Open-Set Heterogeneous Domain Adaptation: Theoretical Analysis and Algorithm

Thai-Hoang Pham[1,2], **Yuanlong Wang**[1,2], **Changchang Yin**[1,2], **Xueru Zhang**[1], **Ping Zhang**[1,2]
[1]Department of Computer Science and Engineering, The Ohio State University, USA
[2]Department of Biomedical Informatics, The Ohio State University, USA
{pham.375,wang.16050,yin.731,zhang.12807,zhang.10631}@osu.edu

Abstract

Domain adaptation (DA) tackles the issue of distribution shift by learning a model from a source domain that generalizes to a target domain. However, most existing DA methods are designed for scenarios where the source and target domain data lie within the same feature space, which limits their applicability in real-world situations. Recently, heterogeneous DA (HeDA) methods have been introduced to address the challenges posed by heterogeneous feature space between source and target domains. Despite their successes, current HeDA techniques fall short when there is a mismatch in both feature and label spaces. To address this, this paper explores a new DA scenario called open-set HeDA (OSHeDA). In OSHeDA, the model must not only handle heterogeneity in feature space but also identify samples belonging to novel classes. To tackle this challenge, we first develop a novel theoretical framework that constructs learning bounds for prediction error on target domain. Guided by this framework, we propose a new DA method called Representation Learning for OSHeDA (RL-OSHeDA). This method is designed to simultaneously transfer knowledge between heterogeneous data sources and identify novel classes. Experiments across text, image, and clinical data demonstrate the effectiveness of our algorithm.

1 Introduction

Machine learning (ML) techniques have achieved unprecedented success over the past decades in numerous areas (LeCun, Bengio, and Hinton 2015). However, ML systems are often built on the assumption that training and testing data are independent and identically distributed, which is commonly violated in real-world applications where the environment changes during model deployment. Existing works have shown that the performance of ML models often deteriorates due to distribution shifts between training and testing data (Ben-David et al. 2010; Quiñonero-Candela et al. 2022). To learn a model robust under distribution shifts, domain adaptation (DA) (Ben-David et al. 2010; Mansour, Mohri, and Rostamizadeh 2009) has been proposed to transfer knowledge from a source domain that possesses abundant labeled data to a different but relevant target domain.

Existing DA methods, however, typically assume a homogeneous scenario where the source and target domains have the same feature and label spaces. Consequently, they may fail when the source and target domain data lie in different spaces. For example, heterogeneous feature space is common in biomedical domains in which medical terms undergoes continuous evolution, leading to the retirement of outdated terms (e.g., ICD-9 coding system) and the introduction of novel ones (e.g., ICD-10 coding system) (Grief et al. 2016). In such cases, acquiring training data that seamlessly aligns with target domain's feature space can be impractical or excessively costly. Heterogeneous domain adaptation (HeDA) methods have emerged to handle the heterogeneity observed in distinct feature spaces, which often vary significantly between domains (Li et al. 2020; Zhao et al. 2022).

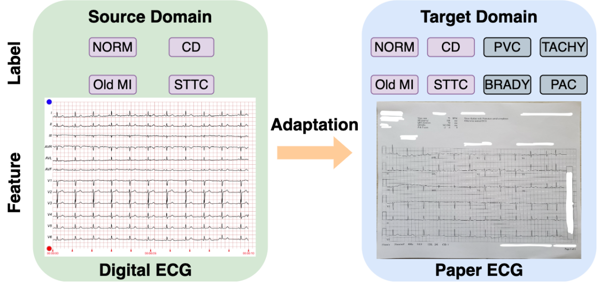

Figure 1: A motivating example about OSHeDA in the context of screening diseases using electrocardiogram (ECG) data. While digital ECGs comprise the majority of labeled data for training ML models for disease screening, physical or paper ECGs remain prevalent worldwide. Thus, the transfer of knowledge from digital ECG datasets is essential to support the training of ML models that analyze paper ECGs. Moreover, ML systems must effectively manage rare abnormalities (indicated with gray boxes), which may not be available in training data, to prevent misdiagnosis.

Despite the significant successes achieved by these HeDA methods, they face a major limitation: current HeDA techniques can only address heterogeneity in feature space and are inadequate when there is a mismatch in both feature and label spaces. This limitation restricts the practical application of HeDA methods in many real-world scenarios because neglecting label mismatch, such as new classes emerging in the target domain, can lead to negative transfer effects from the source to the target domains (Liu et al. 2019).

To overcome this limitation, this study explores a new DA scenario called open-set heterogeneous domain adaptation (OSHeDA). In OSHeDA, ML methods must not only manage heterogeneity in feature space between source and target domains but also identify samples belonging to novel classes in the target domain. Figure 1 illustrates a real example from clinical applications for this novel learning scenario. In this instance, the adaptation process aims to transfer knowledge from digital electrocardiogram (ECG) to paper ECG formats (*heterogeneous*). Moreover, ML models for ECG-based diagnosis must also detect rare abnormalities that were not included in the training data (*open-set*).

To address the challenge of feature and label mismatch in OSHeDA, we first develop a novel theoretical analysis that constructs learning bounds for the prediction error of ML models on the target domain. Guided by this theoretical analysis, we then design a novel representation learning method named **R**epresentation **L**earning for **O**pen-**S**et **He**terogeneous **D**omain **A**daptation (RL-OSHeDA). This method is proposed to transfer knowledge between heterogeneous data sources and identify novel class simultaneously. Unlike existing HeDA methods, RL-OSHeDA transfer knowledge from source to target domains by aligning representations between source and target domains for known classes while also enforcing the representations of novel class in target domains to move apart from the known classes of the source and target domains. To effectively identify samples from novel class within unlabeled data, RL-OSHeDA optimizes a non-negative risk estimator for open-set and employs pseudo labeling to enrich the labeled data.

In summary, the contributions of our work are as follows:

- We conduct a theoretical analysis to establish learning bounds in the OSHeDA scenario. This analysis emphasizes the importance of minimizing the distance between source and target domains for known classes, while maximizing the separation from unknown classes. Moreover, we investigate the impact of pseudo-label and the non-negative risk estimator for open-set in OSHeDA.
- Motivated by the theoretical results, we propose a novel algorithm (RL-OSHeDA) based on representation learning to transfer knowledge from source to target domains.
- We conduct experiments on real data from clinical, computer vision, and natural language processing domains to validate the effectiveness of our method for OSHeDA.

2 Related Works

In this section, we summarize existing research from related areas including heterogeneous domain adaptation, open-set domain adaptation, and open-set semi-supervised learning.

Heterogeneous Domain Adaptation (HeDA). HeDA aims to transfer knowledge across domains with distinct feature spaces and data distributions. Depending on whether unlabeled target data are used in the adaptation process, HeDA approaches are categorized into three types: supervised, semi-supervised, and unsupervised methods. Supervised HeDA methods utilize ample labeled data from both source and target domains for adaptation (Hoffman et al. 2013; Li et al. 2013; Hoffman et al. 2014). In contrast, semi-supervised HeDA methods require only a small number of labeled target domain data and utilize unlabeled instances from the target domain to facilitate transfer (Yao et al. 2020; Li et al. 2020; Fang et al. 2022; Zhao et al. 2022; Yao et al. 2019). Finally, unsupervised HeDA methods operate without any labeled target data, relying solely on unlabeled instances and labeled source data to align cross-domain feature representations (Shen and Guo 2018; Li et al. 2018; Zou et al. 2018). However, successful transfer in unsupervised settings depends on specific assumptions about domain relationships (Liu, Zhang, and Lu 2020).

