Semi-Supervised Domain Adaptation with Instance Constraints

Jeff Donahue[1], Judy Hoffman[1], Erik Rodner[1], Kate Saenko[2], Trevor Darrell[1]
[1]University of California, Berkeley [2]University of Massachusetts Lowell

Problem

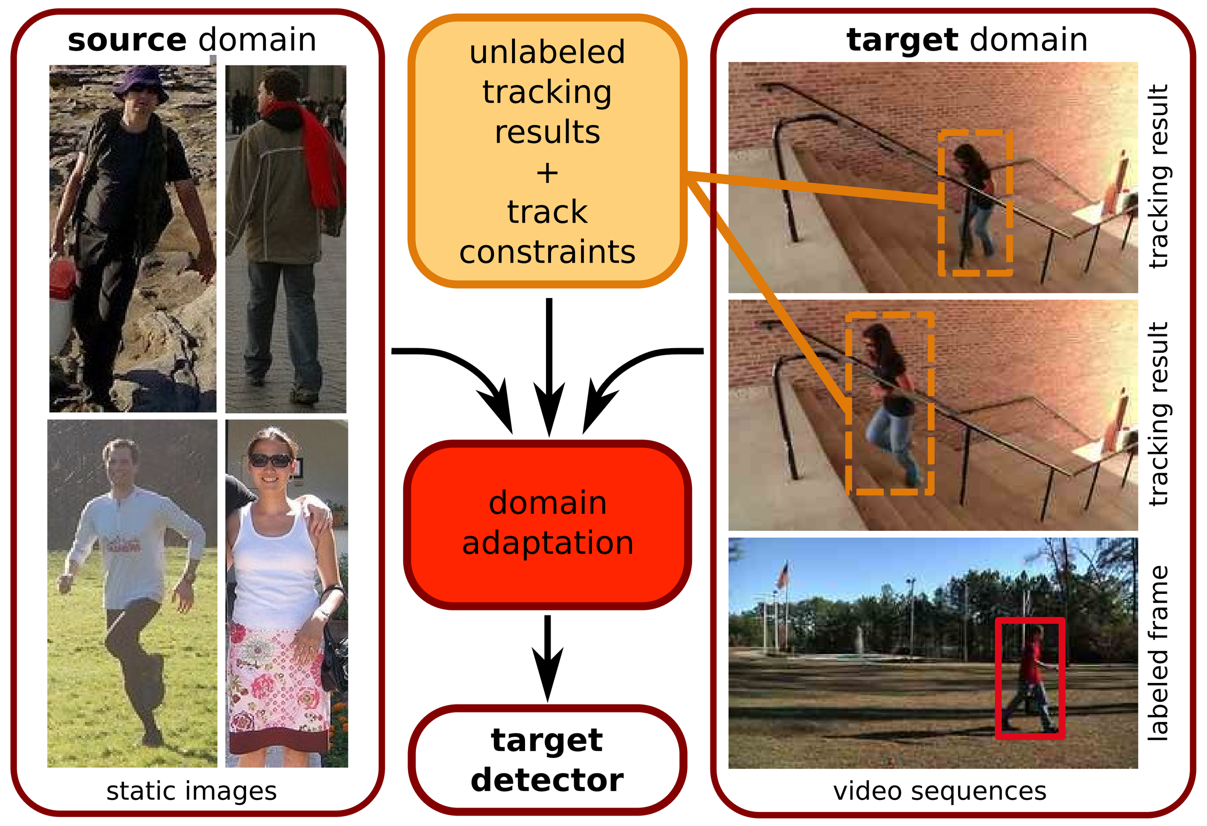

- Classification or detection in a *target domain* with little labeled data, but with a fully labeled *source domain*
- **Auxiliary information about the target domain is available in the form of *instance correspondences***
- We demonstrate this information can be useful in performing adaptation

