

# Domain Adaptation Regularization for Spectral Pruning

Laurent Dillard<sup>1,†</sup>  
laurent.dillard@gmail.com

Yosuke Shinya<sup>2</sup>  
yosuke.shinya.j7r@jp.denso.com

Taiji Suzuki<sup>3,4</sup>  
taiji@mist.i.u-tokyo.ac.jp

<sup>1</sup> Behold.ai  
New York, USA

<sup>2</sup> DENSO CORPORATION  
Tokyo, Japan

<sup>3</sup> The University of Tokyo  
Tokyo, Japan

<sup>4</sup> AIP-RIKEN  
Tokyo, Japan

---

## Abstract

Deep Neural Networks (DNNs) have recently been achieving state-of-the-art performance on a variety of computer vision related tasks. However, their computational cost limits their ability to be implemented in embedded systems with restricted resources or strict latency constraints. Model compression has therefore been an active field of research to overcome this issue. Additionally, DNNs typically require massive amounts of labeled data to be trained. This represents a second limitation to their deployment. Domain Adaptation (DA) addresses this issue by allowing knowledge learned on one labeled source distribution to be transferred to a target distribution, possibly unlabeled. In this paper, we investigate on possible improvements of compression methods in DA setting. We focus on a compression method that was previously developed in the context of a single data distribution and show that, with a careful choice of data to use during compression and additional regularization terms directly related to DA objectives, it is possible to improve compression results. We also show that our method outperforms an existing compression method studied in the DA setting by a large margin for high compression rates. Although our work is based on one specific compression method, we also outline some general guidelines for improving compression in DA setting.

## 1 Introduction

In the past years, Deep Neural Networks have attracted much attention in vision tasks. Indeed, they led to significant improvements in many problems including image classification and object detection [21, 30]. This was mainly achieved by the development of complex new architectures based on Convolutional Neural Networks (CNNs) [12, 32, 41]. These architectures typically involve a huge number of parameters and operations for inference, making them ill-suited to being deployed in constrained environments. Several different strategies have been developed in order to circumvent this issue by compressing the models into smaller ones achieving similar performances. The main methods can be categorized




Figure 1: We propose a compression method with additional regularization terms for Domain Adaptation (DA) setting. Our method can maintain accuracy without fine-tuning after compression, and can achieve higher accuracy than prior work (DALR [26]). DALR is based on low-rank matrix decomposition, while our method is built on Spectral Pruning [39].

into the following schemes: selectively pruning parameters [13, 23], distillation [0, 3, 16], quantization [8, 46], or low-rank factorization and sparsity [17, 35]. Another main challenge for DNNs is to remain efficient on data coming from a target distribution similar yet not identical to the training data. A popular technique is to use a model pre-trained on a large source dataset and fine-tune its parameters on the desired target dataset. However it may be that no labels are available on the target data. In that case, training requires unsupervised DA techniques [0, 0, 53].

In this work, we seek to adapt compression objectives to the specific setting of DA. Concretely, we are provided with both a source and a, possibly unlabeled, target distribution. We are interested in compressing DNNs with high accuracy on the target distribution. Some works have investigated this specific setting. The work in [26] is strongly related to ours but only focuses on the case where a model is fine-tuned on a target dataset. Their method yields great results in this setting because it is data-dependent. That is, compression depends on the input data distribution. In this paper, we focus on a compression method that is also data-dependent. The novelty of our work is to go further in the analysis of how the data affects compression and to add a new regularization term. This regularization term is directly related to a DA objective. It basically favors nodes that are non domain discriminative. This proves useful as the more similar the features extracted are, between source and target distribution, the more discriminative ability learned on the source distribution will also apply to the target distribution. Finally we show that our extended compression method compares favorably to [26] on various cases of fine-tuning a pre-trained model on a target dataset.

## 2 Related work

Here we present a brief review of the main strategies used in model compression. One can find in [9] a more in-depth analysis of the different methods. **(1) Pruning.** These methods focus on finding weights of the network that are not crucial to the network’s performance. Early works used magnitude-based pruning which removed weights with the lowest values [13]. This method was extended by various works yielding significant results [10, 23]. In this work we focus on a pruning method [39], which is based on spectral analysis of the covariance matrix of layers [38, 40]. **(2) Network distillation.** These approaches are based on a teacher-student learning scheme [0, 3, 16]. A compressed, small network will learn to mimic a large network or ensemble of large networks. This is efficient because the student model learns based on a combination of true labels, which are hard encoded, and soft outputs of the teacher model. Those contain valuable information about similarity between classes. Distillation can also train the student network to mimic the behavior of intermediate layers [30]. **(3) Parameter quantization.** Network quantization is based on finding an efficient representation of the network’s parameters by reducing the number of bits required to

represent each weight. Various strategies can achieve this objective: [8, 46] applied k-means scalar quantization to the weight values. Other work use a fixed (8 or 16) bit representation of the weights while maintaining the accuracy of the original model [11, 44]. In [6] a network is constrained to use only binary weights while still achieving high accuracy. (4) **Low-rank factorization.** The idea behind low-rank factorization and sparse techniques is to decompose layers into smaller ones with fewer parameters. For convolutional layers, main techniques rely on the fact that performing a convolution with the original filters can be approached by a linear combination of a set of convolutions with base filters [6, 55]. In [26] truncation of Singular Value Decomposition (SVD) is used to factorize the weight matrices of dense (fully connected) layers.