Open-Set Domain Adaptation (OSDA). OSDA represents a realistic and challenging scenario in DA where the target domain includes instances whose classes are not observed in the source domain, alongside a shift in feature distribution between the two domains. In contrast to the OSHeDA, OSDA assumes a homogeneous feature space between the source and target domains. Existing approaches for OSDA can be categorized into two main groups: adversarial learning and self-supervised learning. Adversarial learning methods employ adversarial networks to detect unknown samples and align the distributions of known samples between the source and target domains (Saito et al. 2018; Luo et al. 2020). On the other hand, self-supervised learning methods utilize techniques like data augmentation to distinguish between known and unknown instances in the target domain (Bucci, Loghmani, and Tommasi 2020; Li et al. 2021).

Open-Set Semi-supervised Learning (OS-SSL). OS-SSL is a SSL scenario that addresses novel classes within unlabeled data during training. Unlike OSDA, OS-SSL requires only a small amount of labeled data. However, it assumes that both labeled and unlabeled data of known classes are drawn from the same distribution, and this setting does not account for novel classes during inference. Methods designed for OS-SSL can be broadly categorized into two types based on how they detect novel classes: criterion-based approaches and detector-based approaches. Criterion-based approaches use heuristic rules to identify novel classes (Chen et al. 2020; Huang, Yang, and Gong 2022; Du et al. 2023; He et al. 2022). In contrast, detector-based approaches employ parameterized detectors to filter outliers (Yu et al. 2020; Huang et al. 2021; Wang et al. 2023; Saito, Kim, and Saenko 2021).

3 Problem Formulation

Notations. Let $\mathcal{X}^d$ and $\mathcal{Y}^d$ denote the feature and label spaces of a domain d associated with a distribution $P_d(X_d, Y_d) : \mathcal{X}^d \times \mathcal{Y}^d \to [0, 1]$ and labeling function $h_d : \mathcal{X}^d \to \Delta\left(\mathcal{Y}^d\right)$ where X_d and Y_d are random variables that take values in $\mathcal{X}^d$ and $\mathcal{Y}^d$, and $\Delta\left(\mathcal{Y}^d\right)$ is a probability simplex over $\mathcal{Y}^d$. Consider a model $h : \mathcal{X}^d \to \Delta\left(\mathcal{Y}^d\right)$, then the *expected error* of h under domain d for some loss function $L : \Delta\left(\mathcal{Y}^d\right) \times \mathcal{Y}^d \to \mathbb{R}_+$ (e.g., 0-1, cross-entropy loss) can be defined as $\mathrm{E}\left(P_d, h\right) = \mathbb{E}_{P_d}\left[L\left(h\left(X_d\right), Y_d\right)\right]$.

Open-Set Heterogeneous Domain Adaptation (OSHeDA) Setup. In DA, we consider $d \in \{s, t\}$ where s and t denote the source and target domains, respectively. Different from conventional DA setup where feature and label spaces remain the same between source and target domains, in OSHeDA, we have $\mathcal{X}^s \neq \mathcal{X}^t$ (*heterogeneous*) and $\mathcal{Y}^s \subset \mathcal{Y}^t$ (*open-set*). Because $\mathcal{Y}^s \subset \mathcal{Y}^t$, we use Y to denote the random variable of label in both source and target domains, and we have $P_s\left(Y \in \mathcal{Y}^t \setminus \mathcal{Y}^s\right) = 0$. Moreover, classes in the sets $\mathcal{Y}^t \setminus \mathcal{Y}^s$ are referred to as unknown in our setting. Given sets of samples $D_s = \{x_i^s, y_i^s\}_{i=1}^{n_s} \overset{i.i.d}{\sim} P_s\left(X_s, Y\right)$ (*source dataset*), $D_{t_l} = \{x_i^t, y_i^t\}_{i=1}^{n_{t_l}} \overset{i.i.d}{\sim} P_t\left(X_t, Y | Y \in \mathcal{Y}^s\right)$ (*labeled target dataset*), and $D_{t_u} = \{x_i^t\}_{i=1}^{n_{t_u}} \overset{i.i.d}{\sim} P_t\left(X_t, Y\right)$ (*unlabeled target dataset*), where n_s, n_{t_l}, n_{t_u} are size of datasets and $n_{t_l} \ll n_s, n_{t_u}$, the goal of OSHeDA is to learn a model $h : \mathcal{X}^t \to \Delta\left(\mathcal{Y}^t\right)$ from D_s, D_{t_l}, D_{t_u} such that the expected error on the target domain $\mathrm{E}\left(P_t, h\right)$ is small.

Representation learning. Representation learning is a common approach for transferring knowledge from a source to a target domain in DA (Zhao et al. 2019; Ganin et al. 2016; Albuquerque et al. 2019; Pham, Zhang, and Zhang 2023), and we will leverage this method in OSHeDA. Specifically, it maps the input spaces $\mathcal{X}^s$ and $\mathcal{X}^t$ of the source and target domains to a shared representation space $\mathcal{Z}$ using two representation mappings: $f_s : \mathcal{X}^s \to \mathcal{Z}$ and $f_t : \mathcal{X}^t \to \mathcal{Z}$. A shared classifier $h : \mathcal{Z} \to \Delta(\mathcal{Y}^t)$ can then be employed to make predictions from this representation space. Notably, h can be utilized for both domains because $\mathcal{Y}^s \subset \mathcal{Y}^t$.

4 Theoretical Analysis

In our analysis, we consider Jensen–Shannon (JS) divergence ($\mathcal{D}_{JS}$) as the statistical distance between two domains. While different distances (Ben-David et al. 2010) were used in domain adaptation literature, we adopt JS divergence because it is aligned with the training objective of adversarial learning (Goodfellow et al. 2014), a technique used in many representation learning-based domain adaptation works (Zhao et al. 2019; Ganin et al. 2016; Pham, Zhang, and Zhang 2023). Next, we present the main theorems, with detailed proofs provided in Appendix A.

4.1 Learning bounds for OSHeDA (infinite case)

To simplify notations used in our following analysis, we denote $P_{t,k}(\cdot) = P_t(\cdot | Y \in \mathcal{Y}^s)$ and $P_{t,u}(\cdot) = P_t(\cdot | Y \notin \mathcal{Y}^s)$ as the distributions of target domain conditioned on known and unknown classes, respectively. We also introduce two distributions P_s^u and P_t^u induced from P_s and P_t by the two mappings f_s^u and f_t^u such that $f_s^u(X^s, Y) = (X^s, unk)$ and $f_t^u(X^t, Y) = (X^t, unk)$ where unk denotes unknown class. In addition, we adopt an assumption commonly used in DA literature (Nguyen et al. 2021; Mansour, Mohri, and Rostamizadeh 2009; Cortes and Mohri 2014) as follows.