Adaptation Formulation

- Traditional semi-supervised domain adaptation inputs:
  - Labeled source data $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}$
  - Labeled target data $\tilde{\mathcal{D}}^L = \{(\tilde{\mathbf{x}}_j, \tilde{y}_j)\}$ (usually, $|\tilde{\mathcal{D}}^L| << |\mathcal{D}|$)
  - Unlabeled target data $\tilde{\mathcal{D}}^U = \{\tilde{\mathbf{x}}_j^U\}$
- Goal: learn linear classifier functions
  - $f(\mathbf{x}) = \boldsymbol{\theta}^T\mathbf{x}$ (source domain)
  - $\tilde{f}(\tilde{\mathbf{x}}) = \tilde{\boldsymbol{\theta}}^T\tilde{\mathbf{x}}$ (target domain)
- Many successful domain adaptation methods have objectives of the form

$$\min_{\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}, \mathbf{A}} \; \mathcal{R}(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}, \mathbf{A}) + C \cdot \mathcal{L}(\mathcal{D}, \boldsymbol{\theta}) + \tilde{C} \cdot \tilde{\mathcal{L}}(\tilde{\mathcal{D}}^L, \tilde{\boldsymbol{\theta}})$$

  - Optional transformation parameter $\mathbf{A}$

Adaptation Methods

- PMT-SVM [1] for adaptation from a single source domain category classifier $\boldsymbol{\theta}$
  - Learns a target hyperplane $\tilde{\boldsymbol{\theta}}$ by regularizing the angle $\alpha_{\tilde{\boldsymbol{\theta}},\boldsymbol{\theta}}$ between the source and target hyperplanes
  - $\mathcal{R}^{\text{pmt}}(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}) = \frac{1}{2}\|\tilde{\boldsymbol{\theta}}\|_2^2 + \frac{\Gamma}{2}\|\tilde{\boldsymbol{\theta}}\|_2^2 \sin^2 \alpha_{\tilde{\boldsymbol{\theta}},\boldsymbol{\theta}}$
- MMDT [4] for multi-category transform-based adaptation
  - Jointly learns a linear domain transform $\mathbf{A}$ and source domain hyperplane classifier parameters $\boldsymbol{\theta}$
  - Implicit target hyperplane $\tilde{\boldsymbol{\theta}} := \mathbf{A}^T\boldsymbol{\theta}$
  - $\mathcal{R}^{\text{mmdt}}(\boldsymbol{\theta}, \mathbf{A}) = \frac{1}{2}\|\boldsymbol{\theta}\|_2^2 + \frac{\gamma}{2}\|\mathbf{A} - \mathbf{I}\|_F^2$

Instance Constraints

- Instance constraints – given as sparse edge weight matrix $\mathbf{B}$, with elements $\beta_{j,j'}$ defining the similarity between each pair of target examples $\tilde{\mathbf{x}}_j, \tilde{\mathbf{x}}_{j'}$

