

Curriculum based Dropout Discriminator for Domain Adaptation

Vinod Kumar Kurmi
vinodkk@iitk.ac.in

Indian Institute of Technology Kanpur
Kanpur, India

Vipul Bajaj
vipulbjj@iitk.ac.in

Venkatesh K Subramanian
venkats@iitk.ac.in

Vinay P Namboodiri
vinaypn@iitk.ac.in

Abstract

Domain adaptation is essential to enable wide usage of deep learning based networks trained using large labeled datasets. Adversarial learning based techniques have shown their utility towards solving this problem using a discriminator that ensures source and target distributions are close. However, here we suggest that rather than using a point estimate, it would be useful if a distribution based discriminator could be used to bridge this gap. This could be achieved using multiple classifiers or using traditional ensemble methods. In contrast, we suggest that a Monte Carlo dropout based ensemble discriminator could suffice to obtain the distribution based discriminator. Specifically, we propose a curriculum based dropout discriminator that gradually increases the variance of the sample based distribution and the corresponding reverse gradients are used to align the source and target feature representations. The detailed results and thorough ablation analysis show that our model outperforms state-of-the-art results.

1 Introduction

Visual recognition has seen vast improvements based mainly on the success of deep learning based models [1]. These models are trained on very large annotated datasets such as Imagenet [2]. The deployment of these generically trained models require them to adapt to work in specific settings (for instance with catalog images in E-commerce websites). This problem is recognized as one of dataset bias and was demonstrated through the work of [3]. However, the requirement of a large annotated dataset becomes a bottleneck for training networks in deep learning frameworks. In this paper, we tackle the problem of adapting classifiers to work on datasets that *do not* have any labeled information. This problem is one of unsupervised domain adaptation in a more general setting.

Ganin and Lempitsky [4] proposed a method to solve unsupervised domain adaptation through back-propagation. In this method, the domain adaptation problem is solved by using a discriminator that ensures domain invariance of learned representations used for classification. There have been several methods [5, 6, 7] proposed for improving the

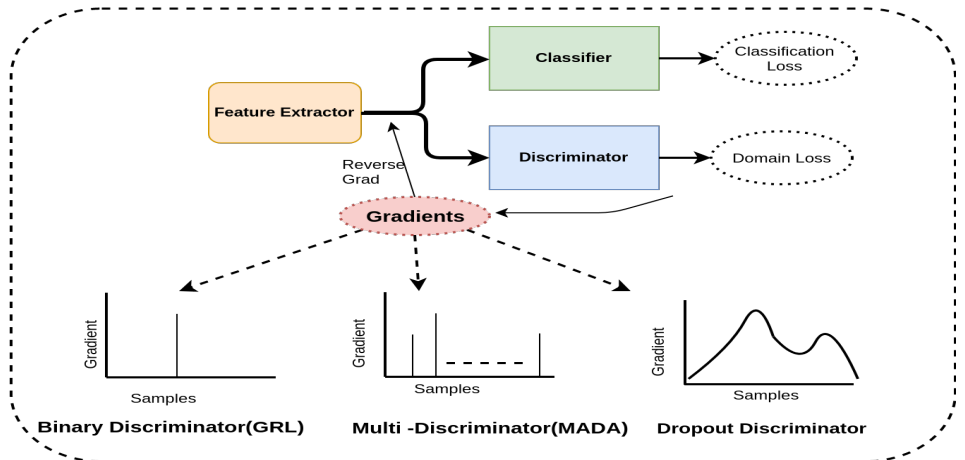


Figure 1: The difference in adversarial learning framework for domain adaption using binary discriminator, multi-discriminator and dropout discriminator. In binary discriminator [14], the feature extractor is trained with one gradient value of the discriminator. In the case of multi-discriminator [50], learning occurs with a fixed number of gradient values, whereas in dropout based discriminator, feature extractor learns a distribution rather than a single value.

