

RESEARCH

Open Access



BarlowTwins-CXR: enhancing chest X-ray abnormality localization in heterogeneous data with cross-domain self-supervised learning

Haoyue Sheng^{1,2,3*}, Linrui Ma^{1,2}, Jean-François Samson³ and Dianbo Liu^{2,4}

Abstract

Background Chest X-ray imaging based abnormality localization, essential in diagnosing various diseases, faces significant clinical challenges due to complex interpretations and the growing workload of radiologists. While recent advances in deep learning offer promising solutions, there is still a critical issue of domain inconsistency in cross-domain transfer learning, which hampers the efficiency and accuracy of diagnostic processes.

This study aims to address the domain inconsistency problem and improve autonomic abnormality localization performance of heterogeneous chest X-ray image analysis, particularly in detecting abnormalities, by developing a self-supervised learning strategy called “BarlowTwins-CXR”.

Methods We utilized two publicly available datasets: the NIH Chest X-ray Dataset and the VinDr-CXR. The BarlowTwins-CXR approach was conducted in a two-stage training process. Initially, self-supervised pre-training was performed using an adjusted Barlow Twins algorithm on the NIH dataset with a Resnet50 backbone pre-trained on ImageNet. This was followed by supervised fine-tuning on the VinDr-CXR dataset using Faster R-CNN with Feature Pyramid Network (FPN). The study employed mean Average Precision (mAP) at an Intersection over Union (IoU) of 50% and Area Under the Curve (AUC) for performance evaluation.

Results Our experiments showed a significant improvement in model performance with BarlowTwins-CXR. The approach achieved a 3% increase in mAP50 accuracy compared to traditional ImageNet pre-trained models. In addition, the Ablation CAM method revealed enhanced precision in localizing chest abnormalities. The study involved 112,120 images from the NIH dataset and 18,000 images from the VinDr-CXR dataset, indicating robust training and testing samples.

Conclusion BarlowTwins-CXR significantly enhances the efficiency and accuracy of chest X-ray image-based abnormality localization, outperforming traditional transfer learning methods and effectively overcoming domain inconsistency in cross-domain scenarios. Our experiment results demonstrate the potential of using self-supervised learning to improve the generalizability of models in medical settings with limited amounts of heterogeneous data. This approach can be instrumental in aiding radiologists, particularly in high-workload environments, offering a promising direction for future AI-driven healthcare solutions.

Keywords Medical image analysis, Chest x-ray, Abnormality localization, Deep learning, Object detection, Self-supervised learning, Transfer learning, Heat map, Area under curve, Mean Average Precision

*Correspondence:

Haoyue Sheng
haoyue.sheng@umontreal.ca

Full list of author information is available at the end of the article



© The Author(s) 2024. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

Introduction

Chest X-ray(CXR) is a fundamental and widespread medical diagnostic tool for diagnosing chest diseases. It is efficient and cost-effective, suitable for preliminary screening and diagnosis [1]. During the 2019 coronavirus pandemic, CXR was widely used for triaging patients and prioritizing the care order due to its convenience and flexibility. Effective mitigation addresses the lack of availability of computed tomography and reduces the risk of transmission in the room with the CT scanner [2]. However, its complex interpretation often requires a highly qualified radiologist to make an accurate diagnosis [1]. As the demand for healthcare increases, the workload of radiologists has significantly increased [3]. It results in less time to analyze each radiographic image, potentially increasing the risk of diagnostic error. In many areas, especially in developing and remote areas, qualified radiologists are insufficient to cope with the increased demand for healthcare. For instance, Europe has 13 radiologists per 100,000 people, while the United Kingdom has 8.5, and Malaysia has approximately 30 per million population [4]. This situation necessitates urgently developing and introducing automated technologies like AI-based image analysis tools to aid radiologists in quicker and more precise CXR image analysis. It will improve the quality of diagnosis and help reduce the workload of doctors.

In recent years, deep learning models have rapidly advanced in various medical image analysis fields of CXR, demonstrating diagnostic accuracy comparable to human experts [5]. Object detection plays a more critical role in medical image analysis because it can identify and precisely locate the types of anomalies in the images, providing doctors with more specific and valuable information. However, training these models requires a large amount of annotated data. These annotations must be performed by experienced radiologists for CXR images, as well as for most medical images, making such annotated data not only costly, but also rare, with only a very limited number of public datasets including bounding box information. Although transfer learning is widely regarded as an effective method to solve the problem of scarce labelling data, its application in medical image analysis still faces limitations. This is mainly due to the significant difference in feature distribution between large datasets(such as ImageNet) used for pre-training models and medical imaging datasets. This cross-domain transfer learning disparity that directly applying these pre-trained weights to medical image analysis might not yield the best outcomes, particularly for specialized medical diagnostic applications [6, 7].

In order to develop an efficient chest x-ray image analysis method while addressing the problem of data scarcity

and domain inconsistency in transfer learning, our study proposed the BarlowTwins-CXR method, and its main contributions are as follows:

1. Proposed a new self-supervised two-phase training strategy to diagnose and locate abnormalities in CXR. Bringing self-supervised learning into chest X-ray anomaly localization solves the problem of cross-domain transfer learning differences.
2. In the first learning phase, our approach leverages self-supervised pre-training with the Barlow Twins algorithm [8], applied to CXR images without annotation. This strategy addresses the challenge of data heterogeneity between pre-trained ImageNet [9] models and CXR images. In the second phase, transfer learning on the VinDr-CXR [10] dataset is applied to fine-tune the model.
3. In the experiment employing ResNet50 [11] as the backbone architecture, we observed that implementing the BarlowTwins-CXR strategy significantly improved model performance. We observed a 3% increase in model accuracy on the mean Average Precision benchmark, surpassing the results achieved by directly performing transfer learning from ImageNet pre-trained weights.
4. Our experiment results demonstrate the potential of self-supervised learning in medical image analysis, especially in the absence of annotated data, paving the way for more efficient and accurate diagnostic aids methods in the future.

The paper is organized as follows: First, the related work section presents an overview of the research related to deep learning, self-supervised learning and transfer learning in the field of medical images. The methods section describes the dual-phase training process with dataset setup and preprocessing. The results section shows and compares the results obtained with different image sizes and backbone weights. We also used Heatmap visualization and Linear Evaluation Protocol to highlight the scheme's effectiveness in this section. Finally, the conclusions of this study and suggestions for future research are discussed in the discussion section.

Related work

In recent years, deep learning techniques have excelled in the field of medical imaging, particularly in analyzing CXR images. For example, in terms of disease classification, ChexNet proposed by Pranav Rajpurkar et al. [12] outperformed radiologists in detecting chest diseases, when benchmarked against the F1 score. Neural network models trained with vast amounts of labelled data are capable of identifying features of various pulmonary

diseases. In anomaly detection tasks, Sun K X et al. used the YOLOv7 object detection framework to effectively identify and locate lesions in CXR images [13]. This achievement is attributed to the advanced image recognition and feature extraction capabilities of neural networks. Additionally, the modified U-net architecture which incorporates attention mechanisms, as proposed by Gusztáv Gaál et al. [14], has made significant strides in accurately segmenting lung structures, thus aiding in detailed analysis and diagnosis of diseases.

