BOSTON UNIVERSITY

OVANet: One-vs-All Network for Universal Domain Adaptation

Kuniaki Saito[1] and Kate Saenko[1,2]

1: Boston University, 2: MIT-IBM Watson AI Lab

(Under Review)

What is a dataset bias ?

Semantic segmentation

Images

Labels

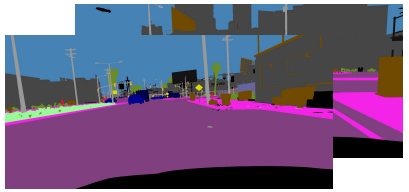

Input

Predicted

Image Classification

Object Detection

Domain Adaptation (DA)

- Goal
  - Transfer knowledge from source (label-rich) to target (label-limited) domain
- Unsupervised Closed Set Domain Adaptation
  - Labeled source and unlabeled target samples
  - Common assumption: categories are completely shared. (Closed Set DA)

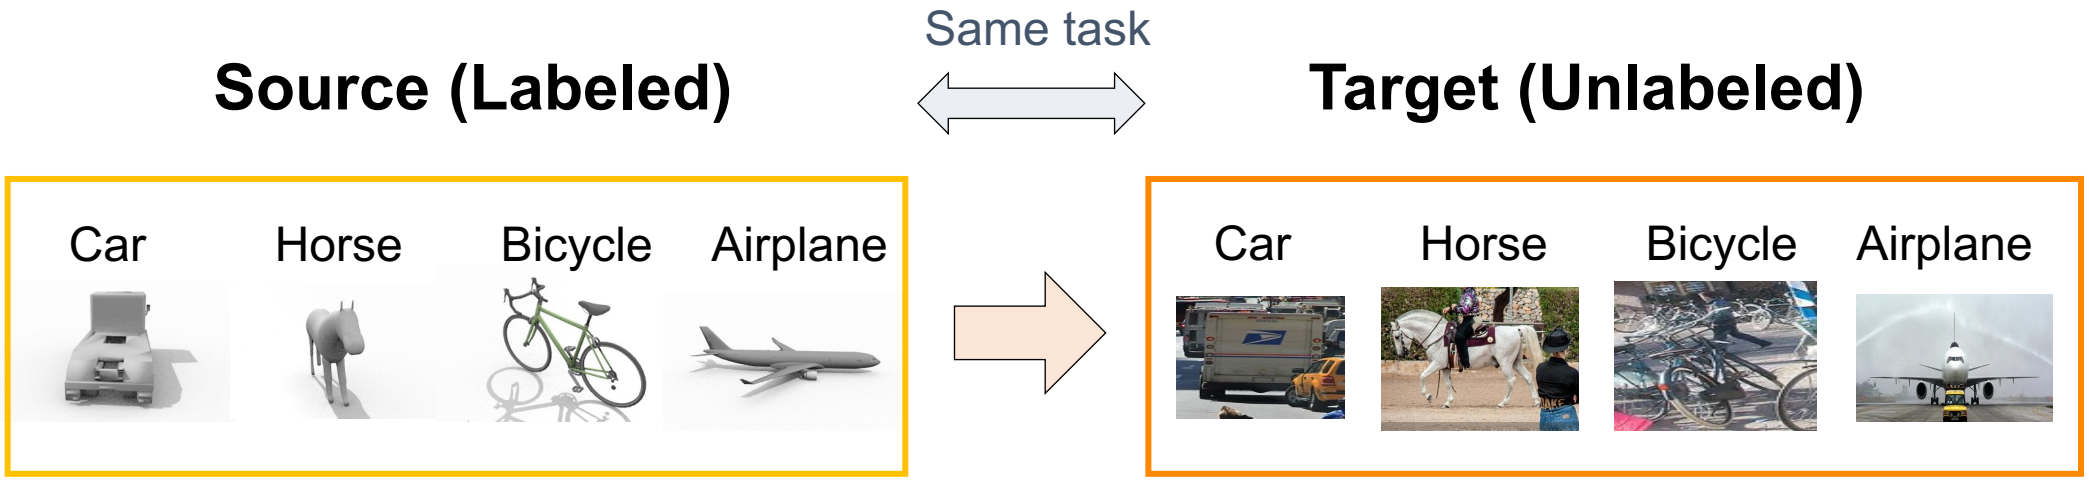

Various category shifts in DA

Source
C_S

Target
C_T

Closed Set DA
$C_S = C_T$

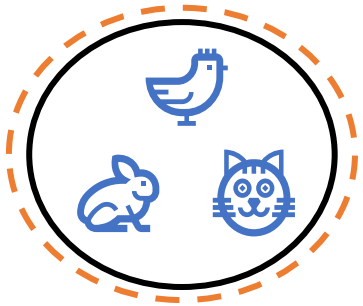

Partial DA
$C_S \supset C_T$

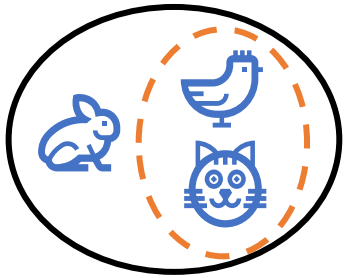

Open Set DA
$C_S \subset C_T$

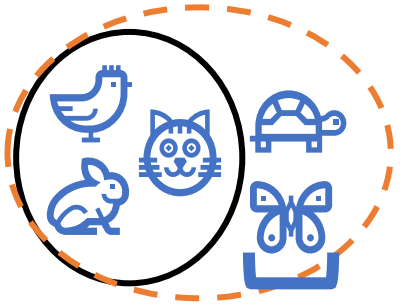

Unknown
Categories

Open-Partial DA

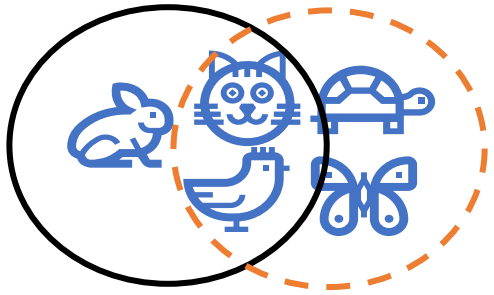

・Open set DA
Known samples => “known categories”
Unknown sample => “unknown!”
・Methods are designed for each setting.

Difficulty
UDA can be any setting.
We cannot know the setting.

