

An Empirical Study of Domain Adaptation: Are We Really Learning Transferable Representations?

by

Nicholas J. Josselyn

A Thesis

Submitted to the Faculty

of the

WORCESTER POLYTECHNIC INSTITUTE

In partial fulfillment of the requirements for the

Degree of Master of Science

in

Data Science

by

October 2022

APPROVED:

Professor Elke A. Rundensteiner, Thesis Advisor

Professor Ziming Zhang, Thesis Co-Advisor

Professor Jacob Whitehill, Thesis Reader

Abstract

Deep learning often relies on the availability of a large amount of high-quality labeled data, which can be very limited in novel domains. To address such data scarcity, domain adaptation is one promising approach that allows for deep networks to leverage large amounts of available data from a source domain to enhance the model's efficacy on the target domain of interest. However, while there is a plethora of alternate models for domain adaptation proposed over many years in the literature, there is a dearth of studies that objectively compare the relative effectiveness of these models in a rigorous, empirical study. To fill this gap, we provide a thorough, unbiased, empirical study of five state-of-the-art (SOTA) deep domain adaptation models proposed over the past 6 years whose codes are publicly available. Models are evaluated on the complex and diverse domain adaptation tasks featured in the DomainNet benchmark dataset as well as the popular Office-31 dataset. Our results suggest that (1) all 5 models perform similarly, on average, and do not even significantly beat the oldest model, and (2) counter to their intended purpose, the transfer loss functions in the literature do not contribute significantly to learning transferable representations. Our observations suggest that domain adaptation research needs to more thoroughly compare newly proposed models against existing works, along with assessing their loss functions' utility thoroughly. **We will release our code and data splits for reproducibility of results by the community.**

Acknowledgements

This research is supported by the US ARMY, ACC-APG-RTP, Cooperative Agreement W911NF-19-2-0112, US. CDCC Army Research Lab and NSF grants CCF-2006738 and NRT-HDR-2021871. Thank you to my advisors, Professor Elke Rundensteiner and Professor Ziming Zhang and to my reader Professor Jacob Whitehill. I would also like to thank Biao Yin and the entire DAISY Lab for their support and help throughout the process of conducting this research. Furthermore, thank you to the Academic & Research Computing group at WPI, particularly James Kingsley. Finally, thank you to my friends and family for their continued support.

Contents

1	Introduction	1
1.1	Background	1
1.2	Applications	2
1.3	Our Approach and Contributions	3
2	Related Work	5
2.1	Domain Adaptation Models	5
2.2	Survey Papers	6
2.3	Experimental Analysis	6
3	Evaluated Models	7
3.1	DANN	8
3.2	JAN	8
3.3	CDAN	9
3.4	AFN	9
3.5	MCC	10
4	Dataset and Experimental Protocol	11
4.1	Datasets and Preparation	11
4.1.1	DomainNet	11
4.1.2	Office-31	13

4.2	Experimental Protocols	14
4.2.1	Baseline Transfer Tasks	14
4.2.2	Model Comparison	14
4.2.3	Hyperparameter Tuning	15
4.2.4	Transfer Loss Analysis	16
5	Experimental Results	18
5.1	Baseline Transfer Tasks	18
5.2	Model Comparison Analysis	19
5.2.1	DomainNet	19
5.2.2	Office-31	22
5.3	Transfer Loss Analysis	25
5.3.1	Loss Curve Analysis	25
5.3.2	Feature space alignment	27
6	Discussion and Conclusion	30
7	Future Work	32
8	Supplementary Material	34
8.1	Office-31 Dataset	34
8.2	DomainNet Dataset	35
8.2.1	DomainNet Examples	35
8.2.2	Dataset Preparation	35
8.3	Experiment Protocol	36
8.3.1	Step-by-Step Procedure	36
8.3.2	Optimal Hyperparameters	37
8.4	Supplemental Results	37

8.4.1	Model Comparisons	39
8.4.2	Transfer Loss Analysis	39
8.4.3	Feature space alignment: T-SNE	40

List of Figures

4.1	Images from the DomainNet dataset. Airplane class for all 6 domains shown.	11
4.2	Images from the Office-31 dataset. Backpack class for all 3 domains shown.	13
5.1	<i>Classification</i> loss curves for varied loss trade-off values (y-axis log-scale). Clipart to painting task for full dataset.	25
5.2	<i>Transfer</i> loss curves for varied loss trade-off values (y-axis log-scale). Clipart to painting task for full dataset.	26
5.3	Target domain test accuracies for varied loss trade-off values. Clipart to painting task for full dataset.	27
5.4	T-SNE plots for the airplane class. Top: DANN, Bottom: MCC. Red is the clipart source domain, blue is the target domain. a-e are 5 of 30 adaptation tasks. (C=clipart, I=infograph, P=painting, Q=quickdraw, R=real, S=sketch)	28
8.1	Images from the Office-31 dataset. Each row is a domain, 3 sample classes (backpack, scissors, trash can) are shown.	34
8.2	Images from the DomainNet dataset. Each row is a domain, 3 sample classes (airplane, dragon, toothpaste) are shown.	35

8.3	T-SNE plots for the airplane class only. Red represents the clipart source domain, blue represents the target domain. a-e represent 5 of the 30 adaptation tasks where clipart is always the source domain. (C=clipart, I=infograph, P=painting, Q=quickdraw, R=real, S=sketch) Each row is a model, top to bottom: DANN, JAN, CDAN, AFN, MCC.	46
8.4	Classification loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to infograph adaptation task for full dataset.	46
8.5	Transfer loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to infograph adaptation task for full dataset.	47
8.6	Classification loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to quickdraw adaptation task for full dataset.	47
8.7	Transfer loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to quickdraw adaptation task for full dataset.	47
8.8	Classification loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to real adaptation task for full dataset.	47
8.9	Transfer loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to real adaptation task for full dataset.	47
8.10	Classification loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to sketch adaptation task for full dataset.	48
8.11	Transfer loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to sketch adaptation task for full dataset.	48
8.12	Target domain test accuracies for varied loss trade-off hyperparameters. Clipart to infograph adaptation task for full dataset.	48
8.13	Target domain test accuracies for varied loss trade-off hyperparameters. Clipart to quickdraw adaptation task for full dataset.	48

8.14 Target domain test accuracies for varied loss trade-off hyperparameters.	
Clipart to real adaptation task for full dataset.	48
8.15 Target domain test accuracies for varied loss trade-off hyperparameters.	
Clipart to sketch adaptation task for full dataset.	49

List of Tables

3.1	5 UDA models studied in this work with publication years.	7
4.1	Amount of <i>training</i> data for each of the 6 domains in DomainNet for each of the 3 data subsets used.	12
5.1	Baseline transfer task test accuracies on the full DomainNet dataset. (Columns = source, rows = target)	19
5.2	Baseline transfer task test accuracies on the Office-31 dataset. (Columns = source, rows = target)	19
5.3	Mean target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the <i>full</i> dataset. Average test accuracies over all 30 tasks reported in bottom right. (Columns = source domain, rows = target domain)	20
5.4	Mean target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the <i>subset-50</i> dataset. Average test accuracies over all 30 tasks reported in bottom right. (Columns = source domain, rows = target domain)	21

5.5	Mean target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the <i>subset-20</i> dataset. Average test accuracies over all 30 tasks reported in bottom right. (Columns = source domain, rows = target domain)	22
5.6	Mean target domain test accuracy over 5-cross-val folds for all models for all 6 adaptation tasks on the Office-31 dataset. Average test accuracies over all 6 tasks reported in far right. (Columns = source domain, rows = target domain)	23
5.7	Standard deviation of validation accuracies achieved for each model and each adaptation task for a variety of learning rate and weight decay hyperparameter combinations tried (A=Amazon, D=DSLR, W=Webcam; A2D is "Amazon source adapted to DSLR target")	24
8.1	Optimal hyperparamters for each model for the DomainNet dataset. . . .	38
8.2	Optimal hyperparamters for each model for the Office-31 dataset. . . .	38
8.3	Mean and standard deviation target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the full dataset. (Columns = source domain, rows = target domain)	41
8.4	Mean and standard deviation target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the subset-50 dataset. (Columns = source domain, rows = target domain)	42
8.5	Mean and standard deviation target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the subset-20 dataset. (Columns = source domain, rows = target domain)	43
8.6	Mean and standard deviation target domain test accuracy over 5-cross-val folds for all models for all 6 adaptation tasks with tuned hyperparameters on the <i>Office-31</i> dataset. (Columns = source domain, rows = target domain)	44

8.7 Mean and standard deviation target domain test accuracy over 5-cross-val
folds for all models for all 6 adaptation tasks with tuned hyperparameters
from the *DomainNet* dataset. (Columns = source domain, rows = target
domain) 45

Chapter 1

Introduction

1.1 Background

A primary influence on the effectiveness of deep learning is the availability of data, specifically high quality labeled data. It could be relatively simple for many industries or organizations to generate a lot of unorganized or unlabeled data, but providing labels for large quantities of data requires costly, both in time and money, human intervention. Worse yet, such human-generated labels are subject to human bias and errors. For each new problem, one would expect that a new model needs to be trained with new data for each task due to dataset bias between the distribution of data used for training *vs.* testing. Over the past few years, transfer learning approaches have thus been developed to exploit large labeled datasets (such as ImageNet) to learn high-level features of images and thereafter fine-tune models for new problems by additional training on new targeted task-specific datasets. However, it is often difficult to have enough data to properly tune a model for specific tasks [5, 27].

Domain adaptation (DA), a subfield in transfer learning, primarily assumes that there are two or more similar but distinct sets of data (domains) called *source* domain and *target* domain. Generally, the task for each of these datasets is the same (e.g. classification of

planes) but there exists a *domain shift* between the domains (e.g. photos of planes v.s. drawings of planes). A domain shift can exist for several reasons: the data is captured via different modalities (infrared vs visual spectrum images), changes in pose, changes in object background, variations in color, dimensionality differences, or any combination, and more. The goal of DA is two-fold, namely, to learn a model that can both perform a task on the source data well and simultaneously learn a transferable representation of the data such that the source and target domains are indistinguishable. This then would allow for good performance of the task on the target domain as well.

Within the field of DA, we distinguish between different approaches depending on the availability of data: supervised, semi-supervised, and unsupervised. In this work, we focus on unsupervised DA (UDA), arguably the most popular in literature. In UDA, during training, both labeled source examples and unlabeled target examples are utilized. We focus on five core models spanning the 6 recent years in the literature. These models all have two core loss components, namely, a classification loss and a transfer loss. The latter is introduced to align the source and target domains.

1.2 Applications

Domain adaptation has potential utility in many application areas, with two important ones being self-driving cars [24, 26, 36] and medical imaging [9, 22, 32], as described below.

For self-driving cars, unforeseen environmental factors, different road layouts, and robust pedestrian identification are some of the many hurdles to overcome. Training data for self-driving cars may not be able to capture every potential environmental scenario, roads from one country may differ from another, and pedestrians being in unforeseen, dangerous scenarios in the real-world. All these are examples of *domain shifts* between training and testing and ultimately deployment. Thus, domain adaptation may be a potential solution.

For medical imaging, consensus for diagnoses across patients' anatomy and medical devices (scanners, MRI, CT) is critical. Depending on the cohort of training data, identifying anatomical structures is complicated due to the variety of anatomy present in each dataset. Further, depending on a hospital's scanning parameters and device choices, model predictions are subject to variability across devices. Finally, if a model is trained on data from one modality (MRI), but it is desired to then perform a diagnosis on another (CT), it is difficult to achieve reliable results. These challenges could potentially be addressed with domain adaptation.

1.3 Our Approach and Contributions

We note that a plethora of UDA models have been proposed in the literature, each making claims of better performance than the previous and proposing alternate variations of transfer loss functions to align domains. While categorizations of these DA models have been provided in survey papers [2, 9, 18, 25, 31], there is a scarcity of unbiased, comprehensive analyses of DA models needed to gain an objective understanding of their relative effectiveness. Given the importance of this task for the community at large, both future model developers and practitioners, we provide an in-depth evaluation in this work. In particular, we conduct a comprehensive experimental study of 5 SOTA UDA models.

Our empirical evaluation study rests upon several important pillars. *First*, we conduct an unbiased, thorough experimental study of popular UDA models on varying amounts of training data (from big to small, more limited data sets) and for a diverse set of domain types and tasks. *Second*, we utilize open-source, benchmark datasets to assure open access to our models and experimental data. *Third*, we provide an analysis of the utility of the proposed transfer loss functions in the DA literature, assessing their relative contribution (or lack thereof) to learning an effective classifier. Our results are two-fold. (1) *All 5 models, regardless of when they were released, perform similarly, on average.* Interestingly, we find

that none of the newer models significantly beats even the oldest model. (2) *Surprisingly, the proposed transfer loss functions across these models do not significantly contribute to learning transferable representations.* This result is in contradiction to their proposed functionality. Given our findings, we recommend to the DA research community that a more careful and fair comparison of newly proposed models against the existing literature be conducted, and in particular, proof of domain alignment is shown. Further, we suggest to practitioners looking to adopt DA models that well-established older models are worth to consider and try out on their applications, as they may be more robust to different data scenarios.

