden* and *hidden-to-output* connection weights¹ respectively, $\mathbf{h} \in \mathbb{R}^{d_h}$ is the hidden node vector, and $\sigma_{\text{enc}}(\cdot) = [s_{\text{enc}}(z_1), \dots, s_{\text{enc}}(z_{d_h})]^{\top}$, $\sigma_{\text{dec}}(\cdot) = [s_{\text{dec}}(z_1), \dots, s_{\text{dec}}(z_{d_x})]^{\top}$ are element-wise non-linear activation functions, and s_{enc} and s_{dec} are not necessarily identical. Popular choices for the activation function $s(\cdot)$ are, *e.g.*, the sigmoid $s(a) = (1 + \exp(-a))^{-1}$ and the rectified linear (ReLU) $s(a) = \max(0, a)$.

Let $\Theta = \{\mathbf{W}, \mathbf{V}\}$ be the autoencoder parameters and $\{\mathbf{x}_i\}_{i=1}^N$ be a set of N input data. Learning corresponds to minimizing the following objective

$$\hat{\Theta} := \arg \min_{\Theta} \sum_{i=1}^N \mathcal{L}(f_{\Theta}(\mathbf{x}_i), \mathbf{x}_i) + \eta \mathcal{R}(\Theta), \quad (2)$$

where $\mathcal{L}(\cdot, \cdot)$ is the loss function, usually in the form of *least square* or *cross-entropy* loss, and $\mathcal{R}(\cdot)$ is a regularization term used to avoid overfitting. The objective (2) can be optimized by the backpropagation algorithm [29]. If we apply autoencoders to raw pixels of visual object images, the weights \mathbf{W} usually form visually meaningful “filters” that can be interpreted qualitatively.

To create a discriminative model using the learnt autoencoder model, either of the following options can be considered: 1) the feature map $\phi(\mathbf{x}) := \sigma_{\text{enc}}(\tilde{\mathbf{W}}^{\top} \mathbf{x})$ is extracted and used as an input to supervised learning algorithms while keeping the weight matrix $\tilde{\mathbf{W}}$ fixed; 2) the learnt weight matrix $\tilde{\mathbf{W}}$ is used to initialize a neural network model and is updated during the supervised neural network training (*fine-tuning*).

Recently, several variants such as denoising autoencoders (DAE) [37] and contractive autoencoders (CAE) [27] have been proposed to extract features that are more robust to small changes of the input. In DAEs, the objective is to reconstruct a *clean* input \mathbf{x} given its *corrupted* counterpart $\tilde{\mathbf{x}} \sim Q(\tilde{\mathbf{x}}|\mathbf{x})$. Commonly used types of corruption are zero-masking, Gaussian, and salt-and-pepper noise. Features extracted by DAE have been proven to be more discriminative than ones extracted by AE [37].

¹While the bias terms are incorporated in our experiments, they are intentionally omitted from equations for the sake of simplicity.




Figure 1. The Multi-task Autoencoder (MTAE) architecture, which consists of three layers with multiple separated outputs; each output corresponds to one task/domain.

3.2. Multi-task Autoencoders

We refer to our proposed domain generalization algorithm as *Multi-task Autoencoder (MTAE)*. From an architectural viewpoint, MTAE is an autoencoder with multiple output layers, see Fig. 1. The input-hidden weights represent *shared* parameters and the hidden-output weights represent *domain-specific* parameters. The architecture is similar to the supervised multi-task neural networks proposed by Caruana [8]. The main difference is that the output layers of MTAE correspond to different domains instead of different class labels.

The most important component of MTAE is the training strategy, which constructs a generalized denoising autoencoder that learns invariances to naturally occurring transformations. Denoising autoencoders focus on the special case where the transformation is simply noise. In contrast, MTAE training treats a specific perspective on an object as the “corrupted” counterpart of another perspective (*e.g.*, a rotated digit 6 is the noisy pair of the original digit). The autoencoder objective is then reformulated along the lines of multi-task learning: the model aims to *jointly achieve good reconstruction of all source views given a particular view*. For example, applying the strategy to handwritten digit images with several views, MTAE learns representations that are invariant across the source views, see Section 4.

Two types of reconstruction tasks are performed during MTAE training: 1) *self-domain* reconstruction and 2) *between-domain* reconstruction. Given M source domains, there are $M \times M$ reconstruction tasks, of which M task are self-domain reconstructions and the remaining $M \times (M - 1)$ tasks are between-domain reconstructions. Note that the self-domain reconstruction is identical to the standard autoencoder reconstruction (1).

