

# A Broad Study of Pre-training for Domain Generalization and Adaptation

Donghyun Kim<sup>2</sup>, Kaihong Wang<sup>1</sup>, Stan Sclaroff<sup>1</sup>, and Kate Saenko<sup>1,2</sup>

<sup>1</sup>Dept. CS, Boston University, <sup>2</sup>MIT-IBM Watson AI Lab  
{donhk, kaiwkh, sclaroff, saenko}@bu.edu

**Abstract.** Deep models must learn robust and transferable representations in order to perform well on new domains. While domain transfer methods (*e.g.*, domain adaptation, domain generalization) have been proposed to learn transferable representations across domains, they are typically applied to ResNet backbones pre-trained on ImageNet. Thus, existing works pay little attention to the effects of pre-training on domain transfer tasks. In this paper, we provide a broad study and in-depth analysis of pre-training for domain adaptation and generalization, namely: network architectures, size, pre-training loss, and datasets. We observe that simply using a state-of-the-art backbone outperforms existing state-of-the-art domain adaptation baselines and set new baselines on Office-Home and DomainNet improving by 10.7% and 5.5%. We hope that this work can provide more insights for future domain transfer research.

**Keywords:** Transfer Learning; Pre-training; Domain Generalization; Domain Adaptation

## 1 Introduction

It is well-known that deep models often perform poorly on out-of-distribution test data [20]. Domain transfer has been an active research topic for years, aiming to learn more robust feature representations that generalize from training data (source domains) to novel data distributions (target domains). There has been significant progress in domain transfer for visual recognition tasks, such as image classification [12], semantic segmentation [55] and object detection [47].

Domain transfer consists of two steps: 1) *pre-training*, where a model is first pre-trained on an upstream task with a massive supervised dataset, *e.g.*, ImageNet, and 2) *transfer (adaptation)*, where the model is fine-tuned on downstream multi-domain data, see Fig. 1-(a). In the latter step, *Domain Adaptation (DA)* tunes on both a labeled source and an unlabeled target domain, while *Domain Generalization (DG)* tunes only on labeled source data. While many DA and DG methods (*e.g.*, adversarial learning [12, 34, 56], entropy optimization [34, 44] or clustering [22]) have been proposed and studied extensively in prior work, little attention has been paid to pre-training for domain transfer. In this paper, we provide comprehensive experiments and an in-depth analysis of pre-training.

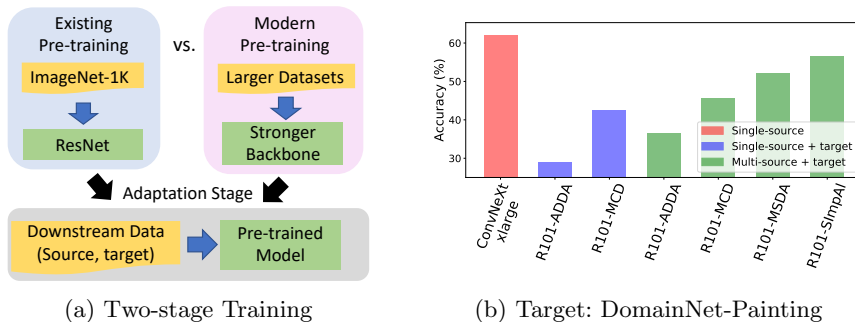


Fig. 1: (a) Existing domain transfer approaches use a ResNet pre-trained on ImageNet-1K and focus on the adaptation stage. We analyze the effect of pre-training on the adaptation stage. (b) Simply using strong pre-training outperforms all domain adaptation (DA) baselines. We compare accuracy on the target domain (*Painting*) with a SOTA architecture, *ConvNeXt-XL*, pre-trained on ImageNet 22k vs. the standard *ResNet backbones* pre-trained on ImageNet-1K. **Red** bars represent the accuracy of model trained only on a single source (*Real* domain). **Blue** bars represent DA models trained on the single source and the unlabeled target domain. **Green** bars represent DA models trained on multiple source domains and the unlabeled target domain. *ConNeXt-XL* fine-tuned only on a single source domain outperforms all the existing DA baselines

Pre-training is a very successful transfer learning technique for many visual tasks, including domain transfer tasks, as it provides a strong initial representation [9, 27]. Pre-training is especially useful when annotations are limited. We decompose pre-training into three parts: (a) network architecture (backbone), (b) dataset, and (c) loss function. It is a common practice of most domain transfer methods to use a ResNet backbone pre-trained on ImageNet-1K with a supervised loss function (*i.e.*, cross-entropy loss). We argue that this evaluation standard is outdated and ignores the effect of modern large-scale pre-training on domain transfer. To illustrate the potential impact of pre-training, Fig. 1-(b) shows an experiment that compares the performance of different backbones to SOTA results on the DomainNet [39] DA benchmark. Simply using a recent backbone [33] pre-trained with ImageNet-22K with no adaptation outperforms existing domain transfer methods. This raises the question, will SOTA DA methods still provide similar gains if applied to the stronger backbones?

To fully explore these issues, in this paper we pose the following questions:

1. **What is the effect of network architecture?** ResNet-based backbones [17] are commonly used in domain generalization [15], single source [45, 12, 34] and multi-source DA [39, 59]. Since larger and more powerful backbones such as Swin-Transformer [32] or ConvNext [33] have been recently proposed, we ask whether they may be more robust to domain shift. Transformers were recently shown to be more robust than CNNs to image corruptions and ad-

versarial examples [1]. We thus conduct an extensive analysis of the impact of network size and architectures, including state-of-the-art Transformers and CNNs on domain transfer tasks.

2. **What is the effect of pre-training dataset?** Several datasets that are larger than the standard ImageNet-1K could potentially improve transfer: ImageNet-22K [41], JFT-300M [19] and Conceptual Captions [8, 49]. These datasets have been very effective for diverse downstream visual tasks (*e.g.*, [10, 30]), but not well explored for domain transfer tasks. We therefore study the effect of a wider range of pre-training datasets, including ImageNet-21K, JFT-300M, and language-vision datasets, on domain transfer.
3. **Supervised vs. Self-supervised Pre-training.** In terms of loss functions, self-supervised learning (*e.g.*, [7, 9, 63]) has obtained powerful performance on diverse visual tasks and often outperforms its supervised counterparts on downstream problems [3, 7, 9]. We therefore compare self-supervised and supervised pre-training for domain transfer.
4. **Domain Adaptation with SOTA Pre-training.** Finally, we investigate a fundamental research question in domain transfer: With the help of the state-of-the-art pre-trained models, do we still need sophisticated domain adaptation methods? We explore the applicability of several existing DA methods to our more advanced pre-training setting.