Self-supervised learning has recently gained popularity in the field of medical imaging [15] and provides an efficient method for utilizing unlabeled data. Initially proposed by Bengio et al., this approach allows models to learn from unlabeled data and extract useful feature representations by training deep networks on unsupervised data [16]. Such learning strategy promotes models to capture the intrinsic structure and relationships in data by designing innovative pretext tasks, such as image reconstruction (e.g., Context encoder [17]), contrastive learning (e.g., SimCLR [18]), or prediction tasks (e.g., rotation prediction [19]). In the field of medical imaging, Shekoofeh Azizi et al. used large-scale images for self-supervised learning to improve accuracy and convergence speed significantly in downstream tasks, achieving better performance than models pre-trained on ImageNet [20]. Sowrirajan H et al. proposed a pre-trained model based on Momentum Contrast to enhance the representativeness and portability of CXR models [21].

In terms of transfer learning, applying models trained in one domain to another has led to notable success in medical image analysis. Research indicates that well-processed transfer results from ImageNet can improve

model performance in the medical imaging domain [6]. However, studies by Christos Matsoukas et al. have shown that due to the significant difference in feature distribution between medical and natural images, features learned from natural images may not always be broadly applicable to medical images [22]. Various cross-domain adaptive transfer learning methods have been developed to address these challenges, such as unsupervised and semi-supervised learning and sequential domain adaptation techniques. By tuning model parameters, these methods can be better adapted to the characteristics of medical images, improving the performance and accuracy of models in medical image analysis [6].

Methods

Dataset selection

This study utilized two publicly available chest X-ray datasets: the NIH-CXR [23] dataset and the VinDr-CXR dataset. The NIH dataset comprises 112,120 posterior-anterior (PA) or anterior-posterior (AP) CXR images from 30,805 patients, covering 14 diseases with image-level annotations, including disease location annotations in some images. The distribution of the NIH-CXR dataset is illustrated in Fig. 1.

Meanwhile, the VinDr-CXR dataset is the largest publicly available dataset for adult CXR object detection, which includes 18,000 PA CXR scans. These scans encompass 14 diseases with detailed instance-level bounding box annotations, making it ideal for the fine-tuning phase.

The VinDr-CXR dataset exhibits a distinct labelling process for its test and training sets. The training set, consisting of 15,000 images, was annotated independently

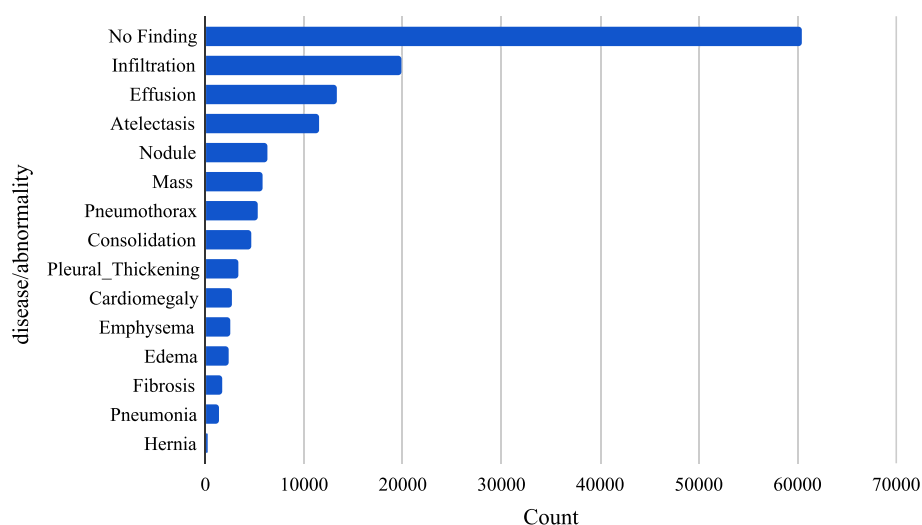


Fig. 1 Image-level label distribution of the NIH-CXR dataset

by three radiologists per image. In contrast, the test set, comprising 3,000 images, underwent a more rigorous annotation process. Initially, each image was independently annotated by three radiologists. This is followed by a secondary review phase where these initial annotations are reviewed by two other more experienced radiologists, they communicated with each other to resolve any disagreements and reach a consensus on the final labelling. This meticulous process for the test set created a potential disparity in data distribution compared to the training set. To eliminate any bias it might introduce in our study, we resplit the original training set into new training, validation, and test sets for our experiments.

To improve the quality of the training data, a Weighted Box Fusion (WBF) [24] preprocessing technique was applied to the VinDr-CXR training set. The WBF involves calculating the weighted average of each set of duplicate

bounding boxes to create a single fused bounding box. Such a preprocessing step is crucial for reducing annotation redundancy and improving target area representation in the dataset. Figure 2 shows the data distribution of VinDr-CXR before and after WBF preprocessing.

We chose the VinDr-CXR dataset not only because it is the largest publicly available dataset for adult CXR object detection, but also because of the high level of diversity and richness of its data.

Dual-phase training process

Our training encompasses two primary phases: self-supervised pre-training and subsequent supervised fine-tuning, as shown in Fig. 3. Initially, we commenced with a Resnet50 model pre-trained on ImageNet.

In the self-supervised pre-training phase, we applied the Barlow Twins method to the NIH-CXR Dataset