discriminator. However, most of these involve an increase in the number of parameters. For instance a recent work by Pei *et al* (MADA) [50] addresses this issue through class-specific discriminators. This leads to a linear increase in the number of parameters with the number of classes in dataset. In contrast, we propose the use of curriculum-based dropout discriminator to obtain improved performance of the domain adaptation task without increasing the number of parameters. It makes our model’s applicability comprehensive as it can also adapt to datasets with a large number of classes. Specifically, in this paper, we propose Curriculum based Dropout Discriminator for Domain Adaptation (CD³A) and compare it with a variant, Dropout Discriminator for Domain Adaptation (D³A). It’s a novel approach that solves the above problem through an adversarial dynamic dropout based ensemble of discriminators, where we consider dropout as being a source of an ensemble of domain classifiers [13]. The proposed model also enables the discriminators reduce the prediction variance, remove overfitting, and average out the bias. The idea for this discriminator is illustrated in Figure 1. The initial discriminator by Ganin and Lempitsky [14] suggests the use of a single binary discriminator and MADA [50] extends it to class-specific cues. In contrast, CD³A obtains a discriminator distribution that provides a much-improved feedback for improving the feature extractor. The performance of any adversarial learning method largely depends upon the capability of the discriminator network. The ensemble method [13] improves the discriminator’s performance and makes it robust. We show that this indeed helps in an improved domain adaptation (around 5.3% improvement in Amazon-DSLRL adaptation) with much fewer parameters (~59M) than MADA (~98M). More importantly, our method does not increase the number of parameters as the number of classes increase, making it scalable to datasets with a large number of classes. Through this paper we make the following main contributions:

- We propose a method to obtain a dropout based discriminator that provides a distribution based discrimination for every sample ensuring a more robust feature adaptation

- We adopt a curriculum based dropout model, CD³A, that ensures gradual increase in the number of samples as the adaptation progresses to ensure better adaptation in contrast to a fixed number of samples based dropout distribution (D³A).
- We provide a thorough empirical analysis of the method (including statistical significance, discrepancy distance) and evaluate our approach against the state-of-the-art approaches.

2 Related Work

Domain Adaptation: A large number of methods have been proposed to tackle the domain adaptation problem. The basic common structure that has been followed is the Siamese architecture [5] with two streams, representing the source and target models. It is trained with a combination of a classification loss and the other being one of discrepancy loss or an adversarial loss. The classification loss depends on the source data label, while the discrepancy loss reduces the shift between the two domains. A discrepancy based deep learning method is that of deep domain confusion (DDC) [46]. The loss between a single FC (fully connected) layer of source and target feature extractor network is used to minimize the maximum mean discrepancy (MMD) between the source and the target. This approach is further extended by deep adaptation network (DAN) [22]. Recently, a number of other methods have been proposed which use discrepancy of domain [24, 51, 36, 39, 42, 43, 44, 53].

Adversarial Learning: In the domain adaptation setting, an adversarial network provides domain invariant representations by making the source and target domain indistinguishable by the discriminator. Adversarial Discriminative Domain Adaptation [48] uses an inverted label GAN loss to split the optimization into two independent objectives. One such method is the domain confusion based model proposed in [47] that considers a domain confusion objective. Domain-Adversarial Neural Networks (DANN) [10] integrates a gradient reversal layer into the standard architecture to promote the emergence of the learned representations that are discriminative for the main learning task on the source domain and non-discriminative concerning the shift between the domains. Recently, some works have been proposed which use an adversarial discriminative approach in solving the domain adaptation problem [3, 6, 16, 19, 21, 35, 51]. Similarly, the model proposed in [9, 7] exploits GANs with the aim to generate source-domain images such that they appear as if they were drawn from the target domain distribution. The closest related work to our approach is the work by [30] that extends the gradient reversal method by a class-specific discriminator.

Ensemble and Curriculum learning: Ensemble methods [20] can capture the uncertainty of the neural network (NN). Gal *et.al.* [11] use dropout to obtain the predictive uncertainty and apply Markov chain Monte Carlo [28] also known as MCMC at the test time to deal with intractable posterior. In discriminator based approaches, ensembles can be considered as multi-discriminator or multi-generator architecture. Multi discriminator approach has also been proposed by [9, 22, 29] to learn the data distribution more effectively. In Bayesian GAN [33], dropout in the discriminator is used which can be interpreted as an ensemble model [10]. The curriculum learning [2] enhances model's performance and its generalization capability. The performance of the GAN is also improved through the curriculum learning of the discriminator [57]. It has been shown that dropout can also work with curriculum learning [27]. In domain adaptation, a curriculum style learning approach has been applied in [52] to minimize the domain gap in semantic segmentation. The curriculum

domain adaptation first solves easy tasks such as estimating label distributions, then infers the necessary properties about the target domain. The theoretical framework for curriculum learning in transfer learning is proposed in [50]. Recently other curriculum learning based domain adaptation methods have been proposed in Transferable Curriculum Learning [40].

In contrast to the previous works, the main contribution of the present work is to propose a curriculum based dropout discriminator. We show that through the proposed method, we are able to outperform state of the art domain adaptation techniques in a scalable way by using fewer number of parameters as compared to techniques such as MADA [30] and similar number of parameters as GRL [14].

