

MLNet: Mutual Learning Network with Neighborhood Invariance for Universal Domain Adaptation

Yanzuo Lu¹, Meng Shen¹, Andy J Ma^{1,2,3*}, Xiaohua Xie^{1,2,3}, Jian-Huang Lai^{1,2,3,4}

¹School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China

²Guangdong Province Key Laboratory of Information Security Technology, China

³Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China

⁴Pazhou Lab (HuangPu), Guangzhou, China

{luyz5, shenm25}@mail2.sysu.edu.cn, {majh8, xiexiaoh6, stsljh}@mail.sysu.edu.cn

Abstract

Universal domain adaptation (UniDA) is a practical but challenging problem, in which information about the relation between the source and the target domains is not given for knowledge transfer. Existing UniDA methods may suffer from the problems of overlooking intra-domain variations in the target domain and difficulty in separating between the similar known and unknown class. To address these issues, we propose a novel **Mutual Learning Network (MLNet)** with neighborhood invariance for UniDA. In our method, confidence-guided invariant feature learning with self-adaptive neighbor selection is designed to reduce the intra-domain variations for more generalizable feature representation. By using the cross-domain mixup scheme for better unknown-class identification, the proposed method compensates for the misidentified known-class errors by mutual learning between the closed-set and open-set classifiers. Extensive experiments on three publicly available benchmarks demonstrate that our method achieves the best results compared to the state-of-the-arts in most cases and significantly outperforms the baseline across all the four settings in UniDA. Code is available at <https://github.com/YanzuoLu/MLNet>.

Introduction

Deep neural networks have made significant progress in computer vision, natural language processing and many other research areas in recent years (LeCun, Bengio, and Hinton 2015). With the rapid development of hardware, existing architectures are expanding with strong data-fitting capacity. Nevertheless, supervised learners may not generalize well when testing on novel domains with distribution shift. Unsupervised domain adaptation (UDA) is proposed to mitigate this by transferring knowledge from the labeled source domain to the unlabeled target domain (Kouw and Loog 2021). Denote the label sets of the source and target domains as \mathcal{C}_s and \mathcal{C}_t , respectively. Though existing works on closed-set domain adaptation (CDA) (Chen et al. 2022a; Rangwani et al. 2022) achieve promising results under the assumption that $\mathcal{C}_s = \mathcal{C}_t$, the validity of this assumption cannot be verified without labels in target domain.

*Corresponding author.

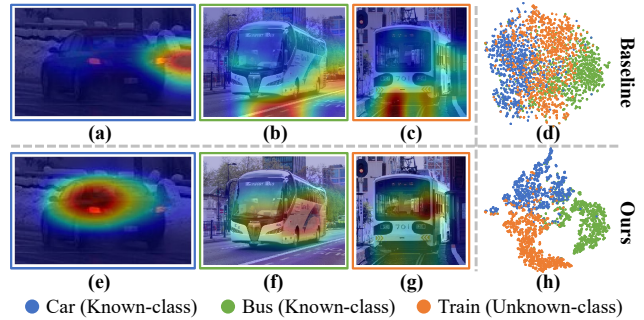


Figure 1: CAM (Selvaraju et al. 2017) and t-SNE (Van der Maaten and Hinton 2008) visualizations on the VisDA (Peng et al. 2017) dataset. Our method outperforms the OVANet baseline by learning more discriminative features for both known and unknown classes even though they are similar. More t-SNE examples are provided in supplementary.

Without the assumption of identical label sets across domains in CDA, several variants of UDA have been developed to solve the practical problem of category shift. In open-set domain adaptation (ODA) (Saito et al. 2018; Jing, Liu, and Ding 2021) \mathcal{C}_s is a subset of \mathcal{C}_t , while \mathcal{C}_s is a superset of \mathcal{C}_t in partial domain adaptation (PDA) (Lin et al. 2022; Cao et al. 2018). If $\mathcal{C}_s \cap \mathcal{C}_t \neq \emptyset$, $\mathcal{C}_s - \mathcal{C}_t \neq \emptyset$ and $\mathcal{C}_t - \mathcal{C}_s \neq \emptyset$, it becomes open partial domain adaptation (OPDA) (Busto and Gall 2017; Kundu et al. 2020). More difficultly but practically, the relation between \mathcal{C}_s and \mathcal{C}_t is unknown, giving rise to the problem of universal domain adaptation (UniDA) (You et al. 2019; Shen et al. 2023). The objective of UniDA is to classify target samples into either one of the known classes in the source domain or an unknown class, without prior knowledge on the relation between label sets across domains.

In UniDA, it is challenging to identify unknown-class samples in the target domain, since the labeled source domain does not contain the unknown class. In several works (You et al. 2019; Fu et al. 2020), a fixed threshold is selected to reject a certain proportion of target samples for learning the concept of the unknown. Due to the difficulty in determining a proper decision threshold, universal classifier is designed for both the known classes and the additional unknown class (Yang et al. 2022; Chen et al. 2022d),

where samples are classified into one of the $|\mathcal{C}_s| + 1$ categories. To address the problem caused by category-shift, Chen et al. construct mixup features within the source domain to simulate the potential unknown-class samples in the target domain. Although the source mixup features obtained by combining samples of different classes will not belong to any of the known classes, they may fail to represent the target unknown class because of the distribution shift.

Another approach to avoid determining the threshold is the OVANet (Saito and Saenko 2021) which trains $|\mathcal{C}_s|$ one-vs-all (open-set) classifiers for each known class. In order to train these classifiers, Saito and Saenko leverage the inter-class distances learned in the source domain and adapt them to the target domain through entropy minimization (Grandvalet and Bengio 2004). Target samples are first classified into one of the known classes (e.g., class l) by a $|\mathcal{C}_s|$ -way closed-set classifier, and then decided as known or unknown by the scores from the l -th open-set sub-classifier. Despite the great success, the feature space may not be discriminative enough for separating the known and unknown classes. As shown in Figure 1(a)(b), the OVANet may wrongly pay attention to the background area even for target samples of the known classes. Moreover, the unknown class may be mistakenly identified as known by simple entropy minimization, so that the attention to the unknown-class samples are misleading as illustrated in Figure 1(c). With the difficulty in distinguishing between known and unknown classes, features of similar objects in the target domain are scattered and mixed together (Figure 1(d)).