Chapter 2

Related Work

2.1 Domain Adaptation Models

There exists a large body of literature on domain adaptation models – all proposing new ways in which we can tackle the problem of domain shift via learning similar representations between domains with new loss functions. These works reflect several areas of domain adaptation, namely, single-source closed-set DA [6, 7, 13–16, 19, 28–30, 33, 35, 38], multi-source DA [12, 20, 34, 39], multi-target DA [21], open-set DA [17], partial DA [1, 37], and other subset fields of domain adaptation. In our work, we focus on one of the foundational settings of domain adaptation: *unsupervised, single-source, single-target, closed-set domain adaptation with two loss components*. When a new model is proposed in the literature and compared against baselines, they are subject to unintentional bias. In our work, we have no preference to any one model and thus can be fully objective in our analysis.

2.2 Survey Papers

This work is similar to survey-style work [2, 9, 18, 25, 31] as it discusses existing models. However, in these works, they only catalogue a large history of domain adaptation and transfer learning models. Results reported in these works are generally taken from existing manuscripts, are not thorough model investigations under varied data scenarios, nor are comparative in nature. In this work, we take an in-depth, empirical approach and provide a thorough evaluation of a set of SOTA domain adaptation models on a diverse set of adaptation tasks. We further deeply investigate the effectiveness of transfer loss functions proposed in the literature.

2.3 Experimental Analysis

Other than survey work, there is one notable work that compares UDA models [3]. However, a large portion of this work is spent on the design of a new model and a new dataset, with evaluation on their new dataset. Further, they do not explicitly compare UDA models in the literature. Instead, they restrict their comparison to shallow and deep DA models under different weight sharing strategies.

Chapter 3

Evaluated Models

Model/Paper	Publication year
Domain-Adversarial Training of Neural Networks (DANN) [6, 7]	2016
Deep Transfer Learning with Joint Adaptation Networks (JAN) [16]	2017
Conditional Adversarial Domain Adaptation (CDAN) [14]	2017
Larger Norm More Transferable: An Adaptive Feature Norm Approach for UDA (AFN) [35]	2019
Minimum Class Confusion for Versatile Domain Adaptation (MCC) [11]	2020

Table 3.1: 5 UDA models studied in this work with publication years.

In this section, we introduce the 5 UDA models investigated in this study, as seen in Table 3.1, and point out their respective uniqueness as well as commonalities between them. The 5 models, introduced over the past 6 years, were selected based on the following criteria: (1) highly cited, recent or established models in literature, (2) models that have been compared against in a large majority of the domain adaptation literature, and (3) models that match our foundational scenario of being single-source, single-target, closed-set DA with two primary loss components. To implement these models, we use a public DA library [10]. We utilize this framework and 5 of its models and update codes according to the datasets we use and metrics we record. Code for the models we have updated will be provided along with access to data for reproducibility ¹

¹<https://github.com/njosselyn13/Empirical-Study-Domain-Adaptation>

3.1 DANN

Domain-Adversarial Training of Neural Networks [7] is a highly cited UDA model introduced in 2016 that is commonly used as a representative baseline. DANN introduces a gradient reversal layer to promote alignment of the source and target domain representations with respect to a domain discriminator. DANN consists of a label predictor that predicts the class label of the source data during training and subsequently target data during testing. Further, DANN incorporates a domain discriminator that discriminates between source and target data during training. The training procedure aims to minimize the loss of the label predictor and maximize the domain confusion loss of the discriminator via an adversarial approach with the gradient reversal layer. With this, there are two main loss components: a classification loss for the label predictor and a transfer loss for the domain discriminator.

3.2 JAN

Deep Transfer Learning with Joint Adaptation Networks [16], a UDA model introduced in 2017, aims to reduce the shifts in joint distributions across domains. Typically, the Maximum Mean Discrepancy [8] is used to measure the discrepancy in marginal distributions between domains. For this, JAN formulates a function to measure the discrepancy between the joint distributions between domains across multiple CNN layers. In measuring the discrepancy of the joint distributions, the authors argue that they can overcome the residual joint distribution shifts not addressed in other domain adaptation models. In this model, a classification loss for assessing the labeled source data remains. The proposed joint distribution alignment corresponds to their metric for transfer loss.

3.3 CDAN

Conditional Adversarial Domain Adaptation [14], a UDA model introduced in 2017, was inspired by the advances in conditional generative adversarial networks (CGANs). CGANs make use of discriminative features between real and fake data and incorporate them in a conditional manner into the generator and discriminator networks. In CDAN, a similar approach is taken by conditioning the domain discriminator with the cross-covariance of domain-specific feature representations and classifier predictions. Additionally, the discriminator is conditioned based on the uncertainty of the classifier predictions – thus allowing the discriminator to prioritize easy-to-transfer examples. In this model, there is still a classification loss for assessing the labeled source data. Also, the proposed transfer loss is fairly similar to DANN and its discriminator, simply now with the inclusion of conditioning.

3.4 AFN

Larger Norm More Transferable: An Adaptive Feature Norm Approach for Unsupervised Domain Adaptation [35], a UDA model introduced in 2019, proposes a novel method for aligning domains based on statistical criterion different from the literature. That is, they suggest that a better way to align domains is via finding a shared, average feature norm (length of the feature vector) between the two domains. They note that the target domain feature norms are typically much smaller than the source feature norms. They conjecture that this may complicate the adaptation. By adapting the feature norms of both domains to a large range of scalars, they expect they can achieve better adaptation. There is again a classification loss, while transfer loss is the feature norm loss.

3.5 MCC

Minimum Class Confusion for Versatile Domain Adaptation [11], a UDA model introduced in 2020, does not aim to explicitly align two domain feature spaces, but instead it aims to reduce class confusion in the label space. MCC is versatile because its approach is suitable for a variety of DA settings such as: closed-set, partial-set, multi-source, and multi-target DA. In our work focused on comparing single-source/single-target solutions, we leverage MCC just for this particular single-source and single-target setup. MCC claims to outperform prior models including the models we study here, namely, DANN, JAN, CDAN, and AFN. In MCC, there are two main loss components, classification loss for the source data and transfer loss designed to minimize the number of misclassifications in the target domain.

Chapter 4

Dataset and Experimental Protocol

In this section, we provide details on the experimental methodology to compare the 5 UDA models, including datasets, data preparation, training setup, and more.

4.1 Datasets and Preparation

In this work, we focus on two benchmark domain adaptation image datasets: the popular Office-31 dataset [23], and the largest, most diverse and complex dataset, DomainNet [20].

4.1.1 DomainNet

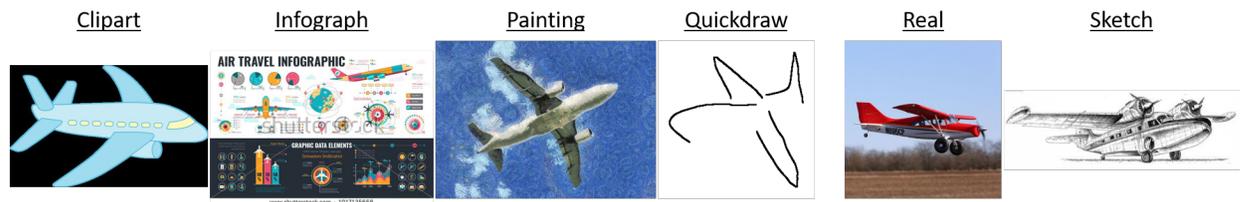


Figure 4.1: Images from the DomainNet dataset. Airplane class for all 6 domains shown.

Published in 2019, DomainNet is the largest domain adaptation benchmark dataset containing 586,575 images across 6 distinct and uniquely challenging domains with 345

	Clipart	Infograph	Painting	Quickdraw	Real	Sketch	Total
Subset-20	5,516	5,953	8,286	19,320	19,455	7,851	66,381
Subset-50	13,485	14,483	20,270	48,300	48,432	19,354	164,324
Full	26,820	28,818	40,358	96,600	96,725	38,570	327,891

Table 4.1: Amount of *training* data for each of the 6 domains in DomainNet for each of the 3 data subsets used.

classes in each domain. The 6 domains are: clipart, infograph, painting, quickdraw, real, and sketch images. This allows for 30 adaptation tasks to be evaluated (single-source, single-target DA setup). In each domain, diverse classes such as airplane, toothpaste, dragon, rabbit, ear, etc. exist. An example of images for each domain are shown in Figure 4.1. More examples can be found in the supplementary material.

In our study, we distinguish between 3 training scenarios based on the amount of training data provided, namely, large, medium, and small datasets. We investigate varying training data scenarios in order to assess the robustness of each model to data availability. Using the original split for DomainNet defined in [20] and the VisDA-2019 competition (with small modifications due to some classes not being assigned images in the original release), 30% of the data is partitioned into a held-out test set. Details on rectifying these small inconsistencies can be found in the supplementary material. The remaining 70% of data is designated for training (and validation). We break this remaining 70% of data into 3 distinct datasets: the full amount of data (full), 50% of the data per class per domain (subset-50), and 20% of the data per class per domain (subset-20). For this data sub-setting, a hard minimum of 6 images per class was enforced to allow for cross-validation to be performed. Then, for each of the 3 subsets of training data, 5 stratified cross-validated folds of training and validation data are generated with 20% being held for validation and 80% for training. In Table 4.1 we provide statistics on how many images are available for training for each of the 6 domains over all 345 classes. All data subsets will be released to

assure reproducibility.

4.1.2 Office-31



Figure 4.2: Images from the Office-31 dataset. Backpack class for all 3 domains shown.

The Office-31 dataset consists of 3 domains of office-space image data taken from Amazon, a DSLR camera, and a webcam camera, allowing for 6 adaptation tasks (single-source, single-target DA setup). There are a total of 4,110 images across all 3 domains. In each domain there are 31 classes, including, back packs, scissors, trash cans, etc. Examples can be seen in Figure 4.2. Additional images can be seen in the supplementary material. In the Amazon domain, there are 2,817 images. The Amazon domain provides images with clean backgrounds and a uniform scale. In the DSLR domain, there is a total of 498 images. The DSLR domain provides images with low noise and high resolution. The webcam domain, with a total of 795 images, presents images with low resolution, significant noise, and color and white balance artifacts. The Office-31 dataset is considered a small-scale domain adaptation dataset with small domain shifts, particularly between webcam and DSLR domains.

To assure a comprehensive analysis of model performance, we assess all 5 models on the full Office-31 dataset, which differs widely in tasks and data types from DomainNet. Unlike DomainNet, we do not subset the Office-31 dataset to smaller training scenarios as it is already a smaller-scale dataset.

We split the Office-31 dataset into training, validation, and testing sets. We split

15% of the data to a held-out test set with the remaining 85% of data being split into 5 cross-validated folds; 70% held for training and 30% held for validation.

4.2 Experimental Protocols

In this section, we first outline our methodology for establishing a lower bound for domain adaptation performance, i.e., the performance of models without domain adaptation on a set of transfer tasks. Then, using the 3 subsets of DomainNet data and the Office-31 dataset, we describe our experimental methodology for comparing and tuning the 5 UDA models and for the assessment of the impact transfer loss has on learning.

4.2.1 Baseline Transfer Tasks

As a first step, we conduct baseline experiments to confirm that indeed a classifier trained on one domain of data with no domain adaptation, when applied to a similar but different domain of data, will perform poorer than when using domain adaptation models.

For this, we first train 6 classifiers on the full DomainNet dataset using classical single-domain models. That is, no domain adaptation is used. Similarly, we train 3 classifiers on the Office-31 dataset using classical single-domain models. For each of these classifiers, we tune the learning rate and weight decay and select the model with highest validation accuracy. We sample learning rate and weight decay values of 0.1, 0.01, and 0.001. Once an optimal model for each classifier is chosen, it is then applied to the 5 (for DomainNet) or 2 (for Office-31) other target domain *test* data sets.

4.2.2 Model Comparison

For each of the 5 UDA models, 5-fold cross-validation experiments are conducted on the Office-31 dataset and each of the 3 subsets of the DomainNet dataset. Experiments are conducted for all 30 domain adaptation tasks in the DomainNet dataset and all 6 adaptation

tasks in the Office-31 dataset. When a domain (e.g. clipart) is selected as the source domain for a model and dataset pair (e.g. JAN and full DomainNet dataset), the labels are made available to the model at train time. When a domain (e.g. infograph) is selected as the target domain for a model and dataset pair, the labels are not available to the model during training. The target domain data for the corresponding validation and test sets are then used to be tuned and evaluated on, respectively.

For each model and dataset experiment pair, and each adaptation task, test set accuracies are reported over the 5 cross-validation folds. Additionally, we compute for each model and dataset pair an average test accuracy over all adaptation tasks. An additional step-by-step procedure is included in the supplementary material. Code and data for reproducibility will be made available.

