Class Conditional Alignment for Partial Domain Adaptation

Mohsen Kheirandishfard, Fariba Zohrizadeh, and Farhad Kamangar
University of Texas at Arlington

Introduction

Unsupervised domain adaptation (UDA) aims to transfer knowledge from a labeled source domain to an unlabeled target domain by learning domain-invariant feature representations.

Partial domain adaptation (PDA) is a new adaptation task that particularly investigates the scenarios in which the source dataset is large and diverse compared to the target dataset.

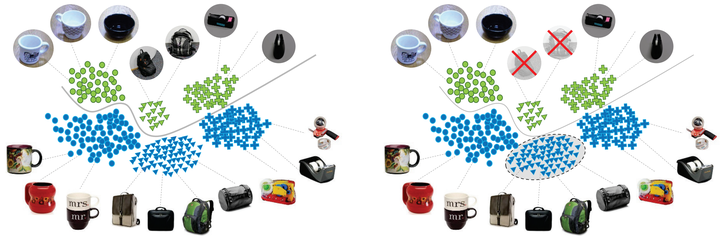

Unsupervised domain adaptation Partial domain adaptation

PDA Approaches:

- **Goal:** promote learning from the shared classes between the domains to enhance target domain classification performance.
- **Assumption:** target domain label space is a subset of the source domain label space.
- **Challenge:** avoid learning from *outlier classes*, i.e. the source domain classes that do not appear in the target domain.

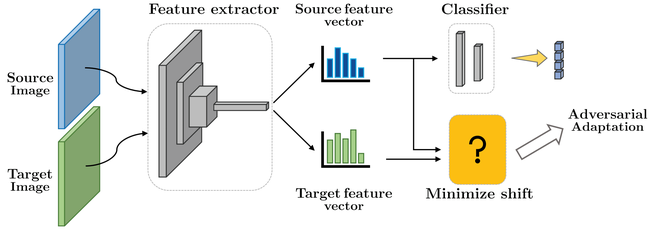

Contribution

Leveraging the same network architecture as partial adversarial domain adaptation (PADA) [1], we jointly align the marginal and class-conditional distributions in the shared label space by minimaxing a novel multi-class adversarial loss function. Furthermore, we incorporate effective regularization terms to encourage selecting the most relevant subset of source domain classes.

Proposed Method

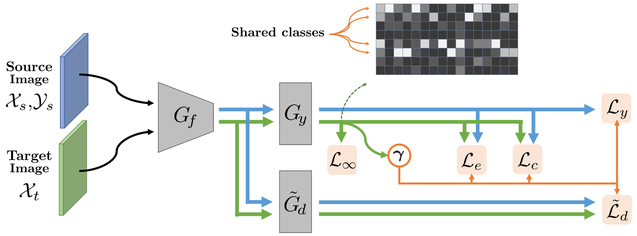

Let $\{(\boldsymbol{x}_s^i, \boldsymbol{y}_s^i)\}_{i=1}^{n_s}$ be the source samples and their labels from $\mathcal{C}_s$ classes, and $\{\boldsymbol{x}_t^i\}_{i=1}^{n_t}$ denote the target samples. PADA minimizes the loss function

$$\max_{\tilde{\boldsymbol{\theta}}_d} \min_{\boldsymbol{\theta}_y, \boldsymbol{\theta}_f} \frac{1}{n_s} \sum_{\boldsymbol{x}^i \in \mathcal{X}_s} \gamma_{c_i} L_y(G_y(G_f(\boldsymbol{x}^i)), \boldsymbol{y}^i) + \lambda \tilde{\mathcal{L}}_d(\boldsymbol{\theta}_f, \tilde{\boldsymbol{\theta}}_d)$$
$$+ \mathcal{L}_c(\boldsymbol{\theta}_f, \boldsymbol{\theta}_y) + \mu \mathcal{L}_\infty(\boldsymbol{\theta}_f, \boldsymbol{\theta}_y) + \zeta \mathcal{L}_e(\boldsymbol{\theta}_f, \boldsymbol{\theta}_y)$$

Weighted centroid alignment regularization $\mathcal{L}_c$ mitigates the adverse effect of falsely-pseudo-labeled target samples by aligning labeled source centroids and pseudo-labeled target centroids in the feature space.

Row-sparsity regularization $\mathcal{L}_\infty$ promotes the selection of a small subset of classes that are in common between the source and target domains.

Minimum entropy regularization $\mathcal{L}_e$ utilizes the output of G_y to down-weight the relative importance of irrelevant samples from both domains.

Experiments

Experiments on Office-31 object dataset consists of $4,652$ images from 31 classes across three domains: *Amazon* (**A**), *Webcam* (**W**), and *DSLR* (**D**).

Amazon (**A**) Digital camera (**D**) Webcam (**W**)

Method	A → W	D → W	W → D	A → D	D → A	W → A	Avg
ResNet	75.59	96.27	98.09	83.44	83.92	84.97	87.05
DANN	73.56	96.27	98.73	81.53	82.78	86.12	86.50
ADDA	75.67	95.38	99.85	83.41	83.62	84.25	87.03
RTN	78.98	93.22	85.35	77.07	89.25	89.46	85.56
IWAN	89.15	99.32	99.36	90.45	95.62	94.26	94.69
SAN	93.90	99.32	99.36	94.27	94.15	88.73	94.96
PADA	86.54	99.32	**100.0**	82.17	92.69	95.41	92.69
ETN	94.52	**100.0**	**100.0**	95.03	**96.21**	94.64	96.73
CCPDA$_\infty$	95.12	99.32	100.0	93.21	96.03	95.19	96.48
CCPDA$_e$	97.45	96.64	100.0	96.47	94.92	93.86	96.56
CCPDA$_{d,c}$	93.42	97.62	100.0	90.43	93.45	95.53	95.07
CCPDA	**99.66**	**100.0**	**100.0**	**97.45**	95.72	**95.71**	**98.09**

Table 1: Accuracy of PDA tasks on *Office-31* (ResNet-50).
Please refer to the paper for more experimental results.

Conclusion

- Our approach adopts a multi-class adversarial loss function to jointly align the marginal and class-conditional distributions across the shared classes between the source and target domains.
- The regularization terms reduce the effects of outlier classes and can be directly incorporated into many adversarial architectures.
- Experiments on a benchmark dataset demonstrate the high potential of our approach for different PDA tasks.

References

[1] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *ECCV*, 2018.