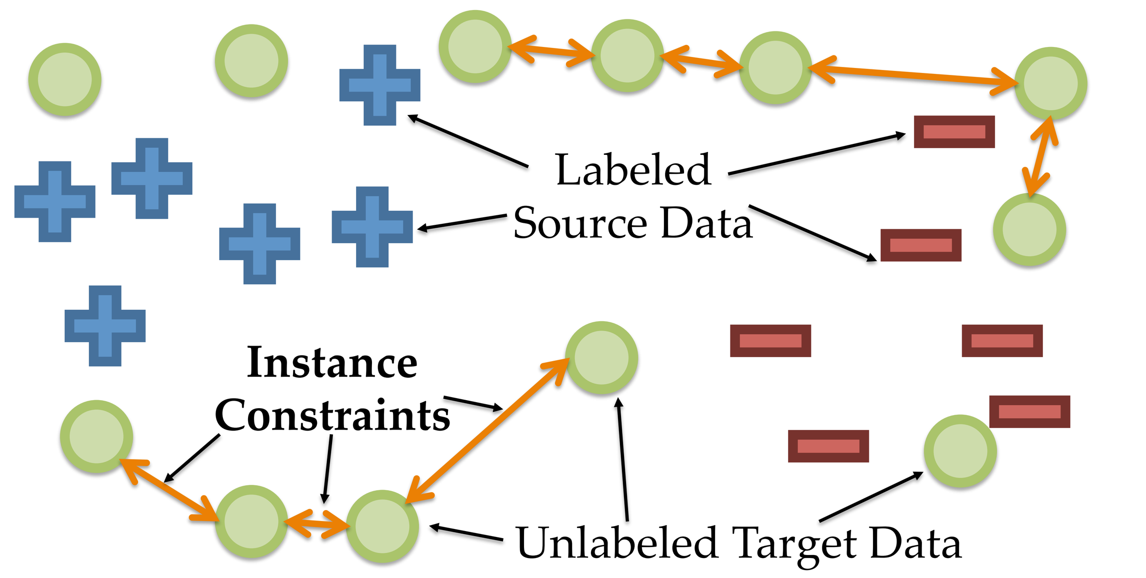

- Define Laplacian matrix $\mathbf{L} = \mathbf{D} - \mathbf{B}$, where $\mathbf{D}$ is the diagonal matrix that contains the row sums of $\mathbf{B}$
- (Convex) regularization term incorporating similarity constraints defined as

$$\begin{aligned} r(\tilde{\mathcal{D}}^U, \tilde{f}_{\tilde{\boldsymbol{\theta}}}) &= \tfrac{1}{2}\tilde{\mathbf{f}}^T\mathbf{L}\tilde{\mathbf{f}} \\ &= \tfrac{1}{2}\left(\tilde{\boldsymbol{\theta}}^T\tilde{\mathbf{X}}^U\right)^T \mathbf{L} \left(\tilde{\boldsymbol{\theta}}^T\tilde{\mathbf{X}}^U\right) \\ &= \sum_{j \neq j'} \beta_{j,j'} \left(\tilde{f}(\tilde{\mathbf{x}}_j^U) - \tilde{f}(\tilde{\mathbf{x}}_{j'}^U)\right)^2 \end{aligned}$$

DA with Correspondences

- Integrate the Laplacian regularizer into the domain adaptation framework

$$\min_{\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}, \mathbf{A}} \; \mathcal{R}(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}, \mathbf{A}) + C \cdot \mathcal{L}(\mathcal{D}, \boldsymbol{\theta}) + \tilde{C} \cdot \tilde{\mathcal{L}}(\tilde{\mathcal{D}}^L, \tilde{\boldsymbol{\theta}}) + r(\tilde{\mathcal{D}}^U, \tilde{f}_{\tilde{\boldsymbol{\theta}}})$$

- Classifiers learned from this objective
  - Maximize the margin
  - Classify the labeled data accurately
  - **Classify similar unlabeled target data similarly**