4.2.3 Hyperparameter Tuning

Extensive hyperparameter tuning was conducted to give each model a fair chance at maximizing their performance on the DomainNet or Office-31 datasets. All experiments run for 30 epochs. The test accuracy where the max validation accuracy is observed is reported, thus implementing an early stopping approach. We experimented with longer training times (100 epochs) with a subset of 5 adaptation tasks from DomainNet (clipart as source). We did not observe significant improvement in accuracy to warrant running longer than 30 epochs. Additionally, we tuned each model’s learning rate, weight decay, and the loss trade-off hyperparameters. The loss trade-off weighs how much to focus on classification loss vs each model’s respective transfer loss. This trade-off is seen in Equation 4.1.

$$loss = cls_loss + transfer_loss * trade_off \quad (4.1)$$

For each model, we tune hyperparameters on the full DomainNet dataset, for a subset of 5 adaptation tasks (clipart as source). For the Office-31 dataset, we perform tuning

for each model over all 6 adaptation tasks. Learning rate and weight decay combinations explored for each model were: 0.1, 0.01, 0.001, 0.0001, and original code defaults. Loss trade-off values explored for each model were: 0.25, 0.5, 0.75, 2, 5, 10, 20, 50, and original code defaults. Larger values mean giving more relative weight to the transfer loss. All final hyperparameters are specified in the supplementary material.

Finally, for the most recent MCC model, after much experimentation, we found it to be one of the poorer performing models on DomainNet. Thus, we reached out to the MCC authors and took their advice to tune the temperature hyperparameter with suggested values of 0.25, 0.5, 1, 2, and 3.

4.2.4 Transfer Loss Analysis

To assess the utility of proposed transfer loss functions in DA models and their contribution to learning an effective target domain classifier, we examine training loss curves for each model for each loss trade-off hyperparameter we worked with. We observe how the classification and transfer losses evolve over the 30 training epochs independently. We then conclude whether any amount of weighting for the transfer loss (via the loss trade-off hyperparameter) leads to significant learning during training with respect to the transfer loss. We show, for each model, the lack of learning occurring with respect to the transfer loss and discuss potential reasons for this observation.

For these experiments, we use optimal hyperparameters determined previously on the DomainNet dataset. We only vary the loss trade-off value. We observe loss curves for the full DomainNet dataset and the first cross-validation fold.

Next, we use the trained models with optimal hyperparameters, from the first cross-validation fold, to generate T-SNE plots. These plots show the learned distributions of each model-adaptation task pair for one class (airplane). We extract the features before the final classification layer for either the source or target data and reduce them to 2 dimensions via

T-SNE. This allows us to visualize how much the representations between two domains have been transformed to be similar.

Chapter 5

Experimental Results

In this section, we first establish a lower bound baseline (no domain adaptation) for the DomainNet and Office-31 datasets’ transfer tasks. Then, we discuss results for our two proposed contributions in Chapter 1: (1) model comparisons and robustness, and (2) the utility of transfer loss functions proposed in the DA literature.

5.1 Baseline Transfer Tasks

Results using the full DomainNet dataset with no domain adaptation are presented in Table 5.1. On average, over all 30 adaptation tasks, this baseline approach achieves only 23.4% test accuracy. We see later in Table 5.3 (test results also using the full DomainNet dataset) that this performance is lower than that for all 5 domain adaptation models. In fact, this performance is lower than all 5 models when only 50% of the training data is used. Furthermore, in Table 5.2 we show baseline results on Office-31 and achieve an average performance over all 6 tasks of 72.8% test accuracy. As we later see in Table 5.6, this is less than all 5 domain adaptation models as well. *This confirms what has been observed in the literature: domain adaptation outperforms the simple application of trained models to related but different domains.*

	clp	inf	pnt	qdr	rel	skt
clp	-	29.2	29.7	8.1	44.4	46.3
inf	15.9	-	15.7	0.6	19.0	13.4
pnt	35.3	27.4	-	1.0	45.1	32.4
qdr	9.1	2.6	2.9	-	4.1	11.1
rel	49.6	43.9	44.4	2.4	-	45.3
skt	37.0	22.8	27.2	5.9	32.0	-

Table 5.1: Baseline transfer task test accuracies on the full DomainNet dataset. (Columns = source, rows = target)

	Amazon	DSLR	Webcam
Amazon	-	55.3	53.7
DSLR	72.0	-	96.0
Webcam	65.0	95.0	-

Table 5.2: Baseline transfer task test accuracies on the Office-31 dataset. (Columns = source, rows = target)

5.2 Model Comparison Analysis

We present results obtained on the DomainNet dataset for all 5 UDA models, followed by results on Office-31.

5.2.1 DomainNet

Addressing the first objective of model robustness in our work, we first present 5-fold cross-validation results on the test dataset (for each target domain) for all 5 UDA models for a diverse set of 30 adaptation tasks from the DomainNet dataset in Tables 5.3, 5.4, and 5.5. Each column designates a source domain and each row a target domain. Table 5.3 reports results on the full dataset, Table 5.4 on the subset-50 dataset, and Table 5.5 on the subset-20 dataset. Additionally, each table presents the average accuracy over all adaptation tasks for each model as is typically done in the DA literature to provide a summary of model performance overall. We include additional tables in the supplementary

DANN [7]	clp	inf	pnt	qdr	rel	skt	JAN [16]	clp	inf	pnt	qdr	rel	skt
clp	-	27.0	37.3	18.7	47.6	50.8	clp	-	28.1	38.5	15.3	46.6	48.5
inf	19.1	-	18.6	3.9	22.7	19.3	inf	17.2	-	17.3	2.4	20.6	16.5
pnt	34.3	25.9	-	7.9	47.7	41.4	pnt	33.9	26.6	-	6.7	46.6	39.2
qdr	12.7	5.5	5.8	-	6.5	12.4	qdr	11.4	4.7	4.8	-	4.6	10.1
rel	50.4	36.9	51.8	12.7	-	50.8	rel	49.5	39.4	51.5	11.7	-	47.0
skt	40.7	22.6	33.5	13.0	37.4	-	skt	40.2	24.0	35.7	10.9	36.5	-
CDAN [14]	clp	inf	pnt	qdr	rel	skt	AFN [35]	clp	inf	pnt	qdr	rel	skt
clp	-	26.6	37.2	19.9	49.3	51.0	clp	-	29.1	38.9	17.3	43.9	49.6
inf	18.7	-	18.4	4.0	22.7	18.8	inf	15.8	-	16.5	2.8	17.6	15.6
pnt	35.0	25.2	-	7.1	49.0	42.0	pnt	36.0	32.5	-	6.6	46.6	42.6
qdr	10.1	4.7	4.7	-	4.5	9.3	qdr	11.3	3.1	4.7	-	4.8	11.1
rel	52.6	37.0	51.6	14.4	-	51.1	rel	51.9	45.7	53.9	11.8	-	51.9
skt	40.9	22.0	34.8	13.8	38.7	-	skt	41.2	25.0	36.4	9.3	34.4	-
MCC [11]	clp	inf	pnt	qdr	rel	skt	Averages:						
clp	-	31.1	43.0	11.6	51.2	53.6	DANN (2016)	27.2					
inf	13.5	-	14.8	1.1	17.7	13.8	CDAN (2017)	27.2					
pnt	33.4	28.7	-	2.6	46.1	39.8	AFN (2019)	26.9					
qdr	12.9	2.7	3.7	-	4.0	11.7	JAN (2017)	26.2					
rel	49.0	43.7	52.5	7.0	-	49.1	MCC (2020)	26.0					
skt	39.3	22.7	36.1	7.9	35.1	-							

Table 5.3: Mean target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the *full* dataset. Average test accuracies over all 30 tasks reported in bottom right. (Columns = source domain, rows = target domain)

material that reports standard deviations for each adaptation task for each model as well.

In Table 5.3, results for all models on the *full* dataset are shown. We observe that for 9/30 adaptation tasks, DANN performs best, for 9/30 tasks AFN performs best, for 8/30 tasks CDAN performs best, for 5/30 tasks MCC performs best, and for 0 tasks JAN performs best. Note that for the real-to-infograph task, DANN and CDAN tie. Over all 30 adaptation tasks, for all models, the standard deviations over the 5 cross-validation folds ranges from 0.1 to 1.1. With respect to averages over all 30 adaptation tasks, DANN and CDAN beat the other 3 models with 27.2% test accuracy, with AFN as a close second with 26.9%. We note that *the oldest method (DANN) is performing best for a diverse set of challenging adaptation tasks; with the most recently published method (MCC) performing the poorest*. This is of particular interest as it raises concerns over the current trend in proposing new domain adaptation methods, and the rigor with which new models should be tested against baselines to assure their robustness in performance in practice.

DANN [7]	clp	inf	pnt	qdr	rel	skt	JAN [16]	clp	inf	pnt	qdr	rel	skt
clp	-	22.9	34.1	17.9	45.3	45.9	clp	-	24.0	35.4	14.8	44.6	44.4
inf	17.1	-	17.0	3.9	21.6	17.5	inf	15.4	-	16.0	2.6	19.5	14.7
pnt	32.0	23.0	-	7.5	46.2	38.7	pnt	31.5	23.3	-	6.9	45.5	36.2
qdr	11.8	5.2	5.3	-	6.6	11.2	qdr	10.5	4.3	4.2	-	4.6	9.5
rel	46.8	33.6	49.7	12.8	-	48.1	rel	46.6	35.2	50.1	12.1	-	44.6
skt	37.3	19.8	30.9	12.5	35.4	-	skt	36.5	20.0	33.1	11.2	34.6	-
CDAN [14]	clp	inf	pnt	qdr	rel	skt	AFN [35]	clp	inf	pnt	qdr	rel	skt
clp	-	22.4	34.7	19.1	47.1	46.2	clp	-	27.0	36.5	17.3	43.2	47.2
inf	16.9	-	17.1	4.0	21.8	17.1	inf	14.8	-	15.8	2.8	17.2	15.0
pnt	32.2	21.7	-	6.9	47.8	39.0	pnt	34.2	29.6	-	6.7	45.9	41.0
qdr	8.4	4.3	3.9	-	3.6	8.5	qdr	10.5	3.2	4.2	-	4.5	10.6
rel	49.1	33.9	50.3	13.8	-	48.5	rel	49.9	42.5	52.8	12.2	-	50.3
skt	37.1	18.7	31.8	13.4	36.4	-	skt	38.2	22.0	34.3	9.2	33.7	-
MCC [11]	clp	inf	pnt	qdr	rel	skt	Averages:						
clp	-	26.4	38.1	11.9	47.2	49.4			AFN (2019)	25.7			
inf	11.7	-	13.2	1.2	16.8	12.4			DANN (2016)	25.2			
pnt	29.6	24.5	-	2.9	43.3	37.0			CDAN (2017)	25.2			
qdr	12.0	2.7	4.0	-	4.7	11.2			JAN (2017)	24.4			
rel	45.1	38.6	49.0	6.6	-	46.1			MCC (2020)	23.8			
skt	34.9	19.0	32.5	7.4	33.2	-							

Table 5.4: Mean target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the *subset-50* dataset. Average test accuracies over all 30 tasks reported in bottom right. (Columns = source domain, rows = target domain)

In Table 5.4, results for all models on the *subset-50* dataset are shown. We observe that for 7/30 adaptation tasks DANN performs best, for 11/30 tasks AFN performs best, for 8/30 tasks CDAN performs best, for 5/30 tasks MCC performs best (tied with DANN for sketch to quickdraw task), and for 0 tasks JAN performs best. Over all 30 adaptation tasks, for all models, the standard deviations over the 5 cross-validation folds ranges from 0.0 to 0.9. With respect to averages over all 30 adaptation tasks, AFN just slightly wins with 25.7% test accuracy, but is closely followed by DANN and CDAN with 25.2%. Again we observe that the oldest of the methods (DANN) is performing near the top, and the newest of the methods (MCC) is performing at the bottom. We begin to observe a *trend of model robustness across training data availability*; DANN, CDAN, and AFN perform well with the full amount of training data and with 50% of the training data.

