

# Deep CORAL: Correlation Alignment for Deep Domain Adaptation

Baochen Sun\* and Kate Saenko\*\*

University of Massachusetts Lowell, Boston University

**Abstract.** Deep neural networks are able to learn powerful representations from large quantities of labeled input data, however they cannot always generalize well across changes in input distributions. Domain adaptation algorithms have been proposed to compensate for the degradation in performance due to domain shift. In this paper, we address the case when the target domain is unlabeled, requiring unsupervised adaptation. CORAL [18] is a simple unsupervised domain adaptation method that aligns the second-order statistics of the source and target distributions with a linear transformation. Here, we extend CORAL to learn a nonlinear transformation that aligns correlations of layer activations in deep neural networks (Deep CORAL). Experiments on standard benchmark datasets show state-of-the-art performance. Our code is available at: <https://github.com/VisionLearningGroup/CORAL>

## 1 Introduction

Many machine learning algorithms assume that the training and test data are independent and identically distributed (i.i.d.). However, this assumption rarely holds in practice as the data is likely to change over time and space. Even though state-of-the-art Deep Convolutional Neural Network features are invariant to low level cues to some degree [16,15,19], Donahue et al. [3] showed that they still are susceptible to domain shift. Instead of collecting labeled data and training a new classifier for every possible scenario, unsupervised domain adaptation methods [17,6,4,20,21,18] try to compensate for the degradation in performance by transferring knowledge from labeled source domains to unlabeled target domains. A recently proposed CORAL method [18] aligns the second-order statistics of the source and target distributions with a linear transformation. Even though it is easy to implement, it works well for unsupervised domain adaptation. However, it relies on a linear transformation and is not end-to-end trainable: it needs to first extract features, apply the transformation, and then train an SVM classifier in a separate step.

In this work, we extend CORAL to incorporate it directly into deep networks by constructing a differentiable loss function that minimizes the difference between source and target correlations—the CORAL loss. Compared to CORAL, our proposed Deep CORAL approach learns a non-linear transformation that is more powerful and also works seamlessly with deep CNNs. We evaluate our method on standard benchmark datasets and show state-of-the-art performance.

---

\* bsun@cs.uml.edu

\*\* saenko@bu.edu

## 2 Related Work

Previous techniques for unsupervised adaptation consisted of re-weighting the training point losses to more closely reflect those in the test distribution [11,9] or finding a transformation in a lower-dimensional manifold that brings the source and target subspaces closer together [7,6,8,4]. Re-weighting based approaches often assume a restricted form of domain shift—selection bias—and are thus not applicable to more general scenarios. Geodesic methods [7,6] bridge the source and target domains by projecting them onto points along a geodesic path [7], or finding a closed-form linear map that transforms source points to target [6]. [8,4] align the subspaces by computing the linear map that minimizes the Frobenius norm of the difference between the top  $n$  eigenvectors. In contrast, CORAL [18] minimizes domain shift by aligning the second-order statistics of source and target distributions.

Adaptive deep neural networks have recently been explored for unsupervised adaptation. DLID [1] trains a joint source and target CNN architecture with two adaptation layers. DDC [23] applies a single linear kernel to one layer to minimize Maximum Mean Discrepancy (MMD) while DAN [13] minimizes MMD with multiple kernels applied to multiple layers. ReverseGrad [5] and Domain-Confusion [22] add a binary classifier to explicitly confuse the two domains.

