A Survey of Unsupervised Deep Domain Adaptation

GARRETT WILSON and DIANE J. COOK, Washington State University Pullman, Washington

Deep learning has produced state-of-the-art results for a variety of tasks. While such approaches for supervised learning have performed well, they assume that training and testing data are drawn from the same distribution, which may not always be the case. As a complement to this challenge, single-source unsupervised domain adaptation can handle situations where a network is trained on labeled data from a source domain and unlabeled data from a related but different target domain with the goal of performing well at test-time on the target domain. Many single-source and typically homogeneous unsupervised deep domain adaptation approaches have thus been developed, combining the powerful, hierarchical representations from deep learning with domain adaptation to reduce reliance on potentially costly target data labels. This survey will compare these approaches by examining alternative methods, the unique and common elements, results, and theoretical insights. We follow this with a look at application areas and open research directions.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Computing methodologies** → **Transfer learning**; *Unsupervised learning*; *Neural networks*; *Adversarial learning*;

Additional Key Words and Phrases: Domain adaptation, deep learning, generative adversarial networks

ACM Reference format:

Garrett Wilson and Diane J. Cook. 2020. A Survey of Unsupervised Deep Domain Adaptation. *ACM Trans. Intell. Syst. Technol.* 11, 5, Article 51 (July 2020), 46 pages.
<https://doi.org/10.1145/3400066>

1 INTRODUCTION

Supervised learning is arguably the most prevalent type of machine learning and has enjoyed much success across diverse application areas. However, many supervised learning methods make a common assumption: the training and testing data are drawn from the same distribution. When this constraint is violated, a classifier trained on the *source domain* will likely experience a drop in performance when tested on the *target domain* due to the differences between domains [182]. Single-source *domain adaptation* refers to the goal of learning a concept from labeled data in a source domain that performs well on a different but related target domain [74, 81, 180]. *Unsupervised* domain adaptation specifically addresses the situation where there are labeled source data and only unlabeled target data available for use during training [74, 147].

Because of its ability to adapt labeled data for use in a new application, domain adaptation can reduce the need for costly labeled data in the target domain. As an example, consider the problem of semantically segmenting images. Each real image in the Cityscapes dataset required approximately

This material is based upon work supported by the National Science Foundation under Grant Nos. 1543656 and 1734558. Authors' addresses: G. Wilson and D. J. Cook, Washington State University, School of Electrical Engineering and Computer Science, Pullman, WA, 99164; emails: {garrett.wilson, djcook}@wsu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

2157-6904/2020/07-ART51 \$15.00

<https://doi.org/10.1145/3400066>

1.5 hours to annotate for semantic segmentation [46]. In this case, human annotation time could be spared by training an image semantic segmentation model on synthetic street view images (the source domain) since these can be cheaply generated, then adapting and testing for real street view images (the target domain, here the Cityscapes dataset).

An undeniable trend in machine learning is the increased usage of deep neural networks. Deep networks have produced many state-of-the-art results for a variety of machine learning tasks [74, 81] such as image classification, speech recognition, machine translation, and image generation [80, 81]. When trained on large amounts of data, these many-layer neural networks can learn powerful, hierarchical representations [81, 147, 182, 225] and can be highly scalable [77]. At the same time, these networks can also experience performance drops due to domain shifts [73, 225]. Thus, much research has gone into adapting such networks from large labeled datasets to domains where little (or possibly no) labeled training data are available (for a list, see [256]). These single-source and typically homogeneous unsupervised deep domain adaptation approaches, which combine the benefit of deep learning with the very practical use of domain adaptation to remove the reliance on potentially costly target data labels, will be the focus of this survey.

A number of surveys have been created on the topic of domain adaptation [12, 24, 42, 48, 49, 121, 122, 162, 182, 226, 245, 284] and more generally transfer learning [45, 128, 152, 180, 215, 231, 234, 251, 272], of which domain adaptation can be viewed as a special case [182]. Previous domain adaptation surveys lack depth of coverage and comparison of unsupervised deep domain adaptation approaches. In some cases, prior surveys do not discuss domain mapping [48, 49, 121], normalization statistic-based [48, 49, 121, 284], or ensemble-based [48, 49, 121, 245, 284] methods. In other cases, they do not survey deep learning approaches [12, 122, 162, 182]. Still others are application-centric, focusing on a single use case such as machine translation [24, 42]. One earlier survey focuses on the multi-source scenario [226], while we focus on the more prevalent single-source scenario. Transfer learning is a broader topic to cover, thus surveys provide minimal coverage and comparison of the deep learning methods that have been designed for unsupervised domain adaptation [152, 180, 215, 231, 251, 272], or they focus on tasks such as activity recognition [45] or reinforcement learning [128, 234]. The goal of this survey is to discuss, highlight unique components, and compare approaches to single-source homogeneous unsupervised deep domain adaptation.

We first provide background on where domain adaptation fits into the more general problem of transfer learning. We follow this with an overview of generative adversarial networks (GANs) to provide background for the increasingly widespread use of adversarial techniques in domain adaptation. Next, we investigate the various domain adaptation methods, the components of those methods, and the results. Then, we overview domain adaptation theory and discuss what we can learn from the theoretical results. Finally, we look at application areas and identify future research directions for domain adaptation.

2 BACKGROUND

2.1 Transfer Learning

The focus of this survey is domain adaptation. Because domain adaptation can be viewed as a special case of transfer learning [182], we first review transfer learning to highlight the role of domain adaptation within this topic. Transfer learning is defined as the learning scenario where a model is trained on a source domain or task and evaluated on a different but related target domain or task, where either the tasks or domains (or both) differ [62, 81, 180, 251]. For instance, we may wish to learn a model on a handwritten digit dataset (e.g., MNIST [130]) with the goal of using it to recognize house numbers (e.g., SVHN [175]). Or, we may wish to learn a model on a synthetic,

cheap-to-generate traffic sign dataset [168] with the goal of using it to classify real traffic signs (e.g., GTSRB [223]). In these examples, the source dataset used to train the model is related but different from the target dataset used to test the model—both are digits and signs respectively, but each dataset looks significantly different. When the source and target differ but are related, then transfer learning can be applied to obtain higher accuracy on the target data.

2.1.1 Categorizing Methods. In a transfer learning survey paper, Pan et al. [180] defined two terms to help classify various transfer learning techniques: “domain” and “task.” A domain consists of a feature space and a marginal probability distribution (i.e., the features of the data and the distribution of those features in the dataset). A task consists of a label space and an objective predictive function (i.e., the set of labels and a predictive function that is learned from the training data). Thus, a transfer learning problem might be either transferring knowledge from a source domain to a different target domain or transferring knowledge from a source task to a different target task (or a combination of the two) [62, 180, 251].

By this definition, a change in domain may result from either a change in feature space or a change in the marginal probability distribution. When classifying documents using text mining, a change in the feature space may result from a change in language (e.g., English to Spanish), whereas a change in the marginal probability distribution may result from a change in document topics (e.g., computer science to English literature) [180]. Similarly, a change in task may result from either a change in the label space or a change in the objective predictive function. In the case of document classification, a change in the label space may result from a change in the number of classes (e.g., from a set of 10 topic labels to a set of 100 topic labels). Similarly, a change in the objective predictive function may result from a substantial change in the distribution of the labels (e.g., the source domain has 100 instances of class A and 10,000 of class B, whereas the target has 10,000 instances of A and 100 of B) [180].

To classify transfer learning algorithms based on whether the task or domain differs between source and target, Pan and Yang [180] introduced three terms: “inductive,” “transductive,” and “unsupervised” transfer learning. In inductive transfer learning, the target and source tasks are different, the domains may or may not differ, and some labeled target data are required. In transductive transfer learning, the tasks remain the same while the domains are different, and both labeled source data and unlabeled target data are required. Finally, in unsupervised transfer learning, the tasks differ as in the inductive case, but there is no requirement of labeled data in either the source domain or the target domain.

2.1.2 Domain Adaptation. One popular type of transfer learning is *domain adaptation*, which will be the focus of our survey. Domain adaptation is a type of transductive transfer learning. Here, the target task remains the same as the source, but the domain differs [56, 180, 182]. Homogeneous domain adaptation is the case where the domain feature space also remains the same, and heterogeneous domain adaptation is the case where the feature spaces differ [182].

In addition to the previous terminology, machine learning techniques are often categorized based on whether or not labeled training data are available. Supervised learning assumes labeled data are available, semi-supervised learning uses both labeled data and unlabeled data, and unsupervised learning uses only unlabeled data. However, domain adaptation assumes data comes from both a source domain and a target domain. Thus, prepending one of these three terms to “domain adaptation” is ambiguous since it may refer to labeled data being available in the source or target domains.

Authors apply these terms in various ways to domain adaptation [55, 111, 180, 203, 251]. In this article, we will refer to “unsupervised” domain adaptation as the case in which both labeled source data and unlabeled target data are available, “semi-supervised” domain adaptation as the case in

which labeled source data in addition to some labeled target data are available, and “supervised” domain adaptation as the case in which both labeled source and target data are available [12]. The distinction between these categories describes the target domain, but only describe situations in which labeled data are available for the source domain. These definitions are commonly used in the methods surveyed in this article as well as others [27, 74, 77, 147, 203, 225].

2.1.3 Related Problems. Multi-domain learning [62, 113] and multi-task learning [29] are related to transfer learning and domain adaptation. In contrast to transfer learning, the goal of these learning approaches is obtaining high performance on all specified domains (or tasks) rather than just on a single target domain (or task) [180, 260]. For example, often it is assumed that the training data are drawn in an independent and identically distributed (i.i.d.) fashion, which may not be the case [113]. One such example is the task of developing a spam filter for users who disagree on what is considered spam. If all the users’ data are combined, the training data will be drawn from multiple domains. While each individual domain may be i.i.d., the aggregated dataset may not be. If the data are split by user, then there may be too little data to learn a model for each user. Multi-domain learning can take advantage of the entire dataset to learn individual user preferences [62, 113]. Some researchers have developed adversarial strategies to tackle this multi-domain learning challenge [90, 212].

When working with multiple tasks, instead of training models separately for different tasks (e.g., one model for detecting shapes in an image and one model for detecting text in an image), multi-task learning will learn these separate but related tasks simultaneously so that they can mutually benefit from the training data of other tasks through a (partially) shared representation [29]. If there are both multiple tasks and domains, then these approaches can be combined into multi-domain multi-task learning, as is described by Yang and Hospedales [260].

Another related problem is domain generalization, in which a model is trained on multiple source domains with labeled data and then tested on a separate target domain that was not seen during training [173]. This contrasts with domain adaptation where target examples (possibly unlabeled) are available during training. Some approaches related to those surveyed in this article have been designed to address this situation. Examples include an adversarial method introduced by Zhao et al. [283] and an autoencoder approach by Ghifary et al. [76] discussed in Section 7.4.

2.2 Generative Adversarial Networks

Many deep domain adaptation methods that we will discuss in the next section incorporate adversarial training. We use the term *adversarial training* broadly to refer to any method that utilizes an adversary or an adversarial process during training. Before other adversarial methods were developed, the term was narrowly applied to training designed to improve the robustness of a model by utilizing adversarial examples, e.g., image inputs with small worst-case perturbations that lead to misclassification [83, 229]. Subsequently, other techniques have arisen that also utilize an adversary during training, including generative-adversarial training of generative adversarial networks (GANs) [82] and domain-adversarial training of domain adversarial neural networks (DANN) [74], both of which have been used for domain adaptation. To provide background for the domain adaptation methods utilizing these techniques, we will first discuss GANs and later when discussing DANN note the differences.

In recent years, there has been a large and growing interest in GANs. Pitting two well-matched neural networks against each other (hence “adversarial”), playing the roles of a data discriminator and a data generator, the pair is able to refine each player’s abilities in order to perform functions such as synthetic data generation. Goodfellow et al. [82] proposed this technique in 2014. Since that time, hundreds of papers have been published on the topic [92, 270]. GANs have traditionally

Fig. 1. Realistic but entirely synthetic images of human faces generated by a GAN trained on the CelebA-HQ dataset [116].

been applied to synthetic image generation, but recently researchers have been exploring other novel use cases such as domain adaptation.

GANs are a type of deep generative model [82]. For synthetic image generation, a training dataset of images must be available. Popular datasets include human faces (CelebA [146]), handwritten digits (MNIST [130]), bedrooms (LSUN [267]), and sets of other objects (CIFAR-10 [123] and ImageNet [57, 199]). After training, the generative model will be able to generate synthetic images that resemble those in the training data. For example, a generator trained with CelebA will generate images of human faces that look realistic but are not images of real people, as shown in Figure 1. To learn to do this, GANs utilize two neural networks competing against each other [82]. One network represents a generator. The generator accepts a noise vector as input, which contains random values drawn from some distribution such as normal or uniform. The goal of the generator network is to output a vector that is indistinguishable from the real training data. The other network represents a discriminator, which accepts as input either a real sample from the training data or a fake sample from the generator. The goal of the discriminator is to determine the probability that the input sample is real. During training, these two networks play a minimax game, where the generator tries to fool the discriminator and the discriminator attempts to not be fooled.

Using the notation from Goodfellow et al. [82], we define a value function $V(G, D)$ employed by the minimax game between the two networks:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))] \quad (1)$$

Here, $x \sim p_{data}(x)$ draws a sample from the real data distribution, $z \sim p_z(z)$ draws a sample from the input noise, $D(x; \theta_d)$ is the discriminator, and $G(z; \theta_g)$ is the generator. As shown in the equation, the goal is to find the parameters θ_d that maximize the log probability of correctly discriminating between real (x) and fake ($G(z)$) samples while at the same time finding the parameters θ_g that minimize the log probability of $1 - D(G(z))$. The term, $D(G(z))$, represents the probability that generated data $G(z)$ is real. If the discriminator correctly classifies a fake input, then $D(G(z)) = 0$. Equation (1) minimizes the quantity $1 - D(G(z))$. This occurs when $D(G(z)) = 1$, or when the discriminator misclassifies the generator's output as a real sample. Thus, the discriminator's mission is to learn to correctly classify the input as real or fake while the generator tries to fool the discriminator into thinking that its generated output is real. This process is illustrated in Figure 2.

Fig. 2. Illustration of the GAN generator G and discriminator D networks. The dashed line between the D networks indicates that they share weights (or are the same network). In the top row, a real image from the training data (horses $\leftrightarrow$ zebras dataset by Zhu et al. [289]) is fed to the discriminator, and the goal of D is to make $D(x) = 1$ (correctly classify as real). In the bottom row, a fake image from the generator is fed to the discriminator, and the goal of D is to make $D(G(z)) = 0$ (correctly classify as fake), which competes with the goal of G to make $D(G(z)) = 1$ (misclassify as real).

2.2.1 Training. In recent years, there have been impressive results from GANs. At the same time, this research faces some challenges since training a GAN can encounter problems such as difficulty converging [6, 80], mode collapse where the generator only learns to generate realistic samples for a few specialized modes of the data distribution [80], and vanishing gradients [82]. Many methods have been proposed to resolve these training challenges using a variety of tricks [82, 91, 179, 206, 218, 228], network architecture choices [116, 190, 206], objective modifications [5, 15, 67, 88, 112, 120, 161, 163, 165, 176–178, 282], mixtures or ensembles [6, 64, 79, 94, 117, 170, 181, 236, 269], maximum mean discrepancy (MMD) [16, 65, 135, 139, 227], making a connection to reinforcement learning [68, 186], or a combination of these modifications [91, 166, 268]. For an in-depth discussion of these techniques, there are a number of survey papers directed at GAN variants that include a discussion of training challenges and work [93, 100, 158]. These techniques can be employed in the domain adaptation methods that utilize GANs [20, 21, 41, 97, 143, 160, 207, 218, 244, 249]. While these training stability methods could similarly be applied to other adversarial domain adaptation approaches, they are not typically needed for the non-GAN methods surveyed here.

2.2.2 Evaluation. Once successfully trained, a GAN model can be difficult to evaluate and compare with other models. Multiple approaches and measures have been introduced to evaluate GAN performance. Often researchers have evaluated their models through visual inspection [209] such as performing user studies where participants mark which images they think look more realistic [206]. However, ideally a more automated metric could be found. Past generative models were evaluated by computing log-likelihood [235], but this is not necessarily tractable in GANs [80]. A proxy for log-likelihood is a Parzen window estimate, which was used for early GAN evaluation [82, 156, 177, 235], but in high dimensions (such as images), this could be far from the actual log-likelihood and not even rank models correctly [86, 235]. Thus, there has been much work proposing various evaluation methods for GANs: methods for detecting memorization [15, 61, 82, 156, 190, 235], determining diversity [7, 91, 179, 209], measuring realism [16, 91, 145, 206], and approximating log-likelihood [252]. Xu et al. [257] and Borji [19] survey and compare many of these GAN evaluation methods.

These techniques can be used for evaluating domain adaptation methods used for image translation (a form of image generation but conditioned on an input image) from one domain to another

[14, 41, 196, 263, 264, 289]. However, many domain adaptation methods (even those that are adversarial such as those using GANs) are not used for generation but rather for tasks with more easily defined loss functions, making these techniques largely not needed for adversarial domain adaptation methods. For example, accuracy [14, 21, 22, 41, 70, 74, 97, 143, 240] or AUC scores [189] can be used to evaluate classification, intersection over union or pixel accuracy can be used to evaluate image segmentation [14, 70, 97, 138, 185], and absolute difference can be used to evaluate regression [218].

3 METHODS

In recent years, numerous new unsupervised domain adaptation methods have been proposed, with a growing emphasis on neural network-based approaches. Distinct lines of research have emerged. These include aligning the source domain and target domain distributions, mapping between domains, separating normalization statistics, designing ensemble-based methods, or focusing on making the model target discriminative by moving the decision boundary into regions of lower data density. In addition, others have explored combinations of these approaches. We will describe each of these categories together with recent methods that fall into these categories.

In this survey, we will focus on homogeneous domain adaptation consisting of one source and one target domain, as is most commonly studied. Another case is multi-source domain adaptation, where there are multiple source domains but still only one target domain. Sun et al. [226] survey multi-source domain adaptation, and since then a number of other methods [28, 89, 96, 157, 184, 191, 255, 258, 280] have been developed for this case. It is also possible to perform multi-target domain adaptation [78], though this case is even more rarely studied. Similarly, we focus on homogeneous domain adaptation due to its prevalence, though some heterogeneous methods have been developed [63, 106, 137, 243, 262, 288].

3.1 Domain-Invariant Feature Learning

Most recent domain adaptation methods align source and target domains by creating a domain-invariant feature representation, typically in the form of a feature extractor neural network. A feature representation is domain-invariant if the features follow the same distribution regardless of whether the input data are from the source or target domain [279]. If a classifier can be trained to perform well on the source data using domain-invariant features, then the classifier may generalize well to the target domain since the features of the target data match those on which the classifier was trained. However, these methods assume that such a feature representation exists and the marginal label distributions do not differ significantly (Section 6).

The general training and testing setup of these methods is illustrated in Figure 3. Methods differ in how they align the domains (the Alignment Component in the figure). Some minimize divergence, some perform reconstruction, and some employ adversarial training. In addition, they differ in weight-sharing choices, which will be discussed in Section 4.3. We discuss the various alignment methods below.

3.1.1 Divergence. One method of aligning distributions is through minimizing a divergence that measures the distance between the distributions. Choices for the divergence measure include maximum mean discrepancy, correlation alignment, contrastive domain discrepancy, the Wasserstein metric, and a graph matching loss.

