Solving Medical Multi-Label Domain Adaptation via Wasserstein Adversarial Learning with Class-Level Alignment

Wenjie Liu[1], Fuyou Miao[1(✉)], Xu Wang[1]

University of Science and Technology of China
lwj1217@mail.ustc.edu.cn, mfy@ustc.edu.cn, worm@mail.ustc.edu.cn

Abstract. In medical imaging, domain adaptation (DA) enables the transfer of knowledge from models trained on labeled source domains to unlabeled target domains that exhibit distribution shifts. In real world, medical images often contain multiple disease-related labels. However, existing multi-label domain adaptation (MLDA) algorithms face two primary challenges in addressing multi-label domain shifts: inadequate capture of disease features and insufficient integration of information from each individual class. To tackle these challenges, we propose a novel approach, Wasserstein Adversarial Learning with Class-Level Alignment, designed to align feature distributions for medical MLDA. By utilizing adversarial learning guided by Wasserstein distance, our approach captures more complete domain-invariant representations of lesion region. Additionally, we introduce a class-level alignment loss that leverages individual class information to further reduce domain discrepancies. Extensive experiments on real medical datasets demonstrate that our method significantly enhances medical multi-label domain adaptation and outperforms existing state-of-the-art algorithms.

Keywords: Medical imaging domain adaptation · Multi-label · Wasserstein adversarial learning · Class-level alignment.

1 Introduction

Nowadays, deep learning has demonstrated its remarkable performance in classification and detection for medical imaging, which attracts wide attention in research community [2,24,28]. However, medical images acquired from different hospitals and healthy centers often exhibit distribution shifts due to variations in scanning parameters set by devices [6]. As illustrated in Figure 1, distribution shifts can stem from differences in the orientation of the scans, for example, between anterior-posterior (AP), posterior-anterior (PA), and lateral scans, in chest X-rays [10]. In addition, the costs associated with annotating medical images and training multiple models for diverse datasets are relatively expensive [13]. These factors make it challenging for a well-trained deep learning model to achieve similarly high performance on other medical datasets with limited annotations [26]. To overcome this issue, the knowledge acquired from the annotated

source domain needs to be transferred to the target domain, which is known as **domain adaptation** (**DA**) [13].

Most deep DA methods mainly align the domain distributions through adversarial processes, which reflect features into a similar space across domains [4,11,15], or minimizing the discrepancy of feature distributions [13,16]. These DA methods are mainly applied to multi-class classification, where each data sample is associated with a single label. However, in real-world medical scenarios, a single radiographic image may express information of multiple diseases simultaneously, rather than just one disease. For example, in the X-ray image of a patient, if *Atelectasis* is diagnosed, there is a significant probability that *Effusion* and *Infiltration* are also present [25]. Furthermore, multi-label deep learning models can significantly assist physicians in reducing the probability of missed diagnoses by effectively capturing the co-occurrence patterns of multiple diseases [23]. Obviously, multi-class DA models mainly focus on assigning a single label among multiple labels to each sample, which limits their abilities to capture co-occurrence of interrelated labels and to couple dependencies between multi disease indicators. Thus, most existing multi-class DA algorithms are not well-suited to perform optimally in the multi-label domain shift for medical imaging [14].

For the aforementioned multi-label medical image classification, **multi-label domain adaptation** (**MLDA**) [12] algorithms are crucial for effectively transferring the knowledge of multiple labels across domains. This topic has gradually attracted the interest of researchers in fields such as medical imaging [14], remote sensing [9] and bearing compound fault diagnosis [3]. Similar to DA methods, many studies [12,14,20] still employ a domain discriminator to adversarially learn the projection of features in a shared space. Except for domain discriminator, adversarial critic with coupling classification and discrimination tasks [19], can also help mitigate the effects of domain shifts. Furthermore, reducing the discrepancy of feature distributions remains a common technique to obtain domain-invariant features in MLDA, which includes multiple kernel maximum mean discrepancy [3] and cosine similarity [9].

Existing MLDA methods rarely incorporate the information specific to medical image to refine alignment strategies. In addition, these algorithms primarily focus on aligning feature spaces, similar to multi-class DA, which treat each domain as a whole. Therefore, these approaches fail to fully exploit individual label information in the multi-label setting to perform more granular alignment.