To address these challenges, we propose a novel Mutual Learning Network (MLNet) with neighborhood invariance for UniDA. Due to the distribution shift, the intra-domain variations in the target domain differs from the source domain in different degrees. Without taking the intra-domain variations into account, existing UniDA methods cannot learn feature representations well generalized to the target domain. In our method, we propose to learn more generalizable features in the target domain by neighborhood invariance. For this purpose, reliable neighbors of each target sample are retrieved and pairwise similarities between the input and its neighbors are maximized to reduce the intra-domain variations. The neighborhood invariance means that samples within the neighborhood have a high probability to share the same category. A self-adaptive neighbor search strategy based on the nearest neighbor is designed to overcome limitations of different sample sizes within or in different datasets. Moreover, we measure the confidence of each neighbor by the additional information about the neighborhood relationship to guide the invariant feature learning and obtain a more robust embedding space in the target domain.

Beyond learning better feature representation with neighborhood invariance, we propose to further mitigate the category shift by cross-domain manifold mixup (CMM) for unknown-class discovery and mutual learning via consistency constraint. Inter-class distances in the source domain are observed to be not trustworthy when adapted to target domain. Thus, a novel cross-domain manifold mixup scheme is designed to explicitly simulate unknown-class samples by leveraging the arbitrary intermediate states of

mixups across domains to smooth the transfer. Though the CMM greatly improves the unknown-class identification, the decision space of the known class will become smaller such that a number of known-class samples are misidentified as unknown. To trade off this error, we further propose an innovative consistency constraint between the closed-set and open-set classifiers. The misidentified known-class samples with inconsistent decisions are optimized to correct the wrong predictions, while the decisions of unknown-class samples are inherently consistent so that the ability to identify the unknown can be well preserved. As visualized in the bottom row of Figure 1, our method can better separate similar objects of both the known and unknown classes by paying attention to the object areas correctly.

Main contributions of this paper can be summarized as:

- We develop a novel Mutual Learning Network with the ability to identify unknown-class samples more accurately by cross-domain manifold mixup. To the best of our knowledge, we are the first to exploit the relation between closed-set and open-set classifiers for better identifying both known and unknown classes in the literature.
- We propose to learn more generalizable features by neighborhood invariance which maximizes pairwise similarities between each target sample and its neighbors to reduce the intra-class variations in the target domain.
- Our proposed method achieves the best results in most cases compared to the state-of-the-arts and outperforms the baseline by a large margin across all the settings in UniDA on three publicly available benchmarks.

Related Work

Universal domain adaptation (UniDA). UniDA is a practical but challenging problem that assumes no prior knowledge about the relation between the label sets of source and target domains for adaptation. To address the issue of category-shift, early works (You et al. 2019; Fu et al. 2020; Saito et al. 2020; Liang et al. 2021) measure the similarity of target samples with known classes and adopt a manually set threshold (e.g. with a margin) to discriminate between known and unknown classes. DCC (Li et al. 2021) and follow-ups (Chang et al. 2022; Chen et al. 2022b) utilize clustering algorithms to learn virtual prototypes for known classes and make decisions based on the distance from clusters, but are sensitive to the choices of hyperparameters. GATE (Chen et al. 2022d) and OneRing (Yang et al. 2022) propose $(|\mathcal{C}_s| + 1)$ -way universal classifiers to classify all unknown classes into the same one, while OVANet (Saito and Saenko 2021) trains $|\mathcal{C}_s|$ one-vs-all classifiers for each known classes. We argue that the simple entropy minimization in the OVANet to adapt inter-class distances from source domain to target domain can lead to a non-discriminative feature space and confusing predictions. In this paper, we first present the neighborhood invariance learning to better separate different classes in the target domain, and then give details on the novel MLNet with cross-domain mixup to obtain more accurate decision boundaries.

Mixup scheme. For unsupervised domain adaptation, several works (Mao et al. 2019; Yan et al. 2020; Wu, Inkpen,

and El-Roby 2020) take online predictions from classifiers as pseudo-labels for target samples and introduce the vanilla mixup (Zhang et al. 2018) within the target domain to promote the prediction stability. To denoise incorrect pseudo-labels, a semi-supervised approach DeCoTa (Yang et al. 2021) proposes to mix the unlabeled target sample with a labeled one from the source or target domain, such that the obtained mixups at least contain a portion of correct labels. DACS (Tranheden et al. 2021) follows the idea of pseudo-labels but replaces the vanilla mixup with ClassMix (Olsson et al. 2021) that is tailored for semantic segmentation to combine pixels of different positions and classes. Unfortunately, none of the above efforts can be applied to UniDA. Because of the category-shift issue we need to simulate unknown-class samples in the target domain instead of known classes. The existing work most relevant to our proposal is the manifold mixup (Verma et al. 2019) within source domain utilized in GATE (Chen et al. 2022d), which is also designed for the UniDA task. We would like to highlight that the mixups obtained within source domain may fail to represent the unknown due to distribution shift. Therefore, a novel cross-domain manifold mixup is proposed in this paper. We derive that the mixup belongs to a certain known class with low probability, so it is beneficial to unknown-class identification in UniDA.

Method

In universal domain adaptation (UniDA), we are given a labeled source domain $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{N_s}$ and an unlabeled target domain $\mathcal{D}_t = \{(x_j^t)\}_{j=1}^{N_t}$. Let \mathcal{C}_s and \mathcal{C}_t be the label sets of the source and target domains, respectively. Our goal is to accurately classify target samples into either one of the known classes in \mathcal{C}_s or identify them as the unknown, without prior knowledge of the relation between \mathcal{C}_s and \mathcal{C}_t .

