Supervised Domain Adaptation using Graph Embedding

AARHUS UNIVERSITY

Lukas Hedegaard Morsing*, Omar Ali Sheikh-Omar*, and Alexandros Iosifidis
Department of Electrical and Computer Engineering, Aarhus University, Denmark
lh@eng.au.dk, sheikhomar@mailbox.com, ai@eng.au.dk

ICPR 2020
25th INTERNATIONAL CONFERENCE ON PATTERN RECOGNITION
Milan, Italy 10 | 15 January 2021

Abstract

Getting deep convolutional neural networks to perform well requires a large amount of training data. When the available labelled data is small, it is often beneficial to use transfer learning to leverage a related larger dataset (source) in order to improve the performance on the small dataset (target). Among the transfer learning approaches, domain adaptation methods assume that distributions between the two domains are shifted and attempt to realign them. In this paper, we consider the domain adaptation problem from the perspective of multi-view graph embedding and dimensionality reduction. Instead of solving the generalised eigenvalue problem to perform the embedding, we formulate the graph-preserving criterion as a loss in the neural network and learn a domain-invariant feature transformation in an end-to-end fashion. We show that the proposed approach leads to a powerful Domain Adaptation framework which generalises the prior methods CCSA [1] and *d*-SNE [2], and enables simple and effective loss designs; an LDA-inspired [3] instantiation of the framework leads to performance on par with the state-of-the-art on the most widely used Domain Adaptation benchmarks, Office31 and MNIST to USPS datasets.

References

[1] S. Motiian, M. Piccirilli, D. A. Adjeroh, and G. Doretto, "Unified deep supervised domain adaptation and generalization," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5715–5725

[2] X. Xu, X. Zhou, R. Venkatesan, G. Swaminathan, and O. Majumder, "d-sne: Domain adaptation using stochastic neighborhood embedding," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2497–2506.

[3] S. Yan, D. Xu, B. Zhang, H.-J. Zhang, Q. Yang, and S. Lin, "Graph embedding and extensions: A general framework for dimensionality reduction," *IEEE trans- actions on pattern analysis and machine intelligence*, vol. 29, no. 1, pp. 40–51, 2006.

Domain Adaptation Using Graph Embedding (DAGE)

Source and target features

$$\mathbf{\Phi} = [\varphi(\mathbf{X}_\mathcal{S}) \quad \varphi(\mathbf{X}_\mathcal{T})]$$

Encode (dis)similarity $(\mathbf{W}_p)$ $\mathbf{W}$

Graph Laplacians $\mathbf{L} = \sum_j \mathbf{W}^{(i,j)} - \mathbf{W}$

$$\mathbf{B} = \sum_j \mathbf{W}_p^{(i,j)} - \mathbf{W}_p$$

Minimise within-class spread

$$\sum_{i=1}^{N}\sum_{j=1}^{N} \left\| \mathbf{\Phi}^{(i)} - \mathbf{\Phi}^{(j)} \right\|_2^2 \mathbf{W}^{(i,j)} = \mathrm{Tr}(\mathbf{\Phi}\mathbf{L}\mathbf{\Phi}^\top)$$

Maximise between-class spread

$$\sum_{i=1}^{N}\sum_{j=1}^{N} \left\| \mathbf{\Phi}^{(i)} - \mathbf{\Phi}^{(j)} \right\|_2^2 \mathbf{W}_\mathrm{p}^{(i,j)} = \mathrm{Tr}(\mathbf{\Phi}\mathbf{B}\mathbf{\Phi}^\top)$$

Jointly $\mathcal{L}_\mathrm{DAGE} = \dfrac{\mathrm{Tr}(\mathbf{\Phi}\mathbf{L}\mathbf{\Phi}^\top)}{\mathrm{Tr}(\mathbf{\Phi}\mathbf{B}\mathbf{\Phi}^\top)}$

DAGE-LDA

Inspired by Linear Discriminant Analysis

$$\mathbf{W}^{(i,j)} = \begin{cases} 1, & \ell_i = \ell_j \wedge \mathcal{D}_i \neq \mathcal{D}_j \\ 0, & \text{otherwise} \end{cases}$$

$$\mathbf{W}_p^{(i,j)} = \begin{cases} 1, & \ell_i \neq \ell_j \wedge \mathcal{D}_i \neq \mathcal{D}_j \\ 0, & \text{otherwise} \end{cases}$$