3 Motivation

In the adversarial domain adaptation problem, the previous methods have used classical statistical inference in the discriminator. A single discriminator learns the source and target domain classification. Our hypothesis is that it may lead to overconfident inference and decisions which in turn may lead to challenges in learning invariant features. In the domain adaptation problem, data is generally structured in a multimodal distribution. Thus, a multiple discriminator approach is compelling [30], due to its capacity to capture multiple modes of the dataset. It also leads to solving the perennial problem of mode collapse (which GANs are infamous for) as multiple discriminators now learn to distinguish classes with different modes. The diversity of an ensemble of such discriminators reduces the random errors in prediction. The performance of an ensemble model rests on the number of entities in the ensemble. However, as the number of entities increase, the model parameters and complexity will increase. This is one of the primary bottlenecks of the ensemble based methods. The number of parameters in an algorithm is a significant factor in determining model efficiency.

To tackle the above problems, we propose a novel and efficient discriminator architecture by using Monte Carlo (MC) sampling [44]. We incorporate Bernoulli dropout in a multi-adversarial network, by dropping out a certain number of neurons from our discriminator with some probability d . This gives rise to a set of dynamic discriminators for every data sample. The main idea behind our method is to construct a training regime for the feature extractor in domain adaptation that consists of increasingly challenging tasks to generate domain invariant features. This allows the sophistication of the feature extractor to gradually increase throughout training, rather than aiming for full sophistication at the outset. This method is similar to that of curriculum in supervised learning, where one orders the training examples to be presented to a learning algorithm according to some measure of difficulty [2]. Despite the conceptual similarity, the methods are quite different. Under our approach, it is not the difficulty of the training examples presented to either network, but rather the *capacity*, and hence strength, of the discriminator network that is increased as the training progresses. The idea behind the use of a curriculum based dropout discriminator is to exploit the characteristics of several independent discriminators by consolidating them in order to achieve higher performance.

We do a curriculum based learning on these dropout discriminators. As the training proceeds, the number of discriminators sampled, increase, thereby boosting the variance of our model's prediction. The proposed approach enforces the feature extractor network not to constrain the learned representations to satisfy a single discriminator, but, instead, to satisfy an ensemble of dynamic discriminators (composition is different across different discriminators). Instead of learning a point estimate (in case of MADA [30]), the feature extractor

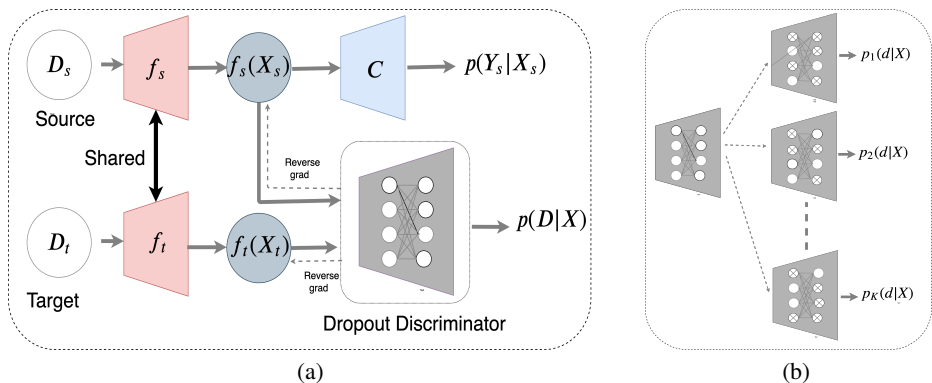


Figure 2: (a) Proposed model includes the source and target feature extractor(shared), classifier network using the fully connected layers and the dynamic ensemble of discriminators using the Bernoulli dropout network (b) Dropout based discriminator architecture

network of our proposed model learns a distribution, due to the ensemble effect of feedback from a set of dynamic discriminators. This approach leads to a more generalized feature extractor, promoting resemblance in learned representations of a class from different domains. The instinct behind incorporating dropout in our model is to warrant that neurons are not exclusively reliant on a precise set of other neurons to determine their outputs. Instead, each neuron relies on the agglomerate behavior of several other neurons, promoting generalization. By applying dropout on the discriminator, we obtain a set of entirely dynamic discriminators and hence the feature extractor cannot use the trick of relying on a specific type of discriminator or ensemble of discriminators to learn to generate representations to deceive the discriminator. Instead, it will now have to genuinely learn domain invariant representations. Thus, the feature extractor network is now guided by diverse feedback given to it by an ensemble of dynamic discriminators. All this increase in performance is obtained without compromising on the scalability and complexity front through our proposed model.

