

Joint Progressive Knowledge Distillation and Unsupervised Domain Adaptation

Le Thanh Nguyen-Meidine, Eric Granger, Madhu Kiran, Jose Dolz
Laboratoire d'imagerie, de vision et d'intelligence artificielle (LIVIA)
Dept. of Systems Engineering, École de technologie supérieure
Montreal, Canada
{le-thanh.nguyen-meidine.1, madhu.kiran.1}@ens.etsmtl.ca
{eric.granger, jose.dolz}@etsmt.ca

Louis-Antoine Blais-Morin
Genetec Inc.
Montreal, Canada
lablaismorin@genetec.com

Abstract—Currently, the divergence in distributions of design and operational data, and large computational complexity are limiting factors in the adoption of CNNs in real-world applications. For instance, person re-identification systems typically rely on a distributed set of cameras, where each camera has different capture conditions. This can translate to a considerable shift between source (e.g. lab setting) and target (e.g. operational camera) domains. Given the cost of annotating image data captured for fine-tuning in each target domain, unsupervised domain adaptation (UDA) has become a popular approach to adapt CNNs. Moreover, state-of-the-art deep learning models that provide a high level of accuracy often rely on architectures that are too complex for real-time applications. Although several compression and UDA approaches have recently been proposed to overcome these limitations, they do not allow optimizing a CNN to simultaneously address both. In this paper, we propose an unexplored direction – the joint optimization of CNNs to provide a compressed model that is adapted to perform well for a given target domain. In particular, the proposed approach performs unsupervised knowledge distillation (KD) from a complex teacher model to a compact student model, by leveraging both source and target data. It also improves upon existing UDA techniques by progressively teaching the student about domain-invariant features, instead of directly adapting a compact model on target domain data. Our method is compared against state-of-the-art compression and UDA techniques, using two popular classification datasets for UDA – Office31 and ImageClef-DA. In both datasets, results indicate that our method can achieve the highest level of accuracy while requiring a comparable or lower time complexity.

Index Terms—Deep Learning, Convolutional Neural Networks, Domain Adaptation, Knowledge Distillation, Visual Recognition.

I. INTRODUCTION

Deep learning (DL) models, and in particular convolutional neural networks (CNNs) can achieve state-of-the-art performance in a wide range of visual recognition applications, such as classification, object detection, and semantic segmentation [1]–[3]. In practice, a main drawback associated with these models is their scalability and computationally complexity, which poses a challenge for many real-time applications, as found in video analytics and surveillance [4].

Currently, the compact DL models that can provide fast inference generally lack the high accuracy of deeper, more complex models. One alternative to overcome this issue is

to compress complex high accuracy models into smaller or simpler models while preserving the same level of accuracy. Several approaches have recently been proposed to accelerate and compress CNNs, include quantization, low-rank approximation, knowledge distillation, compact network design and network pruning. For instance, network compression methods for channel pruning [5]–[10] and knowledge distillation (KD) [11]–[14] have become very popular due to the exponential increase in the complexity of architectures, which may have millions of parameters.

Another limitation is the poor generalization of CNNs across domains, particularly when there is a considerable domain shift between source and target data distributions. This is the case for instance of applications like video-surveillance over a distributed network of camera, where variations in camera viewpoint and capture conditions (e.g., illumination, occlusion and background) introduce a shift between data from source and target domains. The accuracy of CNNs degrades when there is a considerable divergence between the data capture conditions in the model development and operational environments. To alleviate this problem, domain adaptation techniques are commonly proposed, either in a supervised or unsupervised setting [15], [16]. For applications in video-surveillance, it is, however, costly to collect and annotate videos from each camera viewpoint and capture condition to fine-tune a CNN.

In this paper, we focus on DL models for unsupervised domain adaptation (UDA) to allow adapting CNN embeddings based on unlabeled data. The main body of literature on UDA techniques focuses on learning domain invariant features by using adversarial loss [17], [18] to encourage domain confusion, or by minimizing a distance or discrepancy between two the data distributions [19], or both [20]. Another popular paradigm is to learn a mapping between source and target images such that images captured in different domains have a similar appearance. This reconstruction-based approach mimics standard supervised learning [21], [22].

While compression and UDA techniques can respectively provide CNNs with a high level of performance and efficiency, research that explores the benefits of joint model compression and UDA remains scarce. For example, recent research has

combined KD and UDA to improve the performance in a UDA context [23], [24]. Despite the improvement shown in terms of the DA task, they neglect the reduction of model complexity. Another important problem when employing KD is the potentially large gap between the capacity of teacher and student models to learn complex mappings, which can degrade student performance. Distilling knowledge directly into the student model from a trained teacher model may present some challenges. For example, in the early stages of training, the decision boundaries of the student and teacher modes may differ considerably. To bridge this gap some researchers have proposed to either add a teaching assistant [14] or distilling knowledge at several intermediate layers of a CNN [13].

