Comparing Robustness and Curriculum-based strategies for Active Visual Domain Adaptation

Viraj Prabhu[1], Judy Hoffman[1]

[1] Georgia Tech

1 Overview

Motivation

- Often, we want to deploy model trained on data-rich source domain in low-resource target
- We might be able to get some data annotated on our target domain, but all data is not equal
- How do we pick datapoints to get labeled optimally?

Problem Setup

Active Visual Domain Adaptation

- Labeled source data
- Unlabeled target data
- Small target budget

Active Learning (AL) + Domain Adaptation (DA)

Goal: Develop best target model

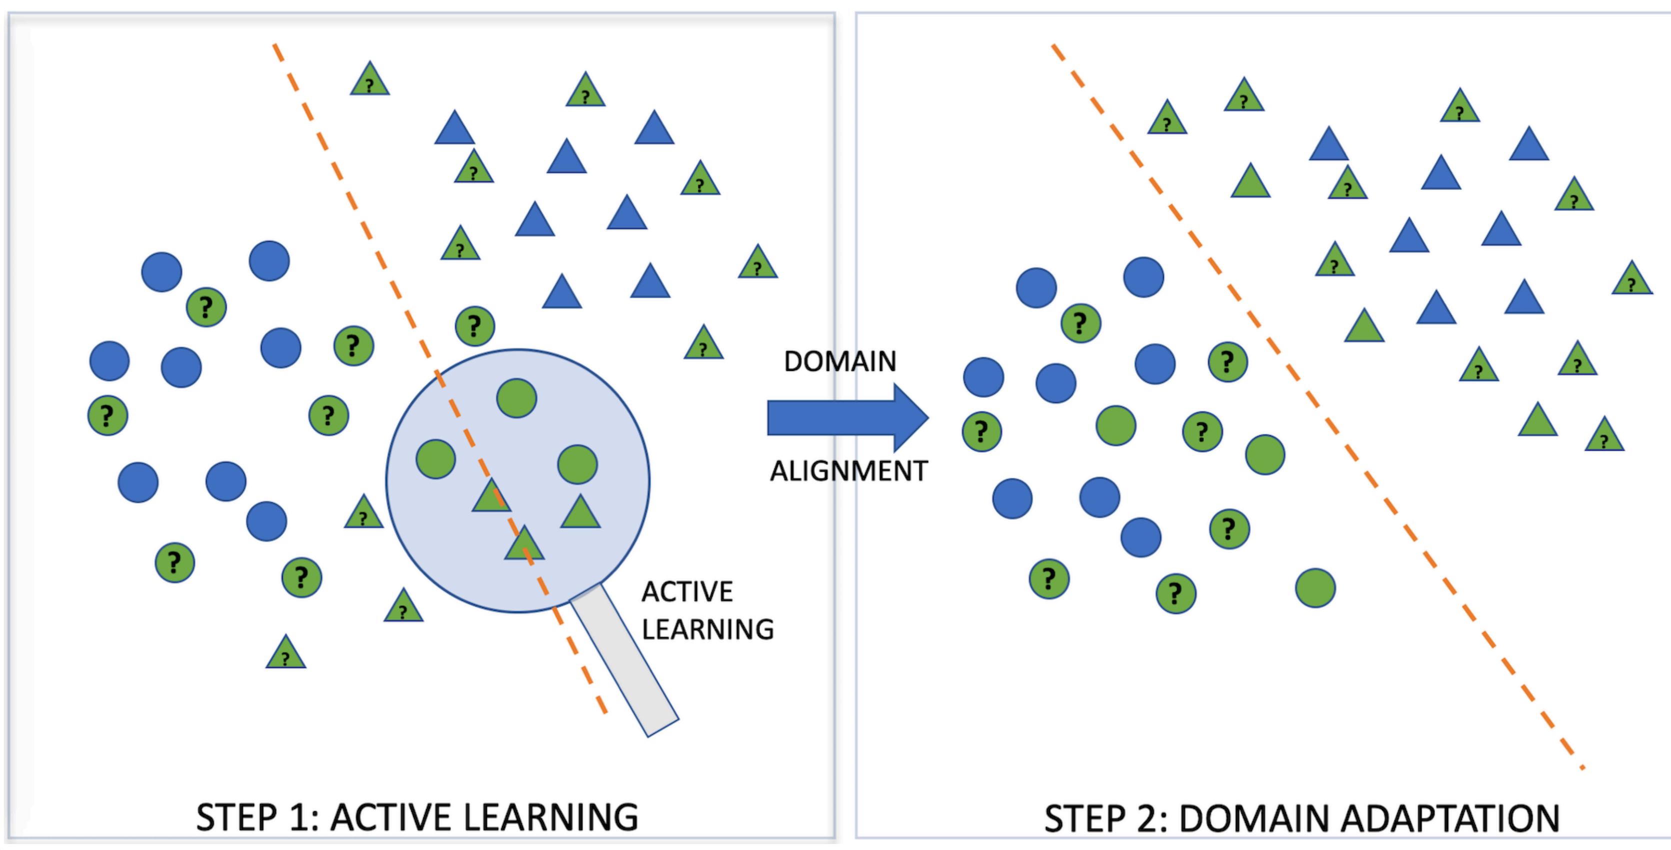

2 Setup

Dataset

- PACS [1]: Photo, Art, Cartoon, Sketch
- 7 way-classification
- We focus on single source adaptation from Photo->Art