To address the above challenges for MLDA in medical imaging, we propose a novel approach named as **Wasserstein Adversarial Learning** with **Class-Level Alignment** (**WAL-CLA**). This approach leverages the adversarial learning guided by the Wasserstein distance [21] to align the feature distributions. The Wasserstein distance helps deep learning models extract the domain-invariant representations more completely [27] and locate pathological organs more accurately [7] across medical imaging domains. It is worth noting that, unlike conventional DA, our Wasserstein adversarial learning framework is specifically designed for multi-label tasks. To further integrate multi-label information, we

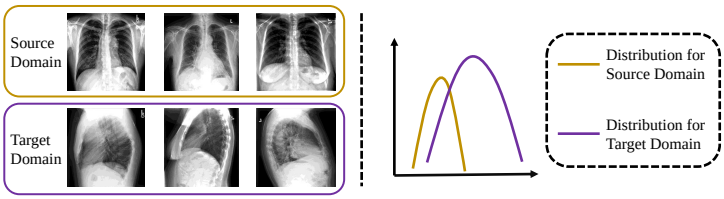

Fig. 1. Illustration of medical images of two domains according to their scan orientations. Feature distributions of two domains are shown in the middle part of the figure.

introduce a class-level alignment strategy that explicitly aligns feature distributions at the label level. This ensures a more precise alignment across domains for each class while preserving the intrinsic label dependencies in medical imaging.

Our contributions can be concluded as follows:

1. We introduce a Wasserstein adversarial learning process guided by the Wasserstein distance to enhance model's ability of capturing more intrinsic characteristics of medical imaging in MLDA.
2. We combine the class-level alignment loss that explicitly aligns feature distributions for each class to integrate individual label information better.
3. Extensive experiments conducted on real X-ray datasets validate the effectiveness of our proposed method, demonstrating its advantages over several state-of-the-art MLDA approaches.

2 Preliminaries

In this section, we briefly review the definitions of multi-label domain adaptation and Wasserstein distance.

Multi-label Domain Adaptation. In multi-label domain adaptation, we define dataset from the source domain as $D_S = \{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}_{i=1}^{n_s}$, where $\mathbf{x}_i^s \in \mathbb{R}^d$ is the i-th medical instance and $\mathbf{y}_i^s = \{y_i^1, \cdots, y_i^C\}$ is the multi-label of total C classes respecting to the instance (for any j-th class, $y_i^j \in \{0, 1\}$). Similarly, we can define the target domain $D_T = \{\mathbf{x}_i^t\}_{i=1}^{n_t}$, which consists of instances without annotated labels provided by doctors. To transfer effective knowledge from the source domain to the target domain, most existing MLDA models contain two main components: a feature extractor $\mathcal{E}$ and a classifier $\mathcal{C}$. The feature extractor $\mathcal{E}$ is responsible for extracting the image features, i.e., $\mathbf{e_i} = \mathcal{E}(\mathbf{x}_i)$, while the classifier $\mathcal{F}$ outputs C logits, $\mathbf{p}_i = \mathcal{F}(\mathbf{e}_i)$, to predict the multi-label output $\hat{\mathbf{y}}_i$ over C classes, where $\hat{\mathbf{y}}_i = \{\hat{y}_i^1, \cdots, \hat{y}_i^C\}$. By leveraging multi-label learning in the source domain and unlabeled information from the target domain, the model can effectively adapt to multi-label classification tasks in the target domain.

Wasserstein Distance. Wasserstein distance [21] is a widely used measure for quantifying the difference between two probability distributions by computing the cost required to transport one distribution to another. In medical imaging, this measure helps models capture a more comprehensive, domain-invariant representation and facilitates more precise localization of lesion region across multiple domains [7,27]. The definition of 1-Wasserstein distance between two domains can be formalized as follows:

$$W(P_S, P_T) = \inf_{\tau \in \Pi(P_S, P_T)} \int \|\mathbf{x}^s - \mathbf{x}^t\| d\tau(\mathbf{x}^s, \mathbf{x}^t), \tag{1}$$

where P_S and P_T denote probability distributions of two domains, and $\|\cdot\|$ denotes the 1-norm Euclidean distance. $\Pi(P_S, P_T)$ represents the set of their all joint distributions and $\tau(\mathbf{x}^s, \mathbf{x}^t)$ represents the "mass" that transports from $\mathbf{x}^s$ to $\mathbf{x}^t$. Since Equation (1) is inherently challenging to compute directly, we derive a more tractable formulation based on Kantorovich-Rubinstein duality [22]:

$$W(P_S, P_T) = \sup_{\|f\|_{Lip} \leq 1} \mathbb{E}_{\mathbf{x}^s \sim P_S}\big[f(\mathbf{x}^s)\big] - \mathbb{E}_{\mathbf{x}^t \sim P_T}\big[(f(\mathbf{x}^t))\big], \tag{2}$$

where the supremum result is computed under the assumption that the learning function f satisfies 1-Lipschitz smoothness condition.