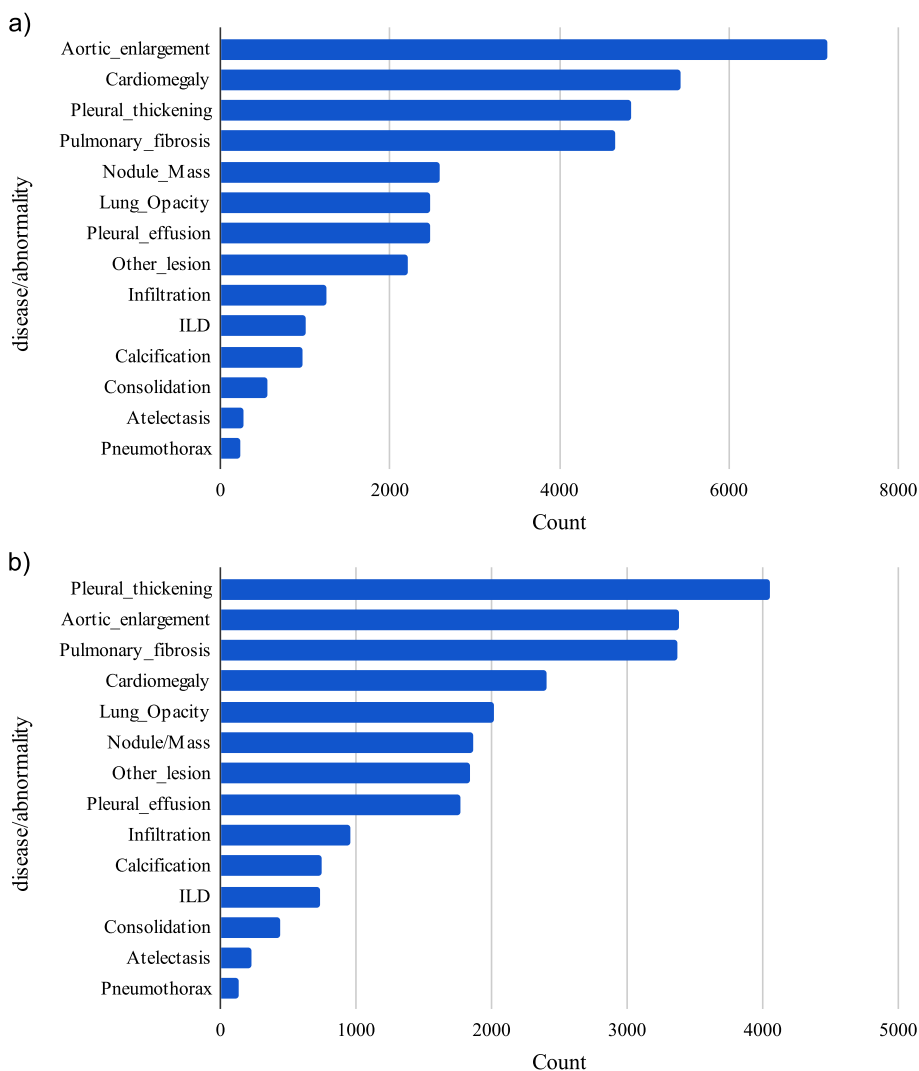


Fig. 2 Instance-level annotation distribution of VinDr-CXR dataset before (a) and after (b) WBF preprocessing

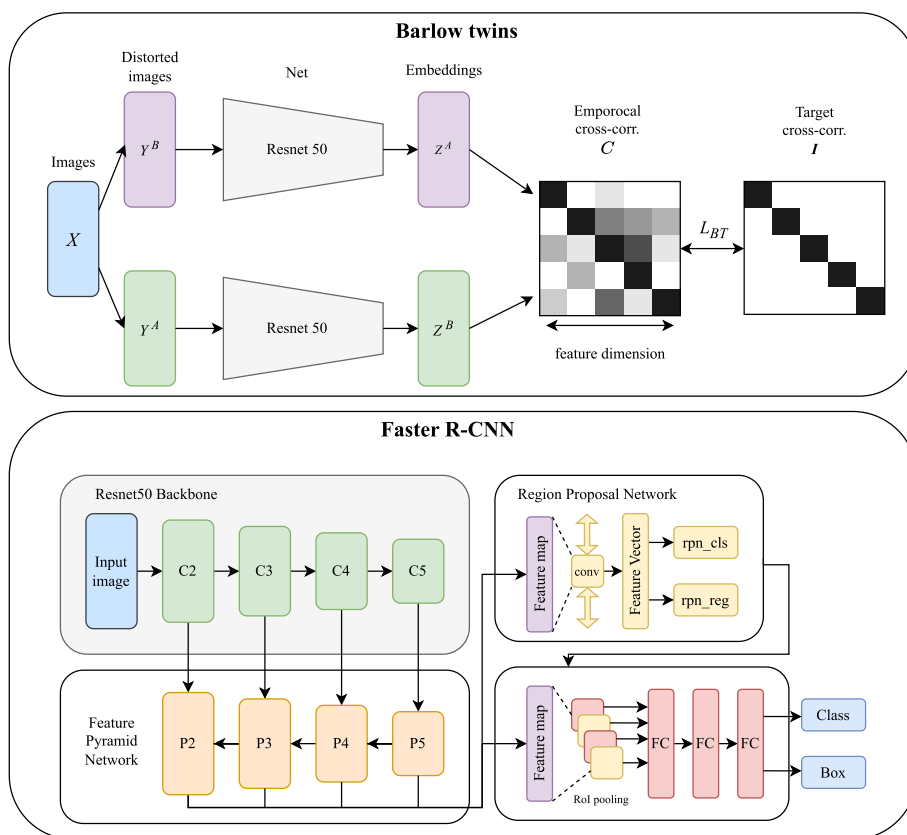


Fig. 3 Schematic Overview of the Dual-phase Training Framework. The upper panel illustrates the Barlow Twins method in Phase One, where pairs of distorted images are processed through a shared ResNet50 network to produce embeddings. These are then compared using an empirical cross-correlation matrix C , striving for the identity matrix I to minimize redundancy in feature dimensions, and optimizing the loss function L_{BT} . In Phase Two (lower panel), the pre-trained ResNet50 backbone from Phase One is integrated into a Faster R-CNN architecture. It starts with multi-scale feature extraction through the Feature Pyramid Network (FPN), followed by the Region Proposal Network (RPN) that generates object region proposals. The features are then pooled and processed by fully connected (FC) layers to output the final class labels and bounding box coordinates for object detection tasks

to exploit its ability to enhance feature independence, thereby significantly reducing feature redundancy in medical imaging. This approach refined the ImageNet pre-trained model by updating its backbone weights.

Subsequently, in the supervised fine-tuning phase, we utilize this refined backbone within a Faster R-CNN framework, chosen for its efficacy in precise localization and classification within complex images, to the VinDr-CXR dataset. This step aims to improve the model’s task-specific performance further, explicitly enhancing its capabilities in localized anomalies in CXR images.

Self-supervised pre-training

We chose the Barlow Twins method for the first stage of our training to ensure that visual features in medical images remain independent from each other, a crucial consideration for improving the performance of medical imaging analysis. As outlined by Zbontar et al. [8], this approach represents a shift from conventional contrastive

learning, introducing a self-supervised learning framework primarily focused on diminishing redundancy. The Barlow Twins method operates on a straightforward yet potent principle: it learns distinctive features by reducing the representational differences between two differently distorted images from the same source as processed by the network. This strategy is instrumental in enabling the model to identify unique and rich features in each image while concurrently minimizing the overlap in features. The process involves generating two distinct variants of an image through data augmentation, followed by their simultaneous processing via two deep neural networks that share identical weights. The objective is to align the network’s weights to enhance the similarity in the high-level representations of these image pairs yet ensure that the individual features remain distinct and independent.

The Barlow Twins method might be particularly useful for medical imaging because it extracts features by minimizing the redundancy between representations of

perturbed images. In CXR imaging, subtle differences might indicate important health information, and the Barlow Twins can effectively capture these subtle but clinically important features. In contrast to other contrastive learning algorithms like MoCo [25] and SimCLR, which construct similarity matrices at the batch dimension, Barlow Twins works at the feature dimension. It aims to assign an independent meaning to each feature dimension. This could lead to a richer feature representation, potentially better adapted to variations in CXR images (e.g., different imaging conditions and pathological states). Moreover, compared to self-supervised learning methods requiring negative samples or complex contrastive mechanisms like SimCLR, Barlow Twins offers a more straightforward training framework, which is particularly important in situations with limited computational resources.