4 Proposed Adaptation Model

In the unsupervised domain adaptation problem, we consider that the *source dataset* \mathcal{D}_s has access to all its labels while there are no labels for the *target dataset* \mathcal{D}_t at the training time. We assume that \mathcal{D}_s comes from a source distribution \mathcal{S} and \mathcal{D}_t comes from a target distribution \mathcal{T} . We assume that there are N_s source data points and N_t unlabeled target data points. So $\mathcal{D}_s = (x_i^s, y_i^s)_{i=1}^{N_s} \in \mathcal{S}$ has N_s labeled examples and the target domain $\mathcal{D}_t = (x_i^t)_{i=1}^{N_t} \in \mathcal{T}$ has N_t unlabeled examples. Our underlying assumption is that both distributions are complex and unknown. Our model provides a deep neural network that enables learning of transferable feature representations $f(x)$ and an adaptive classifier $y = C(f(x))$ to reduce the shift in the joint distributions across domains, such that the target risk $Pr_{(x,y) \sim q}[C(f(x)) \neq y]$ is minimized by jointly minimizing source risk and distribution discrepancy by adversarial domain adaptation where q is assumed to be the joint distribution of target samples.

In this work, we employ a variant of GRL[14], where discriminator is modeled as an MC-dropout based ensemble. The feature extractor network consists of convolution layers to produce image embeddings. Both source and target feature extractors share the same

parameters. The classifier network consists of fully connected layers. Only source embeddings are forwarded to the classifier network to predict the class label. The classifier network parameters (θ_c) are updated only by the loss from source data samples. The discriminator receives both source and target embeddings. The parameters of the MC-dropout discriminator (θ_d) are updated with domain classification loss. The feature extractor parameters (θ_f) are updated by the gradients from the classifier network as well as by the reverse gradient of both source and target data samples from the dynamic set of the ensemble of discriminators. Detailed architecture is presented in Figure 2.

For the adaptation task, the feature extractor learns domain-invariant features with the help of MC-Dropout based discriminator. For each data sample that goes to the discriminator, we obtain the domain classification loss. These losses are backpropagated through respective Monte Carlo sampled dropout discriminators followed by gradient reversal layer. Hence, for every input, we obtain a distribution of gradients. The feature extractor is updated by a gradient from this distribution to generate domain invariant features. In a binary discriminator [□], we obtain a point estimate of the gradient for specific input. In the case of multi discriminator [□], we obtain an ensemble of the point estimates of gradients. The advantage of obtaining a distribution of gradients is that we get generalized learned representations robustly leading to domain invariant features. We propose Curriculum based Dropout Discriminator (CD^3A), where we increase the number of MC samples as training proceeds in a paradigm similar to curriculum learning. However, in the other variant (D^3A), we maintain a fixed number of MC sampled discriminators throughout the training.

4.1 Curriculum based Dropout Discriminator for Domain Adaptation (CD^3A)

In CD^3A , the distribution of gradients is obtained through a curriculum fashioned training, i.e., we increase the number of MC samples as training proceeds. The motivation behind increasing the number of MC samples is that, in the initial phase of the adaptation, the feature extractor learns the domain invariant features without considering the multi-mode structure of data. For this purpose, only a small number of discriminators is required. As the training advances, we expect the network to learn the domain invariant features along with its multi-modal structure. Thus, in the proposed model, we increase the MC samples of discriminator as training progress to obtain the domain invariant feature without losing its multi-mode structure. Given an input sample x_i , we obtain feature embedding $f(x_i)$, by passing it through a feature extractor f . These embeddings are further used to obtain the classification score $C(f(x_i))$ and the domain classification score for j^{th} samples of discriminator $D_j(f(x_i))$, where $j = 1, \dots, K$. The curriculum learning of the discriminator does not rely on the difficulty of the training examples presented to either network, but rather the capacity, and hence strength, of the discriminator that is increased throughout the training. We construct an ordered set of sets of samples of discriminator increasing in numbers. More formally the set of discriminators is $\mathcal{D} = \{\{D_1\}, \{D_1, D_2\}, \dots, \{D_1, D_2, \dots, D_K\}\}$, where D_j is a MC sampled discriminator. We can clearly see that the \mathcal{D} is a ordered set in terms of the capacity, where capacity of $\{D_1\} \subseteq \{D_1, D_2\}$.

