

DiRA: Discriminative, Restorative, and Adversarial Learning for Self-supervised Medical Image Analysis

Fatemeh Haghighi^{1*} Mohammad Reza Hosseinzadeh Taher^{1*} Michael B. Gotway²
 Jianming Liang¹
¹Arizona State University ²Mayo Clinic
 {fhaghigh, mhossei2, jianming.liang}@asu.edu Gotway.Michael@mayo.edu

Abstract

Discriminative learning, restorative learning, and adversarial learning have proven beneficial for self-supervised learning schemes in computer vision and medical imaging. Existing efforts, however, omit their synergistic effects on each other in a ternary setup, which, we envision, can significantly benefit deep semantic representation learning. To realize this vision, we have developed DiRA, the first framework that unites discriminative, restorative, and adversarial learning in a unified manner to collaboratively glean complementary visual information from unlabeled medical images for fine-grained semantic representation learning. Our extensive experiments demonstrate that DiRA (1) encourages collaborative learning among three learning ingredients, resulting in more generalizable representation across organs, diseases, and modalities; (2) outperforms fully supervised ImageNet models and increases robustness in small data regimes, reducing annotation cost across multiple medical imaging applications; (3) learns fine-grained semantic representation, facilitating accurate lesion localization with only image-level annotation; and (4) enhances state-of-the-art restorative approaches, revealing that DiRA is a general mechanism for united representation learning. All code and pretrained models are available at <https://github.com/JLiangLab/DiRA>.

1. Introduction

Self-supervised learning (SSL) aims to learn generalizable representations without using any expert annotation. The representation learning approaches in the SSL paradigm can be categorized into three main groups: (1) *discriminative learning*, which utilizes encoders to cluster instances of the same (pseudo) class and distinguish instances from different (pseudo) classes; (2) *restorative learning*, which utilizes generative models to reconstruct

*Equal contributors ordered alphabetically.




Figure 1. Despite the critical contributions of discriminative, restorative, and adversarial learning to SSL performance, yet no SSL method simultaneously employs all three learning ingredients. Our proposed DiRA, a novel SSL framework, unites discriminative, restorative, and adversarial learning in a unified manner to collaboratively glean complementary visual information from unlabeled data for fine-grained semantic representation learning.

original images from their distorted versions; and (3) *adversarial learning*, which utilizes adversary models to enhance restorative learning. In computer vision, discriminative SSL approaches, especially contrastive learning [8, 12, 13, 15, 21, 24, 27, 27, 34, 44, 53], currently offer state-of-the-art (SOTA) performance, surpassing standard supervised ImageNet models in some tasks. In medical imaging, however, restorative SSL methods [10, 25, 26, 43, 55, 57] compared to discriminative approaches [3, 56] presently reach a new height in performance. Naturally, we contemplate: *What contributes to the popularity differences between discriminative and restorative methods in computer vision and in medical imaging?* Furthermore, from our extensive literature review, we have discovered that *no* SSL method exploits all three learning components simultaneously; therefore, we ponder: *Can discriminative, restorative, and adversarial learning be seamlessly integrated into*




Figure 2. Photographic images typically have large foreground objects with apparent discriminative parts, whereas medical images contain consistent anatomical structures with semantic information dispersed over the entire images. As a result, recognition tasks in photographic images are mainly based on high-level features, while medical tasks demand holistic fine-grained discriminative features captured throughout images.

a single framework to foster collaborative learning for deep semantic representation, yielding more powerful models for a broad range of applications? In seeking answers to the two questions, we have gained the following insights.

Computer vision and medical imaging tasks embrace the spirit of evil in opposite ways, originating from the *marked* differences between photographic and medical images. Photographic images, particularly those in ImageNet, have large foreground objects with apparent discriminative parts, residing in varying backgrounds (e.g. zebra and daisy images in Fig. 2). Thus, object recognition tasks in photographic images are primarily based on high-level features captured from discriminative regions. In contrast, medical images generated from a particular imaging protocol exhibit consistent anatomical structures (e.g. chest anatomy in Fig. 2), with clinically relevant information dispersed over the entire image [26]. In particular, high-level structural information, i.e., anatomical structures and their relative spatial orientations, are essential for the identification of normal anatomy and various disorders. Importantly, medical tasks require much stronger attention to fine-grained details within images as identifying diseases, delineating organs, and isolating lesions rely on subtle, local variations in texture [29]. Therefore, recognition tasks in medical images desire complementary high-level and fine-grained discriminative features captured throughout images.

According to our systematical analysis, we have gained the following understandings: (1) discriminative learning excels in capturing high-level (global) discriminative features, (2) restorative learning is good at conserving fine-grained details embedded in local image regions, and (3) adversarial learning consolidates restoration by conserving

more fine-grained details. Putting these understandings and fundamental differences between photographic and medical images together would explain why restorative learning is preferred in medical imaging while discriminative learning is preferred in computer vision. More importantly, we have acquired a new and intriguing insight into *trio* of discriminative, restorative, and adversarial learning to excavate effective features required for medical recognition tasks—not only high-level anatomical representations but also fine-grained discriminative cues embedded in the local parts of medical images.

Based on the insights above, we have designed a novel self-supervised learning framework, called DiRA, by uniting **discriminative learning**, **restorative learning**, and **adversarial learning** in a unified manner to glean complementary visual information from unlabeled medical images. Our extensive experiments demonstrate that (1) DiRA encourages collaborative learning among three learning components, resulting in more generalizable representation across organs, diseases, and modalities (see Fig. 4); (2) DiRA outperforms fully supervised ImageNet models and increases robustness in small data regimes, thereby reducing annotation cost in medical imaging (Tab. 1 and Tab. 2); (3) DiRA learns fine-grained representations, facilitating more accurate lesion localization with only image-level annotations (Fig. 5); and (4) DiRA enhances SOTA restorative approaches, showing that DiRA is a general framework for united representation learning (Tab. 3).