Formal description. Let $\{\mathbf{x}_i^l\}_{i=1}^{n_l}$ be a set of d_x -dimensional data points in the l^{th} domain, where $l \in \{1, \dots, M\}$. Each domain’s data points are combined into a matrix $\mathbf{X}^l \in \mathbb{R}^{n_l \times d_x}$, where $\mathbf{x}_i^{l\top}$ is its i^{th} row, such

that $(\mathbf{x}_i^1, \mathbf{x}_i^2, \dots, \mathbf{x}_i^M)$ form a category-level correspondence. This configuration enforces the number of samples in a category to be the same in every domain. Note that such a configuration is necessary to ensure that the *between-domain* reconstruction works (we will discuss how to handle the case with unbalanced samples in Section 3.3). The input and output pairs used to train MTAE can then be written as concatenated matrices

$$\begin{aligned} \bar{\mathbf{X}} &= [\mathbf{X}^1; \mathbf{X}^2; \dots; \mathbf{X}^M], \\ \bar{\mathbf{X}}^l &= [\mathbf{X}^l; \mathbf{X}^l; \dots; \mathbf{X}^l] \end{aligned} \quad (3)$$

where $\bar{\mathbf{X}}, \bar{\mathbf{X}}^l \in \mathbb{R}^{N \times d_x}$ and $N = \sum_{l=1}^M n_l$. In words, $\bar{\mathbf{X}}$ is the matrix of data points taken from all domains and $\bar{\mathbf{X}}^l$ is the matrix of replicated data sets taken from the l^{th} domain. The replication imposed in $\bar{\mathbf{X}}^l$ constructs input-output pairs for the autoencoder learning algorithm. In practice, the algorithm can be implemented efficiently – without replicating the matrix in memory.

We now describe MTAE more formally. Let $\bar{\mathbf{x}}_i^\top$ and $\bar{\mathbf{x}}_i^{l\top}$ be the i^{th} row of matrices $\bar{\mathbf{X}}$ and $\bar{\mathbf{X}}^l$, respectively, the feed-forward MTAE reconstruction is

$$\begin{aligned} \mathbf{h}_i &= \sigma_{\text{enc}}(\mathbf{W}^\top \bar{\mathbf{x}}_i), \\ f_{\Theta^{(l)}}(\bar{\mathbf{x}}_i) &= \sigma_{\text{dec}}(\mathbf{V}^{(l)\top} \mathbf{h}_i), \end{aligned} \quad (4)$$

where $\Theta^{(l)} = \{\mathbf{W}, \mathbf{V}^{(l)}\}$ contains the matrices of shared and individual weights, respectively.

The MTAE training is achieved as follows. Let us define the loss function summed over the datapoints

$$J(\Theta^{(l)}) = \sum_{i=1}^N \mathcal{L}(f_{\Theta^{(l)}}(\bar{\mathbf{x}}_i), \bar{\mathbf{x}}_i^l). \quad (5)$$

Given M domains, training MTAE corresponds to minimizing the objective

$$\hat{\Theta}^{(l)} := \arg \min_{\Theta^{(l)}} \sum_{l=1}^M J(\Theta^{(l)}) + \eta \mathcal{R}(\Theta^{(l)}), \quad (6)$$

where $\mathcal{R}(\Theta^{(l)})$ is a regularization term. In this work, we use the standard l_2 -norm weight penalty $\mathcal{R}(\Theta^{(l)}) = \|\mathbf{W}\|_2^2 + \|\mathbf{V}^{(l)}\|_2^2$. Stochastic gradient descent is applied on each reconstruction task to achieve the objective (6). Once training is completed, the optimal *shared* weights $\hat{\mathbf{W}}$ are obtained. The stopping criterion is empirically determined by monitoring the average loss over all reconstruction tasks during training – the process is stopped when the average loss stabilizes. The detailed steps of MTAE training is summarized in Algorithm 1.

The training protocol can be supplemented with a denoising criterion as in [37] to induce more robust-to-noise features. To do so, simply replace $\bar{\mathbf{x}}_i$ in (4) with its corrupted counterpart $\tilde{\bar{\mathbf{x}}}_i \sim Q(\tilde{\bar{\mathbf{x}}}_i | \bar{\mathbf{x}}_i)$. We name the MTAE model after applying the denoising criterion the *Denoising Multi-task Autoencoder (D-MTAE)*.

Algorithm 1 The MTAE feature learning algorithm.

Input:

- Data matrices based on (3): $\bar{\mathbf{X}}$ and $\bar{\mathbf{X}}^l, \forall l \in \{1, \dots, M\}$;
 - Source labels: $\{y_i^l\}_{i=1}^{n_l}, \forall l \in \{1, \dots, M\}$;
 - The learning rate: α ;
- 1: Initialize $\mathbf{W} \in \mathbb{R}^{d_x \times d_h}$ and $\mathbf{V}^{(l)} \in \mathbb{R}^{d_h \times d_x}, \forall l \in \{1, \dots, M\}$ with small random real values;
 - 2: **while** not end of epoch **do**
 - 3: Do RAND-SEL as described in Section 3.3 to balance the number of samples per categories in $\bar{\mathbf{X}}$ and $\bar{\mathbf{X}}^l$;
 - 4: **for** $l = 1$ **to** M **do**
 - 5: **for all** row of $\bar{\mathbf{X}}$ **do**
 - 6: Do a forward pass based on (4);
 - 7: Update \mathbf{W} and $\mathbf{V}^{(l)}$ to achieve the objective (6) with respect to the following rules

$$V_{ij}^{(l)} \leftarrow V_{ij}^{(l)} - \alpha \frac{\partial J(\{\mathbf{W}, \mathbf{V}^{(l)}\})}{\partial V_{ij}^{(l)}},$$

$$W_{ij} \leftarrow W_{ij} - \alpha \frac{\partial J(\{\mathbf{W}, \mathbf{V}^{(l)}\})}{\partial W_{ij}};$$