3 Our method

In this section, we introduce our proposed Wasserstein adversarial learning with class-level alignment in detail.

Wasserstein Adversarial Learning. To further improve computational efficiency in aligning feature distributions, we relax the learning function f in Equation (2) to be K-Lipschitz, where the distance is scaled by a factor of K, leading to $K \cdot W(P_S, P_T)$ [1]. Following the principles of Wasserstein generative adversarial networks (WGAN) [1], we first define the adversarial learning loss function based on Wasserstein distance as follows:

$$\mathcal{L}_W = \mathbb{E}_{\mathbf{x}^s \sim P_S}\big[f(\mathbf{x}^s)\big] - \mathbb{E}_{\mathbf{x}^t \sim P_T}\big[(f(\mathbf{x}^t))\big]. \tag{3}$$

The above adversarial learning mainly focus on normal DA setting, which has been widely applied in [17,18,27]. In MLDA scenario, it is crucial to account for the multi-label nature of the task. Thus, we formalize the loss function for Wasserstein adversarial learning as follows:

$$\mathcal{L}_W = \frac{1}{C}\sum_{j=1}^{C}\Big[\frac{1}{n_s}\sum_{i=1}^{n_s}\mathcal{F}^j\big(\mathcal{E}(\mathbf{x}_i^s)\big) - \frac{1}{n_t}\sum_{i=1}^{n_t}\mathcal{F}^j\big(\mathcal{E}(\mathbf{x}_i^t)\big)\Big], \tag{4}$$

where $\mathcal{F}^j$ denotes the j-th logit value of the classifier output. We adapt the adversarial loss to operate directly from classifier rather than a separate discriminator in WGAN, which allows for a natural adaptation to the multi-label setting during the adversarial learning process. Furthermore, it can decrease the number of model parameters and save computational resources.

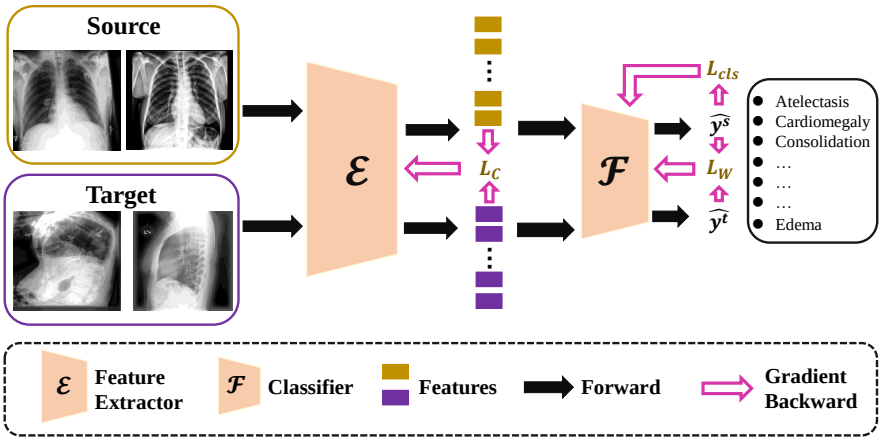

Fig. 2. The architecture of proposed network combining with Wasserstein adversarial learning and class-level alignment for medical multi-label domain adaptation. The black arrows indicate the forward propagation of data through the network for prediction, while the pink arrows represent the back-propagation of gradients. On the right part, the labels depict the set of multi-label annotations for the entire dataset.