Classification

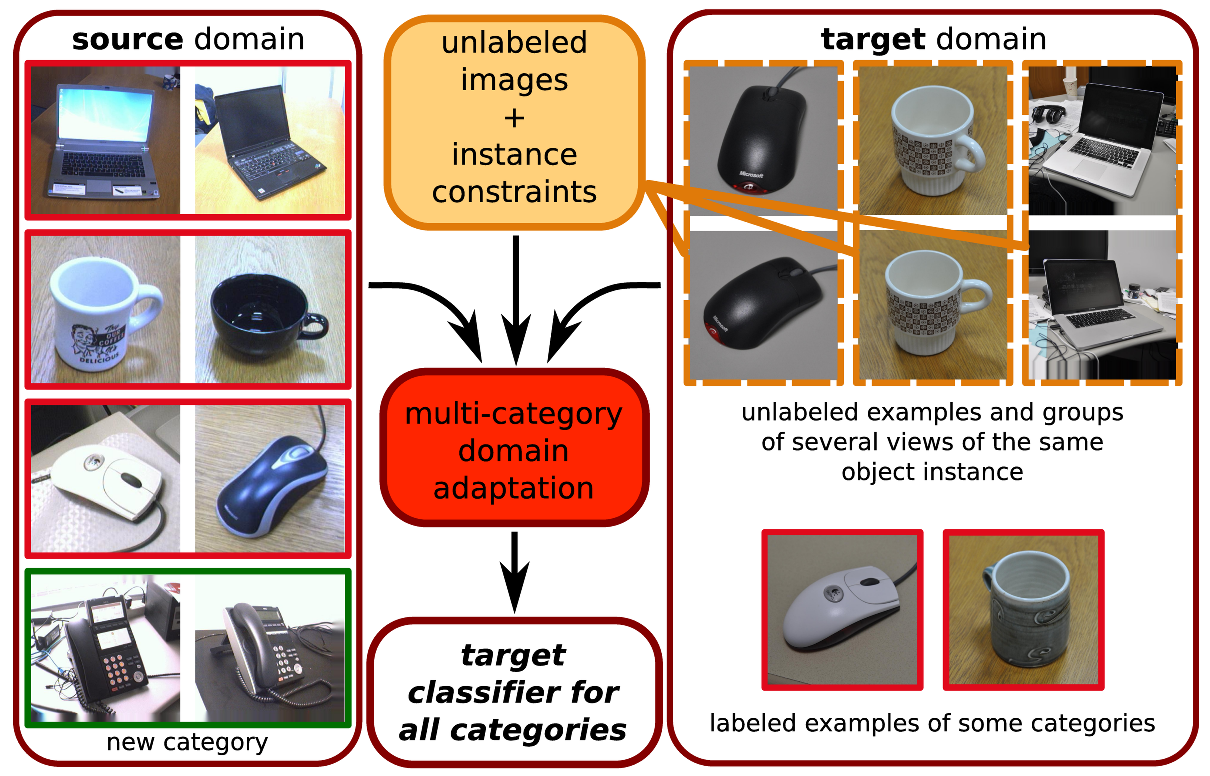

- **Multi-category classification** setting
  - Can use weights $\beta_{j,j'}$ to encode confidence that two images show the same object instance
- Experiments on the *Office* dataset [5]
- Use multi-category MMDT objective [4] as baseline (**da only**); Laplacian-extended version as our method (**da + lap-sim**)
- $\beta_{j,j'} = 1$ if $\tilde{\mathbf{x}}_j$ and $\tilde{\mathbf{x}}_{j'}$ show the same object instance; 0 otherwise
- 8 labeled examples per category in source domain; 1 in target domain

svm$_S$	**svm$_{S \cup T}$**	**da only**	**da + lap-sim**
45.80 ± 2.2	50.79 ± 2.4	54.57 ± 2.2	$\mathbf{56.15 \pm 2.7}$

Detection

- **Video detection** setting (pedestrians)
  - Weights $\beta_{j,j'}$ can be used to encode confidence that a pair of examples belong to the same (unlabeled) track
- Experiments on the *Pascal VOC* [2] (source) and *VisInt* [6] (target) datasets
- Framework integrated into DPM [3] detector
- Use PMT-SVM objective [1] as baseline (**da only**); Laplacian-extended version as our method (**da + lap-sim**)
- Let $\delta_{j,j'}$ be the frame index difference between examples $\tilde{\mathbf{x}}_j$ and $\tilde{\mathbf{x}}_{j'}$ if the examples come from the same track; 0 otherwise
- $\beta_{j,j'} = 1/\delta_{j,j'}$ if $\delta_{j,j'} > 0$; 0 otherwise
- 5011 labeled source images; N_f target

N_f	**dpm$_S$**	**dpm$_{S \cup T}$**	**da only**	**da + lap-sim**
10	0.3220	**0.3502**	0.1121	*0.3317*
20	0.3220	0.3473	*0.3530*	**0.3913**
30	0.3220	0.3816	**0.4306**	**0.4303**
40	0.3220	0.4314	*0.4538*	**0.4584**

da only **da + lap-sim**

References

[1] Y. Aytar and A. Zisserman. Tabula rasa: Model transfer for object category detection. In *ICCV*, 2011.

[2] M. Everingham, L. Van Gool, C. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes (VOC) challenge. *IJCV*, 88(2):303–338, 2010.

[3] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part based models. *PAMI*, 32(9):1627–1645, 2010.

[4] J. Hoffman, E. Rodner, J. Donahue, T. Darrell, and K. Saenko. Efficient learning of domain-invariant image representations. In *ICLR*, 2013.

[5] B. Kulis, K. Saenko, and T. Darrell. What you saw is not what you get: Domain adaptation using asymmetric kernel transforms. In *CVPR*, 2011.

[6] K. Saenko, B. Packer, C.-Y. Chen, S. Bandla, Y. Lee, Y. Jia, J.-C. Niebles, D. Koller, L. Fei-Fei, K. Grauman, and T. Darrell. Mid-level features improve recognition of interactive activities. Technical Report UCB/EECS-2012-209, EECS Department, University of California, Berkeley, 2012.