- 8: **end for**
- 9: **end for**
- 10: **end while**

Output:

- MTAE learnt weights: $\hat{\mathbf{W}} \forall l \in \{1, \dots, M\}$;
-

3.3. Handling unbalanced samples per category

MTAE requires that every instance in a particular domain has a category-level corresponding pair in every other domain. MTAE’s apparent applicability is therefore limited to situations where the number of source samples per category is the same in every domain. However, unbalanced samples per category occur frequently in applications. To overcome this issue, we propose a simple *random selection* procedure applied in the *between-domain* reconstructions, denoted by RAND-SEL, which is simply balancing the samples per category while keeping their category-level correspondence.

In detail, the RAND-SEL strategy is as follows. Let m_c be the number of subsamples in the c -th category, where $m_c = \min(n_{1c}, n_{2c}, \dots, n_{Mc})$ and n_{lc} is the number of samples in the c -th category of domain $l \in \{1, \dots, M\}$. For each category c and each domain l , select m_c samples randomly such that $n_{lc} = n_{2c} = \dots n_{Mc} = m_c$. This procedure is executed in every iteration of the MTAE algorithm, see Line 3 of Algorithm 1.

4. Experiments and Results

We conducted experiments on several real world object datasets to evaluate the domain generalization ability of our proposed system. In Section 4.1, we investigate the behaviour of MTAE in comparison to standard single-task autoencoder models on raw pixels as proof-of-principle. In Section 4.2, we evaluate the performance of MTAE against several state-of-the-art algorithms on mod-

ern object datasets such as the Office [31], Caltech [20], PASCAL VOC2007 [14], LabelMe [30], and SUN09 [10].

4.1. Cross-recognition on the MNIST and ETH-80 datasets

In this part, we aim to understand MTAE’s behavior when learning from multiple domains that form physically reasonable object transformations such as roll, pitch rotation, and dilation. The task is to categorize objects in views (domains) that were not presented during training. We evaluate MTAE against several autoencoder models. To perform the evaluation, a variety of object views were constructed from the MNIST handwritten digit [24] and ETH-80 object [25] datasets.

Data setup. We created four new datasets from MNIST and ETH-80 images: 1) MNIST-r, 2) MNIST-s, 3) ETH80-p, and 4) ETH80-y. These new sets contain multiple domains so that every instance in one domain has a pair in another domain. The detailed setting for each dataset is as follows.

MNIST-r contains six domains, each corresponding to a degree of roll rotation. We randomly chose 1000 digit images of ten classes from the original MNIST training set to represent the *basic* view, i.e., 0 degree of rotation;² each class has 100 images. Each image was subsampled to a 16×16 representation to simplify the computation. This subset of 1000 images is denoted by M . We then created 5 rotated views from M with 15° difference in counterclockwise direction, denoted by $M_{15^\circ}, M_{30^\circ}, M_{45^\circ}, M_{60^\circ}$, and M_{75° . The **MNIST-s** is the counterpart of MNIST-r, where each domain corresponds to a dilation factor. The views are denoted by $M, M_{*0.9}, M_{*0.8}, M_{*0.7}$, and $M_{*0.6}$, where the subscripts represent the dilation factors with respect to M .

The **ETH80-p** consists of eight object classes with 10 subcategories for each class. In each subcategory, there are 41 different views with respect to pose angles. We took five views from each class denoted by $E_{p0^\circ}, E_{p22^\circ}, E_{p45^\circ}, E_{p68^\circ}$, and E_{p90° , which represent the horizontal poses, i.e., pitch-rotated views starting from the top view to the side view. This makes the number of instances only 80 for each view. We then greyscaled and subsampled the images to 28×28 . The **ETH80-y** contains five views of the ETH-80 representing the vertical poses, i.e., yaw-rotated views starting from the right-side view to the left-side view denoted by $E_{+y90^\circ}, E_{+y45^\circ}, E_{y0^\circ}, E_{-y45^\circ}$, and E_{-y90° . Other settings such as the image dimensionality and preprocessing stage are similar to ETH80-p. Examples of the resulting views are depicted in Fig. 2.

Baselines. We compared the classification performance of our models with several single-task autoencoder mod-

²Note that the rotation angle of the basic view is not perfectly 0° since the original MNIST images have varying appearances.

Table 1. The *leave-one-domain-out* classification accuracies % on the MNIST-r and MNIST-s. Bold-red and bold-black indicate the best and second best performance.