Some works explore pruning methods in DA (or transfer learning) setting [25, 27, 36, 42, 48, 49]. Unlike our work, most existing methods depend on iterative fine-tuning. Although each fine-tuning step may mitigate information loss caused by each pruning step, strong dependence on fine-tuning increases the risk of overfitting (see [24] for more discussion).

In contrast to the iterative fine-tuning approach, the Domain Adaptive Low Rank Matrix Decomposition (DALR) method [26] can maintain accuracy without fine-tuning after compression. DALR and our method are data-dependent, which makes them suited for DA setting. Since DALR uses SVD, both are based on spectral analysis. An important difference between DALR and our method is their structural modification. We describe the difference, considering that uncompressed two layers have weight matrices of dimension  $m \times n$  and  $n \times p$  as shown in Fig. 1. For compressing the first layer, DALR produces two smaller layers that have weight matrices of dimension  $m \times k$  and  $k \times n$  ( $k$  is kept rank). First, it is noted that, to actually reduce the number of parameters in the network, the parameter  $k$  must verify the following condition:  $k(m+n) < mn$ . DALR affects only the layer being compressed since the output shape remains unchanged. Conversely, our method does not create additional layers but also affects the input dimension of the following layer. Second, DALR is only designed to compress dense layers while our method can be applied to both dense and convolutional layers. Finally, in [26], the authors only mention using  $k$  as a given input to compression. Therefore compression and validation with various  $k$  will be needed unless a prior knowledge of how much the layer is compressible is known. Our method helps to avoid this issue, using an information retention ratio parameter defined in Section 3.1. It should be noticed however that it is possible to reproduce a similar parametrization in DALR by computing the ratio of kept singular values over all singular values of the decomposition.

### 3 Spectral pruning with DA regularization

Our work was built on a layer-wise compression method [59], which is based on spectral analysis. The method does not require any regularization during training, but is applicable to a pre-trained model. We will first detail this method then introduce our DA regularizer.

#### 3.1 Principle

The method compresses the network layer per layer and can be applied to both dense and convolutional layers. It aims at reducing the dimension of the output of a layer such that the information of the input of the next layer remains mostly unchanged. To compress a layer, it selects the most informative nodes (neurons) or feature map channels (corresponding to convolutional filters). See Fig. 3 in [12] and Fig. 2 in [15] for visual explanations of pruning

dense and convolutional layers, respectively. For brevity, we will avoid the distinction of the layer types in the remainder of this paper and refer only to nodes for both cases.

More formally, let us denote  $x$  the input to the network and suppose that we are given a training data  $(x_i, y_i)_{i=1}^n$  with size  $n$ . Let us consider a model that is a sequence of convolutional and fully connected layers. We denote by  $f^l$  and  $\eta^l$  respectively the mapping function and activation function of a layer  $l$ . Finally we denote by  $\phi^l(x) \in \mathbb{R}^{m_l}$  the representation obtained at layer  $l$  where  $m_l$  is the number of nodes in the layer  $l$ . Therefore we have  $\phi^l(x) = \eta^l(f^l(\phi^{l-1}(x)))$ . The objective of the method is to find for each layer  $l$ , the subset of nodes  $J \subset \{1, \dots, m_l\}$  that minimizes the following quantity:

$$\hat{\mathbb{E}}[\|\phi^l(x) - \hat{A}_J \phi_J^l(x)\|^2], \quad (1)$$

where  $\hat{\mathbb{E}}[\cdot]$  denotes the empirical expectation,  $\phi_J^l(x)$  denotes the vector composed of the components of  $\phi^l(x)$  indexed by  $J$ , and  $\hat{A}_J$  is a matrix to recover the whole vector  $\phi^l$  from its sub-vector  $\phi_J^l$  as well as possible.