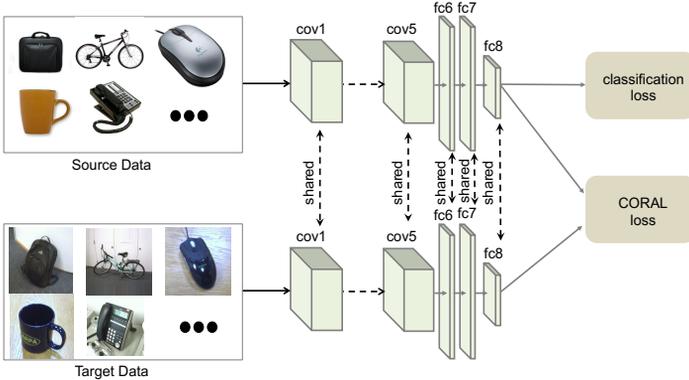
Our proposed Deep CORAL approach is similar to DDC, DAN, and ReverseGrad in the sense that a new loss (CORAL loss) is added to minimize the difference in learned feature covariances across domains, which is similar to minimizing MMD with a polynomial kernel. However, it is more powerful than DDC (which aligns sample means only), much simpler to optimize than DAN and ReverseGrad, and can be integrated into different layers or architectures seamlessly.

## 3 Deep CORAL

We address the unsupervised domain adaptation scenario where there are no labeled training data in the target domain, and propose to leverage both the deep features pre-trained on a large generic domain (*e.g.*, ImageNet [2]) and the labeled source data. In the meantime, we also want the final learned features to work well on the target domain. The first goal can be achieved by initializing the network parameters from the generic pre-trained network and fine-tuning it on the labeled source data. For the second goal, we propose to minimize the difference in second-order statistics between the source and target feature activations—the CORAL loss. Figure 1 shows a sample Deep CORAL architecture using our proposed correlation alignment layer for deep domain adaptation. We refer to Deep CORAL as any deep network incorporating the CORAL loss for domain adaptation.

### 3.1 CORAL Loss

We first describe the CORAL loss between two domains for a single feature layer. Suppose we are given source-domain training examples  $D_S = \{\mathbf{x}_i\}$ ,  $\mathbf{x} \in \mathbb{R}^d$  with



**Fig. 1.** Sample Deep CORAL architecture based on a CNN with a classifier layer. For generalization and simplicity, here we apply the CORAL loss to the *fc8* layer of AlexNet [12]. Integrating it into other layers or network architectures is also possible.

labels  $L_S = \{y_i\}$ ,  $i \in \{1, \dots, L\}$ , and unlabeled target data  $D_T = \{\mathbf{u}_i\}$ ,  $\mathbf{u} \in \mathbb{R}^d$ . Suppose the numbers of source and target data are  $n_S$  and  $n_T$  respectively. Here both  $\mathbf{x}$  and  $\mathbf{u}$  are the  $d$ -dimensional deep layer activations  $\phi(I)$  of input  $I$  that we are trying to learn. Suppose  $D_S^{ij}$  ( $D_T^{ij}$ ) indicates the  $j$ -th dimension of the  $i$ -th source (target) data example and  $C_S$  ( $C_T$ ) denote the feature covariance matrices.

We define the CORAL loss as the distance between the second-order statistics (covariances) of the source and target features:

$$\mathcal{L}_{CORAL} = \frac{1}{4d^2} \|C_S - C_T\|_F^2 \quad (1)$$

where  $\|\cdot\|_F^2$  denotes the squared matrix Frobenius norm. The covariance matrices of the source and target data are given by:

$$C_S = \frac{1}{n_S - 1} (D_S^\top D_S - \frac{1}{n_S} (\mathbf{1}^\top D_S)^\top (\mathbf{1}^\top D_S)) \quad (2)$$

$$C_T = \frac{1}{n_T - 1} (D_T^\top D_T - \frac{1}{n_T} (\mathbf{1}^\top D_T)^\top (\mathbf{1}^\top D_T)) \quad (3)$$

where  $\mathbf{1}$  is a column vector with all elements equal to 1.

The gradient with respect to the input features can be calculated using the chain rule:

$$\frac{\partial \mathcal{L}_{CORAL}}{\partial D_S^{ij}} = \frac{1}{d^2(n_S - 1)} ((D_S^\top - \frac{1}{n_S} (\mathbf{1}^\top D_S)^\top \mathbf{1}^\top)^\top (C_S - C_T))^{ij} \quad (4)$$

$$\frac{\partial \mathcal{L}_{CORAL}}{\partial D_T^{ij}} = -\frac{1}{d^2(n_T - 1)} ((D_T^\top - \frac{1}{n_T} (\mathbf{1}^\top D_T)^\top \mathbf{1}^\top)^\top (C_S - C_T))^{ij} \quad (5)$$

We use batch covariances and the network parameters are shared between the two networks.

