Neighborhood-Consistent Binary Transformation for Domain-Invariant Chest X-ray Diagnosis

Zelong Liu*, Huachao Zhu*, Zhichao Sun, Yuda Zou, Yuliang Gu, Bo Du(✉), and Yongchao Xu(✉)

School of Computer Science, Wuhan University, Wuhan, China
{dubo,yongchao.xu}@whu.edu.cn

Abstract. Chest X-ray diagnosis models face domain generalization challenges due to cross-institutional variations in imaging protocols and scanner specifications, which degrade diagnostic accuracy on unseen domains. To address this, we propose a domain-invariant learning framework leveraging the inherent anatomical consistency of medical imaging. Our method first applies a Neighborhood-Consistent Binarization Transformation (NCBT) to convert grayscale images into topology-preserving high-dimensional binary tensors, encoding pixel intensity relationships within local neighborhoods to strip device-specific textures while retaining anatomical structures. These tensors are then reconstructed into an intermediate domain via an Intermediate Domain Style-preserving Autoencoder (IDSP-AE), decoupling structural information from domain-specific features. Crucially, our framework aligns domains without requiring target domain data during training, leveraging anatomical consistency. Experiments on four public datasets show superior generalization and improved diagnostic accuracy compared to state-of-the-art methods. The source code is available at `https://github.com/LZL501/NCBT`.

Keywords: Chest X-ray · Domain Generalization · Relative Intensity Relationship.

1 Introduction

Chest X-ray remains the most widely used clinical imaging technique for pulmonary disease detection [21,16,26] thanks to its efficiency, low cost, and accessibility. Recent advances in deep learning technology have enabled automated disease diagnosis from chest X-ray images with promising results. Yet, deep learning models trained on data from one institution often suffer from performance degradation when applied to data from new institutions, primarily due to domain shifts caused by differences in imaging devices and protocols across hospitals. This challenge significantly limits the real-world deployment of these models. To address this issue, previous research has developed two primary approaches: Domain Adaptation (DA) [3,5,19] and Domain Generalization (DG) [31,17,23].

* equal contribution

Domain adaptation(DA) aims to adapt models trained on a source domain to a target domain with unknown annotations. Techniques like Generative Adversarial Networks (GANs) [4] and their variants, such as CycleGAN [30], are commonly used to either learn domain-invariant features or translate image styles between source and target domains. However, GAN-based DA methods [13,25] rely on access to target domain data during training, which means they primarily handle known target domains, limiting their flexibility. Additionally, during the image synthesis process, GAN-based DA methods may alter pathological features [15], potentially leading to misdiagnosis. On the other hand, domain generalization seeks to improve a model's ability to generalize to unseen domains during training, without requiring target-domain data. Recent approaches [28,20] have explored style manipulation and neural encoding to overcome domain shifts, demonstrating improvements in robustness. However, DG methods still face challenges in addressing significant distribution differences across diverse medical imaging data sources. First, aggressive style alteration may corrupt diagnostically critical anatomical patterns. Furthermore, existing medical DG approaches focus excessively on global statistical alignment, overlooking the fundamental local anatomical relationships that remain consistent across imaging devices and play a pivotal role in enhancing generalization capabilities [14].

Based on the consistency of relative intensity relationships between neighboring pixels in chest X-rays, we propose a novel transformation called Neighborhood-Consistent Binary Transformation (NCBT) to alleviate domain-specific style information. Unlike conventional style normalization techniques that rely on heuristic assumptions, NCBT operates by exhaustively encoding each pixel's intensity relative to all neighbors within a sliding window, generating high-dimensional binary tensors where domain-sensitive intensity distributions are discarded while anatomical topology is preserved. This transformation achieves an unification: multi-institutional chest X-rays with heterogeneous styles (*e.g.*, varying contrast/brightness) are mapped to a single, domain-agnostic representation space, inherently mitigating domain shifts without requiring adversarial training or target domain data.

To harness NCBT's domain-invariant representations while retaining clinically interpretable visuals, we further design the Intermediate Domain Style-Preserving Autoencoder (IDSP-AE). Trained solely on a style-neutral intermediate domain, this network reconstructs NCBT's binary tensors into diagnostically actionable images that exhibit consistent intermediate-domain styling. Crucially, when deployed on both source and unseen target domains, the IDSP-AE acts as a universal translator: it reprojects all inputs into this unified style space, effectively neutralizing device-specific biases while preserving pathology-critical intensity contrasts. Experimental results on four public chest X-ray datasets for pulmonary disease diagnosis demonstrate that our method significantly improves baseline performance, validating its effectiveness.