We chose to apply Barlow Twins pre-training on the ImageNet pretrained ResNet50 model. Since the ImageNet pre-trained model weights can be easily obtained from the Torchvision library, this step brings no additional cost. We used images from the training set portion of the NIH-CXR dataset for this training phase, with the input image size set to 224*224 pixels. The training was executed on an NVIDIA A100 80G GPU, setting the batch size to 768 to maximize the utilization of this graphics card's capabilities over 600 epochs.

Fine-tuning phase

In our fine-tuning/transfer learning stage, we utilized the Faster R-CNN [26] with Feature Pyramid Network (FPN) [27] as our object detector and trained it on the VinDr-CXR dataset. Faster R-CNN, a widely-used object detection framework, comprises two main components: the Region Proposal Network (RPN) [27] and the Fast R-CNN detector. First, RPN generates candidate regions for objects, and then the Fast R-CNN detector employs these regions to detect and classify targets. This architecture renders Faster R-CNN particularly efficient in processing complex images. The Feature Pyramid Network (FPN), an architecture frequently employed in object detection, particularly enhances performance with multi-scale targets. It integrates high-level semantic information from deeper layers with detailed information from shallower layers, producing feature maps of varied scales that effectively detect differently sized targets.

We employed the MMDetection [28] machine learning toolbox as the platform for Faster R-CNN, utilizing a number of classical image augmentation techniques and maintaining consistent hyperparameters across all experiments. Two different input sizes, 224*224 pixels and 640*640 pixels, were chosen to assess the impact of image size on the model's performance with the pre-trained

models. In addition, for comparison, we also conducted experiments using ImageNet pre-trained weights directly.

We implemented a linear evaluation protocol [29, 30] on the NIH-CXR dataset to comprehensively evaluate the self-supervised learning model's performance in medical imaging. This method examines the model's feature transfer capability - its ability to adapt learned representations to new tasks. We first resplit the test set of the NIH dataset into two parts: 80% as an evaluation training set for training a linear classifier and the remaining 20% as an evaluation test set for assessing model performance.

We adopted two distinct strategies during the evaluation: freezing the backbone weights or fine-tuning the weights. In the freezing backbone strategy, we kept the parameters of the backbone network (i.e., the feature extraction layers) obtained from self-supervised pre-training unchanged. We updated only the weights of the final linear layer. Conversely, under the fine-tuning strategy, we updated parameters across the entire network, encompassing both the self-supervised trained feature extraction layers and the newly added linear classifier layer. We used 100%, 10%, and 1% of the evaluation training set data for training the linear classifier, allowing us to assess the model's performance across different scales of training data.

When evaluating the representation transfer ability of a self-supervised learning model, it is necessary to ensure that the ratio of individual labels in the training and test sets is consistent. We used the Iterative stratification for the multi-label data method [31, 32] to ensure that the proportions of each label in the evaluation training and test sets were roughly similar. This helped prevent biases due to uneven label distribution, making our evaluation results more reliable and convincing.

Results analysis process

For the analysis of results, we employed the mean Average Precision (mAP) at an Intersection over Union (IoU) of 50% as the benchmark for evaluating the performance of our object detection models. mAP is a widely recognized and effective metric in object detection, calculated by averaging precision scores across various object detection confidence thresholds. Specifically, mAP is the mean of the average precision scores for each class. The proportion of correct predictions relative to all predictions for a specific class across different detection confidence thresholds determines the precision score. In the context of CXR abnormality localization, utilizing mAP at an IoU of 50% is beneficial for capturing clinically significant lesion detections while allowing for a reasonable degree of positional deviation, which is practical for actual clinical applications.

Moreover, we utilized the Area Under the Curve (AUC) as a metric for the linear evaluation protocol. AUC, a standard metric in medical image analysis, balances precision and recall, making it an especially appropriate performance indicator for this field. The AUC metric represents the area under the Receiver Operating Characteristic (ROC) curve, accounting for the model's True Positive Rate (TPR) and False Positive Rate (FPR) at various thresholds. This assessment method balances the model's sensitivity and specificity, enhancing detection rates while controlling false positives. Medical image analysis often deals with imbalanced data, and AUC is robust for imbalanced datasets as it does not rely directly on classification thresholds.

Beyond using mAP and AUC for quantitative analysis, our study also utilized the Ablation CAM (Class Activation Mapping) method to create heat maps for qualitative evaluation. Ablation CAM systematically abates features in the model's final convolutional layer and observes the impact on the output class scores. This process reveals

the most influential regions for the model's decision-making. The resulting heat maps delineate areas of interest in CXR images, providing intuitive visual evidence of how our BarlowTwins-CXR model focuses on and recognizes abnormalities.

Results

Transfer learning on VinDr abnormality localization

In this experiment, we examined the efficacy of the ResNet backbone pre-trained by the Barlow Twins-CXR method for abnormality localization on the VinDr-CXR dataset, using two different input resolutions. Consistent hyperparameter settings were maintained across all experiments, ensuring that the performance changes were attributable only to the merits of the pretraining method itself. We visualized the performance of different models such as Barlow twins-CXR pre-training and ImageNet pre-training on the validation set in Fig. 4, and tabulated the corresponding mAP performance in Table 1. As depicted in the figure, the baseline model with an

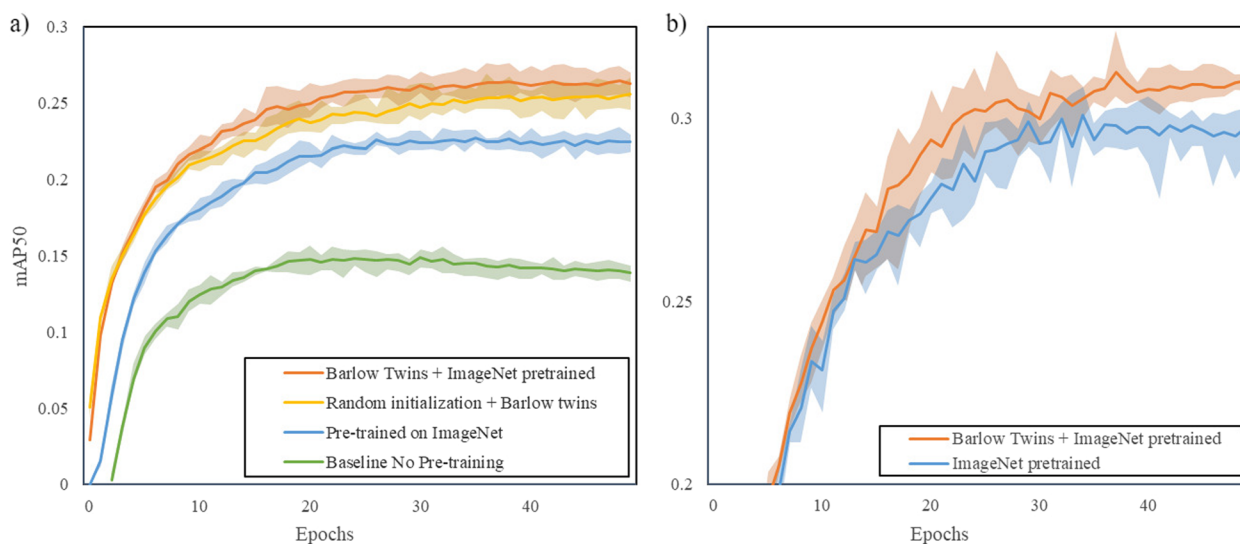


Fig. 4 Evolution of mAP50 across epochs for different ResNet50 backbones on the VinDr-CXR dataset at 224*224(left) and 640*640(right) resolution. The darker lines represent the average mAP50 of four(left) and five(right) trials with different random seeds, with shaded areas indicating the range between the lowest and highest value