In a standard training scenario, the CNN model would be adapted to a new target after having been compressed, although this can reduce its capacity to generalize on target data since over-parametrization is often important for generalization [25]. This can be overcome by learning from a teacher instead of learning directly through the UDA loss. In contrast, another scenario involves adapting a complex CNN model to the target domain, and then compressing it, although unsupervised KD remains a challenge since ground truth data is needed to assure that the student model to remain consistent with the dataset. To address the aforementioned limitations, we argue that by adapting the teacher to the target domain while the student is being trained, the student can adapt progressively to the domain instead of learning directly from a previously targeted domain. The student would thereby learn the steps needed to adapt itself to the target domain, instead of learning directly from a model adapted to the target, i.e. adapted teacher. By jointly exploiting unsupervised KD and DA, it is possible to overcome the lack of ground truth data for KD, as needed to train the student model that is consistent with source features.

Unlike recent work in literature, this paper provides the first attempt to simultaneously address two key problems with CNNs – domain shift and model complexity – through joint progressive KD and UDA. Particularly, the proposed approach learns a compact model with a feature embedding that can provide a high level of accuracy in the target domain. It leverages progressive KD to adapt the student model in a step-by-step manner, using knowledge from the teacher to learn domain invariant features of the source and target domains. In order to ensure the validity of student model w.r.t the UDA loss and the target domain, we introduce a consistency loss that ensures consistency on the student model by learning source domain features from the teacher. We validate our approach with different training scenarios of KD and UDA: (1) UDA then KD, (2) KD then UDA and (3) UDA directly on compact model. Empirical evaluations show that our joint progressive KD and UDA approach facilitates domain adaptation and compression of deep CNNs, and can outperform representative state-of-the-art approaches on the widely used Office31 and ImageClef-DA benchmark datasets. In particular, our approach is general and model-agnostic, and can be combined with different UDA and KD techniques to improve performance.

II. RELATED WORK

A. Compression techniques:

Approaches for compressing CNNs can be mainly categorized in: (1) pruning techniques [5], [7], [9], [10], [26], (2) quantization [27]–[29], (3) decomposition [30], [31], and (4) KD [11]–[14]. Pruning techniques focus on removing non-useful weights or filters in order to reduce the computational complexity. Quantization techniques focus on reducing the representation of weights into lower precision since, for example, 8-bit integer precision provides much faster computation than the floating point computation. Decomposition techniques provide faster computation by decomposing tensor in lower rank approximation as vectors products. Lastly, KD techniques transfer knowledge from a teacher (usually a large model) to a student (a smaller model). Since we will employ this technique to reduce the complexity of the model, we will focus on KD onward. For a comprehensive survey on compression techniques for CNNs, we refer the reader to [32].

There exist several ways of distilling knowledge from a teacher to a student. A well-known technique is to employ the teacher output as the soft label for the student [11]. In this work [11], the temperature value was employed to generate softer versions of the teacher outputs. Another popular solution is to minimize the features differences at intermediate layers between the teacher and student network in order to maximize the information transfer between the teacher and the model [13]. Feature similarity can be enforced by minimizing a partial L2 distance, which is equivalent to L2 norm except that if the value of the student is smaller than the teacher and both are negative then the result is zero, between student and teacher after using a Margin ReLU (use of a margin m instead of 0) which can take in negative values of a feature map. To solve the gap issue between a converged teacher model and a student some other researchers have proposed the integration of a teaching assistant [14]. In this work, an intermediate model with a lesser gap is chosen, this model is then used as a teacher to a smaller model, by progressively reducing the gap with an intermediate teaching assistant, thus, limit the performance degradation of the student model.

B. Unsupervised domain adaptation:

Unsupervised Domain Adaptation (UDA) techniques try to adapt models when a domain shift between source and target dataset exists, where only the source data is labeled. Current main UDA techniques [15] include: finding domain-invariant features [19], domain mapping [33], [34], ensemble learning [35], statistic normalization [36] and target discriminate methods [37]. The first category learns domain invariant features either by domain confusion [17], [18] or minimizing a distance between distribution [19]. Domain confusion can be achieved by employing a domain classifier (or discriminator) [17], [18]. While [17] employs a gradient reversal layer in order to maximize the domain classification loss, the work in [18] uses an adversarial loss on the discriminator. Domain mapping focuses on finding a mapping either from the source

domain to the target domain or vice-versa [33], [34]. Currently, most of the domain mapping based techniques rely on generative adversarial networks (GAN). Hoffman et al. [33] propose to use a pair of discriminator and generator in order to map a source image into the target domain distribution. This mapping is learned at the same time as a task-specific loss (i.e. classification loss) on both transformed image and non transformed-images and the overall optimization is done alternatively between generator and discriminator-task. The paper in [38] goes further by integrating adaptation at feature-level. Ensemble methods use either multiple models or the same model at different times (typically referred to as self-ensembling) in order to produce more reliable pseudo-labels on unlabelled data [35]. Others methods like statistic normalization assume that the task knowledge is learned and the only adaptation needs to be done is on the batch norm statistics [36]. Last, target discriminate methods work with the assumption that data points are distributed in separate clusters and the decision boundary lies in lower density regions. Thus, these methods work by trying to push the decision boundary to lower density regions by adding adversarial losses [37].