In Table 5.5, results for all models on the *subset-20* dataset are shown. We observe that for 10/30 adaptation tasks, DANN performs best, for 15/30 tasks, AFN performs

DANN [7]	clp	inf	pnt	qdr	rel	skt	JAN [16]	clp	inf	pnt	qdr	rel	skt
clp	-	17.5	29.0	16.4	41.0	39.0	clp	-	18.5	30.5	13.7	40.3	37.6
inf	14.1	-	14.9	3.8	19.9	14.8	inf	12.8	-	14.0	2.6	17.8	12.3
pnt	27.1	18.3	-	7.2	43.0	34.2	pnt	26.7	19.9	-	6.6	43.0	31.6
qdr	9.9	3.9	5.0	-	6.1	9.6	qdr	7.9	3.1	4.1	-	3.9	7.7
rel	42.2	28.3	46.3	12.4	-	43.5	rel	40.8	31.0	47.3	11.3	-	39.5
skt	30.8	15.0	27.1	11.2	32.7	-	skt	29.7	15.4	28.4	9.9	30.8	-
CDAN [14]	clp	inf	pnt	qdr	rel	skt	AFN [35]	clp	inf	pnt	qdr	rel	skt
clp	-	16.7	30.3	16.1	42.3	38.9	clp	-	21.4	32.7	17.0	40.7	41.2
inf	13.7	-	14.9	3.6	20.0	14.0	inf	13.2	-	14.7	2.8	16.7	13.0
pnt	25.3	18.9	-	6.2	44.6	34.1	pnt	30.3	24.8	-	6.2	44.7	36.8
qdr	4.8	2.2	2.5	-	3.3	5.4	qdr	8.9	2.4	3.3	-	3.8	9.5
rel	42.1	29.8	47.8	12.2	-	44.1	rel	45.3	35.9	50.7	11.8	-	45.9
skt	28.1	12.3	27.4	11.7	33.0	-	skt	32.5	17.5	30.4	8.7	31.8	-
MCC [11]	clp	inf	pnt	qdr	rel	skt	Averages:						
clp	-	19.0	30.6	11.6	40.6	40.3	AFN (2019)	23.2					
inf	8.7	-	10.7	1.2	15.0	9.6	DANN (2016)	22.1					
pnt	22.5	18.4	-	3.1	39.6	31.0	CDAN (2017)	21.6					
qdr	9.6	2.9	4.3	-	4.9	9.8	JAN (2017)	21.3					
rel	37.8	30.6	44.3	7.2	-	39.6	MCC (2020)	19.8					
skt	26.4	13.5	26.1	7.9	27.9	-							

Table 5.5: Mean target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the *subset-20* dataset. Average test accuracies over all 30 tasks reported in bottom right. (Columns = source domain, rows = target domain)

best, for 1 task, MCC performs best, and JAN and CDAN never perform best. For the painting-to-infograph task, DANN and CDAN tie. Over all 30 adaptation tasks, for all models, the standard deviation over the 5 cross-validation folds ranges from 0.1 to 0.9. With respect to averages over all 30 adaptation tasks, AFN performs best with 23.2% test accuracy, with DANN coming in second again with 22.1%. Again we see MCC perform towards the bottom of the model list. We further observe the *same trend that AFN and the oldest model DANN perform best for a wide variety of adaptation tasks and for all 3 sizes of training data.*

5.2.2 Office-31

To further evaluate model robustness, we present results on the Office-31 dataset. We show 5-fold cross-validation results on the test set (for each target domain) for all 5 models for all 6 adaptation tasks in Table 5.6. Each column denotes a source domain, and each row

DANN [7]	Amazon	DSLRL	Webcam	JAN [16]	Amazon	DSLRL	Webcam	Averages:	
Amazon	-	69.8	70.0	Amazon	-	70.1	68.6	MCC (2020)	89.5
DSLRL	91.7	-	99.7	DSLRL	92.5	-	100.0	AFN (2019)	86.8
Webcam	89.2	97.5	-	Webcam	90.0	97.5	-	JAN (2017)	86.5
CDAN [14]	Amazon	DSLRL	Webcam	AFN [35]	Amazon	DSLRL	Webcam	DANN (2016)	86.3
Amazon	-	69.9	65.9	Amazon	-	70.3	70.2	CDAN (2017)	85.8
DSLRL	91.7	-	100.0	DSLRL	93.1	-	100.0		
Webcam	89.7	97.3	-	Webcam	89.2	98.2	-		
MCC [11]	Amazon	DSLRL	Webcam						
Amazon	-	74.7	77.2						
DSLRL	94.7	-	100.0						
Webcam	92.2	98.3	-						

Table 5.6: Mean target domain test accuracy over 5-cross-val folds for all models for all 6 adaptation tasks on the Office-31 dataset. Average test accuracies over all 6 tasks reported in far right. (Columns = source domain, rows = target domain)

a target domain. We present the average performance over all 6 tasks for each model to provide a summary of the model performance. We again include additional tables in the supplementary material that reports standard deviations for each adaptation task for each model as well.

In Table 5.6 we observe that for all 6 adaptation tasks, the MCC model performs best and has the highest overall average performance of 89.5%. For all 5 models and for all 6 adaptation tasks, the standard deviations over the 5 cross-validation folds range from 0.0 to 2.9. We observe that with Office-31, MCC is now the top performing model and DANN is toward the bottom. *However, all 5 models still perform close on average, with the AFN model still a top performing model.*

This relative change in model performance order points at the importance of needing to tune the MCC model significantly. As another set of experiments, we also run all 5 models for all 6 adaptation tasks from the Office-31 dataset without any tuning of hyperparameters. Instead, we simply take the best hyperparameters from the DomainNet tuning and apply them when using the Office-31 dataset. When doing this, we observe that the DANN, JAN, CDAN, and AFN models still obtain fairly comparable average test performances of: 86.9%, 86.5%, 88.1%, and 86.4%, respectively. However, the MCC model only achieved

	A2D	A2W	D2A	D2W	W2A	W2D
DANN	20.43	22.97	18.59	22.12	18.24	3.00
JAN	19.61	18.29	16.80	6.66	16.20	3.21
CDAN	22.09	22.32	22.59	24.79	19.12	25.29
AFN	35.05	34.15	25.81	14.79	22.01	7.50
MCC	27.64	27.36	26.44	28.92	26.35	28.18

Table 5.7: Standard deviation of validation accuracies achieved for each model and each adaptation task for a variety of learning rate and weight decay hyperparameter combinations tried (A=Amazon, D=DSLR, W=Webcam; A2D is "Amazon source adapted to DSLR target")

an average test accuracy of 24.1%; drastically less than the tuned average test accuracy. A full set of these results on Office-31 using DomainNet hyperparameters can be found in the supplementary material.

Additionally, over all learning rates and weight decay combinations explored for each model and all 6 adaptation tasks, we calculate standard deviations of the validation accuracies achieved to get a sense of model sensitivity with respect to hyperparameters; standard deviations are found in Table 5.7. We observe that for 4 out of 6 of the tasks, the MCC model has the highest variability in performance when varying learning rates and weight decays. Note that these standard deviations are fairly high in all cases because there are a few hyperparameter combinations that fail drastically for all models such as learning rate and weight decay of 0.1. *These two pieces of evidence lead us to conclude that the MCC model is fairly unstable across datasets and requires careful, time-consuming hyperparameter tuning for each new dataset* in order to assure that it outperforms other models in the case of Office-31. However, in the case of DomainNet as previously noted, even with careful hyperparameter tuning, MCC performs the poorest in our study.

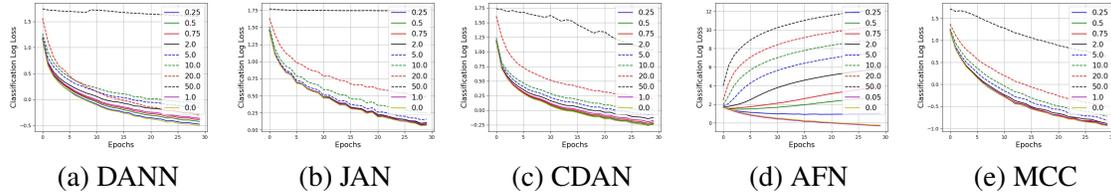


Figure 5.1: *Classification* loss curves for varied loss trade-off values (y-axis log-scale). Clipart to painting task for full dataset.

5.3 Transfer Loss Analysis

Next, we present results on classification loss, transfer loss, and classification accuracy curves for all 5 models. Then, we visualize learned feature spaces for a set of adaptation tasks.

5.3.1 Loss Curve Analysis

Addressing the second objective of transfer loss utility in our work, we present classification and transfer loss curves in Figures 5.1 and 5.2, respectively, for a variety of loss trade-off hyperparameter values as mentioned in Section 4.2. Loss curves are shown for just the clipart-to-painting adaptation task for the full dataset. Four additional adaptation task loss curves are shown in the supplementary material (all with clipart as the source domain) and all show similar trends. The default trade-off value for all models except AFN is 1.0, AFN had a default value of 0.05. All default values are plotted in magenta in Figures 5.1 and 5.2. Loss values are plotted on a log-scale. In Figure 5.2, for the AFN model, the y-axis is all 0.693 as changes in the transfer loss vary only to the seventh decimal place over the 30 epochs; meaning, very little learning occurring. In Figure 5.3, we also present the target domain (painting) test accuracy plots for each model for each loss trade-off value.

In Figure 5.1, across all 5 models, we generally see the classification loss steadily decreasing and flattening out during training. The exception is with trade-off value of 50 (heavily weighting transfer loss) for JAN and CDAN where the classification loss increases

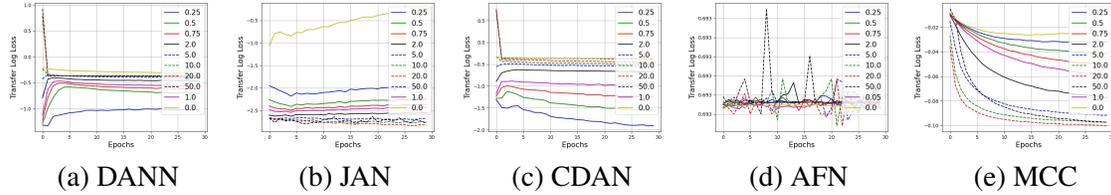


Figure 5.2: *Transfer* loss curves for varied loss trade-off values (y-axis log-scale). Clipart to painting task for full dataset.

or stays flat. For AFN, for the majority of trade-off values, the classification loss increases except for the default trade-off value and small trade-off values. Even for larger trade-off values that weigh transfer loss more of 10.0 and 20.0 for DANN, JAN, CDAN, and MCC and 50.0 for MCC, we still observe a decreasing classification loss, indicating learning with respect to the classification loss.

An interesting observation we make is with respect to the transfer loss curves for each model depicted in Figure 5.2. *For JAN, CDAN, AFN, and DANN we observe little to no learning occurring with respect to the transfer losses proposed for each model.* Each model claims that their transfer loss function would be aligning domains. We should observe this as an improvement in learning via the transfer loss curve. However, even though we vary the weighting of the two losses during training, we consistently observe that transfer loss changes are minimal and sometimes even experience increases, even when heavily weighting it with values up to 50. Given this observation, for a variety of loss trade-off values and models, *the utility of DA transfer loss functions may not be achieving the goal they are designed for: a full alignment of domains.* Given that the inclusion of transfer loss functions in domain adaptation models leads to improvement over traditional transfer learning, *we speculate that these transfer loss functions may be simply acting as a regularization technique to generalize better to the target domain data.*

The only model whose transfer loss curves appear to be decreasing to some degree is the MCC model. However, the actual decrease over 30 epochs is approximately 0.08, a small fraction of the overall loss decrease that occurs during training. Although, even

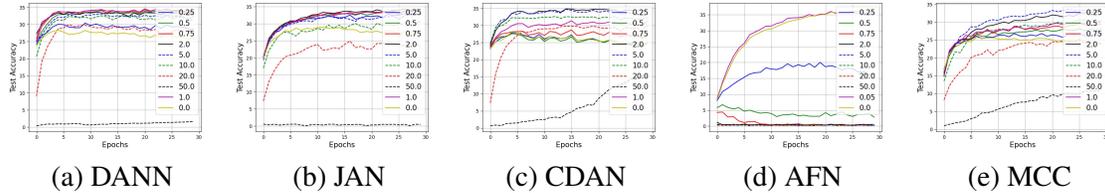


Figure 5.3: Target domain test accuracies for varied loss trade-off values. Clipart to painting task for full dataset.

with the transfer loss decreasing only slightly for MCC, we have consistently seen MCC perform the poorest.

In Figures 5.1 and 5.2 we demonstrate that classification loss is the dominating factor in learning, while transfer loss is not contributing to learning during training. *This raises concern over the utility of the proposed transfer loss functions in that they are not helping much to learn a transferable representation between source and target domains.*

Further, when we compare the loss curves to what we observe with respect to the target domain test accuracy curves in Figure 5.3, we see that for a majority of trade-off values, the test accuracies continue to increase even when the corresponding transfer loss curves remain unchanged or even start to increase. This further supports the observation that the classification loss, regardless of trade-off weight, is the dominating factor in learning. To summarize, given that the inclusion of a transfer loss has been shown to still improve target domain classification over traditional transfer learning, this leads us to suggest that these transfer loss functions may be acting simply as regularization techniques.