4.2 Fixed sampling based Dropout Discriminator for Domain Adaptation (D^3A)

In this variant, we fix the number of MC sampled discriminators during the training. In this scenario, we obtain an ensemble of discriminators. We call this variant as a Dropout

Discriminator for Domain Adaptation (D³A). This modification can be considered a more efficient version of the multi discriminator model. We experimented with different sampling values (details are reported in supplementary) and obtained the best results when the number of samples is chosen close to the number of classes in the target dataset.

4.3 Loss Function

Our loss function is composed of classification loss and domain classification loss. Our classifier takes learned representations as input and predicts its label. Classification loss function \mathcal{L}_c is a cross-entropy loss. Dropout discriminator is expected to label (output) the source domain images as 0 and target domain images as 1. Domain classification loss \mathcal{L}_d is a binary cross entropy loss between the output of discriminator and the expected output. It is summed over the number of MC-sampled discriminators K . K is increased as the training proceeds in the case of CD³A model, whereas it is fixed for D³A model.

$$\mathcal{L}(\theta_f, \theta_c, \theta_d) = \frac{1}{N_s} \sum_{x_i \in \mathcal{D}_s} \mathcal{L}_c(C(f(x_i)), y_i) - \frac{\lambda}{N} \sum_{j=0}^K \sum_{x_i \in \mathcal{D}_s \cup \mathcal{D}_t} \mathcal{L}_d(D_j(f(x_i)), d_i) \quad (1)$$

$$(\hat{\theta}_f, \hat{\theta}_c) = \arg \min_{\theta_f, \theta_c} \mathcal{L}(\theta_f, \theta_c, \hat{\theta}_d) \quad \hat{\theta}_d = \arg \min_{\theta_d} \mathcal{L}(\hat{\theta}_f, \hat{\theta}_c, \theta_d) \quad (2)$$

where $d_i = 0$ if $x_i \in \mathcal{D}_s$ and $d_i = 1$ if $x_i \in \mathcal{D}_t$. The function f is the feature extractor network with shared weights for source and target data (f_s and f_t are denoted by common shared network f). λ is the trade-off parameter between the two objectives. C is the classifier network and D_j is the j^{th} MC-sampled dropout discriminator. \mathcal{D}_s and \mathcal{D}_t represent source and target domains respectively.

We generate the entities of ensemble via dropout. In contrast, previous works [50] use multiple discriminators; their number being equal to the number of classes in the dataset. It leads to an increase in the number of parameters employed in the discriminator which makes it unsuitable for datasets with a large number of classes. Also, due to our model’s parameters being significantly less, the data requirements are also quite low. This has been shown in supplementary material, where we remove half of the source data and still obtain good accuracy. Also, MADA uses the predicted label probabilities to weigh the discriminator’s response. This is a drawback as it can lead to misleading corrections of the feature extractor network in case of wrong predictions by the label predictor (classifier). Our model doesn’t have such constraints making our discriminator even more powerful leading to better learning of domain invariant features by the feature extractor network. The implementation details are provided in the supplementary material, and other details are provided on the project page ¹.

5 Results and Analysis

5.1 Datasets

Office-31 Dataset: Office-31 [54] is a benchmark dataset for domain adaptation, comprising 4,110 images in 31 classes collected from three distinct domains: Amazon (A), Webcam (W)

¹<https://delta-lab-iitk.github.io/CD3A/>

and DSLR (D). To enable unbiased evaluation, we evaluate all the methods on all 6 transfer tasks $A \rightarrow W$, $D \rightarrow A$, $W \rightarrow A$, $A \rightarrow D$, $D \rightarrow W$ and $W \rightarrow D$.

ImageCLEF Dataset: ImageCLEF-2014 dataset consists of 3 domains: Caltech-256 (C), ILSVRC 2012 (I), and Pascal-VOC 2012 (P). There are 12 common classes, and each class has 50 samples. There is a total of 600 images in each domain. We evaluate models on all 6 transfer tasks: $I \rightarrow P$, $P \rightarrow I$, $I \rightarrow C$, $C \rightarrow I$, $C \rightarrow P$, and $P \rightarrow C$.