Assumption 1 (Bounded loss) *Assume loss function L defined on input space $\mathcal{X}$ and output space $\mathcal{Y}$ is upper bounded by a constant C, i.e., $\forall x \in \mathcal{X}, y \in \mathcal{Y}$, $h \in \mathcal{H}$, we have $L(h(x), y) \leq C$.*

We note that this assumption is indeed reasonable rather than stringent. For example, while Assumption 1 does not hold for the cross-entropy loss typically utilized in classification, we can adjust this loss to ensure that it satisfies Assumption 1 (Pham, Zhang, and Zhang 2024). Based on this assumption, we then provide an upper bound for prediction error on the target domain in OSHeDA as follows.

Theorem 1 *Given a loss function L satisfying Assumption 1, then for any $h \in \mathcal{H}, f_s \in \mathcal{F}_s, f_t \in \mathcal{F}_t$, we have:*

$$\begin{aligned}\mathrm{E}\left(P_t, h \circ f_t\right) \leq\ & \underbrace{\lambda \mathrm{E}\left(P_s, h \circ f_s\right)}_{\textbf{\textit{source error}}} \\ & + \underbrace{\mathrm{E}\left(P_t^u, h \circ f_t\right) - \lambda \mathrm{E}\left(P_s^u, h \circ f_s\right)}_{\textbf{\textit{open-set difference}}} \\ & \underbrace{+ \sqrt{2}\lambda C \Big(\left(\mathcal{D}_{JS}\left(P_s(Z) \parallel P_{t,k}(Z)\right)\right)^{\frac{1}{2}} + \left(\mathcal{D}_{JS}\left(P_s(Z,Y) \parallel P_{t,k}(Z,Y)\right)\right)^{\frac{1}{2}}\Big)}_{\textbf{\textit{domain distance}}}\end{aligned}$$

where $\lambda = P_t(Y \in \mathcal{Y}^s)$, $\mathcal{H}$, $\mathcal{F}_s$, $\mathcal{F}_t$ are hypothesis classes for h, f_s, f_t, and $P_s(Z)$ and $P_{t,k}(Z)$ are the distributions induced from $P_s(X_s)$ and $P_{t,k}(X_t)$ by f_s and f_t, respectively.

Remark 1 *The upper bound in Theorem 1 shed a light on achieving good accuracy on target domain. Specifically, to minimize $\mathrm{E}\left(P_t, h \circ f_t\right)$, the model need to optimize three terms: **(i)** the source error $\mathrm{E}\left(P_s, h \circ f_s\right)$, **(ii)** the open-set difference $\mathrm{E}\left(P_t^u, h \circ f_t\right) - \lambda \mathrm{E}\left(P_s^u, h \circ f_s\right)$, and **(iii)** the distances of marginal and joint distributions between source domain and target domain conditioned on known labels $\mathcal{D}_{JS}\left(P_s(Z) \parallel P_{t,k}(Z)\right)$ and $\mathcal{D}_{JS}\left(P_s(Z,Y) \parallel P_{t,k}(Z,Y)\right)$.*

We want to emphasize that minimizing the distance of the joint distribution between the source and target domains, $\mathcal{D}_{JS}\left(P_s(Z,Y) \parallel P_{t,k}(Z,Y)\right)$, requires knowledge of the label distribution in the target domain $P_{t,k}(Y)$. Therefore, access to labeled target data during training is essential to avoid negative transfer. Note that the concept of open-set difference is not exclusive to OSHeDA. This term also appears in existing works for OSDA (Fang et al. 2020) and positive-unlabeled learning (Kiryo et al. 2017) which are special cases of our setting. Thus, this demonstrates the consistency between our work and the existing literature. Next, we present a lower bound for OSHeDA.

Proposition 1 *Given a loss function L satisfying Assumption 1, then for any $h \in \mathcal{H}, f_s \in \mathcal{F}_s, f_t \in \mathcal{F}_t$, we have:*

$$\begin{aligned}\mathrm{E}\left(P_t, h \circ f_t\right) \geq\ & \lambda \mathrm{E}\left(P_{t,k}, h \circ f_t\right) + (1-\lambda) \mathrm{E}\left(P_s^u, h \circ f_s\right) \\ & - \sqrt{2}(1-\lambda) C \left(\mathcal{D}_{JS}\left(P_s(Z) \parallel P_{t,u}(Z)\right)\right)^{\frac{1}{2}}\end{aligned}$$

where $P_{t,u}(Z)$ is distribution induced from $P_{t,u}(X_t)$ by f_t.

Remark 2 *Theorem 1 shows the necessity of reducing $\mathrm{E}\left(P_s, h \circ f_s\right)$ to achieve high accuracy on target domain. However, it may unavoidably increase $\mathrm{E}\left(P_s^u, h \circ f_s\right)$. This observation, combined with Proposition 1, suggests that to avoid the large lower bound for the target error $\mathrm{E}\left(P_t, h \circ f_t\right)$, we should increase the distance of the marginal distribution between the source domain and the unknown data in target domain, $\mathcal{D}_{JS}\left(P_s(Z) \parallel P_{t,u}(Z)\right)$. In other words, we should segregate the representations of known classes from those of unknown class.*

4.2 Learning bound for OSHeDA (finite case)

The learning bounds discussed in Section 4.1 are only applicable for the setting when we have access to unlimited data from source and target domains. In such cases, minimizing JS divergence of data distribution between these domains is equivalent to achieving invariant representations through adversarial learning (Goodfellow et al. 2014). However, we only work with finite data in practice. Thus, we present the following result, which provides a guarantee for using adversarial learning to optimize JS divergence from finite data.

Proposition 2 (Adapted from Biau et al. (2020)) *The error in minimizing JS divergence of data distributions between source and target domains in representation space, using finite data, is up to* $\mathcal{O}\left(\left(1/\sqrt{n_s} + 1/\sqrt{n_t}\right)\right)$.

where n_s and n_t are the size of source and target datasets.

Remark 3 *Proposition 2 states that the performance of minimizing JS divergence from finite data is proportional to the dataset size. Note that in OSHeDA, we only have access to limited label data from target domain which then results in significant error in estimating JS divergence only from labeled source and target data. In essence, this result underscores the need for the development of an effective approach to utilize unlabeled target data for estimating the JS divergence, which involves techniques like pseudo-labeling.*

Therefore, we apply pseudo-labeling on unlabeled data to enrich labeled target data. Let g be pseudo-label model and denote $\mathrm{N}(P_{t,k}, g) = \mathbb{E}\left[\mathcal{D}_{JS}\left(P_{t,k}(g(Z)) \parallel P_{t,k}(Y|Z)\right)\right]$ as the noise of g with respect to the target domain conditioned on known labels. Then, the impact of pseudo-labeled data can be illustrated in a new bound for OSHeDA as follows.