Class-Level Alignment. Except for intrinsic characteristics of medical image captured by Wasserstein adversarial learning, it is also necessary to fully exploit the information provided by multi-label annotations. To address this, we propose a class-level alignment strategy incorporating multi-label information to further mitigate discrepancy across domains for each individual class. First, we assign pseudo labels $\hat{\mathbf{y}}_i^t = \{\mathbb{1}[\mathcal{F}^1\big(\mathcal{E}(\mathbf{x}_i^t)\big) > \gamma], \cdots, \mathbb{1}[\mathcal{F}^C\big(\mathcal{E}(\mathbf{x}_i^t)\big) > \gamma]\}$ to each target domain data, where $\mathcal{F}^j\big(\mathcal{E}(\mathbf{x}_i^t)\big)$ denotes the j-th logit value of classifier output and γ denotes the threshold for predicting pseudo label as 1. Then, combining the labels of the source domain, the loss function for class-level alignment is defined as follows:

$$\mathcal{L}_C = -\log\Big(\sum_{j=1}^{C}\exp\big(\frac{1}{T}\frac{\overline{\mathbf{x}}_j^s \cdot \overline{\mathbf{x}}_j^t}{\|\overline{\mathbf{x}}_j^s\|_2\|\overline{\mathbf{x}}_j^t\|_2}\big)\Big), \tag{5}$$

where T is temperature hyperparameter. Here, $\overline{\mathbf{x}}_j^s = \frac{1}{n_s}\sum_{i=1}^{n_s}\mathbb{1}[y_i^j = 1]\mathcal{E}(\mathbf{x}_i^s)$ and $\overline{\mathbf{x}}_j^t = \frac{1}{n_t}\sum_{i=1}^{n_t}\mathbb{1}[\hat{y}_i^j = 1]\mathcal{E}(\mathbf{x}_i^t)$. The class-level loss refines the model parameters of the feature extractor, ensuring a more direct alignment of the embedded feature distributions for each individual class.

Framework. Combining with the Wasserstein adversarial learning and class-level alignment, the general framework of our proposed WAL-CLA is shown in Figure 2. In this framework, X-ray images from frontal and lateral scans are treated as the source and target domains. These images are processed through

the feature extractor $\mathcal{E}$ and the classifier $\mathcal{F}$ to generate multi-label predictions. We employ ResNet34 [8] as the $\mathcal{E}$ and a two-layer linear network as the $\mathcal{F}$. To simplify the adversarial learning process, we apply a gradient reversal layer (GRL) [4] between network $\mathcal{E}$ and $\mathcal{F}$ to invert the gradient during back-propagation for $\mathcal{L}_W$. The total loss used for updating training parameters $\theta_\mathcal{E}$ and $\theta_\mathcal{F}$ of modules $\mathcal{E}$ and $\mathcal{F}$ are summarized as follows:

$$\begin{aligned} \theta_\mathcal{F}^* &= \arg\min_{\theta_\mathcal{F}} \ \mathcal{L}_{cls} - \lambda_W \mathcal{L}_W, \\ \theta_\mathcal{E}^* &= \arg\min_{\theta_\mathcal{E}} \ \mathcal{L}_{cls} + \lambda_W \mathcal{L}_W + \lambda_C \mathcal{L}_C. \end{aligned} \tag{6}$$

Here, $\mathcal{L}_{cls}$ represents the binary cross-entropy loss of the source domain samples, which is commonly used for multi-label classification tasks:

$$\mathcal{L}_{cls} = -\frac{1}{n_s}\sum_{i=1}^{n_s}\Big[\frac{1}{C}\sum_{j=1}^{C} y_i^j \cdot \log\Big(\mathcal{F}^j\big(\mathcal{E}(\mathbf{x}_i^s)\big)\Big) + (1-y_i^j)\cdot\log\Big(1-\mathcal{F}^j\big(\mathcal{E}(\mathbf{x}_i^s)\big)\Big)\Big]. \tag{7}$$

4 Experiments

Dataset. The proposed method is evaluated on two real X-ray datasets of chest, ChestX-ray14 [25] and CheXpert [10]. ChestX-ray14 contains 112,120 X-ray images from 30,805 unique patients with 14 disease labels, while CheXpert contains 224,316 chest radiographs from 65,240 patients with 14 categories of labels. Given the inconsistency in the majority of labels between these two datasets, we partition them based on their respective scanning parameters, defining corresponding source and target domains. Specifically, for ChestX-ray14, we separate the dataset into the posterior-anterior (PA) domain and anterior-posterior (AP) domain based on different scan orientations. Additionally, we define the Male (M) and Female (F) domains based on patient gender. With different views of the chest in contrast to ChestX-ray14, we split CheXpert dataset into frontal radiographs (Fron) and lateral radiographs (Lat) domain. Similarly, we also create Male (M) and Female (F) domains for CheXpert based on patient gender. This partitioning strategy allows us to examine the effects of domain shifts due to both scan orientations and demographic factors for both datasets.