Preliminaries

The network architecture of our method is the same as the OVANet (Saito and Saenko 2021). It consists of three modules: (1) a feature extractor \mathcal{F} that maps RGB image of input x into an embedding feature $z = \mathcal{F}(x) \in \mathbb{R}^D$, (2) a closed-set classifier \mathcal{C} to classify input x into one of K known classes, where $p_c(l|x)$ denotes the l -th class probability, and (3) K open-set classifiers \mathcal{O}^l that outputs positive and negative scores to decide whether x is from the l -th known class or the unknown. For simplicity, we only record the positive score as $p_o(l|x)$ for the l -th sub-classifier.

During testing, the known class with maximum probability is first picked by the closed-set classifier. Then, the positive score from corresponding open-set sub-classifier is used to compare with the decision threshold of 0.5 to determine known or unknown, as illustrated in Figure 2. Note that the probabilities from closed-set and open-set classifiers are normalized by Softmax function on different directions.

The training objective of the baseline OVANet can be divided by closed-set and open-set classifiers. The closed-set classifier is trained on the labeled source domain using cross-entropy loss, denoted as $\mathcal{L}_{cls}(x_i^s, y_i^s)$. For the open-set classifiers, Saito and Saenko first introduce hard-negative

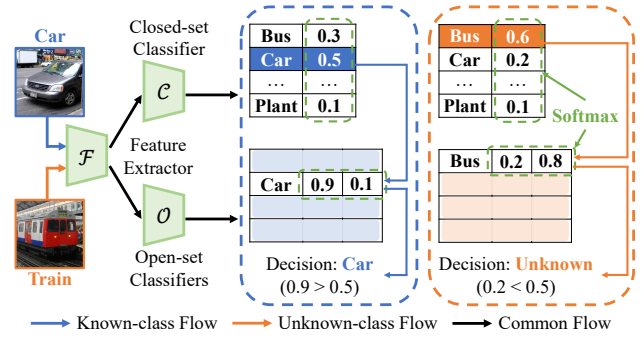


Figure 2: During testing, decisions are made by comparing the positive score with the decision threshold (i.e., 0.5). Only the open-set classifier scores of the maximum-probability known class are used to decide between known or unknown.

classifier sampling (HNCS) to maximise the inter-class margins in the source domain. Then, target samples are aligned to either known or unknown class through open-set entropy minimization (OEM). For HNCS, the one-vs-all loss is computed by a positive class and the hardest negative class, i.e.,

$$\mathcal{L}_{ova}(x_i^s, y_i^s) = -\log p_o(y_i^s | x_i^s) - \min_{l \neq y_i^s} \log(1 - p_o(l | x_i^s)). \quad (1)$$

OEM is a variant of entropy minimization (Grandvalet and Bengio 2004) that performs on each sub-classifier, i.e.,

$$\mathcal{L}_{ent}(x_j^t) = -\frac{1}{K} \sum_{l=1}^K p_o(l | x_j^t) \log p_o(l | x_j^t) + (1 - p_o(l | x_j^t)) \log(1 - p_o(l | x_j^t)). \quad (2)$$

The overall training loss of the baseline is summarized as,

$$\mathcal{L}_{base} = \mathbb{E}_{(x_i^s, y_i^s) \sim \mathcal{D}_s} [\mathcal{L}_{cls}(x_i^s, y_i^s) + \mathcal{L}_{ova}(x_i^s, y_i^s)] + \gamma \mathbb{E}_{x_j^t \sim \mathcal{D}_t} \mathcal{L}_{ent}(x_j^t), \quad (3)$$

where the weight γ is set to 0.1 as the same in OVANet.

Neighborhood Invariance Learning (NIL)

Invariance learning can help to learn discriminative features for unsupervised person re-identification (Zhong et al. 2019; Ding et al. 2020), but has not been exploited in UniDA as far as we know. If the feature embedding is appropriately trained on the source and target domains, target samples within a local neighborhood share the same class with high probability, which we call invariant neighborhood. Thus, we resort to this neighborhood invariance for reducing the intra-domain variations as illustrated in Figure 3(a)(b), such that discriminative features can be learned to better separate different classes in the target domain.

Our method maintains a memory bank $\mathcal{M} \in \mathbb{R}^{N_t \times D}$ for features of all the target samples, where each slot $\mathbf{m}_j \in \mathbb{R}^D$ stores the feature of corresponding sample x_j^t . The memory is updated in place every mini-batch and l_2 -normalized, i.e.,

$$\mathbf{m}_j \leftarrow z_j^t / \|z_j^t\|_2, \text{ where } z_j^t = \mathcal{F}(x_j^t). \quad (4)$$