For a fixed  $J$ , it is easy to minimize the objective (1) with respect to  $\hat{A}_J$  as  $\hat{A}_J = \hat{\Sigma}_{F,J} \hat{\Sigma}_{J,J}^{-1}$ , where  $\hat{\Sigma} := \hat{\mathbb{E}}[\phi^l(x) \phi^l(x)^\top] (= \frac{1}{n} \sum_{i=1}^n \phi^l(x_i) \phi^l(x_i)^\top)$  is the empirical covariance matrix and  $F := \{1, \dots, m_l\}$  is the full set of indexes. Then  $\phi^l(x)$  can be approximated by  $\hat{A}_J \phi_J^l(x)$  which is equivalent to pruning the nodes of layer  $l$  that do not belong to  $J$  and reshaping the weights of the next layer. Let us denote  $W_{l+1}$  the weight matrix if the next layer is a fully connected layer,  $T_{l+1}$  the weight tensor if it is a convolutional layer. In that case we denote by the four dimensional tuple  $(o_c, i_c, W, H)$  the shape of the weight tensor where  $o_c, i_c, W, H$  stand respectively for output channels, input channels, width, height. The reshaping of the weights of the next layer is done the following way. For fully connected layers, we obtain the reshaped weight matrix as  $\tilde{W}^{l+1} = W^{l+1} \hat{A}_J$ . For convolutional layers, it is done by first reshaping  $T_{l+1}$  to a  $(W, H, o_c, i_c)$  tensor. Therefore each point of the filter grid is associated to a filter matrix  $M_{w,h} \in \mathbb{R}^{o_c \times i_c}$  that will be reshaped by the operation:  $\tilde{M}_{w,h} = M_{w,h} \hat{A}_J$ . The tensor is then reshaped back to its original form  $(o_c, i_c, W, H)$  where  $i_c$  has been modified by the process.

Once we obtained the optimal  $\hat{A}_J$  for a fixed  $J$ , then we minimize the objective (1) with respect to  $J$ . By substituting the optimal  $\hat{A}_J$  to the objective function (1), it can be re-written as  $\min_J \text{Tr}[\hat{\Sigma}_{F,F} - \hat{\Sigma}_{F,J} \hat{\Sigma}_{J,J}^{-1} \hat{\Sigma}_{J,F}]$ . The minimand in the right hand side is zero for  $J = F$  and  $\text{Tr}[\hat{\Sigma}_{F,F}]$  for  $J = \emptyset$ . Hence we may consider the following ‘‘ratio’’ of residual as an information retention ratio  $r$  to measure the relative goodness of  $J$ :

$$r = \frac{\text{Tr}(\hat{\Sigma}_{F,J} \hat{\Sigma}_{J,J}^{-1} \hat{\Sigma}_{J,F})}{\text{Tr}(\hat{\Sigma}_{F,F})}. \quad (2)$$

The higher the ratio the more information computed by the layer will be retained. It is no greater than 1 since the denominator is the best achievable value with no cardinality constraint. Compression can therefore be parametrized by a required information retention ratio parameter  $\alpha$  and the final optimization problem translates to  $\min_{J \subset F} |J|$  s.t.  $r \geq \alpha$ . To solve the optimization problem, the authors of [B9] proposed a greedy algorithm where  $J$  is constructed by sequentially adding the node that maximizes the information retention ratio.

It is worth noticing that though the method is aimed at compressing each layer sequentially, using  $\alpha$  allows more flexibility than methods with a fixed compression rate for each layer. Since we do not impose constraints on the cardinality of  $J$ , compression will adapt to the inherent compressibility of the layer. We argue that this makes  $\alpha$  an easy to tune and intuitive hyperparameter and, if necessary, can easily be combined with cardinality constraints.

## 3.2 Domain Adaptation regularization

In DA setting, informative nodes could be different between source and target domains. To adjust this difference, we propose an additional regularization term to select better nodes.

As shown by [1], one of the main challenges of DA is to find a representation discriminative with respect to the classification task but not discriminative between domains. The general idea is that, by matching the two distributions, the classification knowledge learned on the labeled source data can be transferred to the target data. Therefore, we propose to use a measure of the alignment of the source and target features distributions for selecting nodes during compression. We call our method *Moment Matching Spectral Pruning (MoMaSP)*.

Previous work involving covariance alignment achieved high accuracy in DA tasks [17]. Following those observations, we define our regularization term as the following

$$R_J = \|\hat{\mathbb{E}}_{x \sim X^s}[\phi_J^l(x)] - \hat{\mathbb{E}}_{x \sim X^t}[\phi_J^l(x)]\| + \|S_{J,J} \odot (C_{J,J}^s - C_{J,J}^t)\|_F, \quad (3)$$

where  $X^s, X^t$  denotes the source and target distributions,  $C^s, C^t$  respectively the source and target empirical covariance matrices,  $S$  a scaling matrix defined by  $S_{ij} = (C_{i,i}^t C_{j,j}^t)^{-\frac{1}{4}}$ ,  $\odot$  denotes the element wise multiplication, and  $\|\cdot\|_F$  means the Frobenius norm of a matrix. This quantity measures the discrepancy between the source and target distributions by using the up-to second order statistics. The first term measures the discrepancy of mean and the second term measures the discrepancy of the (scaled) second order moment. Hence, this regularization term induces a *moment matching* effect and the two distributions become similar on the selected nodes. We may use the MMD criterion (kernel-based discrepancy measure) [9] to capture not only the first and second moments but also all higher moments instead, but we found that it takes too much computational time. Although the kernelized alignment loss [14] could speed-up computation, source and target samples per class need to be few.

