

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

Kuniaki Saito¹, Kate Saenko^{1,2}

¹Boston University, ²MIT-IBM Watson AI Lab

[keisaito, saenko]@bu.edu

Abstract

Universal Domain Adaptation (UNDA) aims to handle both domain-shift and category-shift between two datasets, where the main challenge is to transfer knowledge while rejecting “unknown” classes which are absent in the labeled source data but present in the unlabeled target data. Existing methods manually set a threshold to reject “unknown” samples based on validation or a pre-defined ratio of “unknown” samples, but this strategy is not practical. In this paper, we propose a method to learn the threshold using source samples and to adapt it to the target domain. Our idea is that a minimum inter-class distance in the source domain should be a good threshold to decide between “known” or “unknown” in the target. To learn the inter- and intra-class distance, we propose to train a one-vs-all classifier for each class using labeled source data. Then, we adapt the open-set classifier to the target domain by minimizing class entropy. The resulting framework is the simplest of all baselines of UNDA and is insensitive to the value of a hyper-parameter, yet outperforms baselines with a large margin. Implementation is available at <https://github.com/VisionLearningGroup/OVANet>.

1. Introduction

Deep neural networks can learn highly discriminative representations for image recognition tasks [6, 30, 14, 26, 12] given a large amount of training data, but do not generalize well to novel domains. Collecting a large amount of annotated data in novel domains incurs a high annotation cost. To tackle this issue, domain adaptation transfers knowledge from a label-rich training domain to a label-scarce novel domain [1]. Traditional unsupervised domain adaptation (UDA) assumes that the source domain and the target domain completely share the sets of categories, *i.e.*, closed-set DA. But, this assumption does not often hold in practice. There are several possible situations: the target domain contains categories absent in the source (unknown categories), *i.e.*, open-set DA (ODA) [4, 29]; the source domain includes categories absent in the target (source pri-

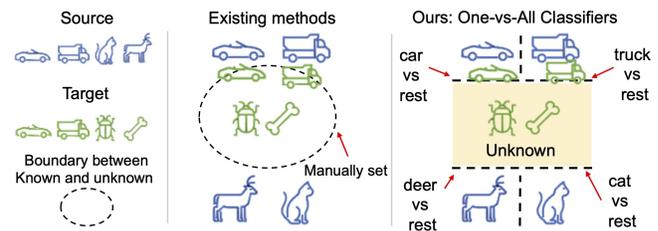


Figure 1: Existing open-set or universal domain adaptation methods handle unknown samples by manually setting a threshold to reject them, either by validation or prior knowledge about the target domain. If set incorrectly, it can mistakenly reject known classes as shown here, *e.g.*, car and truck. Instead, we propose to learn the threshold by training one-vs-all classifiers for each class.

vate categories); *i.e.*, partial DA (PDA) [5]; a mixture of ODA and PDA, called open-partial DA (OPDA). Many approaches have been tailored for a specific setting, but a true difficulty is that we cannot know the category shift in advance.

The task of universal domain adaptation (UNDA) was proposed [37, 28] to account for the uncertainty about the category-shift. The assumption is that the label distributions of labeled and unlabeled data can be different, but we cannot know the difference in advance. Since estimating the label distributions of unlabeled data is very hard in real applications, the setting is very practical. Although we focus on a domain shift problem in this paper, the setting also applies to a semi-supervised learning problems [11].

In this task, our goal is to have a model that can categorize target samples into either one of the correct known labels or the unknown label. The main technical difficulty is that no supervision is available to distinguish unknown samples from known ones; that is, we do not know how many of them are unknown nor the properties of the unknown instances. Obtaining this prior knowledge without manual labeling is hard in practice.

To allow a model to learn the concept of unknown, existing UNDA and ODA methods employ various techniques: rejecting a certain ratio of target samples [3], validating a threshold to decide unknown by using labeled tar-

get samples [7], and synthetically generating unknown instances [15]. Rejecting some ratio of target samples works well if the ratio is accurate. But, estimating the ratio is hard without having labeled target samples. Validation with labeled target samples violates the assumption of UNDA. Synthesized unknown instances define the concept of the unknown for a learned model, but tuning the generation process requires validation with labeled samples since the generated data is not necessarily similar to real unknown data. In summary, as the center of Fig. 1 describes, these existing methods manually define the threshold to reject unknown instances. To achieve a practical solution, we need an approach that does not need the ratio of unknown samples nor any validation to set the threshold.

We cast a question to solve the problem: can we leverage the inter-class distance between source categories to learn the threshold? We assume that the minimum inter-class distance is a good threshold to determine whether a sample comes from the class since it defines a minimum margin from other classes. If the distance between a sample and a class is smaller than the margin, the sample should belong to the class. If the sample does not lie within the margin for any classes, it should be unknown. Fig. 1 illustrates the idea. Car and truck share similar features but belong to different classes. If a model knows the margin between the two classes, it can distinguish unknown classes, *e.g.*, bug and bone, from car and truck.

Given this insight, we explore a simple yet powerful idea: training a one-vs-all (OVA) classifier for every class in the source domain. We train the classifier to categorize inputs other than the corresponding class as negatives. The classifier learns a boundary between positive and negative classes, *i.e.*, employs inter-class distance to learn the boundary. If all of the classifiers regard the input as negative, we assume the input is from unknown classes. Therefore, a model can learn the threshold to reject unknown classes by using source samples. In addition, we propose novel hard-negative classifier sampling, which updates open-set classifiers of a positive and a hard negative class for each source sample, to efficiently learn the minimum inter-class distance for each class. The technique makes a model scalable to a large number of classes. For unlabeled target samples, we propose to apply open-set entropy minimization (OEM), where the entropy of the one-vs-all classifiers is minimized. This allows a model to align unlabeled target samples to either known or unknown classes. Our method is significantly simpler than existing methods since it has only one unique hyper-parameter that controls the trade-off between the classification loss of source samples and OEM, yet shows great robustness to different label distributions of the target domain.