We conduct the study on four standard multi-domain benchmarks. While we find that with better pre-training, DA methods still improve performance compared to a source-only trained model, an outdated DA method outperforms state-of-the-art DA methods. This raises serious fundamental research questions about the current evaluation protocol.

In summary, our work’s main contribution is to provide the field with a broad comparison of modern pre-training approaches for domain transfer tasks. To our knowledge, this is the first work to do such an in-depth analysis. One of our key findings is that SOTA pre-training outperforms SOTA domain transfer methods by a large margin even without access to a target domain, as shown in Fig. 1-(b). We also observe network architectures, sizes, and pre-training datasets play a big role but in a domain-dependent way. Finally, we show that SOTA DA methods work less than older DA methods under modern pre-training. We hope our work will modernize current domain transfer benchmarks and provide helpful and practical insights for future domain adaptation research.

## 2 Related Work

**Domain Transfer.** Domain transfer tasks aim to improve generalization and mitigate domain shift between source and target domains. We study generalization to *natural* data shifts caused by the changes in visual styles, background, lighting, etc. [18, 26, 42] as opposed to artificial corruptions [14]. In this problem setup, we are given a single source domain or multi-source labeled domains. The key is how to learn transferable features that will be useful for the unlabeled target domain. Depending on the specific setup, this task can be categorized into

two: (1) domain adaptation (DA) where we can access the target domain and (2) domain generalization (DG) where we do not have access to a target domain. Depending on the number of labeled source domains, each category can be further divided into single-source or multi-source DA (or DG). Typically, there are two stages: (1) pre-training and (2) adaptation. Most of these methods focus on the second adaptation stage for domain alignment with adversarial domain classifier [12, 34], entropy optimization [34, 44, 25], minimizing maximum discrepancy across domain distributions [48, 69], maximum mean discrepancy [35], or optimal transport [5]. While domain alignment methods have been proposed actively in recent years for the adaptation stage, the importance of the pre-training stage has not been well explored. Pre-training can provide strong weight initialization by learning a general transferable representation that can be useful for diverse downstream tasks [27]. While typical DA or DG methods use ResNet backbones pre-trained on ImageNet-1K, we focus on the pre-training stage and provide an in-depth analysis of its effects on domain transfer tasks.

**Network Architectures and Datasets for Pre-training.** Since the transferability of the model is closely correlated with the performance on downstream tasks as shown in [27], having a strong pre-trained model is important. In terms of architectures, convolutional neural networks (CNN) has been standard and state-of-the-art models in many visual tasks for years. After the introduction of AlexNet [29] with ImageNet-1K [41], new CNN-based architectures have been proposed with deeper, wider, and more effective convolutional layers, *e.g.*, VGGNet [50], ResNe(X)t [17, 65], SENet [21], EfficientNet [52], and ConvNeXt [33]. A newer line of work uses self-attention layers or Transformers for vision. Inspired by the Transformers for NLP [58], transformers for vision have been introduced in Vision Transformers (ViT) [10] and shows encouraging results by training with larger training sets than ImageNet-1K such as JFT-300M [51] or ImageNet-22K. DeiT [54] propose an efficient training strategy to train ViT. Swin Transformers employ a hierarchical transformer with a sliding window strategy where self-attention is performed within a local window. Swin Transformers achieve state-of-the-art performance in a range of computer vision tasks including object detection and segmentation. Bai *et al.* [1] show that Transformers can improve the generalization capability on out-of-distribution samples compared to CNNs. However, Liu *et al.* [33] propose ConvNeXt, which modernizes the ResNet architecture and show that a CNN can still outperform Transformers in vision and more robust to distribution shift. In addition to network architectures, it is shown that larger pre-training datasets such as ImageNet-22K, JFT-300M, or image-text pairs can further improve the transferability [10, 23, 30, 33, 40, 64]. Inspired by these observations, we further study the effect of backbones and pre-training datasets on domain transfer evaluation benchmarks.

**Self-supervised learning.** Self-supervised learning [11, 13, 37, 63] devises pretext tasks with self-supervisory signals without requiring human annotations. These pretext tasks allow a model to learn discriminative and transferable representations with only unlabeled data for later use in downstream tasks. Representative methods include: solving a jigsaw puzzle [37], rotation prediction [13],

Instance Discrimination (ID) [9, 16, 63], contrasting cluster assignments [6], self knowledge distillation [7], and masked image modeling [2]. Instance Discrimination [63] learns an embedding that maps visually similar images closer to each other and far from dissimilar images by classifying an image as its unique class. Some of these self-supervised methods outperform the supervised pre-training on several downstream tasks. For example, SwAV outperforms its supervised pre-training on object detection and image classification tasks on VOC [57] and INaturalist [57]. It is notable that in the VisDA-2021 competition for universal domain adaptation [3], the self-supervised masked image modeling approach with a transformer backbone [2] is the first place solution. We further investigate the effect of self-supervised pre-training approaches for domain transfer tasks.

### 3 Analysis Setup

Our goal is to analyze the effect of pre-training on domain transfer tasks. We assume a single source domain  $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$  with  $N_s$  images  $x$  and labels  $y$  and an unlabeled target domain  $\mathcal{D}_t = \{\mathbf{x}_i^t\}_{i=1}^{N_t}$ . Given a pre-trained model  $f$ , we evaluate two types of domain transfer tasks: 1) domain generalization, *i.e.* fine-tune  $f$  on  $\mathcal{D}_s$  and test on  $\mathcal{D}_t$ , and 2) domain adaptation, *i.e.* fine-tune  $f$  on  $\mathcal{D}_s, \mathcal{D}_t$  and test on  $\mathcal{D}_t$ .

**Pre-training Datasets.** Typically, ImageNet-1K is widely used for pre-training. ImageNet-1K contains 1.2M images of mutually exclusive 1000 classes. ImageNet-22K (the superset of ImageNet-1K) is also used for pre-training (*e.g.*, [32, 33]), which contains 14.1M images of 22K classes. In addition, Xie *et al.* [64] use a larger dataset JFT-300M to further improve the accuracy. Recently, language-vision models [23, 40, 30] can be used in image classification using image and text description pairs. We choose ALBEF [30], which achieves the-state-of-the-art performance and uses publicly available language-vision datasets. In total, ALBEF is pre-trained on ImageNet-1K, web crawled datasets (Conceptual Captions [8, 49], SBU Captions [38]) and two human annotated datasets (COCO[31] and Visual Genome [28]). We explore models pre-trained on these datasets.