Table 1 mAP50 scores in validation and test sets for models with varying pre-training methods at different input resolutions

BackBone_weight	Input_size	mAP50 (val set)	mAP50 (test set)
Baseline No Pre-training	224	0.1388 (0.1337, 0.1439)	0.1342 (0.1291, 0.1393)
Pre-trained on ImageNet		0.2245 (0.2187, 0.2303)	0.2210 (0.2187, 0.2233)
Random initialization + Barlow twins		0.2555 (0.2456, 0.2655)	0.2448 (0.2400, 0.2496)
Barlow Twins + ImageNet pretrained		0.2625 (0.2544, 0.2706)	0.2502 (0.2465, 0.2539)
ImageNet pretrained	640	0.2973 (0.2876, 0.3061)	0.2800 (0.2729, 0.2876)
Barlow Twins + ImageNet pretrained		0.3102 (0.3076, 0.3139)	0.2890 (0.2787, 0.2993)

Scores are presented with 95% confidence intervals

untrained ResNet50 backbone reached a final mAP50 score of 0.1342 (95% CI 0.1291, 0.1393), setting a performance baseline without pre-training benefits.

A significant advancement was observed with the ImageNet pre-trained ResNet50, which attained a mAP50 of 0.2210 (95% CI 0.2187, 0.2233), underscoring the value of pre-training in feature representation across disparate image domains.

More strikingly, incorporating the Barlow Twins-CXR strategy on top of a randomly initialized ResNet50 led to a rapid performance ascent, achieving a mAP50 of 0.2448 (95% CI 0.2400, 0.2496). It marked an expedited training trajectory and a significant increase in detection performance.

When further enhanced by pre-training from ImageNet, the Barlow Twins-CXR approach yielded the best performance, recording a mAP of 0.2502 (95% CI 0.2465, 0.2539), evidencing the synergetic effect of combining pre-training methodologies.

The heat maps generated from the study present a compelling visualization of the performance of the BarlowTwins-CXR method compared to the traditional ImageNet weights approach. We generated heat maps of the first few CXR images of the training and test sets in Fig. 5. In each image, our method’s heat maps show a more focused alignment with the actual lesion areas marked by the Ground Truth Bbox. This indicates a higher precision in localizing and identifying pathological features

with BarlowTwins-CXR, potentially offering more targeted information for clinical diagnoses. Notably, in cases of cardiomegaly and lung opacity, the concentration and localization of the heatmaps from BarlowTwins-CXR are visibly superior to those derived from ImageNet weights, further affirming the efficacy of our approach in enhancing CXR image analysis.

Upon escalating the input resolution to 640 * 640 pixels, both ImageNet and Barlow Twin-CXR weighted models saw performance improvements due to the increased detail in the CXR images. Nonetheless, the performance differential between the two narrowed, indicating that the higher resolution somewhat mitigates the distinct advantages of self-supervised pre-training.

This points to intriguing future research avenues, such as refining image resolution parameters during pre-training and fine-tuning phases and investigating whether higher-resolution pre-training could elevate model performance. It also accentuates the necessity of tailoring deep learning model design to specific tasks, considering factors like image resolution and feature granularity.

Overall, implementing the Barlow Twins-CXR method on the VinDr dataset resulted in substantial gains despite its data limitations and the inherent challenges of CXR abnormality localization. An 11.5% performance enhancement over the baseline and a 2.8% increment over ImageNet pre-trained models were observed on the mAP50 metric. Such marked improvements confirm

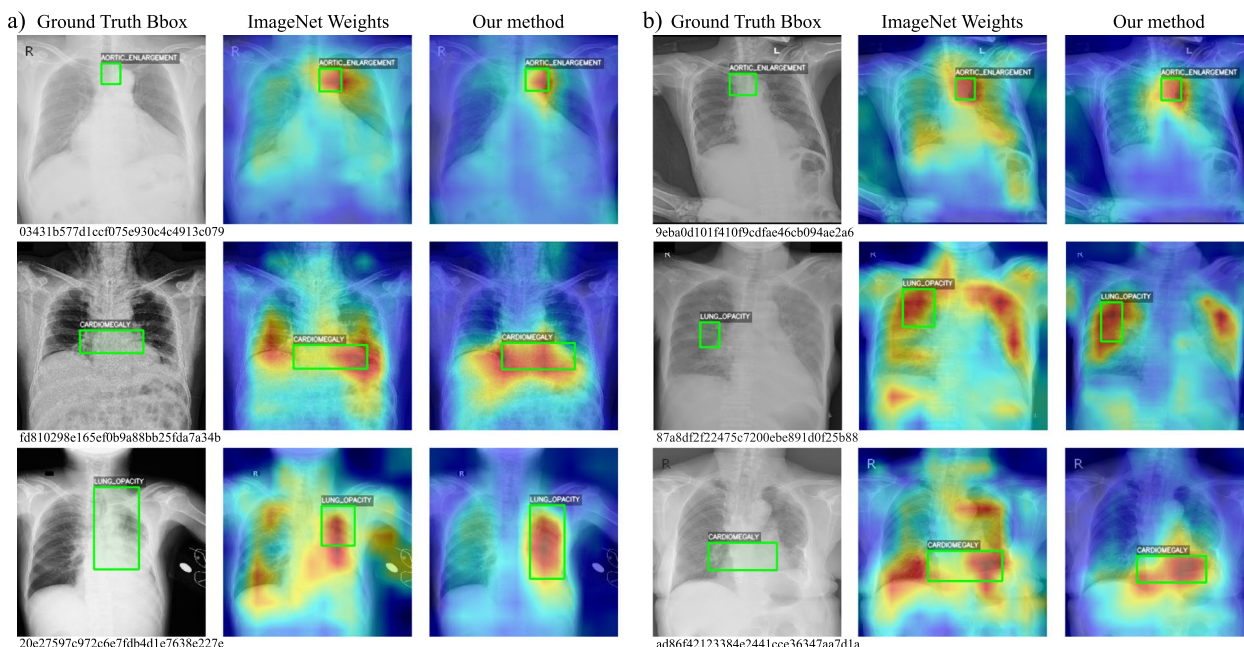


Fig. 5 Heatmaps were generated from the initial images of the training set(left) and test set(right), indicating successful Bbox predictions by the BarlowTwins-CXR model. Each heatmap corresponds to one accurately predicted bbox, despite multiple bboxes present in each CXR image. Serial numbers below the heatmaps refer to the image numbers in the dataset

the Barlow Twins-CXR strategy’s prowess in addressing domain inconsistencies, thereby fine-tuning naturally derived image weights for better applicability in CXR image analysis and beyond in medical imaging.