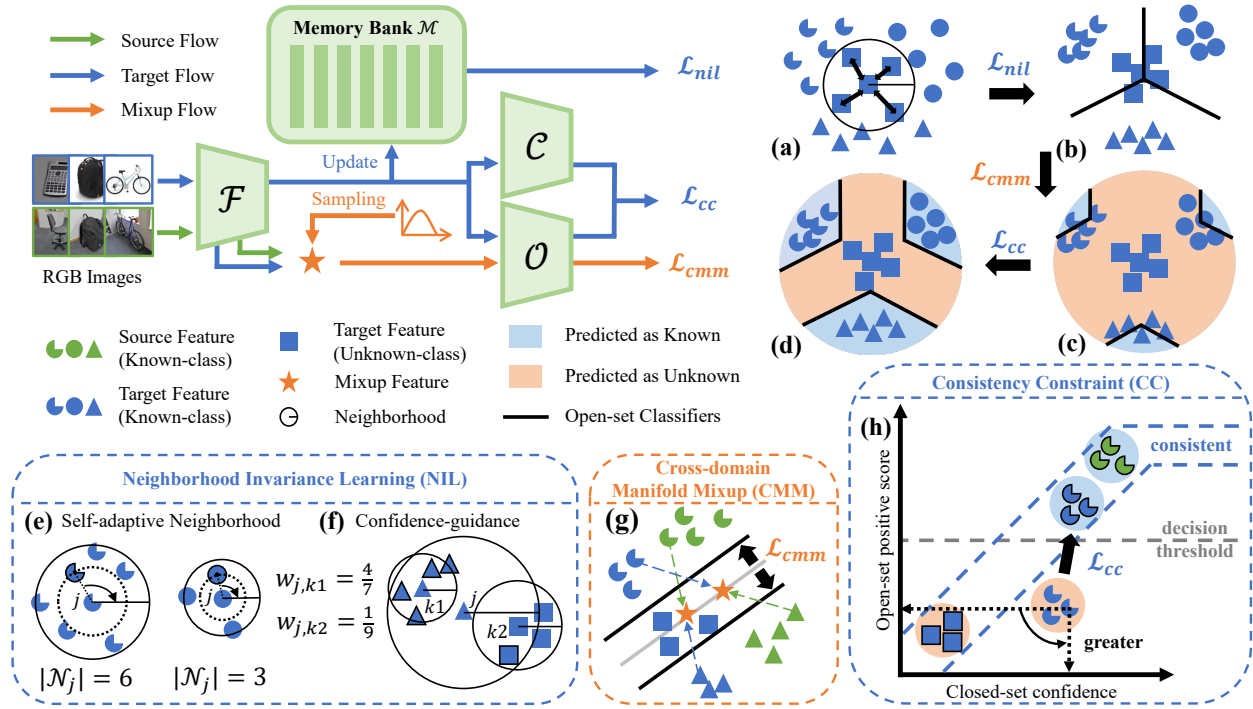


Figure 3: Schematics of MLNet. (a-d) Illustrating the effect of different losses. (e) Neighborhood is adaptive to the sample size of different classes. (f) Neighbors with more similar neighborhoods have higher confidence. (g) We simulate unknown-class samples across domains and utilize them to supervise the open-set classifiers for better unknown-class identification. (h) By optimizing \mathcal{L}_{cc} , misidentified known-class samples are corrected while unknown-class identification is not affected.

Then, the pairwise similarity between the input query x_j^t and the k -th sample in the target domain can then be obtained as,

$$s_{j,k} = \mathbf{m}_j \mathbf{m}_k. \quad (5)$$

Following common practices, we normalize this similarity along all the target instances in the memory using Softmax function. In this way, the memory bank is behaving like a non-parametric classifier or fully connected layer (Wu et al. 2018; Zhong et al. 2019), so that the probability of x_k^t sharing the same class with x_j^t can be estimated, i.e.,

$$p_{jk} = \frac{\exp(\tau \cdot s_{j,k})}{\sum_{n=1}^{N_t} \exp(\tau \cdot s_{j,n})}, \quad (6)$$

where τ is a scaling factor that modulates the sharpness of probability distribution. After that, we reduce the intra-domain variations by maximizing the probabilities between each input sample and its neighbors. However, the quality of neighbor search is crucial that heavily affects whether the outcome embedding space is neighborhood-invariant or not.

Self-adaptive Neighbor Search. A widely used strategy is to select \mathcal{K} -nearest neighbors, but it turned out to be severely limited by the sample size, i.e., the number of samples per class. It is expensive to pick an appropriate \mathcal{K} -value for each dataset or even each class, considering there is also an imbalance of sample sizes within a dataset. Inspired by the idea of self-paced learning (Luo, Song, and Zhang 2022), we turn to a more robust and adaptive approach based on the nearest neighbor, as shown in Figure 3(e).

Concretely, the neighborhood of input x_j^t is defined as,

$$\mathcal{N}_j = \{k \neq j | s_{j,k} > \epsilon \cdot s_{j, \mathcal{N}_j[0]}\}, \quad (7)$$

where its range is based on the relative similarity ratio to the nearest neighbor $\mathcal{N}_j[0]$. This design is advantageous for self-adaptation in two aspects. First, we find ϵ to be highly transferable which avoids setting different \mathcal{K} -value for different datasets. Second, this can mitigate the imbalance of sample sizes within a dataset and promotes efficiency¹.

Confidence-guided Representation Learning. Another technique we employ to improve the robustness is assigning different weights to each neighbor in \mathcal{N}_j , as indicated in Figure 3(f). Instead of directly computing the confidence by the feature embeddings of x_k^t and x_j^t as in Equations (5)(6), the additional information about the neighborhood relationship between each other is utilized. In practice, we use Jaccard distance to measure the confidence that k -th instance in the memory is a neighbor of input sample x_j^t , i.e.,

$$w_{jk} = \frac{|\mathcal{N}_j \cap \mathcal{N}_k|}{|\mathcal{N}_j \cup \mathcal{N}_k|}. \quad (8)$$

Finally, our proposed confidence-guided invariant feature learning objective in the target domain can be written as,

$$\mathcal{L}_{nil}(x_j^t) = -\frac{1}{|\mathcal{N}_j|} \sum_{k \in \mathcal{N}_j} w_{jk} \log p_{jk}. \quad (9)$$

¹Proof is provided in supplementary.

By minimizing the loss function (9), reliable neighbors with high confidence are constrained to have similar feature representations for reducing the intra-domain variations.

Cross-domain Manifold Mixup (CMM)

Though feature representations become more discriminative with neighborhood invariance learning, the problem of category-shift remains. Since only inter-class distances in the source domain are considered by HNCS, a number of target unknown-class samples would be wrongly classified as one of the known classes by OEM in the baseline. We need to equip open-set classifiers with the ability to identify samples close to the decision boundary of known classes as unknown. To achieve this and prevent from distribution shift, unknown-class samples are simulated by a novel cross-domain mixup scheme in a smoother way across domains based on the manifold mixup (Verma et al. 2019).

Concretely, we first randomly sample an interpolation factor from a Beta distribution, i.e., $\lambda \sim \text{Beta}(\alpha, \alpha)$, where α is set to 2.0 for the Beta distribution. The continuous feasible range of the interpolation factor $\lambda \in [0, 1]$ allows our network to leverage arbitrary intermediate states of the mixups across domains to smooth the open-set classifiers. Given the input source and target sample x_i^s and x_j^t , the mixup feature vector across domains is obtained by,

$$z_{i,j,\lambda}^m = \lambda \mathcal{F}(x_i^s) + (1 - \lambda) \mathcal{F}(x_j^t), \quad (10)$$

as illustrated in Figure 3(g). It can be proved that z^m belongs to a certain known class with low probability², so the mixup features are good simulations of potential unknown-class samples. We define the loss based on them to supervise the open-set classifier of corresponding source class, i.e.,

$$\mathcal{L}_{cmm}(x_i^s, y_i^s, x_j^t) = -\log(1 - p_o(y_i^s | z_{i,j,\lambda}^m)), \quad (11)$$

where $p_o(l|z)$ denotes the open-set positive score of input feature $z \in \mathbb{R}^D$ for the l -th class. By minimizing \mathcal{L}_{cmm} , the positive scores of simulated mixup samples become smaller, such that unknown-class samples can be better identified.