Source	Target	Raw	AE	DAE	CAE	uDICA	MTAE	D-MTAE
MNIST-r leave-one-roll-rotation-out								
$M_{15^\circ}, M_{30^\circ}, M_{45^\circ}, M_{60^\circ}, M_{75^\circ}$	M	52.40	74.20	76.90	72.10	67.20	77.90	82.50
$M, M_{30^\circ}, M_{45^\circ}, M_{60^\circ}, M_{75^\circ}$	M_{15°	74.10	93.20	93.20	95.30	87.80	95.70	96.30
$M, M_{15^\circ}, M_{45^\circ}, M_{60^\circ}, M_{75^\circ}$	M_{30°	71.40	89.90	91.30	92.60	88.80	91.20	93.40
$M, M_{15^\circ}, M_{30^\circ}, M_{60^\circ}, M_{75^\circ}$	M_{45°	61.40	82.20	81.10	81.50	77.80	77.30	78.60
$M, M_{15^\circ}, M_{30^\circ}, M_{45^\circ}, M_{75^\circ}$	M_{60°	67.40	90.00	92.80	92.70	84.20	92.40	94.20
$M, M_{15^\circ}, M_{30^\circ}, M_{45^\circ}, M_{60^\circ}$	M_{75°	55.40	73.80	76.50	79.30	69.50	79.90	80.50
Average		63.68	83.88	85.30	85.58	79.22	85.73	87.58
MNIST-s leave-one-dilation-out								
$M_{*0.9}, M_{*0.8}, M_{*0.7}, M_{*0.6}$	M	54.00	67.50	71.80	75.80	75.80	74.50	76.00
$M, M_{*0.8}, M_{*0.7}, M_{*0.6}$	$M_{*0.9}$	80.40	95.10	94.00	94.90	88.60	97.80	98.00
$M, M_{*0.9}, M_{*0.7}, M_{*0.6}$	$M_{*0.8}$	82.60	94.60	92.90	94.90	86.60	96.30	96.40
$M, M_{*0.9}, M_{*0.8}, M_{*0.6}$	$M_{*0.7}$	78.20	93.70	91.60	92.50	87.40	95.80	94.90
$M, M_{*0.9}, M_{*0.8}, M_{*0.7}$	$M_{*0.6}$	64.70	74.80	76.10	77.50	75.30	78.00	78.30
Average		71.98	85.14	85.28	87.12	82.74	88.48	88.72




Figure 2. Some image examples from the MNIST-r, MNIST-s, and ETH80-p.

els: Descriptions of the methods and their hyperparameter settings are provided below.

- **AE** [5]: the standard *autoencoder* model trained by stochastic gradient descent, where all object views were concatenated as one set of inputs. The number of hidden nodes was fixed at 500 on the MNIST dataset and at 1000 on the ETH-80 dataset. The learning rate, weight decay penalty, and number of iterations were empirically determined at 0.1, 3×10^{-4} , and 100, respectively.
- **DAE** [37]: the *denoising autoencoder* with zero-masking noise, where all object views were concatenated as one set of input data. The corruption level was fixed at 30% for all cases. Other hyper-parameter values were identical to AE.
- **CAE** [27]: the autoencoder model with the Jacobian matrix norm regularization referred to as the *contractive autoencoder*.

The corresponding regularization constant λ was set at 0.1.

- **MTAE**: our proposed multi-task autoencoder with identical learning settings as AE, except for the learning rate set at 0.03, which was also chosen empirically. This value provides a lower reconstruction error for each task and visually clearer first layer weights.
- **D-MTAE**: MTAE with a denoising criterion. The learning rate was set the same as MTAE; other hyper-parameters followed DAE.
- **uDICA**: the unsupervised Domain-Invariant Component Analysis [26].

We also did experiments using the supervised variant, DICA. Surprisingly, the peak performance of uDICA is consistently higher than DICA. A possible explanation is that the Dirac kernel function measuring the label similarity is less appropriate in this application.

We normalized the raw pixels to a range of $[0, 1]$ for autoencoder-based models and l_2 -unit ball for uDICA. We evaluated the classification accuracies of the learnt features using multi-class SVM with linear kernel (L-SVM) [12]. Using a linear kernel keeps the classifier simple – since our main focus is on the feature extraction process. The LIBLINEAR package [15] was used to run the L-SVM.

Cross-domain recognition results. We evaluated the object classification accuracies of each algorithm by *leave-one-domain-out* test, *i.e.*, taking one domain as the test set and the remaining domains as the training set. For all autoencoder-based algorithms, we repeated the experiments on each *leave-one-domain-out* case 30 times and reported the average accuracies. The standard deviations are not reported since they are small (± 0.1).

The detailed results on the MNIST-r and MNIST-s can be seen in Table 1. On average, MTAE has the second best classification accuracies, and in particular outperforms

single-task autoencoder models. This indicates that the multi-task feature learning strategy can provide better discriminative features than the single-task feature learning w.r.t. unseen object views.