Theorem 2 *Given a loss function* L *satisfying Assumption 1, for any* $0 < \delta < 1$*, with probability at least* $1 - \delta$*, the following holds for all* $h \in \mathcal{H}, f_s \in \mathcal{F}_s, f_t \in \mathcal{F}_t$:

$$
\begin{aligned}
&\mathrm{E}\left(P_t, h \circ f_t\right) \leq \lambda \widehat{\mathrm{E}}\left(P_s, h \circ f_s\right) + \widehat{\mathrm{E}}\left(P_t^u, h \circ f_t\right) \\
&- \lambda \widehat{\mathrm{E}}\left(P_s^u, h \circ f_s\right) + \sqrt{2}\lambda C\left(\left(\mathcal{D}_{JS}\left(P_s(Z) \parallel P_{t,k}(Z)\right)\right)^{\frac{1}{2}}\right. \\
&\left.+ \left(\mathcal{D}_{JS}\left(P_s(Z,Y) \parallel P_{t,k}(Z, g(Z))\right)\right)^{\frac{1}{2}} + \left(\mathrm{N}(P_{t,k}, g)\right)^{\frac{1}{2}}\right) \\
&+ \mathcal{O}\left(\lambda C \sqrt{\frac{d_s \log n_s + d_s \log |\mathcal{Y}^t| + \log \frac{1}{\delta}}{n_s}}\right. \\
&\left.+ C \sqrt{\frac{d_t \log n_t + d_t \log |\mathcal{Y}^t| + \log \frac{1}{\delta}}{n_t}}\right)
\end{aligned}
$$

where $\widehat{\mathrm{E}}\left(P_s, h \circ f_s\right)$, $\widehat{\mathrm{E}}\left(P_t^u, h \circ f_t\right)$, $\widehat{\mathrm{E}}\left(P_s^u, h \circ f_s\right)$ *are empirical errors calculated on samples from distributions* P_s, P_t^u, P_s^u, $n_t = n_{t_l} + n_{t_u}$*, and* d_s, d_t *are Natarajan dimension (Natarajan 1989) of hypothesis classes* $\mathcal{H} \circ \mathcal{F}_s$, $\mathcal{H} \circ \mathcal{F}_t$.

Theorem 2 shows that the error in the target domain depends on the quality of the pseudo-label model g, with higher-quality g being more effective at reducing noise. Additionally, the bound emphasizes the importance of aligning the joint distributions between the source and target domains in OSHeDA. This makes OSHeDA more challenging compared to homogeneous DA (HoDA), where source and target data lie on the same space. In contrast, HoDA methods can attain good performance under certain conditions by solely aligning the marginal distributions of representations between source and target domains. We will illustrate this contrast through the bound for HoDA in Section 4.3.

4.3 Learning bound for HoDA

Before constructing the learning bound for HoDA, we introduce an assumption about the representation Z as follows.

Assumption 2 (Sufficient representation) *Let* $I_s(\cdot, \cdot)$ *be the mutual information between two random variables in the source domain. We assume* $I_s(Z, Y) = I_s(X_s, Y)$*. In particular,* $I_s(Z, Y) = \mathcal{D}_{KL}\left(P_s(Z, Y) \parallel P_s(Z) \otimes P_s(Y)\right)$ *and* $I_s(X_s, Y) = \mathcal{D}_{KL}\left(P_s(X_s, Y) \parallel P_s(X_s) \otimes P_s(Y)\right)$ *where* D_{KL} *is KL divergence between two distributions.*

Note that Assumption 2 is reasonable because we have access to labeled data of source domain and the dimension of $\mathcal{Y}$ is often smaller than that of $\mathcal{Z}$. Based on this, we establish the learning bound in HoDA under the covariate shift below.

Proposition 3 *Suppose Assumptions 1 and 2 hold and the distribution shift between source and target domains is covariate shift (i.e.,* $P_s(X) \neq P_t(X), P_s(Y|X) = P_t(Y|X)$*), then for any* $h \in \mathcal{H}$ *and* $f \in \mathcal{F}$*, we have:*

$$
\mathrm{E}\left(P_t, h \circ f\right) \leq \mathrm{E}\left(P_s, h \circ f\right) + \sqrt{2} C\left(\mathcal{D}_{JS}\left(P_s(Z) \parallel P_t(Z)\right)\right)^{\frac{1}{2}}
$$

In HoDA, due to the homogeneity of the input space, we can utilize a single representation mapping f for both the source and target domains. Note that the bound in Proposition 3 depends solely on the distance of the marginal distributions between the source and target domains, $\mathcal{D}_{JS}\left(P_s(Z) \parallel P_t(Z)\right)$, which can be effectively minimized even without access to labeled data in the target domain. Clearly, covariate shift assumption is only reasonable in HoDA, where the source and target data share the same feature and label spaces.

5 Methodology

Motivated by theoretical results presented in Section 4, we introduce RL-OSHeDA, a representation learning method specifically designed for OSHeDA. Our method aims to simultaneously optimize both the upper bound in Theorem 2 and the lower bound in Proposition 1. RL-OSHeDA features two distinct representation mappings, f_s and f_t, which map heterogeneous source and target feature spaces to a shared representation space, along with a classifier h that makes predictions based on these representations. Figure 2 presents the overall architecture of RL-OSHeDA, while pseudo code describing training process can be found in Appendix B.2.

5.1 Objective function

To improve predictive performance in OSHeDA, our method targets the following: (i) minimizing prediction errors on both source and labeled target data, (ii) minimizing the distances of marginal and label-conditioned representation distributions for known classes between source and target

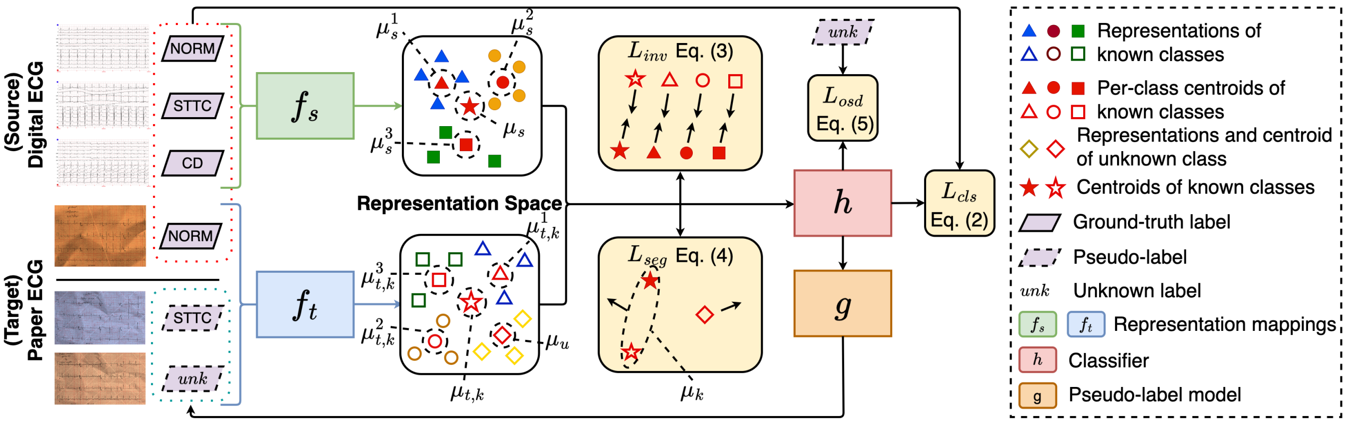

Figure 2: Overall architecture of RL-OSHeDA is illustrated with a motivating example from ECG-based diagnosis application. We leverage 2-stage learning process to update model parameters. In stage 1, model parameters are updated by optimizing L_{cls}. In stage 2, model parameters are updated by optimizing L_{cls}, L_{inv}, L_{seg}, and L_{osd} with the help from pseudo-label model g.

data, (iii) maximizing the distances of marginal representation distributions between known and unknown classes, and (iv) minimizing the open-set difference. Specifically, RL-OSHeDA optimizes the following objective function:

$$L = L_{cls} + L_{inv} - L_{seg} + L_{osd} \tag{1}$$

where L_{cls} is the classification error computed from source and labeled target datasets D_s and D_{t_l}, defined as follows:

$$L_{cls} = \frac{\lambda}{n_s}\sum_{i=1}^{n_s} \text{CE}\left(h\left(f_s\left(x_i^s\right)\right), y_i^s\right) + \frac{1}{n_{t_l}}\sum_{i=1}^{n_{t_l}} \text{CE}\left(h\left(f_t\left(x_i^t\right)\right), y_i^t\right) \tag{2}$$

Here CE is the cross-entropy loss.