In experiments, we extensively evaluate our method on universal domain adaptation benchmarks and vary the pro-

Method	No. of HP	Threshold
UAN [37]	2	Validated
CMU [7]	3	Validated
USFDA [15]	3	Synthesize unknown samples
ROS [3]	4	Reject 50% of target data
DANCE [28]	3	Decide by No. of classes
OVANet	1	Learned by source

Table 1: **Comparison of open-set and universal DA methods.** HP denotes hyper-parameter. Note that USFDA [15] leverages synthetically generated negatives, which requires a complicated process to generate them.

portion of shared and unknown classes. This simple method outperforms various baselines that explicitly or implicitly employ the ratio of unknown samples. Moreover, the proposed way of detecting unknown classes is effective to set a threshold to reject unknown classes in semi-supervised learning.

2. Related Work

Domain Adaptation. Unsupervised domain adaptation (UDA) [27] aims to learn a good classifier for a target domain given labeled source and unlabeled target data. Let L_s and L_t denote the label space of a source and a target domain respectively. A closed-set domain adaptation ($L_s = L_t$) is a popular task in UDA, and distribution alignment [9, 32, 17] is one of the popular approaches. Open-set DA (presence of target-private classes, $|L_t - L_s| > 0, |L_t \cap L_s| = |L_s|$) [22, 29], and partial DA (presence of source-private classes $|L_s - L_t| > 0, |L_t \cap L_s| = |L_t|$) [5] are proposed to handle the category mismatch problem. Universal DA (UNDA) [37] is proposed to handle the mixture of these settings. Saito *et al.* [28] emphasize the importance of measuring the robustness of a model to various category-shifts since we cannot know the detail of the shifts in advance. Prior works [37, 7, 28] compute a confidence score for known classes, and samples with a score lower than a threshold are regarded as unknown. Fu *et al.* [7] seem to validate the threshold using labeled data, which is not a realistic solution. Bucci *et al.* [3] set the mean of the confidence score as the threshold, which implicitly rejects about half of the target data as unknown. Saito *et al.* [28] set a threshold decided by the number of classes in the source, which does not always work well. We summarize how our approach, OVANet, is different from existing methods in Table 1. Our approach trains a model to learn the threshold by using source samples and attempts to adapt the threshold to a target domain. Our model is trained in an end-to-end manner, requires only one hyper-parameter, and is not sensitive to its value.

Open-set recognition. Open-set recognition [2] handles both known and unknown samples in the test phase given known samples during training. Many methods fo-

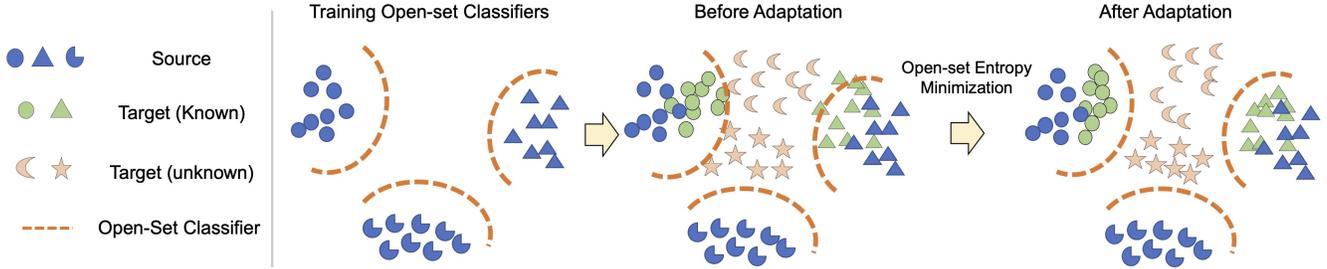


Figure 2: **Conceptual overview of OVANet.** Open-set classifiers are trained on labeled source samples (leftmost), and we attempt to increase the confidence in the predictions by open-set entropy minimization on unlabeled target samples (middle and rightmost). We show the training procedure in a step-by-step manner for better understanding but employ end-to-end training in practice.

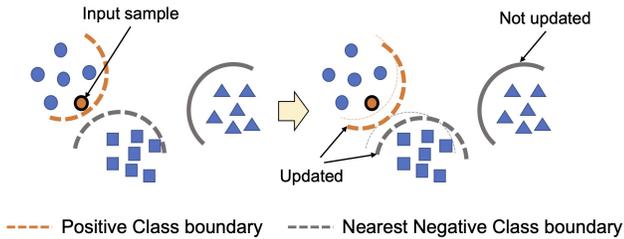


Figure 3: **Overview of the open-set classifier training.** For each sample, a positive and a nearest negative class boundary are updated. This is to let the classifier learn the distance between the positive and nearest negative class. We hypothesize that the distance should be a good threshold to reject “unknown” examples.

focus on how to build a better measurement of anomaly [2] or to let a model learn features effective to distinguish known or unknown samples [36, 25, 20]. Recent works show that contrastive learning is effective to learn representations suitable to detect out of distribution samples [31, 35]. These methods have a common issue with existing UNDA and ODA methods. The threshold to determine unknown is validated or pre-determined [10]. These approaches provide a metric or method of training effective to calibrate the uncertainty score, but it is still necessary to determine the threshold by validation. By contrast, we introduce a technique that learns the threshold using labeled samples without a need for validation.

A One-vs-All (OVA) classifier is employed to apply a binary classifier to multi-class classification [8], aggregating the output of the binary classifier. Padhy *et al.* [21] propose to handle out-of-distribution detection by training neural OVA classifiers. A key difference from this work is that we train open-set classifiers by letting them focus on hard-negative samples. Open-set classifiers need to distinguish positive and nearest negative samples to effectively identify unknown samples. Our proposed hard-negative sampling is effective in building a threshold as shown in Sec. 3.1 and Table 4.