Maximum mean discrepancy (MMD) [84, 85] is a two-sample statistical test of the hypothesis that two distributions are equal based on observed samples from the two distributions. The test is computed from the difference between the mean values of a smooth function on the two domains' samples. If the means are different, then the samples are likely not from the same distribution.

Fig. 3. General network setup for domain adaptation methods learning domain-invariant features. (a) Methods differ in regard to how the domains are aligned during training (the Alignment Component) and whether the feature extractors used on each domain share none, some, or all of the weights between domains. (b) The target data are fed to the domain-invariant feature extractor and then to the task classifier.

The smooth functions chosen for MMD are unit balls in characteristic reproducing kernel Hilbert spaces (RKHS) since it can be proven that the population MMD is zero if and only if the two distributions are equal [85].

To use MMD for domain adaptation, the alignment component can be another classifier similar to the task classifier. MMD can then be computed and minimized between the outputs of these classifiers' corresponding layers (a slightly different setup than that in Figure 3). Rozantsev et al. [198] employ MMD, Long et al. [147] investigate a multiple kernel variant of MMD (MK-MMD), and later Long et al. [151] develop a joint MMD (JMMD) method. Bousmalis et al. [22] also tried MMD but found using an adversarial objective performed better in their experiments.

Correlation alignment (CORAL) [224] is similar to MMD with a polynomial kernel, computed from the distance between second-order statistics (covariances) of the source and target features. For domain adaptation, the alignment component consists of computing the CORAL loss between the two feature extractors' outputs (in order to minimize the distance). A variety of distances have been used: Sun et al. [225] use a squared matrix Frobenius norm in Deep CORAL, Zhang et al. [277] use a Euclidean distance in mapped correlation alignment (MCA), others have used log-Euclidean distances in LogCORAL [248] and Log D-CORAL [172], and Morerio et al. [171] use geodesic distances. Zhang et al. [278] generalize correlation alignment to possibly infinite-dimensional covariance matrices in RKHS. Chen et al. [34] align statistics beyond the first and second orders.

Contrastive domain discrepancy (CCD) [114] is based on MMD but looks at the conditional distributions in order to incorporate label information (unlike CORAL or ordinary MMD). When minimizing CCD, intra-class discrepancy is minimized while inter-class margin is maximized. This has the problem of requiring target labels though, so Kang et al. [114] propose contrastive adaptation networks (CAN) that minimize cross-entropy loss on the labeled target data while alternating between estimating labels for target samples (via clustering) with adapting the feature extractor with the now-computable CCD (using the clusters). This approach outperforms the other methods on the Office dataset as shown in Table 3.

A problem known as "optimal transport" was originally proposed for studying resource allocation such as finding an optimal way to move material from mines to factories [169, 192], but it can also be used to measure the distances between distributions. If the cost of moving each point is a norm (e.g., Euclidean), then the solution to a discrete optimal transport problem can be viewed as a distance: the Wasserstein distance [51] (also known as the earth mover's distance). To align feature and label distributions with this distance, Courty et al. [47] propose joint distribution optimal transport (JDOT). To incorporate this into a neural network, Damodaran et al. [51] propose DeepJDOT.

Another divergence measure arises from graph matching: the problem of finding an optimal correspondence between graphs [259]. A feature extractor's output on a batch of samples can be viewed as an undirected graph (in the form of an adjacency matrix), where similar samples in the batch are connected. Given the graph from a batch of source data fed through the feature extractor and similarly a graph from a batch of target data, then the cost of aligning these graphs can be used as a divergence, as proposed by Das et al. [52–54].

3.1.2 Reconstruction. Rather than minimizing a divergence, Ghifary et al. [77] and Bousmalis et al. [22] hypothesize that alignment can be accomplished by learning a representation that both classifies the labeled source domain data well and can be used to reconstruct either the target domain data (Ghifary et al.) or both the source and target domain data (Bousmalis et al.). The alignment component in these setups is a reconstruction network – the opposite of the feature extractor network – that takes the feature extractor output and recreates the feature extractor's input (in this case, an image). Ghifary et al. [77] propose deep reconstruction-classification networks (DRCN), using a pair-wise squared reconstruction loss. Bousmalis et al. [22] propose domain separation networks (DSN), using a scale-invariant mean squared error reconstruction loss.

3.1.3 Adversarial. Several varieties of feature-level adversarial domain adaptation methods have been introduced in the literature. In most, the alignment component consists of a domain classifier. In one article, this component is instead represented by a network learning an approximate Wasserstein distance, and in another article the component is a GAN.

A domain classifier is a classifier that outputs whether the feature representation was generated from source or target data. Recall that GANs include a discriminator that tries to accurately predict whether a sample is from the real data distribution or from the generator. In other words, the discriminator differentiates between two distributions, one real and one fake. A discriminator could similarly be designed to differentiate two distributions which instead represent a source distribution and a target distribution, as is done with a domain classifier. Note though that an adversarial domain classifier is used for adaptation, whereas a GAN is used for data generation. The domain classifier is trained to correctly classify the domain (source or target). In this scenario, the feature extractor is trained such that the domain classifier is unable to classify from which domain the feature representation originated. This is a type of zero-sum two-player game [279] as in a GAN (Section 2.2). Typically, these networks are adversarially trained by alternating between these two steps. The feature extractor can be trained to make the domain classifier perform poorly by negating the gradient from the domain classifier with a *gradient reversal layer* [73] when performing back propagation to update the feature extractor weights (e.g., in DANN [1, 73, 74] and VRADA [189]), maximally confusing the domain classifier (when it outputs a uniform distribution over binary labels [239]), or inverting the labels (in ADDA [240]). Because data distributions are often multi-modal, results may be improved by conditioning the domain classifier on a multilinear map of the feature representation and the task classifier predictions, which takes into account the multi-modal nature of the distributions [148].

Shen et al. [216] created WDGR, a modification of DANN, by replacing the domain classifier with a network that learns an approximate Wasserstein distance. This distance is then minimized between source and target domains, which they found to yield an improvement. This method is similar to the divergence methods except here the divergence is learned with a network rather than computed based on statistics (e.g., using mean in MMD or covariance in CORAL). This method outperforms the other methods on the Amazon review dataset as shown in Table 4.

Sankaranarayanan et al. [207] propose Generate to Adapt that uses a GAN as the alignment component. The feature extractor output is both fed to a classifier trained to predict the label (if the input is from the source domain) and also to a GAN trained to generate source-like images

Fig. 4. Synthetic vehicle driving image (left) adapted to look realistic (right) [97].

(regardless of if the input is source or target). For training stability, they use an AC-GAN [179]. They note one downside of using a GAN for adaptation is that it requires a large training dataset, but a common strategy is to use a pretrained network on a large dataset such as ImageNet. Using this pretraining, even on small datasets (e.g., Office) where the generated images are poor, the network still learns adaptation satisfactorily. Sankaranarayanan et al. [208] develop a similar approach for semantic segmentation.

3.2 Domain Mapping

An alternative to creating a domain-invariant feature representation is mapping from one domain to another. The mapping is typically created adversarially and at the pixel level (i.e., pixel-level adversarial domain adaptation), but not always, as discussed at the end of this section. This mapping can be accomplished with a conditional GAN. The generator performs adaptation at the pixel level by translating a source input image to an image that closely resembles the target distribution. For example, the GAN could change from a synthetic vehicle driving image to one that looks realistic as shown in Figure 4 [41, 97, 196, 264, 289]. A classifier can then be trained on the source data mapped to the target domain using the known source labels [218] or jointly trained with the GAN [21, 97]. We will first discuss how a conditional GAN works followed by the ways it can be employed for domain adaptation.

3.2.1 Conditional GAN for Image-to-Image Translation. The original formulation of a GAN was unconditional, where a GAN only accepted a noise vector as input. Conditional GANs, on the other hand, accept as input other information such as a class label, image, or other data [60, 75, 82, 164]. In the case of image generation, this means that a particular type of image to generate can be specified. One such example is to generate an image of a particular class within an image dataset such as “cat” rather than a random object from the dataset. Another example is conditioning on an input image such as in Figure 4, mapping an input driving image from one domain (synthetic) to an output image in another domain (realistic). Other uses include: transferring style (e.g., make a photo look like a Van Gogh painting) [118, 263, 289], colorizing images [109], generating satellite images from Google Maps data (or vice versa) [109, 263, 289], generating images of clothing from images of people wearing the clothing [264], generating cartoon faces from real faces [196, 230], converting labels to photos (e.g., semantic segmentation output to a photo) [109, 263, 289], learning disentangled representations [37], improving GAN training stability [179], and domain adaptation, which will be discussed in Section 3.2.2.

GANs conditioned on an input image can be used to perform image-to-image translation. These networks can be trained with varying levels of supervision: the dataset may contain corresponding images in the domains (supervised [109, 264]), only a few corresponding images (semi-supervised [72]), or no corresponding images (unsupervised [118, 263, 289]). A popular and general-purpose supervised method is pix2pix, developed by Isola et al. [109]. A commonly used unsupervised

method is CycleGAN [289], which is based on pix2pix, or methods similar to CycleGAN including DualGAN [263] and DiscoGAN [118].

Numerous modifications to these approaches have been proposed: one that is multi-modal is MUNIT, a multi-modal unsupervised image-to-image translator [105]. By assuming a decomposition into style (domain-specific) and content (domain-invariant) codes, MUNIT can generate diverse outputs for a given input image (e.g., multiple possible output images corresponding to the same input image). A modification to CycleGAN explored by Li et al. [136] uses separate batch normalization for each domain (an idea similar to AdaBN discussed in Section 3.3). Mejjati et al. [3] and Chen et al. [38] improve results with attention, learning which areas of the images on which to focus. Shang et al. [214] improve results by feeding the mapped images into a denoising autoencoder. While CycleGAN and similar approaches use two generators, one for each mapping direction, Benaim et al. [14] developed a method for one-sided mapping that maintains distances between pairs of samples when mapped from the source to the target domain rather than (or in addition to) using a cycle consistency loss, and Fu et al. [70] developed an alternative one-sided mapping using a geometric constraint (e.g., vertical flipping or 90 degree rotation). Royer et al. [196] propose XGAN, a dual adversarial autoencoder capable of handling large domain shifts, where possibly an image in the source domain may correspond to multiple images in the target domain or vice versa. They tested mapping human faces to cartoon faces, which was a shift larger than CycleGAN could adequately handle. Choi et al. [41] propose StarGAN, a method for handling multiple domains with a single GAN. Approaches like CycleGAN need a separate generator (or two, one for each direction) for each pair of domains, which is not a scalable solution to many domains. StarGAN, on the other hand, only needs a single generator. This has the added benefit of allowing the generator to learn using all the available data rather than only the data in a specific pair of domains. During training they randomly pick a target domain at each iteration so the generator learns to generate images in all the domains. Anoosheh et al. [4] propose an approach designed for the same purpose as StarGAN but using one generator per domain.

3.2.2 Image-to-Image Translation for Domain Adaptation. While the above approaches map images from one domain to another without the explicit purpose of performing domain adaptation, they can also be used for domain adaptation. For example, the original CycleGAN paper was application agnostic, but others have experimented with applying CycleGAN to domain adaptation [14, 70, 97]. It is important to note though that these image-to-image translation approaches assume that the domain differences are primarily low-level [20, 21, 240].

If unsupervised domain adaptation is performed for classification, adaptation can be accomplished by training an image-to-image translation GAN to map data from source to target, training a classifier on the mapped source images with known labels, and then subsequently testing by feeding unlabeled target through this target-domain classifier [20, 138, 218], as done in SimGAN [218] and illustrated in Figure 5(a). Alternatively, rather than learning a mapping from source to target, the opposite could be done: learn a mapping from target to source, train a classifier on the source images with known labels, and test by feeding target images to the image-to-image translation model (to make them look like source images) followed by the source-domain classifier [32], as illustrated in Figure 5(b).

In either of these approaches, if the mapping and the classification models are learned independently, the class assignments may not be preserved. For instance, class 1 may end up being “renamed” to class 2 after the mapping since the mapping was learned ignoring the class labels. This issue can be resolved by incorporating a semantic consistency loss (see Section 4.1) and training the mapping and classification models jointly [22, 97], as done in PixelDA [21].

Fig. 5. Two possible configurations using image-to-image translation for domain adaptation. The conditional GAN and classifier can be trained separately or jointly. Method 1 is the most common. Method 2 is used by one article. A combination of methods 1 and 2 is used in one article. The dashed lines between networks indicate that they share weights (or are the same network). NOTE: This figure does not illustrate the many variants of the conditional GAN component, which often train a generator in each direction (one source to target and one target to source) and use additional losses such as cycle consistency.

If there is a way to perform hyperparameter tuning, a third option is possible (combination of Figures 5(a) and 5(b)): train a target-domain classifier on the source-to-target GAN (for which the GAN is not used during testing) and a source-domain classifier on the target-to-source GAN (for which the GAN is used during testing). The algorithm may then output a linear combination of the prediction results from the two classifiers [200]. While this approach does improve results, it requires a method of hyperparameter training (see Section 4.7).

All of the above approaches perform pixel-level mapping. An alternative approach is to perform feature-level mapping. Hong et al. [99] use a conditional GAN to learn to make the source features look more like the target features (a distinctly different idea than making the features domain invariant, which was discussed in Section 3.1). They found this particularly helpful for structured domain adaptation (e.g., semantic segmentation, in their case).

Up to this point, these domain mapping methods have used image-to-image translation to map images (or in one case features) from one domain to another and thereby improve domain adaptation performance. Another line of research using pixel-level image generation for domain adaptation is to use a GAN to generate corresponding images in multiple domains and then employ all but the last layer of the discriminator as a feature extractor for a classifier [143, 160]. Liu and Tuzel [143] train a pair of GANs called CoGAN on two domains of images. Mao and Li [160] propose RegCGAN using only one generator and discriminator but including a domain label prepended to the input noise vector.

3.3 Normalization Statistics

Normalization layers such as batch norm [108] are used in most neural networks [210]. These have benefits including allowing for higher learning rates and thus faster training [108], reducing

initialization sensitivity [108], smoothing the optimization landscape and making the gradients more Lipschitz [210], and allowing for deeper networks to converge [81, 253]. Each batch norm layer normalizes its input to have zero mean and unit variance. At test time, running averages of the batch norm parameters can be used. Alternatives have been developed including instance norm allowing use in recurrent neural networks [10] and group norm removing the dependence on batch size [253]. However, none of these normalization techniques were developed with domain adaptation in mind. In the case of domain adaptation, the normalization statistics for each domain likely differ. Another line of domain adaptation research involves using per-domain batch normalization statistics.

Li et al. [141] assume that the neural net layer weights learn task knowledge and the batch norm statistics learn domain knowledge. If this is the case, then domain adaptation can be performed by modulating all the batch norm layers' statistics from the source to target domain, a technique they call AdaBN. This has the benefit of being simple, parameter free, and complementary to other adaptation methods.

Carlucci et al. [27] propose AutoDIAL, a generalization of AdaBN. In AdaBN, the target data are not used to learn the network weights but only for adjusting the batch norm statistics. AutoDIAL can utilize the target data for learning the network weights by coupling network parameters between source and target domains. They do this through adding domain alignment layers (DA-layers) that differ for source and target input data before each of the batch norm layers. Generally, batch norm computes a moving average of the statistics on a batch of the layer's input data. However, in AutoDIAL, source and target input data to DA-layers are mixed by a learnable amount before feeding this to batch norm (meaning that the batch norm statistics are now computed over some source and some target data rather than just source data or just target data). This allows the network to automatically learn how much alignment is needed at various points in the network.

3.4 Ensemble Methods

Given a base model such as a neural network or decision tree, an ensemble consisting of multiple models can often outperform a single model by averaging together the models' outputs (e.g., regression) or taking a vote (e.g., classification) [50, 81]. This is because if the models are diverse then each individual model will likely make different mistakes [81]. However, this performance gain corresponds with an increase in computation cost due to the large number of models to evaluate for each ensemble prediction, making ensembles common for some use cases such as competitions but uncommon when comparing models [81]. Despite the incurred cost, several ensemble-based methods have been developed for domain adaptation either using the ensemble predictions to guide learning or using the ensemble to measure prediction confidence for pseudo-labeling target data.

3.4.1 Self-Ensembling. An alternative to using multiple instances of a base model as the ensemble is using only a single model but "evaluating" (via a history or average) the models in the ensemble at multiple points in time during training—a technique called *self-ensembling*. This can be done by averaging over past predictions for each example (by recording previous predictions) [126] or past network weights (by maintaining a running average) [233]. Since an ensemble requires diverse models, these self-ensembling approaches require high stochasticity in the networks, which is provided by extensive data augmentation, varying the augmentation parameters, and including dropout. These methods were originally developed for semi-supervised learning.

French et al. [69] modify and extend these prior self-ensembling methods for unsupervised domain adaptation. They use two networks: a student network and a teacher network. Input images are fed first to stochastic data augmentation (Gaussian noise, translations, horizontal flips,

affine transforms, etc.) before being input to both networks. Because the method is stochastic, the augmented images fed to the networks will differ. The student network is trained with gradient descent while the teacher network weights are an exponential moving average (EMA) of the student network's weights. This method outperforms the other methods on the datasets in Table 2. Athiwaratkun et al. [9] show that in at least one experiment stochastic weight averaging [110] can further improve these results.

3.4.2 Pseudo-Labeling. Rather than voting or averaging the outputs of the models in an ensemble, the individual model predictions could be compared to determine the ensemble's confidence in that prediction. The more models in the ensemble that agree, the higher the ensemble's confidence in that prediction. In addition, if performing classification on a particular example, an individual model's confidence can be determined by looking at the last layer's softmax distribution: uniform indicates uncertainty whereas one class' probability much higher than the rest indicates higher confidence. Applying this to domain adaptation, a diverse ensemble trained on source data may be used to label target data. Then, if the ensemble is highly confident, those now-labeled target examples can be used to train a classifier for target data.

This is the method Saito et al. [203] developed called asymmetric tri-training (ATT). Two networks sharing a feature extractor are trained on the labeled source data (i.e., the ensemble in this case is of size two). Those two networks then predict the labels for the unlabeled target data, and if the two agree on the label and have high enough confidence on a particular instance, then the predicted label for that example is assumed to be the true label. After the target data are labeled by the first two networks, the third network (also sharing the same feature extractor) can be trained using the assumed-true labels (pseudo-labels). Diversity in the ensemble is handled with an additional loss (see Section 4.1).

Instead of using an ensemble, Zou et al. [290] rely on just the softmax distribution for the confidence measure. When working with semantic segmentation, they found relying on the prediction confidence for pseudo-labeling results in transferring primarily easy classes while ignoring harder classes. Thus, they additionally propose adding a class-wise weighting term when pseudo-labeling to normalize the class-wise confidence levels and thus balance out the class distribution.

3.5 Target Discriminative Methods

One assumption that has led to successes in semi-supervised learning algorithms is the *cluster assumption* [30]: that data points are distributed in separate clusters and the samples in each cluster have a common label [219]. If this is the case, then decision boundaries should lie in low density regions (i.e., should not pass through regions where there are many data points) [30]. A variety of domain adaptation methods have been explored to move decision boundaries into density regions of lower density. These have typically been trained adversarially.