The algorithm with the best performance is on these datasets is D-MTAE. Specifically, D-MTAE performs best on average and also on 9 out of 11 individual cross-domain cases of the MNIST-r and MNIST-s. The closest single-task feature learning competitor to D-MTAE is CAE. This suggests that the denoising criterion strongly benefits domain generalization. The denoising criterion is also useful for single-task feature learning although it does not yield competitive accuracies, see AE and DAE performance.

We also obtain a consistent trend on the ETH80-p and ETH80-y datasets, *i.e.*, D-MTAE and MTAE are the best and second best models. In detail, D-MTAE and MTAE produce the average accuracies of 87.85% and 87.50% on the ETH80-p, and 97% and 96.50% on the ETH80-y.

Observe that there is an anomaly in the MNIST-r dataset: the performance on M_{45° is worse than its neighbors ($M_{30^\circ}, M_{60^\circ}$). This anomaly appears to be related to the geometry of the MNIST-r digits. We found that the most frequently misclassified digits are 4, 6, and 9 on M_{45° , which rarely occurs on other MNIST-r’s domains – typically 4 as 9, 6 as 4, and 9 as 8.

Weight visualization. Useful insight is obtained from considering the qualitative outcome of the MTAE training by visualizing the first layer weights. Figure 4 depicts the weights of some autoencoder models, including ours, on the MNIST-r dataset. *Both MTAE and D-MTAE’s weights form “filters” that tend to capture the underlying transformation across the MNIST-r views, which is the rotation.* This effect is unseen in AE and DAE, the filters of which only explain the *contents* of handwritten digits in the form of Fourier component-like descriptors such as local blob detectors and stroke detectors [37]. This might be a reason that MTAE and D-MTAE features can provide better domain generalization than AE and DAE, since they implicitly capture the relationship among the source domains.

Next we discuss the difference between MTAE and D-MTAE filters. The D-MTAE filters not only capture the object transformation, but also produce features that describe the object contents more distinctively. These filters basically combine both properties of the DAE and MTAE filters that might benefit the domain generalization.

Invariance analysis. A possible explanation for the effectiveness of MTAE relates to the dimensionality of the manifold in feature space where samples concentrate. We hypothesize that if features concentrate near a low-dimensional submanifold, then the algorithm has found simple invariant features and will generalize well.

To test the hypothesis, we examine the singular value

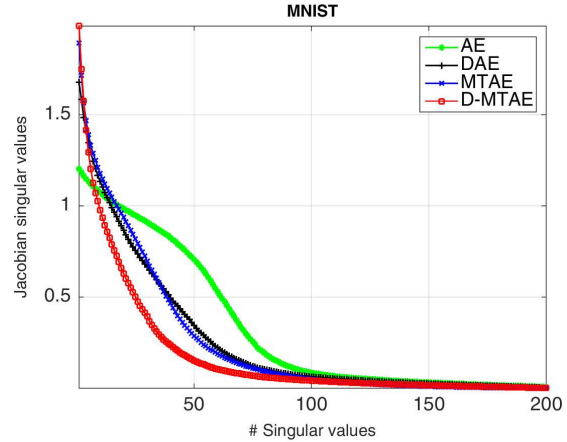


Figure 3. The average singular value spectrum of the Jacobian matrix over the MNIST-r and MNIST-s datasets.

spectrum of the Jacobian matrix $\mathcal{J}_x(\mathbf{z}) = \left[\frac{\partial z_i}{\partial x_j} \right]_{ij}$, where \mathbf{x} and \mathbf{z} are the input and feature vectors respectively [27]. The spectrum describes the local dimensionality of the manifold around which samples concentrate. If the spectrum decays rapidly, then the manifold is locally of low dimension.

Figure 3 depicts the average singular value spectrum on test samples from MNIST-r and MNIST-s. The spectrum of D-MTAE decays the most rapidly, followed by MTAE and then DAE (with similar rates), and AE decaying the slowest. The ranking of decay rates of the four algorithms matches their ranking in terms of empirical performance in Table 1. Figure 3 thus provides partial confirmation for our hypothesis. However, a more detailed analysis is necessary before drawing strong conclusions.

4.2. Cross-recognition on the Office, Caltech, VOC2007, LabelMe, and SUN09 datasets

In the second set of experiments, we evaluated the cross-recognition performance of the proposed algorithms on modern object datasets. The aim is to show that MTAE and D-MTAE are applicable and competitive in the more general setting. We used the Office, Caltech, PASCAL VOC2007, LabelMe, and SUN09 datasets from which we formed two cross-domain datasets. Our general strategy is to extend the generalization of features extracted from the current best deep convolutional neural network [23].