In summary, we make the following contributions:

- The insights that we have gained into the *synergy* of discriminative, restorative, and adversarial learning in a ternary setup, realizing a new paradigm of collaborative learning for SSL.
- The first self-supervised learning framework that seamlessly unites discriminative, restorative, and adversarial learning in a unified manner, setting a new SOTA for SSL in medical imaging.
- A thorough and insightful set of experiments that demonstrate not only DiRA’s generalizability but also its potential to take a fundamental step towards developing *universal* representations for medical imaging.

2. Related works

Discriminative self-supervised learning. Discriminative methods can be divided into *class-level* and *instance-level* discrimination. *Class-level* discrimination methods [7, 8, 17, 22, 35, 54] group images based on certain criteria, assign a pseudo label to each group, and train a model to discriminate the images based on their pseudo labels, such as rotation degrees [22] and cluster assignments [7, 8, 54].

On the other hand, *instance-level* discrimination methods [8, 12, 13, 15, 21, 24, 27, 34, 44, 48, 52, 53] treat each image as a distinct class, and maximize the similarity of representations derived from different views of the same image, seeking to learn transformation invariant representations. *Instance-level* discriminative learning has been investigated in various forms, including contrastive learning [12, 14, 27, 49], asymmetric networks [15, 24], and redundancy reduction [21, 53]. However, both *class-level* and *instance-level* approaches in discriminative learning have shown failures in tasks that require finer-grained features [47, 50, 51]. Our DiRA addresses this limitation by incorporating restorative and adversarial learning, which not only improves discriminative learning but also yields fine-grained representations required for medical imaging tasks.

Restorative and adversarial self-supervised learning.

The key objective for a restorative method is to faithfully reconstruct the distribution of data [36, 48]. In the SSL context, multiple pretext tasks are formulated to reconstruct the perturbed images using generative models [33, 37, 45]. The advance of GANs [23] has led to a new line of research in unsupervised learning, using adversarial learning to generate transferable representations [18, 19]. While recent works [11, 18] have demonstrated impressive results by employing large-scale generative models, it remains unclear to what extent generative models can encapsulate high-level structures. Our DiRA alleviates this limitation by bringing the advantages of discriminative learning into generative models. Through discriminating image samples, generative models are encouraged to capture global discriminative representations rather than superficial representations, leading to a more pronounced embedding space.

Self-supervised learning in medical imaging. Due to the lack of large annotated datasets, SSL created substantial interest in medical imaging. Motivated by the success in computer vision, recent discriminative methods concentrate on instance-level discrimination. A comprehensive benchmarking study in [29] evaluated the efficacy of existing instance discrimination methods pre-trained on ImageNet for diverse medical tasks. Several other works adjusted contrastive-based methods on medical images [3, 9, 56]. A large body of work, on the other hand, focuses on restorative approaches, which can be categorized into restorative only [10, 57], restorative and adversarial [43], and discriminative and restorative [26, 30, 55]. Among these groups, the most recent study on TransVW [25, 26] demonstrated superiority by combining discriminative and restorative components into a single SSL framework. DiRA distinguishes itself from all previous works by demonstrating two key advances: (1) employing discriminative, restorative, and adversarial learning simultaneously in a unified framework; and (2) providing a general representation learning framework that is compatible with existing discriminative and

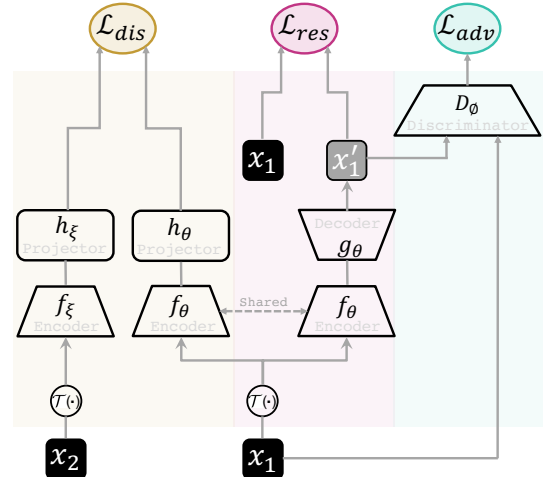


Figure 3. **Our proposed framework.** DiRA consists of three learning components: discriminative, restorative, and adversary. Given two input patches x_1 and x_2 , we perturb them with $\mathcal{T}(\cdot)$ and provide them as input to discrimination and restoration branches. The discrimination branch consists of encoders f_θ and f_ξ , and projectors h_θ and h_ξ , and maximizes the agreement between (high-level) embedding vectors of samples from the same (pseudo) class. The restoration branch consists of encoder f_θ and decoder g_θ , and maximizes the (pixel-level) agreement between original sample x_1 and restored x'_1 . Adversarial discriminator D_ϕ contrasts the original samples with the restored ones, reinforcing the restoration to preserve more fine-grained details.

restorative methods, regardless of their objective functions.

3. DiRA framework

As shown in Fig. 3, DiRA is a SSL framework comprised of three key components: (1) Discrimination (Di) that aims to learn high-level discriminative representations, (2) Restoration (R) that aims to enforce the model to conserve fine-grained information about the image by focusing on more localized visual patterns, and (3) Adversary (A) that aims to further improve feature learning through the restoration component. By integrating these components into a *unified* framework, DiRA captures comprehensive information from images, providing more powerful representations for various downstream tasks. In the following, we first introduce each component by abstracting a common paradigm and then describe the joint training loss.