Linear evaluation protocol

In this experiment, we evaluated the impact of Barlow Twins-CXR pre-training versus traditional ImageNet pre-training on the linear classification performance within the NIH-CXR dataset. We adhered to the linear evaluation protocol, freezing the backbone of the linear classifier and updating only the final linear layer’s weights. This approach was applied across training datasets of varying sizes - 1%, 10%, and 100%, results of these experiments are presented in Fig. 6 and Table 2.

The results show that at a training data size of 1%, the Barlow Twins-CXR pre-trained model demonstrated a significant advantage, achieving an AUC of 0.6586 (95% CI 0.6544, 0.6628) compared to 0.5932 (95% CI 0.5905, 0.5959) for the ImageNet pre-trained model. As the

training data size increased to 10% and 100%, the AUCs for the Barlow Twins-CXR pre-trained model reached 0.7773 (95% CI 0.7749, 0.7797) and 0.8031 (95% CI 0.8026, 0.8036), respectively, while the ImageNet pre-trained model scored 0.6855 (95% CI 0.6808, 0.6903) and 0.7098 (95% CI 0.7085, 0.7111).

Notably, the incremental gains for both pre-training methods diminished with larger data sizes, suggesting that the performance boost provided by additional data becomes marginal when only the linear layer is updated.

These findings highlight the Barlow Twins-CXR pre-training method’s superiority over ImageNet pre-training across various dataset sizes, especially in data-limited scenarios. This demonstrates the promise of self-supervised learning in enhancing medical image analysis, particularly when annotated data is scarce.

End-to-end finetuning

In our end-to-end experiments, where we permitted updates to all model layers, the Barlow Twins-CXR

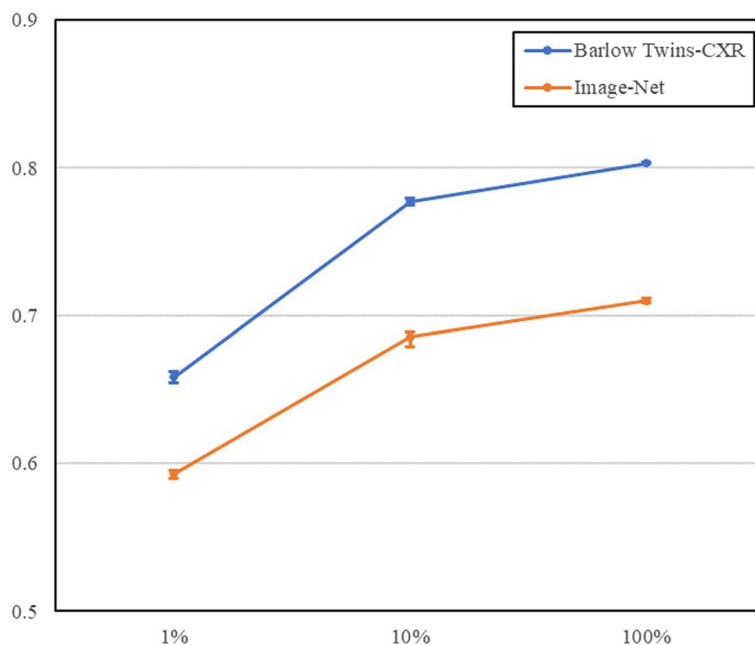


Fig. 6 AUC Scores with Error Bars for NIH-CXR Classification - This figure displays the AUC scores of linear models with Barlow Twins-CXR versus ImageNet weights across various dataset sizes (1%, 10%, 100%). As indicated by higher AUC scores, models using Barlow Twins-CXR consistently outperform those with ImageNet pre-training. Error bars represent the range of scores across five experiments

Table 2 AUC scores in validation and test sets for of linear models with varying pre-training methods at 224 and 640 input resolutions

Model	1%	10%	100%
Barlowtwins-CXR	0.6586 (0.6544, 0.6628)	0.7773 (0.7749, 0.7797)	0.8031 (0.8026, 0.8036)
Image-Net	0.5932 (0.5905, 0.5959)	0.6855 (0.6808, 0.6903)	0.7098 (0.7085, 0.7111)

Scores are presented with 95% confidence intervals

pre-trained ResNet50 backbone consistently outperformed the ImageNet pre-trained equivalent across all training set sizes. The results of these experiments are presented in Fig. 7 and Table 3.

At a 1% training data size, the Barlow Twins-CXR model achieved a 4.2% higher AUC than the ImageNet counterpart.

With 10% and 100% data sizes, the Barlow Twins-CXR model maintained leads of approximately 5.9% and 2.5%, respectively. Notably, the magnitude of improvement over the frozen backbone setup was less marked, suggesting that the wealth of features learned during self-supervised training reduces the margin for additional gains during subsequent fine-tuning.

Overall, these end-to-end fine-tuning results suggest that comprehensive learning across all model layers may elevate the risk of overfitting, particularly when data is scarce. The narrowing performance differential between the two pre-training strategies with increasing data volume indicates that the distinction between domain-specific

(Barlow Twins-CXR) and generalized (ImageNet) pre-training becomes less substantial with larger datasets. This trend implies that the influence of the pre-training strategy on the final performance of models may diminish as the size of the medical image dataset grows.

Discussion

Our study demonstrates that the BarlowTwins-CXR approach effectively utilizes unannotated CXR images for learning valuable representations and enhances transfer learning efficiency from ImageNet, thus addressing issues of domain inconsistency. This leads to quicker training and improved performance on tasks like abnormality detection in the VinDr-CXR dataset. Barlow Twins-CXR excels across various input resolutions, outshining models pre-trained on ImageNet.