Baselines. We compare our proposed WAL-CLA with several existing popular approaches. (1) Source: the fundamental multi-label classification ResNet34 model trained exclusively on the source domain. (2) DANN [4]: a multi-class DA approach which leverage adversarial learning to align distributions by discriminator. (3)GAN-C [14]: a generative adversarial network [5] based method combined with label predicted conditions. (4) MK-MMD [3]: it decreases the discrepancy of distributions by minimizing the measure of multiple kernel maximum mean discrepancy. (5) DDA-MLIC [19]: it utilizes adversarial critic to match two domains by Gaussian Mixture Model, instead of discriminator. (6) Target: the baseline model being same with the network architecture of "Source" and trained on target domain directly.

Implementation details. Our proposed WAL-CLA was implemented by Pytorch and trained on Nvidia 4090 GPU. We employed ResNet34 pretrained on ImageNet as the feature extractor backbone. We set batch size as 64, learning rate as 0.01 and trained the model with 20000 steps. For hyperparamters, we set $\lambda_W = 0.1$, $\lambda_C = 0.05$, $\tau = 0.8$ and $T = 0.5$. For MLDA, we evaluated models performance on target domain by 3 metrics: the predicted multiple label accuracy (ACC), overall average area under the ROC curve (oAUC), and class average area under the ROC curve (cAUC). Higher values of these metrics indicate better model performance.

Results. Table 1 demonstrates the compared experimental results for ChestX-ray14 dataset. We designate AP and PA as one domain pair for mutual adaptation, while constitute M and F as another domain pair for mutual adaptation. For example, PA→AP represents that source domain is PA and target domain is AP. In all domain adaptation cases, it is evident that our proposed method outperforms other baselines across all evaluation metrics. Similarly, Table 2 shows the evaluation results on the CheXpert dataset, where we also employ 4 domain adaptation cases. The reason of the method working less well on the Fron→Lat and the Lat→Fron is mainly due to the severe data imbalance between two domains when they are split by view, where the Fron domain contains over 190k samples and the Lat domain includes only about 30k samples. This large imbalance poses a significant challenge to the Wasserstein adversarial alignment, which relies on sufficient and balanced data distributions for effective alignment. In contrast, the other domain adaptation settings in our experiments involve source and target domains with more balanced data volumes, allowing for more stable and reliable feature alignment and leading to better adaptation performance. With the exception of the Fron→Lat, our algorithm consistently achieves superior performance across almost all other domain adaptation settings, surpassing existing methods.

In Tables 1 and 2, ResNet34 trained on the source domain ("Source") performs significantly worse when transferred directly to the target domain without any adjustment, compared to the ResNet34 trained directly on target domain ("Target"). This finding further confirms that the adopted domain partitioning strategy induces domain shift in the dataset, thereby providing a valid testbed for multi-label domain adaptation.

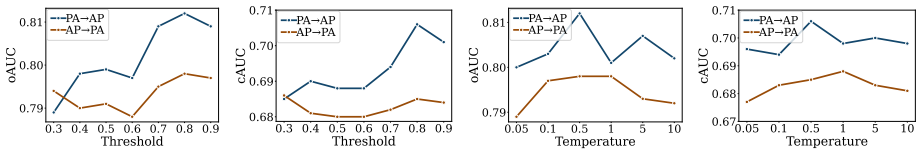

Fig. 3. The oAUC and cAUC results on ChestX-ray14 dataset with the threshold γ varying from 0.3 to 0.9 and the temperature T varying from 0.05 to 10, respectively.

Table 1. Experimental results for ChestX-ray14 dataset, where we evaluate on 4 domain adaptation cases. Except for "Target", the best results of rests are underlined.