3.1. Discriminative learning

Discriminative learning can be thought of as training an encoder to maximize agreement between instances of the same (pseudo) class in the latent space via a discriminative loss. As illustrated in Fig. 3, the discriminator branch is comprised of two twin backbone networks f_θ and f_ξ , and projection heads h_θ and h_ξ . f_θ is a regular encoder, while

f_ξ can be a momentum encoder [24, 27] or share weights with f_θ [15, 26, 53]. Given two patches x_1 and x_2 , which are cropped from the same image or different images, we first apply an augmentation function $\mathcal{T}(\cdot)$ on them. The two augmented patches are then processed by f_θ and f_ξ to generate latent features $y_1 = f_\theta(\mathcal{T}(x_1))$ and $y_2 = f_\xi(\mathcal{T}(x_2))$. The projection heads h_θ and h_ξ projects the latent features to a unit sphere and output projections $z_1 = h_\theta(y_1)$ and $z_2 = h_\xi(y_2)$. The discriminator’s objective is to maximize the similarity between the embedding vectors obtained from two samples of the same (pseudo) class:

$$\mathcal{L}_{dis} = \ell(z_1, z_2) \quad (1)$$

where $\ell(z_1, z_2)$ is the similarity/distance function that measures compatibility between z_1 and z_2 . DiRA is a general framework that allows various choices of discrimination tasks without any constraint. As such, the declaration of class might range from considering every single image as a class (instance discrimination) to clustering images based on a similarity metric (cluster discrimination). Accordingly, x_1 and x_2 can be two views of the same image or two samples from the same cluster. Based on the nature of the discrimination task, the instantiation of \mathcal{L}_{dis} can be cross-entropy [22, 26, 35, 58], contrastive [3, 8, 12, 27], redundancy reduction [21, 53], etc.

3.2. Restorative learning

Our restorative learning branch aims to enhance discrimination learning by leveraging fine-grained visual information. As shown in Fig. 3, the restoration branch is comprise of an encoder f_θ and decoder g_θ , where encoder f_θ is shared with the discrimination branch. Given the input sample x_1 distorted by \mathcal{T} , the f_θ and g_θ aims to map the distorted sample back to the original one, i.e., $f_\theta, g_\theta : (x, \mathcal{T}) \mapsto x$. f_θ and g_θ are trained by minimizing the distance between the original sample and the restored one at pixel-level:

$$\mathcal{L}_{res} = \mathbb{E}_x \text{dist}(x_1, x'_1) \quad (2)$$

where $x'_1 = g_\theta(f_\theta(\mathcal{T}(x_1)))$ denotes the restored image. $\text{dist}(x_1, x'_1)$ presents the distance function that measures similarity between x_1 and x'_1 , such as L_1 or L_2 .

3.3. Adversarial learning

Adversarial learning aims to reinforce f_θ by measuring how realistic the restored images are. As such, the adversarial discriminator D_ϕ is formulated to distinguish (discriminate) the set of training images from the set of synthesized images, guiding encoder f_θ to capture more informative features from images so that g_θ can reproduce the original images effectively. Therefore, the encoder f_θ and decoder g_θ play a minimax game with adversarial discriminator D_ϕ , and are optimized jointly with an adversarial loss [6, 36]:

$$\mathcal{L}_{adv} = \mathbb{E}_x [\log D_\phi(x_1)] + \mathbb{E}_x [\log(1 - D_\phi(x'_1))] \quad (3)$$

3.4. Joint training

Finally, the combined objective for the proposed DiRA framework becomes:

$$\mathcal{L} = \lambda_{dis} * \mathcal{L}_{dis} + \lambda_{res} * \mathcal{L}_{res} + \lambda_{adv} * \mathcal{L}_{adv} \quad (4)$$

where λ_{dis} , λ_{res} , and λ_{adv} are multiplication factors that determine the relative importance of different losses. Through our unified training scheme, DiRA learns a representation that preserves fine-grained details within the samples while being discriminative among the image classes. In particular, the formulation of \mathcal{L}_{dis} encourages the model to capture high-level discriminative features. Moreover, \mathcal{L}_{res} forces the model to encode fine-grained information from the images by focusing on pixel-level visual patterns. This results in more descriptive feature embeddings that elevate the discrimination task. Finally, \mathcal{L}_{adv} elevates restoration based learning through capturing more informative features.

4. Implementations details

4.1. Pre-training protocol

DiRA is a general framework that is compatible with existing self-supervised discriminative and restorative methods, regardless of their objective functions. To assess the effectiveness of our framework, we adapt recent SOTA 2D and 3D self-supervised methods into DiRA, as described in the following. The pretrained models with DiRA are identified as DiRA subscripted by the original method name.

2D image pretraining settings. We apply DiRA to MoCov2 [14], Barlow Twins [53], and SimSiam [15] for 2D image self-supervised learning. All DiRA models are pretrained from scratch on the training set of ChestX-ray14 [46] dataset. For each of these three discrimination tasks [14, 15, 53], we follow the original methods in the formulation of \mathcal{L}_{dis} , projection head architecture, and hyperparameters settings. Furthermore, we optimize the encoder and decoder networks f_θ and g_θ following the optimization setups in [14, 15, 53]. For all methods, we employ a 2D U-Net [38] with a standard ResNet-50 [28] backbone as the f_θ and g_θ . We adopt mean square error (MSE) as the \mathcal{L}_{res} . The adversarial discriminator network D_ϕ consists of four convolutional layers with the kernel size of 3×3 [37], trained using the Adam optimizer with a learning rate of $2e-4$ and $(\beta_1, \beta_2) = (0.5, 0.999)$. We use batch size 256 distributed across 4 Nvidia V100 GPUs. $\lambda_{res}, \lambda_{adv}, \lambda_{dis}$ are empirically set to 10, 0.001, and 1, respectively. Input images are first randomly cropped and resized to 224×224 ; the image augmentation function $\mathcal{T}(\cdot)$ includes random horizontal flipping, color jittering, and Gaussian blurring. Additionally, we apply cutout [16, 37] and shuffling [10] to make the restoration task more challenging. More details are provided in the Appendix.




Figure 4. **Comparison with discriminative self-supervised methods:** We apply our DiRA to three representative SOTA self-supervised methods with different discrimination objectives: MoCo-v2 [14], Barlow Twins [53], and SimSiam [15]. DiRA empowers discriminative methods to capture more fine-grained representations, yielding significant ($p < 0.05$) performance gains on four downstream tasks.