One of the primary limitations of our study is the scarcity of CXR datasets with bounding box. Our reliance on public datasets, due to the absence of a private dataset, may limit the generalizability of our findings. Additionally,

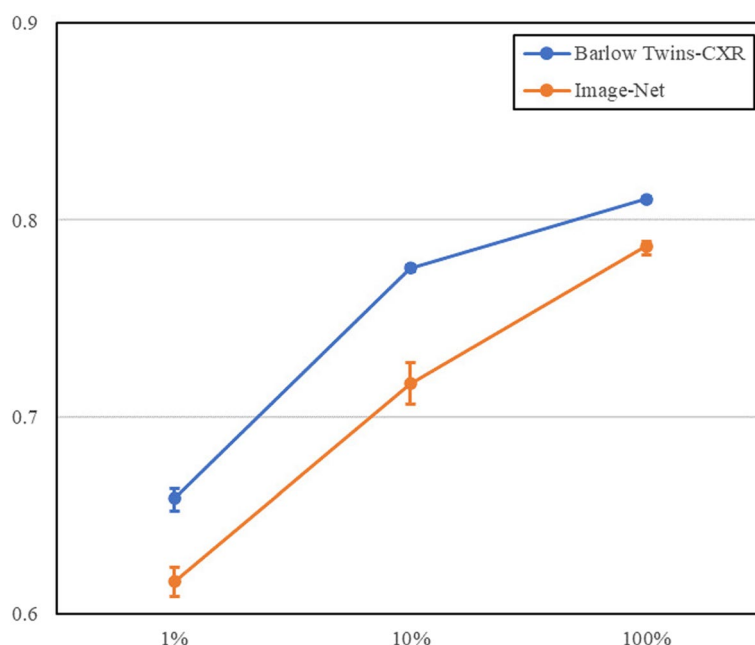


Fig. 7 AUC Scores with Error Bars for NIH-CXR Classification - This figure displays the AUC scores of models fine-tuned end-to-end with Barlow Twins-CXR versus ImageNet weights across various dataset sizes (1%, 10%, 100%). Higher AUC scores indicate that models using Barlow Twins-CXR consistently outperform those with ImageNet pre-training. Error bars represent the range of scores across five experiments

Table 3 AUC scores in validation and test sets for models fine-tuned end-to-end with varying pre-training methods at 224 and 640 input resolutions

Model	1%	10%	100%
Barlowtwins-CXR	0.6585 (0.6527, 0.6644)	0.7756 (0.7740, 0.7772)	0.8107 (0.8095, 0.8119)
Image-Net	0.6163 (0.6087, 0.6239)	0.7168 (0.7062, 0.7274)	0.7866 (0.7833, 0.7899)

Scores are presented with 95% confidence intervals

the computational cost of the BarlowTwins pre-training remains substantial. For a dataset size of 112,120 images with an image size of 224*224 pixels, the training process required two days on an NVIDIA A100 80G GPU. This significant resource requirement constrained our ability to experiment with higher image resolutions, which could potentially enhance the model's performance.

Future work

Our future endeavours include developing a demo interactive system for deployment and testing in emergency rooms. It will allow practical evaluation of the model's effectiveness in a clinical setting and facilitate the collection of a proprietary dataset. Additionally, we plan to explore more advanced self-supervised learning methods, object detection frameworks, and backbone networks to refine our approach further. The continuous evolution of these technologies promises to address some of the current limitations and expand the applicability and accuracy of our model in medical image analysis.

Conclusions

The results of this study provide strong support for the application of self-supervised learning in the field of abnormality detection, especially valuable in environments where radiologists face high workloads but the corresponding data labelling resources are scarce. A critical aspect of this approach is its adaptability to regional variations in CXR image, attributable to differences in imaging equipment, patient demographics, and other locale-specific factors [33, 34]. Such variations often impede the cross-regional applicability of a model, thus limiting its generalizability. By employing the BarlowTwins-CXR strategy, research organizations can transfer pre-trained backbone networks to local datasets tailored to the unique characteristics of their regional data.

Our findings might also have significant implications for clinical practice, suggesting that this strategy could be a game-changer in aiding radiologists to interpret CXR images efficiently. This technology promises to reduce diagnostic times, potentially increasing patients' throughput and improving the overall quality of care. Given its capacity for fine-tuning to specific regional characteristics, our approach holds particular promise in areas where standardization of medical imaging presents challenges.

In summary, the BarlowTwins-CXR approach demonstrates the potential of AI to enhance healthcare delivery. By integrating cutting-edge technology with clinical needs, we aim to pave the way for innovative solutions that benefit healthcare professionals and patients.

Abbreviations

AP	anterior-posterior
AUC	area under the receiver operating characteristic curve
CAM	Class Activation Mapping
CIUSSS	Centre intégré universitaire de santé et de services sociaux
CXR	chest X-ray radiography
FC	Fully connected layer
FPN	Feature Pyramid Network
FPR	False Positive Rate
IoU	Intersection over Union
ROC	receiver operating characteristic
ROI	region of interest
mAP	mean Average Precision
PA	posterior-anterior
TPR	True Positive Rate
WBF	Weighted Box Fusion
YOLO	You Only Look Once

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s12911-024-02529-9>.

Supplementary Material 1.

Acknowledgements

The authors wish to express their gratitude to CIUSSS du centre-sud-de-l'île-de-montréal for the computational resources and support provided, which were essential for the research conducted as part of the graduate internship program. We are especially thankful to our department director, Mathieu Mailhot, for his mentorship and to Chen Cheng for his collaborative efforts and valuable contributions to this project. Their expertise and insights have been greatly appreciated and substantially enhanced this work's quality.

Authors' contributions

HS designed the research methodology, analyzed data, was responsible for experiments and results visualization, and participated in manuscript drafting and revision. LM assisted in developing the research methodology and contributed to the drafting and revision of the manuscript. JFS collected and interpreted data, and provided expertise in statistical analysis. DL contributed to the study design, offered statistical analysis expertise, assisted in interpreting results, and played a significant role in the critical revision of the manuscript.

Funding

No external funding was associated with this research study.

Availability of data and materials

The datasets generated and/or analysed during the current study are available in the VinDr-CXR and NIH-CXRrepository: VinDr-CXR: <https://physionet.org/content/vindr-cxr/1.0.0/> NIH-CXR: <https://nihcc.app.box.com/v/ChestXray-NIHCC>.

Declarations

Ethics approval and consent to participate

All methods were performed under relevant guidelines and regulations (e.g., Declarations of Helsinki). The studies reported in this manuscript used reputable public datasets and did not require any additional data involving human participants, human data, or human tissue.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Author details

¹Département d'informatique et de recherche opérationnelle, Université de Montréal, 2920 chemin de la Tour, Montréal H3T 1J4, QC, Canada. ²Mila - Quebec AI Institute, 6666 Rue Saint-Urbain, Montréal H2S 3H1, QC, Canada. ³Direction des ressources informationnelles, CIUSSS du Centre-Sud-de-l'Île-de-Montréal, 400 Blvd. De Maisonneuve Ouest, Montréal H3A 1L4, QC, Canada. ⁴School of Medicine and College of Design and Engineering, National University of Singapore, 21 Lower Kent Ridge Rd, Singapore 119077, SG, Singapore.