C. Joint unsupervised domain adaptation and knowledge distillation:

Even though jointly exploiting UDA and KD for compression and domain shift problems remains unexplored, there have been few attempts to combine these two techniques in the context of domain adaptation [39], [40]. In [39], the authors propose to use multiple teachers for teaching a single student in order to increase the performance of their model in the context of sentiment analysis. The work in [40] proposes to combine KD and DA for the task of white matter hyperintensities segmentation in magnetic resonance imaging by training the teacher model on the source and trying to minimize the cross-entropy loss between the probability maps of the teacher and the student on the target dataset. Nevertheless, these approaches do not address the problem of reducing model complexity. To the best of our knowledge, there exist only one approach that combines compression and UDA techniques, referred to as TCP [41]. Interaction between the two techniques is done in several steps. First, the model is trained to be adapted to the target domain by minimizing the domain divergence using the maximum mean discrepancy (MMD) [19]. Then, the least important filters are pruned by using a gradient based criterion and the model is continuously refined on the domain adaptation loss. An important limitation of this technique is the need to have an already domain adapted model in order to start the pruning, whereas our technique can directly start from a non-adapted model.

III. PROPOSED METHOD

The main pipeline of the proposed method is depicted in Figure 1. Our method performs domain adaptation of a teacher model by learning domain invariant features between the source and target domains. Meanwhile, it progressively distills its knowledge to a student model on both source and

target features. As shown in Figure 1, DA is performed on the features of the teacher network, while the KD from teacher to student is performed on the result of a temperature-based soft-max on the logits (output of a fully connected layer). Additional details of the proposed UDA and KD techniques are described in the following subsections.

A. Unsupervised domain adaptation:

We start by defining the UDA loss for teacher model, which is based on MMD [19], [41]¹:

$$\mathcal{L}_{MMD} = \left\| \frac{1}{N_s} \sum_{x_i \in D_s^L} \phi_T(x_i) - \frac{1}{N_t} \sum_{x_j \in D_t^U} \phi_T(x_j) \right\|_{\mathcal{H}}^2 \quad (1)$$

With D_s^L the labeled source domain dataset which contains N_s samples and labels, D_t^U the unlabeled target domain dataset of N_t data samples, ϕ_T the teacher feature extractor function that maps an input to a feature map, \mathcal{H} the Reproducing Kernel Hilbert Space (RKHS) with gaussian kernel. As in [19], [41], we incorporate this a supervised loss on the source domain in order to have the final UDA loss for the teacher:

$$\mathcal{L}_{TDA} = \mathcal{L}_{MMD} + \gamma \mathcal{L}_{CE}(T(D_s^L, 1), y_s) \quad (2)$$

\mathcal{L}_{CE} the supervised cross-entropy loss of the teacher model on the source domain, γ a trade-off hyper-parameter that follows the same variations as [41] and T the function that maps an input to the output of the teacher network with a soft-max of temperature 1 (i.e. the regular soft-max).

B. Knowledge distillation for domain knowledge transfer:

The next step is to transfer the target domain knowledge from the teacher to the student, we use a modified version of the KD loss from the work of Hinton:

$$\mathcal{L}_{TKD} = \mathcal{L}_{distill}(S(D_t^U, \tau), T(D_t^U, \tau)) \quad (3)$$

In this equation, S and T represent respectively the output of student network and teacher network with a soft-max based on a temperature τ in order to soften the output and $\mathcal{L}_{distill}$ is a KL divergence loss in our case but can be replaced with a mean squared loss or cross-entropy. This loss differs from the original paper [11] because, we had to remove the cross-entropy loss between the student model output and the ground truth since we are working on UDA. Trivially, this should be enough for joint KD and UDA since we only want to have target domain knowledge, in order to ensure the consistency of the model w.r.t to a common representation, we proposed to add a consistency loss to ensure that the student model can learn a better common representation from source and target domains by distilling on the source data.

$$\mathcal{L}_{SKD} = \mathcal{L}_{distill}(S(D_s^L, \tau), T(D_s^L, \tau)) + \alpha \mathcal{L}_{CE}(S(D_s^L, 1), y_s) \quad (4)$$

¹Note that it can be generalized to most other UDA techniques.




Fig. 1. Illustration of the proposed learning technique for progressive KD with UDA.