Instead of the criterion (3), we also consider the following alternative formulation:

$$R_j = \|\hat{\mathbb{E}}_{x \sim X^s}[\phi_j^l(x)] - \hat{\mathbb{E}}_{x \sim X^t}[\phi_j^l(x)]\| + \|S_{j,F} \odot (C_{j,F}^s - C_{j,F}^t)\|. \quad (4)$$

The first formulation (3) depends on the whole subset  $J$  while the second one computes information more specific to each node. The second formulation is more computationally efficient since the  $R_j$ 's can be computed only once for each candidate index. The first one needs to recompute it for every candidate index at each step of the procedure because of its dependence on  $J$ .

## 3.3 Intuition behind our regularization

The intuition behind our regularization is the following. In DA setting where few or no labels are available on the target distribution, the discriminative knowledge is learned mainly on the source distribution. Thus the feature distributions on both domains should be aligned for the discriminative knowledge learned on the source domain to apply to the target domain. When comparing the information retention ratios of Eq. 2 for different nodes, we observed that many of them had very close values. This means that many nodes capture approximately the same amount of information, in the sense of total variance. Our intuition was that when these relative differences are too small, they are not significant to discriminate between nodes and are more likely the result of noise. Therefore a better criterion should be used to differentiate nodes that capture a same amount of variance. Since our compression method is designed to be applied on models that have been trained using DA methods, it is natural to use a criterion related to DA objectives. We choose to compare nodes based on how well their feature distributions on source and target domains are aligned. Nodes leading to a better alignment

should be favored as they will allow for a better transfer of discriminative knowledge from source to target domain. Our method realizes this via a moment matching type regularization.

### 3.4 Practical implementation

We present in Algorithm 1 a pseudo code for a practical implementation of our method using the first formulation of our regularizer. If the second formulation is used, the regularization term for each node is computed only once before entering the while loop as it does not depend on  $J$ . At each step of the algorithm, we denote by  $C$  the set of candidate indexes. To select which node to add to  $J$ , the ratio from Eq. (2) is computed for each candidate index. We denote by  $V$  the vector where each coordinate corresponds to the ratio value of a candidate index. Similarly, we denote by  $R_J$  the vector where each coordinate correspond to the regularizing term associated to a candidate index.

Without regularization, the index to add to  $J$  is simply given by  $i = \arg \max_{j \in C} V_j$ . Using regularization, the index to add to  $J$  is chosen as

$$i = \arg \max_{j \in C} \left\{ V_j - \lambda \cdot \sigma(V) \cdot \frac{R_{J_j}}{\max_{j'} (R_{J_{j'}})} \right\}, \quad (5)$$

where  $\sigma(V)$  is the standard deviation of  $(V_j)_{j \in C}$ . The values of  $R_{J_j}$  are rescaled to be in the same range as  $V_j$ 's by max normalization. Multiplying by  $V$ 's standard deviation,  $\sigma(V)$ , ensures that the regularization will be applied at the right scale. It should only slightly modify the relative ordering of the values in  $V$  as the main criteria for selection must remain the information ratio. Indeed, only taking into account the regularization term would favor nodes that output a constant value across both distributions. The hyperparameter  $\lambda$  allows for more control over the trade-off between those two terms. We use 1 as the default  $\lambda$  in our experiments. The max normalization and scaling make  $\lambda$  easier to tune. This compression is applied after training and can therefore be combined with any DA method.

## 4 Experiments on digits images

In this section, we conduct experiments with a model trained on digits datasets, using a DA technique [63]. The source dataset is the SVHN dataset [28] and the target dataset is the MNIST dataset [27], which are standard for DA with adversarial training [7, 63, 64].