3. OVANet

Fig. 2 introduces the conceptual overview of OVANet. Our open-set classifier employs the learned distance between categories to identify unknown samples (Sec. 3.1). The learned classifiers are adapted to the target domain by open-set entropy minimization (Sec. 3.2).

Notation. We are given a labeled source domain $\mathcal{D}_s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$ with “known” categories L_s and an unlabeled target domain $\mathcal{D}_t = \{(\mathbf{x}_i^t)\}_{i=1}^{N_t}$ which contains “known” categories and “unknown” categories, where L_s and L_t denote the label spaces of the source and target respectively. Our goal is to label the target samples with either one of the L_s labels or the “unknown” label. We train the model on $\mathcal{D}_s \cup \mathcal{D}_t$ and evaluate on \mathcal{D}_t . To handle the open-set classification, we introduce two classifiers: (1) open-set classifiers, O , to detect unknown samples and (2) a closed-set classifier, C , to classify samples into L_s labels. C is trained to classify source samples with a standard classification loss while O is trained with *hard negative classifier sampling* explained below. In the test phase, C is used to identify the nearest known class while O is used to determine whether the sample is known or unknown.

3.1. Hard Negative Classifier Sampling (HNCS)

Our idea is to train a classifier to learn a boundary between in-liers and outliers for each class. Then, we train a linear classifier for each class. For each classifier, the class is trained to be positive while other classes are negative. The key is how to pick the negative samples, *i.e.*, we propose to pick samples different from the class but similar to it (hard negative class). The overview is illustrated in Fig. 3.

Our open-set classifiers consist of $|L_s|$ sub-classifiers, *i.e.*, $O^j, j \in \{1, \dots, |L_s|\}$. Each classifier is trained to distinguish if the sample is an in-lier for the corresponding class. The sub-classifier outputs a 2-dimensional vector, where each dimension shows the probability of a sample being an in-lier and outlier respectively.

Padhy *et al.* [21] propose to train all $(|L_s| - 1)$ negative classifiers given a training sample, but we observe that

learned classifiers are not useful when the number of classes is large, which is consistent with the observation of [21]. The reason is that the open-set classifiers can classify many negatives too easily. Take an example of three classes: a cat, a dog, and a turtle. We assume that the features of the turtle are very different from those of the cat and the dog. To learn an effective boundary for the cat, a model should focus on the dog rather than the turtle since the cat should be closer to the dog in a feature space. But, if the turtle is sampled very frequently, the learned boundary can be in the middle of the cat and the turtle, accepting too many unknown samples as known. Therefore, each sub-classifier needs hard negative samples to learn a boundary between in-liers and outliers. Considering this insight, we propose to train two one-vs-all classifiers given a sample: a classifier of the positive and that of the nearest negative class (Fig. 3). By picking the nearest negative class for each sample, we can let the corresponding classifier learn an effective boundary to identify unknown instances.

We leverage a linear classifier for each sub-classifier. The open-set classifier is employed on top of an extracted feature, namely, $z^k = \mathbf{w}^k G_\theta(\mathbf{x}) \in \mathbb{R}^2$, where G_θ and \mathbf{w}^k denote a feature extractor and a weight of an open-set classifier for class k respectively. Each dimension of $z^k \in \mathbb{R}^2$ denotes the score for known and unknown respectively. Let $p(\hat{y}^k|\mathbf{x})$ denote the output probability that the instance \mathbf{x} is an in-lier for the class k : $p(\hat{y}^k|\mathbf{x}) = \sigma(z^k)_0$, where σ denotes softmax activation function.

We denote $L_{ova}(\mathbf{x}^s, y^s)$ as the open-set classification loss for a sample (\mathbf{x}^s, y^s) :

$$\mathcal{L}_{ova}(\mathbf{x}^s, y^s) = -\log(p(\hat{y}^{y^s}|\mathbf{x}^s)) - \min_{j \neq y^s} \log(1 - p(\hat{y}^j|\mathbf{x}^s)). \quad (1)$$

This computes the loss on a positive class and the hardest negative class, as we explain above. We call the technique hard negative classifier sampling (HNCS). The process of computing the loss is illustrated at the top of Fig. 4.

3.2. Open-set Entropy Minimization (OEM)

Given the open-set classifiers trained on the source domain, we propose to enhance the low-density separation for the unlabeled target domain. Since the target samples have different characteristics from the source, they can be classified incorrectly with respect to both closed-set and open-set categorization. To handle the issue, we propose a novel entropy minimization method that adapts the open-set classifiers to the target domain.

Our idea is to increase the confidence in the prediction with regard to open-set classification, i.e., known or unknown. Specifically, we apply entropy minimization training for all open-set classifiers for each $\mathbf{x}^t \in \mathcal{D}_t$. We compute entropy of all the classifiers and take the average and train a model to minimize the entropy, as illustrated in the