Eq. 4 is the student KD loss, with hyper-parameter α to balance between the KD and the cross entropy loss of the output of the student model and the ground truth on the source domain. The Figure 1 illustrates all these losses and also the proposed techniques. The final loss of our models, is then:

$$\mathcal{L} = (1 - \beta)\mathcal{L}_{TDA} + \beta(\mathcal{L}_{TKD} + \mathcal{L}_{SKD}) \quad (5)$$

We added the β hyper-parameter in order to balance out the importance between UDA and KD. Since we are performing jointly KD and DA, in the beginning, the teacher would still be learning from the DA. This means that there is not much to be learned for the student model, besides the source representation which can be learned from the KD loss. In light of this, we propose to start by giving more importance to UDA in the beginning and gradually transfer the importance to KD basing β on an exponential growth function between $[b, f]$, with b the starting value of β and f the end value. In order to define as exponential growth, we need to calculate a growth rate based on b and f :

$$g = \frac{\log(\frac{f}{b})}{epochs} \quad (6)$$

With $epochs$ the number of epochs and g the growth rate. Once we have the growth rate, β at epoch t can be found as:

$$\beta_t = b * e^{gt} \quad (7)$$

The overall algorithm is described in Algorithm 1 using an alternate optimization scheme of our algorithm. The details of this implementation can be found in Section IV(C).

IV. EXPERIMENTAL METHODOLOGY

In this section, we detail the experimental methodology employed to validate the proposed method. First, we describe

Algorithm 1: KD-UDA

input : A teacher model M_T , a student model M_S , a source dataset D_s^{sup} , a target dataset D_t^U

output : A target adapted student model

- 1 **for** $epoch \leftarrow 1$ **to** $epochs$ **do**
- 2 **for** $data_s$ in D_s^L and $data_t$ D_t^U **do**
- 3 Obtain the feature map F_s of M_t on $data_s$ and F_t on $data_t$ using ϕ_T
- 4 Optimize the teacher model M_T with $(1 - \beta)\mathcal{L}_{TDA}$ using F_s and F_t
- 5 Obtain the logits o_s^T of M_T on $data_s$ and o_t^T of M_T on $data_t$
- 6 Obtain the logits o_s^S of M_S on $data_s$ and o_t^S of M_S on $data_t$
- 7 Applying soft-max of temperature τ on the logits and optimize $\beta(\mathcal{L}_{TKD} + \mathcal{L}_{SKD})$
- 8 Update β following Eq.7
- 9 **end**
- 10 Evaluate the model
- 11 **end**

the datasets and baselines methods. Then, we provide implementation details to facilitate the reproducibility of the reported results.

A. Datasets:

a) *Office31*: : This dataset contains three subsets of dataset which are Webcam (W), DSLR (D) and Amazon (A) with 31 classes. These subsets contains images from amazon.com (A) or office environment with changes in lighting,

poses using a DSLR camera (D) or a webcam (W). We evaluate our results based on six scenarios: $A \rightarrow W$, $W \rightarrow A$, $D \rightarrow W$, $W \rightarrow D$, $D \rightarrow A$, $A \rightarrow D$.



Fig. 2. Examples of images selected from the Office31 (top) and ImageClef-DA (bottom) datasets.

b) ImageClef-DA: This dataset for UDA contains four subsets which are taken from Imagenet (I), Pascal-Voc (P), Caltech (C) and Bing (B). Each of these subsets contains a total of 600 images for 12 classes. For this dataset, we compare to others techniques using six scenarios: $I \rightarrow P$, $P \rightarrow I$, $I \rightarrow C$, $C \rightarrow I$, $C \rightarrow P$, $P \rightarrow C$. For the evaluation of our algorithm, we chose to use a popular backbone architecture, ResNet50, with our algorithm.

B. Baselines methods:

In order to validate our joint progressive method, we propose to evaluate and compare our method with 3 baseline scenarios:

- 1) UDA, and then KD,
- 2) KD, and then UDA, and
- 3) UDA directly on compact model.

With the first baseline, $UDA \rightarrow KD$, we start with a model of ResNet50 for UDA, then we start KD on this model with a modified version (Eq.3) of [11] since most KD algorithms does not handle unsupervised KD. The student model used for this first baseline is similar to the one we chose for our methods, i.e. ResNet18 and ResNet34. For the second baseline $KD \rightarrow UDA$, we start by training a teacher model, ResNet50, on the source dataset then we apply KD with this teacher using labeled data of the source dataset and a student of one of target model, ResNet18 or ResNet34, and we finish with UDA on this student model. As for the last baseline, we take a student model, ResNet18 and ResNet34, and directly apply UDA algorithm on it [19], [41].

Another aspect of our method is knowledge distillation which is meant for model reduction and acceleration. In order to validate our findings, we also measure the difference in terms of FLOPs and parameters between our method and TCP. Since our models are predetermined architecture, we measure FLOPS and the number of parameters on ResNet50, ResNet34 and ResNet18, as for TCP, we report the number of FLOPS that was reported in the original paper.