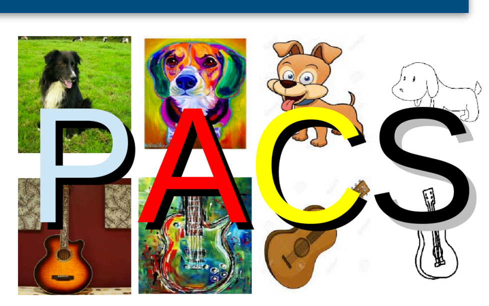

9991 images
90-10 train/test
Challenging domain shift

Approach

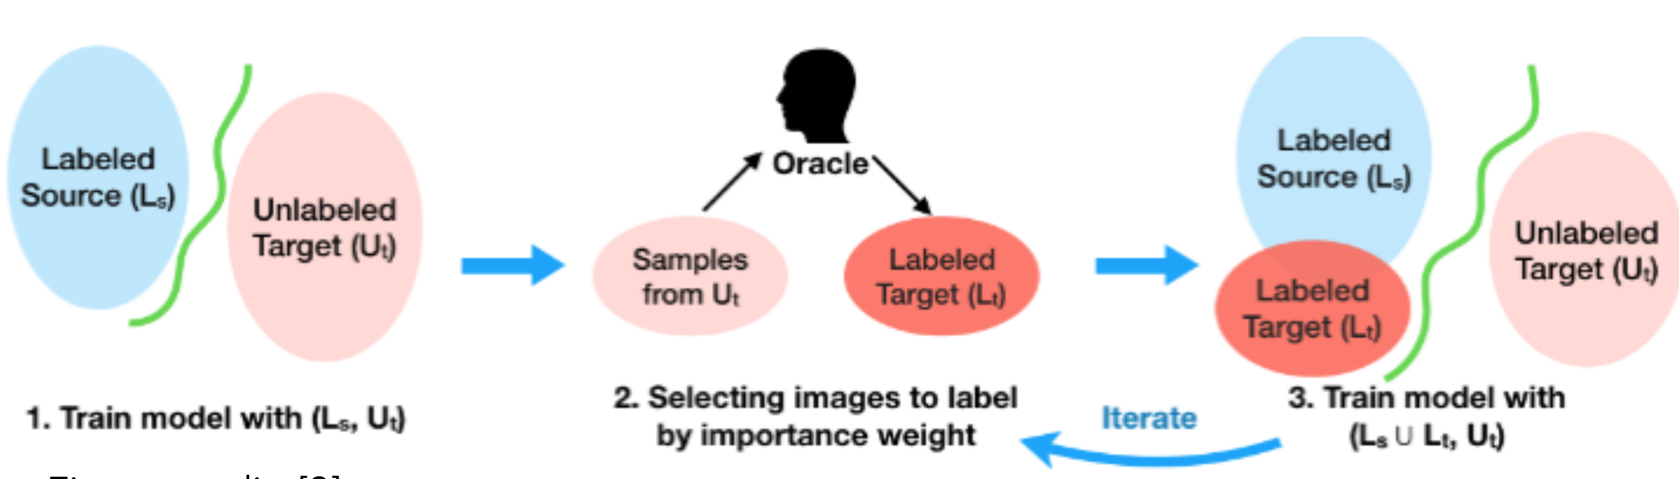

Figure credit: [3]

We study two orthogonal directions for active adaptation

- Is adversarial robustness a good AL heuristic?
- Do curriculums over heuristics help?

Idea #1: Adversarial Robustness

- Adversarial robustness indicates local distributional smoothness
- Are easily perturbable points good candidates for AL?
- Caveat: No access to GT labels for adversarial attack
- Solution: Virtual adversarial attack [2], i.e. find perturbation that maximally changes current output distribution

$$\arg\max_{x} \underbrace{D_{KL}\left(P_\theta(x) \| P_\theta\left(x+\epsilon_L\right)\right)}_{\text{low robustness at large perturbation}} - \underbrace{D_{KL}\left(P_\theta(x) \| P_\theta\left(x+\epsilon_S\right)\right)}_{\text{high robustness at low perturbation}}$$

Idea #2: Curriculum Strategies

- Unsupervised DA often leads to misalignments
- Can curriculums that transition from easier to harder examples over rounds help alignment?

3 Results

Baselines

- Traditional AL: Uniform, Geometry, Top-K Entropy (K=2, N)
- Current SoTA: Active Adversarial Domain Adaptation (AADA) [3]
  - Heuristic: Entropy x "Targetness"

Ours

- robustness: Pick least robust examples at small ϵ
- robustness+curriculum: Anneal ϵ over rounds
- topk-entropy+curriculum: Vary K from 2 to N over rounds

Model: ResNet-18, trained on source train, unsupervised DA [4] to target

At each round t

- Active learning: Pick 50 examples to get labeled from target train
- Domain Adaptation: Adapt to target using labeled+unlabeled data