Method	$A \rightarrow W$	$D \rightarrow W$	$W \rightarrow D$	$A \rightarrow D$	$D \rightarrow A$	$W \rightarrow A$	Average
Alexnet [17]	60.6	95.0	99.5	64.2	45.5	48.3	68.8
MMD[46]	61.0	95.0	98.5	64.9	47.2	49.4	69.3
RTN[23]	73.3	96.8	99.6	71.0	50.5	51.0	74.1
DAN[22]	68.5	96.0	99.0	66.8	50.0	49.8	71.7
GRL [10]	73.0	96.4	99.2	72.3	52.4	50.4	74.1
JAN [25]	75.2	96.6	99.6	72.8	57.5	56.3	76.3
CDAN[26]	77.9	96.9	100.0	74.6	55.1	57.5	77.0
MADA[60]	78.5	99.8	100.0	74.1	56.0	54.5	77.1
IDDA[18]	82.2	99.8	100.0	82.4	54.1	52.5	78.5
$D^3A(31)$	79.0	97.7	100.0	79.4	58.2	55.3	78.3
CD^3A	82.3	99.8	100.0	81.1	58.2	55.6	79.5

Table 1: Classification accuracy (%) on Office-31 dataset for unsupervised domain adaptation on AlexNet[17] pretrained network. Our model is CD^3A and D^3A with the number in bracket indicating the number of Monte Carlo samples

Method	$I \rightarrow P$	$P \rightarrow I$	$I \rightarrow C$	$C \rightarrow I$	$C \rightarrow P$	$P \rightarrow C$	Avg
AlexNet [17]	66.2	70.0	84.3	71.3	59.3	84.5	73.9
DAN[22]	67.3	80.5	87.7	76.0	61.6	88.4	76.9
GRL [10]	66.5	81.8	89.0	79.8	63.5	88.7	78.2
RTN[23]	67.4	82.3	89.5	78.0	63.0	90.1	78.4
MADA [60]	68.3	83.0	91.0	80.7	63.8	92.2	79.8
$D^3A(12)$	69.1	80.9	91.0	81.5	66.2	90.0	79.8
CD^3A	69.3	81.5	91.3	81.6	65.9	90.2	80.0

Table 2: Classification accuracy (%) on ImageCLEF dataset for unsupervised domain adaptation (AlexNet [17])

5.2 Results

We use pre-trained Alexnet [17] architecture following the typical setting in unsupervised domain adaption for our base model. Table 1 summarizes results on Office31 dataset, and Table 2 and Table 3 have results for the ImageClef dataset for AlexNet and ResNet networks respectively. The results on Office-Home [49] along with the implementation details are provided in the supplementary material. We obtained state-of-the-art results on all the datasets. It is noteworthy that the proposed model boosts the classification accuracies substantially on hard transfer tasks, e.g., $A \rightarrow D$, $A \rightarrow W$, etc. where the source and target domains are substantially different. On average, we obtain considerably improved accuracies and statistically significant results as shown further.

Method	I → P	P → I	I → C	C → I	C → P	P → C	Average
ResNet [12]	74.8	83.9	91.5	78.0	65.5	91.2	80.7
DAN [22]	75.0	86.2	93.3	84.1	69.8	91.3	83.3
RTN [23]	75.6	86.8	95.3	86.9	72.7	92.2	84.9
GRL [10]	75.0	86.0	96.2	87.0	74.3	91.5	85.0
JAN [24]	76.8	88.0	94.7	89.5	74.2	91.7	85.8
MADA [60]	75.0	87.9	96.0	88.8	75.2	92.2	85.8
CDAN [26]	77.2	88.3	98.3	90.7	76.7	94.0	87.5
CD³A	77.5	88.7	96.8	93.2	78.3	94.7	88.2

Table 3: Classification accuracy (%) on *ImageCLEF* dataset for unsupervised domain adaptation (ResNet-50 [12])




Figure 3: (a) and (b) figures show t-SNE visualizations of the CNN’s activation (a) in case when adapted through[10] and (b) when adapted through proposed model. Blue and red points correspond to the source domain(A), and the target domain (W) respectively. Sub figures (c) and (d) show the proxy distance for A→D and A→W.

5.3 Analysis

Curriculum v/s Fixed sampling: We have plotted the accuracy as a function of the number of MC samples for both the models, curriculum-based sampling (CD³A) and fixed sampling (D³A) in Figure 4 (a). We can clearly observe that in the case of D³A, the performance increases as we increase the number of MC sampled discriminators, but after some samples, the performance starts to deteriorate. While in case of CD³A, the performance of model saturates after certain epochs. We can also see that CD³A outperforms D³A.