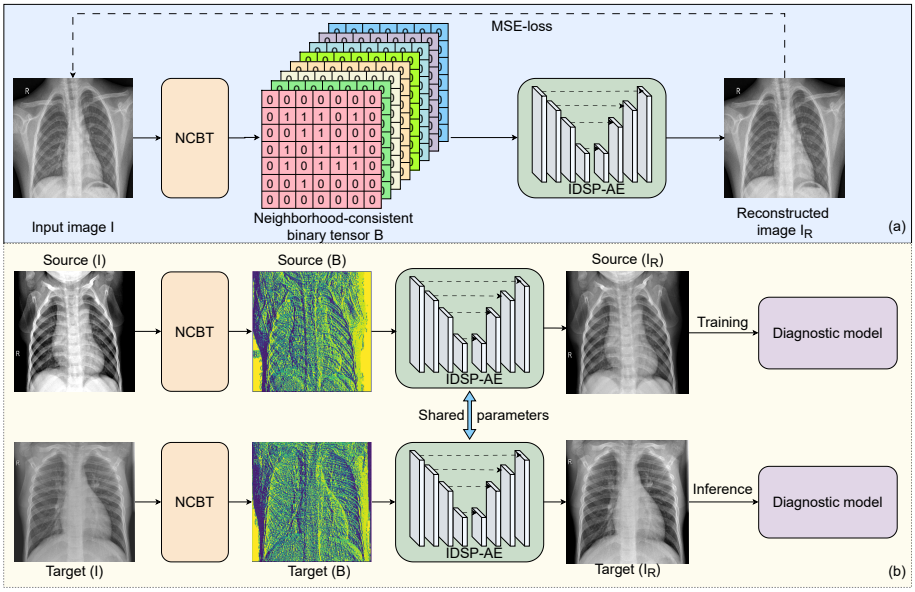

Fig. 1. (a) Training process of the proposed IDSP-AE. (b) Diagram of our method combined with other diagnostic models. In the diagram, Source (**I**), Source (**B**), and Source ($\mathbf{I_R}$) represent the original source domain image, the multi-channel tensor obtained via NCBT, and the reconstructed image with intermediate domain style, respectively. Target (**I**), Target (**B**), and Target ($\mathbf{I_R}$) follow the same structure. Note that in diagram (b), **B** represents the visualization of the multi-channel binary tensor after converting it to decimal and normalizing, which is conceptually identical to **B** in (a).

2 Method

2.1 Overview

Inherent differences in imaging devices lead to significant stylistic variations among chest X-ray datasets from different medical centers, causing pulmonary disease diagnosis models that perform well on source domains to often fail generalizing to unseen target domains. To address this, we propose the Neighborhood-Consistent Binary Transformation (NCBT) – a novel transformation that alleviates domain-specific style information from both source and target domains. The processed multi-channel tensors generated by NCBT are then reconstructed into intermediate-domain-style single-channel images using an Intermediate Domain Style-Preserving Autoencoder (IDSP-AE). This unified representation effectively mitigates cross-domain discrepancies. As shown in Figure 1(b), our framework serves as a plug-and-play preprocessing module for standard diagnostic workflows: both training and test sets undergo NCBT and IDSP-AE reconstruction, producing style-harmonized images aligned with the intermediate domain. Subsequent diagnosis models trained/tested on these normalized images