C. Implementation details:

In this paper, we consider two ways of implementing the optimization process. The first is an end-to-end training by optimizing the loss of Eq.5 with a unique single optimizer. While this approach works and provide reasonable performance, we found that having two optimizers allows more flexibility in terms of scheduling the learning rate since there is can be a difference between optimizing domain adaptation and KD, i.e. different learning rates for each loss or we do not want to update the teacher during KD. Thus, we implement a second alternative approach where we use one optimizer for $(1 - \beta)\mathcal{L}_{TDA}$ with a different learning rate schedule than another optimizer used for $\beta(\mathcal{L}_{TKD} + L_{SKD})$. We would like to point out that it's also possible to implement the second approach with one single optimizer by handling the learning rate scheduling ourselves. For our experiments on Office31 and ImageClef-DA, we use two student models architecture, which are: ResNet34 (12% FLOPS reduction from ResNet50) and ResNet18 (56% FLOPS reduction from ResNet50) which are closely similar to the FLOPS reduction of TCP [41].

In these experiments, the images are cropped to a fixed resolution of 224×224 . Regarding the hyper-parameters, we use a starting β value of 0.1 and an end value of 0.9. For the KD hyper-parameters, we use a temperature $\tau = 20$ and $\alpha = 0.8$. Overall, we use a learning rate starting at 0.001 for UDA optimizer in Office31 and 0.0001 for ImageClef-DA,

TABLE I
ACCURACY OF PROPOSED AND BASELINE METHODS ON OFFICE31 DATASET WHEN RESNET34 IS THE DESIRED MODEL.

Training methods	A \rightarrow W	W \rightarrow A	D \rightarrow W	W \rightarrow D	D \rightarrow A	A \rightarrow D	Average
Baseline 1: UDA \rightarrow KD from ResNet50	25.4	7.1	28.5	50.0	9.7	30.7	25.2
Baseline 2: KD \rightarrow UDA from ResNet50	75.7	61.2	97.8	99.7	59.6	81.1	79.1
Baseline 3: UDA only on ResNet34	67.2	52.3	93.6	96.6	52.2	71.6	72.2
TCP: 12% pruned from ResNet50	81.8	55.5	98.2	99.8	50	77.9	77.2
Ours: ResNet34 from ResNet50	85.7	62.3	97.1	100	61.8	82.1	81.5
Ours: ResNet34 from ResNet101	87.5	62.9	98.1	100	60.8	85.7	82.5

TABLE II
ACCURACY OF PROPOSED AND BASELINE METHODS ON OFFICE31 DATASET WHEN RESNET18 IS THE DESIRED MODEL

Training methods	A \rightarrow W	W \rightarrow A	D \rightarrow W	W \rightarrow D	D \rightarrow A	A \rightarrow D	Average
Baseline 1: UDA \rightarrow KD from ResNet50	28.8	5.8	33.7	51.1	7.8	25.1	25.3
Baseline 2: KD \rightarrow UDA from ResNet50	69.0	57.3	96.2	100	56.3	73.6	75.4
Baseline 3: UDA only on ResNet18	60.2	49.2	93.7	97.7	47.6	66.4	69.1
TCP: 45% pruned from ResNet50	77.4	46.3	96.3	100	36.1	72.0	71.3
Ours: ResNet18 from ResNet50	78.9	56.8	93.8	100	56.0	81.7	77.8
Ours: ResNet18 from ResNet101	79.2	58.1	94.2	100	57.2	79.9	78.1

0.001 for KD optimizer for both datasets, a momentum of 0.9 and 400 epochs.

Our implementation can be found on-line at: https://github.com/LIVIAETS/KD_UDA

V. RESULTS AND DISCUSSION

A. Results on Office31:

For the Office31 dataset, our results outperform most of the baseline and the current existing techniques. From Tables I and II, we see that, the third baseline "UDA only" performs better than UDA \rightarrow KD. This can be explained by the fact that there is no label for target dataset and the distillation loss alone is not sufficient, as there exist a need for supervising the cross-entropy loss. The result of the second baseline is better than the rest of the baselines and TCP, which can be explained because the student model resulting from KD is already trained on a labeled source data distribution. Finally, our techniques using both teacher ResNet50 and ResNet101 perform better than TCP. In both ResNet34 and ResNet18 settings, the difference between the average of our results and TCP is considerable. This shows that combining UDA and KD in a progressive setting brings additional benefits. Between ResNet34 with teacher ResNet50 and ResNet101, there exists a slight difference. We notice that the student model with a larger teacher model performs slightly better. This is expected, since ResNet101 has a better generalization capability than ResNet50. Lastly, our technique also seems to improve UDA since our methods perform better than the second baseline, which only performs UDA on a compact model.