3D volume pretraining settings. We apply DiRA to TransVW [26], the SOTA method for 3D self-supervised learning in medical imaging. We adapt TransVW in DiRA by adding an adversarial discriminator D_ϕ into its training scheme. For fair comparisons, we follow the publicly available TransVW code for setting instance discrimination and restoration tasks. Moreover, similar to publicly released TransVW, DiRA models are pre-trained from scratch using 623 chest CT scans in the LUNA [40] dataset. We use 3D U-Net [20] as the encoder-decoder network and a classification head including fully-connected layers. The adversarial discriminator D_ϕ includes four convolutional blocks with the kernel size $3 \times 3 \times 3$. λ_{res} , λ_{adv} , λ_{dis} are empirically set to 100, 1, and 1, respectively. f_θ , g_θ , and D_ϕ are optimized for 200 epochs using Adam with a learning rate of $1e-3$ and batch size of 8. More details are provided in the Appendix.

4.2. Transfer learning protocol

Target tasks and datasets. We evaluate the effectiveness of DiRA’s representations in transfer learning to a diverse suite of 9 common but challenging 2D and 3D medical imaging tasks, including: ChestX-ray14, CheXpert [31], SIIM-ACR [1], and NIH Montgomery [32] for 2D models, and LUNA, PE-CAD [41], LIDC-IDRI [2], LiTS [5], and BraTS [4] for 3D models (see Appendix for dataset details). These tasks encompass various label structures (multi-label classification and pixel-level segmentation), diseases (brain tumors and thoracic diseases, such as lung nodules, pulmonary emboli, and pneumothorax), organs (lung, liver, brain), and modalities (X-ray, CT, MRI). Moreover, these tasks contain many hallmark challenges encountered when working with medical images, such as imbalanced classes, limited data, and small-scanning areas for the pathology of interest [3, 29]. We use the official data split of these datasets when available; otherwise, we randomly divide the data into 80%/20% for training/testing.

Fine-tuning settings. We transfer the pre-trained (1) encoder (f_θ) to the classification tasks, and (2) encoder and

decoder (f_θ and g_θ) to segmentation tasks. We evaluated the generalization of DiRA representations by fine-tuning all the parameters of downstream models. We use the AUC (area under the ROC curve), and the IoU (Intersection over Union) and Dice coefficient for evaluating classification and segmentation performances, respectively. Following [29], we strive to optimize each downstream task with the best performing hyperparameters (details in Appendix). We employ the early-stop mechanism using 10% of the training data as the validation set to avoid over-fitting. We run each method ten times on each downstream task and report the average, standard deviation, and statistical analysis based on an independent two-sample t -test.

5. Results

We conduct extensive experiments to better understand not only the properties of our framework but also its generalizability across 9 downstream tasks. Through the following groups of experiments, we establish that DiRA (1) enriches existing discriminative approaches, capturing a more diverse visual representation that generalizes better to different tasks; (2) addresses the annotation scarcity challenge in medical imaging, providing an annotation-efficient solution for medical imaging; (3) learns fine-grained features, facilitating more accurate lesion localization with only image-level annotation; and (4) improves SOTA restorative approaches, demonstrating that DiRA is a general framework for united representation learning.

5.1. DiRA enriches discriminative learning

Experimental setup: To study the flexibility and efficacy of our proposed self-supervised framework, we apply DiRA to three recent SOTA self-supervised methods with diverse discrimination objectives: MoCo-v2, Barlow Twins, and SimSiam. To evaluate the quality of our learned representations and ascertain the generality of our findings, we follow [29] and consider a broader range of four target tasks, covering classification (ChestX-Ray14 and CheXpert) and

Method	ChestX-ray14 [AUC (%)]			CheXpert [AUC (%)]			Montgomery [Dice (%)]		
	Label fraction			Label fraction			Label fraction		
	1%	25%	50%	1%	25%	50%	1%	25%	50%
MoCo-v2 [14]	52.99	74.89	76.71	76.87	81.70	83.23	63.69	96.44	97.60
DiRA _{MoCo-v2}	59.39 (↑ 6.4)	77.55 (↑ 2.6)	78.74 (↑ 2.0)	78.43 (↑ 1.5)	87.12 (↑ 5.4)	87.31 (↑ 4.0)	72.53 (↑ 8.8)	97.06 (↑ 0.62)	98.14 (↑ 0.5)
Barlow Twins [53]	62.43	76.23	77.59	82.85	83.74	84.66	86.79	97.49	97.68
DiRA _{Barlow Twins}	62.51 (↑ 0.08)	77.18 (↑ 0.9)	78.46 (↑ 0.8)	83.12 (↑ 0.2)	84.20 (↑ 0.4)	85.32 (↑ 0.6)	87.25 (↑ 0.4)	97.62 (↑ 0.1)	98.15 (↑ 0.4)
SimSiam [15]	51.07	73.05	75.20	65.39	80.05	81.46	48.20	94.86	97.21
DiRA _{SimSiam}	53.42 (↑ 2.3)	74.38 (↑ 1.3)	76.43 (↑ 1.2)	70.46 (↑ 5.0)	81.03 (↑ 1.0)	82.70 (↑ 1.2)	61.86 (↑ 13.6)	96.61 (↑ 1.7)	97.91 (↑ 0.7)

Table 1. **Transfer learning under different downstream label fractions:** DiRA models combat overfitting in low data regimes and provide stronger representations for downstream tasks with limited annotated data. For each downstream task, we report the average performance over multiple runs. (↑) shows the improvement of DiRA models compared with the underlying discriminative method.