### 4.1 Model

Considering the relative simplicity of the task, we used a custom model composed of a 4 layers feature generator (3 convolutional + 1 dense) and a 2 layers classifier (see Supplementary Material for details). To train the model we used the DA technique presented in [63] which utilizes "maximum classifier discrepancy." Briefly, the adaptation is realized by adversarial training of a feature generator network and two classifier networks. The adversarial training

---

#### Algorithm 1: MoMaSP

---

**Input** :  $\Sigma_{F,F}$ : Empirical covariance matrix,  
 $\alpha$ : Required information retention ratio  
**Output**:  $J$ : Subset of selected nodes  
**Function** FindSubset( $\Sigma_{F,F}, \alpha$ )  
 $J \leftarrow \emptyset$   
 $T \leftarrow \text{Tr}(\Sigma_{F,F})$   
 $r \leftarrow 0$   
 $C \leftarrow F$   
**while**  $r < \alpha$  **do**  
  **for**  $j \in C$  **do**  
     $J_j \leftarrow J \cup j$   
    Compute  $R_{J_j}$   
     $V_j \leftarrow \frac{\text{Tr}(\Sigma_{F,J_j} \Sigma_{J_j}^{-1} \Sigma_{J_j, F})}{T}$   
  **end**  
   $i \leftarrow \arg \max_{j \in C} \left\{ V_j - \lambda \cdot \sigma(V) \cdot \frac{R_{J_j}}{\max_{j'} (R_{J_{j'}})} \right\}$   
   $J \leftarrow J \cup \{i\}$   
   $r \leftarrow V_i$   
   $C \leftarrow C \setminus \{i\}$   
**end**  
**end**

---

forces the feature generator to align the source and target distribution. Contrary to other methods relying on adversarial training with a domain discriminator, this method considers the decision boundary while aligning the distributions, ensuring that the generated features are discriminative with respect to the classification task at hand.




Figure 2: Target accuracy over compression rate for three different settings of data used as input. Results are presented as mean  $\pm$  standard deviation based on 10 iterations.




Figure 3: Comparison using only target data or each formulation of our DA regularization. Results are presented as mean  $\pm$  standard deviation based on 10 iterations.

Test data	First layer nodes specificity		
	Source	Target	None
Source	0.79	0.77	0.72
Target	0.65	0.84	0.71

Table 1: Mean activation rates.

C.R. (%)	96.0	96.5	97.0	97.5	98.0	98.5
Target (%)	94.6	94.0	92.4	90.2	84.4	68.7
Reg $J$ (%)	95.0	94.1	93.7	<b>91.1</b>	83.3	72.7
Reg $j$ (%)	<b>95.2</b>	<b>94.9</b>	<b>94.3</b>	90.6	<b>89.3</b>	<b>77.5</b>

Table 2: Accuracy on target test dataset for different compression rate (C.R.).

## 4.2 Data choice for compression

The uncompressed model was trained using both train splits and evaluated on the test split of MNIST. It reached an accuracy of 96.64%, similar to the results obtained in the original paper [53]. We then apply compression on the trained model. During compression, only data of the train splits was used and the compressed models were evaluated on the MNIST test split. Since the method is data-dependent, the choice of data to use to compute the empirical covariance matrix should be taken care of. Three different settings were tested: (1) Target: Using 14,000 samples from the target distribution; (2) Target/Source: Using 7,000 samples from the target distribution and 7,000 samples from the source distribution; and (3) Target+Source: Using 14,000 samples from the target distribution and 7,000 additional samples from the source distribution. The results obtained are presented in Fig. 2. Using only target samples to compute the covariance matrix shows the best performance.

To give a better understanding of this result, we conducted an analysis about how the nodes activation pattern depends on data distribution as follows: we compressed the first layer of a trained network using either only target data or only source data. We then compared the activation of nodes that were selected only in one of the two cases, in other words, nodes specific to target or source data. We show the results in Table 1.

As expected, nodes selected using source data had a significantly higher activation rate on the source distribution and conversely for target specific nodes. As a control case, we added activation of nodes that were selected in neither of the two settings. Those do not show any significant difference in their activation. Interestingly, this difference was no longer appearing when comparing activation of the last fully connected layer, because DA training aligns the distributions of extracted features. This experiment sheds light on how the input data affects the compression of data-dependent methods. In case of a model trained on different

distributions with DA, early layers contain nodes that are specific to each distribution. Thus it is critical for compression to use data coming exclusively from the distribution the model will be applied to. Partially using source data forces compression in the early layers to select source specific nodes to the expense of target specific nodes leading to poorer results.

### 4.3 Regularization for compression

Finally we compared the results of the best baseline, using only target data, with adding the regularization term introduced in Section 3. The results are presented in Fig. 3 and summarized in Table 2. It appears that the second formulation of our regularizer gives the best performance. This is probably due to the fact it focuses on node level information and not on the whole subset  $J$ , giving better ability to discriminate against specific nodes.

Compared to the baseline, our regularizer leads to significant improvements in compression performance. We observe up to 9% bumps in accuracy. Yet, we notice that the baseline performs slightly better for very high compression rates. We conjecture that reducing information loss becomes more important for such compression rates and our regularization with  $\lambda=1$  is too strong. There is room for improvement by tuning the hyperparameter  $\lambda$  depending on the compression rate or developing methods for adaptive regularization.