5.3.2 Feature space alignment

We present T-SNE plots showing feature distributions for 5 adaptation tasks for 2 models in Figure 5.4. We select the DANN model, a consistently top performing model, and the MCC model, the newest albeit poorest performing model to show in Figure 5.4. We show the 5 adaptation tasks with clipart used as the source domain and all other domains used as target. Additional plots for other models can be found in the supplementary material. For

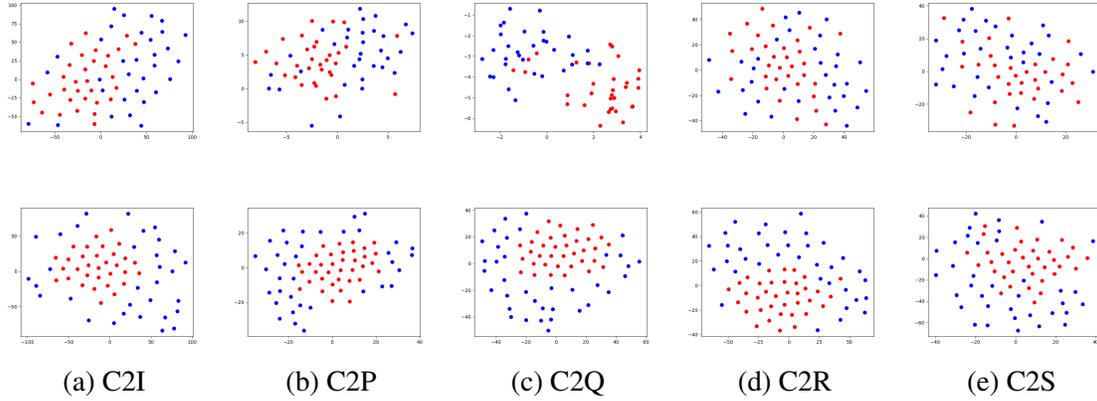


Figure 5.4: T-SNE plots for the airplane class. Top: DANN, Bottom: MCC. Red is the clipart source domain, blue is the target domain. a-e are 5 of 30 adaptation tasks. (C=clipart, I=infograph, P=painting, Q=quickdraw, R=real, S=sketch)

each adaptation task, we plot the learned representations of data for each domain, with red indicating source domain, and blue the target domain, for a single class (airplane).

When transferring knowledge between domains, we anticipate overlapping distributions of data representations for the same class. With Figure 5.4, we can check on the effectiveness of each transfer loss function for a variety of transfer tasks by observing if the two domains have overlapping distributions. *For a majority of the tasks we see a visual separation of the source (red) and target (blue) domain representations.* The test accuracies for these 5 tasks, left to right, for DANN were: 19.1%, 34.3%, 12.7%, 50.4%, and 40.7%. And for MCC they were: 13.5%, 33.4%, 12.9%, 49.0%, and 39.3%. We see in particular that for the clipart-to-quickdraw (C2Q) transfer task, a fairly noticeable separation between domains remains. Correspondingly, we see fairly low accuracies of $\sim 12\%$. For the clipart-to-infograph (C2I) task, DANN has a higher performance of 19.1% compared to MCC’s 13.5%. Correspondingly in Figure 5.4 there is a slightly larger overlap of distributions for DANN than observed for MCC for the C2I task. *With this separation of domains remaining even after applying domain adaptation, this is further evidence that supports that the transfer loss functions are not necessarily achieving their goal of aligning domains.* However, since accuracies are boosted with domain adaptation, transfer

loss functions are indeed having some impact, possibly as regularization.

Chapter 6

Discussion and Conclusion

In this work, we provide a detailed, unbiased empirical evaluation of 5 state-of-the-art deep unsupervised DA models to assess model robustness across different benchmark domain adaptation tasks and datasets. This work provides a valuable objective empirical analysis for model choice for a variety of data scenarios. Surprisingly, we observe that AFN and DANN, the oldest model, are consistently top performers for all 3 training data availability scenarios on the DomainNet benchmark. Interestingly, one of the oldest UDA methods, DANN, outperforms newer models; most notably the 2020 MCC model, with the latter performing at or near the bottom for a majority of adaptation tasks. This observation suggests that future DA research thoroughly evaluate baseline models under a rich variety of conditions when comparing them to newly proposed models.

Furthermore, we show that for Office-31, while MCC becomes the top-performing model, all 5 models on average remain close in performance to one another. We observe, however, that the MCC model is more sensitive to hyperparameter choice across datasets, requires careful tuning for new datasets, and is not guaranteed to be the best model even when tuned (i.e. on DomainNet). This suggests that the newest model is not the most robust choice for practitioners and more careful analysis needs to be done when proposing new models.

Another important take-away of our study is that the dominating factor to learning is the classification loss, even when weighting the two losses (classification and transfer loss) with a wide range of weight values. The transfer loss curves indicate that little to no learning occurs during training. We thus speculate that their effectiveness in target domain classification is due more to them acting as a regularizer instead of truly aligning domains. Future work is thus needed to confirm the usefulness of the proposed transfer loss functions.

Further, as illustrated by visual inspection of T-SNE plots, the learned representations for a (source, target) pair of domains are not mapped to a significantly overlapping feature space - as would have been expected from DA models.

Chapter 7

Future Work

As a potential future direction, we would like to explore comparing domain adaptation models with a two-stage training approach. We would take the same proposed domain adaptation models in literature, but first train them to solely be good classifiers, freeze that model, and then subsequently attempt to learn to align the two domains with the respective proposed transfer loss functions in each method. We would want to observe whether the joint training, which involves complex hyperparameter tuning, may cause a disadvantage to model performance.

Additionally, another avenue to pursue in investigating the effectiveness of domain adaptation models is the data imbalance present both between classes of one domain and between the domains. Overfitting to one domain due to an excess of data in one class or domain is a possible hurdle in adapting across domains. We would like to look into possible data imbalance techniques available in literature and apply them to the domain adaptation setup.

Finally, as a result of this study, we not only provide a thorough empirical investigation of domain adaptation models, but we also gain insight into what models may work well for a variety of training scenarios and data types. We observed that the AFN and DANN models are robust models and maintain high performance even under more limited data

scenarios. This observation gives us a starting point for more applied work to be done on real-world datasets with limited data and label availability. One specific application task we would like to explore is in conjunction with the Army Research Lab (ARL) and our mission to automatically assess the level of corrosion present for various anti-corrosive materials that are subjected to two environments (two domains) of testing: indoor and outdoor testing.

Chapter 8

Supplementary Material

8.1 Office-31 Dataset

We provide some additional images in 8.1 from the Office-31 dataset [23] to show a few more of the 31 classes present in the 3 domains: Amazon, DSLR, and Webcam.



Figure 8.1: Images from the Office-31 dataset. Each row is a domain, 3 sample classes (backpack, scissors, trash can) are shown.

8.2 DomainNet Dataset

8.2.1 DomainNet Examples

We provide some additional images in Figure 8.2 from the DomainNet dataset [20] to show a few more of the 345 classes present in the 6 domains: clipart, infograph, painting, quickdraw, real, and sketch.

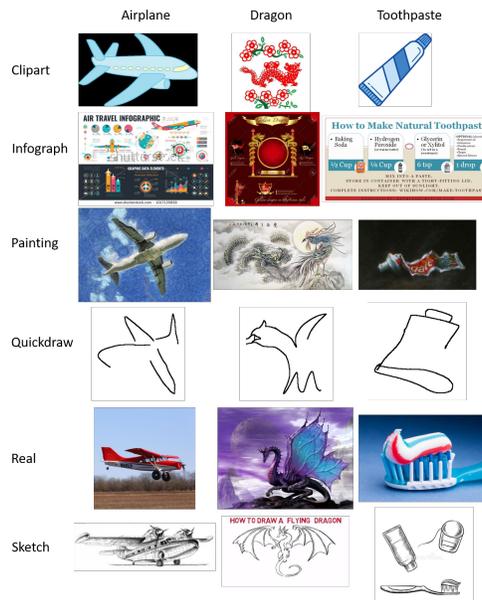


Figure 8.2: Images from the DomainNet dataset. Each row is a domain, 3 sample classes (airplane, dragon, toothpaste) are shown.

8.2.2 Dataset Preparation

There were a few inconsistencies identified with the original split of the DomainNet dataset as originally defined in [20] and the VisDA-2019 competition that had to be rectified for the purpose of the experiments to be conducted. We had to insure that there existed images in each class, something we found to be a problem with the original split released, and had to insure that we had enough images of each class present for each of the data subsets and

train/validation/test sets we were creating for our experimental design.

The inconsistencies specifically occurred in the painting domain. In the painting domain *test* set, no images were assigned to the class "syringe" when images were available in the dataset to be assigned there. Therefore, we assigned the 4 available images to the painting domain syringe class in the test set. In the painting domain *training* set, no images were assigned to the class "t-shirt" when images were available in the dataset to be assigned there. Therefore, we assigned the 10 available images not already assigned to the test set for "t-shirt" class here.

Additionally, in order to run 5-fold cross validation, a 6 image per class minimum was enforced for all 3 subsets of the DomainNet dataset such that when creating the cross-validated folds enough images would be available for splitting. This required adding images previously not used in the defined 70-30 split (yet available in the dataset) to classes that did not have 6 images assigned per class. We added to the original 70% of training data in the painting domain 4 images to the "stereo" class, 3 images to the "underwear" class, 4 images to the "waterslide" class, 10 images to the "toe" class, and an additional 6 images were moved from the test set to the training set for the "syringe" class.

8.3 Experiment Protocol

We provide an additional step-by-step procedure for how experiments were conducted to fairly compare models and then present the optimal hyperparameters found for each model on the DomainNet and Office-31 datasets.

8.3.1 Step-by-Step Procedure

1. Domain adaptation model is chosen (e.x. JAN)
2. Data subset is chosen (full, subset-50, subset-20, or Office-31)

3. Source domain is chosen (clipart, infograph, painting, quickdraw, real, sketch, amazon, dslr, or webcam)
4. Target domain is chosen, different than the source domain (clipart, infograph, painting, quickdraw, real, sketch, amazon, dslr, or webcam)
5. 5-fold cross validation is run
6. Choose new target domain
7. Repeat steps 5 and 6 until all domains are used as target
8. Choose new source domain
9. Repeat steps 4-8 until all domains are used as source
10. Choose new data subset
11. Repeat steps 3-10 until all data subsets chosen
12. Choose new domain adaptation model
13. Repeat steps 2-12 until all domain adaptation models chosen

8.3.2 Optimal Hyperparameters

In Tables 8.1, 8.2 we show the optimal values for the tuned hyperparameters on the DomainNet and Office-31 datasets, respectively. We report: learning rate, weight decay, loss trade-off, and temperature with a few other hyperparameter values. Pre-trained ResNet-50 models are pre-trained on ImageNet [4]. All other default hyperparameter values will be in the released code.

8.4 Supplemental Results

In this section we first show additional tables for all 5 models, for all 3 data subsets, for all 30 adaptation tasks on the DomainNet dataset but with standard deviations over the 5 cross-validation folds for each adaptation task. Then we show additional tables for the Office-31 dataset, for all 6 adaptation tasks but with standard deviations over 5 cross-validation

	DANN	JAN	CDAN	AFN	MCC
Pre-trained (backbone)	ResNet-50	ResNet-50	ResNet-50	ResNet-50	ResNet-50
Batch Size	32	32	32	32	36
Momentum	0.9	0.9	0.9	0.9	0.9
Learning Rate	0.01	0.01	0.01	0.001	0.1
Weight Decay	0.001	0.001	0.003	0.001	0.0001
Loss Trade-off	1.0	1.0	2.0	0.05	5.0
Temperature	-	-	-	-	2.0

Table 8.1: Optimal hyperparamters for each model for the DomainNet dataset.

	DANN	JAN	CDAN	AFN	MCC
Pre-trained (backbone)	ResNet-50	ResNet-50	ResNet-50	ResNet-50	ResNet-50
Batch Size	32	32	32	32	36
Momentum	0.9	0.9	0.9	0.9	0.9
Learning Rate	0.001	0.001	0.001	0.001	0.005
Weight Decay	0.01	0.01	0.001	0.01	0.001
Loss Trade-off	0.5	0.75	1.0	0.05	0.5
Temperature	-	-	-	-	2.0

Table 8.2: Optimal hyperparamters for each model for the Office-31 dataset.

folds. We then include a set of tables with standard deviations over 5 cross-validation folds for the Office-31 dataset but without the tuned hyperparameters on the Office-31 dataset, instead showing the results on Office-31 when using the tuned hyperparameters found on DomainNet; this table is shown to emphasize the sensitivity of the MCC model to hyperparameter tuning. We then show additional classification and transfer loss curves for 4 more adaptation tasks and their corresponding test accuracy curves. Finally, we show additional T-SNE plots for all 5 models for 5 adaptation tasks (clipart as source).

8.4.1 Model Comparisons

In Tables 8.3, 8.4, and 8.5 we show 5-fold cross-validated test set accuracies (means and standard deviations) for each of the 30 adaptation tasks on DomainNet for all 5 models for the full, subset-50, and subset-20 datasets, respectively.

In Table 8.6 we show 5-fold cross-validated test set accuracies (means and standard deviations) for each of the 6 adaptation tasks on Office-31 for all 5 models with the tuned hyperparameters on the *Office-31* dataset.

In Table 8.7 we show 5-fold cross-validated test set accuracies (means and standard deviations) for each of the 6 adaptation tasks on Office-31 for all 5 models with the tuned hyperparameters from the *DomainNet* dataset.

8.4.2 Transfer Loss Analysis

In this section we provide the classification and transfer loss curves for the remaining 4 adaptation tasks where clipart is the source domain (infograph, quickdraw, real, and sketch as target domains) as well as the test accuracy curves for each of the 4 adaptation tasks. These results are also only for the full dataset and cross-validation fold 1. Refer to Figures 8.4, 8.5, 8.6, 8.7, 8.8, 8.9, 8.10, and 8.11 for classification and transfer loss curves. Refer to Figures 8.12, 8.13, 8.14, and 8.15 for test accuracy curves.