**Downstream Datasets.** We choose Office-Home (OH) [60], DomainNet (DN) [39], CUB [62, 61], and iWildCAM2020 (WILD) [4, 26, 43]. Office-Home contains 15K images from 4 domains (Real (Rw), Painting (Pa), Clipart (Cl), Art (Ar)) on 65 classes. DomainNet contains 586K images from 6 domains (Clipart (Cl), Infograph (In), Painting (Pa), Quickdraw (Qu), Real (Rw), Sketch(Sk)) on 345 classes. Office-Home and DomainNet contain many common classes with ImageNet such as a chair, clock, and table. CUB contains 15K images from two domains (Real and Painting) on 200 fine-grained bird classes. For WILD, the source domain and target domain contains 182 different animal species from camera traps in different locations spread across multiple countries in the world. The source domain contains 129K images from 243 camera traps and the target domain contains 14K from 32 camera traps. Then we use the test set with 42K images from 48 different camera traps. The sets of camera traps in source and target domains are disjoint. The images of WILD are all realistic images from

camera traps. ImageNet contains many animal classes, but the annotations are more fine-grained in CUB and WILD. While the main cause of domain-shift is visual styles in Office-Home, DomainNet, CUB, the domain-shift of WILD is mainly caused by location differences between camera traps in the world.

**Backbone.** For CNNs, we investigate the variants of ResNet [17], EfficientNet [52], and ConvNeXt [33]. For Transformers, we explore the variants of ViT [10], DeiT [54] and Swin [32]. Variants include different depths and sizes (*e.g.*, Swin- $\{S,B,L\}$ ).

**Self-supervised Learning for Pre-training.** In addition to the supervised pre-training, we also explore recent self-supervised learning approaches for pre-training. We study SwAV [6], MoCo [16], DINO [7], and BEiT [2].

**Challenges of Evaluating Models.** We use pre-trained models, which are publicly available. One of the big challenges is that it is not possible to fairly compare all possible combinations of pre-training datasets, backbones, and self-supervised learning. For example, SwAV (self-supervised method) only provides pre-trained models on one architecture (*e.g.*, ResNet-50) and ImageNet-1K. Therefore, it is not possible to fairly compare these with self-supervised pre-training and supervised pre-training, which use the state-of-the-art architecture (*e.g.*, ConvNeXt) and larger datasets (*e.g.*, ImageNet-22K).

**Fine-tuning Details.** From each pre-trained model, we fine-tune the model with a downstream dataset. We use source domain data to train a model and keep 20% of the source data as a validation set. We choose the learning rate and optimizer by tuning on the validation set. We test different learning rates (lr=1e-1, 1e-2, 1e-3) of SGD and learning rates (lr=1e-3, 1e-4, 1e-5) of the Adam optimizer. We add a new FC layer for downstream tasks and train it from scratch with a learning rate 10 times that of the pre-trained layers. We use the image size of  $224 \times 224$  with random resized cropping. We also use random color jittering, gray-scaling, and horizontal flipping for augmentation for all models. Additional training details can be found in appendix.

## 4 Experiments

In this section, our goal is to explore the effects of pre-training for domain transfer. We reiterate that most of the prior domain adaptation (DA) or generalization (DG) work use a ResNet backbone pre-trained on ImageNet-1K. A model is denoted by  $X$ - $Y$  where  $X$  represents the name of architecture and  $Y$  represents the size of the backbone. For example, Swin-T, Swin-S, and Swin-B represent the tiny, small, and base model of Swin Transformer. Unless specified otherwise, pre-trained models are trained with a supervised loss (*i.e.*, cross-entropy loss). We now evaluate different pre-trained models in domain transfer tasks. In Sec. 4.1, we investigate single source DG and analyze the architecture, pre-training datasets, and loss functions. We also compare these models with the existing DA works. In Sec. 4.2, we explore the existing DA with new architectures. Lastly, we provide feature analysis in Sec. 4.3. Our code is available at: [https://github.com/VisionLearningGroup/Benchmark\\_Domain\\_Transfer](https://github.com/VisionLearningGroup/Benchmark_Domain_Transfer).

Table 1: Accuracy comparison on architectures in single source domain generalization. Each backbone pre-trained on ImageNet1-K is fine-tuned on a single domain (*Real*) and tested on the remaining domains in each benchmark. Recent architectures achieve higher accuracy than ResNet

Backbone	Pre-train. Data	Params	Office-Home			CUB	WILD	DomainNet					AVG
			Ar	Cl	Pr	Pa	-	Cl	In	Pa	Qu	Sk	
ResNet-50	ImageNet-1K	23 M	66.1	49.0	77.2	42.3	70.7	46.6	17.3	45.2	6.5	35.3	45.6
ConvNeXt-T	ImageNet-1K	27 M	67.4	48.7	77.9	42.5	74.0	<b>52.8</b>	<b>20.6</b>	<b>50.8</b>	<b>7.8</b>	<b>41.2</b>	48.4
DeiT-S	ImageNet-1K	21 M	70.0	<b>51.3</b>	<b>81.2</b>	<b>58.0</b>	73.4	49.5	19.4	49.2	6.9	36.1	<b>49.5</b>
Swin-T	ImageNet-1K	27 M	<b>71.3</b>	49.4	81.1	52.0	<b>74.2</b>	51.8	19.8	49.3	7.3	37.3	49.3
ResNet-101	ImageNet-1K	42 M	68.5	52.4	79.9	46.1	74.0	49.3	19.2	48.6	<b>8.7</b>	38.5	48.5
ConvNeXt-S	ImageNet-1K	49 M	72.2	52.7	80.9	43.7	76.2	54.9	22.2	<b>52.8</b>	8.1	<b>43.0</b>	50.7
Swin-S	ImageNet-1K	48 M	<b>73.8</b>	<b>54.5</b>	<b>84.2</b>	<b>56.5</b>	<b>78.6</b>	<b>55.9</b>	<b>22.5</b>	51.8	8.6	41.4	<b>52.8</b>

#### 4.1 Single Source Domain Generalization

For this experiment, we fine-tune different pre-trained models with only a single source domain. We take the *Real* domain as the source domain on Office-Home, DomainNet, and CUB and treat the remaining domains as target domains. For WILD, we follow the split in [26]. We do not use the target domain data.