Various category shifts in DA

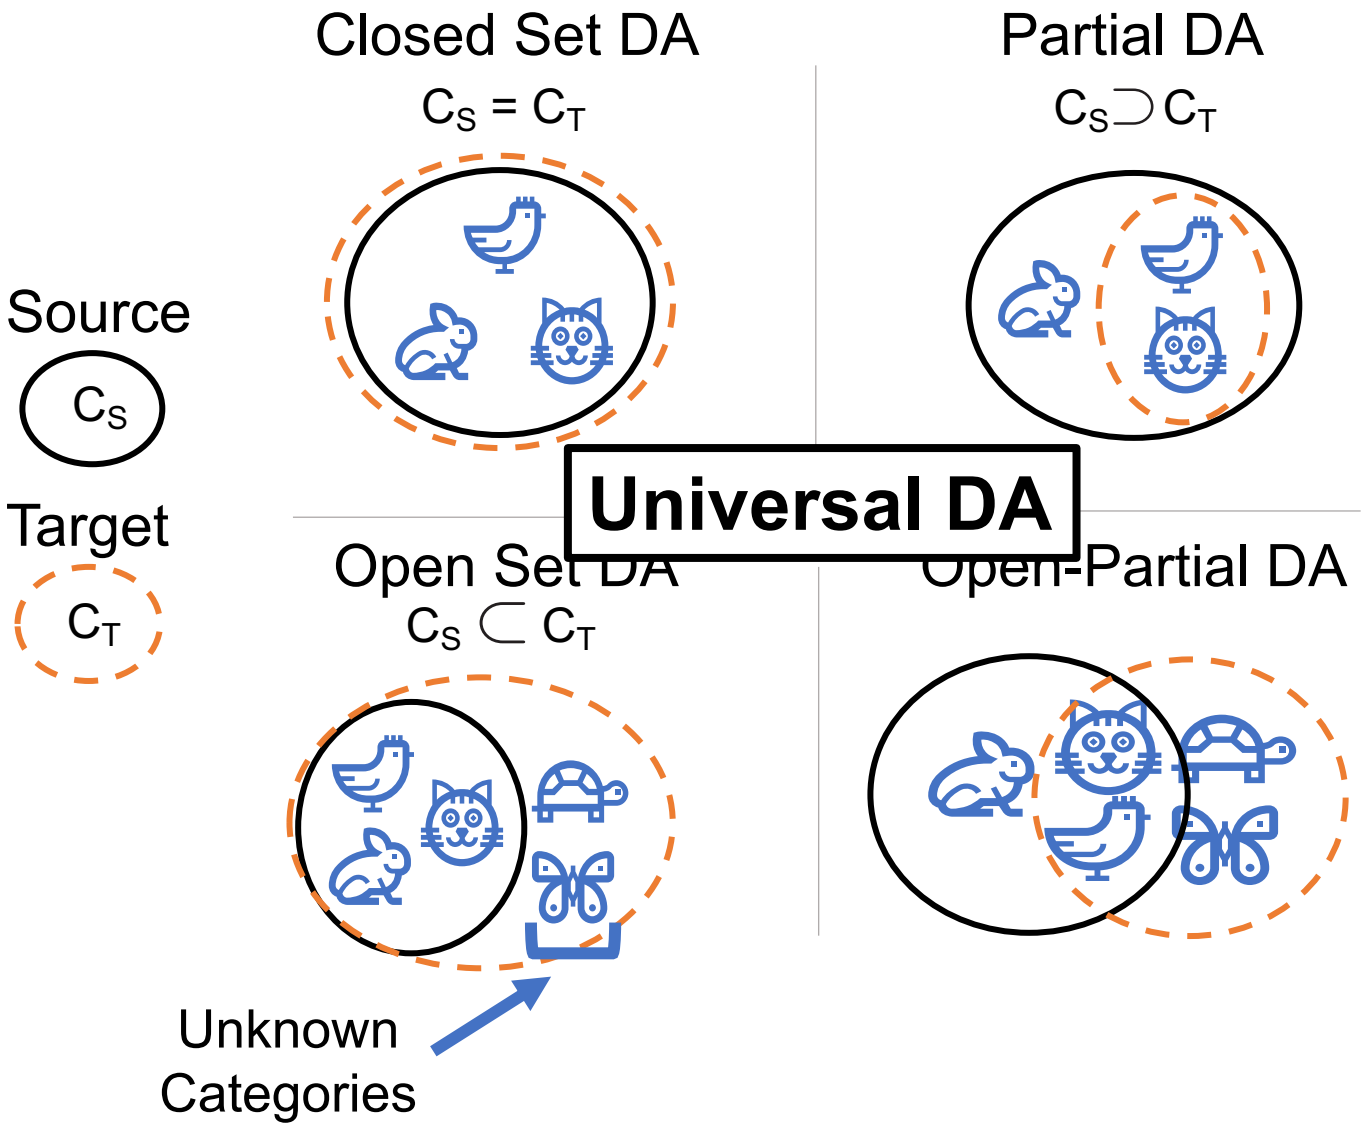

・Open set DA
Known samples => “known categories”
Unknown sample => “unknown!”
・Methods are designed for each setting.

Difficulty
UDA can be any setting.
We cannot know the setting.

Difficulty of Universal Domain Adaptation

- Rejection of unknown target samples
  - No supervision is available, e.g., labeled target samples, ratio of unknown
- Existing approaches: Confidence thresholding
  - Set a threshold by validation or ratio of unknown

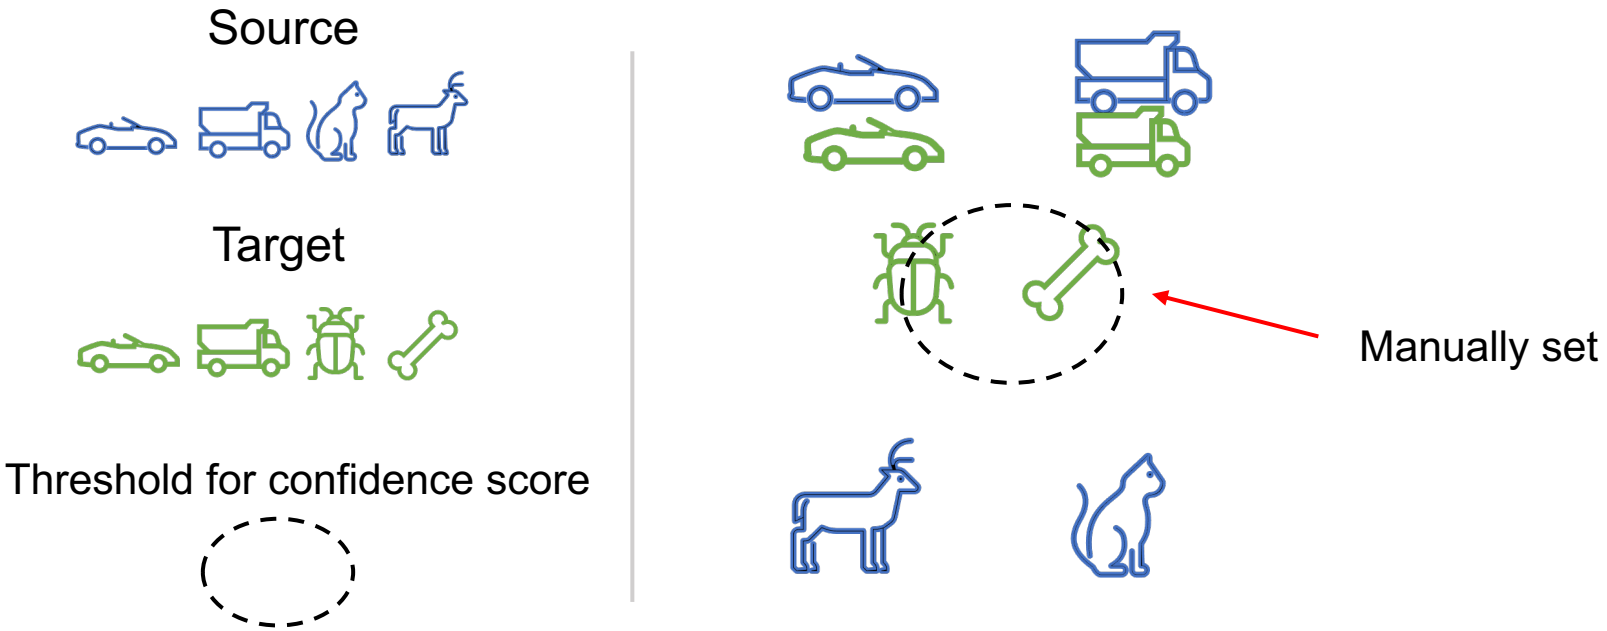

Learn the threshold by source samples

- Distance between classes can be a good threshold.
- Train a one-vs-all classifier for each class.

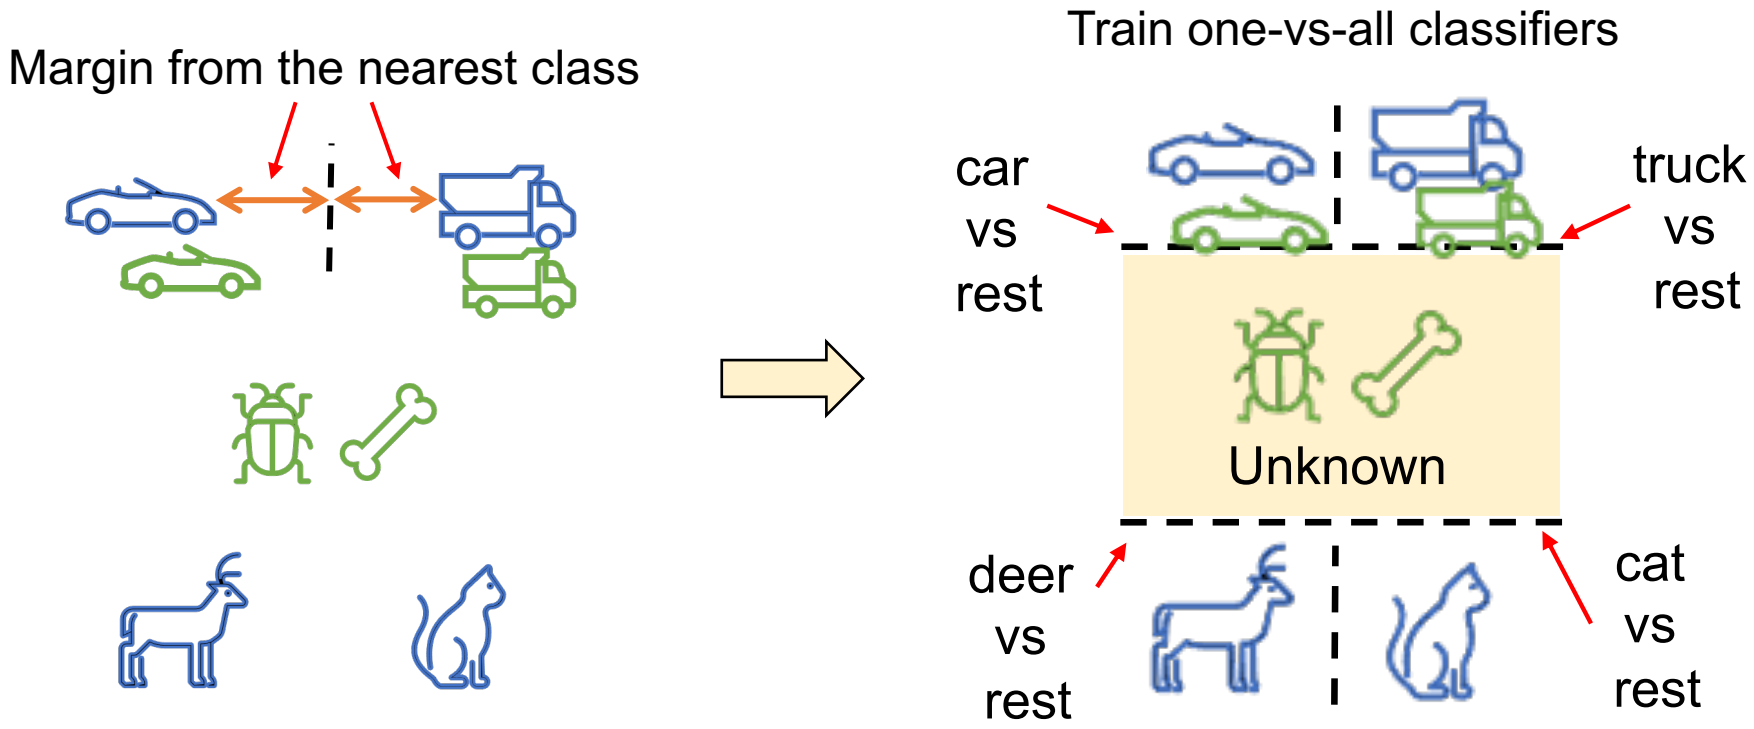

Training One-vs-All Classifiers

- C one-vs-all classifiers (C = number of source classes)
- Learn the margin from the nearest negative class (samples).