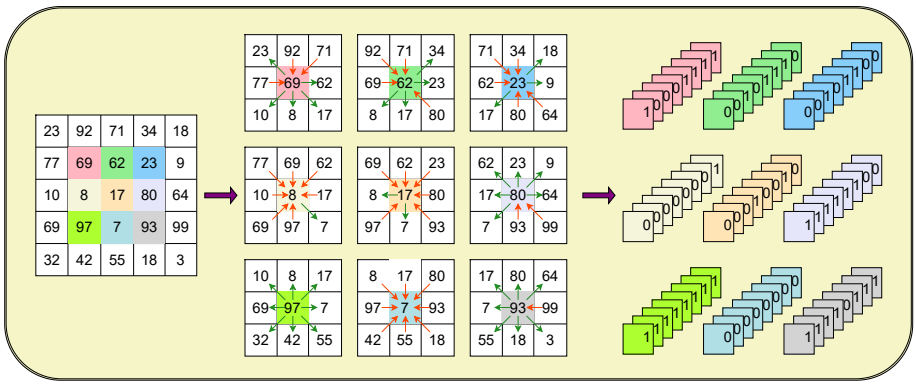

Fig. 2. Illustration of the Proposed Neighborhood-Consistent Binary Transformation. A 3×3 patch (colored pixels in the image) from a chest X-ray image is used as an example, with the neighborhood size $N = 3$. During NCBT, each pixel in the patch undergoes intensity comparisons with its eight neighboring pixels within the sliding 3×3 window (shown as nine 3×3 regions centered on colored pixels). Binary values (0/1) are assigned to new channels based on these comparisons: red arrows indicate higher intensity in neighboring pixels compared to the center pixel, while green arrows denote the reverse.

exhibit enhanced generalization capabilities without architectural modifications. The following sections detail each component of our approach.

2.2 Neighborhood-Consistent Binary Transformation

Although chest X-ray datasets from different medical centers exhibit stylistic variations due to disparities in imaging protocols, the relative intensity relationships between neighboring pixels maintain domain-invariant consistency for identical anatomical structures. Leveraging this property, we propose a novel robust representation transformation named Neighborhood-Consistent Binary Transformation (NCBT). This approach converts input grayscale images into multi-channel binary tensors encoding neighborhood ordinal relationships.

Specifically, as illustrated in Figure 2, given an input grayscale image $\mathbf{I} \in \mathbb{R}^{H \times W}$, NCBT compares each pixel with the $N \times N$ neighboring pixels, excluding the center pixel, to evaluate the relative intensity. If the intensity of the center pixel is greater than that of a neighboring pixel, the binary value 1 is assigned; otherwise, 0 is assigned. This binary value is then assigned to the corresponding channel in the transformed multi-channel binary tensor. Each spatial position in the neighborhood corresponds to a unique channel in the transformed binary tensor. Thus, NCBT transforms the input grayscale image $\mathbf{I} \in \mathbb{R}^{H \times W}$ into a multi-channel binary tensor $\mathbf{B} \in \{0, 1\}^{H \times W \times (N^2 - 1)}$ through local ordinal comparisons. For each pixel (i, j), the transformation is formally defined as:

$$\mathbf{B}(i,j,c_{(k,l)}) = \begin{cases} 1 & \text{if } \mathbf{I}(i,j) > \mathbf{I}(k,l) \\ 0 & \text{otherwise} \end{cases}, \tag{1}$$

where:

- $\mathcal{N}(i,j) = \{(k,l) \mid (k,l) \in \mathcal{W}_N(i,j) \backslash \{(i,j)\}\}$ denotes the $N \times N$ neighborhood window centered at (i,j), excluding the center pixel.
- $c_{(k,l)} \in \{1, ..., N^2 - 1\}$ maps each neighbor position to a unique channel index.
- $\mathcal{W}_N(i,j)$ represents the $N \times N$ sliding window:

$$\mathcal{W}_N(i,j) = \left\{ (i+m, j+n) \;\middle|\; m,n \in \left[-\left\lfloor \frac{N}{2} \right\rfloor, \left\lceil \frac{N}{2} \right\rceil \right] \right\}. \tag{2}$$

As formalized in Equation (1), the proposed NCBT is independent of absolute pixel intensity values, relying solely on relative intensity relationships between each pixel and its neighbors. By discarding domain-sensitive stylistic attributes (*e.g.*, global contrast, brightness), this design effectively mitigates inter-domain discrepancies among input images while preserving diagnostically critical ordinal patterns.