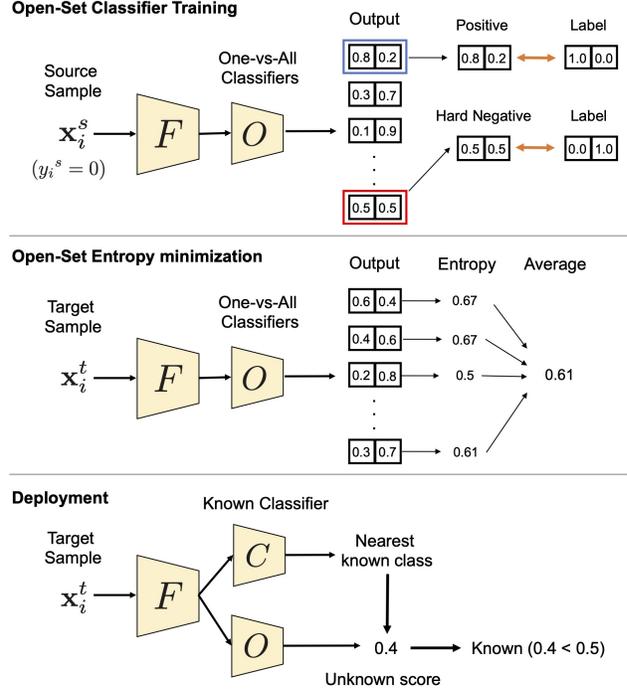


Figure 4: **Overview of training and testing.** O denotes one-vs-all open-set classifiers whereas C denotes a closed-set classifier. F is a feature extractor. **Top:** We propose hard-negative classifier sampling to train one-vs-all classifiers using source samples (Sec. 3.1). **Middle:** We apply entropy minimization with open-set classifiers for unlabeled target samples (Sec. 3.2). **Bottom:** In the test phase, a nearest known class is identified by a closed-set classifier, and the corresponding open-set classifier’s score is leveraged to decide known or unknown (Sec. 3.4).

middle of Fig. 4.

$$\mathcal{L}_{ent}(\mathbf{x}^t) = - \sum_{j=1}^{|L_s|} p(\hat{y}^j|\mathbf{x}^t) \log(p(\hat{y}^j|\mathbf{x}^t)) + (1 - p(\hat{y}^j|\mathbf{x}^t)) \log(1 - p(\hat{y}^j|\mathbf{x}^t))$$

By this entropy minimization, known target samples will be aligned to source samples whereas unknown ones can be kept as unknown. One clear difference from the existing entropy minimization [11] is that we are able to keep unknown instances as unknown since the entropy minimization is performed by open-set classifiers, not by a closed-set classifier. Entropy minimization by a closed-set classifier necessarily aligns the unlabeled samples to known classes because there is no concept of an unknown class. Since our open-set classifiers have the concept of unknown, the model can increase the confidence of it.

3.3. Learning

We combine the open-set classifier and closed-set classifier to learn both open-set and closed-set categorization.

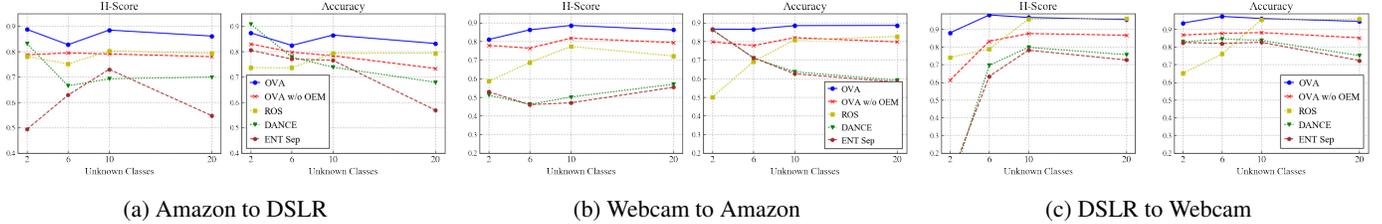


Figure 5: **H-score and accuracy in open-set domain adaptation.** Blue: OVANet (Ours). We vary the number of unknown classes using Office ($|L_s \cap L_t| = 10, |L_s - L_t| = 0$). The left and right side show H-score and accuracy respectively. OVANet shows stable performance across different openness while baselines can much degrade performance in some settings.

For the closed-set classifier, we simply train a linear classifier on top of the feature extractor using cross-entropy loss, which we denote $\mathcal{L}_{cls}(x, y)$. Then, the overall training loss can be computed as follows:

$$\mathcal{L}_{all} = \mathbb{E}_{(\mathbf{x}_i^s, y_i^s) \sim \mathcal{D}_s} \mathcal{L}_{src}(\mathbf{x}_i^s, y_i^s) + \lambda \mathbb{E}_{\mathbf{x}_i^t \sim \mathcal{D}_t} \mathcal{L}_{ent}(\mathbf{x}_i^t), \quad (2)$$

$$\mathcal{L}_{src}(\mathbf{x}_i^s, y_i^s) = \mathcal{L}_{cls}(\mathbf{x}_i^s, y_i^s) + \mathcal{L}_{ova}(\mathbf{x}_i^s, y_i^s). \quad (3)$$

The parameters of F , O , and C are optimized to minimize the loss. Note that OVANet has only one hyper-parameter, λ . This method is much simpler than existing ODA and UNDA methods [28, 7, 3] all of which require setting the threshold manually and/or multiple training phases.

3.4. Inference

In the test phase, we utilize both the trained closed-set and open-set classifier. We first get the closest known class by using the closed-set classifier and take the corresponding score of the open-set classifier. The process is illustrated at the bottom of Fig. 4.

4. Experiments

We evaluate our method in UNDA settings along with ablation studies. To evaluate the robustness to the change of the number of unknown target samples, we vary the number and compare it with other baselines.