Option: Train all one-vs-all classifiers for each input

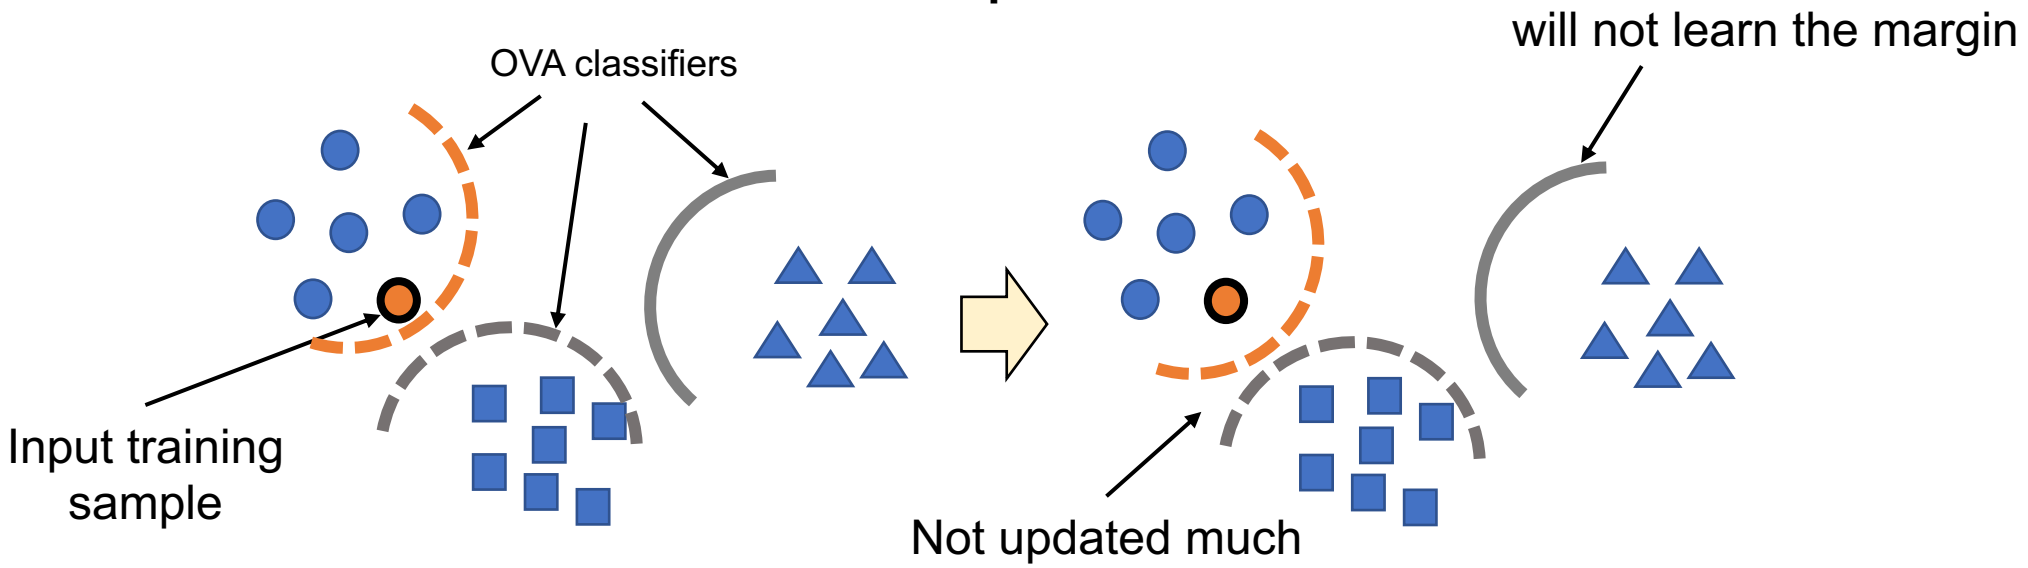

Training One-vs-All Classifiers

- C one-vs-all classifiers (C = number of source classes)
- Learn the margin from the nearest negative class (samples).

Our solution: Train a positive and nearest negative class's one-vs-all classifier