Shu et al. [219] in virtual adversarial domain adaptation (VADA) and Kumar et al. [124] in co-regularized alignment (Co-DA) both use a combination of variational adversarial training (VAT) developed by Miyato et al. [167] and conditional entropy loss. They are used in combination because VAT without the entropy loss may result in overfitting to the unlabeled data points [124] and the entropy loss without VAT may result in the network not being locally-Lipschitz and thus not resulting in moving the decision boundary away from the data points [219]. Shu et al. [219] additionally propose a decision-boundary iterative refinement step with a teacher (DIRT-T) for use after training to further refine the decision boundaries on the target data, allowing for a slight improvement over VADA. An entropy loss was also used in AutoDIAL [27] but without VAT.

In generative adversarial guided learning (GAGL), Wei and Hsu [249] propose to let a GAN move decision boundaries into lower-density regions. Using domain alignment methods that learn

domain-invariant features like DANN (Section 3.1), typically the data fed to the feature extractor is either source or target data. However, Wei and Hsu [249] propose to alternate this with feeding generated (fake) images and appending a “fake” label to the task classifier, thus repurposing the task classifier as a GAN discriminator. They found this to have the effect of moving the decision boundaries in the target domain into areas of lower density with a GAN, promoting target-discriminative features as a result.

Saito et al. [204] propose adversarial dropout regularization. Since dropout is stochastic, when they create two instances of the task classifier containing dropout, the resulting networks may produce different predictions. The difference between these predictions can be viewed as a discriminator. Using this discriminator to adversarially train the feature extractor has the effect of producing target discriminative features. Lee et al. [134] alter adversarial dropout to better handle convolutional layers by dropping channel-wise rather than element-wise.

3.6 Combinations

In recent work, researchers have proposed various combinations of the above methods. Domain mapping has been combined with domain-invariant feature learning methods either trained separately (in GraspGAN [20]) or jointly (in CyCADA [97]). Following AdaBN, many researchers started employing domain-specific batch normalization [20, 69, 114, 124, 136]. Kumar et al. [124] propose co-regularized alignment (Co-DA), an approach in which two separate adversarial domain-invariant feature networks are learned with different feature spaces, drawing on ensemble-based methods. Kang et al. [115] combine domain mapping with aligning the models’ attention by minimizing an attention-based discrepancy. Deng et al. [59] combine target discriminative methods with self-ensembling. Lee et al. [132] combine target discriminative methods and domain-invariant feature learning with a sliced Wasserstein metric.

Multi-adversarial domain adaptation (MADA) [183] combines adversarial domain-invariant feature learning with ensemble methods for the purpose of better handling multi-modal data. This is accomplished by incorporating a separate discriminator for each class and using the task classifier’s softmax probability to weight the loss from each discriminator for unlabeled target samples.

Saito et al. [205] combine elements of adversarial domain-invariant feature learning, ensemble methods, and target discriminative features in their maximum classifier discrepancy (MCD) method. They propose using a shared feature extractor followed by an ensemble (of size two) of task-specific classifiers, where the discrepancy between predictions measures how far outside the support of the source domain the target samples lie. The discriminator in this setup is the combination of the two classifiers. The feature extractor is trained to minimize the discrepancy (i.e., fool the classifiers that the samples are from the source domain) while the classifiers are trained to maximize the discrepancy on the target samples.

4 COMPONENTS

Table 1 summarizes the neural network-based domain adaptation methods we discuss showing components each method uses including what type of adaptation, which loss functions, whether the method uses a generator, and which weights are shared. Below we discuss each of these aspects followed by how the networks are trained, what types of networks can be used, multi-level adaptation techniques, and how to tune the hyperparameters of these methods.

4.1 Losses

4.1.1 Distance. Distance functions play a variety of roles in domain adaptation losses. A distance loss can be used to align two distributions by minimizing a distance function (e.g., MMD) as explained in Section 3.1. If using an ensemble, minimizing a distance function can align the outputs

Table 1. Comparison of Different Neural Network Based Domain Adaptation Methods Based on Method of Adaptation (Domain-Invariant Feature Learning [DI], Domain Mapping [DM], Normalization [N], Ensemble [En], Target Discriminative [TD]), Various Loss Functions (Distance, Promoting Different Features, Cycle Consistency, Semantic Consistency, Task, Feature- or Pixel-Level Adversarial), Usage of a Generator, and Which Weights Are Shared (in the Feature Extractor)

Name	Method	Loss Functions					Adversarial Loss		Generator	Shared Weights
		Distance	Diff.	Cycle	Sem.	Task	Feature	Pixel		
CAN [114]	DI,N	CCD				✓				not BN
French et al. [69]	En,N	sq. diff.				✓				EMA
Co-DA [124] ^a	DI,En,N,TD	L1	✓			✓	✓			optional
VADA [219] ^a	DI,TD					✓	✓			✓
DeepJDOT [51]	DI	JDOT				✓				✓
CyCADA [97]	DI,DM			✓	✓	✓	✓	✓	✓	
Gen. to Adapt [207]	DI					✓	✓		✓	✓
SimNet [187]	DI	prototypes					✓			
MADA [183]	DI,En					✓	✓			✓
MCD [205]	DI,En,TD	✓	✓			✓	✓			
GAGL [249]	DI,TD					✓	✓	✓	✓	✓
SBADA-GAN [200] ^b	DM				✓	✓		✓	✓	
MCA [277]	DI	MCA				✓				✓
CCN ⁺⁺ [102]	DI	clusters					✓			✓
M-ADDA [127]	DI	clusters				✓	✓			
Rozant. et al. [198]	DI	MMD				✓				regularize
XGAN [196]	DM			✓			✓	✓	✓	some
StarGAN [41]	DM			✓				✓	✓	✓
PixelDA [21]	DM				✓	✓		✓	✓	✓
AutoDIAL [27]	N,TD					✓				not BN
AdaBN [145]	N									not BN
JAN-A [151]	DI	JMMD				✓	✓			✓
LogCORAL [248]	DI	logCOR, mean				✓				✓
Log D-CORAL [172]	DI	logDCOR				✓				✓
VRADA [189]	DI					✓	✓			✓
ATT [203]	En		✓			✓				✓
SimGAN [218]	DM							✓	✓	N/A ^c
ADDA [240]	DI					✓	✓			
CycleGAN [289]	DM			✓				✓	✓	^d
RegCGAN [160]	DM					✓	✓		✓	✓
Sener et al. [213]	DI	<i>k</i> -NN								✓
DSN [22]	DI		✓	✓		✓	✓			some
DRCN [77]	DI			✓		✓				✓
CoGAN [143]	DM					✓	✓		✓	some
Deep CORAL [225]	DI	CORAL				✓				✓
DANN [1, 73, 74]	DI					✓	✓			✓

(Continued)

Table 1. Continued

Name	Method	Loss Functions				Adversarial Loss			Generator	Shared Weights
		Distance	Diff.	Cycle	Sem.	Task	Feature	Pixel		
DAN [147]	DI	MK-MMD				✓				low
Tzeng et al. [239] ^e	DI					✓	✓			✓

^aAlso incorporate virtual adversarial training [167].

^bAlso a self-labeled classification loss (learn label on source images, pseudo-label mapped target to source).

^cMaps to target domain so only have feature extractor for target (part of the classifier).

^dUnspecified; originally not applied to domain adaptation, but later used for this [14, 70, 97].

^eSemi-supervised for some classes, i.e., requires some labeled target data for some of the classes.

of the ensemble's models: an L1 loss of the difference in predicted target class probabilities from two networks in Co-DA [124] or a squared difference between the predictions of the student and teacher networks in self-ensembling [69]. (Note the squared difference loss is confidence thresholded, i.e., if the max predicted output is below a certain threshold then the squared difference loss is set to zero.)

Some of the described methods have been altered replacing the task loss with one of similarity. Laradji and Babanezhad [127] propose M-ADDA, a metric-learning modification to ADDA but with the goal of maximizing the margin between clusters of data points' embeddings. Based on DANN, Pinheiro [187] proposes SimNet, classifying based on how close an embedding is to the embeddings of a random subset of source images for each class. Hsu et al. [102] propose CCN++ incorporating a pairwise similarity network (trained with the same class is similar and different classes are dissimilar).

4.1.2 Promote Differences. Methods that rely on multiple networks learning different features (such as to make an ensemble diverse) do so by promoting differences between the networks. Saito et al. [203] train the two classifiers labeling unlabeled data to use different features by adding a norm of the product of the two classifiers' weights. Bousmalis et al. [22] promote different features between two private feature extractors with a soft subspace orthogonality constraint, which is similarly used by Liu et al. [144] for text classification. Kumar et al. [124] train the feature extractors to be different by pushing minibatch means apart. Saito et al. [205] maximize the discrepancy between two classifiers using a fixed, shared feature extractor to promote using different features.

4.1.3 Cycle Consistency / Reconstruction. A cycle consistency loss or reconstruction loss is commonly used in domain mapping methods to avoid requiring a dataset of corresponding images to be available in both domains. This is how CycleGAN [289], DualGAN [263], and DiscoGAN [118] can be unsupervised. This means that after translating an image from one domain (e.g., horses) to another (e.g., zebras), the new image can be translated back to reconstruct the original image, as illustrated in Figure 6(a). Some variants of this have been proposed such as an L1 loss with a transformation function (e.g., identity, image derivatives, mean of color channels) [218], a feature-level cycle-consistency loss (mapping from source to embedding to target then back to embedding resulting in the same embeddings) [196], or using the loss in one [41] or both directions [97, 196]. Sener et al. [213] enforce cycle consistency in their k -nearest neighbors (k -NN) approach by requiring the distance between any source and target point labeled the same to be less than the distance between any source and target point labeled differently and derive a rule they can solve with stochastic gradient descent.

4.1.4 Semantic Consistency. A semantic consistency loss can be used to preserve class assignments as illustrated in Figure 6(b) (a segmentation example). The semantic consistency loss

Fig. 6. (a) Illustration of a cycle-consistency loss using the horses $\leftrightarrow$ zebras dataset by Zhu et al. [289]. The difference between the original source image and the reconstructed image (source to target and back to source) is minimized. (b) Example semantic segmentation situation in which the class names are swapped between the input image and the mapped image that would be prevented by including a semantic-consistency loss. The semantic-consistency loss requires that the class assignments are preserved.

requires that a classifier output (or semantic segmentation labeling) from the original source image is the same as the same classifier's output on the pixel-level mapped target output.

4.1.5 Task. Nearly all of the domain adaptation methods include some form of task loss that helps the network learn to perform the desired task. For example, for classification, the goal is to output the ground truth source label, or for semantic segmentation, to label each pixel with the correct ground truth source label. The task loss used is generally a cross-entropy loss, or more specifically the negative log likelihood of a softmax distribution [81] when using a softmax output layer. The exceptions not including a task loss are SimNet [187] that classify based on distance to prototypes of each class, the work by Sener et al. [213] that uses k nearest neighbors, and AdaBN [141] that only adjusts the batch norm layers to the target domain. In addition, the image-to-image translation methods are application agnostic unless trained jointly for domain adaptation.

4.1.6 Adversarial. A variety of methods use a discriminator (or critic) for learning domain-invariant features, realistic image generation, or promoting target discriminative features by forcing a network (either a feature extractor or generator) to produce outputs indistinguishable between two domains (source and target or real and fake). This loss is different than the other losses discussed in this section because this *adversarial loss* is learned [80, 109] (where learning is more than a hyperparameter search) rather than being provided as a predefined function. During training, gradients from the discriminator are used to train the feature extractor or generator (e.g., negated by a gradient reversal layer, Section 3.1.3). This alternates with updating the discriminator itself to make the correct domain classification.

4.1.7 Additions for Specific Problems. Some research focusing on specific problems has resulted in additional losses. For semantic segmentation, Li et al. [138] develop a loss making segmentation boundaries sharper to help when the mapped image-to-image translation images will be used for

segmentation, Chen et al. [40] develop a distillation loss in addition to performing location-aware alignment (e.g., “road” is usually at the bottom of each image), Hoffman et al. [98] develop a class-aware constrained multiple instance loss, Zhang et al. [275] develop a curriculum where after learning some high-level properties on easy tasks the segmentation network is forced to follow those properties (interpretations include student-teacher setup or posterior regularization), and Perone et al. [185] apply the self-ensembling method [69] replacing the cross-entropy loss with a consistency loss. For object detection, Chen et al. [39] use two domain classifiers (one on an image-level representation and the other on an instance-level representation) with a consistency regularization between them. For adaptation from synthetic images where it is known which pixels are foreground in the source images, Bousmalis et al. [21] and Bak et al. [11] mask certain losses to only penalize foreground pixel differences. For person re-identification, Wei et al. [250] include a person identity-keeping constraint in their domain mapping GAN.

4.2 Low-Confidence or Low-Relevance Rejection

Given a measure of confidence, performance may increase if we can reject data points for training the target classifier that are not of sufficient confidence. This, of course, assumes our confidence measurement is accurate enough. Saito et al. [203] used the label agreement of an ensemble combined with the softmax distribution output (uniform is not confident, one probability much higher than the rest is confident). Sener et al. [213] used the label agreement of the k nearest source data points. If the confidence is too low, then the example is rejected and not used in training until, when it is re-evaluated in subsequent iterations, it is determined to be sufficiently confident. Inoue et al. [107] used an object detector’s prediction probability as a measure of confidence, only using high-confidence detections for fine-tuning an object detection network. Similarly, a rejection approach could be used if we have a measure of relevance. For text classification, Zhang et al. [274] weighed examples by their relevance to their target aspect based on a small set of positive and negative keywords (a form of weak supervision).

4.3 Weight Sharing

Methods employ different amounts of sharing network weights between domains or regularizing the weights to be similar. Most methods completely share weights between the feature extractors used on the source and target domains (as shown in Table 1). However, some techniques do not. Since deep networks consist of many layers, allowing them to represent hierarchical features, Long et al. [147] propose copying the lower layers from a network trained on the source domain and adapting higher layers to the target domain with MK-MMD since higher layers do not transfer well between domains. In CoGAN, Liu et al. [143] share the first few layers of the generators and the last few layers of the discriminators, making the assumption that the domains share high-level representations. In AdaBN, Li et al. [141] assume domain knowledge is stored in the batch norm statistics, so they share all weights except for the batch norm statistics. French et al. [69] define the teacher network as an exponential moving average of the student network’s weights (a type of ensemble). Instead of sharing weights, Rozantsev et al. [197, 198] propose two variants: regularizing weights to be similar but not penalizing linear transformations and transforming the weights from the source network to the target network with small residual networks. Bousmalis et al. [22] propose domain separation networks (DSN): learning source-specific, target-specific, and shared features where the “shared” source domain encoder and “shared” target domain encoder do share weights, but the “private” source domain encoder and “private” target domain encoders do not. Others have similarly explored this idea of shared vs. specific features [25, 144, 193].

4.4 Training Stages

Some have trained networks for domain adaptation in stages. Tzeng et al. [240] train a source classifier first followed by adaptation. Taigman et al. [230] use a pre-trained encoder during adaptation. Bousmalis et al. [20] in GraspGAN first train the domain-mapping network followed by the domain-adversarial network. Hoffman et al. [97] in CyCADA train their many components in stages because it would not all fit into GPU memory at once.

Other methods train the domain adaptation networks jointly, which using an adversarial approach is done by alternating between training the discriminator and the rest of the networks (Sections 2.2 and 3.1.3). However, variations exist for some other methods. Saito et al. [203] in ATT cycle through training the source networks, generating pseudo-labels, and training the target network. Zou et al. [290] alternate between pseudo-labeling the target data and re-training the model using the labels (a form of self-training). Wei et al. [249] in GAGL alternate between feeding in real source and target data and the fake images generated by a GAN. Sener et al. [213] alternate between k -nearest neighbors and performing gradient descent.

4.5 Multi-Level

Some adaptation methods perform adaptation at more than one level. As discussed in Section 3.6, GraspGAN [20] and CyCADA [97] perform pixel-level adaptation with domain mapping and feature-level adaptation with domain-invariant feature learning. Hoffman et al. [97] found that performing both levels of adaptation significantly improves accuracy: using domain mapping to capture low-level image domain shifts and learning domain-invariant features to handle larger domain shifts than what pure domain mapping methods can support. Following this idea, Tsai et al. [238] make semantic segmentation predictions and perform domain-invariant feature learning at multiple levels in their semantic segmentation network, and Zhang et al. [273] perform domain-invariant feature learning at multiple levels while automatically learning how much to align to each level. Chen et al. [39] perform domain-invariant feature learning at both image and instance levels for object detection but also include a consistency regularization between the two domain classifiers.

4.6 Types of Networks

Nearly all of the surveyed approaches focus on learning from image data and use convolutional neural networks (CNNs) such as ResNet-50 or Inception (Table 3). Wang et al. [246] explore the use of attention networks, Kang et al. [115] a combination of CNNs and attention, Ma et al. [154] graph convolutional networks, and Kurmi et al. [125] Bayesian neural networks. In the case of time-series data, Purushotham et al. [189] propose instead using a variational recurrent neural network (RNN) [43] or LSTM (a type of RNN) [95] rather than a CNN. The RNN learns the temporal relationships while adversarial training is used to achieve domain adaptation. For text classification (a type of natural language processing), Liu et al. [144] also use LSTMs while Zhang et al. [274] found a CNN to work just as well as RNNs or bi-LSTMs in their experiments. For relation extraction (another type of natural language processing), Fu et al. [71] also use a CNN. For time-series speech recognition, Zhao et al. [281] use bi-LSTMs while Hosseini-Asl et al. [101] used a combination of CNNs and RNNs. In the related problem of domain generalization, a combination of CNNs and RNNs have been used for handling a radio spectrogram changing through time to identify sleep stages [283].

4.7 Hyperparameter Tuning

Normal supervised learning-based hyperparameter tuning methods do not carry over to unsupervised domain adaptation [22, 74, 149, 150, 171, 185, 244]. A common supervised learning approach

is to split the training data into a smaller training set and a validation set. After repeatedly altering the hyperparameters, retraining the model, and testing on this validation set for each set of hyperparameters, the model yielding the highest validation set accuracy is selected. Another option is cross-validation. However, in unsupervised domain adaptation, there are now two domains, and the data for the target domain may not include any labels. When evaluating domain adaptation approaches on common datasets, generally the target data does contain labels, so work by some groups [22, 27, 124, 200, 219, 244, 249] do use some labeled target data (or all of it [149, 216]) for hyperparameter tuning, which can be interpreted as an upper bound on how well the method could perform [244]. For example, some [27, 150] tuned for Office on one W labeled example per class on the $A \rightarrow W$ task, while others [200, 249] tuned with a validation set of 1000 randomly sampled target examples. Using any labeled target data is not ideal because real-world testing will not include labels for tuning (unless it is semi-supervised, in which case semi-supervised learning is recommended in Section 6).

One tuning method not requiring labeled target data is *reverse validation* [74], which is a variant of *reverse cross validation* [285]. For a set of hyperparameters, the *reverse validation risk* can be estimated by first splitting source (labeled) and target (unlabeled) data into training and validation sets. Then, the labeled source and unlabeled target data are used to learn a classifier (as is normally done). Next, this forward classifier is used to label the target data and a new reverse classifier is learned (with the same algorithm) using the pseudo-labeled target data (as “source”) and unlabeled source data (as “target”, i.e., ignoring the known labels). This reverse classifier is evaluated on the source validation data to measure the reverse validation risk. Ganin et al. [74] found this method works better if the reverse classifier is initialized with the weights of the forward classifier and if using early stopping on the source validation set and a pseudo-labeled target validation set. Finally, hyperparameters are selected (e.g., grid search, random search, Bayesian optimization, or other gradient-free optimization methods) that minimize this reverse validation risk.