Model complexity comparison with MADA: The proposed CD³A model uses one discriminator(ensemble using dropout) whereas MADA uses as many discriminators as are the number of classes. Therefore, CD³A has very few parameters as compared to MADA even for datasets with a small number of classes. For instance, in case of Office-31 dataset, MADA has 31 discriminators compared to CD³A, which has only one discriminator. MADA has ~98M parameters, while CD³A has ~59M parameters for Office-31 dataset. If we further increase the class size, the number of parameters in MADA increases(by ~1.3M for every class label), but CD³A will have constant number of parameters (~59M).

Feature visualization: The adaptability of target to source features can be visualized using the t-SNE embeddings of image features. We follow similar setting as in [10] to plot t-SNE embeddings for A→W adaptation task in Figure 3 (a) and (b). From the plot, we observe that adapted features(CD³A) are more domain invariant than the features adapted with GRL.

Statistical significance analysis: We analyzed statistical significance [8] for our CD³A model against GRL[10] and source only method for the domain adaptation task. The Criti-

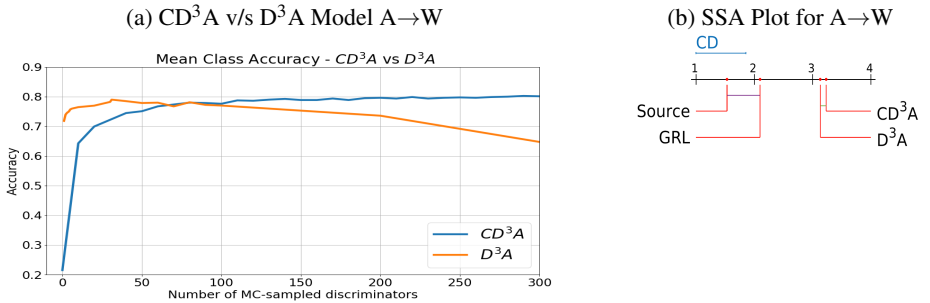


Figure 4: (a) Accuracy v/s Number of MC samples for D^3A and CD^3A model. Note that in D^3A model, each model is trained separately and reported accuracy after the training, while in CD^3A model the accuracy is calculated with single training process (b) Analysis of statistically significant difference for $A \rightarrow W$ in Binary label Discriminator (GRL) [17], proposed model and Source only methods, with a significance level of 0.05.

cal Difference (CD) for Nemenyi test depends upon the given confidence level (0.05 in our case) for average ranks and number of tested datasets. If the difference in the rank of the two methods lies within CD (our case $CD = 0.6051$), then they are not significantly different. Figure 4(b) visualizes the posthoc analysis using the CD diagram for $A \rightarrow W$. From the figures, it is clear that our CD^3A model is better and significantly different from other methods.

Proxy- \mathcal{A} -Distance \mathcal{A} -distance as a measure of cross domain discrepancy [17], which, together with the source risk, will bound the target risk. The proxy \mathcal{A} -distance is defined as $d_{\mathcal{A}} = 2(1 - 2\varepsilon)$, where ε is the generalization error of a classifier (e.g. kernel SVM) trained on the binary task of discriminating source and target. Figure 3(c) and (d) shows $d_{\mathcal{A}}$ on tasks $A \rightarrow D$ and $A \rightarrow W$, with features of source only model [17], GRL [17], MADA [30] and proposed model CD^3A . We observe that $d_{\mathcal{A}}$ calculated using CD^3A model features is much smaller than calculated using source only model, GRL and MADA features, which suggests that representations learned via CD^3A can reduce the cross-domain gap more effectively.

6 Conclusion

In this paper, we provide a simple approach to obtain an improved discriminator for adversarial domain adaptation. We specifically show that the use of sampling-based ensemble results in an improved discriminator without increasing the number of parameters. The main reason for this improvement is that the features are made domain invariant based on a distribution of observations as against a single point estimate. Our approach based on curriculum dropout suggests that we are able to obtain an improved discriminator that is stable and improves the feature invariance learnt. We compare our method with standard baselines and provide a thorough empirical analysis of the method. We further observe through visualization that domain adapted features do result in domain invariant feature representations. Using the discriminator obtained through curriculum based dropout to solve domain adaptation is a promising direction, which we have initiated through this work.

Acknowledgment: We acknowledge the resource support from Delta Lab, IIT Kanpur. Vinod Kurmi acknowledges support from TCS Research Scholarship Program.