Data Setup. The first cross-domain dataset consists of images from PASCAL VOC2007 (V), LabelMe (L), Caltech-101 (C), and SUN09 (S) datasets, each of which represents one domain. C is an object-centric dataset, while V, L, and S are scene-centric. This dataset, which we abbreviate as VLCS, shares five object categories: ‘bird’, ‘car’, ‘chair’, ‘dog’, and ‘person’. Each domain in the VLCS dataset was divided into a training set (70%) and a test set (30%) by random selection from the overall dataset. The




Figure 4. The 2D visualization of 100 randomly chosen weights after pretraining on the MNIST-r dataset. Each patch corresponds to a row of the learnt weight matrix \mathbf{W} that represents a “filter”. The weight value $w_{ij} \geq 3$ is depicted with white, $w_{ij} \leq -3$ is depicted with black, otherwise it is gray.

detailed training-test configuration for each domain is summarized in Table 2. Instead of using the raw features directly, we employed the DeCAF₆ features [13] as inputs to the algorithms. These features have dimensionality of 4,096 and are publicly available.³

The second cross-domain dataset is referred to as the **Office+Caltech** [31, 19] dataset that contains four domains: Amazon (A), Webcam (W), DSLR (D), and Caltech-256 (C), which share ten common categories. This dataset has 8 to 151 instances per category per domain, and 2,533 instances in total. We also used the DeCAF₆ features extracted from this dataset, which are also publicly available.⁴

Table 2. The number of training and test instances for each domain in the VLCS dataset.

Domain	VOC2007	LabelMe	Caltech-101	SUN09
#training	2,363	1,859	991	2,297
#test	1,013	797	424	985

Table 3. The groundtruth L-SVM accuracies % on the standard training-test evaluation. The left-most column indicates the training set, while the upper-most row indicates the test set.

Training/Test	VOC2007	LabelMe	Caltech-101	SUN09
VOC2007	66.34	34.50	65.09	52.49
LabelMe	44.03	68.76	43.87	41.02
Caltech-101	52.81	32.37	95.99	39.29
SUN09	52.42	42.03	40.33	74.21

Training protocol. On these datasets, we utilized the MTAE or D-MTAE learning as *pretraining* for a fully-connected neural network with one hidden layer (1HNN). The number of hidden nodes was set at 2,000, which is less than the input dimensionality. In the pretraining stage, the number of output layers was the same as the number of source domains – each corresponds to a particular source domain. The sigmoid activation and linear activation functions were used for $\sigma_{\text{enc}}(\cdot)$ and $\sigma_{\text{dec}}(\cdot)$.

³http://www.cs.dartmouth.edu/~chenfang/proj_page/FXR_iccv13/index.php

⁴http://vc.sce.ntu.edu.sg/transfer_learning_domain_adaptation/

The MTAE pretraining was run with the learning rate at 5×10^{-4} , the number of epochs at 500, and the batch size at 10, which were empirically determined w.r.t. the smallest average reconstruction loss. D-MTAE has the same hyper-parameter setting as MTAE except the additional zero-masking corruption level at 20%. After the pretraining is completed, we then performed back-propagation *fine-tuning* using 1HNN with softmax output, where the first layer weights were initialized by either the MTAE or D-MTAE learnt weights. The supervised learning hyper-parameters were tuned using 10-fold cross validation (10FCV) on source domains. We denote the overall models by **MTAE+1HNN** and **D-MTAE+1HNN**.

Baselines. We compared our proposed models with six baselines:

1. **L-SVM**: an SVM with linear kernel.
2. **1HNN**: a single hidden layer neural network without pretraining.
3. **DAE+1HNN**: a two-layer neural network with denoising autoencoder pretraining (DAE+1HNN).
4. **Undo-Bias** [21]: a multi-task SVM-based algorithm for undoing dataset bias. Three hyper-parameters (λ, C_1, C_2) require tuning by 10FCV.
5. **UML** [16]: a structural metric learning-based algorithm that aims to learn a less biased distance metric for classification tasks. The initial tuning proposal for this method was using a set of weakly-labeled data retrieved from querying class labels to search engine. However, here we tuned the hyperparameters using the same strategy as others (10FCV) for a fair comparison.
6. **LRE-SVM** [38]: a non-linear exemplar-SVMs model with a nuclear norm regularization to impose a low-rank *likelihood matrix*. Four hyper-parameters ($\lambda_1, \lambda_2, C_1, C_2$) were tuned using 10FCV.

The last three are the state-of-the-art domain generalization algorithms for object recognition.

Table 4. The cross-recognition accuracy % on the VLCS dataset.

Source	Target	L-SVM	1HNN	DAE+1HNN	Undo-Bias	UML	LRE-SVM	MTAE+1HNN	D-MTAE+1HNN
L,C,S	V	58.86	59.10	62.00	54.29	56.26	60.58	61.09	63.90
V,C,S	L	52.49	58.20	59.23	58.09	58.50	59.74	59.24	60.13
V,L,S	C	77.67	86.67	90.24	87.50	91.13	88.11	90.71	89.05
V,L,C	S	49.09	57.86	57.45	54.21	58.49	54.88	60.20	61.33
Avg.		59.93	65.46	67.23	63.52	65.85	65.83	67.81	68.60

Table 5. The cross-recognition accuracy % on the Office+Caltech dataset.