**Analysis of Network Architectures.** We first compare generalization performance of architectures in Table 1. All models are pre-trained on ImageNet-1K. ConvNeXt and Transformers (DeiT, Swin) outperform their ResNet counterparts. In this experiment, Transformer models achieve the highest accuracy on average. Swin-T outperforms ResNet-50 by 3.7%. This improvement becomes larger in the deeper model. Swin-S outperforms ResNet-101 by 4.3%. We further analyze the effect of depth in a later section. The big difference between ConvNeXt and Transformers is in the CUB experiment. While DeiT-S significantly improves the accuracy by 15.7%, ConvNeXt-T could not improve much compared to ResNet. However, ConvNeXt attains slightly higher accuracy on DomainNet compared to Swin. This suggests that CNN and Transformers may be robust to different types of domain shift. We put additional results of larger networks (*e.g.*, DeiT-B) in the appendix.

**Analysis of Pre-training Datasets.** We now analyze the effect of additional datasets during pre-training as shown in Fig. 2. Due to the availability of pre-trained models (see *Challenges of Evaluating Models* in the above section), we use different backbones to compare the effect of pre-training datasets. In Fig. 2-(a), we compare the accuracy between ImageNet-1K and ImageNet-22K on Swin-B and ConvNeXt-B for Office-Home and DomainNet. We report the accuracy averaged over all settings in each benchmark. In both architectures, pre-training with ImageNet-22K boosts the accuracy for all benchmarks. Especially, there are significant boosts in accuracy on Office-Home and CUB. To be specific, the accuracy of ConvNeXt-B on CUB increases by 19.8%. In Fig. 2-(b), we study the effect of JFT-300M. Since a supervised pre-trained model on JFT-300M is not released publicly, we use the publicly available self-trained EfficientNet-B7 [64] on ImageNet-1K and JFT-300M. While it shows similar accuracy improvements on Office-Home, the accuracy improvements on CUB are smaller than ImageNet-

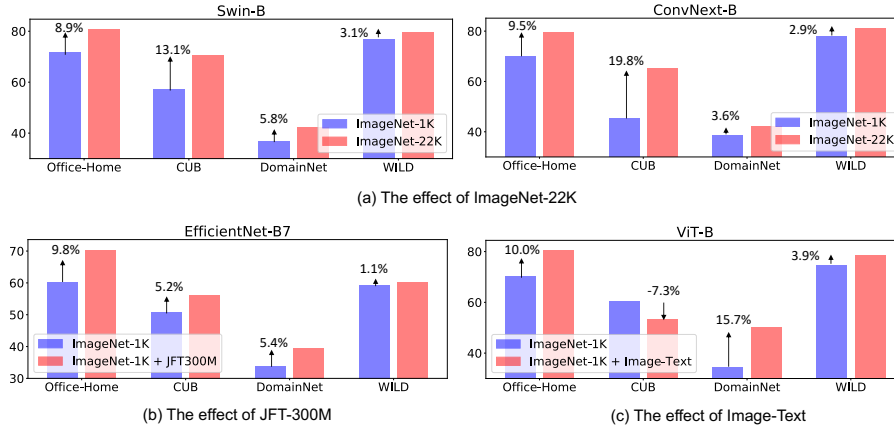


Fig. 2: The effect of additional datasets in pre-training on single source domain generalization. We investigate the following pre-training datasets, (a) ImageNet-22K, (b) JFT-300M, and (c) Image-Text pairs

22K. This could be because self-training uses pseudo-labeling and the number of classes is still limited to 1K. In Fig. 2-(c), we study the effect of vision-language datasets containing image-text pairs. We study ALBEF, which uses ViT-B as an image encoder and BERT<sub>base</sub> as a text encoder. We made one modification for ALBEF. We first extract the sentence representation from the text encoder with the prompt template “A photo of a {label}” following [40]. Then we initialize the weights of the last FC layer with the sentence representations. While it shows similar behavior on Office-Home, it seriously hurts the performance of CUB. In contrast, ALBEF obtains the highest accuracy on DomainNet compared to all the other models. This indicates that improvement depends on both the pre-training dataset and the downstream task at the same time.

**Analysis of Network Depth.** We investigate the accuracy gained from the deeper layers. In Fig. 3, we report the accuracy of ConvNeXt- $\{B,L,XL\}$ , Swin- $\{B,L\}$ , and ViT- $\{S,B,L\}$  pre-trained on ImageNet-22K. In general, ViT shows bigger changes in accuracy according to its depth. The accuracy of Swin and ConvNeXt increases slightly on the benchmarks except for CUB. In WILD, all the architectures show minimal changes according to their depth. When comparing the adaptation methods, we often choose a shallow and light model to reduce computational costs. For example, most of the domain generalization methods [15] employ ResNet-18 rather than ResNet-101. However, the results suggest that we should be careful when choosing a backbone. If we compare adaptation methods and want to use a shallow model for efficiency, it is more desirable to use Swin or ConvNeXt than ViT.

**Supervised vs. Self-supervised Learning.** Recently, several self-supervised learning (*e.g.*, [7, 9, 63]) methods outperformed their supervised counterparts on various downstream tasks [2, 3, 6, 7, 9]. In Table 2, we investigate the effect of






Fig. 3: The effect of architecture depth (number of layers) on single source domain generalization. While ConvNeXt and Swin tend to have slight improvements in accuracy, but the accuracy of ViT is more sensitive to its depth

Backbone	Training	Dataset			
		OH	CUB	WILD	DN
<b>(a) Pre-training Data: ImageNet-1K</b>					
ResNet-50	Sup.	<b>64.1</b>	<b>42.3</b>	<b>70.7</b>	30.2
ResNet-50	MoCo	53.7	37.1	67.9	28.0
ResNet-50	SwAV	56.5	39.2	68.4	<b>30.7</b>
ResNet-50	DINO	56.6	37.2	67.6	23.1
DeiT-S	Sup.	<b>58.0</b>	<b>58.0</b>	73.4	<b>32.2</b>
ViT-S	DINO	56.3	57.9	<b>75.1</b>	29.8
<b>(b) Pre-training Data: ImageNet-22K</b>					
ViT-L	Sup.	<b>82.3</b>	76.5	77.7	43.3
ViT-L	Sup.+BEiT	80.9	<b>78.4</b>	<b>79.7</b>	<b>45.1</b>