Alternatively, given some domain knowledge, one may devise relevant measures of similarity between the domains and tune parameters to increase the similarity. For example, French et al. [69] were able to improve performance on the challenging problem of MNIST $\rightarrow$ SVHN by tuning data augmentation hyperparameters for MNIST to match pixel intensities apparent in the SVHN dataset. By doing this, they were able to improve the state-of-the-art to 97.0% (Table 2).

5 RESULTS

Tables 2 through 5 summarize the results of evaluating many of these methods on datasets used for image classification as well as sentiment analysis. Care must be taken in the extent to which conclusions are drawn from comparing published numbers in different papers since the provided accuracies are for different network architectures, hyperparameters, amount of data augmentation, random initializations (or averages over a number of them), etc. and the methods may perform differently in other application areas. However, interestingly, at least one method in each of the categories surveyed gives promising results on at least one of the datasets.

With domain-invariant feature learning with the contrastive domain discrepancy, CAN [114] has the highest performance on the Office dataset (Table 3). By using adversarial domain-invariant feature learning, WDGRl generally outperforms the other methods on the Amazon review dataset (Table 4) and Generate to Adapt is second highest of the methods evaluated on the Office dataset. By using adversarial pixel-level domain mapping, SBADA-GAN [200] obtains the highest accuracy on MNIST $\rightarrow$ MNIST-M (Table 2). AutoDIAL [27], a normalization statistics method, does on-par with CAN and Generate to Adapt in two of Office adaptation tasks. The self-ensembling method by French et al. [69] outperforms all other methods on the datasets in Table 2, and Co-DA [124] comes close using an ensemble (of size two) of adversarial domain-invariant feature networks.

Table 2. Classification Accuracy (Source $\rightarrow$ Target, Mean $\pm$ std %) of Different Neural Network Based Domain Adaptation Methods on Various Computer Vision Datasets (Only Including Those Used in > 2 Articles)

Name	MNIST and USPS		MNIST and SVHN		MNIST[-M]	Synthetic to Real	
	MN $\rightarrow$ US	US $\rightarrow$ MN	SV $\rightarrow$ MN	MN $\rightarrow$ SV	MN $\rightarrow$ MN-M	SYN _N $\rightarrow$ SV	SYN _S $\rightarrow$ GTSRB
Target only (i.e., if we had the target labels)	96.3 $\pm$ 0.1 [97] 96.5 [21]	99.2 $\pm$ 0.1 [97]	99.2 $\pm$ 0.1 [97] 99.5 [22] 99.51 [73]		96.4 [21] 98.7 [22] 98.91 [73]	92.44 [73] 92.4 [22]	99.87 [73] 99.8 [22]
French et al. [69]	98.2	99.5	99.3	37.5 97.0 ^a		97.1	99.4
Co-DA [124] ^{b*}			98.6	81.7	97.5	96.0	
DIRT-T [219] ^{b*}			99.4	76.5	98.7	96.2	99.6
VADA [219] ^{b*}			94.5	73.3	95.7	94.9	99.2
DeepJDOT [51]	95.7	96.4	96.7		92.4		
CyCADA [97]*	95.6 $\pm$ 0.2	96.5 $\pm$ 0.1	90.4 $\pm$ 0.4				
Gen. to Adapt [207]*	92.8 $\pm$ 0.9	90.8 $\pm$ 1.3	92.4 $\pm$ 0.9				
SimNet [187]*	96.4	95.6			90.5		
MCD [205]*	96.5 $\pm$ 0.3	94.1 $\pm$ 0.3	96.2 $\pm$ 0.4				94.4 $\pm$ 0.3
GAGL [249] ^{b*}			96.7	74.6	94.9	93.1	97.6
SBADA-GAN [200] ^{b*}	97.6	95.0	76.1	61.1	99.4		96.7
MCA [277]			96.6		96.8	89.0	
CCN ⁺⁺ [102]*			89.1				
M-ADDA [127]*	98	97					
Rozantsev et al. [198]	60.7	67.3					
PixelDA [21]*	95.9				98.2		
ATT [203]			85.0	52.8	94.0	92.9	96.2
ADDA [240]*	89.4 $\pm$ 0.2	90.1 $\pm$ 0.8	76.0 $\pm$ 1.8				
RegCGAN [160]*	93.1 $\pm$ 0.7	89.5 $\pm$ 0.9					
DTN [230]*			84.4				
Sener et al. [213]			78.8	40.3	86.7		
DSN [22] ^{b*}	91.3 [21]		82.7		83.2	91.2	93.1
DRCN [77]	91.80 $\pm$ 0.09	73.67 $\pm$ 0.04	81.97 $\pm$ 0.16	40.05 $\pm$ 0.07			
CoGAN [143]*	91.2 $\pm$ 0.8	89.1 $\pm$ 0.8			62.0 [21]		
DANN [73, 74]*	85.1 [21]		71.07 70.7 [22] 71.1 [203] 73.6 [97]	35.7 [203]	81.49 77.4 [22] 81.5 [203]	90.48 90.3 [22, 203]	88.66 88.7 [203] 92.9 [22]
DAN [147]	81.1 [21]		71.1 [22]		76.9 [22]	88.0 [22]	91.1 [22]
Source only (i.e., no adaptation)	78.9 [21] 82.2 $\pm$ 0.8 [97]	69.6 $\pm$ 3.8 [97]	59.19 [73] 59.2 [22] 67.1 $\pm$ 0.6 [97]		56.6 [22] 57.49 [73] 63.6 [21]	86.65 [73] 86.7 [22]	74.00 [73] 85.1 [22]

Adversarial approaches denoted by *.

^aProblem-specific hyperparameter tuning of data augmentation to match pixel intensities of target domain images.

^bHyperparameter tuned on some labeled target data.

CyCADA increases accuracy from 54% to 82% for a synthetic season adaptation dataset [97] by combining both adversarial domain-invariant feature learning and domain mapping.

A number of these promising methods use adversarial techniques, which may be a key ingredient in solving domain adaptation problems. Adversarial approaches may be helpful on certain

Table 3. Classification Accuracy (Source $\rightarrow$ Target, Mean $\pm$ std %) of Different Neural Network Based Domain Adaptation Methods on the Office Computer Vision Dataset

Name	Office (Amazon, DSLR, Webcam)					
	A $\rightarrow$ W	D $\rightarrow$ W	W $\rightarrow$ D	A $\rightarrow$ D	D $\rightarrow$ A	W $\rightarrow$ A
CAN [114] ^a	94.5 $\pm$ 0.3	99.1 $\pm$ 0.2	99.8 $\pm$ 0.2	95.0 $\pm$ 0.3	78.0 $\pm$ 0.3	77.0 $\pm$ 0.3
Gen. to Adapt [207] ^{a*}	89.5 $\pm$ 0.5	97.9 $\pm$ 0.3	99.8 $\pm$ 0.4	87.7 $\pm$ 0.5	72.8 $\pm$ 0.3	71.4 $\pm$ 0.4
SimNet [187] ^{a*}	88.6 $\pm$ 0.5	98.2 $\pm$ 0.2	99.7 $\pm$ 0.2	85.3 $\pm$ 0.3	73.4 $\pm$ 0.8	71.8 $\pm$ 0.6
MADA [183] ^{a*}	90.0 $\pm$ 0.1	97.4 $\pm$ 0.1	99.6 $\pm$ 0.1	87.8 $\pm$ 0.2	70.3 $\pm$ 0.3	66.4 $\pm$ 0.3
AutoDIAL [27] ^{b,c}	84.2	97.9	99.9	82.3	64.6	64.2
CCN ⁺⁺ [102] ^{d*}	78.2	97.4	98.6	73.5	62.8	60.6
Rozantsev et al. [198]	76.0	96.7	99.6			
AdaBN [145] ^b	74.2	95.7	99.8	73.1	59.8	57.4
JAN-A [151] ^{a*}	86.0 $\pm$ 0.4	96.7 $\pm$ 0.3	99.7 $\pm$ 0.1	85.1 $\pm$ 0.4	69.2 $\pm$ 0.4	70.7 $\pm$ 0.5
LogCORAL [248]	70.2 $\pm$ 0.6	95.5 $\pm$ 0.1	99.5 $\pm$ 0.3	69.4 $\pm$ 0.5	51.2 $\pm$ 0.3	51.6 $\pm$ 0.5
Log D-CORAL [172]	68.5	95.3	98.7	62.0	40.6	40.6
ADDA [240] ^{a*}	75.1	97.0	99.6			
Sener et al. [213]	81.1	96.4	99.2	84.1	58.3	63.8
DRCN [77]	68.7 $\pm$ 0.3	96.4 $\pm$ 0.3	99.0 $\pm$ 0.2	66.8 $\pm$ 0.5	56.0 $\pm$ 0.5	54.9 $\pm$ 0.5
Deep CORAL [225]	66.4 $\pm$ 0.4	95.7 $\pm$ 0.3	99.2 $\pm$ 0.1	66.8 $\pm$ 0.6	52.8 $\pm$ 0.2	51.5 $\pm$ 0.3
DANN [73, 74] [*]	67.3 $\pm$ 1.7	94.0 $\pm$ 0.8	93.7 $\pm$ 1.0			
	72.6 $\pm$ 0.3 [77]	96.4 $\pm$ 0.1 [77]	99.2 $\pm$ 0.3 [77]	67.1 $\pm$ 0.3 [77]	54.5 $\pm$ 0.4 [77]	52.7 $\pm$ 0.2 [77]
	73.0 [198, 240]	96.4 [198, 240]	99.2 [198, 240]			
DAN [147]	68.5 $\pm$ 0.4	96.0 $\pm$ 0.3	99.0 $\pm$ 0.2			
	63.8 $\pm$ 0.4 [225]	94.6 $\pm$ 0.5 [225]	98.6 [198]	67.0 $\pm$ 0.4	54.0 $\pm$ 0.4	53.1 $\pm$ 0.3
	64.5 [198]	95.2 [198]	98.8 $\pm$ 0.6 [225]	65.8 $\pm$ 0.4 [225]	52.8 $\pm$ 0.4 [225]	51.9 $\pm$ 0.5 [225]
	68.5 [240]	96.0 [240]	99.0 [240]			
Tzeng et al. [239] ^{e*}	59.3 $\pm$ 0.6	90.0 $\pm$ 0.2	97.5 $\pm$ 0.1	68.0 $\pm$ 0.5	43.1 $\pm$ 0.2	40.5 $\pm$ 0.2
Source only (i.e., no adaptation)	62.6 [240] ^a	96.1 [240] ^a	98.6 [240] ^a			

Adversarial approaches denoted by *.

^aWith ResNet-50 network.

^bWith Inception-based network.

^cHyperparameter tuned on one W labeled example per class on $A \rightarrow W$ task (see [150]).

^dWith ResNet-18 network.

^eSemi-supervised for some classes, but evaluated on 16 hold-out categories for which the labels were not seen during training.

datasets (e.g., WDGRL on the Amazon review dataset on Office), certain types of data (e.g., VRADA was developed for time series data rather than image data), or may not require as extensive of tuning (e.g., Co-DA on MNIST $\rightarrow$ SVHN). Or adversarial training may be an additional tool to incorporate into existing non-adversarial methods. For instance, promising non-adversarial methods such as AutoDIAL and by French et al. [69] could be combined with adversarial methods (see Section 8.3). In fact, Long et al. [151] develop both JAN and then the adversarial version JAN-A, and JAN-A on average outperformed JAN on the Office dataset. CAN [114], which presently is the highest on the Office dataset, might also be improved by incorporating an adversarial component to it as in Long et al. [151].

Interestingly, French et al. by far outperform all other methods on MNIST $\rightarrow$ SVHN, though this requires a problem-specific data augmentation and hyperparameter tuning. This may indicate

Table 4. Classification Accuracy Comparison for Domain Adaptation Methods for Sentiment Analysis (Positive or Negative Review) on the Amazon Review Dataset [18]^a with Domains Books (B), DVD (D), Electronics (E), and Kitchen (K)

Source → Target	DANN [74] ^{b*}	DANN [74] ^{c*}	CORAL [224] ^d	ATT [203] ^c	WDGRL [216] ^{c e*}	No Adapt. [224] ^f
B → D	82.9	78.4		80.7	83.1	
B → E	80.4	73.3	76.3	79.8	83.3	74.7
B → K	84.3	77.9		82.5	85.5	
D → B	82.5	72.3	78.3	73.2	80.7	76.9
D → E	80.9	75.4		77.0	83.6	
D → K	84.9	78.3		82.5	86.2	
E → B	77.4	71.3		73.2	77.2	
E → D	78.1	73.8		72.9	78.3	
E → K	88.1	85.4	83.6	86.9	88.2	82.8
K → B	71.8	70.9		72.5	77.2	
K → D	78.9	74.0	73.9	74.9	79.9	72.2
K → E	85.6	84.3		84.6	86.3	

Adversarial approaches denoted by *.

^a<http://www.cs.jhu.edu/~mdredze/datasets/sentiment/>.

^bUsing 30,000-dimensional feature vectors from marginalized stacked denoising autoencoders (mSDA) by Chen et al. [36], which is an unsupervised method of learning a feature representation from the training data.

^cUsing 5000-dimensional unigram and bigram feature vectors.

^dUsing bag-of-words feature vectors including only the top 400 words, but suggest using deep text features in future work.

^eThe best results on target data for various hyperparameters.

^fUsing bag-of-words feature vectors.

that for some problems, maybe in particular the more challenging domain adaptation problems, hyperparameter tuning for a specific dataset may be of utmost importance. Possibly if other domain adaptation methods similarly were tuned appropriately, they would also experience large improvements. This is an area of research requiring further work (see Section 8.2). However, Co-DA [124] is not far behind on SVHN→MNIST and MNIST→MNIST-M and is the closest on MNIST→SVHN, achieving 81.7% compared with 97.0%. A great advantage of Co-DA is that it does not require highly problem-specific tuning on MNIST→SVHN as required by French et al. (without, they only achieved 37.5%). Possibly some components of Co-DA such as the adversarial domain adaptation or virtual adversarial training may be partially responsible for the decrease in hyperparameter sensitivity.

6 THEORY

Having surveyed domain adaptation methods, we now address the question of when adaptation may be beneficial. Ben-David et al. [13] develop a theory answering this in terms of an ideal predictor on both domains, Zhao et al. [279] further this theory by removing the dependence on a joint ideal predictor while focusing on domain-invariant feature learning methods, and Le et al. [129] develop theory looking beyond domain-invariant methods. These theoretical results can help answer two questions: (1) when will a classifier (or other predictor) trained on the source data perform well on the target data and (2) given a small number of labeled target examples, how can they best be used during training to minimize target test error?

Answering the first question, labeled source data and unlabeled target data are both required (unsupervised). Answering the second question, additionally some labeled target data are required (semi-supervised). We will first review the theoretical bounds followed by a discussion of what

Table 5. List and Description of Computer Vision Datasets from Tables 2 and 3

Computer Vision Datasets used for Domain Adaptation	
MNIST [130] ^a	This is a binary (mostly black and white, but actually grayscale due to anti-aliasing) handwritten digit dataset (digits 0–9), which stands for “modified NIST.” It is based on the National Institute of Standards and Technology’s (NIST) Special Databases 1 and 3, one of which was easier than the other, so MNIST is a combination of the two that are size normalized to fit in a 20x20 box preserving the aspect ratio and centered in a 28x28 pixel image.
MNIST-M [74] ^b	This is a modification of MNIST where the digits are blended with random patches from BSDS500 dataset color photos.
USPS [131] ^c	This is another handwritten digit dataset (digits 0–9). It consists of handwritten zipcodes scanned and segmented by the U.S. Postal Service (USPS). They were size normalized to 16x16 pixels preserving the aspect ratio. The values are normalized to be between –1 and 1.
SVHN [175] ^d	The Streetview House Numbers (SVHN) consists of single digits extracted from images of urban house numbers in Google Street View. The digits have been size normalized to 32x32 pixels.
SYN_N [74] ^b	Ganin et al. [74] used Microsoft Windows fonts to create a synthetic digit dataset (“Syn Numbers”) consisting of 1–3 digit numbers with various positions, orientation, background color, stroke color, and amount of blur.
SYN_S [168] ^e	This is a synthetic sign dataset created from modifications to Wikipedia pictograms of traffic signs. It consists of 100,000 images and 43 classes of signs.
GTSRB [223] ^f	The German Traffic Signs Recognition Benchmark (GTSRB) is a dataset created from video taken driving around Germany. It consists of about 50,000 images and 43 classes of signs.
Office [201] ^g	This dataset consists of 31 classes of objects in three different domains: Amazon (taken from its online website; medium resolution and studio lighting), DSLR (taken with a digital SLR camera; high resolution and in a real-world environment), and Webcam (taken with a 640x480 computer webcam; have noise, artifacts, and white balance issues). Note: due to Office’s small size, some networks [74, 198, 225] were pre-trained on ImageNet.

^a <http://yann.lecun.com/exdb/mnist/>.^b See Ganin’s website <http://yaroslav.ganin.net/> for links to download.^c This can be found on various sites and some Github repositories. One such place: <https://web.stanford.edu/~hastie/ElemStatLearn/data.html>.^d <http://ufldl.stanford.edu/housenumbers>.^e The synthetic dataset linked to on: <http://graphics.cs.msu.ru/en/research/projects/imagerecognition/trafficsign>.^f <http://benchmark.ini.rub.de/?section=gtsrb&subsection=dataset>.^g <http://ai.bu.edu/adaptation.html>.

insights these bounds provide into answering the above two questions. Ben-David et al. [13] also address the case of multiple source domains, as do Mansour et al. [159]. In this article, we have focused on the cases containing only one source and one target (as is common in the methods we survey).