4.1. Setup

Datasets. We utilize popular datasets in DA: Office [27], OfficeHome [34], VisDA [24], and DomainNet [23]. Unless otherwise noted, we follow existing protocols [7, 37, 29] to split the datasets into source-private ($|L_s - L_t|$), target-private ($|L_t - L_s|$) and shared categories ($|L_s \cap L_t|$). For fairness, Saito *et al.* [28] propose to evaluate universal DA methods on the various number of unknown and known classes, which can reveal methods tailored for a specific setting. We follow their policy and provide experimental results varying the number of unknown and known classes. Since many existing methods are optimized to handle a specific benchmark, we aim to fairly evaluate methods with respect to the sensitivity to diverse settings. To briefly

Method	Office (10 / 10 / 11)						Avg
	A2D	A2W	D2A	D2W	W2D	W2A	
UAN [37]	59.7	58.6	60.1	70.6	71.4	60.3	63.5
CMU [7]	68.1	67.3	71.4	79.3	80.4	72.2	73.1
DANCE [28]	78.6	71.5	79.9	91.4	87.9	72.2	80.3
DCC [16]	88.5	78.5	70.2	79.3	88.6	75.9	80.2
ROS [3]	71.4	71.3	81.0	94.6	95.3	79.2	82.1
USFDA [15]	85.5	79.8	83.2	90.6	88.7	81.2	84.8
OVANet w/o OEM	69.6	63.1	79.9	85.9	88.7	80.6	77.9
OVANet	85.8	79.4	80.1	95.4	94.3	84.0	86.5

Table 2: **Open-partial domain adaptation using Office (H-score).**

describe the method for splitting categories in each experiment, each table shows $|L_s \cap L_t|/|L_s - L_t|/|L_t - L_s|$, *i.e.*, (shared, source private, and target private classes). See our supplementary material for more details.

Evaluation Metric. Considering the trade-off between the accuracy of known and unknown classes is important in evaluating universal or open-set DA methods. To this end, we evaluate methods using H-score [3]. H-score is the harmonic mean of the accuracy on common classes (acc_c) and accuracy on the “unknown” classes acc_t as:

$$H_{score} = \frac{2acc_c \cdot acc_t}{acc_c + acc_t}. \quad (4)$$

The evaluation metric is high only when both the acc_c and acc_t are high. So, H-score can measure both accuracies of UNDA methods well. Unless otherwise noted, we show H-score in tables and graphs. The drawback of this metric is that the importance of recognizing known and unknown classes is always equal. If the number of unknown instances is small, the metric puts too much weight on the unknown class. Thus, when the number is small, we also report the instance-wise accuracy over all samples.

Implementation. Following previous works [28, 37], we employ ResNet50 [13] pre-trained on ImageNet [6] as our backbone network. Evaluation using VGGNet [30] is also performed in analysis. We replace the last linear classification layer with a new linear layer. We follow [28] and train our models with inverse learning rate decay scheduling. Note that we set $\lambda = 0.1$ across all settings. The

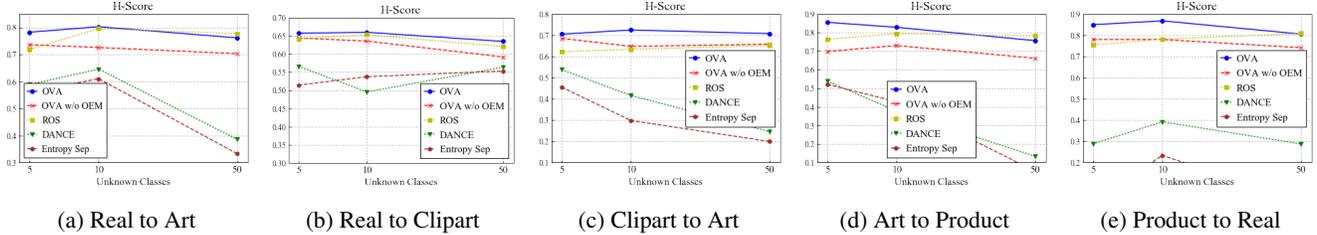


Figure 6: **H-score of open-partial domain adaptation.** Blue: OVANet (Ours). We vary the number of unknown classes in OfficeHome ($|L_s \cap L_t| = 10, |L_s - L_t| = 5$). OVANet shows stable performance across different openness while baselines can much degrade performance in some settings.

Method	DomainNet (150 / 50 / 145)						Avg	VisDA (6 / 3 / 3)	OfficeHome (15 / 5 / 50)
	P2R	R2P	P2S	S2P	R2S	S2R			
DANCE [28]	21.0	47.3	37.0	27.7	46.7	21.0	33.5	4.4	49.2
UAN [37]	41.9	43.6	39.1	38.9	38.7	43.7	41.0	30.5	56.6
CMU [7]	50.8	52.2	45.1	44.8	45.6	51.0	48.3	34.6	61.6
DCC [16]	56.9	50.3	43.7	44.9	43.3	56.2	49.2	43.0	70.2
OVANet	56.0	51.7	47.1	47.4	44.9	57.2	50.7	53.1	71.8

Table 3: **H-score of open-partial DA using DomainNet, VisDA and OfficeHome.** Note that CMU [7] and DCC [16] use different hyper-parameters for different datasets while OVANet uses the same hyper-parameter throughout all settings.

value is determined by the result of open-set DA using Office (Amazon to DSLR) following DANCE [28].

Baselines. We aim to compare methods of universal domain adaptation, which are able to reject unknown samples, such as DANCE [28], UAN [37], ROS [3] and CMU [7]. To see the difference from using closed-set classifier’s entropy as a threshold, we employ Entropy Separation (Ent Sep) [28] in several experiments. Note that these baselines are unfair in that they utilize validated thresholds or heuristically decided thresholds. We decide not to include the results of standard domain alignment baselines such as DANN [9], CDAN [18] since existing works have already shown that these methods significantly worsen the performance in rejecting unknown samples. Since CMU [7] does not publish complete code to reproduce the results, we rely on their reported results to compare with the method.

Overview of Results. In summary, our method is superior or comparable to baseline methods across all diverse settings without optimizing the hyper-parameter for each setting. The fact is verified by diverse settings using 4 benchmark datasets.