8.4.3 Feature space alignment: T-SNE

Here we show in Figure 8.3 T-SNE plots for all 5 models for 5 adaptation tasks where clipart is always the source domain.

DANN	clp	inf	pnt	qdr	rel	skt
clp	-	27.0 ± 0.2	37.3 ± 0.2	18.7 ± 0.2	47.6 ± 0.4	50.8 ± 0.3
inf	19.1 ± 0.2	-	18.6 ± 0.2	3.9 ± 0.2	22.7 ± 0.1	19.3 ± 0.1
pnt	34.3 ± 0.2	25.9 ± 0.2	-	7.9 ± 0.3	47.7 ± 0.2	41.4 ± 0.2
qdr	12.7 ± 0.5	5.5 ± 0.3	5.8 ± 0.3	-	6.5 ± 0.8	12.4 ± 0.2
rel	50.4 ± 0.2	36.9 ± 0.4	51.8 ± 0.2	12.7 ± 0.3	-	50.8 ± 0.3
skt	40.7 ± 0.1	22.6 ± 0.4	33.5 ± 0.5	13.0 ± 0.3	37.4 ± 0.3	-
JAN	clp	inf	pnt	qdr	rel	skt
clp	-	28.1 ± 0.2	38.5 ± 0.5	15.3 ± 0.6	46.6 ± 0.3	48.5 ± 0.3
inf	17.2 ± 0.3	-	17.3 ± 0.2	2.4 ± 0.1	20.6 ± 0.1	16.5 ± 0.3
pnt	33.9 ± 0.3	26.6 ± 0.9	-	6.7 ± 0.3	46.6 ± 0.1	39.2 ± 0.5
qdr	11.4 ± 0.2	4.7 ± 0.2	4.8 ± 0.3	-	4.6 ± 0.2	10.1 ± 0.4
rel	49.5 ± 0.4	39.4 ± 0.4	51.5 ± 0.3	11.7 ± 0.7	-	47.0 ± 0.3
skt	40.2 ± 0.2	24.0 ± 0.1	35.7 ± 0.2	10.9 ± 0.3	36.5 ± 0.5	-
CDAN	clp	inf	pnt	qdr	rel	skt
clp	-	26.6 ± 0.1	37.2 ± 0.2	19.9 ± 0.6	49.3 ± 0.3	51.0 ± 0.1
inf	18.7 ± 0.2	-	18.4 ± 0.3	4.0 ± 0.1	22.7 ± 0.2	18.8 ± 0.1
pnt	35.0 ± 0.2	25.2 ± 0.3	-	7.1 ± 0.2	49.0 ± 0.2	42.0 ± 0.2
qdr	10.1 ± 0.3	4.7 ± 0.4	4.7 ± 0.2	-	4.5 ± 0.5	9.3 ± 0.3
rel	52.6 ± 0.2	37.0 ± 0.3	51.6 ± 0.2	14.4 ± 0.2	-	51.1 ± 0.2
skt	40.9 ± 0.2	22.0 ± 0.4	34.8 ± 0.4	13.8 ± 0.5	38.7 ± 0.4	-
AFN	clp	inf	pnt	qdr	rel	skt
clp	-	29.1 ± 0.5	38.9 ± 0.2	17.3 ± 0.6	43.9 ± 0.5	49.6 ± 0.4
inf	15.8 ± 0.2	-	16.5 ± 0.2	2.8 ± 0.1	17.6 ± 0.2	15.6 ± 0.2
pnt	36.0 ± 0.2	32.5 ± 0.5	-	6.6 ± 0.3	46.6 ± 0.2	42.6 ± 0.4
qdr	11.3 ± 0.2	3.1 ± 0.2	4.7 ± 0.1	-	4.8 ± 0.2	11.1 ± 0.2
rel	51.9 ± 0.3	45.7 ± 0.3	53.9 ± 0.4	11.8 ± 0.4	-	51.9 ± 0.4
skt	41.2 ± 0.4	25.0 ± 0.1	36.4 ± 0.1	9.3 ± 0.3	34.4 ± 0.3	-
MCC	clp	inf	pnt	qdr	rel	skt
clp	-	31.1 ± 0.4	43.0 ± 0.4	11.6 ± 0.6	51.2 ± 0.6	53.6 ± 0.4
inf	13.5 ± 0.4	-	14.8 ± 0.4	1.1 ± 0.1	17.7 ± 0.3	13.8 ± 0.3
pnt	33.4 ± 0.6	28.7 ± 0.1	-	2.6 ± 0.4	46.1 ± 0.2	39.8 ± 0.3
qdr	12.9 ± 0.8	2.7 ± 0.4	3.7 ± 1.0	-	4.0 ± 0.7	11.7 ± 0.5
rel	49.0 ± 0.4	43.7 ± 0.6	52.5 ± 0.3	7.0 ± 0.4	-	49.1 ± 0.5
skt	39.3 ± 0.1	22.7 ± 0.7	36.1 ± 0.5	7.9 ± 0.8	35.1 ± 0.3	-

Table 8.3: Mean and standard deviation target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the full dataset. (Columns = source domain, rows = target domain)

DANN	clp	inf	pnt	qdr	rel	skt
clp	-	22.9 ± 0.4	34.1 ± 0.2	17.9 ± 0.3	45.3 ± 0.5	45.9 ± 0.4
inf	17.1 ± 0.2	-	17.0 ± 0.2	3.9 ± 0.3	21.6 ± 0.1	17.5 ± 0.1
pnt	32.0 ± 0.3	23.0 ± 0.3	-	7.5 ± 0.2	46.2 ± 0.2	38.7 ± 0.1
qdr	11.8 ± 0.3	5.2 ± 0.5	5.3 ± 0.5	-	6.6 ± 0.2	11.2 ± 0.4
rel	46.8 ± 0.3	33.6 ± 0.5	49.7 ± 0.1	12.8 ± 0.2	-	48.1 ± 0.2
skt	37.3 ± 0.3	19.8 ± 0.4	30.9 ± 0.4	12.5 ± 0.5	35.4 ± 0.5	-
JAN	clp	inf	pnt	qdr	rel	skt
clp	-	24.0 ± 0.4	35.4 ± 0.1	14.8 ± 0.4	44.6 ± 0.3	44.4 ± 0.5
inf	15.4 ± 0.2	-	16.0 ± 0.3	2.6 ± 0.1	19.5 ± 0.3	14.7 ± 0.2
pnt	31.5 ± 0.4	23.3 ± 0.2	-	6.9 ± 0.3	45.5 ± 0.2	36.2 ± 0.4
qdr	10.5 ± 0.4	4.3 ± 0.2	4.2 ± 0.2	-	4.6 ± 0.3	9.5 ± 0.6
rel	46.6 ± 0.3	35.2 ± 0.4	50.1 ± 0.3	12.1 ± 0.7	-	44.6 ± 0.1
skt	36.5 ± 0.4	20.0 ± 0.4	33.1 ± 0.5	11.2 ± 0.3	34.6 ± 0.4	-
CDAN	clp	inf	pnt	qdr	rel	skt
clp	-	22.4 ± 0.2	34.7 ± 0.2	19.1 ± 0.5	47.1 ± 0.2	46.2 ± 0.3
inf	16.9 ± 0.2	-	17.1 ± 0.1	4.0 ± 0.2	21.8 ± 0.1	17.1 ± 0.3
pnt	32.2 ± 0.7	21.7 ± 0.3	-	6.9 ± 0.5	47.8 ± 0.2	39.0 ± 0.2
qdr	8.4 ± 0.5	4.3 ± 0.2	3.9 ± 0.3	-	3.6 ± 0.5	8.5 ± 0.7
rel	49.1 ± 0.3	33.9 ± 0.8	50.3 ± 0.2	13.8 ± 0.4	-	48.5 ± 0.2
skt	37.1 ± 0.2	18.7 ± 0.2	31.8 ± 0.5	13.4 ± 0.5	36.4 ± 0.2	-
AFN	clp	inf	pnt	qdr	rel	skt
clp	-	27.0 ± 0.2	36.5 ± 0.4	17.3 ± 0.5	43.2 ± 0.4	47.2 ± 0.1
inf	14.8 ± 0.1	-	15.8 ± 0.2	2.8 ± 0.1	17.2 ± 0.2	15.0 ± 0.2
pnt	34.2 ± 0.3	29.6 ± 0.2	-	6.7 ± 0.2	45.9 ± 0.1	41.0 ± 0.2
qdr	10.5 ± 0.3	3.2 ± 0.2	4.2 ± 0.2	-	4.5 ± 0.1	10.6 ± 0.1
rel	49.9 ± 0.3	42.5 ± 0.3	52.8 ± 0.2	12.2 ± 0.3	-	50.3 ± 0.1
skt	38.2 ± 0.4	22.0 ± 0.2	34.3 ± 0.3	9.2 ± 0.3	33.7 ± 0.2	-
MCC	clp	inf	pnt	qdr	rel	skt
clp	-	26.4 ± 0.4	38.1 ± 0.5	11.9 ± 0.4	47.2 ± 0.2	49.4 ± 0.3
inf	11.7 ± 0.2	-	13.2 ± 0.2	1.2 ± 0.1	16.8 ± 0.3	12.4 ± 0.3
pnt	29.6 ± 0.4	24.5 ± 0.7	-	2.9 ± 0.0	43.3 ± 0.2	37.0 ± 0.4
qdr	12.0 ± 0.5	2.7 ± 0.3	4.0 ± 0.4	-	4.7 ± 0.4	11.2 ± 0.5
rel	45.1 ± 0.6	38.6 ± 0.7	49.0 ± 0.3	6.6 ± 0.5	-	46.1 ± 0.5
skt	34.9 ± 0.5	19.0 ± 0.5	32.5 ± 0.3	7.4 ± 0.9	33.2 ± 0.4	-

Table 8.4: Mean and standard deviation target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the subset-50 dataset. (Columns = source domain, rows = target domain)

DANN	clp	inf	pnt	qdr	rel	skt
clp	-	17.5 ± 0.5	29.0 ± 0.4	16.4 ± 0.6	41.0 ± 0.3	39.0 ± 0.4
inf	14.1 ± 0.1	-	14.9 ± 0.2	3.8 ± 0.1	19.9 ± 0.1	14.8 ± 0.3
pnt	27.1 ± 0.3	18.3 ± 0.8	-	7.2 ± 0.2	43.0 ± 0.2	34.2 ± 0.3
qdr	9.9 ± 0.1	3.9 ± 0.2	5.0 ± 0.4	-	6.1 ± 0.6	9.6 ± 0.2
rel	42.2 ± 0.5	28.3 ± 0.3	46.3 ± 0.2	12.4 ± 0.4	-	43.5 ± 0.3
skt	30.8 ± 0.3	15.0 ± 0.3	27.1 ± 0.4	11.2 ± 0.2	32.7 ± 0.2	-
JAN	clp	inf	pnt	qdr	rel	skt
clp	-	18.5 ± 0.7	30.5 ± 0.5	13.7 ± 0.5	40.3 ± 0.4	37.6 ± 0.4
inf	12.8 ± 0.2	-	14.0 ± 0.2	2.6 ± 0.2	17.8 ± 0.1	12.3 ± 0.2
pnt	26.7 ± 0.3	19.9 ± 0.3	-	6.6 ± 0.2	43.0 ± 0.3	31.6 ± 0.3
qdr	7.9 ± 0.5	3.1 ± 0.1	4.1 ± 0.3	-	3.9 ± 0.4	7.7 ± 0.5
rel	40.8 ± 0.5	31.0 ± 0.3	47.3 ± 0.3	11.3 ± 0.6	-	39.5 ± 0.6
skt	29.7 ± 0.3	15.4 ± 0.5	28.4 ± 0.3	9.9 ± 0.4	30.8 ± 0.3	-
CDAN	clp	inf	pnt	qdr	rel	skt
clp	-	16.7 ± 0.2	30.3 ± 0.7	16.1 ± 0.7	42.3 ± 0.2	38.9 ± 0.3
inf	13.7 ± 0.1	-	14.9 ± 0.1	3.6 ± 0.1	20.0 ± 0.1	14.0 ± 0.1
pnt	25.3 ± 0.2	18.9 ± 0.6	-	6.2 ± 0.3	44.6 ± 0.1	34.1 ± 0.5
qdr	4.8 ± 0.6	2.2 ± 0.4	2.5 ± 0.3	-	3.3 ± 0.2	5.4 ± 0.5
rel	42.1 ± 0.6	29.8 ± 0.7	47.8 ± 0.3	12.2 ± 0.6	-	44.1 ± 0.4
skt	28.1 ± 0.5	12.3 ± 0.3	27.4 ± 0.4	11.7 ± 0.3	33.0 ± 0.3	-
AFN	clp	inf	pnt	qdr	rel	skt
clp	-	21.4 ± 0.4	32.7 ± 0.2	17.0 ± 0.6	40.7 ± 0.3	41.2 ± 0.2
inf	13.2 ± 0.3	-	14.7 ± 0.2	2.8 ± 0.1	16.7 ± 0.1	13.0 ± 0.2
pnt	30.3 ± 0.4	24.8 ± 0.1	-	6.2 ± 0.3	44.7 ± 0.2	36.8 ± 0.3
qdr	8.9 ± 0.3	2.4 ± 0.3	3.3 ± 0.2	-	3.8 ± 0.3	9.5 ± 0.1
rel	45.3 ± 0.4	35.9 ± 0.6	50.7 ± 0.3	11.8 ± 0.3	-	45.9 ± 0.3
skt	32.5 ± 0.2	17.5 ± 0.3	30.4 ± 0.2	8.7 ± 0.3	31.8 ± 0.0	-
MCC	clp	inf	pnt	qdr	rel	skt
clp	-	19.0 ± 0.5	30.6 ± 0.4	11.6 ± 0.3	40.6 ± 0.3	40.3 ± 0.6
inf	8.7 ± 0.3	-	10.7 ± 0.2	1.2 ± 0.1	15.0 ± 0.3	9.6 ± 0.2
pnt	22.5 ± 0.4	18.4 ± 0.5	-	3.1 ± 0.3	39.6 ± 0.3	31.0 ± 0.6
qdr	9.6 ± 0.2	2.9 ± 0.6	4.3 ± 0.3	-	4.9 ± 0.6	9.8 ± 0.6
rel	37.8 ± 0.3	30.6 ± 0.4	44.3 ± 0.4	7.2 ± 0.9	-	39.6 ± 0.3
skt	26.4 ± 0.7	13.5 ± 0.5	26.1 ± 0.3	7.9 ± 0.5	27.9 ± 0.5	-