2.3 Intermediate Domain Style-Preserving Autoencoder

Building upon NCBT's generation of domain-agnostic binary tensors $\mathbf{B}$ described in Section 2.2, we introduce the Intermediate Domain Style-Preserving Autoencoder (IDSP-AE) to reconstruct these tensors into style-normalized images $\mathbf{I_R} \in \mathbb{R}^{H \times W}$. As depicted in Figure 1(a), IDSP-AE employs a U-Net [12] backbone with skip connections to preserve fine anatomical details. During training, the network learns to decode $\mathbf{B}$ into images that conform to to a style-neutral intermediate domain by minimizing pixel-wise reconstruction error:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{HW} \sum_{i=1}^{H} \sum_{j=1}^{W} \left(I_R(i,j) - I(i,j) \right)^2, \tag{3}$$

where H and W respectively represent the height and width of $\mathbf{I}$.

A critical design principle lies in our intermediate domain construction: the IDSP-AE is trained on a dedicated, large-scale chest X-ray dataset disjoint from both source and target domains used in downstream diagnostic tasks. This strategic isolation ensures the intermediate domain serves as a stylistic 'neutral ground' devoid of domain-specific biases. Through training, IDSP-AE learns to reconstruct NCBT-generated binary tensors $\mathbf{B}$ into images $\mathbf{I_R}$ that impose intermediate-domain styling – effectively acting as a universal style filter.

During deployment, both source domain training data and unseen target domain test data first undergo NCBT transformation to remove device-dependent styles, then pass through the frozen IDSP-AE to synthesize images exhibiting

unified intermediate-domain features. This two-step process bridges the domain gap by ensuring all input images share identical stylistic characteristics while preserving NCBT-encoded anatomical accuracy. Consequently, downstream diagnostic models encounter only anatomy-centered representations, alleviating the need for domain-invariant feature learning and significantly boosting cross-domain generalization.

3 Experiments

3.1 Datasets and Implementation Details

Datasets. We employ the ChestX-ray8 dataset [22] (108,948 chest X-ray images) as the training set for the IDSP-AE, constituting the intermediate domain described in Section 2.2. For baseline comparison, we adopt CheXzero [18]– a state-of-the-art vision-language model for chest X-ray diagnosis – following its official training protocol. The MIMIC-CXR dataset [7] (377,110 image-text pairs) serves as the source domain for training the vision-language diagnostic model. To evaluate cross-domain generalization, we utilize four publicly available benchmark datasets for testing: CheXpert [6], ZhangLab [8], VinDr-CXR [10], and RSNA Pneumonia Detection Challenge [1], which are referred to as the unseen target domains.

Implementation Details. We choose the previously well-performing visual language model CheXzero [18] for chest X-ray disease diagnosis as the baseline. Our framework is integrated as a plug-and-play preprocessing module into CheXzero without modifying its architecture. To ensure fair comparison, all DG methods (including our augmented CheXzero) strictly follow CheXzero's original training protocol. Additionally, the proposed IDSP-AE is a simple U-Net [12] network, trained for 30 epochs with a batch size of 128. The Adam optimizer is used to optimize the network parameters, with a learning rate set to 1e-4. The neighborhood window size N for the NCBT transformation is set to 15. All experiments are conducted using four GeForce RTX 3090 GPUs.

3.2 Main Results

Table 1 and Table 2 present the quantitative evaluation results on four public datasets. It is important to note that the Zhanglab [8] and RSNA [1] datasets in Table 1 are binary classification datasets (containing only pneumonia and normal classes), while the CheXpert [6] and VinDr-CXR [10] datasets in Table 2 contain multiple disease categories. On the Zhanglab dataset, compared to the baseline model CheXzero, our method improves AUC, ACC, and F1 by 7.5%, 7.1%, and 5.0%, respectively. On the RSNA dataset, the improvements are 2.4%, 0.7%, and 0.5%, which are higher than other **plug-and-play methods** [28,24,2,27,9,11] designed to enhance network generalization.

For multi-class datasets in Table 2, previous methods like MixStyle [28] degrade performance because their implicit style mixing, which is effective for

Table 1. Quantitative Evaluation Results on the Zhanglab [8] and RSNA [1] Datasets. The '+' indicates that the baseline model CheXzero is combined with the plug-and-play method following the '+'. The highest result in each column is highlighted in bold.