L_{inv} denotes the distances of marginal and label-conditioned representation distributions for known classes between source and target datasets. Note that, we minimize the distance of the label-conditioned representation distribution $P(Z|Y)$, rather than the joint distribution $P(Z, Y)$, as noted by Pham, Zhang, and Zhang (2023). As shown in Proposition 2, L_{inv} can be defined based on JS divergence and minimized through adversarial learning. However, the number of discriminators required for this approach scales linearly with the number of classes, leading to instability in training when the dataset has a large number of classes. To address this issue, we implement L_{inv} using maximum mean discrepancy (MMD) defined as follows:

$$L_{inv} = \left\|\mu_s - \mu_{t,k}\right\|_2^2 + \sum_{m=1}^{|\mathcal{Y}^s|} \left\|\mu_s^m - \mu_{t,k}^m\right\|_2^2 \tag{3}$$

where μ_s (resp. $\mu_{t,k}$) is centroid of representations from source data (resp. target data belonging to known classes), and μ_s^m (resp. $\mu_{t,k}^m$) is centroid of representations from source data (resp. target data) belonging to known class m. Note that $\mu_{t,k}$ and $\mu_{t,k}^m$ are computed using both instances with ground-truth labels from labeled target data and those with high-quality pseudo-labels (see Section 5.2) from unlabeled target data to provide a more accurate estimation.

L_{seg} is the distances between marginal representation distributions of known and unknown classes. Similarly, we implement L_{seg} with MMD as follows:

$$L_{seg} = \left\|\mu_k - \mu_u\right\|_2^2 \tag{4}$$

where μ_k (resp. μ_u) are centroids of representations from both source and target datasets belonging to ground-truth and pseudo known (resp. unknown) classes.

L_{osd} represents the open-set difference, as detailed in Theorem 2. The optimal value for the open-set difference is 0. However, due to the flexibility of deep neural networks, this term can become excessively negative during training and adversely affect model performance. To address this issue, we implement L_{osd} as a non-negative risk estimator:

$$L_{osd} = \max\left(0, \frac{1}{n_t}\sum_{i=1}^{n_t} \text{CE}\left(h\left(f_t\left(x_i^t\right)\right), unk\right) - \frac{\lambda}{n_s}\sum_{i=1}^{n_s} \text{CE}\left(h\left(f_s\left(x_i^s\right)\right), unk\right)\right) \tag{5}$$

where $n_t = n_{t_l} + n_{t_u}$ is the size of the target dataset.

5.2 Pseudo-labeling using 2-stage learning

The accuracy of L_{inv} and L_{seg} highly depends on the quality of pseudo-labels. Traditionally, the pseudo-label model g is derived by modifying the classifier h (e.g., using hard labels calculated from h's outputs as pseudo-labels), which creates a coupling between g and h. Specifically, g is defined as $a \circ h$, where a is an operator applied to the output of h (e.g., $a := \arg\max$). When the distributions of the source and target domains are well-aligned, this coupling is harmless, as the optimal solution for g also aligns with that for h. However, at the beginning of the training process, when the distributions of the source and target domains are not aligned, g and h have completely different objective functions, resulting in a trade-off between them. To address this issue, we propose a 2-stage learning approach as follows:

- **Stage 1** ($epoch < T$): Update f_s, f_t, h using L_{cls}.
- **Stage 2** ($epoch \geq T$): Update f_s, f_t, h using L.

	CIFAR10 & ILSVRC2012			ImageCLEF-DA			Multilingual Reuters Collection			NUSWIDE & ImageNet		
	HOS	OS^*	UNK	HOS	OS^*	UNK	HOS	OS^*	UNK	HOS	OS^*	UNK
DS3L	61.49±0.74	59.04±1.00	64.40±1.06	58.62±2.04	52.87±2.70	66.74±2.77	59.35±0.94	52.92±1.27	67.57±1.24	67.61±1.65	66.17±2.22	69.20±2.30
KPG	57.27±0.50	54.73±0.00	60.30±1.11	40.68±0.94	34.60±0.00	50.79±2.91	11.27±0.07	8.59±0.00	17.04±0.96	55.18±1.17	52.60±0.00	58.10±2.45
OPDA	53.30±0.77	48.22±1.00	60.26±1.11	53.17±1.83	45.28±2.13	65.76±2.76	55.85±0.96	48.47±1.23	65.94±1.23	71.06±1.44	66.60±1.98	76.38±2.09
PL	42.75±0.52	37.12±0.49	52.31±1.10	39.20±1.62	31.93±1.66	54.34±2.91	42.85±0.81	34.56±1.00	57.86±1.29	42.43±0.26	34.05±0.00	61.15±2.10
SCT	59.61±0.75	57.35±1.00	62.33±1.08	58.76±2.05	53.09±2.71	66.71±2.76	61.17±0.94	**54.96±1.30**	69.00±1.21	70.42±1.49	68.00±2.20	73.10±1.99
SSAN	60.38±0.73	59.01±1.00	62.01±1.08	58.61±2.05	53.18±2.74	66.14±2.74	58.25±0.93	51.99±1.25	66.26±1.24	67.98±1.49	66.25±2.04	69.85±2.21
STN	61.59±0.72	58.80±0.98	64.87±1.05	56.25±2.06	49.80±2.69	65.84±2.76	59.21±0.96	52.91±1.31	67.24±1.23	67.75±1.23	64.80±1.42	71.08±2.16
SL	60.74±0.74	58.29±1.00	63.67±1.08	58.59±2.05	52.84±2.70	66.71±2.76	58.53±0.96	52.14±1.32	66.74±1.21	69.41±1.64	66.63±2.26	72.57±2.23
RL-OSHeDA	**72.33±0.70**	**67.88±0.98**	**77.81±0.97**	**63.98±2.04**	**56.63±2.72**	**74.80±2.51**	**65.39±0.91**	54.47±1.21	**81.97±0.96**	**80.01±1.30**	**74.65±2.01**	**86.35±0.81**