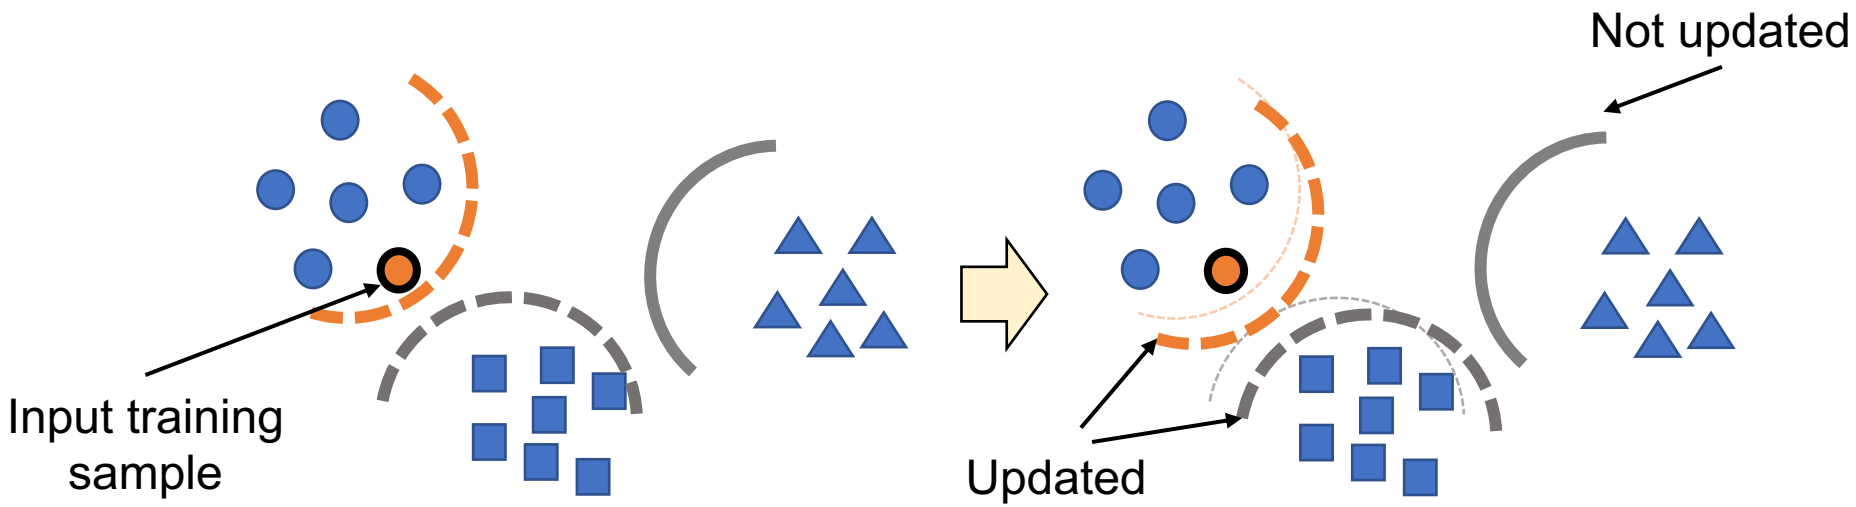

Training One-vs-All Classifiers

1. Train a one-vs-all classifier of the corresponding class

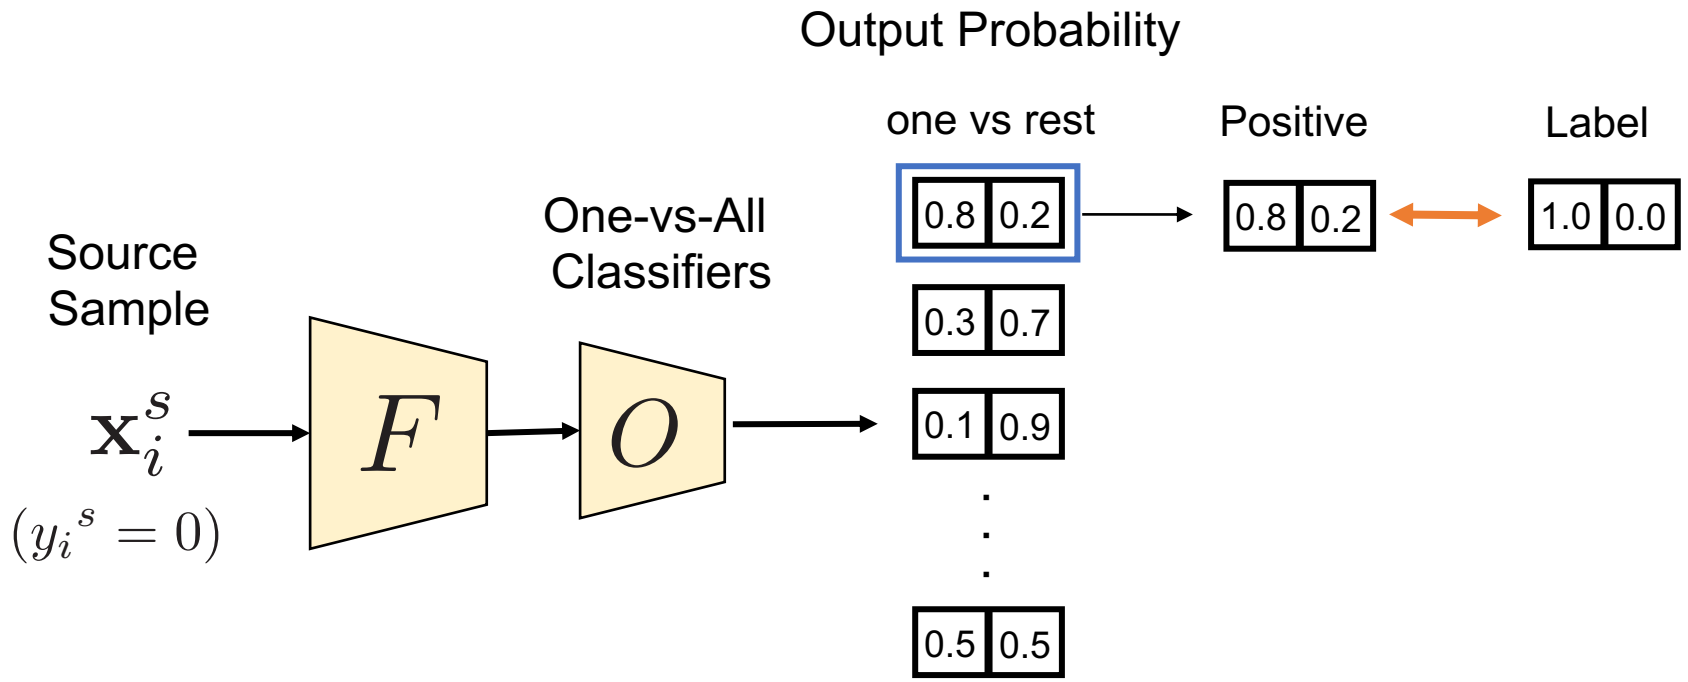

Training One-vs-All Classifiers

2. Train a one-vs-all classifier of the nearest negative class

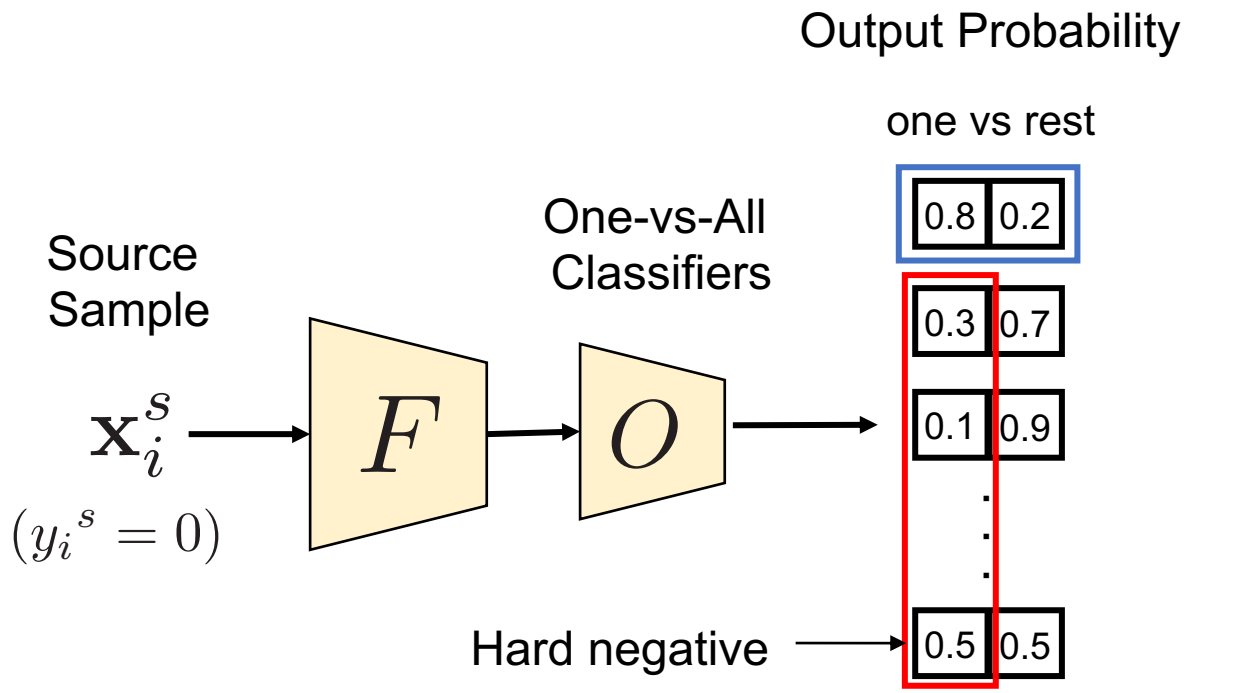

Which class is the nearest?