### 3.2 End-to-end Domain Adaptation with CORAL Loss

We describe our method by taking a multi-class classification problem as the running example. As mentioned before, the final deep features need to be both discriminative enough to train a strong classifier and invariant to the difference between source and target domains. Minimizing the classification loss itself is likely to lead to overfitting to the source domain, causing reduced performance on the target domain. On the other hand, minimizing the CORAL loss alone might lead to degenerated features. For example, the network could project all of the source and target data to a single point, making the CORAL loss trivially zero. However, no strong classifier can be constructed on these features. Joint training with both the classification loss and CORAL loss is likely to learn features that work well on the target domain:

$$\mathcal{L} = \mathcal{L}_{CLASS} + \sum_{i=1}^t \lambda_i \mathcal{L}_{CORAL} \quad (6)$$

where  $t$  denotes the number of CORAL loss layers in a deep network and  $\lambda$  is a weight that trades off the adaptation with classification accuracy on the source domain. As we show below, these two losses play counterparts and reach an *equilibrium* at the end of training, where the final features are discriminative and generalize well to the target domain.

## 4 Experiments

We evaluate our method on a standard domain adaptation benchmark—the Office dataset [17]. The Office dataset contains 31 object categories from an office environment in 3 image domains: *Amazon*, *DSLR*, and *Webcam*.

We follow the standard protocol of [6,13,3,23,5] and use all the labeled source data and all the target data without labels. Since there are 3 domains, we conduct experiments on all 6 shifts (5 runs per shift), taking one domain as the source and another as the target.

In this experiment, we apply the CORAL loss to the last classification layer as it is the most general case—most deep classifier architectures (*e.g.*, convolutional neural networks, recurrent neural networks) contain a fully connected layer for classification. Applying the CORAL loss to other layers or other network architectures is also possible.

The dimension of the last fully connected layer (*fc8*) was set to the number of categories (31) and initialized with  $\mathcal{N}(0, 0.005)$ . The learning rate of *fc8* was set to 10 times the other layers as it was training from scratch. We initialized the other layers with the parameters pre-trained on ImageNet [2] and kept the original layer-wise parameter settings. In the training phase, we set the batch size to 128, base learning rate to  $10^{-3}$ , weight decay to  $5 \times 10^{-4}$ , and momentum to 0.9. The weight of the CORAL loss ( $\lambda$ ) is set in such way that at the end of training the classification loss and CORAL loss are roughly the same. It seems be a reasonable choice as we want to have a feature representation that is both discriminative and also minimizes the distance between the source and

target domains. We used Caffe [10] and BVLC Reference CaffeNet for all of our experiments.

We compare to 7 recently published methods: CNN [12] (no adaptation), GFK [6], SA [4], TCA [14], CORAL [18], DDC [23], DAN [13]. GFK, SA, and TCA are manifold based methods that project the source and target distributions into a lower-dimensional manifold and are not end-to-end deep methods. DDC adds a domain confusion loss to AlexNet [12] and fine-tunes it on both the source and target domain. DAN is similar to DDC but utilizes a multi-kernel selection method for better mean embedding matching and adapts in multiple layers. For direct comparison, DAN in this paper uses the hidden layer *fc8*. For GFK, SA, TCA, and CORAL, we use the *fc7* feature fine-tuned on the source domain (*FT7* in [18]) as it achieves better performance than generic pre-trained features, and train a linear SVM [4,18]. To have a fair comparison, we use accuracies reported by other authors with exactly the same setting or conduct experiments using the source code provided by the authors.