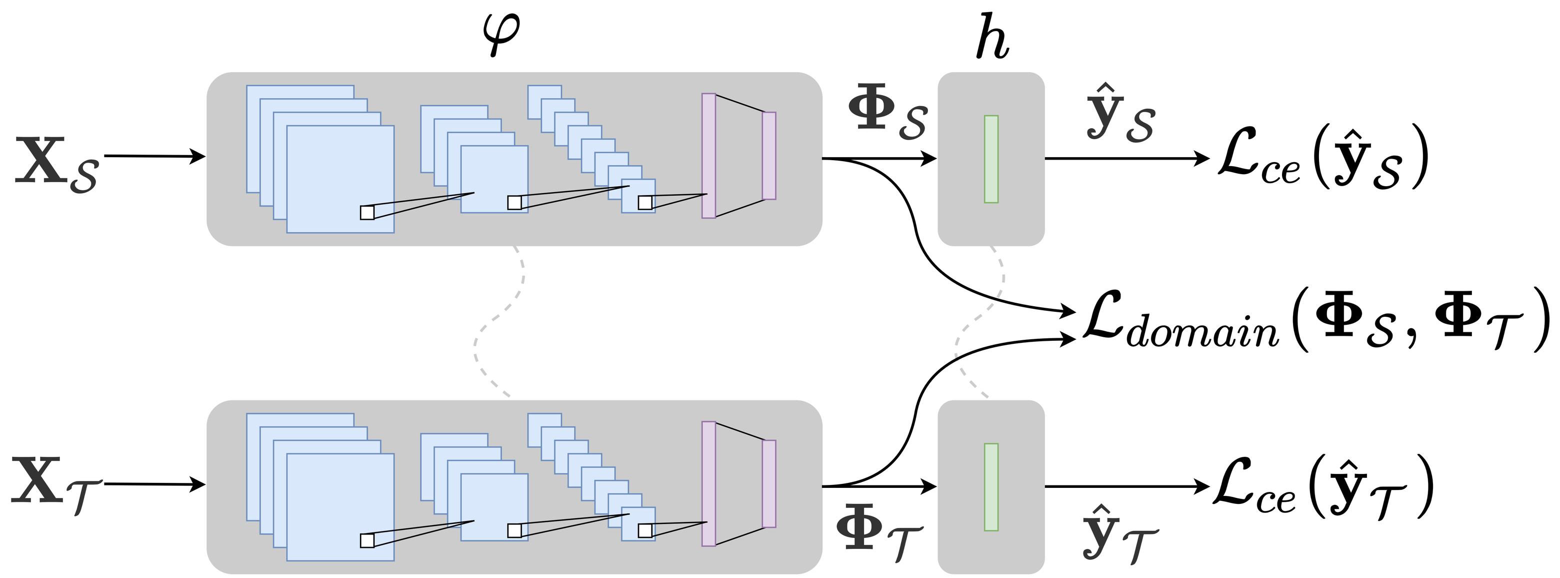

Figure 1: Two-stream network architecture for domain adaptation.

Experiments and Results

	$\mathcal{A} \rightarrow \mathcal{D}$	$\mathcal{A} \rightarrow \mathcal{W}$	$\mathcal{D} \rightarrow \mathcal{A}$	$\mathcal{D} \rightarrow \mathcal{W}$	$\mathcal{W} \rightarrow \mathcal{A}$	$\mathcal{W} \rightarrow \mathcal{D}$	Avg.
FT-Source	66.6 ± 3.0	59.8 ± 2.1	42.8 ± 5.2	92.3 ± 2.8	44.0 ± 0.7	98.5 ± 1.2	67.4
FT-Target	71.4 ± 2.0	74.0 ± 4.9	56.2 ± 3.6	95.9 ± 1.2	50.2 ± 2.6	99.1 ± 0.8	74.5
CCSA	84.8 ± 2.1	87.5 ± 1.5	$\mathbf{66.5 \pm 1.9}$	97.2 ± 0.7	64.0 ± 1.6	98.6 ± 0.4	83.1
d-SNE	$\mathbf{86.5 \pm 2.5}$	$\mathbf{88.7 \pm 1.9}$	65.9 ± 1.1	97.6 ± 0.7	63.9 ± 1.2	99.0 ± 0.5	**83.6**
DAGE-LDA	85.9 ± 2.8	87.8 ± 2.3	66.2 ± 1.4	$\mathbf{97.9 \pm 0.6}$	$\mathbf{64.2 \pm 1.2}$	$\mathbf{99.5 \pm 0.5}$	**83.6**

Table I: Macro average classification accuracy (%) for Office-31 using a VGG16 network pretrained on ImageNet. For each run, 3 samples/class were drawn for the target data, and 8(20) from source data for DSLR, Webcam, (Amazon). The reported results are the mean and standard deviation across five runs.

Samples/class	1	3	5	7
CCSA	$\mathbf{75.6 \pm 2.1}$	$\mathbf{85.0 \pm 1.4}$	87.8 ± 0.7	89.1 ± 0.7
d-SNE	69.0 ± 1.7	80.4 ± 1.7	86.1 ± 0.9	87.7 ± 0.9
DAGE-LDA	67.0 ± 1.9	82.7 ± 1.7	$\mathbf{89.0 \pm 0.8}$	$\mathbf{90.7 \pm 0.5}$

Table II: Classification accuracy (%) for MNIST → USPS with a varying number of target samples per class.

Conclusion

DAGE provides a flexible and interpretable framework for domain adaptation. A simple instantiation, DAGE-LDA, achieves performance matching prior stat-of-the-art methods.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 871449

OpenDeep Robotics