Received: 16 February 2024 Accepted: 7 May 2024

Published online: 16 May 2024

References

- Satia I, Bashagha S, Bibi A, et al. Assessing the accuracy and certainty in interpreting chest X-rays in the medical division. *Clin Med*. 2013;13:349–52. PMID: 23908502.
- Rubin GD, Ryerson CJ, Haramati LB, et al. The role of chest imaging in patient management during the COVID-19 pandemic: a multinational consensus statement from the Fleischner Society. *Radiology*. 2020;296:172–80. PMID: 32275978.
- Lantsman DC, Barash Y, Klang E, Guranda L, Konen E, Tau N. Trend in radiologist workload compared to number of admissions in the emergency department. *Eur J Radiol*. 2022;149:110195. PMID: 35149337.
- <https://www.rsna.org/news/2022/may/Global-Radiologist-Shortage>. Accessed 2022.
- Seah JCY, Tang CHM, Buchlak QD, et al. Effect of a comprehensive deep-learning model on the accuracy of chest x-ray interpretation by radiologists: a retrospective, multireader multicase study. *Lancet Digit Health*. 2021;3:e496–506. PMID: 34219054.
- Morid MA, Borjali A, Del Fiol G. A scoping review of transfer learning research on medical image analysis using ImageNet. *Comput Biol Med*. 2021;128:104115. <https://doi.org/10.1016/j.combiomed.2020.104115>.
- Kim HE, Cosa-Linan A, Santhanam N, et al. Transfer learning for medical image classification: a literature review. *BMC Med Imaging*. 2022;22:69. <https://doi.org/10.1186/s12880-022-00793-7>.
- Zbontar J, Jing L, Misra I, et al. Barlow twins: Self-supervised learning via redundancy reduction. In: Proceedings of the International Conference on Machine Learning. PMLR; 2021. pp. 12310–12320. <https://doi.org/10.48550/arXiv.2103.03230>.
- Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L. ImageNet: A large-scale hierarchical image database. In: Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition, Miami, FL, USA. IEEE; 2009. pp. 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>.
- Nguyen HQ, Lam K, Le LT, et al. VinDr-CXR: An open dataset of chest X-rays with radiologist's annotations. *Sci Data*. 2022;9:429. <https://doi.org/10.1038/s41597-022-01498-w>.
- He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Las Vegas; 2016. pp. 770–778. <https://doi.org/10.1109/CVPR.2016.90>.
- Rajpurkar P, Irvin J, Zhu K, et al. CheXnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. 2017. <https://doi.org/10.48550/arXiv.1711.05225>. ArXiv preprint arXiv:1711.05225.
- Sun KX, Cong C. Research on chest abnormality detection based on improved YOLOv7 algorithm. In: Proceedings of the 2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). Las Vegas; 2022. pp. 3884–3886. <https://doi.org/10.1109/BIBM55620.2022.9995687>.
- Gaál G, Maga B, Lukács A. Attention u-net based adversarial architectures for chest x-ray lung segmentation. 2020. <https://doi.org/10.48550/arXiv.2003.10304>. ArXiv preprint arXiv:2003.10304.
- Shurrab S, Duwairi R. Self-supervised learning methods and applications in medical imaging analysis: A survey. *PeerJ Comput Sci*. 2022;8:e1045. <https://doi.org/10.7717/peerj-cs.1045>.
- Bengio Y, Lamblin P, Popovici D, et al. Greedy layer-wise training of deep networks. In: Proceedings of the 19th International Conference on Neural Information Processing Systems (NIPS'06). Cambridge; 2006. pp. 153–160. <https://doi.org/10.5555/2976456.2976476>.
- Pathak D, Krahenbuhl P, Donahue J, et al. Context encoders: Feature learning by inpainting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. Las Vegas; 2016. pp. 2536–2544. <https://doi.org/10.1109/CVPR.2016.278>.
- Chen T, Kornblith S, Norouzi M, et al. A simple framework for contrastive learning of visual representations. In: Proceedings of the International Conference on Machine Learning. 2020. pp. 1597–1607. <https://doi.org/10.5555/3524938.3525087>.
- Gidaris S, Singh P, Komodakis N. Unsupervised representation learning by predicting image rotations. 2018. <https://doi.org/10.48550/arXiv.1803.07728>. ArXiv preprint arXiv:1803.07728.
- Azizi S, Mustafa B, Ryan F, et al. Big self-supervised models advance medical image classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. Montreal; 2021. pp. 3478–3488. <https://doi.org/10.1109/ICCV48922.2021.00346>.
- Sowrirajan H, Yang J, Ng AY, Rajpurkar P. MoCo pretraining improves representation and transferability of chest x-ray models. In: Medical Imaging with Deep Learning. 2021. pp. 728–744. <https://doi.org/10.48550/arXiv.2010.05352>.
- Matsoukas C, Haslum J, Sorkhei M, Soderberg M, Smith K. What makes transfer learning work for medical images: Feature reuse & other factors. In: Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). New Orleans; 2022. pp. 9215–9224. <https://doi.org/10.1109/CVPR52688.2022.00901>.
- Wang X, Peng Y, Lu L, et al. ChestX-Ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Honolulu; 2017. pp. 3462–3471. <https://doi.org/10.1109/CVPR.2017.369>.
- Solovyeve R, Wang W, Gabruseva T. Weighted boxes fusion: Ensembling boxes from different object detection models. *Image Vis Comput*. 2021;107:104117. <https://doi.org/10.1016/j.imavis.2021.104117>.
- He K, Fan H, Wu Y, et al. Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. Seattle; 2020. pp. 9729–9738. <https://doi.org/10.1109/CVPR42600.2020.00975>.
- Girshick R. Fast R-CNN. In: Proceedings of the IEEE International Conference on Computer Vision. Santiago; 2015. pp. 1440–1448. <https://doi.org/10.1109/ICCV.2015.169>.
- Lin TY, Dollár P, Girshick R, et al. Feature pyramid networks for object detection. In: Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Honolulu; 2017. pp. 936–944. <https://doi.org/10.1109/CVPR.2017.106>.
- Chen K, Wang J, Pang J, Cao Y, et al. MMDetection: open mmlab detection toolbox and benchmark. 2019. <https://doi.org/10.48550/arXiv.1906.07155>. ArXiv preprint arXiv:1906.07155.
- Bachman P, Hjelm RD, Buchwalter W. Learning representations by maximizing mutual information across views. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems. Red Hook; 2019. pp. 15535–15545. <https://doi.org/10.5555/3454287.3455679>.
- Kornblith S, Shlens J, Le QV. Do better imagenet models transfer better? In: Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Long Beach; 2019. pp. 2656–2666. <https://doi.org/10.1109/CVPR.2019.00277>.
- Sechidis K, Tsoumakas G, Vlahavas I. On the stratification of multi-label data. In: Gunopulos D, Hofmann T, Malerba D, Vazirgiannis M, editors. Machine Learning and Knowledge Discovery in Databases. Berlin: Springer; 2011. p. 145–58.
- Szymański P, Kajdanowicz T. A network perspective on stratification of multi-label data. Proceedings of the First International Workshop on Learning with Imbalanced Domains: Theory and Applications. 2017. <https://doi.org/10.48550/arXiv.1704.08756>.
- Van Ryn M, Burke J. The effect of patient race and socio-economic status on physicians' perceptions of patients. *Soc Sci Med*. 2000;50:813–28. PMID: 10695979.
- Waite S, Scott J, Colombo D. Narrowing the gap: imaging disparities in radiology. *Radiology*. 2021;299:27–35. PMID: 33560191.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.