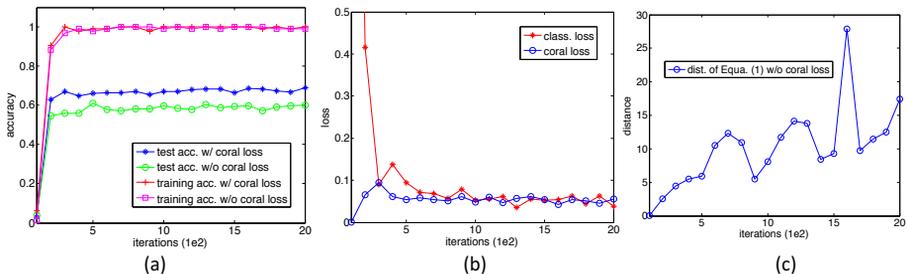
From Table 1 we can see that Deep CORAL (D-CORAL) achieves better average performance than CORAL and the other 6 baseline methods. In 3 out of 6 shifts, it achieves the highest accuracy. For the other 3 shifts, the margin between D-CORAL and the best baseline method is very small ( $\leq 0.7$ ).

	A→D	A→W	D→A	D→W	W→A	W→D	AVG
GFK	52.4±0.0	54.7±0.0	43.2±0.0	92.1±0.0	41.8±0.0	96.2±0.0	63.4
SA	50.6±0.0	47.4±0.0	39.5±0.0	89.1±0.0	37.6±0.0	93.8±0.0	59.7
TCA	46.8±0.0	45.5±0.0	36.4±0.0	81.1±0.0	39.5±0.0	92.2±0.0	56.9
CORAL	65.7±0.0	64.3±0.0	48.5±0.0	<b>96.1±0.0</b>	48.2±0.0	<b>99.8±0.0</b>	70.4
CNN	63.8±0.5	61.6±0.5	51.1±0.6	95.4±0.3	49.8±0.4	99.0±0.2	70.1
DDC	64.4±0.3	61.8±0.4	52.1±0.8	95.0±0.5	<b>52.2±0.4</b>	98.5±0.4	70.6
DAN	65.8±0.4	63.8±0.4	<b>52.8±0.4</b>	94.6±0.5	51.9±0.5	98.8±0.6	71.3
D-CORAL	<b>66.8±0.6</b>	<b>66.4±0.4</b>	<b>52.8±0.2</b>	95.7±0.3	51.5±0.3	99.2±0.1	<b>72.1</b>

**Table 1.** Object recognition accuracies for all 6 domain shifts on the standard Office dataset with deep features, following the standard unsupervised adaptation protocol.

To get a better understanding of Deep CORAL, we generate three plots for domain shift A→W. In Figure 2(a) we show the training (source) and testing (target) accuracies for training with v.s. without CORAL loss. We can clearly see that adding the CORAL loss helps achieve much better performance on the target domain while maintaining strong classification accuracy on the source domain.

In Figure 2(b) we visualize both the classification loss and the CORAL loss for training w/ CORAL loss. As the last fully connected layer is randomly initialized with  $\mathcal{N}(0, 0.005)$ , in the beginning the CORAL loss is very small while the classification loss is very large. After training for a few hundred iterations, these two losses are about the same and reach an *equilibrium*. In Figure 2(c) we show



**Fig. 2.** Detailed analysis of shift A→W for training w/ v.s. w/o CORAL loss. (a): training and test accuracies for training w/ v.s. w/o CORAL loss. We can see that adding CORAL loss helps achieve much better performance on the target domain while maintaining strong classification accuracy on the source domain. (b): classification loss and CORAL loss for training w/ CORAL loss. As the last fully connected layer is randomly initialized with  $\mathcal{N}(0, 0.005)$ , CORAL loss is very small while classification loss is very large at the beginning. After training for a few hundred iterations, these two losses are about the same. (c): CORAL distance for training w/o CORAL loss (setting the weight to 0). The distance is getting much larger ( $\geq 100$  times larger compared to training w/ CORAL loss).

the CORAL distance between the domains for training w/o CORAL loss (setting the weight to 0). We can see that the distance is getting much larger ( $\geq 100$  times larger compared to training w/ CORAL loss). Comparing Figure 2(b) and Figure 2(c), we can see that even though the CORAL loss is not always decreasing during training, if we set its weight to 0, the distance between source and target domains becomes much larger. This is reasonable as fine-tuning without domain adaptation is likely to overfit the features to the source domain. Our CORAL loss constrains the distance between source and target domain during the fine-tuning process and helps to maintain an *equilibrium* where the final features work well on the target domain.