Methods	Zhanglab Dataset [8]			RSNA Dataset [1]		
	AUC(%) ↑	ACC(%)↑	F1(%)↑	AUC(%)↑	ACC(%)↑	F1(%)↑
CheXzero [18]	87.5	80.6	84.9	94.8	85.9	87.3
+pAdaIN [11]	94.0	85.4	85.4	95.5	84.5	86.0
+EFDMix [24]	94.0	86.1	88.9	96.5	86.4	87.5
+DSU [9]	94.5	86.7	89.6	96.3	86.0	86.9
+TriD [2]	94.2	86.8	87.7	96.9	84.7	86.4
+MixStyle [29]	92.9	84.5	86.5	96.5	85.9	**87.8**
+Morestyle [27]	92.4	85.3	86.7	95.4	84.9	86.8
+Ours	**95.0**	**87.7**	**89.9**	**97.2**	**86.6**	**87.8**

Table 2. Quantitative Evaluation Results on the CheXpert [6] and VinDr-CXR [10] Datasets. The highest result in each column is highlighted in bold.

Methods	CheXpert Dataset [6]			VinDr-CXR Dataset [10]		
	AUC(%) ↑	ACC(%)↑	F1(%)↑	AUC(%)↑	ACC(%)↑	F1(%)↑
CheXzero [18]	88.2	77.7	51.0	65.4	72.1	22.0
+pAdaIN [11]	85.0	71.4	44.9	69.2	74.5	22.6
+EFDMix [24]	84.0	72.3	50.3	65.8	74.9	21.9
+DSU [9]	85.3	73.6	45.5	70.5	66.0	19.1
+TriD [2]	88.5	77.8	49.6	67.0	74.5	20.6
+MixStyle [29]	84.1	**78.5**	47.5	67.4	74.0	22.1
+Morestyle [27]	86.2	71.6	50.9	68.3	73.7	22.0
+Ours	**88.8**	78.3	**54.1**	**71.4**	**75.6**	**22.7**

natural images, blurs medically critical boundaries in chest X-rays. Disease features (*e.g.*, localized opacities, effusions) are inherently tied to intensity patterns, which style randomization inadvertently distorts. In contrast, our NCBT-based approach alleviates style variability while preserving pathological semantics, enabling consistent multi-class diagnosis across domains.

As demonstrated in our experiments, compared to other plug-and-play DG methods, our approach substantially reduces domain shifts across chest X-ray datasets by unifying the style space between source domains and unseen target domains. This leads to robust improvements and state-of-the-art performance on all four downstream datasets, confirming the effectiveness and strong generalization capability of our method.

3.3 Ablation Study

Ablation Study on the Neighborhood Size N in NCBT. We conduct an ablation study on the neighborhood size N involved in NCBT (*i.e.*, N in Equation (2)) using the Zhanglab dataset. As Table 3 demonstrates, enlarging N broadens the receptive field for ordinal intensity comparisons, enhancing local

Table 3. Ablation study on Neighborhood Size N in NCBT on Zhanglab dataset [8].

Neighborhood Size N	AUC(%) ↑	ACC(%) ↑	F1(%) ↑
3	89.2	81.4	84.9
7	90.0	83.4	86.7
11	92.1	85.5	88.4
15	**95.0**	**87.7**	**89.9**

Table 4. Ablation study on the proposed NCBT on the Zhanglab dataset [8].

NCBT	IDSP-AE	AUC(%) ↑	ACC(%) ↑	F1(%) ↑
-	-	87.5	80.6	84.9
-	✓	88.0	81.1	85.0
✓	✓	**95.0**	**87.7**	**89.9**

structural encoding in $\mathbf{B}$ and improving diagnostic metrics. However, increasing N also increases the number of parameters in the IDSP-AE and results in greater computational cost. When N exceeds 15, both the training cost and inference speed of IDSP-AE become prohibitively expensive. Thus, we select $N = 15$ as the optimal trade-off between performance and efficiency for practical deployment.

The Necessity of NCBT. We validate the necessity of NCBT through an ablation study. We train the IDSP-AE using original images from ChestX-ray8 [22] to reconstruct the input images, then fed raw images from the source and target domains into the trained IDSP-AE for intermediate-domain style transfer. The results are shown in the second row of Table 4. Compared to our proposed full pipeline (NCBT + IDSP-AE), using IDSP-AE alone yielded suboptimal results, achieving only marginal improvements over the baseline. This gap arises because raw images retain scanner-specific textures (*e.g.*, high-frequency noise, contrast variations) that IDSP-AE cannot fully normalize. NCBT's neighborhood ordinal encoding fundamentally removes these domain-specific features, making it indispensable for effective style unification.