## 5 Experiments on natural images

In this section, we compare our method (MoMaSP) with the factorization based compression method (DALR) in [26]. We reproduce the same setting as in their experiment: a VGG19 model pre-trained on the ImageNet dataset [82] is fine-tuned on different target datasets.

### 5.1 Experimental settings

We establish our comparison based on the following three datasets used in [26] experiments: Oxford 102 Flowers [29], CUB-200-2011 Birds [45], and Stanford 40 Actions [47].

For all three datasets, we first trained uncompressed models by fine-tuning from an ImageNet pre-trained VGG19 model. In the fine-tuning before compression, we trained the weights of the fully connected layers while keeping the weights of the convolutional layers frozen to their pre-trained value on ImageNet like [26].

We then compressed the models. The input data to compression was always composed of 4,000 samples of the train split, except for the Oxford 102 Flowers where the whole train split was used. Additional 4,000 randomly sampled images from the ImageNet train split were used for DA regularization. Note that our method and DALR [26] do not need the labels of target datasets in this compression phase. After that we optionally fine-tuned the fully connected layers of compressed models with target labels.

We used the Adam optimizer [18] with a batch size of 50, a learning rate of  $10^{-4}$ , and a weight decay of  $5 \times 10^{-4}$  for training models. See Supplementary Material for details.

### 5.2 VGG19 fc7 compression

We first evaluated the compression on the last fully connected layer (fc7) of the model, containing 4,096 nodes, because the fc7 compression is a main evaluation setting in [26]. For each trial we report the results both with and without fine-tuning of the classifier layers



$k$	4	8	16	32	64	128
params (%)	0.20	0.39	0.78	1.56	3.13	6.25
SVD w/o FT.	5.2	12.8	30.2	54.6	65.2	67.6
SVD w/ FT.	13.7	36.7	56.6	65.9	70.0	70.8
DALR w/o FT.	8.6	31.1	55.1	66.9	71.0	<b>72.8</b>
DALR w/ FT.	18.9	48.2	64.4	<b>70.3</b>	70.7	<b>72.3</b>
Ours w/o FT.	<b>68.9</b>	<b>69.3</b>	<b>70.0</b>	<b>71.2</b>	<b>72.0</b>	72.5
Ours w/ FT.	<b>67.3</b>	<b>68.9</b>	<b>70.3</b>	69.9	<b>70.8</b>	70.6

Table 3: VGG19 fc7 compression on Oxford 102 Flowers. Test dataset accuracy results (%). Original accuracy: 73.1%.

$k$	4	8	16	32	64	128
params (%)	0.20	0.39	0.78	1.56	3.13	6.25
SVD w/o FT.	4.1	12.6	30.3	46.9	54.8	60.5
SVD w/ FT.	17.3	39.1	51.2	57.2	58.8	58.5
DALR w/o FT.	5.6	23.1	49.3	57.8	<b>60.1</b>	<b>60.9</b>
DALR w/ FT.	19.0	40.9	54.0	<b>59.8</b>	<b>59.3</b>	59.5
Ours w/o FT.	<b>58.3</b>	<b>58.2</b>	<b>58.2</b>	<b>58.8</b>	59.3	60.2
Ours w/ FT.	<b>57.9</b>	<b>57.3</b>	<b>57.7</b>	58.2	58.2	<b>59.7</b>

Table 4: VGG19 fc7 compression on CUB-200-2011 Birds. Test dataset accuracy results (%). Original accuracy: 61.5%.

$k$	4	8	16	32	64	128
params (%)	0.20	0.39	0.78	1.56	3.13	6.25
SVD w/o FT.	17.1	31.0	54.5	67.3	72.2	72.7
SVD w/ FT.	45.4	62.0	68.6	72.8	73.1	73.7
DALR w/o FT.	20.7	43.3	65.0	73.0	<b>74.3</b>	<b>74.7</b>
DALR w/ FT.	46.8	63.3	69.6	<b>73.5</b>	72.9	72.9
Ours w/o FT.	<b>69.0</b>	<b>70.2</b>	<b>71.8</b>	<b>73.1</b>	73.9	73.8
Ours w/ FT.	<b>70.5</b>	<b>70.4</b>	<b>71.7</b>	72.9	<b>73.7</b>	<b>74.0</b>

Table 5: VGG19 fc7 compression on Stanford 40 Actions. Test dataset accuracy results (%). Original accuracy: 74.6%.