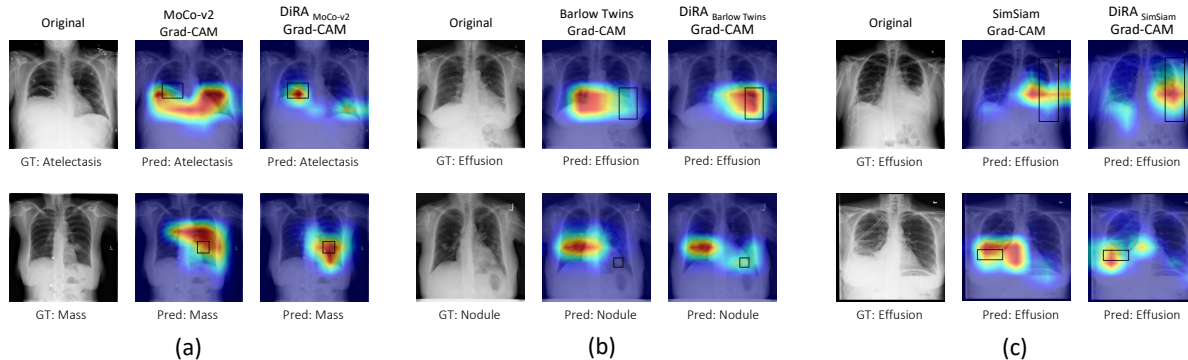


Figure 5. **Visualization of Grad-CAM heatmaps** for (a) MoCo-v2 vs. DiRA_{MoCo-v2}, (b) Barlow Twins vs. DiRA_{Barlow Twins}, and (c) SimSiam vs. DiRA_{SimSiam}. Ground truth bounding box annotations are shown in black. Training with DiRA leads to improvements in weakly-supervised disease localization. While both DiRA and underlying models predict the correct disease label on the test images, DiRA models capture the diseased locations more precisely than the baselines which attune to larger regions of the image (e.g. (c), second row) or provide inaccurate localization with no overlap with the ground truth (e.g. (b), second row).

segmentation (SIIM-ACR and Montgomery).

Results: As seen in Fig. 4, utilizing our DiRA framework consistently enhances its underlying discriminative method across all tasks (1) ChestX-ray14, (2) CheXpert, (3) SIIM-ACR, and (4) NIH Montgomery. Compared with the original methods, DiRA_{MoCo-v2} showed increased performance by 0.76%, 1.17%, 1.35%, and 0.21%, respectively; Similarly, DiRA_{Barlow Twins} showed increased performance by 0.43%, 0.60%, 0.16%, and 0.03%. Finally, DiRA_{SimSiam} showed increased performance by 0.82%, 2.22%, 1.18%, and 0.45%. These results imply that DiRA is a comprehensive representation learning framework that encourages existing self-supervised instance discriminative approaches to retain more fine-grained information from images, enriching their visual representation and allowing them to generalize to different medical tasks more effectively.

5.2. DiRA improves robustness to small data regimes

Experimental setup: We investigate the robustness of representations learned with DiRA in small data regimes to determine if the learned representation can serve as a

proper foundation for fine-tuning. We randomly select 1%, 25%, and 50% of training data from ChestX-ray14, CheXpert, and Montgomery, and fine-tune the self-supervised pre-trained models on these training-data subsets.

Results: As shown in Tab. 1, our DiRA pre-trained models outperform their counterparts' original methods in all subsets, 1%, 25%, and 50%, across ChestX-ray14, CheXpert, and Montgomery. In particular, the average of improvement for MoCo-v2 and SimSiam across all three downstream tasks in each underlying subset garnering: (1) 5.6 % and 7% when using 1%, (2) 2.9 % and 1.3% when using 25%, and (3) 2.2 % and 1% when using 50%. As seen in 1%, DiRA outperforms its counterparts MoCo-v2 and SimSiam by a large margin, demonstrating our framework's potential for combating overfitting in extreme low data regimes. Although the Barlow Twins is more resistant to low data regimes than the previous two approaches, DiRA still improves its performance by 0.5%, 0.5%, and 0.6% on average across all three datasets when using 1%, 25%, and 50% of labeled data, respectively. In summary, our results in the low-data regimes demonstrate our framework's superiority for providing more robust and transferable repre-

Method	Pretraining Dataset	Classification [AUC (%)]		Segmentation [Dice (%)]	
		ChestX-ray14	CheXpert	SIIM-ACR	Montgomery
Random	-	80.31±0.10	86.62±0.15	67.54±0.60	97.55±0.36
Supervised	ImageNet	81.70±0.15	87.17±0.22	67.93±1.45	98.19±0.13
Supervised	ChestX-ray14	-	87.40±0.26	68.92±0.98	98.16±0.05
DiRA _{MoCo-v2}	ChestX-ray14	81.12±0.17	87.59±0.28 ↑ ↑	69.24±0.41 ↑ *	98.24±0.09 * ↑
DiRA _{Barlow Twins}	ChestX-ray14	80.88±0.30	87.50±0.27 ↑ *	69.87±0.68 ↑ ↑	98.16±0.06 * *
DiRA _{SimSiam}	ChestX-ray14	80.44±0.29	86.04±0.43	68.76±0.69 * *	98.17±0.11 * *

Table 2. **Comparison with fully-supervised transfer learning:** DiRA models outperform fully-supervised pre-trained models on ImageNet and ChestX-ray14 in three downstream tasks. The best methods are bolded while the second best are underlined. ↑ and ↑ present the statistically significant ($p < 0.05$) improvement compared with supervised ImageNet and ChestX-ray14 baselines, respectively, while * and * presents the statistically equivalent performances accordingly. For supervised ChestX-ray14 model, transfer learning to ChestX-ray14 is not applicable since pre-training and downstream tasks are the same, denoted by “-”.

sentations that can be harnessed for downstream tasks with limited amounts of data, thereby reducing annotation costs.

5.3. DiRA improves weakly-supervised localization

Experimental setup: We investigate our DiRA framework in a weakly supervised setting, comparing its applicability for localizing chest pathology to underlying discriminative methods. Given this goal, we follow [46] and use the ChestX-ray14 dataset, which contains bounding box annotations for approximately 1,000 images. For training, we initialize models with our DiRA pre-trained models, and train downstream models using only image-level disease labels. Following [39, 46], bounding boxes are only used as ground truth to evaluate disease localization accuracy in the testing phase. To generate heatmaps, we leverage Grad-CAM [39]. Heatmaps indicate the spatial location of a particular thoracic disease.

Results: As seen in Fig. 5, our framework learns more fine-grained representations, enabling it to localize diseases more accurately. In particular, heatmaps generated by MoCo-v2, Barlow Twins, and SimSiam models are highly variable, whereas DiRA models consistently achieve more robust and accurate localization results over each corresponding original method. Through the production of more interpretable activation maps, our DiRA framework demonstrates possible clinical potential for post-hoc interpretation by radiologists. Quantitative disease localization results are provided in the Appendix.