Mutual Learning via Consistency Constraint (CC)

Despite the strong ability of CMM to identify unknown-class samples, the decision space of the known class will shrink or even degenerate into the trivial solution to identifying all the target samples as unknown. This is because the mixups obtained using unknown-class samples in ODA and OPDA will overly squeeze the decision space for open-set classifiers to identify known classes. As a result, the performance for known-class identification degrades significantly as illustrated in Figure 3(c). To overcome this limitation, we propose to distinguish the misidentified known-class samples by the relation between the closed-set and open-set classifiers, which has never been studied in the literature.

Let's say a known-class sample x_j^t with groundtruth label $l \in \mathcal{C}_s \cap \mathcal{C}_t$ is misidentified as unknown. We observe that, it generally has high confidence for the maximum-probability known class l in the closed-set classifier, but is with low

positive score in the corresponding open-set sub-classifier. Differently, for unknown-class samples, they usually have smaller closed-set confidence due to both the distribution and category shifts. This means the closed-set confidence of an unknown-class sample is consistent with its low open-set positive score, as illustrated in Figure 3(h). Hence, we can conclude that target samples with inconsistent closed-set confidence and open-set positive score are probably from the known classes misidentified as unknown.

Based on the above observations, we propose a novel consistency constraint for mutual learning between the closed-set and open-set classifiers, i.e.,

$$\mathcal{L}_{cc}(x_j^t) = -\frac{1}{K} \sum_{l=1}^K p_c(l|x_j^t) \cdot p_o(l|x_j^t). \quad (12)$$

The effect of minimizing the consistency constraint (12) is analyzed as follows. Since the closed-set confidence $p_c(l|x_j^t)$ is positive, the partial derivative w.r.t. the open-set positive score $p_o(l|x_j^t)$ is negative, i.e., $-p_c(l|x_j^t)/K < 0$. Therefore, the open-set positive score $p_o(l|x_j^t)$ increases by a magnitude proportional to the closed-set confidence $p_c(l|x_j^t)$ via gradient descend. In this way, the wrong predictions for known-class samples can be corrected due to their high closed-set confidence. For the unknown-class samples, both $p_c(l|x_j^t)$ and $p_o(l|x_j^t)$ are consistently lower, so that the increment is slight and they will not be wrongly identified as the known-class. Finally, the open-set classifiers and decision boundaries are optimized to better separate both the known and unknown classes as shown in Figure 3(d).

Optimization

The overall training objective of our method is given by,

$$\begin{aligned} \mathcal{L}_{all} = & \mathcal{L}_{base} + \mathbb{E}_{\substack{(x_i^s, y_i^s, x_j^t) \\ \sim (\mathcal{D}_s \cup \mathcal{D}_t)}} [\beta_1 \mathcal{L}_{nil}(x_j^t) \\ & + \beta_2 \mathcal{L}_{cmm}(x_i^s, y_i^s, x_j^t) + \eta \mathcal{L}_{cc}(x_j^t)], \end{aligned} \quad (13)$$

where β_1, β_2, η are weights that trade off different losses. They are determined by experiments on Office (Saenko et al. 2010) and generalized to others. In our experiments, we fix β_1 and perform a grid search on the values of β_2 and η .

Experiments

Experimental Setup

Datasets. We conduct experiments on three UniDA benchmarks including Office (Saenko et al. 2010), Office-Home (Venkateswara et al. 2017) and VisDA (Peng et al. 2017). Office consists of 31 categories and contains around 4,700 images across three domains, namely Amazon(A), DSLR(D) and Webcam(W). OfficeHome includes 15,500 images of 65 categories across four domains, while VisDA is a large-scale dataset that contains around 150,000 synthetic and 50,000 real-world images of 12 categories. We use $(|\mathcal{C}_s \cap \mathcal{C}_t| / |\mathcal{C}_s - \mathcal{C}_t| / |\mathcal{C}_t - \mathcal{C}_s|)$ to represent the split for datasets. For more details, please refer to the supplementary.

Evaluation protocols. Following previous studies (Li et al. 2021; Chen et al. 2022d), we use the same evaluation protocols as those used in different settings. In CDA and PDA,

²Proof is provided in supplementary.

Method	Office(10/10/11)								OfficeHome(10/5/50)												VisDA (6/3/3)
	A2D	A2W	D2A	D2W	W2A	W2D	Avg	A2C	A2P	A2R	C2A	C2P	C2R	P2A	P2C	P2R	R2A	R2C	R2P	Avg	
UAN	59.7	58.6	60.1	70.6	60.3	71.4	63.5	51.6	51.7	54.3	61.7	57.6	61.9	50.4	47.6	61.5	62.9	52.6	65.2	56.6	30.5
CMU	68.1	67.3	71.4	79.3	72.2	80.4	73.1	56.0	56.6	59.2	67.0	64.3	67.8	54.7	51.1	66.4	68.2	57.9	69.7	61.6	34.6
DANCE	79.6	75.8	82.9	90.9	77.6	87.1	82.3	61.0	60.4	64.9	65.7	58.8	61.8	73.1	61.2	66.6	67.7	62.4	63.7	63.9	42.8
DCC	88.5	78.5	70.2	79.3	75.9	88.6	80.2	58.0	54.1	58.0	74.6	70.6	77.5	64.3	73.6	74.9	81.0	75.1	80.4	70.2	43.0
GATE	87.7	81.6	84.2	94.8	83.4	94.1	87.6	63.8	75.9	81.4	74.0	72.1	79.8	74.7	70.3	82.7	79.1	71.5	81.7	75.6	56.4
TNT	85.7	80.4	83.8	92.0	79.1	91.2	85.4	61.9	74.6	80.2	73.5	71.4	79.6	74.2	69.5	82.7	77.3	70.1	81.2	74.7	55.3
UniOT	87.0	88.5	88.4	98.8	87.6	96.6	91.2	67.3	80.5	86.0	73.5	77.3	84.3	75.5	63.3	86.0	77.8	65.4	81.9	76.6	57.3
CPR	84.4	81.4	85.5	93.4	91.3	96.8	88.8	59.0	77.1	83.7	69.7	68.1	75.4	74.6	56.1	78.9	80.5	63.0	81.0	72.3	58.2
NCAL	85.3	85.3	88.0	94.0	87.9	95.5	89.3	59.1	88.3	87.3	72.1	73.2	81.0	76.3	57.4	88.4	81.1	62.0	85.4	75.9	62.9
Baseline	83.9	78.5	81.3	95.4	83.2	96.4	86.5	62.2	79.0	80.0	69.2	70.5	76.5	70.9	59.5	80.9	76.8	62.8	79.6	72.3	49.6
Ours	90.4	93.7	89.7	96.2	88.4	98.3	92.8	68.2	83.8	85.0	73.6	78.2	82.2	75.2	64.7	85.1	78.8	69.9	83.9	77.4	69.9