	Layers		FT.	C.R. (%)	#params	Accuracy (%)		
	conv	fc				Oxford	CUB	Stanford
DALR	✓	✓	✓	84.9	21.1 M	48.1	43.0	57.7
						62.2	52.8	67.4
Ours	✓	✓	✓	85.0	21.1 M	48.1	37.1	61.0
						63.9	55.1	69.0
Ours	✓	✓	✓	<b>85.4</b>	<b>20.4 M</b>	<b>60.6</b>	<b>51.4</b>	<b>69.2</b>
						<b>71.4</b>	<b>58.3</b>	<b>72.6</b>

Table 6: VGG19 full compression on each dataset.

after compression. We also report the results of compression using basic SVD on fc7 weight matrix as a baseline to further illustrate the advantage of using data-dependent methods. To compare our method with DALR, a dimension  $k$  was set for the compression using DALR then  $k'$ , the dimension to keep in our method resulting in an equal number of parameters, was determined accordingly. See Supplementary Material for details.

The results are presented in Tables 3, 4, and 5. We show the relative numbers of parameters in the fc7 layer compressed by DALR in the “params” rows in the tables as in the DALR paper [26]. In all experiments, our method maintains high accuracy even for high compression rates. In such cases it outperforms DALR by a large margin. However, for lower compression rates DALR consistently compares favorably to our method though the difference is arguably small. In most cases, fine-tuning does not improve much the performance of any of the two methods, except for DALR for high compression rates. This phenomenon implies that the learning rate and/or the dropout rate we used are too high especially for our method (see additional fine-tuning results in Supplementary Material for details).

### 5.3 VGG19 full compression

It is important to notice that, contrary to DALR, our method is able to compress both convolutional and dense layers. To further demonstrate this advantage, we proceeded to the comparison of fully compressing the VGG19 network using both methods. We conducted the experiment on all three datasets. DALR was applied to the fully connected layers. The dropout layers of models compressed by our method were disabled in fine-tuning, because our method reduce the input dimension of the dropout layers unlike DALR, resulting in performance degradation. See Supplementary Material for other details.

Results are presented in Table 6. For all three datasets, our method consistently achieves a better compression rate (C.R.) while reaching test accuracy 5.2%–12.5% higher than DALR. If we apply our method to the fully connected layers only, the difference in accuracy is smaller but still favorable to our method in many cases. Although further improvement

would be preferable to keep representations for many classes (CUB-200 without fine-tuning), in practice, it is not necessary to overcompress the fully connected layers by our method.

## 5.4 Discussion

Our method is drastically better for high compression rates, while DALR is slightly better for low compression rates. We conjecture that this is mainly due to the fact that the two methods modify differently the network (see Fig. 1). DALR affects only one layer. Therefore if the compression is too important, it reaches a critical point where not enough information can possibly be retained. Our method affects two successive layers therefore spreading the effect of compression and avoiding this critical point for high compression rates. On the other hand, our method needs to consider nonlinearity between two layers and uses a greedy algorithm [39]. Thus only affecting one layer is better to maintain high accuracy for lower compression rates, because the output of the one layer can be optimally approximated.

The optimization problem to solve in DALR admits a closed form solution, making it a very fast compression method. Compared to DALR, it is a limitation of our method to require an iterative optimization. However, the extra computational time is usually a few hours on 1 GPU (Tesla V100) in our experiments. Furthermore, an iterative process is used in the DALR paper [26] to determine compression rate pairs for fc6 and fc7, and it needs iterative accuracy evaluation. Pruning methods based on iterative fine-tuning approach also take time for pruning, fine-tuning, and evaluation. Therefore our method is practical enough.

## 6 Conclusions

In this paper, we investigated compression of DNNs in the DA setting. As shown by [26], using a data-dependent method is crucial in order to achieve good results. In that matter, our work shows that the input data to calculate information retention ratio for compression should only come from the distribution on which the model will be applied to, the target distribution in our case. This is because adding samples from another distribution will force compression in the early layers to select nodes that are specific to this distribution. However, we show that source data can still be used for regularization to improve nodes selection. This is done by comparing the first and second order statistics of the node’s feature distributions on each of the source and target data. This criterion serves as a measure of the alignment of the two distributions which directly relates to DA objectives. Therefore we denote this measure as a DA regularizer. We evaluated this regularization on a spectral pruning method introduced in [39] and obtained significant improvements on its compression results. Finally we compared our regularized compression method with the factorization based method of [26] on real world image datasets. Our method compares favorably on all three datasets, leading to significant improvements in retained accuracy for high compression rates.

Although our work focused on one compression method, we argue that using first and second order statistics of feature distributions to measure the alignment between source and target features and using it as a criterion for compression can be applied to other methods. This work can therefore serve as a first example and practical implementation of this idea. Applying our method to unsupervised DA setting [20, 48] is an interesting future direction.

**Acknowledgements** TS was partially supported by JSPS KAKENHI (26280009, 15H05707, and 18H03201), Japan Digital Design, and JST-CREST.