Metrics: Accuracy on target test set vs % labels from target train set

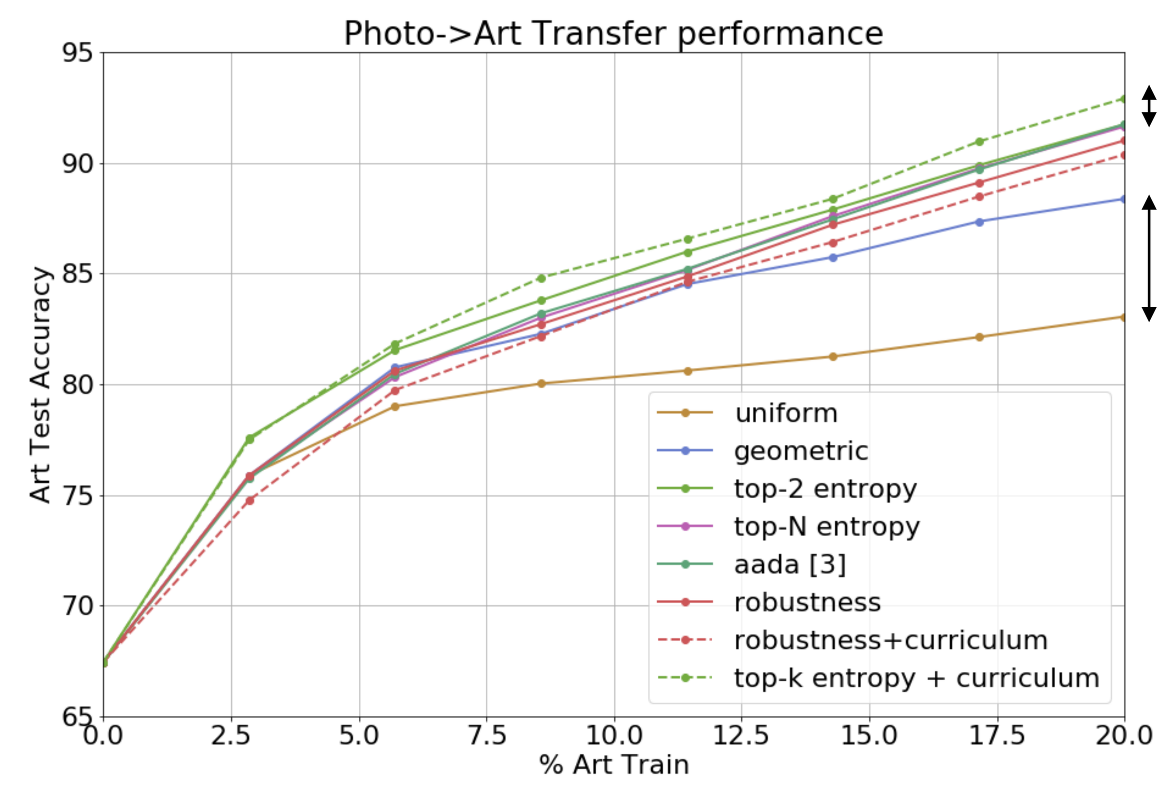

Curriculum over Top-K entropy outperforms all baselines

Results

- Robustness-based heuristic underperforms under margin, entropy; adding a curriculum on top does not help
- A simple curriculum over entropy however achieves SoTA performance

Analysis of topk entropy + curriculum at Round 1

- Person class is the most misaligned
- Top-K entropy transitions from easy to hard examples as desired
- Label histograms show directed strategies