## 5 Conclusion

In this work, we extended CORAL, a simple yet effective unsupervised domain adaptation method, to perform end-to-end adaptation in deep neural networks. Experiments on standard benchmark datasets show state-of-the-art performance. Deep CORAL works seamlessly with deep networks and can be easily integrated into different layers or network architectures.

## 6 Acknowledgments

This research was supported by NSF awards IIS-1451244 and IIS-1212928. The Tesla K40 and Titan X GPUs used for this research were donated by the NVIDIA Corporation.

## References

1. Chopra, S., Balakrishnan, S., Gopalan, R.: Dlid: Deep learning for domain adaptation by interpolating between domains. In: ICML Workshop (2013)
2. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
3. Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., Darrell, T.: Decaf: A deep convolutional activation feature for generic visual recognition. In: ICML (2014)
4. Fernando, B., Habrard, A., Sebban, M., Tuytelaars, T.: Unsupervised visual domain adaptation using subspace alignment. In: ICCV (2013)
5. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: ICML (2015)
6. Gong, B., Shi, Y., Sha, F., Grauman, K.: Geodesic flow kernel for unsupervised domain adaptation. In: CVPR (2012)
7. Gopalan, R., Li, R., Chellappa, R.: Domain adaptation for object recognition: An unsupervised approach. In: ICCV (2011)
8. Harel, M., Mannor, S.: Learning from multiple outlooks. In: ICML (2011)
9. Huang, J., Smola, A.J., Gretton, A., Borgwardt, K.M., Schölkopf, B.: Correcting sample selection bias by unlabeled data. In: NIPS (2006)
10. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding. arXiv preprint arXiv:1408.5093 (2014)
11. Jiang, J., Zhai, C.: Instance Weighting for Domain Adaptation in NLP. In: ACL (2007)
12. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
13. Long, M., Cao, Y., Wang, J., Jordan, M.I.: Learning transferable features with deep adaptation networks. In: ICML (2015)
14. Pan, S.J., Tsang, I.W., Kwok, J.T., Yang, Q.: Domain adaptation via transfer component analysis. In: IJCAI (2009)
15. Peng, X., Sun, B., Ali, K., Saenko, K.: Learning deep object detectors from 3d models. In: ICCV (2015)
16. Peng, X., Sun, B., Ali, K., Saenko, K.: What do deep cnns learn about objects? In: ICLR Workshop Track (2015)
17. Saenko, K., Kulis, B., Fritz, M., Darrell, T.: Adapting visual category models to new domains. In: ECCV (2010)
18. Sun, B., Feng, J., Saenko, K.: Return of frustratingly easy domain adaptation. In: AAAI (2016)
19. Sun, B., Peng, X., Saenko, K.: Generating large scale image datasets from 3d cad models. In: CVPR'15 Workshop on The Future of Datasets in Vision (2015)
20. Sun, B., Saenko, K.: From virtual to reality: Fast adaptation of virtual object detectors to real domains. In: BMVC (2014)
21. Sun, B., Saenko, K.: Subspace distribution alignment for unsupervised domain adaptation. In: BMVC (2015)
22. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous deep transfer across domains and tasks. In: International Conference in Computer Vision (ICCV) (2015)
23. Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., Darrell, T.: Deep domain confusion: Maximizing for domain invariance. CoRR abs/1412.3474 (2014), <http://arxiv.org/abs/1412.3474>