Table 1: H-score comparison in the OPDA setting. Some results are referred to previous work (Chen et al. 2022d). Due to page limit, the comparisons with state-of-the-arts under other settings (i.e., CDA, ODA and PDA) are given in the supplementary.

Dataset	Method	Accuracy			H-Score		
		CDA	PDA	Avg	ODA	OPDA	Avg
Office	Baseline	82.7	90.0	86.4	91.6	86.5	89.1
	Ours	86.9	94.4	90.7	92.8	92.8	92.8
	(Inc.)	+4.2	+4.4	+4.3	+1.2	+6.3	+3.7
Office Home	Baseline	63.9	65.1	64.5	64.1	72.3	68.2
	Ours	69.2	78.1	73.7	67.0	77.4	72.2
	(Inc.)	+5.3	+13.0	+9.2	+2.9	+5.1	+4.0
VisDA	Baseline	58.5	60.8	59.7	61.6	49.6	55.6
	Ours	73.0	80.4	76.7	63.9	69.9	66.9
	(Inc.)	+14.5	+19.6	+17.0	+2.3	+20.3	+11.3

Table 2: Overall comparison against baseline. Average performances count for average classification accuracy in CDA and PDA, and H-score in ODA and OPDA.

we calculate the classification accuracy over all target samples as the performance metric. In ODA and OPDA, we compute the H-score to account for the trade-off between known-class and unknown-class identification, i.e. the harmonic mean of the average of per-class accuracies on known classes and accuracy on the unknown class.

Implementation details. We adopt the same network architecture and optimization setting as the baseline. The feature extractor \mathcal{F} is a ResNet-50 (He et al. 2016) pretrained on ImageNet (Deng et al. 2009). The SGD optimizer with nesterov momentum and inverse scheduler is employed and the batch size is set to 36. The learning rates for the feature extractor \mathcal{F} and all the other components are 0.001 and 0.01, respectively. We train 50 epochs for all datasets and settings, and \mathcal{L}_{nil} is disabled in the first epoch. The scaling factor τ is set to 10.0 as the same in previous work (Luo, Song, and Zhang 2022). The neighborhood relative ratio $\epsilon = 0.875$ and loss weights $\beta_1 = 0.5$, $\beta_2 = 0.1$ are fixed across all datasets and settings. Through extensive experiments, we find that the difficulty in achieving consistency is related to the sample size as detailed in ablation studies, so η is set to 0.16 by default and doubled on VisDA.

Comparison with State-of-the-arts

We compare MLNet with state-of-the-arts across all settings, including UAN (You et al. 2019), CMU (Fu et al. 2020), DANCE (Saito et al. 2020), DCC (Li et al. 2021), OVANet (Saito and Saenko 2021), GATE (Chen et al. 2022d), MATHS (Chen et al. 2022b), TNT (Chen et al. 2022c), UniOT (Chang et al. 2022), CPR (Hur et al. 2023) and NCAL (Su et al. 2023). For fair comparison, we reproduce the baseline (i.e., OVANet) in the same environment.

Overall comparison. Table 2 presents the overall comparison against the baseline across all settings. These results show that the substantial performance improvement of our proposed MLNet over baseline. Especially in the challenging VisDA dataset, the average improvements are 17.0% and 11.3% for the accuracy and H-score, respectively.

OPDA and ODA settings. From Table 1, we can see that MLNet achieves new state-of-the-art results in the challenging OPDA setting across all datasets. When compared to UniOT which is tailored for OPDA, our method still outperforms it by around 1% on Office and OfficeHome, while the improvement is incredible 12.6% on VisDA. The results in the supplementary for ODA show that our method achieves competitive results on OfficeHome and VisDA, and constantly gives the highest H-score on Office.

PDA and CDA settings. In the supplementary, the comparison under the PDA setting shows that MLNet outperforms the state-of-the-art methods averagely by 4.9% and 4.8% on the OfficeHome and VisDA datasets, respectively. For the CDA setting, results show that our method achieves competitive performance compared to existing works.

Ablation Studies

In Table 3, we present the ablation studies on the Office dataset across all settings. Since H-score is a harmonic mean sensitive to extreme values, it is only high when both accuracies are high. Therefore, removing any loss function (Rows 1-3) will result in a severe imbalance of trade-off between known-class and unknown-class identification.

Effect of Confidence-guidance. We also try removing the confidence w_{ij} in Equation (9), where the results show




Figure 4: Ablation studies for (a) consistency constraint and (b-d) hyperparameter sensitivity on the Office dataset. (a) Each row collects results for all settings with a specific target domain (e.g., Amazon with sample size of 94) when varying the loss weight η . A higher average ranking (up to 1) indicates higher performance across the 8 experiments (4 DA settings and 2 sources).