6.1 Unsupervised

6.1.1 Shared Hypothesis Space. Ben-David et al. [13] propose setting a bound on the target error based on the source error and the divergence between the source and target domains. The

empirical source error is easy to obtain by first training and then testing a classifier. However, the divergence between the domains cannot be directly obtained with standard methods like Kullback-Leibler divergence due to only having a finite number of samples from the domains and not assuming any particular distribution. Thus, an alternative is to measure it using a classifier-induced divergence called $\mathcal{H}\Delta\mathcal{H}$ -divergence. Estimates of this divergence with finite samples converges to the real $\mathcal{H}\Delta\mathcal{H}$ -divergence. This divergence can be estimated by measuring the error when getting a classifier to discriminate between the unlabeled source and target examples; though, it is often intractable to find the theoretically required divergence upper bound. Using the empirical source error $\hat{\epsilon}_S(h)$, the $\mathcal{H}\Delta\mathcal{H}$ -divergence between source and target samples $d_{\mathcal{H}\Delta\mathcal{H}}(\hat{\mathcal{D}}_S, \hat{\mathcal{D}}_T)$, and ideal predictor error λ^* using the optimal hypothesis for the source and target, the target error $\epsilon_T(h)$ can be bounded as shown in Equation (2) (using the form given by Zhao et al. [279]), $\forall h \in \mathcal{H}$ with probability at least $1 - \delta$ for $\delta \in (0, 1)$.

$$\epsilon_T(h) \leq \hat{\epsilon}_S(h) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\hat{\mathcal{D}}_S, \hat{\mathcal{D}}_T) + \lambda^* + O\left(\sqrt{\frac{d \log n + \log(\frac{1}{\delta})}{n}}\right) \quad (2)$$

Zhao et al. [279] develop another upper bound that removes the reliance on λ^* . Let $\mathcal{H} \subseteq [0, 1]^X$, $\tilde{\mathcal{H}} := \{\text{sgn}(|h(x) - h'(x)| - t) \mid h, h' \in \mathcal{H}, 0 \leq t \leq 1\}$, $\langle \mathcal{D}_S, f_S \rangle$ and $\langle \mathcal{D}_T, f_T \rangle$ be the source and target domains (the true distributions, not empirical). The target error can then be bounded by the source error $\epsilon_S(h)$, the discrepancy between marginal distributions $d_{\tilde{\mathcal{H}}}(\mathcal{D}_S, \mathcal{D}_T)$, and the distance between the optimal source and target labeling functions $\forall h \in \mathcal{H}$, as shown in Equation (3).

$$\epsilon_T(h) \leq \epsilon_S(h) + d_{\tilde{\mathcal{H}}}(\mathcal{D}_S, \mathcal{D}_T) + \min\{\mathbb{E}_{\mathcal{D}_S}[|f_S - f_T|], \mathbb{E}_{\mathcal{D}_T}[|f_S - f_T|]\} \quad (3)$$

Zhao et al. [279] also develop an information-theoretic lower bound for target error. Let the labeling function $Y = f(X) \in \{0, 1\}$, the prediction function $\hat{Y} = h(g(X)) \in \{0, 1\}$, and Z be the intermediate representation output by a shared feature extractor used on source and target domain data. If the Jensen-Shannon distance $d_{JS}(\mathcal{D}_S^Y, \mathcal{D}_T^Y) \geq d_{JS}(\mathcal{D}_S^Z, \mathcal{D}_T^Z)$ and the Markov chain $X \xrightarrow{g} Z \xrightarrow{h} \hat{Y}$ holds, then Equation (4) provides a lower bound on the source and target error.

$$\epsilon_S(h \circ g) + \epsilon_T(h \circ g) \geq \frac{1}{2} \left(d_{JS}(\mathcal{D}_S^Y, \mathcal{D}_T^Y) - d_{JS}(\mathcal{D}_S^Z, \mathcal{D}_T^Z) \right)^2 \quad (4)$$

6.1.2 Different Hypothesis Spaces. Le et al. [129] develop an upper bound that allows for different hypothesis spaces for source and target functions, possibly non-deterministic labeling, and any bounded or continuous loss. If l is a bounded or continuous loss, $x \sim \mathbb{P}^s$ (source) and $x \sim \mathbb{P}^t$ (target), $T : X^s \rightarrow X^t$ and $K := T^{-1}$ (bijective mapping), $R(\theta) = \mathbb{E}_{p(x,y)}[l(y, h_\theta(x))]$ for θ parameterizing a hypothesis set $\mathcal{H} = \{h_\theta \mid \theta \in \Theta\}$, $\Delta R(h^s, h^t) := |R^t(h^t) - R^s(h^s)|$, $y \in \{-1, 1\}$, M is the number of labels, $\mathbb{P}^\# := K_\# \mathbb{P}^t$ is the pushforward probability distribution transporting $\mathbb{P}^t$ via K , $\Delta p(y|x) := p^t(y|T(x)) - p^s(y|x)$ for the true source and target labeling functions $p^s(y|x)$ and $p^t(y|x)$, where $WS_c(\mathbb{P}^s, \mathbb{P}^\#)$ denotes the Wasserstein-1 distance between the source and target distributions with a cost function $c(x, x') = 1_{x \neq x'}$ (1 if $x \neq x'$, otherwise 0), then Equation (5) provides an upper bound for the variance between a general loss on the source and target predictions.

$$\Delta R(h^s, h^t) \leq M \left(WS_c(\mathbb{P}^s, \mathbb{P}^\#) + \min\{\mathbb{E}_{\mathbb{P}^\#}[\|\Delta p(y|x)\|_1], \mathbb{E}_{\mathbb{P}^s}[\|\Delta p(y|x)\|_1]\} \right) \quad (5)$$

6.2 Semi-Supervised

In the semi-supervised case, a linear combination of the source and target errors is computed [13], called the α -error. A bound can be calculated on the true α -error based on the empirical α -error. Finding the minimum α -error depends on the empirical α -error, the divergence between source

and target, and the number of labeled source and target examples. Experimentation can be used to empirically determine the values of α that will perform well. Ben-David et al. [13] also demonstrate the process on sentiment classification, illustrating that the optimum uses non-trivial values.

The bound is given in Equation (6). If S is a labeled sample of size m with $(1 - \beta)m$ points drawn from the source distribution and βm from the target distribution, then with at least probability $1 - \delta$ for $\delta \in (0, 1)$:

$$\begin{aligned} \epsilon_T(\hat{h}) \leq & \epsilon_T(h_T^*) + 4\sqrt{\frac{\alpha^2}{\beta} + \frac{(1 - \alpha)^2}{1 - \beta}} \sqrt{\frac{2d \log(2(m + 1)) + 2 \log(\frac{8}{\delta})}{m}} \\ & + 2(1 - \alpha) \left(\frac{1}{2} \hat{d}_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{U}_S, \mathcal{U}_T) + 4\sqrt{\frac{2d \log(2m') + \log(\frac{8}{\delta})}{m'}} + \lambda \right) \end{aligned} \quad (6)$$

Here, $\hat{h} \in \mathcal{H}$ is the empirical minimizer of the α -error on S given by $\hat{\epsilon}_\alpha(h) = \alpha \hat{\epsilon}_T(h) + (1 - \alpha) \hat{\epsilon}_S(h)$ and $h_T^* = \min_{h \in \mathcal{H}} \epsilon_T(h)$ is the target error minimizer.

The optimum α is then:

$$\alpha^*(m_T, m_S; D) = \begin{cases} 1 & m_T \geq D^2 \\ \min\{1, v\} & m_T \leq D^2 \end{cases} \quad (7)$$

Here, $m_S = (1 - \beta)m$ is the number of source examples, $m_T = \beta m$ is the number of target examples, $D = \sqrt{d}/A$, and

$$v = \frac{m_T}{m_T + m_S} \left(1 + \frac{m_S}{\sqrt{D^2(m_S + m_T) - m_S m_T}} \right) \quad (8)$$

$$A = \frac{1}{2} \hat{d}_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{U}_S, \mathcal{U}_T) + 4\sqrt{\frac{2d \log(2m') + \log(\frac{4}{\delta})}{m'}} + \lambda \quad (9)$$

$$B = 4\sqrt{\frac{2d \log(2(m + 1)) + 2 \log(\frac{8}{\delta})}{m}} \quad (10)$$

6.3 Discussion

6.3.1 Unsupervised. Equation (2) indicates that if the optimal predictor error λ^* on both source and target data is large, then there is no good hypothesis from training on the source domain that will work well on the target domain [13, 279]. However, as is more common in the application of domain adaptation, if λ^* is small, then the bound depends on the source error and the $\mathcal{H}\Delta\mathcal{H}$ -divergence [13]. The domain-invariant feature learning methods discussed in Section 3.1 try minimizing these two terms [279]: the source error via a task loss on labeled source data and divergence via a divergence measure such as MMD, with reconstruction, or adversarially. While Section 5 shows that on many datasets these methods work, there is no guarantee that such adaptation will increase performance (these are upper bounds), as shown by simple counterexamples [279]. It may actually decrease performance if the marginal label distributions differ significantly between source and target [279].

Equation (3) shows that the target error upper bound alternatively involves the marginal distributions and Equation (4) shows that the lower bound does too. These indicate the importance of aligning the label distributions. If the marginal label distributions are significantly different, then minimizing the source error and divergence between feature representations will actually increase the error [279]. Thus over-training domain-invariant feature learning methods can increase

target error, and Zhao et al. [279] experimentally verified this. They found on MNIST, USPS, and SVHN adaptation that during training the target accuracy would initially rise rapidly but would eventually decrease again despite increasing source accuracy, an effect even more apparent with larger differences in the marginal label distributions. It is an open problem as to when the label distributions can be aligned without target labels [279].

6.3.2 Semi-Supervised. Equation (6) indicates that when only source or target data are available, that data should be used (as we might expect). If the source and target are the same, then $\alpha^* = \beta$, which implies a uniform weighting of examples. Given enough target data, source data should not be used at all because it might increase the test-time error. Furthermore, without enough source data using it may also not be worthwhile, i.e., $\alpha^* \approx 0$ [13]. In this article, we focus on unsupervised domain adaptation, but these are important considerations if target labels can be obtained. For example, this shows that it may be better to perform semi-supervised adaptation if some labeled target examples are available rather than using the labeled target examples to hyperparameter tune an unsupervised adaptation method.

7 APPLICATIONS

Domain adaptation has been applied in a variety of areas including computer vision, natural language processing, and for time-series data. Using domain adaptation in these various problems can save the human time that would be spent labeling the target data. In some cases, such as image semantic segmentation, providing ground truth is very labor intensive. Each pixel-level annotated image in the Cityscapes dataset took on average 1.5 hours to complete [46]. In addition, similar methods as described in this article have been applied to the related problem of domain generalization and some other problems as well.

7.1 Computer Vision

Most of the methods surveyed in this article are for computer vision tasks such as adapting a model trained on synthetic images to real photos (e.g., from synthetic numbers or signs, Table 2), stock photos to real photos (e.g., Amazon to DSLR on the Office dataset, Table 3), or simple to complex images (e.g., MNIST to SVHN, Table 2). Others have been used in robotics for robot grasping [20], autonomous navigation [265], and lifelong learning [254], for semantic segmentation [40, 99, 103, 133, 153, 208, 238, 242, 290] including when additional information is available from a simulator [133], in a medical context for chest X-ray segmentation [32], 3D CT scans to X-ray segmentation [276], MRI to CT scan segmentation [33], and MRI segmentation [185], in low resource situations (where there are very few target data points) [101], in situations with different label sets for each domain [222], for object detection [39, 98, 107], for person re-identification [11, 58, 74, 142, 250, 286, 287], and for depth estimation [8, 155, 174].

7.2 Natural Language Processing

Domain adaptation has been used in natural language processing such as for sentiment analysis (Table 4, [274, 281]), other text classification [144, 274] including weakly supervised aspect-transfer from one aspect of a dataset to another [274], relation extraction [71], semi-supervised sequence labeling [55], semi-supervised question answering [261], sentence specificity [119], and neural machine translation [23, 31, 42].

7.3 Time Series

For time-series data, domain adaptation has been used for learning temporal latent relationships in health data across different population age groups [189], to perform speech recognition [101,

[217, 281], for predicting driving maneuvers [237], anomaly detection [241], and inertial tracking [35]. In a method addressing the related problem of domain generalization, time-series radio data was used for sleep-stage classification [283]. Finally, a combination of pre-training and fine-tuning was used to solve another transfer learning problem, where the source datasets have a different label space than the target dataset [66].

7.4 Domain Generalization

Domain-invariant feature learning approaches similar to those discussed in Section 3.1 have been used for the related problem of domain generalization, where there are multiple source domains and an unseen target domain [17, 173]. Zhao et al. [283] use an adversarial approach with a domain classifier to learn a model on a dataset collected from a number of people sleeping in various environments that will generalize well to new people and/or new environments (e.g., sleeping in a different room). Ghifary et al. [76] use a reconstruction approach with a denoising autoencoder to improve object recognition generalizability, where the “noise” is different views (domains) of the data (e.g., rotation, change in size, or variation in lighting) and the autoencoder tries to reconstruct corresponding views of the object in other domains. Carlucci et al. [28] propose an adversarial approach combining domain adaptation and generalization while also doing domain mapping. Akuzawa et al. [2] note the domain-invariance objective may compete with the discriminative objective and thus develop a method to find the most domain-invariant representation that does not hurt classification performance. Li et al. [140] note that previous domain-invariant methods typically assume balanced classes and develop a method to handle changes in class proportions.

7.5 Other Problems

Adversarial losses like those used in adversarial domain adaptation methods have also been applied in multiple other settings. Wang et al. [247] created an adversarial spacial dropout network to add occlusions to images to improve the accuracy of object detection algorithms. They also created an adversarial spatial transformer network to add deformations such as rotations to objects to again increase object detection accuracy. Pinto et al. [188] used adversarial agents to improve a robot’s ability to grasp an object via self-supervised learning by employing both shaking and snatching adversaries. Giu et al. [87] used an adversarial loss to predict and demonstrate (i.e., robot will copy) human motion. Rippel et al. [194, 195] used a reconstruction and adversarial loss with an autoencoder for learning higher quality image compression at low bit rates. Sinclair [221] applied adversarial loss to clone a physical model for real-time sound synthesis. Adversarial techniques may also be applied to machine learning security, where the goal is to train a classifier robust to adversarial examples [104, 167].

8 RESEARCH DIRECTIONS

As we have seen, the rapidly growing body of research focused on unsupervised deep domain adaptation now encompasses many novel methods and components. Here we look at what could be explored in future research to further enhance this existing work.

8.1 Bi-Directional Adaptation

The more difficult domain adaptation problems are far from being solved. Tables 2 through 5 indicate that some domain adaptation problems are harder than others and point to the challenge that more work needs to be focused on these harder problems. While accuracy for SVHN→MNIST ranges from 70.7% to 99.3%, for the reverse case of MNIST→SVHN, the highest without highly-problem-specific hyperparameter tuning is 81.7% by Kumar et al. [124] (though tuned on a small amount of labeled target data). This indicates how this reverse problem is much harder [69, 73]. As

a result, few papers offer results for this direction. French et al. [69] were able to vastly improve performance up to 97.0%; however, this required developing a problem-specific unsupervised hyperparameter tuning method. Other methods may similarly benefit from such tuning. Continued work is needed to strengthen general-purpose bi-directional adaptation.

8.2 Hyperparameter Tuning

Some methods such as reverse validation and a problem-specific pixel intensity matching have been applied to hyperparameter tuning without requiring target labels (Section 4.7). While the reverse validation method appears promising, it was not used in most of the methods surveyed (only [74, 183, 187]). This may be because of the increase in computation cost [185] or problems with the reverse validation accuracy not aligning with test accuracy [22]. It is also possible researchers may just be unaware of the method since in the surveyed papers few mention the idea (only [22, 74, 183, 185, 187]). Problem-specific methods such as matching pixel intensity between domains as done by French et al. [69] are possible given some domain knowledge, but hyperparameter tuning methodologies should be developed that will work across a wider range of problems. This remains an open area of research.

8.3 Combining Promising Methods

French et al. [69], Co-DA [124], CAN [114], AutoDIAL [27], Generate to Adapt [207], and WD-GRL [216] are promising approaches based on Tables 2 through 4. French et al. uses a student and teacher network for self-ensembling, Co-DA trains multiple (e.g., two) adaptation networks while requiring diversity and agreement in addition to incorporating virtual adversarial training, CAN alternates between clustering and adaptation through minimizing intra-class discrepancy and maximizing inter-class margin, AutoDIAL adjusts batch normalization layer weights, Generate to Adapt uses an embedding-conditional GAN for adversarial domain adaptation, and WD-GRL performs adversarial domain adaptation similar to DANN by using a domain classifier. These are largely independent ideas that if combined may result in additional performance gains.

For instance, the student network in French et al. that accepts either a source or target augmented image could be replaced by the AutoDIAL network to learn how much adaptation to perform at each level of the network. Or to combine with adversarial methods, the student and teacher networks' outputs (or an intermediate layer's outputs, as is being explored by Wang et al. [244]) could be fed to a gradient reversal layer followed by a domain classifier, in effect adding an adversarial loss term to the existing two terms used by French et al. Or since French et al. is based upon data augmentation, one might try replacing the existing stochastic data augmentation with a GAN since a GAN can be used for data augmentation (given enough unlabeled training data).

Alternatively, key aspects of other methods could be incorporated. While domain adaptation methods commonly align feature distributions, a different line of research aligns the joint or conditional distribution of the feature and label spaces instead [44, 47, 51, 148, 151, 154, 232, 266]. Researchers found aligning in this manner improves results when handling multi-modal data distributions [148] or when label proportions differ between domains [47]. Other domain adaptation strategies may similarly benefit from aligning the joint or conditional distribution rather than merely the feature distribution.

8.4 Balancing Classes

In order to obtain high accuracy on the challenging problem of MNIST→SVHN, French et al. [69] include an additional class-balance term in their loss function, which both improved training stability and helped the network avoid a degenerate local minimum. Though, this term was not

required in their other experiments. Clearly, class balancing is an important concern; although, this depends on the dataset being used. Other methods may similarly benefit from balancing classes.

For instance, Hoffman et al. [97] note that the frequency-weighted intersection over union results in their article were very close to the target-only model accuracy (an approximate upper bound). Thus, they conclude that domain mapping followed by domain-invariant feature learning is very effective for the common classes in the SYNTHIA dataset (season adaptation on a synthetic driving dataset). It is possible then that additional balancing of classes could help the not-as-common classes to perform better. In addition, data augmentation through occluding parts of the images may improve class balancing as would the adversarial spatial dropout network by Wang et al. [247] since the two best classes (road and sky) were likely in almost every image.

8.5 Incorporating Improved Image-to-Image Translation Methods

Bousmalis et al. [21] with PixelDA had difficulty applying their method with large domain differences. However, other image-to-image translation methods like XGAN [196] have been developed that may support larger domain shifts. These methods could be extended to domain adaptation directly or also incorporating a semantic consistency loss (as explained in Section 4.1). This may allow for more substantial differences between domains. Similarly, image-to-image translation methods like StarGAN [41] have been developed for multiple domains, which could be extended for multi-domain adaptation.

8.6 Further Experimental Comparison between Methods

As shown in Table 2, French et al. [69] outperforms all the other methods and Co-DA [124] is quite close behind (with the advantage that it does not require highly-problem-specific tuning on MNIST→SVHN). In Table 3, CAN [114] outperforms the others followed by Generate to Adapt [207]. Finally, in Table 4, WDGRL [216] generally performs the best. However, these methods are not all compared on the same dataset, making a direct comparison difficult. Additional experiments must be performed to see how these methods compare. Similarly, other promising approaches may outperform other methods on some datasets, which could be determined through additional experiments.

These comparisons can be made easier through developing a unified implementation of these various methods. Schneider et al. [211] are developing such an open-source set of implementations of state-of-the-art domain adaptation (and domain generalization) methods. The results provided in individual papers have different hyperparameters, data augmentation, network architectures, etc. that can make direct comparisons challenging. Using a unified implementation of these methods can facilitate more clearly understanding what aspects of a method are responsible for performance gains and also support combining the novel elements from multiple methods.

8.7 Limitations of Datasets

Varying amounts of source and target data are available in different situations. The datasets used for comparisons (the image datasets listed in Table 5 and the Amazon review dataset) are relatively small when compared with the sizes of datasets commonly in use in deep learning, e.g., ImageNet [57, 199] (though ImageNet is often used to pretrain adaptation networks). For example, Sankaranarayanan et al. [207] note how GANs require a lot of training data. This may limit GAN-based methods from being used on too small of source or target datasets. Modifications may need to be developed for such low resource situations, an area explored by Hosseini-Asl et al. [101]. Additionally, most domain adaptation datasets are for computer vision. To spur research in other application areas, other datasets could be created.