Training One-vs-All Classifiers

2. Train a one-vs-all classifier of the nearest negative class

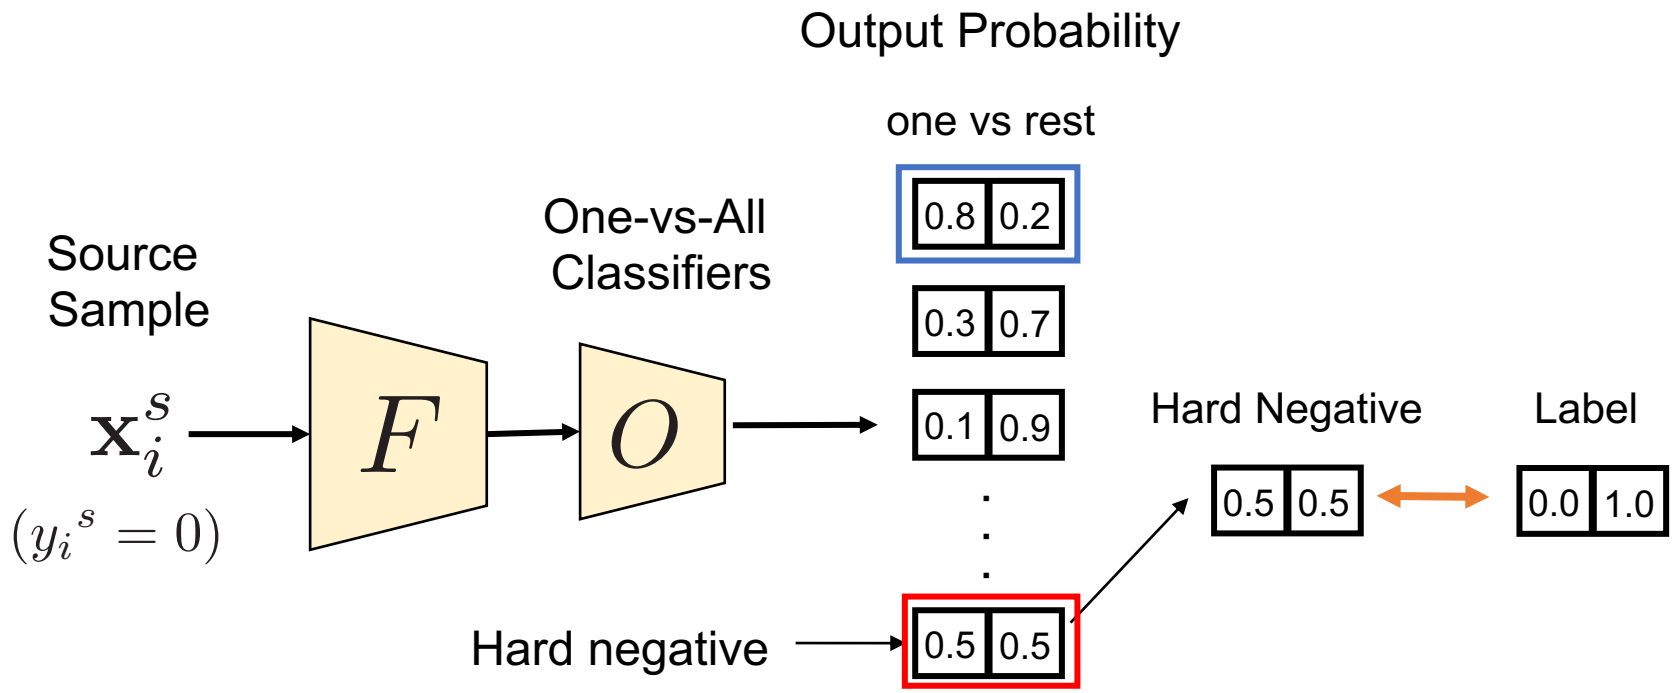

Training One-vs-All Classifiers

Overview

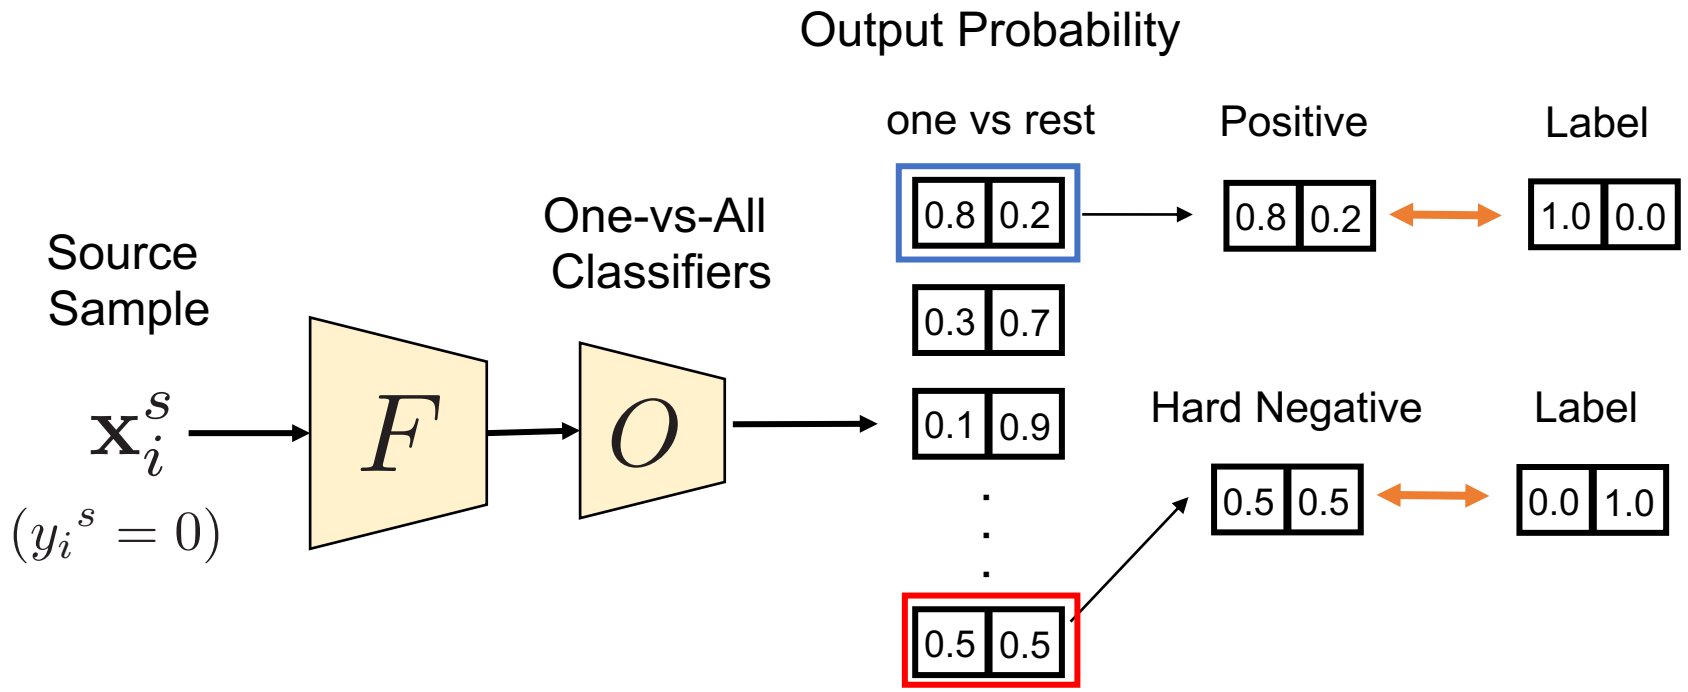

Adaptation to a target domain

- Target samples are from different domains.
- Need to separate from the boundary.

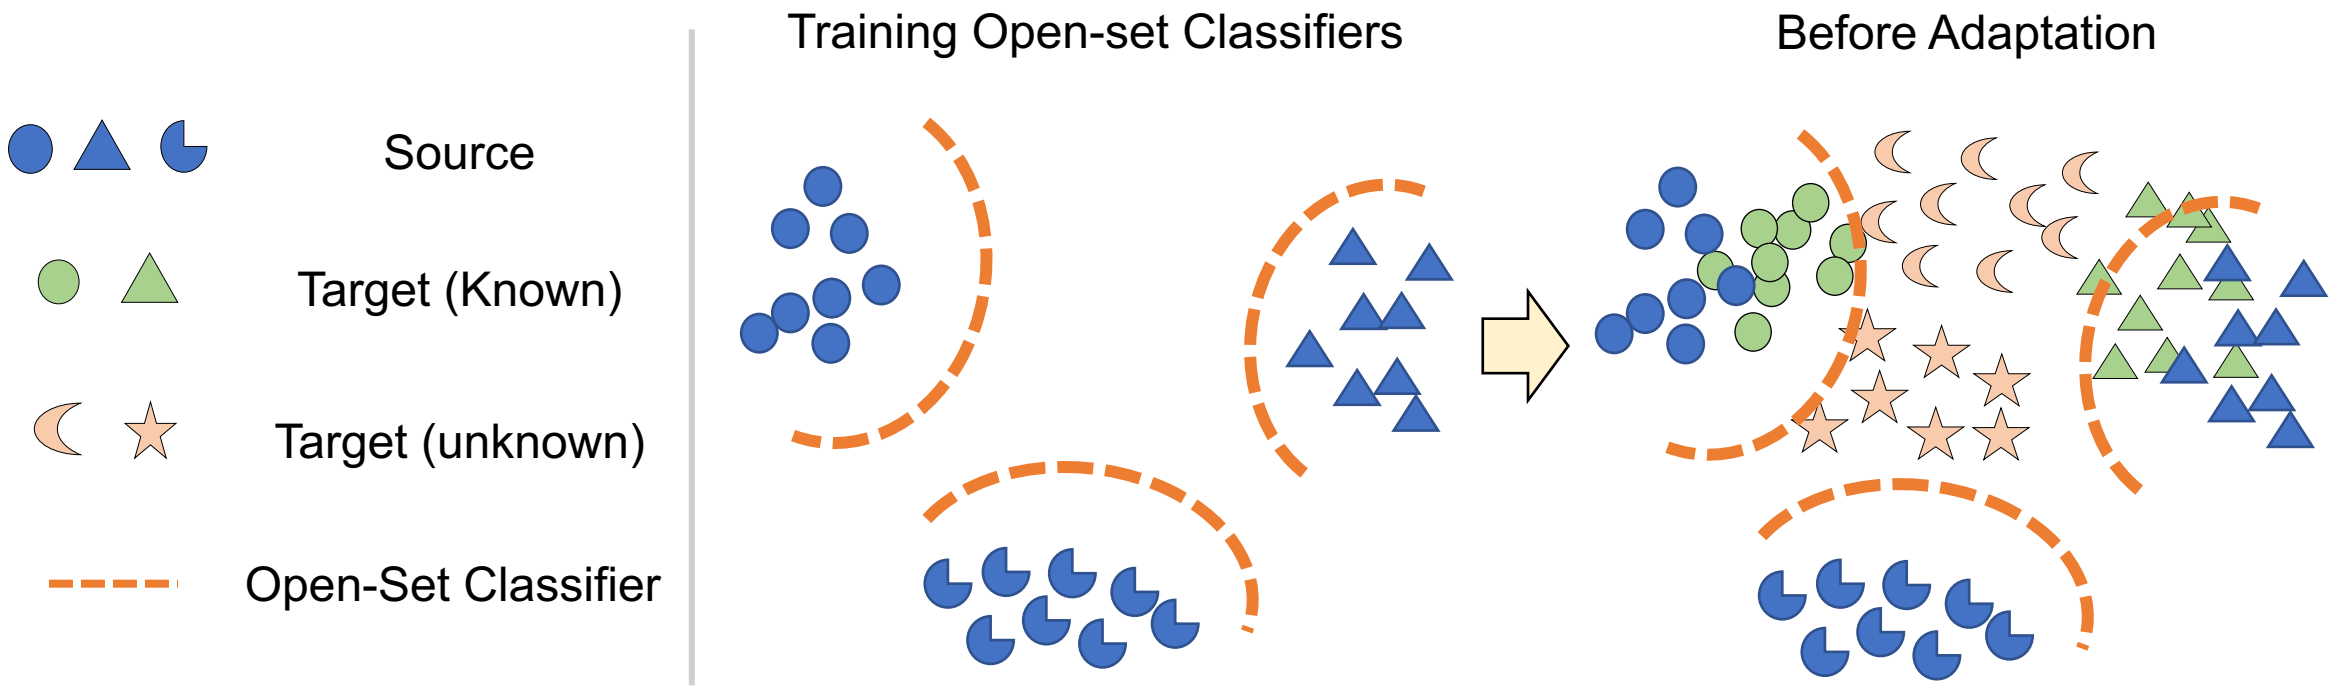

Open-set Entropy Minimization

- Compute the entropy of the open-set classifiers and minimize it.