B. Results on ImageClef-DA:

For this dataset, the results obtained by our methods are shown in Table III and Table IV, which demonstrate that the proposed technique outperforms both the baselines and

current existing techniques. In contrast, the results obtained with ResNet34 are closer to TCP and the baseline than in the previous dataset, whereas our result with ResNet18 shows a bigger gap between the baselines and the proposed techniques. This can be first explained by the fact that ImageClef-DA is a better balanced dataset where each subset has the same number of samples. Secondly, the third baseline of "UDA only on a ResNet18" already performs better than TCP, which may explain why our methods have a better performance. In this table, the difference between teacher models is closer, this shows that a bigger teacher may not help in improving performance since the learning bottleneck is now on the student model.

C. Computational Complexity:

Comparison in terms of complexity is depicted in Table V. While pruned TCP [41] models have fewer parameters than our student models, we achieve the same number of FLOPS on ResNet34, and fewer FLOPs on ResNet18. This means that while TCP prunes more parameters, it may not have a lot of impact on the number of FLOPS since the pruned filters are ranked and pruned globally across the network instead of being pruned at each layer. This also means that TCP prunes away filters that do not impact the FLOPS, but that can impact performance, which may hamper the global objective. Another important point of having more parameters is that, over-parametrization can help generalization, increasing the chances of our student model to have a better generalization than a pruned model with less parameters.

D. Comparison over larger teacher and student:

In this section, we select the scenario of $I \rightarrow C$ and increase the teacher model to ResNet152. Then, we apply our algorithm on this model with several student models going

TABLE III
ACCURACY OF PROPOSED AND BASELINE METHODS ON IMAGECLEF-DA DATASET WHEN RESNET34 IS THE DESIRED MODEL

Training methods	I → P	P → I	I → C	C → I	C → P	P → C	Average
Baseline 1: UDA → KD from ResNet50	48.0	41.0	46.0	39.6	38.8	39.0	42.0
Baseline 2: KD → UDA from ResNet50	76.6	87.3	92.0	80.0	65.6	90.3	81.9
Baseline 3: UDA only on ResNet34	73.3	86.3	92.6	79.3	65.8	87.5	80.8
TCP: 12% pruned from ResNet50	75.0	82.6	92.5	80.8	66.2	86.5	80.6
Ours: ResNet34 from ResNet50	75.6	89.0	92.6	83.8	66.5	92.8	83.3
Ours: ResNet34 from ResNet101	75.0	87.6	93.3	83.5	66.6	91.8	83.2

TABLE IV
ACCURACY OF PROPOSED AND BASELINE METHODS ON IMAGECLEF-DA DATASET WHEN RESNET18 IS THE DESIRED MODEL

Training methods	I → P	P → I	I → C	C → I	C → P	P → C	Average
Baseline 1: UDA → KD from ResNet50	45.1	41.8	42.5	43.1	43.3	34.5	41.7
Baseline 2: KD → UDA from ResNet50	72.1	86.3	91.8	74.6	61.8	90.6	79.5
Baseline 3: UDA only on ResNet18	70.6	83.8	86.1	75.3	62.0	89.1	77.8
TCP: 45% pruned from ResNet50	67.8	77.5	88.6	71.6	57.7	79.5	73.7
Ours: ResNet18 from ResNet50	73.1	88.0	92.1	77.3	65.6	91.0	81.1
Ours: ResNet18 from ResNet101	73.5	88.6	92.3	76.8	64.1	91.1	81.0

TABLE V
COMPUTATIONAL COMPLEXITY OF PROPOSED AND TCP NETWORKS.

Models	no. operations (GFLOPS)	no. parameters (M)	
		Office31	ImageClef
Teacher/Original: ResNet50	4.1	25.5	25.5
TCP:			
- 12% pruned from ResNet50	3.6	15.8	15.9
- 45% pruned from ResNet50	2.2	10.6	10.9
Student: from ResNet50			
- ResNet34	3.6	21.7	21.7
- ResNet18	1.8	11.1	11.1

from ResNet18 to Resnet101. We compare the results from these experiments with our previous ones, along with the result of performing UDA using only Eq.1.

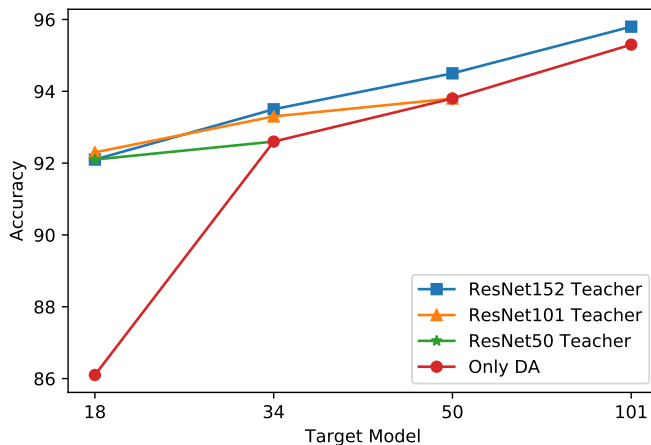


Fig. 3. Comparison of different teacher and student models.