8.8 Other Applications

Other application areas may benefit from performing domain adaptation as have those discussed in Section 7. In particular, only a few methods were applied to time-series data. One time-series application that may benefit from adaptation is activity prediction, e.g., adapting from one type of sensor to another or from one person's data to another's. Some added challenges in this context may be the large differences in feature spaces due to the wide variety of sensors used (e.g., an event stream of fixed motion sensors turning on and off in a smart home vs. sampled motion and location data collected from smart phones or watches) or the difference in labels (e.g., one model may learn a "walk" activity while another learns "exercise" or may learn "read" while another model learns "school"). Applying domain adaptation in new areas may yield novel methods or components applicable in other areas as well.

8.9 Other Domain Adaptation Cases

As mentioned in Section 3, we have surveyed single-source homogeneous unsupervised domain adaptation methods due to this being the most commonly studied case of domain adaptation. However, exploring other cases is warranted. By utilizing data from multiple source domains and/or multiple target domains, additional gains in performance may be achievable. By handling heterogeneous feature spaces or various other levels of supervision (e.g., semi-supervised learning [202] or weakly-supervised learning [220]), domain adaptation may bring performance gains to other problems as well. Finally, another under-studied case of domain adaptation is partial domain adaptation, where the target domain contains only a subset of the source domain's labels [26, 232, 271].

9 CONCLUSIONS

For supervised learning, deep neural networks are in prevalent use, but these networks require large labeled datasets for training. Unsupervised domain adaptation can be used to adapt deep networks to possibly smaller datasets that may not even have target labels. Several categories of methods have been developed for this goal: domain-invariant feature learning, domain mapping, normalization statistics-based, and ensemble-based methods. These various methods have some unique and common elements as we have discussed. Additionally, theoretical results provide some insight into empirical observations. Some methods appear very promising, but further research is required for direct comparisons, novel method combinations, improved bi-directional adaptation, and use for novel datasets and applications.

REFERENCES

- [1] Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, and Mario Marchand. 2014. Domain-adversarial neural networks. *arXiv preprint arXiv:1412.4446* (2014).
- [2] Kei Akuzawa, Yusuke Iwasawa, and Yutaka Matsuo. 2019. Adversarial invariant feature learning with accuracy constraint for domain generalization. *arXiv preprint arXiv:1904.12543* (2019).
- [3] Youssef Alami Mejjati, Christian Richardt, James Tompkin, Darren Cosker, and Kwang In Kim. 2018. Unsupervised attention-guided image-to-image translation. In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 3693–3703. <http://papers.nips.cc/paper/7627-unsupervised-attention-guided-image-to-image-translation.pdf>.
- [4] Asha Anosheh, Eirikur Agustsson, Radu Timofte, and Luc Van Gool. 2018. ComboGAN: Unrestrained scalability for image domain translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- [5] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Doina Precup and Yee Whye Teh (Eds.), Vol. 70. PMLR, 214–223. <http://proceedings.mlr.press/v70/arjovsky17a.html>.
- [6] Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, and Yi Zhang. 2017. Generalization and equilibrium in generative adversarial nets (GANs). In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of*

- Machine Learning Research*), Doina Precup and Yee Whye Teh (Eds.), Vol. 70. PMLR, 224–232. <http://proceedings.mlr.press/v70/arora17a.html>.
- [7] Sanjeev Arora, Andrej Risteski, and Yi Zhang. 2018. Do GANs learn the distribution? Some theory and empirics. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=BJehNFW0->.
 - [8] Amir Atapour-Abarghouei and Toby P. Breckon. 2018. Real-time monocular depth estimation using synthetic data with domain adaptation via image style transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [9] Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. 2019. There are many consistent explanations of unlabeled data: Why you should average. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=rkgKBhA5Y7>.
 - [10] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450* (2016).
 - [11] Slawomir Bak, Peter Carr, and Jean-Francois Lalonde. 2018. Domain adaptation through synthesis for unsupervised person re-identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
 - [12] Oscar Beijbom. 2012. Domain adaptations for computer vision applications. *arXiv preprint arXiv:1211.4860* (2012).
 - [13] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. 2010. A theory of learning from different domains. *Machine Learning* 79, 1 (01 May 2010), 151–175. DOI: <https://doi.org/10.1007/s10994-009-5152-4>
 - [14] Sagie Benaïm and Lior Wolf. 2017. One-sided unsupervised domain mapping. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 752–762. <http://papers.nips.cc/paper/6677-one-sided-unsupervised-domain-mapping.pdf>.
 - [15] David Berthelot, Tom Schumm, and Luke Metz. 2017. Began: Boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717* (2017).
 - [16] Mikołaj Bińkowski, Dougal J. Sutherland, Michael Arbel, and Arthur Gretton. 2017. Demystifying MMD GANs. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=r1lUOzWCW>.
 - [17] Gilles Blanchard, Gyemin Lee, and Clayton Scott. 2011. Generalizing from several related classification tasks to a new unlabeled sample. In *Advances in Neural Information Processing Systems 24*, J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 2178–2186. <http://papers.nips.cc/paper/4312-generalizing-from-several-related-classification-tasks-to-a-new-unlabeled-sample.pdf>.
 - [18] John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, Bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. 440–447.
 - [19] Ali Borji. 2018. Pros and cons of GAN evaluation measures. *arXiv preprint arXiv:1802.03446* (2018).
 - [20] Konstantinos Bousmalis, Alex Irpan, Paul Wohlhart, Yunfei Bai, Matthew Kelcey, Mrinal Kalakrishnan, Laura Downs, Julian Ibarz, Peter Pastor, Kurt Konolige, Sergey Levine, and Vincent Vanhoucke. 2018. Using simulation and domain adaptation to improve efficiency of deep robotic grasping. In *Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA)*. 4243–4250. DOI: <https://doi.org/10.1109/ICRA.2018.8460875>
 - [21] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. 2017. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [22] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. 2016. Domain separation networks. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 343–351. <http://papers.nips.cc/paper/6254-domain-separation-networks.pdf>.
 - [23] Denny Britz, Quoc Le, and Reid Pryzant. 2017. Effective domain mixing for neural machine translation. In *Proceedings of the 2nd Conference on Machine Translation*. 118–126.
 - [24] Lars Bungum and Björn Gambäck. 2011. A survey of domain adaptation in machine translation: Towards a refinement of domain space. In *Proceedings of the India-Norway Workshop on Web Concepts and Technologies*, Vol. 112.
 - [25] Jinming Cao, Oren Katzir, Peng Jiang, Dani Lischinski, Danny Cohen-Or, Changhe Tu, and Yangyan Li. 2018. Disentangled synthesis for domain adaptation. *arXiv preprint arXiv:1805.08019* (2018).
 - [26] Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Michael I. Jordan. 2018. Partial transfer learning with selective adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [27] Fabio Maria Carlucci, Lorenzo Porzi, Barbara Caputo, Elisa Ricci, and Samuel Rota Bulò. 2017. AutoDIAL: Automatic domain alignment layers. In *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE, 5077–5085.

- [28] Fabio M. Carlucci, Paolo Russo, Tatiana Tommasi, and Barbara Caputo. 2018. Agnostic domain generalization. *arXiv preprint arXiv:1808.01102* (2018).
- [29] Rich Caruana. 1997. Multitask learning. *Machine Learning* 28, 1 (01 Jul 1997), 41–75. DOI : <https://doi.org/10.1023/A:1007379606734>
- [30] Olivier Chapelle and Alexander Zien. 2005. Semi-supervised classification by low density separation. In *AISTATS*, Vol. 2005. Citeseer, 57–64.
- [31] Boxing Chen, Colin Cherry, George Foster, and Samuel Larkin. 2017. Cost weighting for neural machine translation domain adaptation. In *Proceedings of the 1st Workshop on Neural Machine Translation*. 40–46.
- [32] Cheng Chen, Qi Dou, Hao Chen, and Pheng-Ann Heng. 2018. Semantic-aware generative adversarial nets for unsupervised domain adaptation in chest X-ray segmentation. In *Machine Learning in Medical Imaging*, Yinghuan Shi, Heung-Il Suk, and Mingxia Liu (Eds.). Springer International Publishing, Cham, 143–151.
- [33] Cheng Chen, Qi Dou, Hao Chen, Jing Qin, and Pheng-Ann Heng. 2019. Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation. *arXiv preprint arXiv:1901.08211* (2019).
- [34] Chao Chen, Zhihang Fu, Zhihong Chen, Sheng Jin, Zhaowei Cheng, Xinyu Jin, and Xian-Sheng Hua. 2019. HoMM: Higher-order moment matching for unsupervised domain adaptation. *arXiv preprint arXiv:1912.11976* (2019).
- [35] Changhao Chen, Yishu Miao, Chris Xiaoxuan Lu, Linhai Xie, Phil Blunsom, Andrew Markham, and Niki Trigoni. 2019. Motion Transformer: Transferring neural inertial tracking between domains. In *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (Jul. 2019), 8009–8016. DOI : <https://doi.org/10.1609/aaai.v33i01.33018009>
- [36] Minmin Chen, Zhixiang Xu, Kilian Weinberger, and Fei Sha. 2012. Marginalized denoising autoencoders for domain adaptation. In *Proceedings of the 29th International Conference on Machine Learning (ICML '12)*, John Langford and Joelle Pineau (Eds.). Omnipress, New York, NY., 767–774.
- [37] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems* 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 2172–2180. <http://papers.nips.cc/paper/6399-infoGAN-interpretable-representation-learning-by-information-maximizing-generative-adversarial-nets.pdf>.
- [38] Xinyuan Chen, Chang Xu, Xiaokang Yang, and Dacheng Tao. 2018. Attention-GAN for object transfiguration in wild images. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [39] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. 2018. Domain adaptive faster R-CNN for object detection in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [40] Yuhua Chen, Wen Li, and Luc Van Gool. 2018. ROAD: Reality oriented adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [41] Yunje Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [42] Chenhui Chu and Rui Wang. 2018. A survey of domain adaptation for neural machine translation. In *Proceedings of the 27th International Conference on Computational Linguistics*. 1304–1319.
- [43] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C. Courville, and Yoshua Bengio. 2015. A recurrent latent variable model for sequential data. In *Advances in Neural Information Processing Systems* 28, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Eds.). Curran Associates, Inc., 2980–2988. <http://papers.nips.cc/paper/5653-a-recurrent-latent-variable-model-for-sequential-data.pdf>.
- [44] Safa Cicek and Stefano Soatto. 2019. Unsupervised domain adaptation via regularized conditional alignment. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [45] Diane Cook, Kyle D. Feuz, and Narayanan C. Krishnan. 2013. Transfer learning for activity recognition: A survey. *Knowledge and Information Systems* 36, 3 (01 Sep 2013), 537–556. DOI : <https://doi.org/10.1007/s10115-013-0665-3>
- [46] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [47] Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. 2017. Joint distribution optimal transportation for domain adaptation. In *Advances in Neural Information Processing Systems* 30, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 3730–3739. <http://papers.nips.cc/paper/6963-joint-distribution-optimal-transportation-for-domain-adaptation.pdf>.
- [48] Gabriela Csurka. 2017. A comprehensive survey on domain adaptation for visual applications. In *Domain Adaptation in Computer Vision Applications*. Springer, 1–35.
- [49] Gabriela Csurka. 2017. Domain adaptation for visual applications: A comprehensive survey. *arXiv preprint arXiv:1702.05374* (2017).

- [50] Mostafa El Habib Daho, Nesma Settouti, Mohammed El Amine Lazouni, and Mohammed El Amine Chikh. 2014. Weighted vote for trees aggregation in random forest. In *Proceedings of the 2014 International Conference on Multimedia Computing and Systems (ICMCS)*. 438–443. DOI : <https://doi.org/10.1109/ICMCS.2014.6911187>
- [51] Bharath Bhushan Damodaran, Benjamin Kellenberger, Rémi Flamary, Devis Tuia, and Nicolas Courty. 2018. Deep-JDOT: Deep joint distribution optimal transport for unsupervised domain adaptation. In *Computer Vision – ECCV 2018*, Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (Eds.). Springer International Publishing, Cham, 467–483.
- [52] Debasmit Das and C. S. George Lee. 2018. Graph matching and pseudo-label guided deep unsupervised domain adaptation. In *Artificial Neural Networks and Machine Learning – ICANN 2018*, Věra Kůrková, Yannis Manolopoulos, Barbara Hammer, Lazaros Iliadis, and Ilias Maglogiannis (Eds.). Springer International Publishing, Cham, 342–352.
- [53] Debasmit Das and C. S. George Lee. 2018. Sample-to-sample correspondence for unsupervised domain adaptation. *Engineering Applications of Artificial Intelligence* 73 (2018), 80–91. DOI : <https://doi.org/10.1016/j.engappai.2018.05.001>
- [54] Debasmit Das and C. S. George Lee. 2018. Unsupervised domain adaptation using regularized hyper-graph matching. In *Proceedings of the 2018 25th IEEE International Conference on Image Processing (ICIP)*. 3758–3762. DOI : <https://doi.org/10.1109/ICIP.2018.8451152>
- [55] Hal Daumé III. 2007. Frustratingly easy domain adaptation. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. 256–263.
- [56] Hal Daumé III and Daniel Marcu. 2006. Domain adaptation for statistical classifiers. *Journal of Artificial Intelligence Research* 26 (2006), 101–126.
- [57] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition*. 248–255. DOI : <https://doi.org/10.1109/CVPR.2009.5206848>
- [58] Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, and Jianbin Jiao. 2018. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [59] Zhijie Deng, Yucen Luo, and Jun Zhu. 2019. Cluster alignment with a teacher for unsupervised domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [60] Emily L. Denton, Soumith Chintala, Arthur Szlam, and Rob Fergus. 2015. Deep generative image models using a Laplacian pyramid of adversarial networks. In *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Eds.). Curran Associates, Inc., 1486–1494. <http://papers.nips.cc/paper/5773-deep-generative-image-models-using-a-laplacian-pyramid-of-adversarial-networks.pdf>.
- [61] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. 2017. Adversarial feature learning. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=BJtNZAFgg>.
- [62] Mark Dredze, Alex Kulesza, and Koby Crammer. 2010. Multi-domain learning by confidence-weighted parameter combination. *Machine Learning* 79, 1 (01 May 2010), 123–149. DOI : <https://doi.org/10.1007/s10994-009-5148-0>
- [63] Lixin Duan, Dong Xu, and Ivor W. Tsang. 2012. Learning with augmented features for heterogeneous domain adaptation. In *Proceedings of the 29th International Conference on International Conference on Machine Learning (ICML '12)*. Omnipress, Madison, WI., 667–674.
- [64] Ishan Durugkar, Ian Gemp, and Sridhar Mahadevan. 2017. Generative multi-adversarial networks. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=Byk-VI9eg>.
- [65] Gintare Karolina Dziugaite, Daniel M. Roy, and Zoubin Ghahramani. 2015. Training generative neural networks via maximum mean discrepancy optimization. In *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence (UAI'15)*. AUAI Press, Arlington, VA, 258–267. <http://dl.acm.org/citation.cfm?id=3020847.3020875>
- [66] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2018. Transfer learning for time series classification. In *Proceedings of the 2018 IEEE International Conference on Big Data (Big Data)*. 1367–1376. DOI : <https://doi.org/10.1109/BigData.2018.8621990>
- [67] William Fedus, Mihaela Rosca, Balaji Lakshminarayanan, Andrew M. Dai, Shakir Mohamed, and Ian Goodfellow. 2018. Many paths to equilibrium: GANs do not need to decrease a divergence at every step. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=ByQpn1ZA->.
- [68] Chelsea Finn, Paul Christiano, Pieter Abbeel, and Sergey Levine. 2016. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *arXiv preprint arXiv:1611.03852* (2016).
- [69] Geoff French, Michal Mackiewicz, and Mark Fisher. 2018. Self-ensembling for visual domain adaptation. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=rkpoTaxA->.
- [70] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, Kun Zhang, and Dacheng Tao. 2018. Geometry-consistent adversarial networks for one-sided unsupervised domain mapping. *arXiv preprint arXiv:1809.05852* (2018).
- [71] Lisheng Fu, Thien Huu Nguyen, Bonan Min, and Ralph Grishman. 2017. Domain adaptation for relation extraction with domain adversarial neural network. In *Proceedings of the 8th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, Vol. 2. 425–429.

- [72] Zhe Gan, Liqun Chen, Weiyao Wang, Yuchen Pu, Yizhe Zhang, Hao Liu, Chunyuan Li, and Lawrence Carin. 2017. Triangle generative adversarial networks. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 5247–5256. <http://papers.nips.cc/paper/7109-triangle-generative-adversarial-networks.pdf>.
- [73] Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Francis Bach and David Blei (Eds.), Vol. 37. PMLR, 1180–1189. <http://proceedings.mlr.press/v37/ganin15.html>.
- [74] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *Journal of Machine Learning Research* 17, 59 (2016), 1–35. <http://jmlr.org/papers/v17/15-239.html>.
- [75] Jon Gauthier. 2014. *Conditional Generative Adversarial Nets for Convolutional Face Generation*. Technical Report. Stanford University.
- [76] Muhammad Ghifary, W. Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. 2015. Domain generalization for object recognition with multi-task autoencoders. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [77] Muhammad Ghifary, W. Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. 2016. Deep reconstruction-classification networks for unsupervised domain adaptation. In *Computer Vision – ECCV 2016*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, Cham, 597–613.
- [78] Behnam Gholami, Pritish Sahu, Ognjen Rudovic, Konstantinos Bousmalis, and Vladimir Pavlovic. 2018. Unsupervised multi-target domain adaptation: An information theoretic approach. *arXiv preprint arXiv:1810.11547* (2018).
- [79] Arnab Ghosh, Viveka Kulharia, Vinay P. Namboodiri, Philip H. S. Torr, and Puneet K. Dokania. 2018. Multi-agent diverse generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8513–8521.
- [80] Ian Goodfellow. 2016. NIPS 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160* (2016).
- [81] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT press.
- [82] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 2672–2680. <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.
- [83] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014).
- [84] Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex J. Smola. 2007. A kernel method for the two-sample-problem. In *Advances in Neural Information Processing Systems*. 513–520.
- [85] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. A kernel two-sample test. *Journal of Machine Learning Research* 13, Mar (2012), 723–773.
- [86] Aditya Grover, Manik Dhar, and Stefano Ermon. 2017. Flow-GAN: Combining maximum likelihood and adversarial learning in generative models. *arXiv preprint arXiv:1705.08868* (2017).
- [87] Liangyan Gui, Kevin Zhang, Yu-Xiong Wang, Xiaodan Liang, José M. F. Moura, and Manuela M. Veloso. 2018. Teaching robots to predict human motion. (2018). preprint on webpage at <http://www.cs.cmu.edu/mmv/papers/18iros-GuiEtAl.pdf>.
- [88] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C. Courville. 2017. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 5767–5777. <http://papers.nips.cc/paper/7159-improved-training-of-wasserstein-gans.pdf>.
- [89] Jiang Guo, Darsh Shah, and Regina Barzilay. 2018. Multi-source domain adaptation with mixture of experts. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 4694–4703.
- [90] Eman T. Hassan, Xin Chen, and David Crandall. 2018. Unsupervised domain adaptation using generative models and self-ensembling. *arXiv preprint arXiv:1812.00479* (2018).
- [91] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs trained by a two-time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 6626–6637. <http://papers.nips.cc/paper/7240-gans-trained-by-a-two-time-scale-update-rule-converge-to-a-local-nash-equilibrium.pdf>.
- [92] Avinash Hindupur. 2018. The GAN Zoo. Retrieved February 25, 2019 from <https://github.com/hindupuravinash/the-gan-zoo>.
- [93] Saifuddin Hitawala. 2018. Comparative study on generative adversarial networks. *arXiv preprint arXiv:1801.04271* (2018).