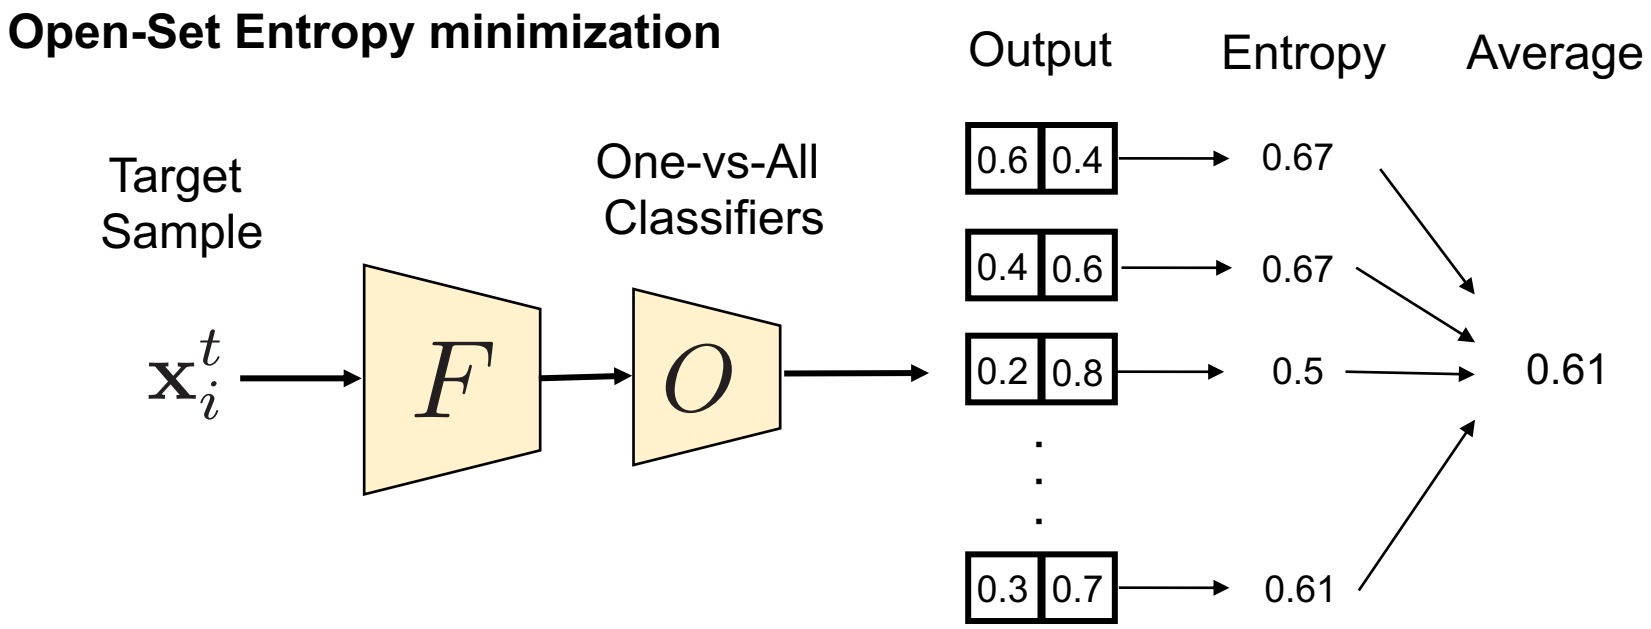

Open-Set Classifier Training

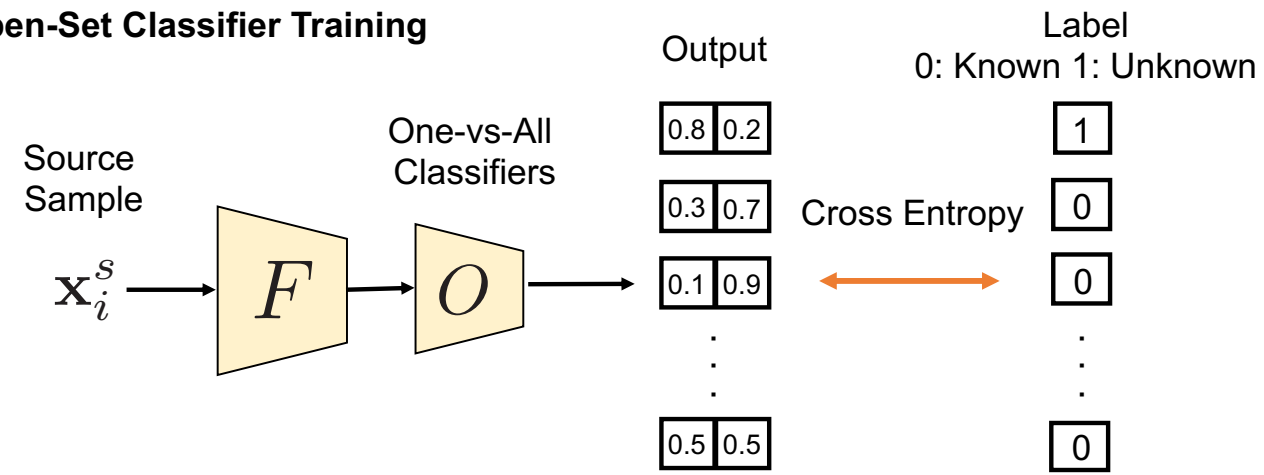

Open-Set Entropy minimization

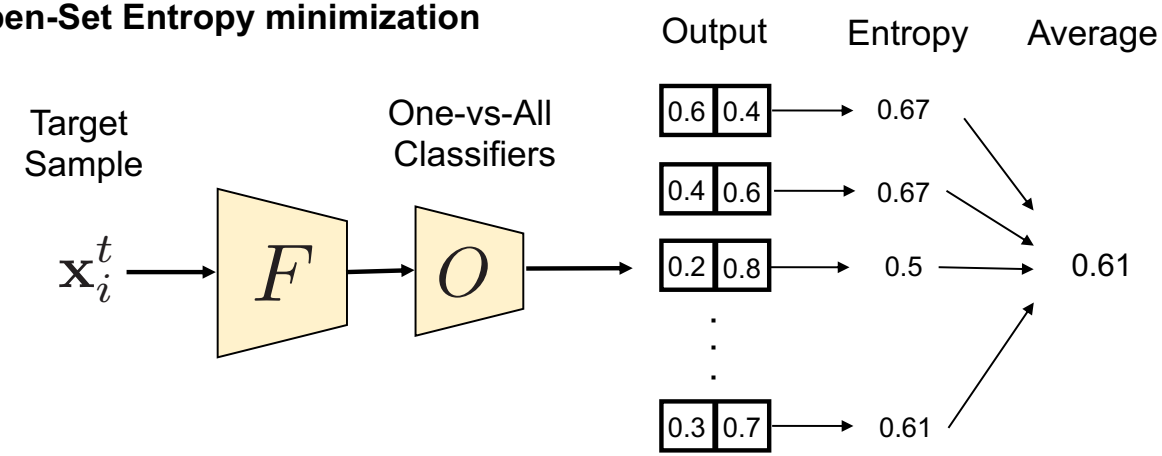

Deployment

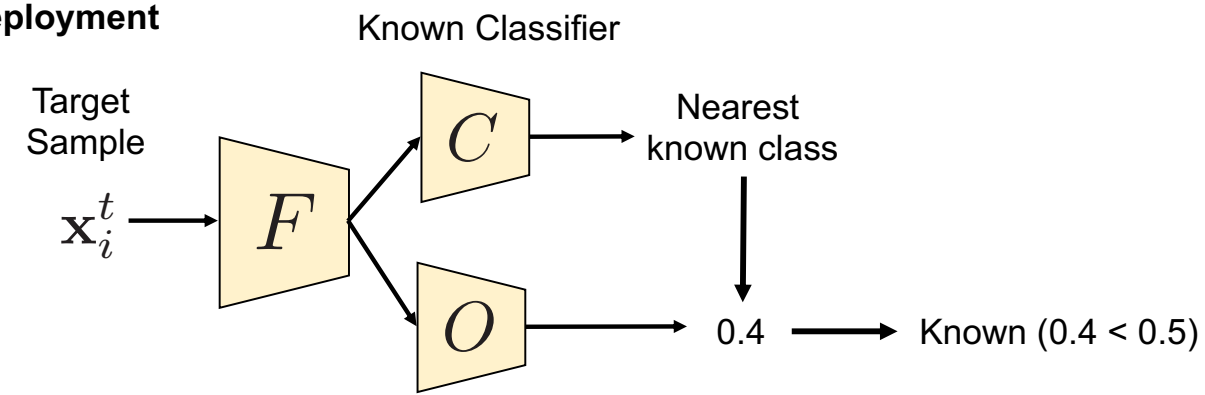

Hyper-parameter

$$\mathcal{L}_{all} = \mathcal{L}_{src} + \lambda \mathcal{L}_{ent}$$

Source loss Entropy loss

Comparison with other methods

Method	Number of hyper-parameters	How is the threshold set?
UAN [You et al., CVPR 2019]	2	Validated
CMU [Fu et al., ECCV 2020]	3	Validated
USFDA [Kundu et al., CVPR 2020]	3	Synthesize unknown samples
ROS [Bucci et al., ECCV 2020]	4	Ratio of unknown target
DANCE [Saito et al., NeurIPS 2020]	3	Set by the number of source classes
OVANet	1	Learned by source