Office. Fig. 5 shows the results of varying the number of unknown classes in the ODA setting on Office. The left and the right of each graph show H-score and instance-wise accuracy respectively. The x-axis denotes the number of unknown classes. The number varies from 2 to 20 while that of known classes is fixed to 10. The performance of OVANet is always better than or comparable to baselines. Even without OEM, our proposed model consistently performs well. Comparison between OVANet and OVANet w/o OEM demonstrates the effectiveness of OEM. Since ROS [3] sets a threshold to reject about half of the target

samples as unknown, it performs poorly when the unknown samples are rare. DANCE [28] is also sensitive to the number of unknown classes. Table 2 shows the result of the OPDA setting on Office, where OVANet outperforms baselines on average.

OfficeHome, VisDA and DomainNet. Results of varying the openness in the OPDA on OfficeHome are described in Fig. 6. We pick 5 adaptation scenarios to cover various domains and vary the number of unknown classes. The trend is similar to experiments on Office. OVANet consistently performs better than baselines. Results of the OPDA on OfficeHome, VisDA and DomainNet are summarized in Table 3, where we follow CMU [7] to split the categories. For OfficeHome, the mean of 12 adaptation scenarios is presented. In VisDA and DomainNet, the number of samples and/or that of classes are very different from those of Office and OfficeHome. OVANet outperforms existing methods with a large margin, more than 10 points in VisDA and OfficeHome. Note that CMU [7] selects the optimal threshold hyper-parameters for each dataset, while OVANet achieves the best H-score without tuning a hyper-parameter. From these results, we observe that OVANet works well across diverse settings.

4.2. Analysis in Domain Adaptation

How effective are OEM and HNCS? Table 4 shows ablation study in ODA on Office (Webcam to Amazon), ODA on OfficeHome (Real to Art), and OPDA on DomainNet (Real to Painting). *Acc close* measures accuracy at recognizing known samples without rejection, i.e, this metric evaluates the ability of closed-set recognition. *UNK* is the accuracy of rejecting unknown samples. *AUROC* measures

Ablation		Office (W2A) (10 / 0 / 11)				OfficeHome (R2A) (20 / 0 / 45)				DomainNet (R2P) (150 / 50 / 145)			
HNCS	OEM	H-score	Acc close	UNK	AUROC	H-score	Acc close	UNK	AUROC	H-score	Acc close	UNK	AUROC
	✓	84.7	94.4	82.7	92.4	63.5	80.7	58.0	76.3	18.4	52.2	11.3	66.1
✓		79.1	93.3	83.1	89.4	64.9	79.9	62.5	75.6	49.7	50.7	57.8	67.2
✓	✓	87.4	94.7	90.3	94.7	67.5	81.3	68.4	77.6	51.9	53.5	74.4	67.6

Table 4: **Ablation study.** We ablate open-set entropy minimization (OEM) and/or hard negative classifier sampling (HNCS). Note that both techniques are necessary to achieve good performance (FULL vs each ablation).

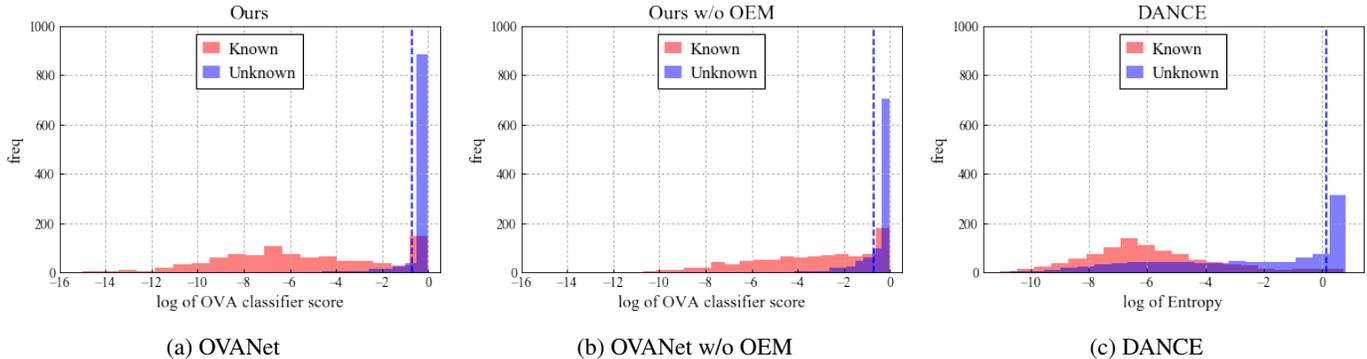


Figure 7: **Histogram of anomaly score in open-set DA** (Webcam to Amazon, $|L_s \cap L_t| = 10$, $|L_s - L_t| = 0$, $|L_t - L_s| = 11$). **Red:** Known samples. **Blue:** Unknown samples. The output of open-set classifiers is used as anomaly score in OVA Net, and the entropy of a closed classifier is used for DANCE. Note that the threshold between known and unknown is a blue dotted line in each graph. OVA Net separates known and unknown samples well whereas DANCE fails.

how well known and unknown samples are separated given the open-set classifier’s output. We have two observations: applying open-set entropy minimization (OEM) is effective to improve both closed-set accuracy and rejecting unknown examples (w/o OEM vs FULL), and hard-negative classifier selection is an appropriate way of training the classifier (w/o HNCS vs FULL). The effectiveness of hard-negative sampling is more evident in DomainNet because the dataset has much more known classes. When the number of known classes is large, there are many useless negative classes to train our open-set classifier. Then, without the sampling, the model cannot learn an effective decision boundary between known and unknown, resulting in a significant degradation in the unknown sample recognition.