4 Conclusion

To address domain shifts in chest X-ray diagnosis caused by cross-institutional imaging heterogeneity, we propose a physiology-driven framework that leverages anatomical consistency across domains. Our Neighborhood-Consistent Binarization Transformation (NCBT) encodes images into domain-invariant binary tensors through localized ordinal intensity comparisons, while an Intermediate Domain style-preserving autoencoder reconstructs intermediate-domain-aligned images. This approach alleviates device-specific stylistic variations without requiring target domain data. Experiments on four public datasets demonstrate superior generalization performance over state-of-the-art methods, validating its effectiveness as a plug-and-play solution for clinically deployable AI diagnostics. The modular design allows integration with existing diagnostic systems without

requiring extra data effort. This work addresses the gap between AI models and clinical deployment by disentangling clinically relevant features from domain-specific features. In the future, we would like to replace the fixed center-based ordinal encoding in current NCBT with other more robust pairwise settings, further enhancing robustness against local noise artifacts while preserving structural semantics.

Acknowledgments. This work was supported partially by the National Key Research and Development Program of China (2023YFC2705700), NSFC 62222112, 62176186, and 62225113, the NSF of Hubei Province of China (2024AFB245), the Special Fund for Central Guidance on Local Science and Technology Development from the Sichuan Provincial Department of Science and Technology (2024ZYD0285), the Key Research and Development Program of Dazhou Science and Technology Bureau (24ZDYF0005).

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

References

1. rsna challenge. www.kaggle.com/c/rsna-pneumonia- detection-challenge/overview (2018)
2. Chen, Z., Pan, Y., Ye, Y., Cui, H., Xia, Y.: Treasure in distribution: A domain randomization based multi-source domain generalization for 2d medical image segmentation. In: Proc. of Intl. Conf. on Medical Image Computing and Computer Assisted Intervention. pp. 89–99 (2023)
3. Fang, Y., Wu, J., Wang, Q., Qiu, S., Bozoki, A., Liu, M.: Source-free collaborative domain adaptation via multi-perspective feature enrichment for functional MRI analysis. Pattern Recognition **157**, 110912 (2025)
4. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Proc. of Advances in Neural Information Processing Systems **27** (2014)
5. Hou, Q., Wang, Y., Cao, P., Cheng, S., Lan, L., Yang, J., Liu, X., Zaiane, O.R.: A collaborative self-supervised domain adaptation for low-quality medical image enhancement. IEEE Trans. on Medical Imaging (2024)
6. 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: Proc. of the AAAI Conf. on Artificial Intelligence. vol. 33, pp. 590–597 (2019)
7. Johnson, A.E., Pollard, T.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Peng, Y., Lu, Z., Mark, R.G., Berkowitz, S.J., Horng, S.: Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. arXiv preprint arXiv:1901.07042 (2019)
8. Kermany, D.S., Goldbaum, M., Cai, W., Valentim, C.C., Liang, H., Baxter, S.L., McKeown, A., Yang, G., Wu, X., Yan, F., et al.: Identifying medical diagnoses and treatable diseases by image-based deep learning. Cell **172**(5), 1122–1131 (2018)
9. Li, X., Dai, Y., Ge, Y., Liu, J., Shan, Y., DUAN, L.: Uncertainty modeling for out-of-distribution generalization. In: Proc. of International Conference on Learning Representations (2022)