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010.
- [2] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48. ACM, 2009.
- [3] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, page 7, 2017.
- [4] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, page 7, 2017.
- [5] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a " siamese " time delay neural network. In *Advances in Neural Information Processing Systems*, pages 737–744, 1994.
- [6] Qingchao Chen, Yang Liu, Zhaowen Wang, Ian Wassell, and Kevin Chetty. Re-weighted adversarial adaptation network for unsupervised domain adaptation. 2018.
- [7] Yunjey Choi, Minje Choi, and Munyoung Kim. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. 2017.
- [8] Janez Demšar. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine learning research*, 7(Jan):1–30, 2006.
- [9] Ishan Durugkar, Ian Gemp, and Sridhar Mahadevan. Generative multi-adversarial networks. *arXiv preprint arXiv:1611.01673*, 2016.
- [10] Yarín Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [11] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning*, pages 1180–1189, 2015.
- [12] Arnab Ghosh, Viveka Kulharia, Vinay Namboodiri, Philip HS Torr, and Puneet K Dokania. Multi-agent diverse generative adversarial networks. *CoRR, abs/1704.02906*, 6:7, 2017.
- [13] Kazuyuki Hara, Daisuke Saitoh, and Hayaru Shouno. Analysis of dropout learning regarded as ensemble learning. In *International Conference on Artificial Neural Networks*, pages 72–79. Springer, 2016.

- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [15] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213*, 2017.
- [16] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alyosha Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. 2018.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 2012.
- [18] Vinod Kumar Kurmi and Vinay P Namboodiri. Looking back at labels: A class based domain adaptation technique. *arXiv preprint arXiv:1904.01341*, 2019.
- [19] Vinod Kumar Kurmi, Shanu Kumar, and Vinay P Namboodiri. Attending to discriminative certainty for domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 491–500, 2019.
- [20] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, pages 6402–6413, 2017.
- [21] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.(CVPR)*, 2018.
- [22] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International Conference on Machine Learning*, pages 97–105, 2015.
- [23] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems*, pages 136–144, 2016.
- [24] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *International Conference on Machine Learning*, pages 2208–2217, 2017.
- [25] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *International Conference on Machine Learning*, pages 2208–2217, 2017.
- [26] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pages 1640–1650, 2018.

- [27] Pietro Morerio, Jacopo Cavazza, Riccardo Volpi, René Vidal, and Vittorio Murino. Curriculum dropout. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3544–3552, 2017.
- [28] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [29] Tu Nguyen, Trung Le, Hung Vu, and Dinh Phung. Dual discriminator generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2670–2680, 2017.
- [30] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. In *Advances in Artificial Intelligence (AAAI)*, 2018.
- [31] Artem Rozantsev, Mathieu Salzmann, and Pascal Fua. Beyond sharing weights for deep domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [32] 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. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- [33] Yunus Saatci and Andrew G Wilson. Bayesian gan. In *Advances in neural information processing systems*, pages 3622–3631, 2017.
- [34] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European conference on computer vision*, pages 213–226. Springer, 2010.
- [35] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial dropout regularization. 2018.
- [36] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3723–3732, 2018.
- [37] Rishi Sharma, Shane Barratt, Stefano Ermon, and Vijay Pande. Improved training with curriculum gans. *arXiv preprint arXiv:1807.09295*, 2018.
- [38] Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Adversarial representation learning for domain adaptation. *arXiv preprint arXiv:1707.01217*, 2017.
- [39] Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. In *AAAI*, 2018.
- [40] Yang Shu, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Transferable curriculum for weakly-supervised domain adaptation. 2019.
- [41] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

- [42] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European Conference on Computer Vision*, pages 443–450. Springer, 2016.
- [43] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *AAAI*, volume 6, page 8, 2016.
- [44] Baochen Sun, Jiashi Feng, and Kate Saenko. Correlation alignment for unsupervised domain adaptation. *Domain Adaptation in Computer Vision Applications*, page 153, 2017.
- [45] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1521–1528. IEEE, 2011.
- [46] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.
- [47] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pages 4068–4076. IEEE, 2015.
- [48] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Computer Vision and Pattern Recognition (CVPR)*, volume 1, page 4, 2017.
- [49] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proc. CVPR*, pages 5018–5027, 2017.
- [50] Daphna Weinshall, Gad Cohen, and Dan Amir. Curriculum learning by transfer learning: Theory and experiments with deep networks. In *International Conference on Machine Learning*, pages 5235–5243, 2018.
- [51] Jing Zhang, Zewei Ding, Wanqing Li, and Philip Ogunbona. Importance weighted adversarial nets for partial domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8156–8164, 2018.
- [52] Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2020–2030, 2017.
- [53] Zhen Zhang, Mianzhi Wang, Yan Huang, and Arye Nehorai. 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*, pages 3437–3445, 2018.