	PA→AP			AP→PA			M→ F			F→ M		
	ACC	oAUC	cAUC	ACC	oAUC	cAUC	ACC	oAUC	cAUC	ACC	oAUC	cAUC
Source	86.55	0.743	0.645	86.86	0.716	0.643	89.35	0.822	0.717	89.29	0.826	0.726
DANN	88.24	0.782	0.681	88.68	0.781	0.673	89.61	0.836	0.754	89.11	0.828	0.741
GAN-C	88.35	0.795	0.690	88.66	0.780	0.673	89.47	0.834	0.754	89.15	0.826	0.738
MK-MMD	87.92	0.772	0.620	88.77	0.765	0.631	88.99	0.804	0.696	88.88	0.808	0.695
DDA-MLIC	80.16	0.771	0.702	78.12	0.775	0.683	79.83	0.830	0.761	81.26	0.822	0.751
WAL-CLA	88.68	0.812	0.706	89.47	0.798	0.685	89.72	0.838	0.761	89.50	0.838	0.752
Target	90.87	0.909	0.866	93.57	0.948	0.931	90.96	0.890	0.842	90.53	0.880	0.816

Table 2. Experimental results for CheXpert dataset, where we evaluate on 4 domain adaptation cases. Except for "Target", the best results of rests are underlined.

	Fron→Lat			Lat→Fron			M→ F			F→ M		
	ACC	oAUC	cAUC	ACC	oAUC	cAUC	ACC	oAUC	cAUC	ACC	oAUC	cAUC
Source	87.46	0.736	0.615	82.17	0.683	0.551	87.68	0.881	0.749	87.73	0.875	0.743
DANN	88.21	0.771	0.649	84.14	0.762	0.627	87.98	0.883	0.751	87.67	0.877	0.748
GAN-C	88.10	0.768	0.658	83.26	0.775	0.634	88.01	0.884	0.754	87.75	0.878	0.750
MK-MMD	86.84	0.744	0.605	83.54	0.804	0.601	86.43	0.858	0.698	86.47	0.853	0.695
DDA-MLIC	62.59	0.691	0.592	72.38	0.687	0.566	72.31	0.879	0.774	68.70	0.814	0.691
WAL-CLA	88.19	0.755	0.630	84.27	0.815	0.628	88.32	0.890	0.769	88.07	0.884	0.766
Target	93.83	0.954	0.935	89.15	0.916	0.830	89.69	0.916	0.832	89.03	0.912	0.837

Ablation Studies. We conduct ablation studies to analyze the contribution of each proposed loss term in the PA→AP and AP→ PA settings on the ChestX-ray14 dataset. As shown in Table 3, each individual loss component contributes to a performance improvement, further validating the effectiveness and rationality of each proposed loss term. When all loss items are combined, the model achieves the best overall performance.

Effectiveness of Hyperparameters. Figure 3 illustrates the effect of varying threshold values γ and temperature T on the oACU and cAUC performance metrics for WAL-CLA on the ChestX-ray14 dataset. It is evident that the performance of our WAL-CLA achieves its near-optimal levels when the threshold is set as 0.8 and the temperature is approximately 0.5.

5 Conclusion

In this paper, we propose a novel Wasserstein adversarial learning with class-level alignment for medical multi-label domain adaptation. On real-world medical imaging datasets, our proposed method demonstrates superior performance compared to existing algorithms. In the future, we aim to extend our approach to multi-modal medical imaging by more deeply integrating modality-specific

Table 3. Effectiveness of each individual component of WAL-CLA on ChestX-ray14.

$\mathcal{L}_{\mathcal{W}}$	$\mathcal{L}_{\mathcal{C}}$	PA→AP			AP→PA		
		ACC	oAUC	cAUC	ACC	oAUC	cAUC
		86.55	0.743	0.645	86.86	0.716	0.643
✓		88.46	0.803	0.697	89.09	0.788	0.679
	✓	88.36	0.803	0.693	89.22	0.794	0.683
✓	✓	88.68	0.812	0.706	89.47	0.798	0.685

characteristics. This will enhance multi-label domain adaptation in medical applications, ultimately aiding clinicians in reducing misdiagnosis rates.