Table 8.5: Mean and standard deviation target domain test accuracy over 5-cross-val folds for all models for all 30 adaptation tasks on the subset-20 dataset. (Columns = source domain, rows = target domain)

DANN	Amazon	DSLR	Webcam
Amazon	-	69.8 +/- 1.0	70.0 +/- 1.7
DSLR	91.7 +/- 2.6	-	99.7 +/- 0.6
Webcam	89.2 +/- 1.3	97.5 +/- 0.6	-
JAN	Amazon	DSLR	Webcam
Amazon	-	70.1 +/- 0.5	68.6 +/- 0.6
DSLR	92.5 +/- 1.2	-	100.0 +/- 0.0
Webcam	90.0 +/- 1.8	97.5 +/- 0.6	-
CDAN	Amazon	DSLR	Webcam
Amazon	-	69.9 +/- 2.9	65.9 +/- 1.0
DSLR	91.7 +/- 1.5	-	100.0 +/- 0.0
Webcam	89.7 +/- 1.5	97.3 +/- 0.4	-
AFN	Amazon	DSLR	Webcam
Amazon	-	70.3 +/- 0.8	70.2 +/- 1.0
DSLR	93.1 +/- 0.6	-	100.0 +/- 0.0
Webcam	89.2 +/- 0.8	98.2 +/- 0.4	-
MCC	Amazon	DSLR	Webcam
Amazon	-	74.7 +/- 1.8	77.2 +/- 0.6
DSLR	94.7 +/- 2.5	-	100.0 +/- 0.0
Webcam	92.2 +/- 1.9	98.3 +/- 0.6	-

Table 8.6: Mean and standard deviation target domain test accuracy over 5-cross-val folds for all models for all 6 adaptation tasks with tuned hyperparameters on the *Office-31* dataset. (Columns = source domain, rows = target domain)

DANN	Amazon	DSLR	Webcam
Amazon	-	72.5 +/- 1.4	74.2 +/- 1.2
DSLR	89.3 +/- 2.5	-	99.7 +/- 0.6
Webcam	88.0 +/- 2.0	97.5 +/- 1.0	-
JAN	Amazon	DSLR	Webcam
Amazon	-	69.2 +/- 0.5	69.9 +/- 0.7
DSLR	89.3 +/- 0.9	-	99.7 +/- 0.6
Webcam	94.0 +/- 2.6	96.7 +/- 0.6	-
CDAN	Amazon	DSLR	Webcam
Amazon	-	72.2 +/- 2.6	71.3 +/- 2.1
DSLR	93.3 +/- 1.9	-	100.0 +/- 0.0
Webcam	93.3 +/- 1.1	98.0 +/- 0.7	-
AFN	Amazon	DSLR	Webcam
Amazon	-	69.9 +/- 0.8	70.3 +/- 1.1
DSLR	92.0 +/- 1.6	-	100.0 +/- 0.0
Webcam	88.3 +/- 0.6	98.2 +/- 0.4	-
MCC	Amazon	DSLR	Webcam
Amazon	-	8.7 +/- 1.7	12.3 +/- 1.0
DSLR	10.7 +/- 1.6	-	56.3 +/- 4.8
Webcam	9.3 +/- 1.4	47.2 +/- 3.9	-

Table 8.7: Mean and standard deviation target domain test accuracy over 5-cross-val folds for all models for all 6 adaptation tasks with tuned hyperparameters from the *DomainNet* dataset. (Columns = source domain, rows = target domain)

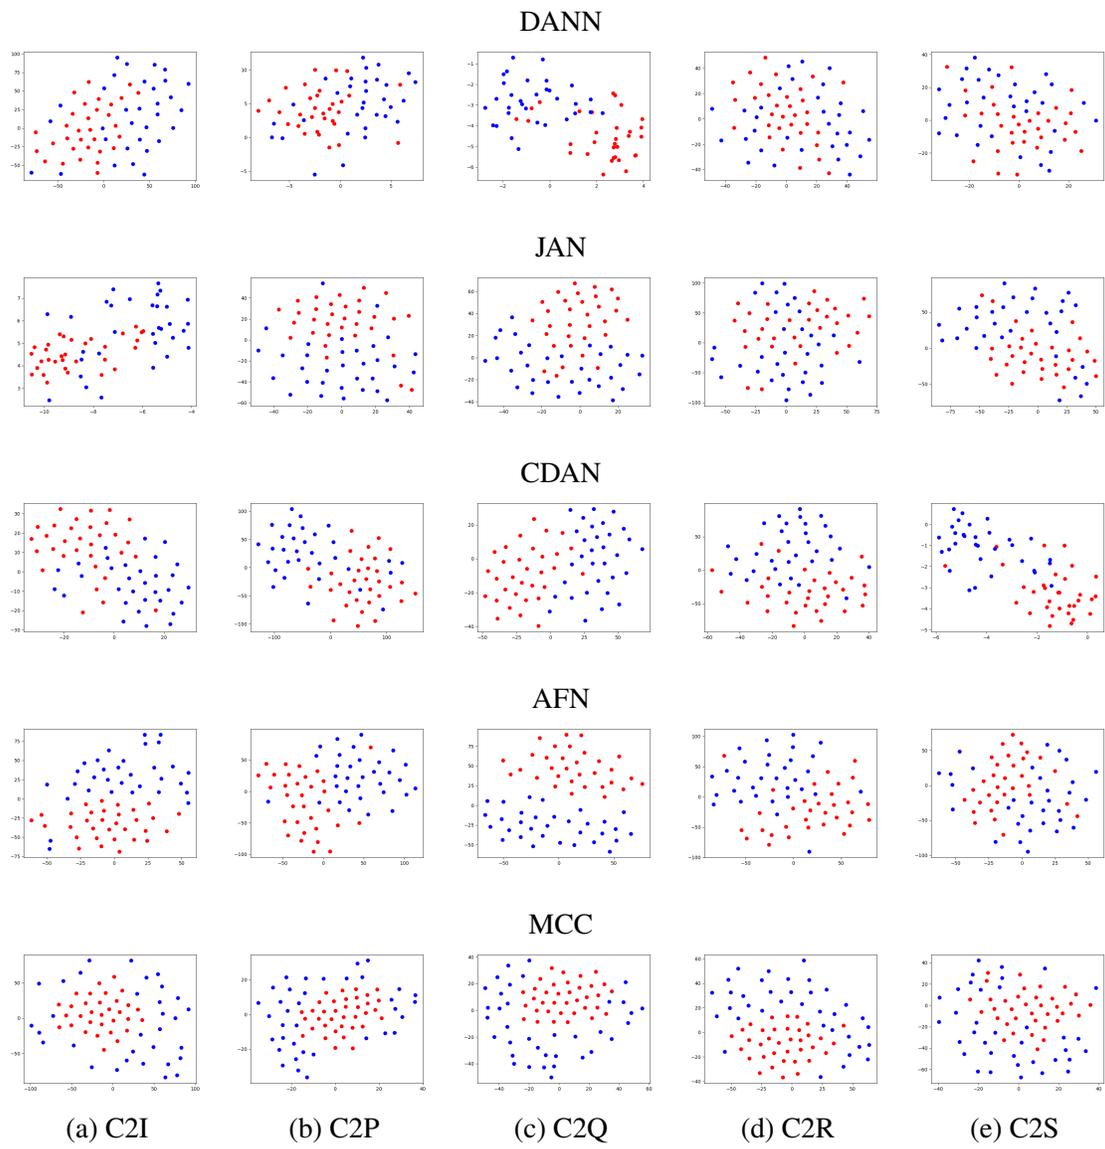


Figure 8.3: T-SNE plots for the airplane class only. Red represents the clipart source domain, blue represents the target domain. a-e represent 5 of the 30 adaptation tasks where clipart is always the source domain. (C=clipart, I=infograph, P=painting, Q=quickdraw, R=real, S=sketch) Each row is a model, top to bottom: DANN, JAN, CDAN, AFN, MCC.

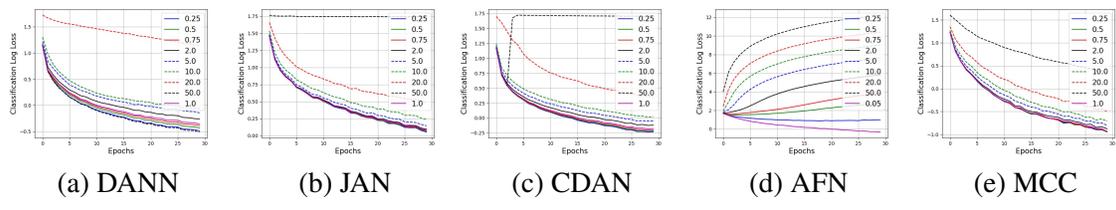


Figure 8.4: Classification loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to infograph adaptation task for full dataset.

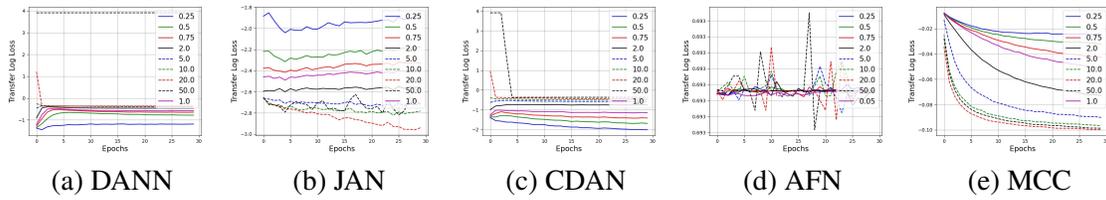


Figure 8.5: Transfer loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to infograph adaptation task for full dataset.

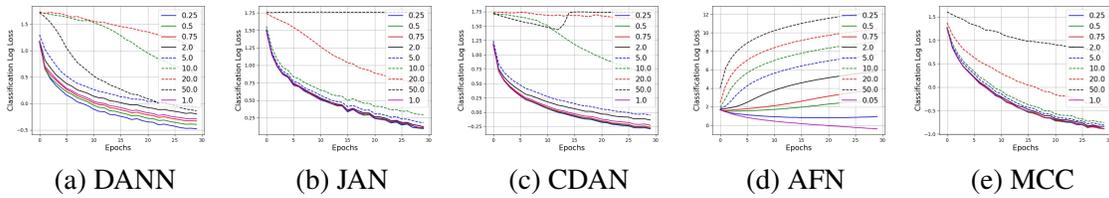


Figure 8.6: Classification loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to quickdraw adaptation task for full dataset.

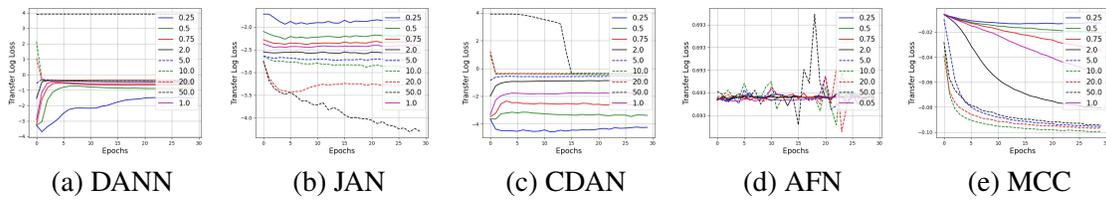


Figure 8.7: Transfer loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to quickdraw adaptation task for full dataset.