Table 2: Accuracy comparison between supervised and self-supervised learning in pre-training. (a) In this domain generalization task, supervised pre-training outperforms the self-supervised pre-training in most cases. (b) When self-supervised learning is combined with supervised learning, the performance can be improved

self-supervised learning approaches in pre-training for domain transfer. In Table 2-(a), we observe that supervised learning (denoted as Sup.) performs better than self-supervised learning in most cases. Especially on Office-Home and CUB, self-supervised learning significantly hurts performance compared to supervised learning. In Table 2-(b), we explore the BEiT, the winner of the VisDA 2021 DA competition [3]. BEiT uses both self-supervised and supervised learning. When self-supervised learning is combined with supervised learning and improves the performance on CUB, DomainNet (denoted as DN), and WILD. While most of the self-supervised approaches focus on only unlabeled data, combining it with labels should be considered to improve further in future research.

**Comparison with Domain Adaptation Baselines.** We provide a performance comparison with the existing DA baselines in Table 3. In this comparison we use Office-Home and DomainNet, which are most extensively explored in prior

Table 3: Accuracy comparison with domain adaptation baselines on (a) Office-Home and (b) DomainNet. (I): For single source adaptation, an unlabeled target domain is used for alignment with a single source domain. (II) For multi source adaptation, all the domains except the target domain in each benchmark are used as the source domains along with the unlabeled target domain. (III) We change the pre-training with the state-of-the-art backbone and larger datasets. Even though these are trained with only one single domain **without any adaptation**, simply using the state-of-the-art pre-training outperforms the existing domain adaptation baselines by up to 10.7% and 5.5% on average on each dataset. \* represents the numbers reported in [24, 39, 53, 59]

(a) Office-Home

Backbone	Pre-training Data	Downstream Data	Adaptation	Target Domain			AVG
				Ar	Cl	Pr	
<b>(I) Single source domain adaptation baselines</b>							
ResNet-50*	ImageNet-1K	Source + Target	DANN [12]	63.2	51.8	76.8	63.9
ResNet-50*	ImageNet-1K	Source + Target	CDAN [34]	70.9	56.7	81.6	69.7
ResNet-50*	ImageNet-1K	Source + Target	AFN [67]	70.9	57.1	81.5	69.8
ResNet-50*	ImageNet-1K	Source + Target	MDD [69]	72.5	60.2	82.3	71.7
ResNet-50*	ImageNet-1K	Source + Target	SRDC [53]	76.3	57.1	85.0	72.8
<b>(II) Multi source domain adaptation baselines</b>							
ResNet-50*	ImageNet-1K	Multi source + Target	MFSAN [70]	72.1	62	80.3	71.5
ResNet-50*	ImageNet-1K	Multi source + Target	SImpAI [59]	70.8	56.3	80.2	69.1
ResNet-101*	ImageNet-1K	Multi source + Target	SImpAI [59]	73.4	62.4	81.0	72.3
<b>(III) Single source only baselines without adaptation</b>							
ResNet-50*	ImageNet-1K	Source	<b>x</b>	53.9	41.2	59.9	46.1
ConvNeXt-XL	ImageNet-22K	Source	<b>x</b>	<b>85.1</b>	74.0	<b>91.4</b>	<b>83.5</b>
ViT-L	ImageNet-22K	Source	<b>x</b>	84.0	73.0	89.9	82.3
Swin-L	ImageNet-22K	Source	<b>x</b>	83.4	<b>74.3</b>	90.9	82.8
ViT-B	ALBEF [30]	Source	<b>x</b>	81.7	72.5	87.2	80.4

(b) DomainNet

Backbone	Pre-training Data	Downstream Data	Adaptation	Target Domain					AVG
				Cl	In	Pa	Qu	Sk	
<b>(I) Single source domain adaptation baselines</b>									
ResNet-101*	ImageNet-1K	Source + target	ADDA [56]	39.5	14.5	29.1	14.9	30.7	25.7
ResNet-101*	ImageNet-1K	Source + target	MCD [48]	42.6	19.6	42.6	3.8	30.8	27.9
<b>(II) Multi source domain adaptation baselines</b>									
ResNet-101*	ImageNet-1K	Multi source + target	ADDA [56]	47.5	11.4	36.7	14.7	33.5	28.8
ResNet-101*	ImageNet-1K	Multi source + target	MCD [48]	54.3	22.1	45.7	7.6	43.5	34.6
ResNet-101*	ImageNet-1K	Multi source + target	DCTN [66]	48.6	23.5	48.8	7.2	47.3	35.1
ResNet-101*	ImageNet-1K	Multi source + target	MSDA [39]	58.6	26.0	52.3	6.3	49.5	38.5
ResNet-101*	ImageNet-1K	Multi source + target	SImpAI [59]	66.4	26.5	56.6	18.9	55.5	44.8
<b>(III) Single source only baselines without adaptation</b>									
ResNet-101*	ImageNet-1K	Source	<b>x</b>	48.4	22.2	49.4	6.4	38.8	33.0
ConvNext-XL	ImageNet-22K	Source	<b>x</b>	67.7	29.7	62.2	11.4	55.5	45.3
ViT-L	ImageNet-22K	Source	<b>x</b>	65.5	27.3	61.3	10.2	52.1	43.3
Swin-L	ImageNet-22K	Source	<b>x</b>	67.2	30.6	62.5	11.2	54.1	45.1
ViT-B	ALBEF [30]	Source	<b>x</b>	<b>73.6</b>	<b>37.3</b>	<b>65.3</b>	<b>12.8</b>	<b>62.2</b>	<b>50.3</b>

Table 4: Single source domain adaptation with the state-of-the-art pre-training. We report the target accuracy on each benchmark. For Office-Home, CUB, and DomainNet, we use the *Real* domain as a source domain. Surprisingly, more recent adaptation methods (AFN, MDD, MCC) underperform CDAN on average