## References

- [1] Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *NIPS*. 2014.
- [2] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 2010.
- [3] Cristian Bucila, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *KDD*, 2006.
- [4] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. Model compression and acceleration for deep neural networks: The principles, progress, and challenges. *IEEE Signal Processing Magazine*, 2018.
- [5] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. BinaryConnect: Training deep neural networks with binary weights during propagations. In *NIPS*, 2015.
- [6] Emily Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *NIPS*. 2014.
- [7] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 2016.
- [8] Yunchao Gong, Liu Liu, Ming Yang, and Lubomir D. Bourdev. Compressing deep convolutional networks using vector quantization. *arXiv:1412.6115*, 2014.
- [9] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *JMLR*, 2012.
- [10] Yiwen Guo, Anbang Yao, and Yurong Chen. Dynamic network surgery for efficient dnns. In *NIPS*. 2016.
- [11] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *ICML*, 2015.
- [12] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *NIPS*. 2015.
- [13] Stephen Jose Hanson and Lorien Y. Pratt. Comparing biases for minimal network construction with back-propagation. In *NIPS*. 1988.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [15] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *ICCV*, 2017.
- [16] Geoffrey E. Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv:1503.02531*, 2015.

- [17] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. In *BMVC*, 2014.
- [18] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [19] Piotr Koniusz, Yusuf Tas, and Fatih Porikli. Domain adaptation by mixture of alignments of second- or higher-order scatter tensors. In *CVPR*, 2017.
- [20] Piotr Koniusz, Yusuf Tas, Hongguang Zhang, Mehrtash Harandi, Fatih Porikli, and Rui Zhang. Museum exhibit identification challenge for the supervised domain adaptation and beyond. In *ECCV*, 2018.
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet classification with deep convolutional neural networks. In *NIPS*. 2012.
- [22] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.
- [23] Guiying Li, Chao Qian, Chunhui Jiang, Xiaofen Lu, and Ke Tang. Optimization based layer-wise magnitude-based pruning for DNN compression. In *IJCAI*, 2018.
- [24] Bingyan Liu, Yao Guo, and Xiangqun Chen. WealthAdapt: A general network adaptation framework for small data tasks. In *ACMMM*, 2019.
- [25] Jiaming Liu, Yali Wang, and Yu Qiao. Sparse deep transfer learning for convolutional neural network. In *AAAI*, 2017.
- [26] Marc Masana, Joost van de Weijer, Luis Herranz, Andrew D. Bagdanov, and Jose M. Alvarez. Domain-adaptive deep network compression. In *ICCV*, 2017.
- [27] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. In *ICLR*, 2017.
- [28] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [29] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics & Image Processing*, 2008.
- [30] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [31] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. FitNets: Hints for thin deep nets. In *ICLR*, 2015.
- [32] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015.

- [33] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 2018.
- [34] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [35] Amos Sironi, Bugra Tekin, Roberto Rigamonti, Vincent Lepetit, and Pascal Fua. Learning separable filters. *TPAMI*, 2015.
- [36] Xavier Suau, Luca Zappella, and Nicholas Apostoloff. Filter distillation for network compression. In *WACV*, 2020.
- [37] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *AAAI*, 2016.
- [38] Taiji Suzuki. Fast generalization error bound of deep learning from a kernel perspective. In *AISTATS*, 2018.
- [39] Taiji Suzuki, Hiroshi Abe, Tomoya Murata, Shingo Horiuchi, Kotaro Ito, Tokuma Wachi, So Hirai, Masatoshi Yukishima, and Tomoaki Nishimura. Spectral pruning: Compressing deep neural networks via spectral analysis and its generalization error. In *IJCAI-PRICAI*, 2020.
- [40] Taiji Suzuki, Hiroshi Abe, and Tomoaki Nishimura. Compression based bound for non-compressed network: unified generalization error analysis of large compressible deep neural network. In *ICLR*, 2020.
- [41] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [42] Frederick Tung, Srikanth Muralidharan, and Greg Mori. Fine-Pruning: Joint fine-tuning and compression of a convolutional network with Bayesian optimization. In *BMVC*, 2017.
- [43] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017.
- [44] Vincent Vanhoucke, Andrew Senior, and Mark Z. Mao. Improving the speed of neural networks on CPUs. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [45] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [46] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *CVPR*, 2016.
- [47] Bangpeng Yao, Xiaoye Jiang, Aditya Khosla, Andy Lai Lin, Leonidas Guibas, and Li Fei-Fei. Human action recognition by learning bases of action attributes and parts. In *ICCV*, 2011.

- [48] Chaohui Yu, Jindong Wang, Yiqiang Chen, and Zijing Wu. Accelerating deep unsupervised domain adaptation with transfer channel pruning. In *IJCNN*, 2019.
- [49] Yang Zhong, Vladimir Li, Ryuzo Okada, and Atsuto Maki. Target aware network adaptation for efficient representation learning. In *ECCV Workshop on Compact and Efficient Feature Representation and Learning*, 2018.