From the Fig.3, we see that, having a larger teacher only has

a slight performance increase compared to the smaller teacher. Furthermore, we also noticed that there is no degradation in terms of performance when the teacher is ResNet152 and the student is ResNet18. Finally, having a ResNet50 as a student model definitely helps increasing the performance, which makes sense since we are using a much larger model. Overall, while having a larger teacher does not seem to have a big impact on accuracy, it can negatively impact the training time and increases the risk of overfitting.

VI. CONCLUSION

In this paper, we proposed a combination of KD and UDA that remains unexplored in literature and tackles both the problem of domain shift and model complexity. Our results suggest that the proposed method is capable of obtaining a compressed model adapted to an unsupervised target domain that performs better than state-of-the-art method and current baselines. Additionally, our progressive technique is capable of having a big gap between the student and the teacher without suffering having a performance degradation on the student. The proposed technique is generic and should be able to work with most of current UDA and KD techniques, in future works, we will evaluate our method on other KD and UDA techniques.

ACKNOWLEDGEMENT

This research was partially supported by the Mathematics of Information Technology and Complex Systems (MITACS) and the Natural Sciences and Engineering Research Council of Canada (NSERC) organizations.

REFERENCES

- [1] Y. Yuan, X. Chen, and J. Wang, "Object-contextual representations for semantic segmentation," 2019.