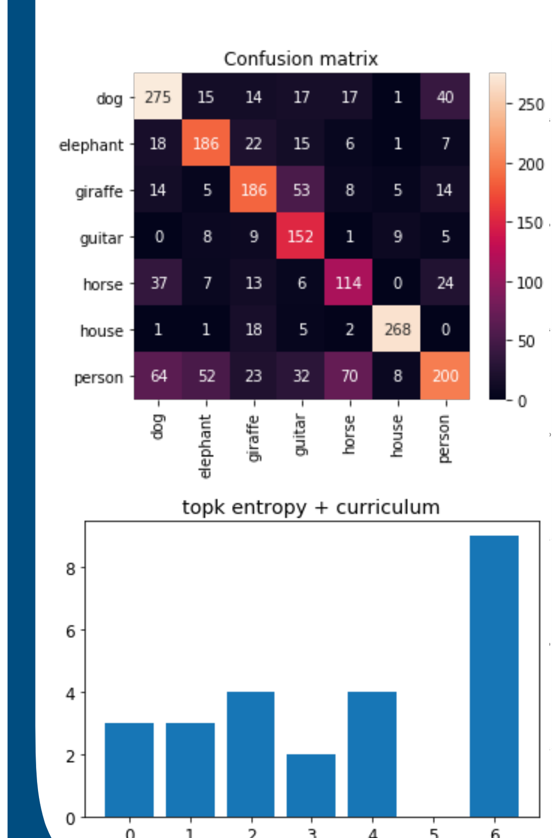

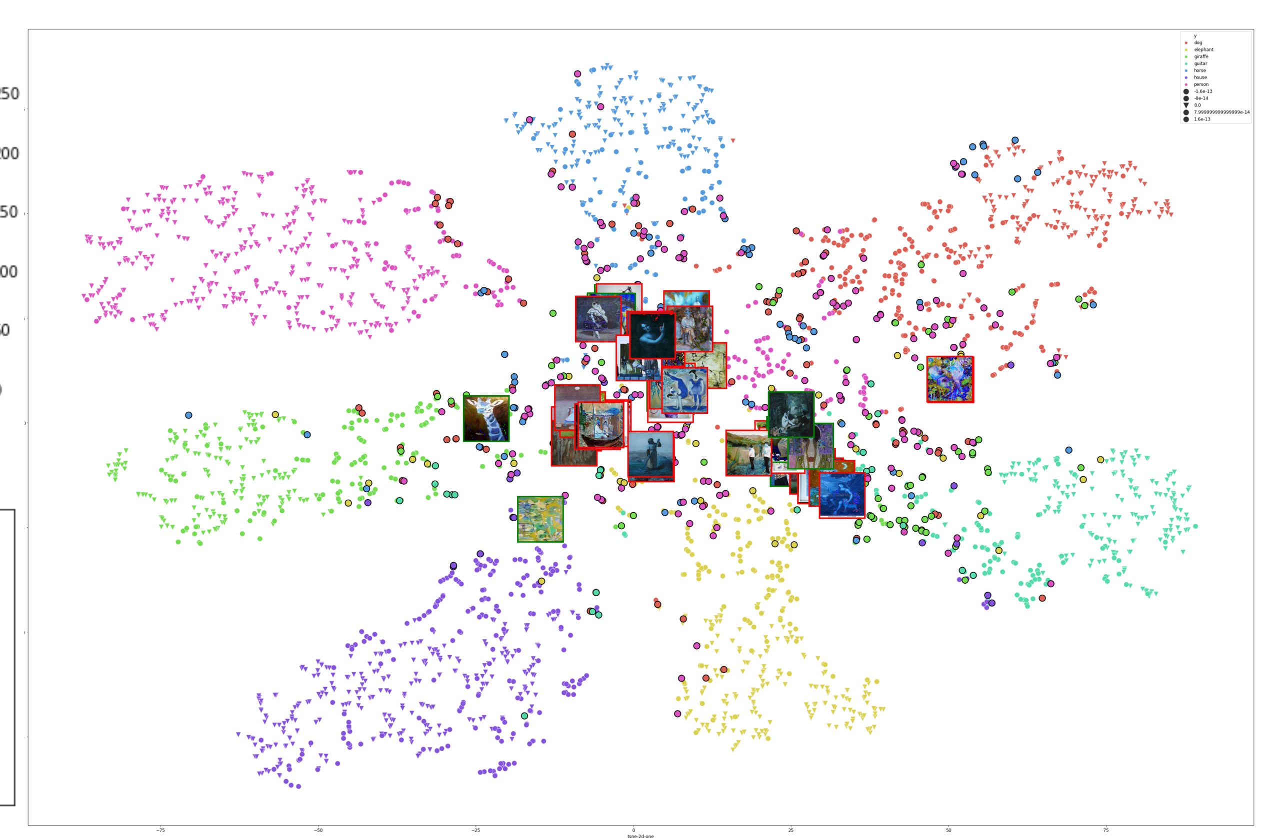

4 Summary

- We experiment with adversarial robustness-based heuristics for active adaptation, but do not see performance gains
- However, curriculums from easier to harder examples over rounds help

Future Work

- Adaptive curriculum strategies
- Label propagation
- Detecting systematic misalignments

References

[1] Li, Da, et al. "Deeper, broader and artier domain generalization." Proceedings of the IEEE International Conference on Computer Vision. 2017.
[2] Miyato, Takeru, et al. "Virtual adversarial training: a regularization method for supervised and semi-supervised learning." IEEE TPAMI 41.8 (2018)
[3] Su, Jong-Chyi, et al. "Active Adversarial Domain Adaptation." *arXiv preprint arXiv: 1904.07848* (2019).
[4] Tzeng, Eric, et al. "Adversarial discriminative domain adaptation." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2017.