Unknown and known samples are well separated by OVA classifiers. Fig. 7 shows histograms of anomaly score, where the x-axis is the anomaly score, and the y-axis shows the frequency of the corresponding range of anomaly score. The blue dotted line denotes the threshold between known and unknown. (b) indicates that the learned threshold works well to separate known and unknown samples. Applying OEM further enhances the separation ((a) vs (b)). (c) DANCE utilizes a manually set threshold and fails in rejecting many unknown samples.

Feature Visualization. Fig. 8 visualizes learned features with corresponding ground-truth labels (a) and predicted labels (b). Although not all unknown samples are clearly separated from known ones, most are correctly classified.

OEM is not sensitive to a hyper-parameter. Fig. 9a shows the sensitivity to the hyper-parameter λ in ODA on

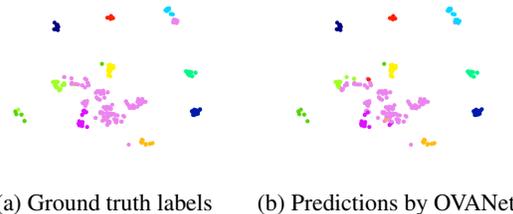


Figure 8: **Feature visualization** with t-SNE [19]. Comparison between (a) true labels and (b) decision by OVA Net (ODA, Amazon to DSLR). Different colors indicate different classes. Plots with pink are unknown samples, others are known samples.

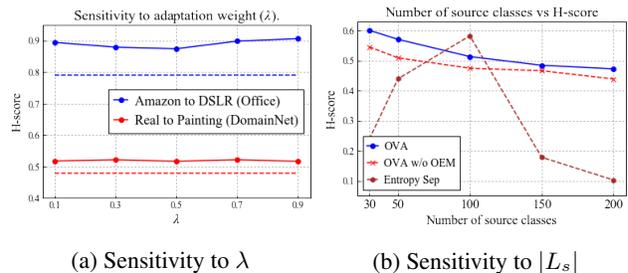


Figure 9: (a) Sensitivity to a hyper-parameter λ in Eq. 3. Dotted lines are results w/o OEM. (b) Sensitivity to the number of source classes, where $|L_s \cap L_t| = 30$, $|L_t - L_s| = 45$ in DomainNet.

Office and OPDA on DomainNet. λ is the only hyper-parameter specific to OVA Net. OVA Net shows stable performance across different values of λ .

Varying known classes. The performance across different numbers of source-private classes ($|L_s|$) is visualized in

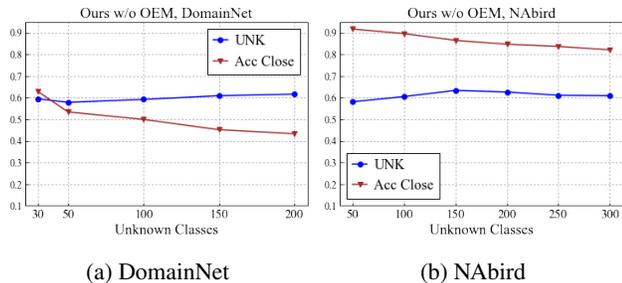


Figure 10: **Increasing number of known classes** with OVANet w/o OEM. **Blue**: Accuracy for known samples. **Brown**: Accuracy for unknown samples. As a model sees a larger number of known classes, the model rejects more unknown samples.

Fig. 9b, As there are more out-lier source classes, correctly classifying target classes can become harder. While Entropy Separation [28] is very sensitive to the number, OVANet shows stable performance.

More known classes lead to a better boundary to reject unknown. Fig. 10 focuses on the results of OVANet w/o OEM in increasing the number of known classes while fixing the number of unknown classes. See appendix for the detail of used datasets. Although the number gets larger, there is more probability of unknown samples being categorized into known classes. However, the accuracy at rejecting unknown classes improves or does not change with the increase of the number. This indicates that the model is more likely to correctly reject unknown samples given more known classes. This observation is consistent with the design of one-vs-all classifiers. It is necessary for the classifiers to see hard negative samples to build a good boundary. When more classes are available, there will be more hard negatives.

OVANet is effective for different networks. Table 5 shows results of ODA on Office using VGGNet as a backbone. Note that ROS [3] changes hyper-parameters to train their model while we do not change the hyper-parameter of OVANet. The idea of employing one-vs-all classifiers is effective for different networks.

OVANet needs both open-set and closed-set classifiers. OVANet employs both open-set and closed-set classifiers, but an alternative way of training is to employ only an open-set classifier. However, we see a large degrade in performance if we do not use a closed-set classifier in training. This is because the open-set classifier is not trained to distinguish between known classes. The training objective considers only distinguish one known class from other classes, which is not enough to give discriminative features to classify known samples.

4.3. Semi-supervised Learning

So far, we have assessed OVANet in the domain adaptation scenario. However, the importance of detecting an

Method	Office (10 / 0 / 11)						Avg
	A2D	A2W	D2A	D2W	W2D	W2A	
OSBP [29]	81.0	77.5	78.2	95.0	91.0	72.9	82.6
ROS [3]	79.0	81.0	78.1	94.4	99.7	74.1	84.4
OVANet	89.5	84.9	89.7	93.7	85.8	88.5	88.7

Table 5: Results of using VGGNet. Open-set DA setting is used.

Method	NAbird (300 / 0 / 255)			
	H-score	UNK	Acc	Acc close
Ent Sep [28]	3.8	1.9	32.2	81.9
DANCE [28]	8.8	4.6	34.0	82.1
OVANet	67.6	63.1	67.4	83.0

Table 6: Results of open-set semi-supervised learning using NAbird [33]. Acc is the instance-wise accuracy for all samples.

unknown class is not limited to domain adaptation. When deploying a recognition model in a real application, a model can encounter samples from classes unseen during training. In semi-supervised learning, a model can encounter unlabeled unknown data during training. Models able to detect unknown samples are desirable. In this section, we provide the analysis of OVANet for the semi-supervised setting.