Experiments

- Datasets: Office, OfficeHome, VisDA, DomainNet
- Metric: Accuracy, H-score

Office: 31 classes

VisDA: 12 classes
Synthetic to Real

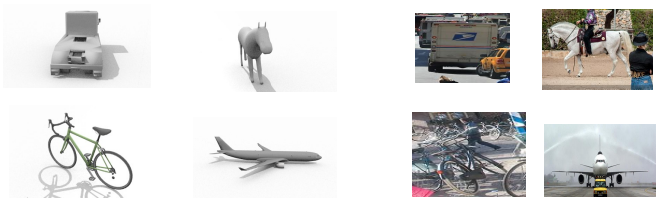

OfficeHome: 65 classes

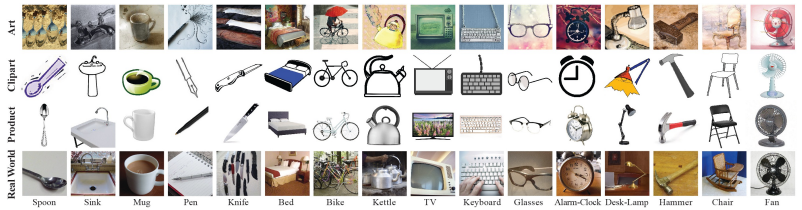

DomainNet: 345 classes

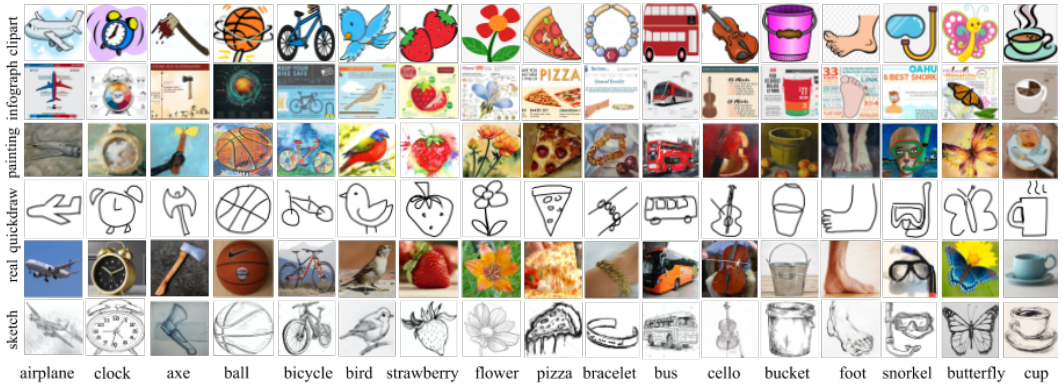

Peng et. al, Moment matching for multi-source domain adaptation

Robustness to the number of unknown target samples

Office, Source: Known 10 classes, Target: Known 10 classes + Unknown classes

Amazon to DSLR

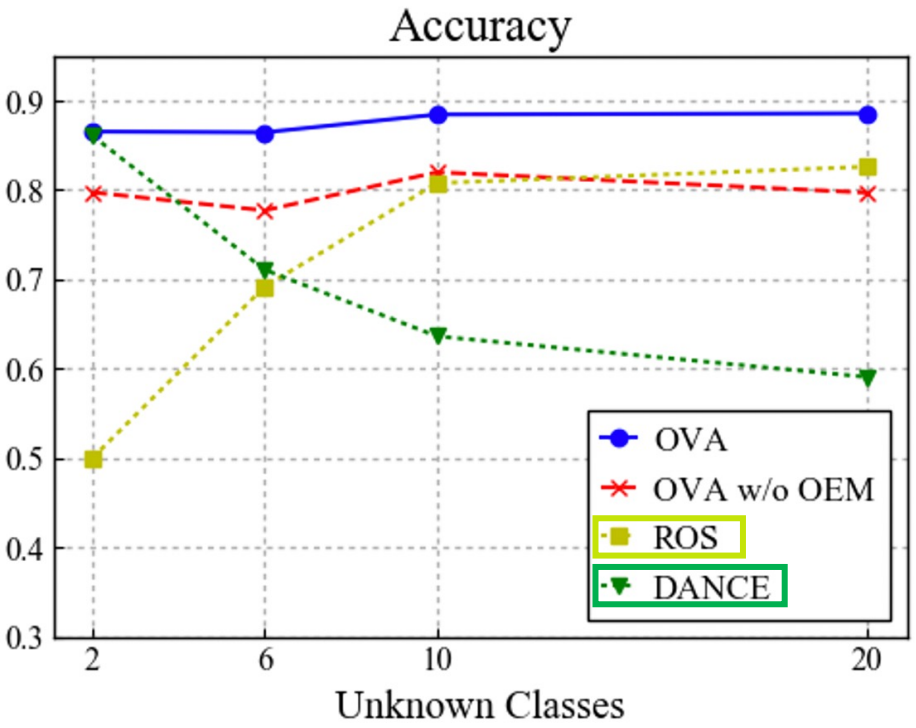

Webcam to Amazon

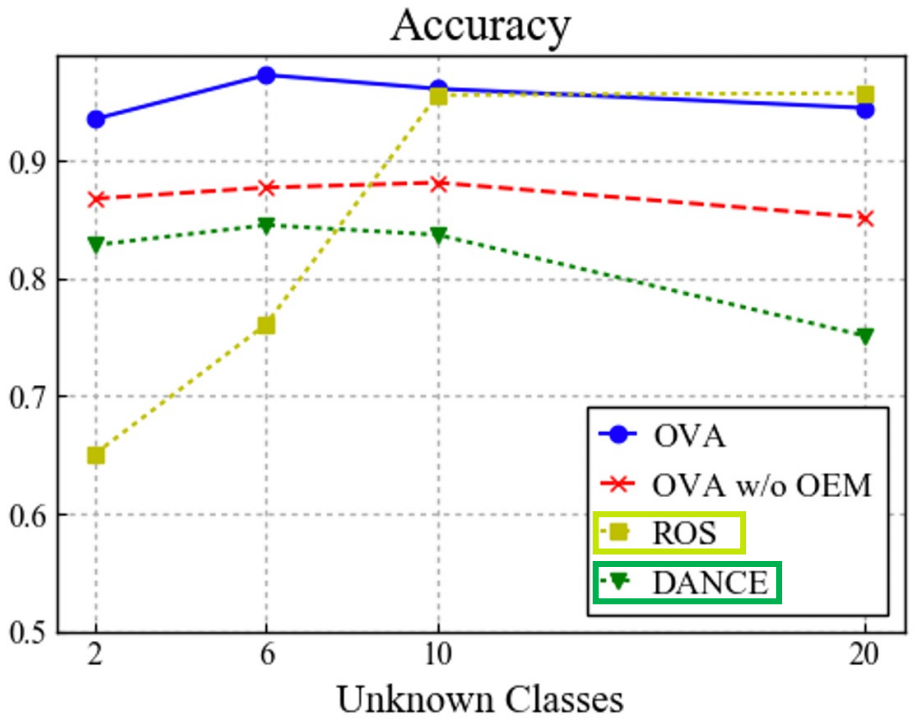

- OVA (ours) shows robustness to different number of unknown samples.
- DANCE selects a threshold in a heuristic way, ROS assumes a specific ratio of unknown samples.

Experiments on large scale datasets

VisDA: 6 shared, 3 source-private, 3 unknown
DomainNet: 150 shared, 50 source-private, 145 unknown

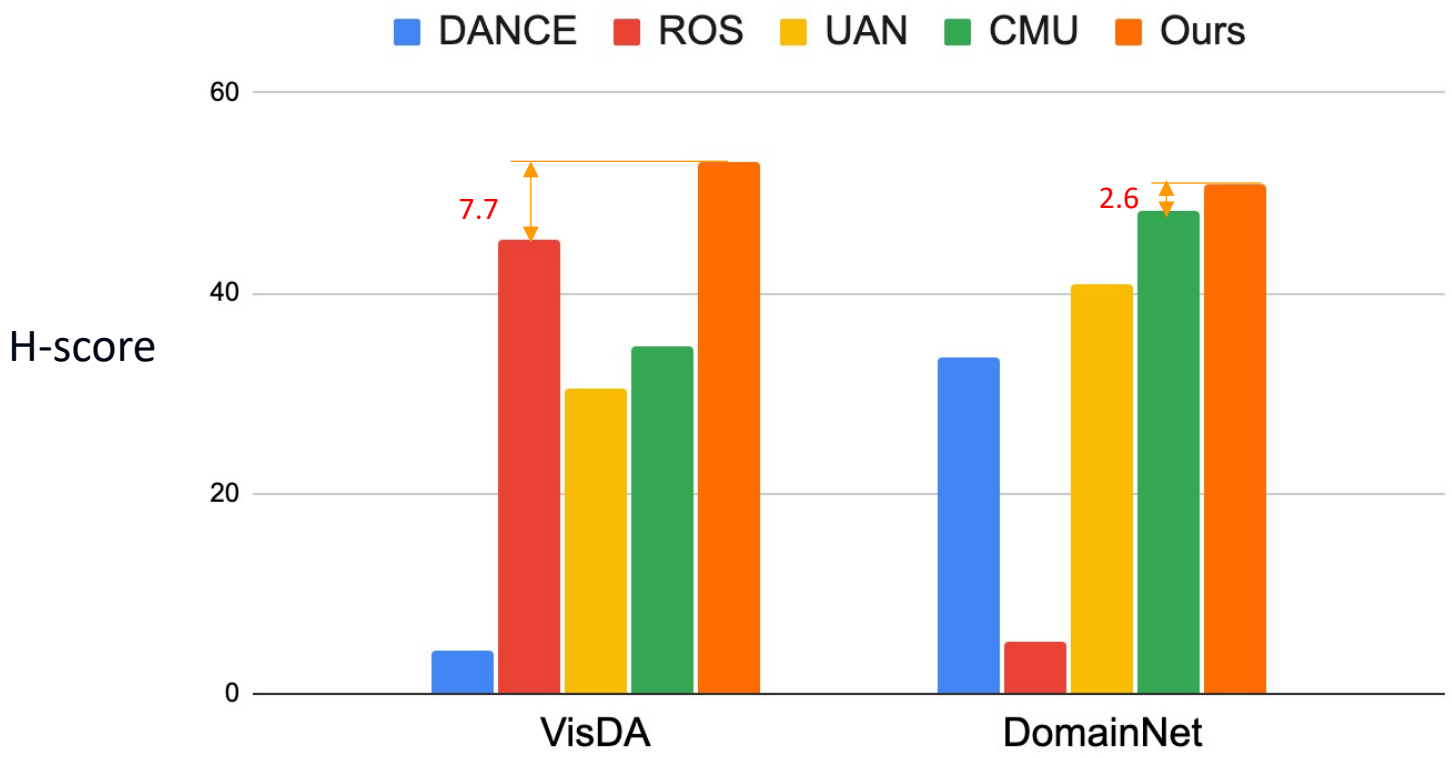

ROS: Reject half of target as unknown.
CMU, UAN: Tune the threshold.