Source	Target	L-SVM	1HNN	DAE+1HNN	Undo-Bias	UML	LRE-SVM	MTAE+1HNN	D-MTAE+1HNN
A,C	D,W	82.08	83.41	82.05	80.49	82.29	84.59	84.23	85.35
D,W	A,C	76.12	76.49	79.04	69.98	79.54	81.17	79.30	80.52
C,D,W	A	90.61	92.13	92.02	90.98	91.02	91.87	92.20	93.13
A,W,D	C	84.51	85.89	85.17	85.95	84.59	86.38	85.98	86.15
Avg.		83.33	84.48	84.70	81.85	84.36	86.00	85.43	86.29

We report the performance in terms of the classification accuracy (%) following Xu *et al.* [38]. For all algorithms that are optimized stochastically, we ran independent training processes using the best performing hyper-parameters in 10 times and reported the average accuracies. Similar to the previous experiment, we do not report the standard deviations due to their small values (± 0.2).

Results on the VLCS dataset. We first conducted the standard training-test evaluation using L-SVM, *i.e.*, learning the model on a training set from one domain and testing it on a test set from another domain, to check the groundtruth performance and also to identify the existence of the dataset bias. The performance is summarized in Table 3. We can see that the bias indeed exists in every domain despite the use of DeCAF₆, the sixth layer features of the state-of-the-art deep convolutional neural network. The performance gap between the best cross-domain performance and the groundtruth is large, with $\geq 14\%$ difference.

We then evaluated the domain generalization performance of each algorithm. We conducted *leave-one-domain-out* evaluation, which induces four cross-domain cases. Table 4 summarizes the algorithms’ accuracy. In general, the dataset bias can be reduced by all algorithms after learning from multiple source domains (compare, *e.g.*, the minimum accuracy over L,C,S \rightarrow V tasks in Table 4 with the maximum cross-recognition accuracy over the VOC2007’s column in Table 3). Caltech-101, which is object-centric, appears to be the easiest dataset to recognize, consistent with an investigation in [35]: scene-centric datasets tend to generalize well over object-centric datasets. Surprisingly, the performance of 1HNN has already achieved competitive accuracy compared to more complicated state-of-the-art algorithms, Undo-Bias, UML, and LRE-SVM. Furthermore, D-MTAE outperforms other algorithms on three out of four cross-domain cases and on average, while MTAE has the second best performance on average.

Results on the Office+Caltech dataset. We report the experiment results on the Office+Caltech dataset. Table 5

summarizes the recognition accuracies of each algorithm over four cross-domain cases. D-MTAE+1HNN has the best performance on two out of four cross-domain cases and ranks second for the remaining cases. On average, D-MTAE+1HNN has better performance than the prior state-of-the-art on this dataset, LRE-SVM [38].

5. Conclusions

We have proposed a new approach to multi-task feature learning that reduces *dataset bias* in object recognition. The main idea is to extract features shared across domains via a training protocol that, given an image from one domain, learns to reconstruct analogs of that image for all domains. The strategy yields two variants: the Multi-task Autoencoder (MTAE) and the Denoising MTAE which incorporates a denoising criterion. A comprehensive suite of cross-domain object recognition evaluations shows that the algorithms successfully learn domain-invariant features, yielding state-of-the-art performance when predicting the labels of objects from unseen target domains.

Our results suggest several directions for further study. Firstly, it is worth investigating whether stacking MTAEs improves performance. Secondly, more effective procedures for handling unbalanced samples are required, since these occur frequently in practice. Finally, a natural application of MTAEs is to streaming data such as *video*, where the appearance of objects transforms in real-time.

The problem of dataset bias remains far from solved: the best model on the VLCS dataset achieved accuracies less than 70% on average. A partial explanation for the poor performance compared to supervised learning is insufficient training data: the class-overlap across datasets is quite small (only 5 classes are shared across VLCS). Further progress in domain generalization requires larger datasets.

6. Acknowledgments

The authors would like to thank Chen Fang for sharing the *Unbiased Metric Learning* code and useful discussions.