- [2] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, "Mask R-CNN," *CoRR*, vol. abs/1703.06870, 2017. [Online]. Available: <http://arxiv.org/abs/1703.06870>
- [3] A. Kolesnikov, L. Beyer, X. Zhai, J. Puigcerver, J. Yung, S. Gelly, and N. Houlsby, "Large scale learning of general visual representations for transfer," 2019.
- [4] L. T. Nguyen-Meidine, E. Granger, M. Kiran, and L. Blais-Morin, "A comparison of cnn-based face and head detectors for real-time video surveillance applications," in *2017 Seventh International Conference on Image Processing Theory, Tools and Applications (IPTA)*, 2017, pp. 1–7.
- [5] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=rJl-b3RcF7>
- [6] H. Masson, A. Bhuiyan, L. T. Nguyen-Meidine, M. Javan, P. Siva, I. B. Ayed, and E. Granger, "A survey of pruning methods for efficient person re-identification across domains," *CoRR*, vol. abs/1907.02547, 2019. [Online]. Available: <http://arxiv.org/abs/1907.02547>
- [7] P. Molchanov, A. Mallya, S. Tyree, I. Frosio, and J. Kautz, "Importance estimation for neural network pruning," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [8] L. T. Nguyen-Meidine, E. Granger, M. Kiran, and L. Blais-Morin, "An improved trade-off between accuracy and complexity with progressive gradient pruning," *CoRR*, vol. abs/1906.08746, 2019. [Online]. Available: <http://arxiv.org/abs/1906.08746>
- [9] X. Dong and Y. Yang, "Network pruning via transformable architecture search," in *Advances in Neural Information Processing Systems 32*, 2019, pp. 759–770. [Online]. Available: <http://papers.nips.cc/paper/8364-network-pruning-via-transformable-architecture-search.pdf>
- [10] Z. Liu, M. Sun, T. Zhou, G. Huang, and T. Darrell, "Rethinking the value of network pruning," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=rJlnB3C5Ym>
- [11] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," in *NIPS Deep Learning and Representation Learning Workshop*, 2015. [Online]. Available: <http://arxiv.org/abs/1503.02531>
- [12] W. Park, D. Kim, Y. Lu, and M. Cho, "Relational knowledge distillation," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [13] B. Heo, J. Kim, S. Yun, H. Park, N. Kwak, and J. Y. Choi, "A comprehensive overhaul of feature distillation," in *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [14] S. Mirzadeh, M. Farajtabar, A. Li, and H. Ghasemzadeh, "Improved knowledge distillation via teacher assistant: Bridging the gap between student and teacher," *CoRR*, vol. abs/1902.03393, 2019. [Online]. Available: <http://arxiv.org/abs/1902.03393>
- [15] G. Wilson and D. J. Cook, "Adversarial transfer learning," *CoRR*, vol. abs/1812.02849, 2018. [Online]. Available: <http://arxiv.org/abs/1812.02849>
- [16] M. Wang and W. Deng, "Deep visual domain adaptation: A survey," *CoRR*, vol. abs/1802.03601, 2018. [Online]. Available: <http://arxiv.org/abs/1802.03601>
- [17] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," 2014.
- [18] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [19] M. Long, Y. Cao, J. Wang, and M. I. Jordan, "Learning transferable features with deep adaptation networks," in *ICML 2015*.
- [20] S. A. McIlraith and K. Q. Weinberger, Eds., *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18)*. AAAI Press, 2018.
- [21] H. Fu, M. Gong, C. Wang, K. Batmanghelich, K. Zhang, and D. Tao, "Geometry-consistent generative adversarial networks for one-sided unsupervised domain mapping," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [22] S. Benaim and L. Wolf, "One-sided unsupervised domain mapping," in *Advances in Neural Information Processing Systems 30*. Curran Associates, Inc., 2017, pp. 752–762. [Online]. Available: <http://papers.nips.cc/paper/6677-one-sided-unsupervised-domain-mapping.pdf>
- [23] S. Ruder, P. Ghaffari, and J. G. Breslin, "Knowledge adaptation: Teaching to adapt," *CoRR*, vol. abs/1702.02052, 2017. [Online]. Available: <http://arxiv.org/abs/1702.02052>
- [24] M. Orbes-Arteaga, J. Cardoso, L. Sørensen, C. Igel, S. Ourselin, M. Modat, M. Nielsen, and A. Pai, "Knowledge distillation for semi-supervised domain adaptation," 2019.
- [25] Z. Allen-Zhu, Y. Li, and Y. Liang, "Learning and generalization in over-parameterized neural networks, going beyond two layers," in *Advances in Neural Information Processing Systems 32*, 2019, pp. 6155–6166.
- [26] Y. He, P. Liu, Z. Wang, Z. Hu, and Y. Yang, "Filter pruning via geometric median for deep convolutional neural networks acceleration," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [27] R. Zhao, Y. Hu, J. Dotzel, C. De Sa, and Z. Zhang, "Improving neural network quantization without retraining using outlier channel splitting," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 97. PMLR, 2019, pp. 7543–7552.
- [28] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding," *CoRR*, vol. abs/1510.00149, 2015.
- [29] Anonymous, "And the bit goes down: Revisiting the quantization of neural networks," in *Submitted to International Conference on Learning Representations*, 2020, under review. [Online]. Available: <https://openreview.net/forum?id=rJehVyrKwH>
- [30] M. Jaderberg, A. Vedaldi, and A. Zisserman, "Speeding up convolutional neural networks with low rank expansions," *CoRR*, vol. abs/1405.3866, 2014. [Online]. Available: <http://arxiv.org/abs/1405.3866>
- [31] W. Wen, C. Xu, C. Wu, Y. Wang, Y. Chen, and H. Li, "Coordinating filters for faster deep neural networks," *CoRR*, vol. abs/1703.09746, 2017. [Online]. Available: <http://arxiv.org/abs/1703.09746>
- [32] Y. Cheng, D. Wang, P. Zhou, and T. Zhang, "Model compression and acceleration for deep neural networks: The principles, progress, and challenges," *IEEE Signal Processing Magazine*, p. 126–136, 2018.
- [33] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan, "Unsupervised pixel-level domain adaptation with generative adversarial networks," *CoRR*, vol. abs/1612.05424, 2016. [Online]. Available: <http://arxiv.org/abs/1612.05424>
- [34] W. Hong, Z. Wang, M. Yang, and J. Yuan, "Conditional generative adversarial network for structured domain adaptation," in *CVPR*, 2018.
- [35] K. Saito, Y. Ushiku, and T. Harada, "Asymmetric tri-training for unsupervised domain adaptation," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 70. PMLR, 2017, pp. 2988–2997.
- [36] Y. Li, N. Wang, J. Shi, J. Liu, and X. Hou, "Revisiting batch normalization for practical domain adaptation," *CoRR*, vol. abs/1603.04779, 2016. [Online]. Available: <http://arxiv.org/abs/1603.04779>
- [37] K.-Y. Wei and C.-T. Hsu, "Generative adversarial guided learning for domain adaptation," in *BMVC*, 2018.
- [38] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell, "CyCADA: Cycle-consistent adversarial domain adaptation," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 2018, pp. 1989–1998.
- [39] S. Ruder, P. Ghaffari, and J. G. Breslin, "Knowledge adaptation: Teaching to adapt," *CoRR*, vol. abs/1702.02052, 2017. [Online]. Available: <http://arxiv.org/abs/1702.02052>
- [40] M. Orbes-Arteaga, M. J. Cardoso, L. Sørensen, C. Igel, S. Ourselin, M. Modat, M. Nielsen, and A. Pai, *Knowledge Distillation for Semi-supervised Domain Adaptation*, 10 2019, pp. 68–76.
- [41] C. Yu, J. Wang, Y. Chen, and Z. Wu, "Accelerating deep unsupervised domain adaptation with transfer channel pruning," *CoRR*, vol. abs/1904.02654, 2019. [Online]. Available: <http://arxiv.org/abs/1904.02654>