Method	CDA	ODA			OPDA			PDA
		K	U	H	K	U	H	
Ours w/o NIL	79.5	11.9	100.0	17.1	27.1	99.3	35.3	85.9
Ours w/o CMM	85.7	97.4	27.2	41.3	96.3	34.6	50.2	93.0
Ours w/o CC	86.7	3.9	99.5	6.9	64.4	96.5	75.7	94.1
Ours w/o Conf.	86.5	88.7	93.3	90.8	93.1	88.5	90.8	91.9
Ours w/ KNN	85.3	51.6	98.6	62.1	68.4	93.3	77.0	79.7
Ours w/ SMM	85.4	97.0	66.2	77.2	95.5	59.2	72.7	92.9
Ours (full)	86.9	89.6	95.9	92.8	92.3	93.5	92.8	94.4

Table 3: Ablation studies of different components on the Office dataset. Note that K , U , H refer to the known accuracy, unknown accuracy and H-score, respectively. KNN and SMM are short for \mathcal{K} -nearest neighbors and source-domain manifold mixup (Chen et al. 2022d), respectively.

a slight decrease in the performance of different settings (Row 4), suggesting that exploiting the relationship between neighborhoods is helpful in improving the robustness.

Effect of Self-adaptive Neighborhood. In Row 5, we instead use \mathcal{K} -nearest neighbors to determine neighborhood, where \mathcal{K} is set to 5 as the same in GATE and the performance degrades severely when category-shift exists, which validates the superiority of our self-adaptive approach.

Effect of Cross-domain Mixup. To evaluate the contribution of our proposed CMM, we conduct ablation experiments using source-domain manifold mixup as utilized in GATE (Row 6). The results indicate that our cross-domain scheme is more effective since we can leverage arbitrary intermediate states across domains to smooth the transfer.

Effect of Consistency Constraint. Table 4 presents the overall comparison of different choices on \mathcal{L}_{cc} , including cosine distance, JS and KL divergence. We can observe that bridging the closed-set and open-set classifiers is indeed beneficial (Rows 2-5) compared to the baseline in CDA and PDA. Only our \mathcal{L}_{cc} can further mitigate the inconsistency between known-class samples and increase the H-score in ODA and OPDA (Row 6). In Figure 4(a), we collect the average rank versus sample sizes across different settings when varying the weight of consistency constraint η . We observe that the larger the sample size, the greater the weights to achieve the better results. To ensure the robustness of our method, we follow the practice in previous works (Chen

Method	Accuracy			H-Score		
	CDA	PDA	Avg	ODA	OPDA	Avg
Baseline	82.7	90.0	86.4	91.6	86.5	89.1
CC via Cosine	86.7	94.8	90.8	85.9	84.9	85.4
CC via JS-Div	86.6	94.5	90.6	86.2	90.4	88.3
CC via KL-Div (C)	86.5	95.1	90.8	87.7	90.8	89.3
CC via KL-Div (O)	86.3	94.8	90.6	86.9	89.2	88.1
Ours	86.9	94.4	90.7	92.8	92.8	92.8

Table 4: Ablation studies of \mathcal{L}_{cc} on the Office dataset. We normalize open-set positive scores across all classes using Softmax to obtain probability distribution for methods in Rows 3-5. Since $D_{KL}(P||Q)$ is asymmetric, we refer to the case when P denotes closed-set probability and normalized open-set positive scores as KL-Div (C) and (O), respectively.

et al. 2022c,d; Chang et al. 2022) to fix the hyperparameter across different settings and datasets by default and simply double it in the large-scale VisDA for better performance.

Hyperparameter Sensitivity. This experiment is performed with $\beta_1 = 0.5$. More results on $\beta_1 = 0.1$ and 1.0 are provided in the supplementary. In Figure 4(b)(c), we conduct a grid search on the loss weights of β_2, η , where the accuracies in CDA and PDA keep almost unchanged. While the average performance of ODA and OPDA is also stable, the H-score is sensitive to the trade-off between known and unknown classes. Figure 4(d) presents the experiments on varying neighborhood ratio ϵ . Similarly, both accuracy and H-score are only slightly varied by less than 1% and 3%, respectively, which shows the robustness of our method.

Conclusion

In this paper, we propose a novel Mutual Learning Network for universal domain adaptation by incorporating neighborhood invariance and cross-domain mixup. To trade off the misidentified known-class samples, a novel consistency constraint is employed between the closed-set and open-set classifiers. Extensive experiments validate that our method outperforms the baseline significantly and achieve state-of-the-art results in most cases across all settings in UniDA.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (U22A2095, 62276281) and the Science and Technology Program of Guangzhou (202002030371).