- [94] Quan Hoang, Tu Dinh Nguyen, Trung Le, and Dinh Phung. 2018. MGAN: Training generative adversarial nets with multiple generators. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=rkmu5b0a->.
- [95] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780. DOI : <https://doi.org/10.1162/neco.1997.9.8.1735> arXiv:<https://doi.org/10.1162/neco.1997.9.8.1735>
- [96] Judy Hoffman, Mehryar Mohri, and Ningshan Zhang. 2018. Algorithms and theory for multiple-source adaptation. In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 8246–8256. <http://papers.nips.cc/paper/8046-algorithms-and-theory-for-multiple-source-adaptation.pdf>.
- [97] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. 2018. CyCADA: Cycle-consistent adversarial domain adaptation. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Jennifer Dy and Andreas Krause (Eds.), Vol. 80. PMLR, 1994–2003. <http://proceedings.mlr.press/v80/hoffman18a.html>.
- [98] Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. 2016. FCNS in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649* (2016).
- [99] Weixiang Hong, Zhenzhen Wang, Ming Yang, and Junsong Yuan. 2018. Conditional generative adversarial network for structured domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [100] Yongjun Hong, Uiwon Hwang, Jaeyoon Yoo, and Sungroh Yoon. 2019. How generative adversarial networks and their variants work: An overview. *ACM Comput. Surv.* 52, 1, Article 10 (Feb. 2019), 43 pages. DOI : <https://doi.org/10.1145/3301282>
- [101] Ehsan Hosseini-Asl, Yingbo Zhou, Caiming Xiong, and Richard Socher. 2019. Augmented cyclic adversarial learning for low resource domain adaptation. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=B1G9doA9F7>.
- [102] Yen-Chang Hsu, Zhaoyang Lv, and Zsolt Kira. 2018. Learning to cluster in order to transfer across domains and tasks. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=ByRWCqvT->.
- [103] Haoshuo Huang, Qixing Huang, and Philipp Krahenbuhl. 2018. Domain transfer through deep activation matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [104] Ling Huang, Anthony D. Joseph, Blaine Nelson, Benjamin I. P. Rubinstein, and J. D. Tygar. 2011. Adversarial machine learning. In *Proceedings of the 4th ACM Workshop on Security and Artificial Intelligence (AISec'11)*. ACM, New York,, 43–58. DOI : <https://doi.org/10.1145/2046684.2046692>
- [105] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. 2018. Multimodal unsupervised image-to-image translation. *arXiv preprint arXiv:1804.04732* (2018).
- [106] Yao-Hung Hubert Tsai, Yi-Ren Yeh, and Yu-Chiang Frank Wang. 2016. Learning cross-domain landmarks for heterogeneous domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [107] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. 2018. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [108] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Francis Bach and David Blei (Eds.), Vol. 37. PMLR, Lille, France, 448–456. <http://proceedings.mlr.press/v37/ioffe15.html>.
- [109] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [110] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. 2018. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407* (2018).
- [111] Jing Jiang. 2008. *Domain Adaptation in Natural Language Processing*. Technical Report. University of Illinois at Urbana-Champaign.
- [112] Alexia Jolicoeur-Martineau. 2018. The relativistic discriminator: A key element missing from standard GAN. *arXiv preprint arXiv:1807.00734* (2018).
- [113] Mahesh Joshi, William W. Cohen, Mark Dredze, and Carolyn P. Rosé. 2012. Multi-domain learning: When do domains matter? In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL'12)*. Association for Computational Linguistics, Stroudsburg, PA,, 1302–1312. <http://dl.acm.org/citation.cfm?id=2390948.2391096>

- [114] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G. Hauptmann. 2019. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [115] Guoliang Kang, Liang Zheng, Yan Yan, and Yi Yang. 2018. Deep adversarial attention alignment for unsupervised domain adaptation: The benefit of target expectation maximization. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [116] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive growing of GANs for improved quality, stability, and variation. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=Hk99zCeAb>.
- [117] Mahyar Khayatkhoei, Maneesh K. Singh, and Ahmed Elgammal. 2018. Disconnected manifold learning for generative adversarial networks. In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 7343–7353. <http://papers.nips.cc/paper/7964-disconnected-manifold-learning-for-generative-adversarial-networks.pdf>.
- [118] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. 2017. Learning to discover cross-domain relations with generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70 (ICML '17)*. JMLR.org, 1857–1865. <http://dl.acm.org/citation.cfm?id=3305381.3305573>
- [119] Wei-Jen Ko, Greg Durrett, and Junyi Jessy Li. 2018. Domain agnostic real-valued specificity prediction. *arXiv preprint arXiv:1811.05085* (2018).
- [120] Naveen Kodali, Jacob Abernethy, James Hays, and Zsolt Kira. 2017. On convergence and stability of GANs. *arXiv preprint arXiv:1705.07215* (2017).
- [121] Wouter M. Kouw. 2018. An introduction to domain adaptation and transfer learning. *arXiv preprint arXiv:1812.11806* (2018).
- [122] Wouter M. Kouw and Marco Loog. 2019. A review of single-source unsupervised domain adaptation. *arXiv preprint arXiv:1901.05335* (2019).
- [123] Alex Krizhevsky. 2009. *Learning Multiple Layers of Features from Tiny Images*. Technical Report. University of Toronto.
- [124] Abhishek Kumar, Prasanna Sattigeri, Kahini Wadhawan, Leonid Karlinsky, Rogerio Feris, Bill Freeman, and Gregory Wornell. 2018. Co-regularized alignment for unsupervised domain adaptation. In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 9345–9356. <http://papers.nips.cc/paper/8146-co-regularized-alignment-for-unsupervised-domain-adaptation.pdf>.
- [125] Vinod Kumar Kurmi, Shanu Kumar, and Vinay P. Namboodiri. 2019. Attending to discriminative certainty for domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [126] Samuli Laine and Timo Aila. 2017. Temporal ensembling for semi-supervised learning. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=BJ6oOfqge>.
- [127] Issam Laradji and Reza Babanezhad. 2018. M-ADDA: Unsupervised domain adaptation with deep metric learning. *arXiv preprint arXiv:1807.02552* (2018).
- [128] Alessandro Lazaric. 2012. *Transfer in Reinforcement Learning: A Framework and a Survey*. Springer, Berlin,, 143–173. DOI : https://doi.org/10.1007/978-3-642-27645-3_5
- [129] Trung Le, Khanh Nguyen, and Dinh Phung. 2018. Theoretical perspective of deep domain adaptation. *arXiv preprint arXiv:1811.06199* (2018).
- [130] Yann LeCun, Corinna Cortes, and Christopher J. C. Burges. 1998. The MNIST database of handwritten digits. Retrieved August 16, 2018 from <http://yann.lecun.com/exdb/mnist/>.
- [131] Yann LeCun, Ofer Matan, Bernhard Boser, John S. Denker, Don Henderson, Richard E. Howard, Wayne Hubbard, L. D. Jacket, and Henry S. Baird. 1990. Handwritten zip code recognition with multilayer networks. In *Proceedings. 10th International Conference on Pattern Recognition*, Vol. ii. 35–40 vol.2. DOI : <https://doi.org/10.1109/ICPR.1990.119325>
- [132] Chen-Yu Lee, Tanmay Batra, Mohammad Haris Baig, and Daniel Ulbricht. 2019. Sliced Wasserstein discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [133] Kuan-Hui Lee, German Ros, Jie Li, and Adrien Gaidon. 2019. SPIGAN: Privileged adversarial learning from simulation. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=rkxoNnCF5Q>.
- [134] Seungmin Lee, Dongwan Kim, Namil Kim, and Seong-Gyun Jeong. 2019. Drop to adapt: Learning discriminative features for unsupervised domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [135] Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. 2017. MMD GAN: Towards deeper understanding of moment matching network. In *Advances in Neural Information Processing Systems*. 2203–2213.

- [136] Jerry Li. 2018. Twin-GAN-unpaired cross-domain image translation with weight-sharing GANs. *arXiv preprint arXiv:1809.00946* (2018).
- [137] Jingjing Li, Ke Lu, Zi Huang, Lei Zhu, and Heng Tao Shen. 2019. Heterogeneous domain adaptation through progressive alignment. *IEEE Transactions on Neural Networks and Learning Systems* 30, 5 (May 2019), 1381–1391. DOI: <https://doi.org/10.1109/TNNLS.2018.2868854>
- [138] Peilun Li, Xiaodan Liang, Daoyuan Jia, and Eric P. Xing. 2018. Semantic-aware GRAD-GAN for virtual-to-real urban scene adaption. *arXiv preprint arXiv:1801.01726* (2018).
- [139] Yujia Li, Kevin Swersky, and Rich Zemel. 2015. Generative moment matching networks. In *Proceedings of the 32nd International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Francis Bach and David Blei (Eds.), Vol. 37. PMLR, Lille, France, 1718–1727. <http://proceedings.mlr.press/v37/li15.html>.
- [140] Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. 2018. Deep domain generalization via conditional invariant adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [141] Yanghao Li, Naiyan Wang, Jianping Shi, Xiaodi Hou, and Jiaying Liu. 2018. Adaptive batch normalization for practical domain adaptation. *Pattern Recognition* 80 (2018), 109–117. DOI: <https://doi.org/10.1016/j.patcog.2018.03.005>
- [142] Yu-Jhe Li, Fu-En Yang, Yen-Cheng Liu, Yu-Ying Yeh, Xiaofei Du, and Yu-Chiang Frank Wang. 2018. Adaptation and re-identification network: An unsupervised deep transfer learning approach to person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- [143] Ming-Yu Liu and Oncel Tuzel. 2016. Coupled generative adversarial networks. In *Advances in Neural Information Processing Systems* 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 469–477. <http://papers.nips.cc/paper/6544-coupled-generative-adversarial-networks.pdf>.
- [144] Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2017. Adversarial multi-task learning for text classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vol. 1. 1–10.
- [145] Shaohui Liu, Yi Wei, Jiwen Lu, and Jie Zhou. 2018. An improved evaluation framework for generative adversarial networks. *arXiv preprint arXiv:1803.07474* (2018).
- [146] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [147] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. 2015. Learning transferable features with deep adaptation networks. In *Proceedings of the 32nd International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Francis Bach and David Blei (Eds.), Vol. 37. PMLR, 97–105. <http://proceedings.mlr.press/v37/long15.html>.
- [148] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I. Jordan. 2018. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems* 31, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 1640–1650. <http://papers.nips.cc/paper/7436-conditional-adversarial-domain-adaptation.pdf>.
- [149] Mingsheng Long, Jianmin Wang, Guiguang Ding, Jianguang Sun, and Philip S. Yu. 2013. Transfer feature learning with joint distribution adaptation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [150] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. 2016. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems* 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 136–144. <http://papers.nips.cc/paper/6110-unsupervised-domain-adaptation-with-residual-transfer-networks.pdf>.
- [151] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. 2017. Deep transfer learning with joint adaptation networks. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Doina Precup and Yee Whye Teh (Eds.), Vol. 70. PMLR, International Convention Centre, Sydney, Australia, 2208–2217. <http://proceedings.mlr.press/v70/long17a.html>.
- [152] Jie Lu, Vahid Behbood, Peng Hao, Hua Zuo, Shan Xue, and Guangquan Zhang. 2015. Transfer learning using computational intelligence: A survey. *Knowledge-Based Systems* 80 (2015), 14–23. DOI: <https://doi.org/10.1016/j.knosys.2015.01.010> 25th anniversary of Knowledge-Based Systems.
- [153] Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. 2018. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. *arXiv preprint arXiv:1809.09478* (2018).
- [154] Xinhong Ma, Tianzhu Zhang, and Changsheng Xu. 2019. GCAN: Graph convolutional adversarial network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [155] Faisal Mahmood, Richard Chen, and Nicholas J. Durr. 2018. Unsupervised reverse domain adaptation for synthetic medical images via adversarial training. *IEEE Transactions on Medical Imaging* 37, 12 (Dec. 2018), 2572–2581. DOI: <https://doi.org/10.1109/TMI.2018.2842767>

- [156] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. 2015. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644* (2015).
- [157] Massimiliano Mancini, Lorenzo Porzi, Samuel Rota Bulò, Barbara Caputo, and Elisa Ricci. 2018. Boosting domain adaptation by discovering latent domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [158] Padala Manisha and Sujit Gujar. 2018. Generative adversarial networks (GANs): What it can generate and What it cannot? *arXiv preprint arXiv:1804.00140* (2018).
- [159] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. 2009. Domain adaptation with multiple sources. In *Advances in Neural Information Processing Systems 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou (Eds.). Curran Associates, Inc., 1041–1048. <http://papers.nips.cc/paper/3550-domain-adaptation-with-multiple-sources.pdf>.
- [160] Xudong Mao and Qing Li. 2018. Unpaired multi-domain image generation via regularized conditional GANs. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*. AAAI Press, 2553–2559. <http://dl.acm.org/citation.cfm?id=3304889.3305015>
- [161] Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, Zhen Wang, and Stephen Paul Smolley. 2017. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*. 2794–2802.
- [162] Anna Margolis. 2011. *A Literature Review of Domain Adaptation with Unlabeled Data*. Technical Report. University of Washington.
- [163] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. 2017. Unrolled generative adversarial networks. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=BydrOicle>.
- [164] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
- [165] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. 2018. Spectral normalization for generative adversarial networks. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=B1QRgziT->.
- [166] Takeru Miyato and Masanori Koyama. 2018. cGANs with projection discriminator. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=ByS1VpgRZ>.
- [167] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: A regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018), 1–1. DOI: <https://doi.org/10.1109/TPAMI.2018.2858821>
- [168] Boris Moiseev, Artem Konev, Alexander Chigorin, and Anton Konushin. 2013. Evaluation of traffic sign recognition methods trained on synthetically generated data. In *Advanced Concepts for Intelligent Vision Systems*, Jacques Blanc-Talon, Andrzej Kasinski, Wilfried Philips, Dan Popescu, and Paul Scheunders (Eds.). Springer International Publishing, Cham, 576–583.
- [169] Gaspard Monge. 1781. Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences de Paris* (1781).
- [170] Gonçalo Mordido, Haojin Yang, and Christoph Meinel. 2018. Dropout-GAN: Learning from a dynamic ensemble of discriminators. *arXiv preprint arXiv:1807.11346* (2018).
- [171] Pietro Morerio, Jacopo Cavazza, and Vittorio Murino. 2018. Minimal-entropy correlation alignment for unsupervised deep domain adaptation. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=rjWechg0Z>.
- [172] Pietro Morerio and Vittorio Murino. 2017. Correlation alignment by Riemannian metric for domain adaptation. *arXiv preprint arXiv:1705.08180* (2017).
- [173] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. 2013. Domain generalization via invariant feature representation. In *Proceedings of the 30th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Sanjoy Dasgupta and David McAllester (Eds.), Vol. 28. PMLR, 10–18. <http://proceedings.mlr.press/v28/muandet13.html>.
- [174] Jogendra Nath Kundu, Phani Krishna Uppala, Anuj Pahuja, and R. Venkatesh Babu. 2018. AdaDepth: Unsupervised content congruent adaptation for depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [175] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. 2011. Reading digits in natural images with unsupervised feature learning. In *Proceedings of the NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, Vol. 2011. 5.
- [176] Anh Nguyen, Jeff Clune, Yoshua Bengio, Alexey Dosovitskiy, and Jason Yosinski. 2017. Plug & play generative networks: Conditional iterative generation of images in latent space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [177] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. 2016. f-GAN: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama,

- U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 271–279. <http://papers.nips.cc/paper/6066-f-gan-training-generative-neural-samplers-using-variational-divergence-minimization.pdf>.
- [178] Augustus Odena, Jacob Buckman, Catherine Olsson, Tom Brown, Christopher Olah, Colin Raffel, and Ian Goodfellow. 2018. Is generator conditioning causally related to GAN performance? In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Jennifer Dy and Andreas Krause (Eds.), Vol. 80. PMLR, 3846–3855. <http://proceedings.mlr.press/v80/odena18a.html>.
 - [179] Augustus Odena, Christopher Olah, and Jonathon Shlens. 2017. Conditional image synthesis with auxiliary classifier GANs. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Doina Precup and Yee Whye Teh (Eds.), Vol. 70. PMLR, 2642–2651. <http://proceedings.mlr.press/v70/odena17a.html>.
 - [180] Sinno Jialin Pan and Qiang Yang. 2010. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* 22, 10 (Oct 2010), 1345–1359. DOI : <https://doi.org/10.1109/TKDE.2009.191>
 - [181] David Keetae Park, Seungjoo Yoo, Hyojin Bahng, Jaegul Choo, and Noseong Park. 2018. MEGAN: Mixture of experts of generative adversarial networks for multimodal image generation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*. AAAI Press, 878–884. <http://dl.acm.org/citation.cfm?id=3304415.3304540>
 - [182] Vishal M. Patel, Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. 2015. Visual domain adaptation: A survey of recent advances. *IEEE Signal Processing Magazine* 32, 3 (May 2015), 53–69. DOI : <https://doi.org/10.1109/MSP.2014.2347059>
 - [183] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. 2018. Multi-adversarial domain adaptation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
 - [184] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. 2018. Moment matching for multi-source domain adaptation. *arXiv preprint arXiv:1812.01754* (2018).
 - [185] Christian S. Perone, Pedro Ballester, Rodrigo C. Barros, and Julien Cohen-Adad. 2018. Unsupervised domain adaptation for medical imaging segmentation with self-ensembling. *arXiv preprint arXiv:1811.06042* (2018).
 - [186] David Pfau and Oriol Vinyals. 2016. Connecting generative adversarial networks and actor-critic methods. *arXiv preprint arXiv:1610.01945* (2016).
 - [187] Pedro O. Pinheiro. 2018. Unsupervised domain adaptation with similarity learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [188] Lerrel Pinto, James Davidson, and Abhinav Gupta. 2017. Supervision via competition: Robot adversaries for learning tasks. In *Proceedings of the 2017 IEEE International Conference on Robotics and Automation (ICRA)*. 1601–1608. DOI : <https://doi.org/10.1109/ICRA.2017.7989190>
 - [189] Sanjay Purushotham, Wilka Carvalho, Tanachat Nilanon, and Yan Liu. 2017. Variational adversarial deep domain adaptation for health care time series analysis. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=rk9eAFcxg>.
 - [190] Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* (2015).
 - [191] Ievgen Redko, Nicolas Courty, Rémi Flamary, and Devis Tuia. 2018. Optimal transport for multi-source domain adaptation under target shift. *arXiv preprint arXiv:1803.04899* (2018).
 - [192] Ievgen Redko, Amaury Habrard, and Marc Sebban. 2017. Theoretical analysis of domain adaptation with optimal transport. In *Machine Learning and Knowledge Discovery in Databases*, Michelangelo Ceci, Jaakko Hollmén, Ljupčo Todorovski, Celine Vens, and Sašo Džeroski (Eds.). Springer International Publishing, Cham, 737–753.
 - [193] Jian Ren, Jianchao Yang, Ning Xu, and David J. Foran. 2018. Factorized adversarial networks for unsupervised domain adaptation. *arXiv preprint arXiv:1806.01376* (2018).
 - [194] Oren Rippel and Lubomir Bourdev. 2017. Real-time adaptive image compression. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Doina Precup and Yee Whye Teh (Eds.), Vol. 70. PMLR, 2922–2930. <http://proceedings.mlr.press/v70/rippel17a.html>.
 - [195] Oren Rippel, Lubomir Bourdev, Carissa Lew, and Sanjay Nair. 2018. Using generative adversarial networks in compression. US Patent App. 15/844,449.
 - [196] Amélie Royer, Konstantinos Bousmalis, Stephan Gouws, Fred Bertsch, Inbar Mosseri, Forrester Cole, and Kevin Murphy. 2017. XGAN: Unsupervised image-to-image translation for many-to-many mappings. *arXiv preprint arXiv:1711.05139* (2017).
 - [197] Artem Rozantsev, Mathieu Salzmann, and Pascal Fua. 2018. Residual parameter transfer for deep domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [198] Artem Rozantsev, Mathieu Salzmann, and Pascal Fua. 2019. Beyond sharing weights for deep domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 4 (April 2019), 801–814. DOI : <https://doi.org/10.1109/TPAMI.2018.2814042>