5.4. DiRA outperforms fully-supervised baselines

Experimental setup: Following the recent transfer learning benchmark in medical imaging [29], we compare the transferability of DiRA models, pre-trained solely on unlabeled images from ChestX-ray14, with two fully-supervised representation learning approaches: (1) supervised ImageNet model, the most common transfer learning pipeline in medical imaging and (2) supervised model pre-trained on ChestX-ray14, the upper-bound in-domain transfer learning baseline. The supervised baselines benefit from

Dataset	Method		
	Random	TransVW [26]	DiRA _{TransVW}
LUNA	94.25±5.07	98.46±0.30	98.87±0.61 (↑ 0.41)
LIDC-IDRI	74.05±1.97	77.33±0.52	77.51±1.36 (↑ 0.18)
LiTS	79.76±5.42	86.53±1.30	86.85±0.81 (↑ 0.32)
BraTS	59.87±4.04	68.82±0.38	69.57±1.13 (↑ 0.75)
PE-CAD	80.36±3.58	87.07±2.83	86.91±3.27

Table 3. **Comparison with restorative self-supervised method:** We apply our DiRA to the TransVW as the SOTA restorative self-supervised method. DiRA enhances TransVW by conserving more fine-grained details, resulting in performance boosts in four 3D downstream tasks.

the same encoder as DiRA, namely ResNet-50. We fine-tune all pre-trained models for 4 distinct medical applications ranging from target tasks on the source dataset to the tasks with comparatively significant domain-shifts in terms of data distribution and disease/object of interest.

Results: As shown in Tab. 2, DiRA models achieves significantly better or on-par performance compared with both supervised ImageNet and ChestX-ray14 models across four downstream tasks. In particular, DiRA_{MoCo-v2} and DiRA_{Barlow Twins}, outperforms both supervised baselines in CheXpert, SIIM-ACR, and Montgomery, respectively. Moreover, DiRA_{SimSiam} outperforms the supervised ImageNet and the ChestX-ray14 pre-trained models in SIIM-ACR and Montgomery, respectively. These results indicate that our framework, with zero annotated data, is capable of providing more generic features for different medical tasks.

5.5. DiRA sets a new state-of-the-art for self-supervised learning in 3D medical imaging

Experimental setup: We further investigate the effectiveness of our framework for enhancing restorative representation learning by applying DiRA to TransVW [26], the state-of-the-art SSL approach for 3D medical imaging. We select TransVW as representative of restorative self-supervised methods because it shows superior performance over discriminative [42, 58], restorative only [10, 57], and restora-

Base	Pretraining dataset	\mathcal{L}_{dis}	\mathcal{L}_{res}	\mathcal{L}_{adv}	Classification [AUC (%)]		Segmentation [Dice (%)]	
					ChestX-ray14	CheXpert	SIIM-ACR	Montgomery
MoCo-v2	ChestX-ray14	✓	×	×	80.36±0.26	86.42±0.42	67.89±1.14	98.03±0.22
		✓	✓	×	80.72±0.29 ↑	86.86±0.37 ↑	68.16±1.07 ↑	98.19±0.08 ↑
		✓	✓	✓	81.12±0.17 ↑	87.59±0.28 ↑	69.24±0.41 ↑	98.24±0.09 ↑
Barlow Twins	ChestX-ray14	✓	×	×	80.45±0.29	86.90±0.62	69.71±0.34	98.13±0.13
		✓	✓	×	80.86±0.16 ↑	87.44±0.33 ↑	69.83±0.29 ↑	98.15±0.14 ↑
		✓	✓	✓	80.88±0.30 ↑	87.50±0.27 ↑	69.87±0.68 ↑	98.16±0.06 ↑
SimSiam	ChestX-ray14	✓	×	×	79.62±0.34	83.82±0.94	67.58±1.89	97.72±0.27
		✓	✓	×	79.41±0.42 ↓	84.45±0.46 ↑	68.35±1.16 ↑	98.02±0.21 ↑
		✓	✓	✓	80.44±0.29 ↑	86.04±0.43 ↑	68.76±0.69 ↑	98.17±0.11 ↑

Table 4. **Ablation study on different components of DiRA:** We study the impact of each component of DiRA, including discrimination, restoration, and adversary, in four downstream tasks. Adding restorative learning (\mathcal{L}_{res}) to discriminative learning leads to consistent performance improvements. Furthermore, equipping models with adversarial learning (\mathcal{L}_{adv}) yields performance boosts across all tasks.

tive and adversarial [43] methods. Following the common evaluation pipeline [26], we evaluate our learned representations by transfer learning to five common and challenging 3D downstream tasks, including classification (LUNA and PE-CAD) and segmentation (LIDC, LiTS, and BraTS).

Results: As shown in Tab. 3, DiRA framework consistently enhances TransVW across all downstream tasks. In particular, DiRA improves TransVW in LUNA, LIDC-IDRI, LiTS, and BraTS, and offers equivalent performance in PE-CAD. These results imply that by utilizing three learning components in tandem, image-based self-supervision approaches capture a more diverse visual representation that generalizes better to different downstream tasks.

6. Ablation study

Experimental setup: We conduct a thorough ablation study to show how each component contributes to DiRA. To do so, we only vary the loss function of DiRA. For each underlying self-supervised method, i.e. MoCo-v2, Barlow Twins, and SimSiam (referred to as the base), we start with the discrimination component and incrementally add restorative and the adversarial learning. When all three components are unified, they represent the completed DiRA models. All models are pretrained on the ChestX-ray14 dataset and fine-tuned for four downstream tasks, including ChestX-ray14, CheXpert, SIIM-ACR, and Montgomery.