References

- Busto, P. P.; and Gall, J. 2017. Open Set Domain Adaptation. In *ICCV*, 754–763.
- Cao, Z.; Ma, L.; Long, M.; and Wang, J. 2018. Partial Adversarial Domain Adaptation. In *ECCV*, 135–150.
- Chang, W.; Shi, Y.; Tuan, H.; and Wang, J. 2022. Unified Optimal Transport Framework for Universal Domain Adaptation. In *NeurIPS*.
- Chen, L.; Chen, H.; Wei, Z.; Jin, X.; Tan, X.; Jin, Y.; and Chen, E. 2022a. Reusing the Task-specific Classifier as a Discriminator: Discriminator-free Adversarial Domain Adaptation. In *CVPR*, 7181–7190.
- Chen, L.; Du, Q.; Lou, Y.; He, J.; Bai, T.; and Deng, M. 2022b. Mutual Nearest Neighbor Contrast and Hybrid Prototype Self-Training for Universal Domain Adaptation. *AAAI*, 36(6): 6248–6257.
- Chen, L.; Lou, Y.; He, J.; Bai, T.; and Deng, M. 2022c. Identical Neighborhood Contrastive Learning for Universal Domain Adaptation. *AAAI*, 36(6): 6258–6267.
- Chen, L.; Lou, Y.; He, J.; Bai, T.; and Deng, M. 2022d. Geometric Anchor Correspondence Mining with Uncertainty Modeling for Universal Domain Adaptation. In *CVPR*, 16134–16143.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*, 248–255.
- Ding, Y.; Fan, H.; Xu, M.; and Yang, Y. 2020. Adaptive Exploration for Unsupervised Person Re-identification. *TOMM*, 16(1): 1–19.
- Fu, B.; Cao, Z.; Long, M.; and Wang, J. 2020. Learning to Detect Open Classes for Universal Domain Adaptation. In *ECCV*, 567–583.
- Grandvalet, Y.; and Bengio, Y. 2004. Semi-supervised Learning by Entropy Minimization. In *NeurIPS*, volume 17.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *CVPR*, 770–778.
- Hur, S.; Shin, I.; Park, K.; Woo, S.; and Kweon, I. S. 2023. Learning Classifiers of Prototypes and Reciprocal Points for Universal Domain Adaptation. In *WACV*, 531–540.
- Jing, T.; Liu, H.; and Ding, Z. 2021. Towards Novel Target Discovery Through Open-Set Domain Adaptation. In *ICCV*, 9322–9331.
- Kouw, W. M.; and Loog, M. 2021. A Review of Domain Adaptation without Target Labels. *TPAMI*, 43(3): 766–785.
- Kundu, J. N.; Venkat, N.; V., R. M.; and Babu, R. V. 2020. Universal Source-Free Domain Adaptation. In *CVPR*, 4544–4553.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature*, 521(7553): 436–444.
- Li, G.; Kang, G.; Zhu, Y.; Wei, Y.; and Yang, Y. 2021. Domain Consensus Clustering for Universal Domain Adaptation. In *CVPR*, 9757–9766.
- Liang, J.; Hu, D.; Feng, J.; and He, R. 2021. UMAD: Universal Model Adaptation under Domain and Category Shift. *arXiv:2112.08553*.
- Lin, K.-Y.; Zhou, J.; Qiu, Y.; and Zheng, W.-S. 2022. Adversarial Partial Domain Adaptation by Cycle Inconsistency. In *ECCV*, 530–548.
- Luo, C.; Song, C.; and Zhang, Z. 2022. Learning to Adapt Across Dual Discrepancy for Cross-Domain Person Re-Identification. *TPAMI*, 45(2): 1963–1980.
- Mao, X.; Ma, Y.; Yang, Z.; Chen, Y.; and Li, Q. 2019. Virtual Mixup Training for Unsupervised Domain Adaptation. *arXiv:1905.04215*.
- Olsson, V.; Tranheden, W.; Pinto, J.; and Svensson, L. 2021. Classmix: Segmentation-based data augmentation for semi-supervised learning. In *WACV*, 1369–1378.
- Peng, X.; Usman, B.; Kaushik, N.; Hoffman, J.; Wang, D.; and Saenko, K. 2017. VisDA: The Visual Domain Adaptation Challenge. *arXiv:1710.06924*.
- Rangwani, H.; Aithal, S. K.; Mishra, M.; Jain, A.; and Babu, R. V. 2022. A Closer Look at Smoothness in Domain Adversarial Training. In *ICML*, 18378–18399.
- Saenko, K.; Kulis, B.; Fritz, M.; and Darrell, T. 2010. Adapting Visual Category Models to New Domains. In *ECCV*, 213–226.
- Saito, K.; Kim, D.; Sclaroff, S.; and Saenko, K. 2020. Universal Domain Adaptation through Self Supervision. In *NeurIPS*, volume 33, 16282–16292.
- Saito, K.; and Saenko, K. 2021. OVA Net: One-vs-All Network for Universal Domain Adaptation. In *ICCV*, 9000–9009.
- Saito, K.; Yamamoto, S.; Ushiku, Y.; and Harada, T. 2018. Open Set Domain Adaptation by Backpropagation. In *ECCV*, 153–168.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 618–626.
- Shen, M.; Lu, Y.; Hu, Y.; and Ma, A. J. 2023. Collaborative Learning of Diverse Experts for Source-free Universal Domain Adaptation. In *ACM MM*, 2054–2065.
- Su, W.; Han, Z.; He, R.; Wei, B.; He, X.; and Yin, Y. 2023. Neighborhood-based credibility anchor learning for universal domain adaptation. *Pattern Recognition*, 142: 109686.
- Tranheden, W.; Olsson, V.; Pinto, J.; and Svensson, L. 2021. DACS: Domain Adaptation via Cross-domain Mixed Sampling. In *WACV*, 1379–1389.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(11).
- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep Hashing Network for Unsupervised Domain Adaptation. In *CVPR*, 5018–5027.

- Verma, V.; Lamb, A.; Beckham, C.; Najafi, A.; Mitliagkas, I.; Lopez-Paz, D.; and Bengio, Y. 2019. Manifold Mixup: Better Representations by Interpolating Hidden States. In *ICML*, 6438–6447.
- Wu, Y.; Inkpen, D.; and El-Roby, A. 2020. Dual Mixup Regularized Learning for Adversarial Domain Adaptation. In *ECCV*, 540–555.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised Feature Learning via Non-parametric Instance Discrimination. In *CVPR*, 3733–3742.
- Yan, S.; Song, H.; Li, N.; Zou, L.; and Ren, L. 2020. Improve Unsupervised Domain Adaptation with Mixup Training. *arXiv:2001.00677*.
- Yang, L.; Wang, Y.; Gao, M.; Shrivastava, A.; Weinberger, K. Q.; Chao, W.-L.; and Lim, S.-N. 2021. Deep Co-Training with Task Decomposition for Semi-Supervised Domain Adaptation. In *ICCV*, 8906–8916.
- Yang, S.; Wang, Y.; Wang, K.; Jui, S.; and van de Weijer, J. 2022. One Ring to Bring Them All: Towards Open-Set Recognition under Domain Shift. *arXiv:2206.03600*.
- You, K.; Long, M.; Cao, Z.; Wang, J.; and Jordan, M. I. 2019. Universal Domain Adaptation. In *CVPR*, 2720–2729.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. In *ICLR*.
- Zhong, Z.; Zheng, L.; Luo, Z.; Li, S.; and Yang, Y. 2019. Invariance Matters: Exemplar Memory for Domain Adaptive Person Re-Identification. In *CVPR*, 598–607.