Histogram comparison

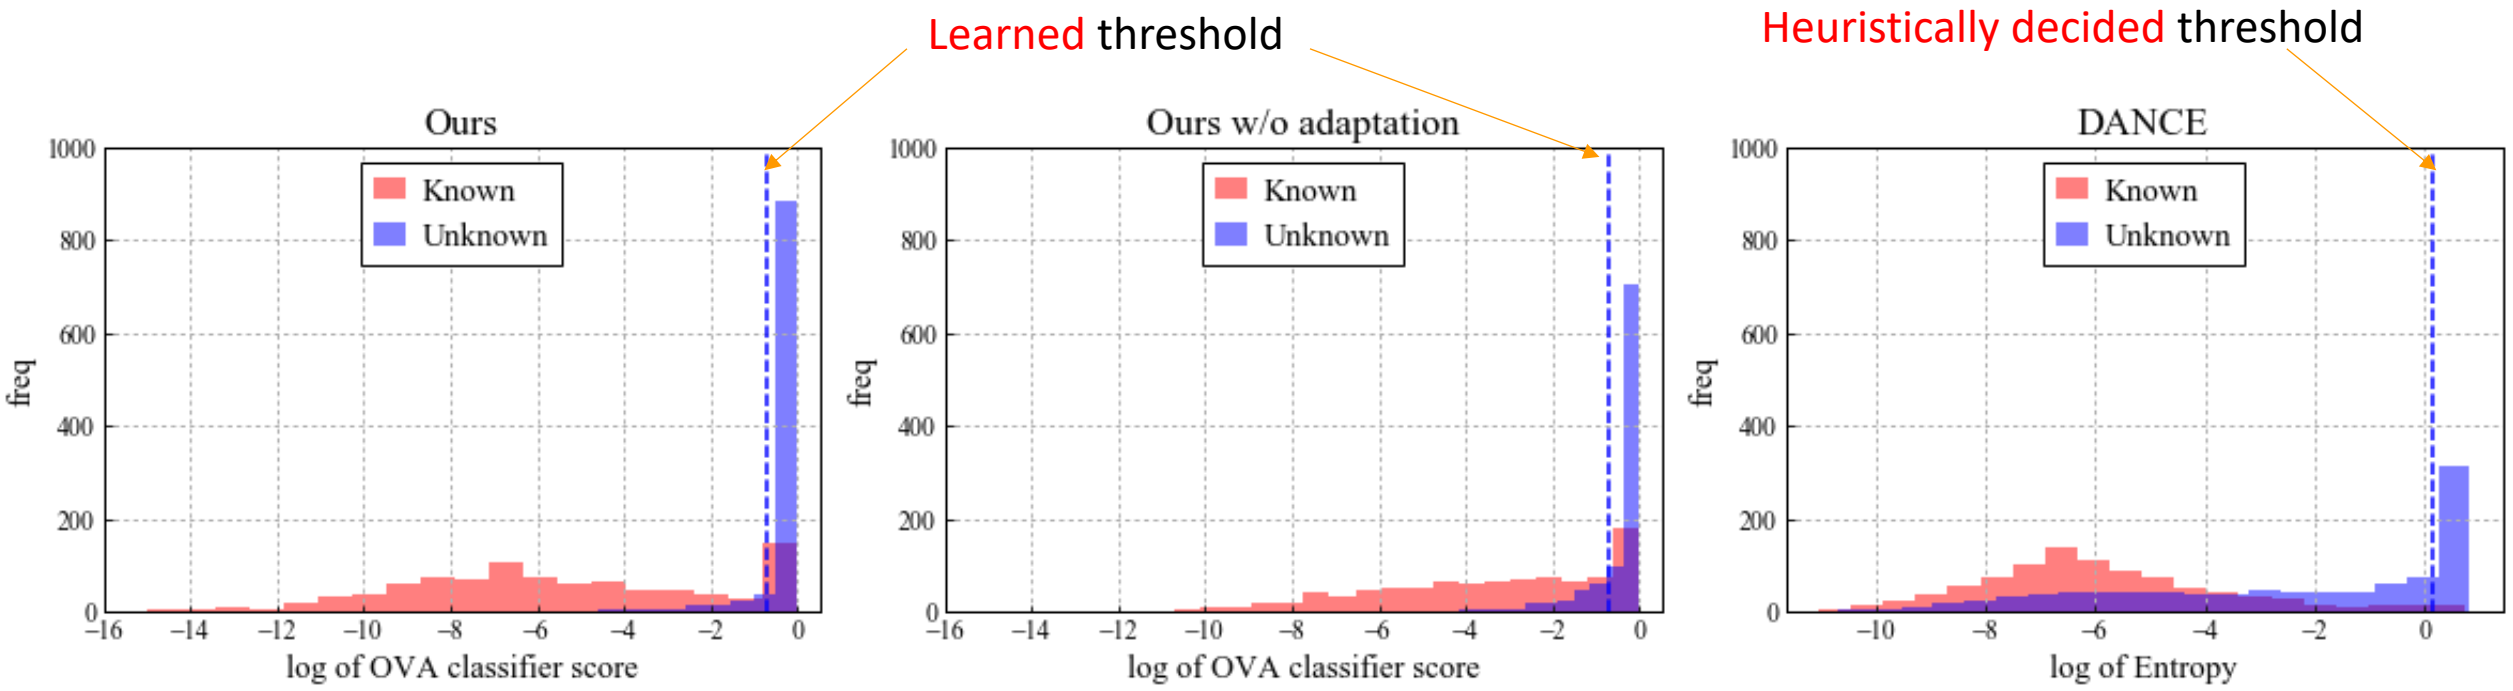

1. Ours separates known and unknown better with the learned threshold.
2. OEM enhances separate between known and unknown instances.

Ablation

Acc close: accuracy for known samples, UNK: accuracy to reject unknown samples

HNCS: hard negative sampling, OEM: open-set entropy minimization

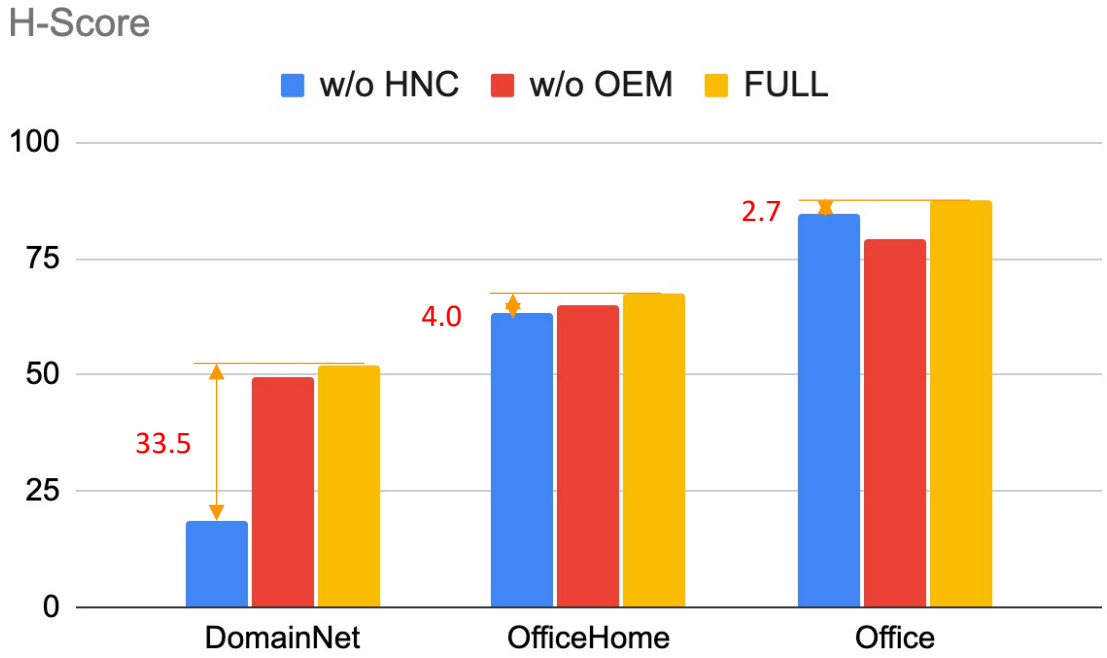

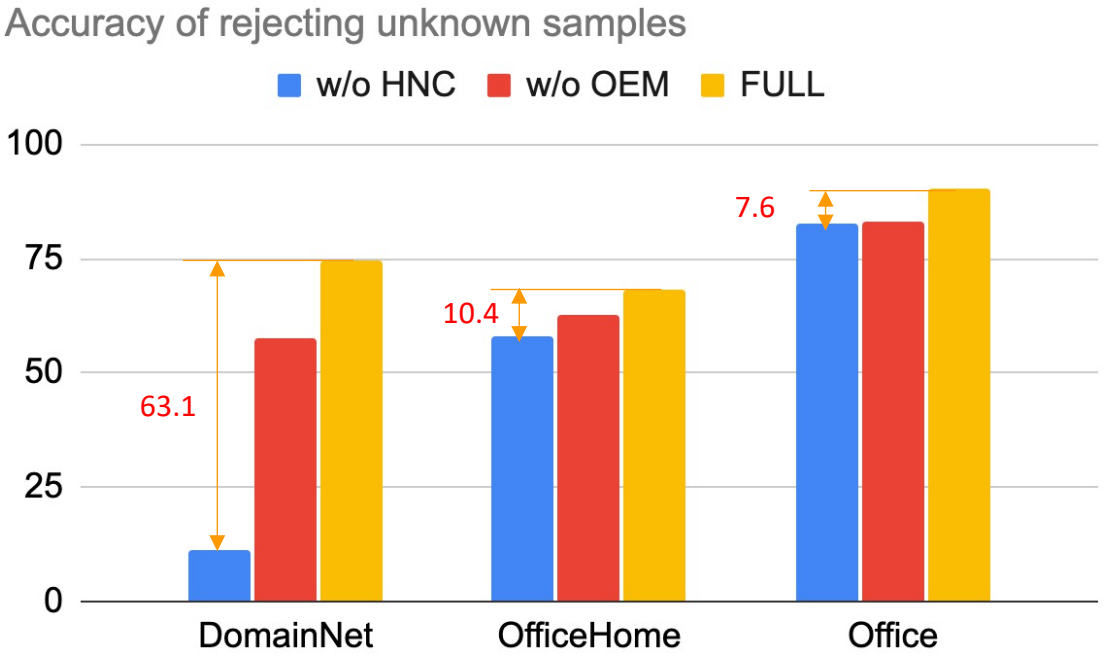

Summary

- One-vs-All Classifier Training for Universal Domain Adaptation
  - Hard-negative sampling
  - Open set entropy minimization
- Effectiveness for Universal DA
  - Better robustness to the number of unknown classes
  - Various datasets
  - Ablation