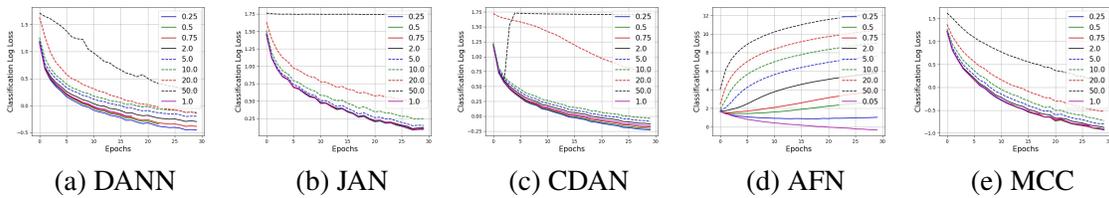


Figure 8.8: Classification loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to real adaptation task for full dataset.

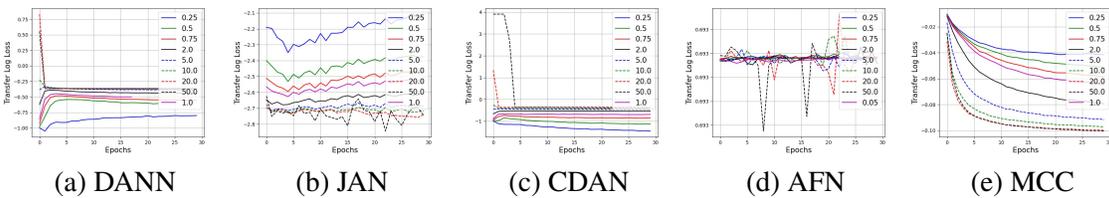


Figure 8.9: Transfer loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to real adaptation task for full dataset.

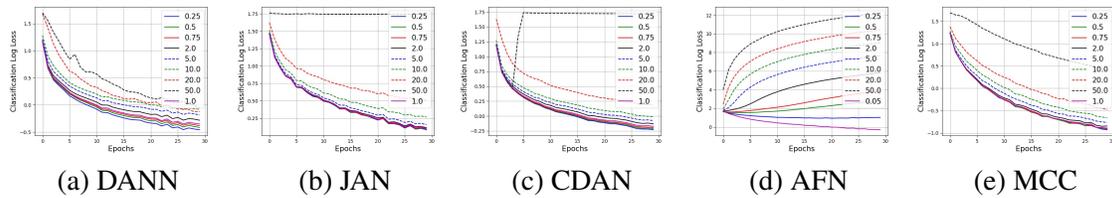


Figure 8.10: Classification loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to sketch adaptation task for full dataset.

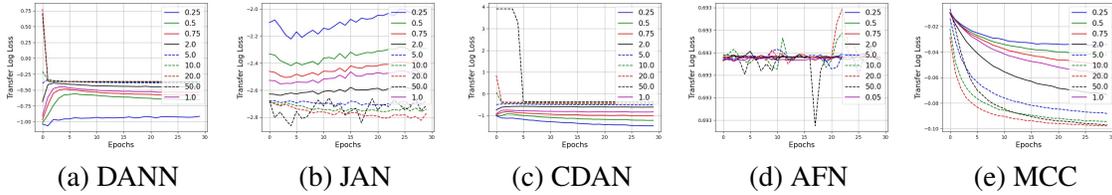


Figure 8.11: Transfer loss curves for varied loss trade-off hyperparameters (y-axis log-scale). Clipart to sketch adaptation task for full dataset.

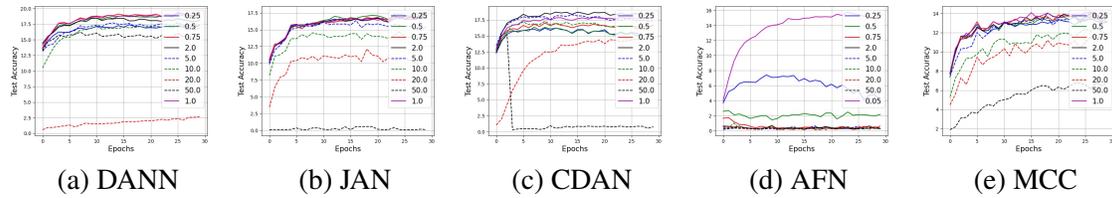


Figure 8.12: Target domain test accuracies for varied loss trade-off hyperparameters. Clipart to infograph adaptation task for full dataset.

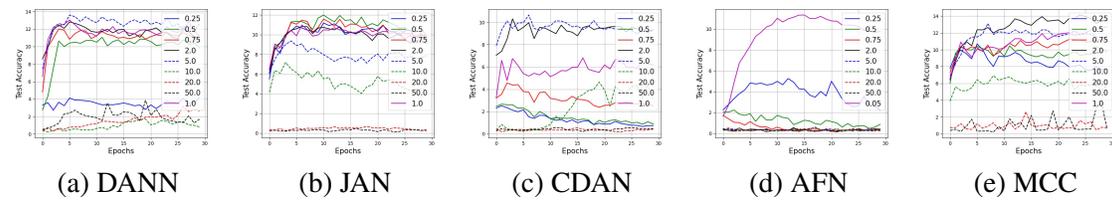


Figure 8.13: Target domain test accuracies for varied loss trade-off hyperparameters. Clipart to quickdraw adaptation task for full dataset.

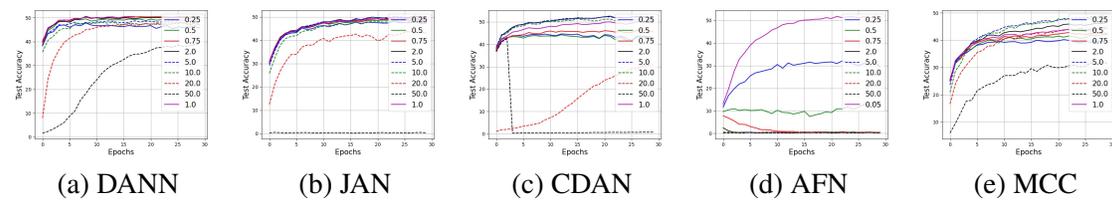


Figure 8.14: Target domain test accuracies for varied loss trade-off hyperparameters. Clipart to real adaptation task for full dataset.

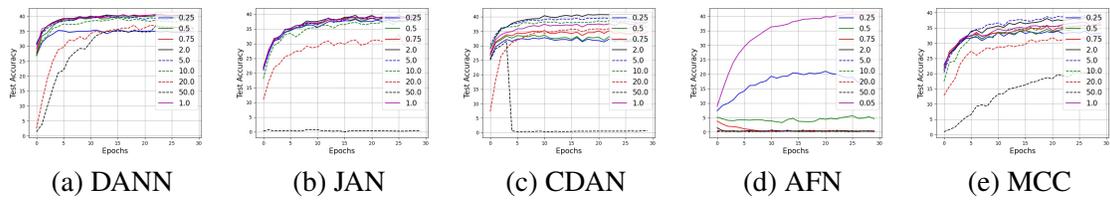


Figure 8.15: Target domain test accuracies for varied loss trade-off hyperparameters. Clipart to sketch adaptation task for full dataset.

Bibliography

- [1] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *Eur. Conf. on Comput. Vision*, pages 135–150, 2018.
- [2] Gabriela Csurka. Domain adaptation for visual applications: A comprehensive survey. *arXiv:1702.05374*, 2017.
- [3] Gabriela Csurka, Fabien Baradel, Boris Chidlovskii, and Stéphane Clinchant. Discrepancy-based networks for unsupervised domain adaptation: a comparative study. In *IEEE Int. Conf. on Comput. Vision Workshops*, pages 2630–2636, 2017.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. on Comput. vision and pattern recognition*, pages 248–255. IEEE, 2009.
- [5] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *Int. Conf. on Mach. Learn.*, pages 647–655. PMLR, 2014.
- [6] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by back-propagation. In *Int. Conf. on Mach. Learn.*, pages 1180–1189. PMLR, 2015.
- [7] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of Mach. Learn. research*, 17(1):2096–2030, 2016.
- [8] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Mach. Learn. Research*, 13(1):723–773, 2012.
- [9] Hao Guan and Mingxia Liu. Domain adaptation for medical image analysis: a survey. *arXiv:2102.09508*, 2021.
- [10] Janguang Jiang, Baixu Chen, Bo Fu, and Mingsheng Long. Transfer-learning-library. <https://github.com/thuml/Transfer-Learning-Library>, 2020.

- [11] Ying Jin, Ximei Wang, Mingsheng Long, and Jianmin Wang. Minimum class confusion for versatile domain adaptation. In *Eur. Conf. on Comput. Vision*, pages 464–480. Springer, 2020.
- [12] Yunsheng Li, Lu Yuan, Yinpeng Chen, Pei Wang, and Nuno Vasconcelos. Dynamic transfer for multi-source domain adaptation. In *IEEE/CVF Comput. Vision and Pattern Recognition*, pages 10998–11007, 2021.
- [13] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learn. transferable features with deep adaptation networks. In *Int. Conf. on Mach. Learn.*, pages 97–105. PMLR, 2015.
- [14] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. *arXiv:1705.10667*, 2017.
- [15] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. *Adv. in neural inform. processing syst.*, 29, 2016.
- [16] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learn. with joint adaptation networks. In *Int. Conf. on Mach. Learn.*, pages 2208–2217. PMLR, 2017.
- [17] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *IEEE Int. Conf. on Comput. Vision*, pages 754–763, 2017.
- [18] Vishal M Patel, Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Visual domain adaptation: A survey of recent adv. *IEEE signal processing magazine*, 32(3):53–69, 2015.
- [19] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. In *AAAI Conf. on Artif. Intell.*, 2018.
- [20] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *IEEE/CVF Int. Conf. on Comput. Vision*, pages 1406–1415, 2019.
- [21] Xingchao Peng, Zijun Huang, Ximeng Sun, and Kate Saenko. Domain agnostic learn. with disentangled representations. In *Int. Conf. on Mach. Learn.*, pages 5102–5112. PMLR, 2019.
- [22] Christian S Perone, Pedro Ballester, Rodrigo C Barros, and Julien Cohen-Adad. Unsupervised domain adaptation for medical imaging segmentation with self-ensembling. *NeuroImage*, 194:1–11, 2019.
- [23] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Eur. Conf. on Comput. Vision*, pages 213–226. Springer, 2010.

- [24] Khaled Saleh, Ahmed Abobakr, Mohammed Attia, Julie Iskander, Darius Nahavandi, Mohammed Hossny, and Saeid Nahvandi. Domain adaptation for vehicle detection from bird’s eye view lidar point cloud data. In *IEEE/CVF Int. Conf. on Comput. Vision Workshops*, Oct 2019.
- [25] Danielle Saunders. Domain adaptation and multi-domain adaptation for neural mach. translation: A survey. *arXiv:2104.06951*, 2021.
- [26] Tao Sun, Mattia Segu, Janis Postels, Yuxuan Wang, Luc Van Gool, Bernt Schiele, Federico Tombari, and Fisher Yu. Shift: A synthetic driving dataset for continuous multi-task domain adaptation. In *IEEE/CVF Conf. on Comput. Vision and Pattern Recognition*, pages 21371–21382, 2022.
- [27] A Torralba and AA Efros. Unbiased look at dataset bias. In *IEEE Conf. on Comput. Vision and Pattern Recognition*, pages 1521–1528, 2011.
- [28] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *IEEE Int. Conf. on Comput. vision*, pages 4068–4076, 2015.
- [29] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *IEEE Conf. on Comput. Vision and Pattern Recognition*, pages 7167–7176, 2017.
- [30] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv:1412.3474*, 2014.
- [31] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neuro-computing*, 312:135–153, 2018.
- [32] Rongguang Wang, Pratik Chaudhari, and Christos Davatzikos. Embracing the disharmony in medical imaging: A simple and effective framework for domain adaptation. *Med. Image Analysis*, 76:102309, 2022.
- [33] Ximei Wang, Liang Li, Weirui Ye, Mingsheng Long, and Jianmin Wang. Transferable attention for domain adaptation. In *AAAI Conf. on Artif. Intell.*, volume 33, pages 5345–5352, 2019.
- [34] Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *IEEE Comput. Vision and Pattern Recognition*, pages 3964–3973, 2018.
- [35] Ruijia Xu, Guanbin Li, Jihan Yang, and Liang Lin. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *IEEE/CVF Int. Conf. on Comput. Vision*, pages 1426–1435, 2019.

- [36] Yurong You, Carlos Andres Diaz-Ruiz, Yan Wang, Wei-Lun Chao, Bharath Hariharan, Mark Campbell, and Kilian Q Weinbergert. Exploiting playbacks in unsupervised domain adaptation for 3d object detection in self-driving cars. In *2022 Int. Conf. on Robotics and Autom.*, pages 5070–5077, 2022.
- [37] Jing Zhang, Zewei Ding, Wanqing Li, and Philip Ogunbona. Importance weighted adversarial nets for partial domain adaptation. In *IEEE Comput. Vision and Pattern Recognition*, pages 8156–8164, 2018.
- [38] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *Int. Conf. on Mach. Learn.*, pages 7404–7413. PMLR, 2019.
- [39] Han Zhao, Shanghang Zhang, Guanhang Wu, José MF Moura, Joao P Costeira, and Geoffrey J Gordon. Adversarial multiple source domain adaptation. *Adv. in neural inform. processing syst.*, 31, 2018.