Backbone	Adaptation	Office-Home				CUB	WILD	DomainNet					
		Ar	Cl	Pr	AVG	Pa	-	Cl	In	Pa	Qu	Sk	AVG
Swin-L	Source-only	74.3	83.4	90.9	82.8	73.0	81.4	67.2	30.6	62.5	11.2	54.1	45.1
ConvNeXt-XL	Source-only	74.0	<b>85.1</b>	91.4	83.5	71.9	81.5	67.7	29.7	62.2	11.4	55.5	45.3
Swin-L	DANN [12]	87.3	79.5	93.0	86.6	82.0	68.4	70.4	36.7	66.6	13.0	60.5	49.4
ConvNext-XL	DANN [12]	87.2	79.8	93.1	86.7	80.4	66.2	70.2	36.7	66.6	8.6	62.1	48.8
Swin-L	CDAN [35]	90.1	81.9	93.1	88.4	<b>84.3</b>	81.5	72.0	<b>39.7</b>	<b>67.5</b>	10.8	61.8	50.4
ConvNext-XL	CDAN [35]	<b>90.2</b>	84.6	93.8	<b>89.5</b>	83.7	<b>82.5</b>	<b>72.1</b>	39.5	67.3	<b>13.9</b>	<b>63.4</b>	<b>51.2</b>
Swin-L	AFN [67]	87.4	77.6	92.0	85.7	79.9	78.9	68.0	34.5	64.7	6.9	58.1	46.4
ConvNext-XL	AFN [67]	86.0	77.7	92.8	85.5	77.7	79.7	68.0	33.4	65.0	7.8	59.2	46.7
Swin-L	MDD [69]	87.8	78.0	93.6	86.5	83.6	81.1	62.8	34.5	60.7	2.8	46.6	41.5
ConvNext-XL	MDD [69]	88.0	77.6	93.6	86.4	78.8	<b>82.5</b>	58.9	29.6	62.8	11.4	51.4	42.8
Swin-L	MCC [24]	89.6	81.1	94.1	88.3	66.7	79.5	71.5	36.4	66.3	3.0	58.2	47.1
ConvNext-XL	MCC [24]	89.5	82.9	<b>94.4</b>	88.9	78.8	79.8	71.3	27.7	66.8	1.9	60.2	45.6

work. In each table, (I) reports the performance of single source DA, where a labeled source domain and unlabeled target domain data is used together for adaptation. In (II), multiple labeled source domains are used in addition to the unlabeled target domain. (I,II) use ResNet backbones pre-trained on ImageNet-1K. (III) only uses one single source (*Real*) domain with the recent backbones and larger pre-training datasets. It is surprising that (III), the state-of-art backbones pre-trained on a larger dataset, significantly outperform DA baselines (I, II) despite being trained only on a single source and *not using any adaptation on the target domain*. While ConvNext-XL obtains the best results and outperforms adaptation baselines by up to 10.7% on Office-Home, ViT-B with ALBEF gains more improvements and outperforms adaptation baselines by up to 5.5% on DomainNet. This observation raises a question, is it still fair and reasonable to use ResNet backbones pre-trained on ImageNet-1K as a standard backbone for the comparison of adaptation methods? From the results, it is clear that the the standard backbone in the existing domain adaptation benchmarks are outdated. The pre-training stage for DA needs to be updated to reflect recent advances in pre-training.

## 4.2 Domain Adaptation with Modern Pre-training

The observation in Sec. 4.1 leads us to the next question, will SOTA DA methods still provide gains when these are applied to the recent stronger architectures and pre-training? We study the transferability of prior adaptation methods with new architectures pre-trained on larger datasets.

**Transferability of Domain Adaptation.** We employ DANN (JMLR’16), CDAN (NeurIPS’18), AFN (ICCV’19), MDD (ICML’19), and MCC(ECCV’20). Table 4 provides the performance of domain adaptation between Swin-L and ConvNeXt-XL pre-trained on ImageNet-22K. First, DA methods still improve the accuracy on average compared to source only (SO) models. However, we




Fig. 4: Analysis on the transferability of hyper-parameters in adaptation methods between a shallow and deep network. We report the accuracy of CNN and Transformers on the *Real*  $\rightarrow$  *Clipart* setting in Office-Home. We observe that the shallow and deep network achieve the highest accuracy on the same value of hyper-parameters in both CDAN and MCC

observe negative transfer in some settings where DA hurts the performances compared SO. To be specific, all the adaptation methods badly affect the performance on *Real*  $\rightarrow$  *Clipart*(*Cl*) in Office-Home. Second, the relative ranking among adaptation methods on new architectures is different from the ranking on ResNet-based backboes. While MCC obtains SOTA accuracy on ResNet, but CDAN outperforms AFN, MDD, and MCC with these new architectures, which was proposed earlier than the others. This behavior raises another practical question, which adaptation method should we consider the state-of-the-art? We certainly want to have a model with strong performance, but the existing adaptation benchmark with outdated pre-training can not choose the strongest model. We argue that adaptation methods should have transferability on various backbones in order to avoid the potential risk of overfitting to a specific backbone, so that it is able to successfully transfer to new architectures in future.

**Analysis on Hyper-parameter.** In this experiment, we analyze the transferability of hyper-parameters to new architectures in each adaptation method. The question is whether the optimal hyper-parameters in a shallow network (*e.g.*, ResNet-50) are still optimal in a deep network (*e.g.*, ConvNeXt-XL). We believe this is a practically important question as training a big network is computationally expensive. Hyper-parameter search in a big network can be prohibitively expensive in terms of computational cost. Therefore, the desired property is that hyper-parameters in adaptation methods are transferable between different architectures and their depth. Fig. 4 shows the analysis of performance depending on the hyper-parameters between a shallow and deep network. We investigate two adaptation methods CDAN and MCC. Following [46], we vary the trade-off ( $\lambda = 0.05, 0.1, 0.5, 1.0, 2.5, 5.0$ ) hyper-parameter in CDAN, which controls the trade-off between domain confusion loss and classification loss on the source domain. The default  $\lambda$  in ResNet is 1.0. For MCC, we vary the temperature hyper-parameter ( $\eta=1.5, 2.0, 2.5, 3.0, 3.5$ ), which affects of classifier’s confusion loss. The default  $\eta$  in ResNet is 2.5. We employ ResNet-50 as a shallow network and ConvNeXt-XL as a deep network for CNN, and Swin-S as a shallow network and Swin-L as a deep network. The shallow networks are pre-trained on ImageNet-1K and deep networks are pre-trained on ImageNet-22K. The accuracy




Fig. 5: t-SNE visualization of pre-trained models. We use the Real (**red**) and Clipart (**blue**) domains in Office-Home. We directly utilize features from each pre-trained model without fine-tuning. Compared to (a) the ResNet-50 pre-trained on ImageNet-1K, (b) ConvNext-XL and (c) Swin-L pre-trained on ImageNet-22K produce more discriminative as well as domain-aligned representations

across the depth (shallow vs. deep) and architectures (CNN vs. Transformers) show similar tendency and obtain the highest accuracy with the same hyper-parameter values. Therefore, we observe that the default hyper-parameters of  $\lambda$  and  $\eta$  in shallow networks are transferable to deep networks. Additionally, the sensitivity of deep networks is small compared to that of shallow networks. We measure the standard deviation of accuracy. Swin-L obtains a standard deviation of 4.2%, but Swin-S obtains a much higher standard deviation of 8.3%.