10. Nguyen, H.Q., Lam, K., Le, L.T., Pham, H.H., Tran, D.Q., Nguyen, D.B., Le, D.D., Pham, C.M., Tong, H.T., Dinh, D.H., et al.: Vindr-cxr: An open dataset of chest x-rays with radiologist's annotations. Scientific Data **9**(1), 429 (2022)
11. Nuriel, O., Benaim, S., Wolf, L.: Permuted adain: Reducing the bias towards global statistics in image classification. In: Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition. pp. 9482–9491 (2021)
12. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Proc. of Intl. Conf. on Medical Image Computing and Computer Assisted Intervention. pp. 234–241 (2015)
13. Sanchez, K., Hinojosa, C., Arguello, H., Kouamé, D., Meyrignac, O., Basarab, A.: Cx-dagan: Domain adaptation for pneumonia diagnosis on a small chest x-ray dataset. IEEE Trans. on Medical Imaging **41**(11), 3278–3288 (2022)
14. Shi, T., Boutry, N., Xu, Y., Géraud, T.: Local intensity order transformation for robust curvilinear object segmentation. IEEE Trans. on Image Processing **31**, 2557–2569 (2022)
15. Shin, H.C., Tenenholtz, N.A., Rogers, J.K.: Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In: SASHIMI 2018 (2018)
16. Sun, Z., Gu, Y., Liu, Y., Zhang, Z., Zhao, Z., Xu, Y.: Position-guided prompt learning for anomaly detection in chest x-rays. In: Proc. of Intl. Conf. on Medical Image Computing and Computer Assisted Intervention. pp. 567–577 (2024)
17. Tian, Y., Wen, C., Shi, M., Afzal, M.M., Huang, H., Khan, M.O., Luo, Y., Fang, Y., Wang, M.: Fairdomain: Achieving fairness in cross-domain medical image segmentation and classification. In: Proc. of European Conference on Computer Vision. pp. 251–271 (2024)
18. Tiu, E., Talius, E., Patel, P., Langlotz, C.P., Ng, A.Y., Rajpurkar, P.: Expert-level detection of pathologies from unannotated chest x-ray images via self-supervised learning. Nature Biomedical Engineering **6**(12), 1399–1406 (2022)
19. Wang, H., Chen, J., Zhang, S., He, Y., Xu, J., Wu, M., He, J., Liao, W., Luo, X.: Dual-reference source-free active domain adaptation for nasopharyngeal carcinoma tumor segmentation across multiple hospitals. IEEE Trans. on Medical Imaging (2024)
20. Wang, H., Xia, Y.: Domain-ensemble learning with cross-domain mixup for thoracic disease classification in unseen domains. Biomedical Signal Processing and Control **81**, 104488 (2023)
21. Wang, X., Li, Y., Wu, W., Jin, J., Rong, Y., Jiang, B., Li, C., Tang, J.: Pre-training on high-resolution x-ray images: an experimental study. Visual Intelligence **3**(1), 1–15 (2025)
22. 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: Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition. pp. 2097–2106 (2017)
23. Yan, S., Yu, Z., Liu, C., Ju, L., Mahapatra, D., Betz-Stablein, B., Mar, V., Janda, M., Soyer, P., Ge, Z.: Prompt-driven latent domain generalization for medical image classification. IEEE Trans. on Medical Imaging (2024)
24. Zhang, Y., Li, M., Li, R., Jia, K., Zhang, L.: Exact feature distribution matching for arbitrary style transfer and domain generalization. In: Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition. pp. 8035–8045 (2022)
25. Zhang, Y., Miao, S., Mansi, T., Liao, R.: Task driven generative modeling for unsupervised domain adaptation: Application to x-ray image segmentation. In: Proc.

of Intl. Conf. on Medical Image Computing and Computer Assisted Intervention. pp. 599–607 (2018)
26. Zhang, Z., Sun, Z., Liu, Z., Zhao, Z., Yu, R., Du, B., Xu, Y.: Spatial-aware attention generative adversarial network for semi-supervised anomaly detection in medical image. In: Proc. of Intl. Conf. on Medical Image Computing and Computer Assisted Intervention. pp. 638–648 (2024)
27. Zhao, H., Dong, W., Yu, R., Zhao, Z., Du, B., Xu, Y.: Morestyle: relax low-frequency constraint of fourier-based image reconstruction in generalizable medical image segmentation. In: Proc. of Intl. Conf. on Medical Image Computing and Computer Assisted Intervention. pp. 434–444 (2024)
28. Zhou, K., Yang, Y., Qiao, Y., Xiang, T.: Domain generalization with mixstyle. In: Proc. of International Conference on Learning Representations (2021)
29. Zhou, K., Yang, Y., Qiao, Y., Xiang, T.: Mixstyle neural networks for domain generalization and adaptation. International Journal of Computer Vision **132**(3), 822–836 (2024)
30. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proc. of IEEE Intl. Conf. on Computer Vision. pp. 2223–2232 (2017)
31. Zunaed, M., Haque, M.A., Hasan, T.: Learning to generalize towards unseen domains via a content-aware style invariant model for disease detection from chest x-rays. IEEE Journal of Biomedical and Health Informatics (2024)