Acknowledgments. The research of this work was supported in part by Innovation Program for Quantum Science and Technology (2021ZD0302902), Hi-tech project(231-08-01) and Anhui Province University Natural Science Research Project (2023AH051102).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: International conference on machine learning. pp. 214–223. PMLR (2017)
2. Chan, H.P., Samala, R.K., Hadjiiski, L.M., Zhou, C.: Deep learning in medical image analysis. Deep learning in medical image analysis: challenges and applications pp. 3–21 (2020)
3. Chu, L., Li, Q., Yang, B., Chen, L., Shen, C., Wang, D.: Exploring the essence of compound fault diagnosis: A novel multi-label domain adaptation method and its application to bearings. Heliyon **9**(3) (2023)
4. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: International conference on machine learning. pp. 1180–1189. PMLR (2015)
5. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems **27** (2014)
6. Guan, H., Liu, M.: Domain adaptation for medical image analysis: a survey. IEEE Transactions on Biomedical Engineering **69**(3), 1173–1185 (2021)
7. He, B., Chen, Y., Zhu, D., Xu, Z.: Domain adaptation via wasserstein distance and discrepancy metric for chest x-ray image classification. Scientific Reports **14**(1), 2690 (2024)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
9. Huang, R., Ma, M., Huang, W.: Domain adaptation for multi-label remote sensing image annotation with contrastive pseudo-label generation. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing (2024)
10. Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., et al.: Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 590–597 (2019)

11. Li, J., Li, G., Shi, Y., Yu, Y.: Cross-domain adaptive clustering for semi-supervised domain adaptation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2505–2514 (2021)
12. Lin, D., Lin, J., Zhao, L., Wang, Z.J., Chen, Z.: Multilabel aerial image classification with unsupervised domain adaptation. IEEE Transactions on Geoscience and Remote Sensing **60**, 1–13 (2021)
13. Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: International conference on machine learning. pp. 97–105. PMLR (2015)
14. Pham, D.D., Koesnadi, S., Dovletov, G., Pauli, J.: Unsupervised adversarial domain adaptation for multi-label classification of chest x-ray. In: 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI). pp. 1236–1240. IEEE (2021)
15. Saito, K., Kim, D., Sclaroff, S., Darrell, T., Saenko, K.: Semi-supervised domain adaptation via minimax entropy. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 8050–8058 (2019)
16. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3723–3732 (2018)
17. She, Q., Chen, T., Fang, F., Zhang, J., Gao, Y., Zhang, Y.: Improved domain adaptation network based on wasserstein distance for motor imagery eeg classification. IEEE Transactions on Neural Systems and Rehabilitation Engineering **31**, 1137–1148 (2023)
18. Shen, J., Qu, Y., Zhang, W., Yu, Y.: Wasserstein distance guided representation learning for domain adaptation. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
19. Singh, I.P., Ghorbel, E., Kacem, A., Rathinam, A., Aouada, D.: Discriminator-free unsupervised domain adaptation for multi-label image classification. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3936–3945 (2024)
20. Singh, I.P., Ghorbel, E., Oyedotun, O., Aouada, D.: Multi-label image classification using adaptive graph convolutional networks: from a single domain to multiple domains. Computer Vision and Image Understanding **247**, 104062 (2024)
21. Villani, C.: Topics in optimal transportation, vol. 58. American Mathematical Soc. (2021)
22. Villani, C., et al.: Optimal transport: old and new, vol. 338. Springer (2008)
23. Wang, J., Yang, L., Huo, Z., He, W., Luo, J.: Multi-label classification of fundus images with efficientnet. IEEE access **8**, 212499–212508 (2020)
24. Wang, R., Lei, T., Cui, R., Zhang, B., Meng, H., Nandi, A.K.: Medical image segmentation using deep learning: A survey. IET image processing **16**(5), 1243–1267 (2022)
25. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2097–2106 (2017)
26. Wang, Z., Ye, M., Zhu, X., Peng, L., Tian, L., Zhu, Y.: Metateacher: Coordinating multi-model domain adaptation for medical image classification. Advances in Neural Information Processing Systems **35**, 20823–20837 (2022)
27. You, C., Yang, J., Chapiro, J., Duncan, J.S.: Unsupervised wasserstein distance guided domain adaptation for 3d multi-domain liver segmentation. In: Interpretable and Annotation-Efficient Learning for Medical Image Computing: Third

International Workshop, IMIMIC 2020, Second International Workshop, MIL3ID 2020, and 5th International Workshop, LABELS 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4–8, 2020, Proceedings 3. pp. 155–163. Springer (2020)
28. Zhou, S.K., Greenspan, H., Shen, D.: Deep learning for medical image analysis. Academic Press (2023)