References

- [1] A. Argyriou, T. Evgeniou, and M. Pontil. Convex Multi-Task Feature Learning. *Machine Learning*, 73(3):243–272, 2008. **2**
- [2] M. Baktashmotlagh, M. T. Harandi, B. C. Lovell, and M. Salzmann. Domain Adaptation on the Statistical Manifold. In *CVPR*, pages 2481–2488, 2014. **2**
- [3] J. Baxter. A Model of Inductive Bias Learning. *Journal of Artificial Intelligence Research*, 12:149–198, 2000. **2**
- [4] Y. Bengio, A. C. Courville, and P. Vincent. Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013. **2**
- [5] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy Layer-Wise Training of Deep Networks. In *NIPS*, pages 153–160, 2007. **1, 2, 5**
- [6] G. Blanchard, G. Lee, and C. Scott. Generalizing from Several Related Classification Tasks to a New Unlabeled Sample. In *NIPS*, volume 1, pages 2178–2186, 2011. **1, 2**
- [7] H. Boullard and Y. Kamp. Auto-Association by Multilayer Perceptrons and Singular Value Decomposition. *Biological Cybernetics*, 59:291–294, 1988. **2**
- [8] R. Caruana. Multitask Learning. *Machine Learning*, 28:41–75, 1997. **1, 2, 3**
- [9] M. Chen, Z. Xu, K. Weinberger, and F. Sha. Marginalized Denoising Autoencoders for Domain Adaptation. In *ICML*, pages 767–774, 2012. **2**
- [10] M. J. Choi, J. J. Lim, A. Torralba, and A. S. Willsky. Exploiting hierarchical context on a large database of object categories. In *CVPR*, pages 129–136, 2010. **4**
- [11] D. Ciresan, U. Meier, and J. Schmidhuber. Multi-column deep neural network for image classification. In *CVPR*, pages 3642–3649, 2012. **1**
- [12] K. Crammer and Y. Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *JMLR*, 2:265–292, 2001. **5**
- [13] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. In *ICML*, pages 647–655, 2014. **2, 7**
- [14] M. Everingham, L. Van-Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results, 2007. **1, 4**
- [15] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A Library for Large Linear Classification. *JMLR*, 9:1871–1874, 2008. **5**
- [16] C. Fang, Y. Xu, and D. N. Rockmore. Unbiased Metric Learning: On the Utilization of Multiple Datasets and Web Images for Softening Bias. In *ICCV*, pages 1657–1664, 2013. **2, 7**
- [17] M. Ghifary, W. B. Kleijn, and M. Zhang. Deep hybrid networks with good out-of-sample object recognition. In *ICASSP*, pages 5437–5441, 2014. **2**
- [18] B. Gong, K. Grauman, and F. Sha. Reshaping Visual Datasets for Domain Adaptation. In *NIPS*, pages 1286–1294, 2013. **1**
- [19] B. Gong, Y. Shi, F. Sha, and K. Grauman. Geodesic Flow Kernel for Unsupervised Domain Adaptation. In *CVPR*, pages 2066–2073, 2012. **7**
- [20] G. Griffin, A. Holub, and P. Perona. Caltech-256 object category dataset. Technical report, California Inst. of Tech., 2007. **4**
- [21] A. Khosla, T. Zhou, T. Malisiewicz, A. Efros, and A. Torralba. Undoing the Damage of Dataset Bias. In *ECCV*, volume I, pages 158–171, 2012. **2, 7**
- [22] A. Krizhevsky. Learning Multiple Layers of Features from Tiny Images. Master’s thesis, Department of Computer Science, University of Toronto, Apr. 2009. **1**
- [23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*, volume 25, pages 1106–1114, 2012. **1, 6**
- [24] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, volume 86, pages 2278–2324, 1998. **4**
- [25] B. Leibe and B. Schiele. Analyzing appearance and contour based methods for object categorization. In *CVPR*, pages 409–415, 2003. **4**
- [26] K. Muandet, D. Balduzzi, and B. Schölkopf. Domain Generalization via Invariant Feature Representation. In *ICML*, pages 10–18, 2013. **1, 2, 5**
- [27] S. Rifai, P. Vincent, X. Muller, X. Glorot, and Y. Bengio. Contractive Auto-Encoders : Explicit Invariance During Feature Extraction. In *ICML*, number 1, pages 833–840, 2011. **2, 5, 6**
- [28] D. Rumelhart, G. Hinton, and R. Williams. *Parallel Distributed Processing. I: Foundations*. MIT Press, 1986. **1**
- [29] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986. **2**
- [30] B. C. Russell, A. Torralba, K. P. Murphy, and W. T. Freeman. LabelMe: A database and web-based tool for image annotation. In *IJCV*, volume 77, pages 157–173, 2008. **4**
- [31] K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting Visual Category Models to New Domains. In *ECCV*, pages 213–226, 2010. **1, 2, 4, 7**
- [32] I. Sutskever, O. Vinyals, and Q. Le. Sequence to Sequence Learning with Neural Networks. In *NIPS*, 2014. **1**
- [33] Y. Tang and C. Eliasmith. Deep networks for robust visual recognition. In *ICML*, pages 1055–1062, 2010. **2**
- [34] S. Thrun. Is learning the n-th thing any easier than learning the first? In *NIPS*, pages 640–646, 1996. **1, 2**
- [35] A. Torralba and A. Efros. Unbiased look at dataset bias. In *CVPR*, pages 1521–1528, 2011. **1, 8**
- [36] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol. Extracting and Composing Robust Features with Denoising Autoencoders. In *ICML*, 2008. **1**
- [37] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol. Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion. *Journal of Machine Learning Research*, 11:3371–3408, 2010. **2, 3, 5, 6**
- [38] Z. Xu, W. Li, L. Niu, and D. Xu. Exploiting Low-Rank Structure from Latent Domains for Domain Generalization. In *ECCV*, pages 628–643, 2014. **2, 7, 8**