Results: We draw the following observations from the results in Tab. 4: (1) Expanding discriminative self-supervised methods by adding a restoration task consistently enhances the original methods. In particular, incorporating \mathcal{L}_{res} into training objectives of MoCo-v2, Barlow Twins, and SimSiam outperforms the corresponding original methods, with the exception of SimSiam in ChestX-ray14, which shows slight performance degradation. Note that this gap later compensates after adding \mathcal{L}_{adv} , which signifies collaborative learning among restorative and adversary components in our framework. (2) The overall trend showcases the advantage of the adversarial discriminator when added to the restoration component, improving

the performance of all methods in four downstream tasks. Our findings indicate that unifying the three components in DiRA models significantly enhances the original self-supervised methods by retaining more fine-grained information from images.

7. Conclusion and discussion

We propose DiRA, the *first* SSL framework that unites discriminative, restorative, and adversarial learning in a unified manner. The key contribution of our DiRA arises from the insights that we have gained into the *synergy* of these three SSL approaches for collaborative learning. Given DiRA’s generalizability, we envisage it will take a fundamental step towards developing universal representations for medical imaging. Our DiRA achieves remarkable performance gains, though we fixed the restorative learning tasks in all experiments when examining various formulations of discriminative learning. In the future, examining various choices of restoration tasks and searching for optimal collaborative learning strategies may lead to even stronger representations for medical imaging. In this paper, we have focused on medical imaging, but we envision that DiRA can also offer outstanding performance for vision tasks that demand fine-grained details.

Acknowledgments: With the help of Zongwei Zhou, Zuwei Guo started implementing the earlier ideas behind “United & Unified”, which has branched out into DiRA. We thank them for their feasibility exploration, especially their initial evaluation on TransVW [26] and various training strategies. This research has been supported in part by ASU and Mayo Clinic through a Seed Grant and an Innovation Grant and in part by the NIH under Award Number R01HL128785. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH. This work utilized the GPUs provided in part by the ASU Research Computing and in part by the Extreme Science and Engineering Discovery Environment (XSEDE) funded by the National Science Foundation (NSF) under grant number ACI-1548562. Paper content is covered by patents pending.

References

- [1] Siim-acr pneumothorax segmentation, 2019. **5**
- [2] Samuel G Armato III, Geoffrey McLennan, Luc Bidaut, Michael F McNitt-Gray, Charles R Meyer, Anthony P Reeves, Binsheng Zhao, Denise R Aberle, Claudia I Henschke, Eric A Hoffman, et al. The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931, 2011. **5**
- [3] Shekoofeh Azizi, Basil Mustafa, Fiona Ryan, Zachary Beaver, Jan Freyberg, Jonathan Deaton, Aaron Loh, Alan Karthikesalingam, Simon Kornblith, Ting Chen, Vivek Natarajan, and Mohammad Norouzi. Big self-supervised models advance medical image classification. *arXiv:2101.05224*, 2021. **1, 3, 4, 5**
- [4] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv:1811.02629*, 2018. **5**
- [5] Patrick Bilic, Patrick Ferdinand Christ, Eugene Vorontsov, Grzegorz Chlebus, Hao Chen, Qi Dou, Chi-Wing Fu, Xiao Han, Pheng-Ann Heng, Jürgen Hesser, et al. The liver tumor segmentation benchmark (lits). *arXiv:1901.04056*, 2019. **5**
- [6] Bing Cao, Han Zhang, Nannan Wang, Xinbo Gao, and Dinggang Shen. Auto-gan: Self-supervised collaborative learning for medical image synthesis. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):10486–10493, Apr. 2020. **4**
- [7] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision*, pages 132–149, 2018. **2**
- [8] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv:2006.09882*, 2021. **1, 2, 3, 4**
- [9] Krishna Chaitanya, Ertunc Erdil, Neerav Karani, and Ender Konukoglu. Contrastive learning of global and local features for medical image segmentation with limited annotations. In *Advances in Neural Information Processing Systems*, volume 33, pages 12546–12558. Curran Associates, Inc., 2020. **3**
- [10] Liang Chen, Paul Bentley, Kensaku Mori, Kazunari Misawa, Michitaka Fujiwara, and Daniel Rueckert. Self-supervised learning for medical image analysis using image context restoration. *Medical image analysis*, 58:101539, 2019. **1, 3, 4, 7**
- [11] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1691–1703. PMLR, 13–18 Jul 2020. **3**
- [12] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 13–18 Jul 2020. **1, 3, 4**
- [13] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners, 2020. **1, 3**
- [14] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning, 2020. **3, 4, 5, 6**
- [15] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15750–15758, June 2021. **1, 3, 4, 5, 6**
- [16] Terrance DeVries and Graham W. Taylor. Improved regularization of convolutional neural networks with cutout, 2017. **4**
- [17] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1422–1430, 2015. **2**
- [18] Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. **3**
- [19] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Alex Lamb, Martín Arjovsky, Olivier Mastropietro, and Aaron C. Courville. Adversarially learned inference. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. **3**
- [20] David G. Ellis and Michele R. Aizenberg. Trialing u-net training modifications for segmenting gliomas using open source deep learning framework. In Alessandro Crimi and Spyridon Bakas, editors, *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, pages 40–49, Cham, 2021. Springer International Publishing. **5**
- [21] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for self-supervised representation learning. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 3015–3024. PMLR, 18–24 Jul 2021. **1, 3, 4**
- [22] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv:1803.07728*, 2018. **2, 4**
- [23] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. **3**
- [24] Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch,

- Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21271–21284. Curran Associates, Inc., 2020. 1, 3, 4
- [25] Fatemeh Haghighi, Mohammad Reza Hosseinzadeh Taher, Zongwei Zhou, Michael B. Gotway, and Jianming Liang. Learning semantics-enriched representation via self-discovery, self-classification, and self-restoration. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, pages 137–147, Cham, 2020. Springer International Publishing. 1, 3
- [26] Fatemeh Haghighi, Mohammad Reza Hosseinzadeh Taher, Zongwei Zhou, Michael B. Gotway, and Jianming Liang. Transferable visual words: Exploiting the semantics of anatomical patterns for self-supervised learning. *IEEE Transactions on Medical Imaging*, 40(10):2857–2868, 2021. 1, 2, 3, 4, 5, 7, 8
- [27] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 3, 4
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 4
- [29] Mohammad Reza Hosseinzadeh Taher, Fatemeh Haghighi, Ruibin Feng, Michael B. Gotway, and Jianming Liang. A systematic benchmarking analysis of transfer learning for medical image analysis. In Shadi Albarqouni, M. Jorge Cardoso, Qi Dou, Konstantinos Kamnitsas, Bishesh Khanal, Islem Rekik, Nicola Rieke, Debodoot Sheet, Sotirios Tsafaris, Daguang Xu, and Ziyue Xu, editors, *Domain Adaptation and Representation Transfer, and Affordable Healthcare and AI for Resource Diverse Global Health*, pages 3–13, Cham, 2021. Springer International Publishing. 2, 3, 5, 7
- [30] Mohammad Reza Hosseinzadeh Taher, Fatemeh Haghighi, Michael B. Gotway, and Jianming Liang. CAiD: a self-supervised learning framework for empowering instance discrimination in medical imaging. *Proceedings of Machine Learning Research*, 2022. 3
- [31] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. *arXiv:1901.07031*, 2019. 5
- [32] Stefan Jaeger, Sema Candemir, Sameer Antani, Yi-Xiáng J Wáng, Pu-Xuan Lu, and George Thoma. Two public chest x-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative imaging in medicine and surgery*, 4(6), 2014. 5
- [33] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. In *CVPR*, 2017. 3
- [34] Junnan Li, Pan Zhou, Caiming Xiong, and Steven C. H. Hoi. Prototypical contrastive learning of unsupervised representations, 2021. 1, 3
- [35] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, pages 69–84. Springer, 2016. 2, 4
- [36] Gaurav Parmar, Dacheng Li, Kwonjoon Lee, and Zhuowen Tu. Dual contradistinctive generative autoencoder. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 823–832, June 2021. 3, 4
- [37] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544, 2016. 3, 4
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. 4
- [39] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 7
- [40] Arnaud Arindra Adiyoso Setio, Alberto Traverso, Thomas De Bel, Moira SN Berens, Cas van den Bogaard, Piergiorgio Cerello, Hao Chen, Qi Dou, Maria Evelina Fantacci, Bram Geurts, et al. Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the luna16 challenge. *Medical image analysis*, 42:1–13, 2017. 5
- [41] Nima Tajbakhsh, Michael B Gotway, and Jianming Liang. Computer-aided pulmonary embolism detection using a novel vessel-aligned multi-planar image representation and convolutional neural networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 62–69. Springer, 2015. 5
- [42] Aiham Taleb, Winfried Loetzsch, Noel Danz, Julius Severin, Thomas Gaertner, Benjamin Bergner, and Christoph Lippert. 3d self-supervised methods for medical imaging. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 18158–18172. Curran Associates, Inc., 2020. 7
- [43] Xing Tao, Yuexiang Li, Wenhui Zhou, Kai Ma, and Yefeng Zheng. Revisiting rubik’s cube: Self-supervised learning with volume-wise transformation for 3d medical image segmentation. In Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, Diana Mateus, Maria A. Zuluaga, S. Kevin Zhou, Daniel Racoceanu, and Leo Joskowicz, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, pages 238–248, Cham, 2020. Springer International Publishing. 1, 3, 8

- [44] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning?, 2020. [1](#), [3](#)
- [45] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, page 1096–1103. Association for Computing Machinery, 2008. [3](#)
- [46] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers. 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*, pages 2097–2106, 2017. [4](#), [7](#)
- [47] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3024–3033, June 2021. [3](#)
- [48] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [3](#)
- [49] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. [3](#)
- [50] Enze Xie, Jian Ding, Wenhai Wang, Xiaohang Zhan, Hang Xu, Peize Sun, Zhenguo Li, and Ping Luo. Detco: Unsupervised contrastive learning for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8392–8401, October 2021. [3](#)
- [51] Zhenda Xie, Yutong Lin, Zheng Zhang, Yue Cao, Stephen Lin, and Han Hu. Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16684–16693, June 2021. [3](#)
- [52] Mang Ye, Xu Zhang, Pong C. Yuen, and Shih-Fu Chang. Unsupervised embedding learning via invariant and spreading instance feature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [3](#)
- [53] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. *arXiv:2103.03230*, 2021. [1](#), [3](#), [4](#), [5](#), [6](#)
- [54] Xiaohang Zhan, Jiahao Xie, Ziwei Liu, Yew-Soon Ong, and Chen Change Loy. Online deep clustering for unsupervised representation learning. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [55] Hong-Yu Zhou, Chixiang Lu, Sibe Yang, Xiaoguang Han, and Yizhou Yu. Preservational learning improves self-supervised medical image models by reconstructing diverse contexts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3499–3509, October 2021. [1](#), [3](#)
- [56] Hong-Yu Zhou, Shuang Yu, Cheng Bian, Yifan Hu, Kai Ma, and Yefeng Zheng. Comparing to learn: Surpassing imagenet pretraining on radiographs by comparing image representations. In Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, Diana Mateus, Maria A. Zuluaga, S. Kevin Zhou, Daniel Racoceanu, and Leo Joskowicz, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, pages 398–407, Cham, 2020. Springer International Publishing. [1](#), [3](#)
- [57] Zongwei Zhou, Vatsal Sodha, Jiaxuan Pang, Michael B. Gotway, and Jianming Liang. Models genesis. *Medical Image Analysis*, 67:101840, 2021. [1](#), [3](#), [7](#)
- [58] Xinrui Zhuang, Yuexiang Li, Yifan Hu, Kai Ma, Yujiu Yang, and Yefeng Zheng. Self-supervised feature learning for 3d medical images by playing a rubik’s cube. In Dinggang Shen, Tianming Liu, Terry M. Peters, Lawrence H. Staib, Caroline Essert, Sean Zhou, Pew-Thian Yap, and Ali Khan, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, pages 420–428, Cham, 2019. Springer International Publishing. [4](#), [7](#)