OVANet is effective in a semi-supervised setting. One interesting application of open-set classification and semi-supervised learning is the classification of animals. The unlabeled images of wild animals can be collected by a monitoring camera or crawling the web. Annotating the images may need expert knowledge. Therefore, the number of annotated images can be limited, and the obtained known label-space may not cover all categories of the dataset. Then, we utilize NAbird [33], a large scale bird image dataset, to evaluate our method, where 300 categories are known and 255 are unknown. Half of the samples are labeled for each known class, and the rest are treated as unlabeled data. All samples of unknown classes are used as unlabeled data. OVANet outperforms the baselines in all metrics (Table 6). This result validates the benefit of OVANet without domain-shift between labeled and unlabeled data.

5. Conclusion

In this paper, we present a new technique, OVANet, which trains a One-vs-All classifier for each class and decides known or unknown by using the output. Our proposed framework is the simplest of all UNDA methods, yet shows strong performance across diverse settings. The extensive evaluation shows OVANet’s applicability to semi-supervised settings.

6. Acknowledgment

This work was supported by Honda, DARPA LwLL and NSF Award No. 1535797.

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79(1-2):151–175, 2010. **1**
- [2] Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1563–1572, 2016. **2, 3**
- [3] Silvia Bucci, Mohammad Reza Loghmani, and Tatiana Tommasi. On the effectiveness of image rotation for open set domain adaptation. In *Eur. Conf. Comput. Vis.*, pages 422–438. Springer, 2020. **1, 2, 5, 6, 8**
- [4] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *Int. Conf. Comput. Vis.*, 2017. **1**
- [5] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *Eur. Conf. Comput. Vis.*, 2018. **1, 2**
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2009. **1, 5**
- [7] Bo Fu, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Learning to detect open classes for universal domain adaptation. In *Eur. Conf. Comput. Vis.* Springer, 2020. **2, 5, 6**
- [8] Mikel Galar, Alberto Fernández, Edurne Barrenechea, Humberto Bustince, and Francisco Herrera. An overview of ensemble methods for binary classifiers in multi-class problems: Experimental study on one-vs-one and one-vs-all schemes. *Pattern Recognition*, 44(8):1761–1776, 2011. **3**
- [9] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Int. Conf. Mach Learn.*, 2014. **2, 6**
- [10] Chuanxing Geng, Sheng-jun Huang, and Songcan Chen. Recent advances in open set recognition: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020. **3**
- [11] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Adv. Neural Inform. Process. Syst.*, 2005. **1, 4**
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Int. Conf. Comput. Vis.*, 2017. **1**
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016. **5**
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Adv. Neural Inform. Process. Syst.*, 2012. **1**
- [15] Jogendra Nath Kundu, Naveen Venkat, and R Venkatesh Babu. Universal source-free domain adaptation. *arXiv preprint arXiv:2004.04393*, 2020. **2, 5**
- [16] Guangrui Li, Guoliang Kang, Yi Zhu, Yunchao Wei, and Yi Yang. Domain consensus clustering for universal domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. **5, 6**
- [17] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. Learning transferable features with deep adaptation networks. In *Int. Conf. Mach Learn.*, 2015. **2**
- [18] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Adv. Neural Inform. Process. Syst.*, 2018. **6**
- [19] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 9(11):2579–2605, 2008. **7**
- [20] Poojan Oza and Vishal M Patel. C2ae: Class conditioned auto-encoder for open-set recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2307–2316, 2019. **3**
- [21] Shreyas Padhy, Zachary Nado, Jie Ren, Jeremiah Liu, Jasper Snoek, and Balaji Lakshminarayanan. Revisiting one-vs-all classifiers for predictive uncertainty and out-of-distribution detection in neural networks. *arXiv preprint arXiv:2007.05134*, 2020. **3, 4**
- [22] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *Int. Conf. Comput. Vis.*, 2017. **2**
- [23] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. *Int. Conf. Comput. Vis.*, 2019. **5**
- [24] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017. **5**
- [25] Pramuditha Perera, Vlad I Morariu, Rajiv Jain, Varun Manjunatha, Curtis Wigington, Vicente Ordonez, and Vishal M Patel. Generative-discriminative feature representations for open-set recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11814–11823, 2020. **3**
- [26] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Adv. Neural Inform. Process. Syst.*, 2015. **1**
- [27] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Eur. Conf. Comput. Vis.*, 2010. **2, 5**
- [28] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, and Kate Saenko. Universal domain adaptation through self supervision. *arXiv preprint arXiv:2002.07953*, 2020. **1, 2, 5, 6, 8**
- [29] Kuniaki Saito, Shohei Yamamoto, Yoshitaka Ushiku, and Tatsuya Harada. Open set domain adaptation by backpropagation. In *Eur. Conf. Comput. Vis.*, 2018. **1, 2, 5, 8**
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv*, 2014. **1, 5**
- [31] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *arXiv preprint arXiv:2007.08176*, 2020. **3**
- [32] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017. **2**
- [33] Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 595–604, 2015. **8**

- [34] Hemant Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017. 5
- [35] Jim Winkens, Rudy Bunel, Abhijit Guha Roy, Robert Stanforth, Vivek Natarajan, Joseph R Ledsam, Patricia MacWilliams, Pushmeet Kohli, Alan Karthikesalingam, Simon Kohl, et al. Contrastive training for improved out-of-distribution detection. *arXiv preprint arXiv:2007.05566*, 2020. 3
- [36] Ryota Yoshihashi, Wen Shao, Rei Kawakami, Shaodi You, Makoto Iida, and Takeshi Naemura. Classification-reconstruction learning for open-set recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4016–4025, 2019. 3
- [37] Kaichao You, Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I. Jordan. Universal domain adaptation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 1, 2, 5, 6