	Office & Caltech256			Wikipedia			PTB-XL			Average over datasets		
	HOS	OS^*	UNK	HOS	OS^*	UNK	HOS	OS^*	UNK	HOS	OS^*	UNK
DS3L	72.06±2.48	67.41±3.68	78.15±2.89	56.00±2.01	50.72±2.36	66.24±2.80	30.30±1.19	34.95±0.58	26.74±1.82	57.92±1.58	54.87±1.97	62.72±2.12
KPG	34.46±0.99	29.18±0.00	45.34±3.58	24.82±0.42	16.40±0.00	52.52±3.18	N/A	N/A	N/A	37.28±0.68	32.68±0.00	47.35±2.36
OPDA	65.23±2.58	57.70±3.43	76.23±3.03	52.66±1.92	45.94±1.81	65.42±3.12	31.47±1.22	36.35±0.63	27.74±1.86	54.67±1.53	49.79±1.74	62.53±2.17
PL	48.92±1.36	40.38±1.13	68.41±3.40	41.87±1.65	35.14±1.48	58.40±3.10	26.18±1.37	36.43±0.55	20.43±1.66	40.60±1.08	35.66±0.90	53.27±2.22
SCT	75.72±2.14	71.05±3.20	81.79±2.54	58.41±2.01	52.86±2.39	68.42±2.74	26.23±1.65	**46.48±1.71**	18.27±1.60	59.89±1.58	57.10±1.93	64.49±1.99
SSAN	72.95±2.36	67.99±3.37	79.67±2.88	58.37±1.76	52.76±1.95	68.36±2.80	25.16±1.47	40.40±0.65	18.27±1.54	57.39±1.54	55.94±1.86	61.51±2.07
STN	72.26±2.28	66.46±3.30	79.84±2.75	57.75±1.91	51.40±2.13	69.00±2.97	27.08±0.96	22.63±0.26	33.72±1.65	57.41±1.45	52.40±1.73	64.51±2.08
SL	72.14±2.54	67.72±3.75	77.89±2.96	57.10±1.97	51.60±2.19	67.04±2.76	25.74±1.55	44.50±0.76	18.11±1.52	57.46±1.64	56.24±2.00	61.82±2.08
RL-OSHeDA	**78.18±2.05**	**73.04±2.91**	**85.25±2.50**	**63.10±1.89**	**57.26±2.45**	**73.04±2.37**	**47.48±1.25**	44.30±1.39	**51.16±1.86**	**67.21±1.45**	**61.18±1.95**	**75.77±1.71**

Table 1: Prediction performances (HOS, OS^*, UNK) of RL-OSHeDA and baselines for OSHeDA scenario on 7 datasets. We report average results over 10 random seeds for each dataset.

where T is a threshold indicating when to switch from stage 1 to stage 2. In stage 1, optimizing L_{cls} partially aligns the source and target domains, thereby reducing the trade-off between g and h during the optimization of L in stage 2. Additionally, rather than simply using the hard labels with the largest logits from h as the output of g, we propose generating pseudo-labels as follows:

- First, select pseudo-labels as $g(x^t) = a'(h(f_t(x^t)))$ where a' is $\arg\max$ operator applied to the logits of the known classes only.
- Then, select $1 - \lambda$ fraction of instances with the smallest maximum logits and assign pseudo-labels unk to them.

The motivation behind this design of the pseudo-label model g is that, at the beginning of stage 2, there is no supervision signal for training the parameters of h related to unknown class. Therefore, relying solely on logits to determine the unknown class is unreliable. Note that this strategy is only used to generate pseudo-labels during the training of stage 2. Once training is complete and h's parameters for unknown class are well-trained by optimizing L_{osd}, we can simply use $\arg\max$ across all classes to generate predictions.

6 Experiments

Next, we empirically evaluate the performance of our methods across clinical, computer vision, and natural language processing applications. We focus on the OsHeDA scenario, characterized by heterogeneity in the feature space between the source and target domains, with the label space of the target domain encompassing both known and unknown classes.

6.1 Experimental setup

Datasets. We conduct our experiments on 7 datasets including CIFAR10 (Krizhevsky 2009) & ILSVRC2012 (Russakovsky et al. 2015); Wikipedia (Rasiwasia et al. 2010); Multilingual Reuters Collection (Amini, Usunier, and Goutte 2009); NUSWIDE (Chua et al. 2009) & ImageNet (Deng et al. 2009); Office (Saenko et al. 2010) & Caltech256 (Griffin et al. 2007); ImageCLEF-DA (Griffin et al. 2007); PTB-XL (Wagner et al. 2020). These datasets results in 56 DA tasks. Detailed descriptions and statistics of these datasets are provided in Appendix C.1.

Baselines. We compare our method with several representative methods from **HeDA** (SSAN (Li et al. 2020), STN (Yao et al. 2019), SCT (Zhao et al. 2022), KPG (Gu et al. 2022)), **OSDA** (OPDA (Saito et al. 2018)), and **OS-SSL** (DS3L (Guo et al. 2020)) literature. For the HeDA methods, they are trained on both source and target data. In contrast, OSDA and OS-SSL methods are trained only on target data as they cannot handle heterogeneous feature spaces. During inference, HeDA and OS-SSL methods classify instances as unk using the same method as our pseudo-label model g (see Section 5.2). Additionally, we explore supervised learning (SL) and pseudo-labeling (PL) methods trained on target data. Among all baselines, only KPG is designed to handle OSHeDA by combining Gromov-Wasserstein distance and partial optimal transport (Xu et al. 2020). Since λ is a required input for most methods in our experiments, we utilize techniques from positive-unlabeled learning (Zeiberg, Jain, and Radivojac 2020) to estimate λ. Detailed architectures of our model and the baselines are in Appendix B.1.

Evaluation method. We utilize HOS, the harmonic mean of OS^* and UNK (Bucci, Loghmani, and Tommasi 2020). OS^* is the class-wise averaged accuracy of known classes, while UNK measures the accuracy for the unknown class. HOS is particularly suitable for OSHeDA because it emphasizes the ability to both correctly classify known classes and detect out-of-distribution instances simultaneously. In particular, this metric increases when the performance in both known and unknown classifications is high.

6.2 Experimental results

OSHeDA benchmark. The prediction performance (HOS) of RL-OSHeDA and the baselines is summarized in Table 1. RL-OSHeDA consistently outperforms all baselines across all datasets, demonstrating its effectiveness in simultaneously addressing heterogeneity in the feature space and open-set in the label space during training. Among the baselines, KPG is specifically designed for OSHeDA by using

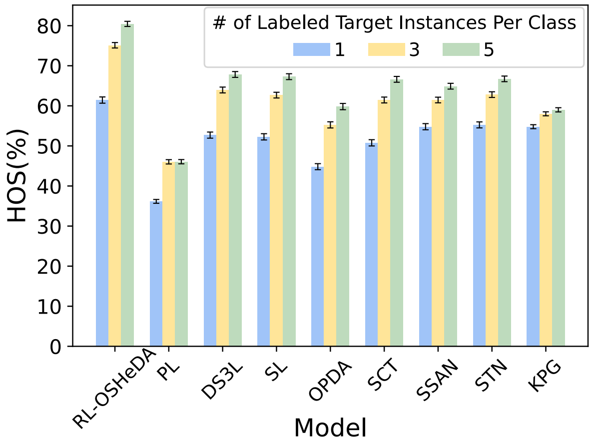

Figure 3: Performances w.r.t. different number of labeled target instances per class on CIFAR10 & ILSVRC2012 dataset.