### 4.3 Feature Analysis

We provide feature analysis in this section. We use features directly obtained from each pre-training without fine-tuning on domain adaptation benchmarks.

**Feature Visualization.** First, we show the t-SNE [36] feature visualization of each pre-trained models in Fig. 5. We compare (a) ResNet-50 pre-trained on ImageNet-1K with (b) ConvNeXt-XL, (c) Swin-L pre-trained on ImageNet-22K, and (d) ViT-B from ALBEF [30]. We directly extract features from the pre-trained models on the Real (colored by **red**) and Clipart (colored by **blue**) domains in Office-Home. While the red and blue dots are highly separated in ResNet-50, these are aligned with each other in ConvNeXt-XL and Swin-L. It is also clear that ConvNeXt-XL and Swin-L obtain better clustered and discriminative representations. (d) ViT-B pre-trained (ALBEF) on image-text pairs shows different patterns, where the blue dots are red dots are somewhat aligned but it does not provide clustered representations. This could be probably due to that the pretext task of ALBEF is to align image and text, but not classification.

**Analysis on Feature Transferability.** We employ LogME [68] to evaluate the transferability of features for downstream tasks. LogME is used to assess pre-trained models, which estimates the maximum value of label evidence from the extracted features of downstream data. A higher value LogME implies better transferability to downstream tasks. We measure LogME of each pre-training for




Fig. 6: Analysis on feature transferability of each pre-trained model. ResNet pre-trained models on ImageNet-1K obtain lower values of LogME than that of state-of-the-art backbones, which implies weak transferability to DA benchmarks

all domains in the benchmarks as shown in Fig. 6. As expected, ResNet backbones pre-trained on ImageNet-1K obtain very low values of LogME compared to the state-of-the-art backbones. We reiterate that the pre-training stage should be modernized according to the recent advances in computer vision.

## 5 Conclusions

Most domain transfer works pay little attention to the importance of the pre-training stage. In this work, we provide an in-depth analysis of the effect of modern pre-training on domain transfer. We summarize some of our key findings:

1. **What makes strong pre-training for domain transfer?** In Sec. 4.1, we observe that many factors, including network architecture, pre-training dataset, and network size contribute to the improvements in domain transfer tasks. However, there is no single winner across all benchmark datasets. The transferability of pre-training depends on the target benchmark, adaptation method, and network depth. Most importantly, we observe that simply using the SOTA pre-training outperforms all the domain adaptation baselines.
2. **Do we still need domain adaptation with modern pre-training?** In Sec. 4.2, while we find that adaptation methods still improve the accuracy with modern pre-training, the relative ranking of domain adaptation methods is not preserved. With modern pre-training, an outdated DA method performs better than more recent DA methods in our experiments.

**Limitations.** Due to the availability of pre-trained models, we could not analyze full ablations (*e.g.*, SOTA backbones pre-trained on JFT-300M). In this work, we use a very simple fine-tuning strategy by adding a single FC layer, but there could be other simple better ways to fine-tune pre-trained models for downstream tasks. In addition to image classification tasks, other computer vision tasks including domain adaptive object detection, segmentation, or video domain adaptation should be explored with modern pre-training in future research. We hope future work should use these results as a new baseline.

**Acknowledgments.** This work was supported by DARPA LwLL and NSF Award No. 1535797

## References

1. Bai, Y., Mei, J., Yuille, A.L., Xie, C.: Are transformers more robust than CNNs? In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 34 (2021)
2. Bao, H., Dong, L., Piao, S., Wei, F.: BEit: BERT pre-training of image transformers. In: *International Conference on Learning Representations (2022)*, <https://openreview.net/forum?id=p-BhZSz59o4>
3. Bashkirova, D., Hendrycks, D., Kim, D., Mishra, S., Saenko, K., Saito, K., Teterwak, P., Usman, B.: Visda-2021 competition universal domain adaptation to improve performance on out-of-distribution data. *arXiv preprint arXiv:2107.11011* (2021)
4. Beery, S., Cole, E., Gjoka, A.: The iwildcam 2020 competition dataset. *arXiv preprint arXiv:2004.10340* (2020)
5. Bhushan Damodaran, B., Kellenberger, B., Flamary, R., Tuia, D., Courty, N.: Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In: *European Conference on Computer Vision (ECCV)*. pp. 447–463 (2018)
6. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 33, pp. 9912–9924 (2020)
7. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: *IEEE International Conference on Computer Vision (ICCV)*. pp. 9650–9660 (2021)
8. Changpinyo, S., Sharma, P., Ding, N., Soricut, R.: Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3558–3568 (2021)
9. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning (ICML)*. pp. 1597–1607. PMLR (2020)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations (ICLR)* (2021), <https://openreview.net/forum?id=YicbFdNTTy>
11. Dosovitskiy, A., Fischer, P., Springenberg, J.T., Riedmiller, M., Brox, T.: Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **38**(9), 1734–1747 (2015)
12. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *Journal of Machine Learning Research (JMLR)* **17**(1), 2096–2030 (2016)
13. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728* (2018)
14. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014)
15. Gulrajani, I., Lopez-Paz, D.: In search of lost domain generalization. In: *International Conference on Learning Representations (2021)*, <https://openreview.net/forum?id=lQdXeXDwTl>
16. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 9729–9738 (2020)

17. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
18. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=HJz6tiCqYm>
19. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015)
20. Hoffman, J., Tzeng, E., Donahue, J., Jia, Y., Saenko, K., Darrell, T.: One-shot adaptation of supervised deep convolutional models. In: International Conference on Learning Representations (ICLR) (2014)
21. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7132–7141 (2018)
22. Huang, J., Dong, Q., Gong, S., Zhu, X.: Unsupervised deep learning by neighbourhood discovery. arXiv preprint arXiv:1904.11567 (2019)
23. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: International Conference on Machine Learning (ICML). pp. 4904–4916. PMLR (2021)
24. Jin, Y., Wang, X., Long, M., Wang, J.: Minimum class confusion for versatile domain adaptation. In: European Conference on Computer Vision (ECCV). pp. 464–480. Springer (2020)
25. Kim, D., Saito, K., Oh, T.H., Plummer, B.A., Sclaroff, S., Saenko, K.: CDS: Cross-domain self-supervised pre-training. In: IEEE International Conference on Computer Vision (ICCV). pp. 9123–9132 (2021)
26. Koh, P.W., Sagawa, S., Marklund, H., Xie, S.M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R.L., Gao, I., et al.: Wilds: A benchmark of in-the-wild distribution shifts. In: International Conference on Machine Learning (ICML). pp. 5637–5664. PMLR (2021)
27. Kornblith, S., Shlens, J., Le, Q.V.: Do better ImageNet models transfer better? In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2661–2671 (2019)
28. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision (IJCV)* **123**(1), 32–73 (2017)
29. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems (NIPS)*. vol. 25 (2012)
30. Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 34 (2021)
31. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: *European Conference on Computer Vision (ECCV)*. pp. 740–755. Springer (2014)
32. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *IEEE International Conference on Computer Vision (ICCV)*. pp. 10012–10022 (2021)
33. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A ConvNet for the 2020s. arXiv preprint arXiv:2201.03545 (2022)



34. Long, M., Cao, Z., Wang, J., Jordan, M.I.: Conditional adversarial domain adaptation. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 1640–1650 (2018)
35. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Unsupervised domain adaptation with residual transfer networks. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 136–144 (2016)
36. Van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *Journal of Machine Learning Research (JMLR)* **9**(11) (2008)
37. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: *European Conference on Computer Vision (ECCV)*. pp. 69–84. Springer (2016)
38. Ordonez, V., Kulkarni, G., Berg, T.: Im2text: Describing images using 1 million captioned photographs. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 24 (2011)
39. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: *IEEE International Conference on Computer Vision (ICCV)*. pp. 1406–1415 (2019)
40. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International Conference on Machine Learning (ICML)*. pp. 8748–8763. PMLR (2021)
41. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)* **115**(3), 211–252 (2015)
42. Saenko, K., Kulis, B., Fritz, M., Darrell, T.: Adapting visual category models to new domains. In: *European Conference on Computer Vision (ECCV)*. pp. 213–226. Springer (2010)
43. Sagawa, S., Koh, P.W., Lee, T., Gao, I., Xie, S.M., Shen, K., Kumar, A., Hu, W., Yasunaga, M., Marklund, H., Beery, S., David, E., Stavness, I., Guo, W., Leskovec, J., Saenko, K., Hashimoto, T., Levine, S., Finn, C., Liang, P.: Extending the WILDS benchmark for unsupervised adaptation. In: *International Conference on Learning Representations (2022)*, <https://openreview.net/forum?id=z7p2V6KR00V>
44. Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised domain adaptation via minimax entropy. In: *IEEE International Conference on Computer Vision (ICCV)* (2019)
45. Saito, K., Kim, D., Sclaroff, S., Saenko, K.: Universal domain adaptation through self supervision. *arXiv preprint arXiv:2002.07953* (2020)
46. Saito, K., Kim, D., Teterwak, P., Sclaroff, S., Darrell, T., Saenko, K.: Tune it the right way: Unsupervised validation of domain adaptation via soft neighborhood density. In: *IEEE International Conference on Computer Vision (ICCV)*. pp. 9184–9193 (2021)
47. Saito, K., Ushiku, Y., Harada, T., Saenko, K.: Strong-weak distribution alignment for adaptive object detection. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 6956–6965 (2019)
48. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3723–3732 (2018)

49. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 2556–2565 (2018)
50. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
51. Sun, C., Shrivastava, A., Singh, S., Gupta, A.: Revisiting unreasonable effectiveness of data in deep learning era. In: IEEE International Conference on Computer Vision (ICCV). pp. 843–852 (2017)
52. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International Conference on Machine Learning (ICML). pp. 6105–6114. PMLR (2019)
53. Tang, H., Chen, K., Jia, K.: Unsupervised domain adaptation via structurally regularized deep clustering. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8725–8735 (2020)
54. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: International Conference on Machine Learning (ICML). pp. 10347–10357. PMLR (2021)
55. Tsai, Y.H., Sohn, K., Schuler, S., Chandraker, M.: Domain adaptation for structured output via discriminative patch representations. In: IEEE International Conference on Computer Vision (ICCV). pp. 1456–1465 (2019)
56. Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. arXiv preprint arXiv:1412.3474 (2014)
57. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8769–8778 (2018)
58. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need **30** (2017)
59. Venkat, N., Kundu, J.N., Singh, D., Revanur, A., et al.: Your classifier can secretly suffice multi-source domain adaptation. In: Advances in Neural Information Processing Systems (NeurIPS). vol. 33, pp. 4647–4659 (2020)
60. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5018–5027 (2017)
61. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD birds-200-2011 dataset (2011)
62. Wang, S., Chen, X., Wang, Y., Long, M., Wang, J.: Progressive adversarial networks for fine-grained domain adaptation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9213–9222 (2020)
63. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3733–3742 (2018)
64. Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10687–10698 (2020)
65. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1492–1500 (2017)

66. Xu, R., Chen, Z., Zuo, W., Yan, J., Lin, L.: Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3964–3973 (2018)
67. Xu, R., Li, G., Yang, J., Lin, L.: Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In: IEEE International Conference on Computer Vision (ICCV). pp. 1426–1435 (2019)
68. You, K., Liu, Y., Wang, J., Long, M.: Logme: Practical assessment of pre-trained models for transfer learning. In: International Conference on Machine Learning (ICML). pp. 12133–12143. PMLR (2021)
69. Zhang, Y., Liu, T., Long, M., Jordan, M.: Bridging theory and algorithm for domain adaptation. In: International Conference on Machine Learning (ICML). pp. 7404–7413. PMLR (2019)
70. Zhu, Y., Zhuang, F., Wang, D.: Aligning domain-specific distribution and classifier for cross-domain classification from multiple sources. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 5989–5996 (2019)