- [199] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)* 115, 3 (2015), 211–252. DOI: <https://doi.org/10.1007/s11263-015-0816-y>
- [200] Paolo Russo, Fabio M. Carlucci, Tatiana Tommasi, and Barbara Caputo. 2018. From source to target and back: Symmetric bi-directional adaptive GAN. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [201] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. 2010. Adapting visual category models to new domains. In *Computer Vision – ECCV 2010*, Kostas Daniilidis, Petros Maragos, and Nikos Paragios (Eds.). Springer Berlin Heidelberg, Berlin,, 213–226.
- [202] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Darrell, and Kate Saenko. 2019. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [203] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. 2017. Asymmetric tri-training for unsupervised domain adaptation. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Doina Precup and Yee Whye Teh (Eds.), Vol. 70. PMLR, 2988–2997. <http://proceedings.mlr.press/v70/saito17a.html>.
- [204] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. 2018. Adversarial dropout regularization. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=HJloJWZCZ>.
- [205] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. 2018. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [206] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 2234–2242. <http://papers.nips.cc/paper/6125-improved-techniques-for-training-gans.pdf>.
- [207] Swami Sankaranarayanan, Yogesh Balaji, Carlos D. Castillo, and Rama Chellappa. 2018. Generate to adapt: Aligning domains using generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [208] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. 2018. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [209] Shibani Santurkar, Ludwig Schmidt, and Aleksander Madry. 2018. A classification-based study of covariate shift in GAN distributions. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Jennifer Dy and Andreas Krause (Eds.), Vol. 80. PMLR, 4487–4496. <http://proceedings.mlr.press/v80/santurkar18a.html>.
- [210] Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. 2018. How does batch normalization help optimization? In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 2483–2493. <http://papers.nips.cc/paper/7515-how-does-batch-normalization-help-optimization.pdf>.
- [211] Steffen Schneider, Alexander S. Ecker, Jakob H. Macke, and Matthias Bethge. 2018. Salad: A toolbox for semi-supervised adaptive learning across domains. <https://openreview.net/forum?id=S1lTifykqm>.
- [212] Alice Schoenauer Sebag, Louise Heinrich, Marc Schoenauer, Michèle Sebag, Lani Wu, and Steven Altschuler. 2019. Multi-domain adversarial learning. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=SkIv5iRqYX>.
- [213] Ozan Sener, Hyun Oh Song, Ashutosh Saxena, and Silvio Savarese. 2016. Learning transferrable representations for unsupervised domain adaptation. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 2110–2118. <http://papers.nips.cc/paper/6360-learning-transferrable-representations-for-unsupervised-domain-adaptation.pdf>.
- [214] Chao Shang, Aaron Palmer, Jiangwen Sun, Ko-Shin Chen, Jin Lu, and Jinbo Bi. 2017. VIGAN: Missing view imputation with generative adversarial networks. In *Proceedings of the 2017 IEEE International Conference on Big Data (Big Data)*. 766–775. DOI: <https://doi.org/10.1109/BigData.2017.8257992>
- [215] Ling Shao, Fan Zhu, and Xuelong Li. 2015. Transfer learning for visual categorization: A survey. *IEEE Transactions on Neural Networks and Learning Systems* 26, 5 (May 2015), 1019–1034. DOI: <https://doi.org/10.1109/TNNLS.2014.2330900>
- [216] Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. 2018. Wasserstein distance guided representation learning for domain adaptation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- [217] Yusuke Shinohara. 2016. Adversarial multi-task learning of deep neural networks for robust speech recognition. In *Proceedings of Interspeech 2016*. 2369–2372. DOI: <https://doi.org/10.21437/Interspeech.2016-879>

- [218] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. 2017. Learning from simulated and unsupervised images through adversarial training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [219] Rui Shu, Hung Bui, Hirokazu Narui, and Stefano Ermon. 2018. A DIRT-T approach to unsupervised domain adaptation. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=H1q-TM-AW>.
- [220] Yang Shu, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. 2019. Transferable curriculum for weakly-supervised domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (Jul. 2019), 4951–4958. DOI : <https://doi.org/10.1609/aaai.v33i01.33014951>
- [221] Stephen Sinclair. 2018. Sounderfeit: Cloning a physical model using a conditional adversarial autoencoder. *arXiv preprint arXiv:1806.09617* (2018).
- [222] Kihyuk Sohn, Wenling Shang, Xiang Yu, and Manmohan Chandraker. 2019. Unsupervised domain adaptation for distance metric learning. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=BklhAj09K7>.
- [223] Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel. 2011. The German traffic sign recognition benchmark: A multi-class classification competition. In *Proceedings of the 2011 International Joint Conference on Neural Networks*. 1453–1460. DOI : <https://doi.org/10.1109/IJCNN.2011.6033395>
- [224] Baochen Sun, Jiashi Feng, and Kate Saenko. 2016. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*. <https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12443>.
- [225] Baochen Sun and Kate Saenko. 2016. Deep CORAL: Correlation alignment for deep domain adaptation. In *Proceedings of Computer Vision – ECCV 2016 Workshops*, Gang Hua and Hervé Jégou (Eds.). Springer International Publishing, Cham, 443–450.
- [226] Shiliang Sun, Honglei Shi, and Yuanbin Wu. 2015. A survey of multi-source domain adaptation. *Information Fusion* 24 (2015), 84–92. DOI : <https://doi.org/10.1016/j.inffus.2014.12.003>
- [227] Dougal J. Sutherland, Hsiao-Yu Tung, Heiko Strathmann, Soumyajit De, Aaditya Ramdas, Alex Smola, and Arthur Gretton. 2016. Generative models and model criticism via optimized maximum mean discrepancy. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=HJWHIKqgl>.
- [228] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [229] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013).
- [230] Yaniv Taigman, Adam Polyak, and Lior Wolf. 2016. Unsupervised cross-domain image generation. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Sk2Im59ex>.
- [231] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. 2018. A survey on deep transfer learning. In *Artificial Neural Networks and Machine Learning – ICANN 2018*, Věra Kůrková, Yannis Manolopoulos, Barbara Hammer, Lazaros Ilidis, and Ilias Maglogiannis (Eds.). Springer International Publishing, Cham, 270–279.
- [232] Hui Tang and Kui Jia. 2019. Discriminative adversarial domain adaptation. *arXiv preprint arXiv:1911.12036* (2019).
- [233] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in Neural Information Processing Systems* 30, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 1195–1204. <http://papers.nips.cc/paper/6719-mean-teachers-are-better-role-models-weight-averaged-consistency-targets-improve-semi-supervised-deep-learning-results.pdf>.
- [234] Matthew E. Taylor and Peter Stone. 2009. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research* 10, Jul (2009), 1633–1685.
- [235] Lucas Theis, Aaron van den Oord, and Matthias Bethge. 2016. A note on the evaluation of generative models. In *Proceedings of the International Conference on Learning Representations*. <https://arxiv.org/abs/1511.01844>
- [236] Ilya O. Tolstikhin, Sylvain Gelly, Olivier Bousquet, Carl-Johann Simon-Gabriel, and Bernhard Schölkopf. 2017. AdaGAN: Boosting generative models. In *Advances in Neural Information Processing Systems* 30, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 5424–5433. <http://papers.nips.cc/paper/7126-adagan-boosting-generative-models.pdf>.
- [237] Michele Tonutti, Emanuele Ruffaldi, Alessandro Cattaneo, and Carlo Alberto Avizzano. 2019. Robust and subject-independent driving manoeuvre anticipation through domain-adversarial recurrent neural networks. *Robotics and Autonomous Systems* 115 (2019), 162–173. DOI : <https://doi.org/10.1016/j.robot.2019.02.007>
- [238] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. 2018. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

- [239] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. 2015. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [240] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. 2017. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [241] Vincent Vercruyssen, Wannes Meert, and Jesse Davis. 2017. Transfer learning for time series anomaly detection. In *Proceedings of the CEUR Workshop*, Vol. 1924. 27–37.
- [242] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Mathieu Cord, and Patrick Pérez. 2018. ADVENT: Adversarial entropy minimization for domain adaptation in semantic segmentation. *arXiv preprint arXiv:1811.12833* (2018).
- [243] Chang Wang and Sridhar Mahadevan. 2011. Heterogeneous domain adaptation using manifold alignment. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*.
- [244] Jiawei Wang, Zhaoshui He, Chengjian Feng, Zhouping Zhu, Qinzhuang Lin, Jun Lv, and Shengli Xie. 2018. Domain confusion with self ensembling for unsupervised adaptation. *arXiv preprint arXiv:1810.04472* (2018).
- [245] Mei Wang and Weihong Deng. 2018. Deep visual domain adaptation: A survey. *Neurocomputing* 312 (2018), 135–153. DOI : <https://doi.org/10.1016/j.neucom.2018.05.083>
- [246] Ximei Wang, Liang Li, Weirui Ye, Mingsheng Long, and Jianmin Wang. 2019. Transferable attention for domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (Jul. 2019), 5345–5352. DOI : <https://doi.org/10.1609/aaai.v33i01.33015345>
- [247] Xiaolong Wang, Abhinav Shrivastava, and Abhinav Gupta. 2017. A-fast-RCNN: Hard positive generation via adversary for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [248] Yifei Wang, Wen Li, Dengxin Dai, and Luc Van Gool. 2017. Deep domain adaptation by geodesic distance minimization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*.
- [249] Kai-Ya Wei and Chiou-Ting Hsu. 2018. Generative adversarial guided learning for domain adaptation. In *Proceedings of the British Machine Vision Conference* (2018).
- [250] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. 2018. Person transfer GAN to bridge domain gap for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [251] Karl Weiss, Taghi M. Khoshgoftaar, and DingDing Wang. 2016. A survey of transfer learning. *Journal of Big Data* 3, 1 (28 May 2016), 9. DOI : <https://doi.org/10.1186/s40537-016-0043-6>
- [252] Yuhuai Wu, Yuri Burda, Ruslan Salakhutdinov, and Roger Grosse. 2017. On the quantitative analysis of decoder-based generative models. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=B1M8JF9xx>.
- [253] Yuxin Wu and Kaiming He. 2018. Group normalization. In *Proceedings of Computer Vision – ECCV 2018*, Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss (Eds.). Springer International Publishing, Cham, 3–19.
- [254] Markus Wulfmeier, Alex Bewley, and Ingmar Posner. 2018. Incremental adversarial domain adaptation for continually changing environments. In *Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA)*. 1–9. DOI : <https://doi.org/10.1109/ICRA.2018.8460982>
- [255] Qizhe Xie, Zihang Dai, Yulun Du, Eduard Hovy, and Graham Neubig. 2017. Controllable invariance through adversarial feature learning. In *Advances in Neural Information Processing Systems* 30, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 585–596. <http://papers.nips.cc/paper/6661-controllable-invariance-through-adversarial-feature-learning.pdf>.
- [256] Zhao Xin. 2019. A collection of AWESOME things about domain adaptation. Retrieved March 20, 2019 from <https://github.com/zhaoxin94/awesome-domain-adaptation>.
- [257] Qiantong Xu, Gao Huang, Yang Yuan, Chuan Guo, Yu Sun, Felix Wu, and Kilian Weinberger. 2018. An empirical study on evaluation metrics of generative adversarial networks. *arXiv preprint arXiv:1806.07755* (2018).
- [258] Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, and Liang Lin. 2018. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [259] Junchi Yan, Xu-Cheng Yin, Weiya Lin, Cheng Deng, Hongyuan Zha, and Xiaokang Yang. 2016. A short survey of recent advances in graph matching. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval (ICMR '16)*. ACM, New York, 167–174. DOI : <https://doi.org/10.1145/2911996.2912035>
- [260] Yongxin Yang and Timothy M. Hospedales. 2015. A unified perspective on multi-domain and multi-task learning. In *Proceedings of the International Conference on Learning Representations*. <https://arxiv.org/abs/1412.7489>
- [261] Zhilin Yang, Junjie Hu, Ruslan Salakhutdinov, and William Cohen. 2017. Semi-supervised QA with generative domain-adaptive nets. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vol. 1. 1040–1050.
- [262] Yuan Yao, Yu Zhang, Xutao Li, and Yunming Ye. 2020. Discriminative distribution alignment: A unified framework for heterogeneous domain adaptation. *Pattern Recognition* (2020), 107165. DOI : <https://doi.org/10.1016/j.patcog.2019.107165>

- [263] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. 2017. DualGAN: Unsupervised dual learning for image-to-image translation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [264] Donggeun Yoo, Namil Kim, Sunggyun Park, Anthony S. Paek, and In So Kweon. 2016. Pixel-level domain transfer. In *Proceedings of the European Conference on Computer Vision*. Springer, 517–532.
- [265] Jaeyoon Yoo, Yongjun Hong, YungKyun Noh, and Sungroh Yoon. 2017. Domain adaptation using adversarial learning for autonomous navigation. *arXiv preprint arXiv:1712.03742* (2017).
- [266] Chaohui Yu, Jindong Wang, Yiqiang Chen, and Meiyu Huang. 2019. Transfer learning with dynamic adversarial adaptation network. *arXiv preprint arXiv:1909.08184* (2019).
- [267] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. 2015. LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365* (2015).
- [268] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. 2018. Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318* (2018).
- [269] Honglun Zhang, Liqiang Xiao, Wenqing Chen, Yongkun Wang, and Yaohui Jin. 2018. Generative warfare nets: Ensemble via adversaries and collaborators. In *Proceedings of IJCAI*. 3075–3081.
- [270] JiChao Zhang. 2019. Adversarial Nets Papers. Retrieved February 25, 2019 from <https://github.com/zhangqianhui/AdversarialNetsPapers>.
- [271] Jing Zhang, Zewei Ding, Wanqing Li, and Philip Ogunbona. 2018. Importance weighted adversarial nets for partial domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [272] Jing Zhang, Wanqing Li, and Philip Ogunbona. 2017. Transfer learning for cross-dataset recognition: A survey. *arXiv preprint arXiv:1705.04396* (2017).
- [273] Weichen Zhang, Wanli Ouyang, Wen Li, and Dong Xu. 2018. Collaborative and adversarial network for unsupervised domain adaptation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [274] Yuan Zhang, Regina Barzilay, and Tommi Jaakkola. 2017. Aspect-augmented adversarial networks for domain adaptation. *Transactions of the Association for Computational Linguistics* 5 (2017), 515–528. DOI : https://doi.org/10.1162/tacl_a_00077 arXiv:https://doi.org/10.1162/tacl_a_00077
- [275] Yang Zhang, Philip David, and Boqing Gong. 2017. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- [276] Yue Zhang, Shun Miao, Tommaso Mansi, and Rui Liao. 2018. Task driven generative modeling for unsupervised domain adaptation: Application to X-ray image segmentation. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, Carlos Alberola-López, and Gabor Fichtinger (Eds.). Springer International Publishing, Cham, 599–607.
- [277] Yun Zhang, Nianbin Wang, Shaobin Cai, and Lei Song. 2018. Unsupervised domain adaptation by mapped correlation alignment. *IEEE Access* 6 (2018), 44698–44706. DOI : <https://doi.org/10.1109/ACCESS.2018.2865249>
- [278] Zhen Zhang, Mianzhi Wang, Yan Huang, and Arye Nehorai. 2018. Aligning infinite-dimensional covariance matrices in reproducing kernel hilbert spaces for domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [279] Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. 2019. On learning invariant representations for domain adaptation. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.), Vol. 97. PMLR, Long Beach, CA, 7523–7532. <http://proceedings.mlr.press/v97/zhao19a.html>.
- [280] Han Zhao, Shanghang Zhang, Guanhong Wu, José M. F. Moura, Joao P. Costeira, and Geoffrey J. Gordon. 2018. Adversarial multiple source domain adaptation. In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 8559–8570. <http://papers.nips.cc/paper/8075-adversarial-multiple-source-domain-adaptation.pdf>.
- [281] Han Zhao, Zhenyao Zhu, Junjie Hu, Adam Coates, and Geoff Gordon. 2017. Principled hybrids of generative and discriminative domain adaptation. *arXiv preprint arXiv:1705.09011* (2017).
- [282] Junbo Zhao, Michael Mathieu, and Yann LeCun. 2017. Energy-based generative adversarial network. In *Proceedings of the International Conference on Learning Representations*. <https://openreview.net/forum?id=ryh9pmcee>.
- [283] Mingmin Zhao, Shichao Yue, Dina Katabi, Tommi S. Jaakkola, and Matt T. Bianchi. 2017. Learning sleep stages from radio signals: A conditional adversarial architecture. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Doina Precup and Yee Whye Teh (Eds.), Vol. 70. PMLR, 4100–4109. <http://proceedings.mlr.press/v70/zhao17d.html>.
- [284] Sicheng Zhao, Bichen Wu, Joseph Gonzalez, Sanjit A Seshia, and Kurt Keutzer. 2018. Unsupervised domain adaptation: From simulation engine to the realworld. *arXiv preprint arXiv:1803.09180* (2018).
- [285] Erheng Zhong, Wei Fan, Qiang Yang, Olivier Verscheure, and Jiangtao Ren. 2010. Cross validation framework to choose amongst models and datasets for transfer learning. In *Machine Learning and Knowledge Discovery in*

- Databases*, José Luis Balcázar, Francesco Bonchi, Aristides Gionis, and Michèle Sebag (Eds.). Springer Berlin Heidelberg, Berlin., 547–562.
- [286] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. 2019. Invariance matters: Exemplar memory for domain adaptive person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [287] Zhun Zhong, Liang Zheng, Zhedong Zheng, Shaozi Li, and Yi Yang. 2018. Camera style adaptation for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 - [288] Joey Tianyi Zhou, Ivor W. Tsang, Sinno Jialin Pan, and Mingkui Tan. 2019. Multi-class heterogeneous domain adaptation. *Journal of Machine Learning Research* 20, 57 (2019), 1–31. <http://jmlr.org/papers/v20/13-580.html>.
 - [289] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
 - [290] Yang Zou, Zhiding Yu, B. V. K. Vijaya Kumar, and Jinsong Wang. 2018. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European Conference on Computer Vision (ECCV)*.

Received March 2019; revised January 2020; accepted May 2020