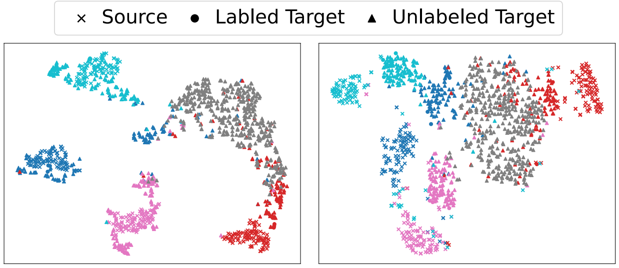

(a) RL-OSHeDA

(b) STN

Figure 4: Visualization of representation spaces learned by RL-OSHeDA and STN for NUSWIDE & ImageNet dataset. Different colors represent different classes, with the unknown class denoted in grey.

Align	Segregate	OSD	2-stage	HOS	OS^*	UNK
✓	✓	✓	✓	**65.39**	**54.47**	**81.97**
✓	✓	✓	✗	59.40	49.47	74.42
✓	✗	✓	✓	61.92	52.63	75.37
✗	✓	✓	✓	58.23	51.13	68.01
✓	✓	✗	✓	59.96	53.10	68.97
✗	✗	✗	✗	58.33	51.86	66.68

Table 2: Ablation study for RL-OSHeDA on Multilingual Reuters Collection dataset. **Align** refers to using L_{inv}; **Segregate** refers to using L_{seg}; **OSD** refers to using L_{osd}; **2-stage** refers to using 2-stage learning approach.

optimal transport. Then, SVM trained on transported source and labeled target data is used to make prediction. However, this method underperforms in our evaluation due to its difficulty in correctly transporting from source to target data. Moreover, this method is not applicable for complex data structures, such as those found in PTB-XL dataset. Other baselines achieve better prediction performances, but their HOS remains suboptimal due to their inability to handle novel classes or heterogeneous source data during training.

To further validate the superiority of RL-OSHeDA across all datasets, we conduct significance testing, including Friedman test followed by Nemenyi test (Demšar 2006). The results (see Appendix C.4) show that our method significantly outperforms the baselines, with P-values < 0.05. Among all the baselines, SCT, SSAN, STN, SL, and DS3L exhibit better prediction performances than KPG, OPDA, and PL. Note that all methods, except OPDA and KPG, utilize our approach to detect the unknown class based on logits of known classes. This result suggests that while this approach can partially address the open-set issue, it cannot fully resolve it. For OPDA, although it is designed to handle open-set issue, its inability to leverage heterogeneous source data limits its performance to adapting with only a small labeled target dataset, resulting in suboptimal performance.

Ablation study. We conduct an ablation study to better understand the contribution of each component in the objective function of our method. As shown in Table 2, removing any component deteriorates model performance. This finding highlights the importance of achieving a good pseudo-label model using 2-stage learning approach as well as aligning the data distribution of known classes between source and target domains while simultaneously detecting and segregating unknown class from known ones for OSHeDA.

Impact of labeled target data. We vary the number of instances per class in the labeled target data to investigate their impact on the DA process. Specifically, we conduct experiments on CIFAR10 & ILSVRC2012 dataset with 1, 3, and 5 instances per class in the labeled target data and visualize the result in Figure 3. Generally, we observe that increasing the number of labeled target instances facilitates better alignment and enhances the performance of all methods. This result demonstrates the importance of labeled target data for DA methods in OSHeDA.

Visualization of representation space. We perform a qualitative analysis to examine the learned representations of RL-OSHeDA and STN for NUSWIDE & ImageNet dataset. Specifically, we use t-SNE (Van der Maaten and Hinton 2008) to project these representations into a 2-dimensional space. As shown in Figure 4, our method effectively aligns representations of the known classes between source and target domains while simultaneously segregating the representations of the unknown class (grey color). This results in improved HOS scores compared to STN.

7 Conclusion

This paper studied a novel domain adaptation scenario called open-set heterogeneous domain adaptation (OSHeDA). We first conducted a theoretical analysis to establish learning bounds in OSHeDA. Based on these theorems, we proposed a representation learning method that aligns the data distribution of known classes between source and target domains while simultaneously detecting and segregating unknown class from known ones. The resulting models trained with the proposed method generalize well to target domains. Experiments on real datasets across diverse domains, including healthcare, natural language processing, and computer vision, demonstrate the effectiveness of our proposed method.

Acknowledgements

This work was funded in part by the National Science Foundation under award number IIS-2145625 and by the National Institutes of Health under award number R01AI188576. We would like to express our sincere gratitude to Dr. Wei-Lun Chao from The Ohio State University for his valuable suggestions, which greatly improved this paper.

References

Albuquerque, I.; Monteiro, J.; Darvishi, M.; Falk, T. H.; and Mitliagkas, I. 2019. Generalizing to unseen domains via distribution matching. *arXiv preprint arXiv:1911.00804*.

Amini, M. R.; Usunier, N.; and Goutte, C. 2009. Learning from multiple partially observed views-an application to multilingual text categorization. *Advances in neural information processing systems*, 22.

Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine learning*, 79: 151–175.

Biau, G.; Cadre, B.; Sangnier, M.; and Tanielian, U. 2020. Some theoretical properties of GANS. *The Annals of Statistics*, 48(3): 1539 – 1566.

Bucci, S.; Loghmani, M. R.; and Tommasi, T. 2020. On the effectiveness of image rotation for open set domain adaptation. In *European conference on computer vision*, 422–438. Springer.

Chen, Y.; Zhu, X.; Li, W.; and Gong, S. 2020. Semi-supervised learning under class distribution mismatch. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 3569–3576.

Chua, T.-S.; Tang, J.; Hong, R.; Li, H.; Luo, Z.; and Zheng, Y. 2009. Nus-wide: a real-world web image database from national university of singapore. In *Proceedings of the ACM international conference on image and video retrieval*, 1–9.

Cortes, C.; and Mohri, M. 2014. Domain adaptation and sample bias correction theory and algorithm for regression. *Theoretical Computer Science*, 519: 103–126.

Demšar, J. 2006. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine learning research*, 7: 1–30.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.

Du, P.; Zhao, S.; Sheng, Z.; Li, C.; and Chen, H. 2023. Semi-Supervised Learning via Weight-aware Distillation under Class Distribution Mismatch. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16410–16420.

Fang, Z.; Lu, J.; Liu, F.; Xuan, J.; and Zhang, G. 2020. Open set domain adaptation: Theoretical bound and algorithm. *IEEE transactions on neural networks and learning systems*, 32(10): 4309–4322.

Fang, Z.; Lu, J.; Liu, F.; and Zhang, G. 2022. Semi-supervised heterogeneous domain adaptation: Theory and algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1): 1087–1105.

Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1): 2096–2030.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.

Grief, S. N.; Patel, J.; Kochendorfer, K. M.; Green, L. A.; Lussier, Y. A.; Li, J.; Burton, M.; and Boyd, A. D. 2016. Simulation of ICD-9 to ICD-10-CM transition for family medicine: simple or convoluted? *The Journal of the American Board of Family Medicine*, 29(1): 29–36.

Griffin, G.; Holub, A.; Perona, P.; et al. 2007. Caltech-256 object category dataset. Technical report, Technical Report 7694, California Institute of Technology Pasadena.

Gu, X.; Yang, Y.; Zeng, W.; Sun, J.; and Xu, Z. 2022. Keypoint-guided optimal transport with applications in heterogeneous domain adaptation. *Advances in Neural Information Processing Systems*, 35: 14972–14985.

Guo, L.-Z.; Zhang, Z.-Y.; Jiang, Y.; Li, Y.-F.; and Zhou, Z.-H. 2020. Safe deep semi-supervised learning for unseen-class unlabeled data. In *International conference on machine learning*, 3897–3906. PMLR.

He, R.; Han, Z.; Lu, X.; and Yin, Y. 2022. Safe-student for safe deep semi-supervised learning with unseen-class unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14585–14594.

Hoffman, J.; Rodner, E.; Donahue, J.; Darrell, T.; and Saenko, K. 2013. Efficient learning of domain-invariant image representations. *arXiv preprint arXiv:1301.3224*.

Hoffman, J.; Rodner, E.; Donahue, J.; Kulis, B.; and Saenko, K. 2014. Asymmetric and category invariant feature transformations for domain adaptation. *International journal of computer vision*, 109: 28–41.

Huang, J.; Fang, C.; Chen, W.; Chai, Z.; Wei, X.; Wei, P.; Lin, L.; and Li, G. 2021. Trash to treasure: Harvesting ood data with cross-modal matching for open-set semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8310–8319.

Huang, Z.; Yang, J.; and Gong, C. 2022. They are not completely useless: Towards recycling transferable unlabeled data for class-mismatched semi-supervised learning. *IEEE Transactions on Multimedia*, 25: 1844–1857.

Kiryo, R.; Niu, G.; Du Plessis, M. C.; and Sugiyama, M. 2017. Positive-unlabeled learning with non-negative risk estimator. *Advances in neural information processing systems*, 30.

Krizhevsky, A. 2009. Learning Multiple Layers of Features from Tiny Images. *Master's thesis, University of Tront.*

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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9757–9766.

Li, J.; Lu, K.; Huang, Z.; Zhu, L.; and Shen, H. T. 2018. Heterogeneous domain adaptation through progressive alignment. *IEEE transactions on neural networks and learning systems*, 30(5): 1381–1391.

Li, S.; Xie, B.; Wu, J.; Zhao, Y.; Liu, C. H.; and Ding, Z. 2020. Simultaneous semantic alignment network for heterogeneous domain adaptation. In *Proceedings of the 28th ACM international conference on multimedia*, 3866–3874.

Li, W.; Duan, L.; Xu, D.; and Tsang, I. W. 2013. Learning with augmented features for supervised and semi-supervised heterogeneous domain adaptation. *IEEE Transactions on Pattern analysis and machine intelligence*, 36(6): 1134–1148.

Liu, F.; Zhang, G.; and Lu, J. 2020. Heterogeneous domain adaptation: An unsupervised approach. *IEEE transactions on neural networks and learning systems*, 31(12): 5588–5602.

Liu, H.; Cao, Z.; Long, M.; Wang, J.; and Yang, Q. 2019. Separate to adapt: Open set domain adaptation via progressive separation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2927–2936.

Luo, Y.; Wang, Z.; Huang, Z.; and Baktashmotlagh, M. 2020. Progressive graph learning for open-set domain adaptation. In *International Conference on Machine Learning*, 6468–6478. PMLR.

Mansour, Y.; Mohri, M.; and Rostamizadeh, A. 2009. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*.

Natarajan, B. K. 1989. On learning sets and functions. *Machine Learning*, 4: 67–97.

Nguyen, A. T.; Tran, T.; Gal, Y.; Torr, P.; and Baydin, A. G. 2021. KL Guided Domain Adaptation. In *International Conference on Learning Representations*.

Pham, T.-H.; Zhang, X.; and Zhang, P. 2023. Fairness and Accuracy under Domain Generalization. In *The Eleventh International Conference on Learning Representations*.

Pham, T.-H.; Zhang, X.; and Zhang, P. 2024. Non-stationary Domain Generalization: Theory and Algorithm. In *The 40th Conference on Uncertainty in Artificial Intelligence*.

Quiñonero-Candela, J.; Sugiyama, M.; Schwaighofer, A.; and Lawrence, N. D. 2022. *Dataset shift in machine learning*. Mit Press.

Rasiwasia, N.; Costa Pereira, J.; Coviello, E.; Doyle, G.; Lanckriet, G. R.; Levy, R.; and Vasconcelos, N. 2010. A new approach to cross-modal multimedia retrieval. In *Proceedings of the 18th ACM international conference on Multimedia*, 251–260.

Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115: 211–252.

Saenko, K.; Kulis, B.; Fritz, M.; and Darrell, T. 2010. Adapting visual category models to new domains. In *Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part IV 11*, 213–226. Springer.

Saito, K.; Kim, D.; and Saenko, K. 2021. Openmatch: Open-set semi-supervised learning with open-set consistency regularization. *Advances in Neural Information Processing Systems*, 34: 25956–25967.

Saito, K.; Yamamoto, S.; Ushiku, Y.; and Harada, T. 2018. Open set domain adaptation by backpropagation. In *Proceedings of the European conference on computer vision (ECCV)*, 153–168.

Shen, C.; and Guo, Y. 2018. Unsupervised heterogeneous domain adaptation with sparse feature transformation. In *Asian conference on machine learning*, 375–390. PMLR.

Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).

Wagner, P.; Strodthoff, N.; Bousseljot, R.-D.; Kreiseler, D.; Lunze, F. I.; Samek, W.; and Schaeffter, T. 2020. PTB-XL, a large publicly available electrocardiography dataset. *Scientific data*, 7(1): 1–15.

Wang, Y.; Qiao, P.; Liu, C.; Song, G.; Zheng, X.; and Chen, J. 2023. Out-of-distributed semantic pruning for robust semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23849–23858.

Xu, R.; Liu, P.; Zhang, Y.; Cai, F.; Wang, J.; Liang, S.; Ying, H.; and Yin, J. 2020. Joint Partial Optimal Transport for Open Set Domain Adaptation. In *IJCAI*, 2540–2546.

Yao, Y.; Zhang, Y.; Li, X.; and Ye, Y. 2019. Heterogeneous domain adaptation via soft transfer network. In *Proceedings of the 27th ACM MM*, 1578–1586.

Yao, Y.; Zhang, Y.; Li, X.; and Ye, Y. 2020. Discriminative distribution alignment: A unified framework for heterogeneous domain adaptation. *Pattern Recognition*, 101: 107165.

Yu, Q.; Ikami, D.; Irie, G.; and Aizawa, K. 2020. Multi-task curriculum framework for open-set semi-supervised learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, 438–454. Springer.

Zeiberg, D.; Jain, S.; and Radivojac, P. 2020. Fast nonparametric estimation of class proportions in the positive-unlabeled classification setting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 6729–6736.

Zhao, H.; Des Combes, R. T.; Zhang, K.; and Gordon, G. 2019. On learning invariant representations for domain adaptation. In *International conference on machine learning*, 7523–7532. PMLR.

Zhao, Y.; Li, S.; Zhang, R.; Liu, C. H.; Cao, W.; Wang, X.; and Tian, S. 2022. Semantic correlation transfer for heterogeneous domain adaptation. *IEEE Transactions on Neural Networks and Learning Systems*.

Zou, Y.; Yu, Z.; Kumar, B.; and Wang, J. 2018. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, 289–305.