

Big Self-Supervised Models Advance Medical Image Classification

Shekoofeh Azizi, Basil Mustafa, Fiona Ryan*, Zachary Beaver, Jan Freyberg, Jonathan Deaton, Aaron Loh, Alan Karthikesalingam, Simon Kornblith, Ting Chen, Vivek Natarajan, Mohammad Norouzi

Google Research and Health[†]

Abstract

Self-supervised pretraining followed by supervised fine-tuning has seen success in image recognition, especially when labeled examples are scarce, but has received limited attention in medical image analysis. This paper studies the effectiveness of self-supervised learning as a pre-training strategy for medical image classification. We conduct experiments on two distinct tasks: dermatology condition classification from digital camera images and multi-label chest X-ray classification, and demonstrate that self-supervised learning on ImageNet, followed by additional self-supervised learning on unlabeled domain-specific medical images significantly improves the accuracy of medical image classifiers. We introduce a novel Multi-Instance Contrastive Learning (MICLe) method that uses multiple images of the underlying pathology per patient case, when available, to construct more informative positive pairs for self-supervised learning. Combining our contributions, we achieve an improvement of 6.7% in top-1 accuracy and an improvement of 1.1% in mean AUC on dermatology and chest X-ray classification respectively, outperforming strong supervised baselines pretrained on ImageNet. In addition, we show that big self-supervised models are robust to distribution shift and can learn efficiently with a small number of labeled medical images.

1. Introduction

Learning from limited labeled data is a fundamental problem in machine learning, which is crucial for medical image analysis because annotating medical images is time-consuming and expensive. Two common pretraining approaches to learning from limited labeled data include: (1) *supervised pretraining* on a large labeled dataset such as ImageNet, (2) *self-supervised pretraining* using contrastive learning (e.g., [16, 8, 9]) on unlabeled data. After pretraining, supervised fine-tuning on a target labeled dataset of interest is used. While ImageNet pretraining is ubiquitous in medical image analysis [46, 32, 31, 29, 15, 20], the use of self-supervised approaches has received limited attention. Self-supervised approaches are attractive because they en-

(1) Self-supervised learning on **unlabeled** natural images



(2) Self-supervised learning on **unlabeled** medical images and **Multi-Instance Contrastive Learning (MICLe)** if multiple images of each medical condition are available



(3) Supervised fine-tuning on **labeled** medical images

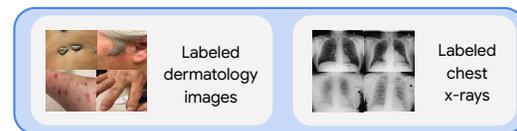


Figure 1: Our approach comprises three steps: (1) Self-supervised pretraining on unlabeled ImageNet using SimCLR [8]. (2) Additional self-supervised pretraining using unlabeled medical images. If multiple images of each medical condition are available, a novel Multi-Instance Contrastive Learning (MICLe) is used to construct more informative positive pairs based on different images. (3) Supervised fine-tuning on labeled medical images. Note that unlike step (1), steps (2) and (3) are task and dataset specific.

able the use of *unlabeled* domain-specific images during pretraining to learn more relevant representations.

This paper studies self-supervised learning for medical image analysis and conducts a fair comparison between self-supervised and supervised pretraining on two distinct medical image classification tasks: (1) dermatology skin condition classification from digital camera images, (2) multi-label chest X-ray classification among five pathologies based on the CheXpert dataset [23]. We observe that self-supervised pretraining outperforms supervised pretraining, even when the full ImageNet dataset (14M images and 21.8K classes) is used for supervised pretraining. We attribute this finding to the domain shift and discrepancy between the nature of recognition tasks in Im-

*Former intern at Google. Currently at Georgia Institute of Technology.

[†]{shekazizi, skornblith, iamtingchen, natviv, mnorouzi}@google.com

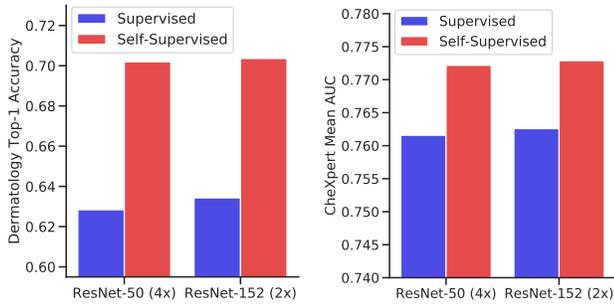


Figure 2: Comparison of supervised and self-supervised pretraining, followed by supervised fine-tuning using two architectures on dermatology and chest X-ray classification. Self-supervised learning utilizes unlabeled domain-specific medical images and significantly outperforms supervised ImageNet pretraining.

ImageNet and medical image classification. Self-supervised approaches bridge this domain gap by leveraging in-domain medical data for pretraining and they also scale gracefully as they do not require any form of class label annotation.

An important component of our self-supervised learning framework is an effective *Multi-Instance Contrastive Learning (MICLe)* strategy that helps adapt contrastive learning to multiple images of the underlying pathology per patient case. Such multi-instance data is often available in medical imaging datasets – *e.g.*, frontal and lateral views of mammograms, retinal fundus images from each eye, *etc.* Given multiple images of a given patient case, we propose to construct a positive pair for self-supervised contrastive learning by drawing two crops from two distinct images of the same patient case. Such images may be taken from different viewing angles and show different body parts with the same underlying pathology. This presents a great opportunity for self-supervised learning algorithms to learn representations that are robust to changes of viewpoint, imaging conditions, and other confounding factors in a direct way. MICLe does not require class label information and only relies on different images of an underlying pathology, the type of which may be unknown.

Fig. 1 depicts the proposed self-supervised learning approach, and Fig. 2 shows the summary of results. Our key findings and contributions include:

- We investigate the use of self-supervised pretraining on medical image classification. We find that self-supervised pretraining on unlabeled medical images significantly outperforms standard ImageNet pretraining and random initialization.
- We propose Multi-Instance Contrastive Learning (MICLe) as a generalization of existing contrastive learning approaches to leverage multiple images per medical condition. We find that MICLe improves the performance of self-supervised models, yielding state-of-the-art results.
- On dermatology condition classification, our self-supervised approach provides a sizable gain of 6.7% in

top-1 accuracy, even in a highly competitive production setting. On chest X-ray classification, self-supervised learning outperforms strong supervised baselines pre-trained on ImageNet by 1.1% in mean AUC.

- We demonstrate that self-supervised models are robust and generalize better than baselines, when subjected to shifted test sets, without fine-tuning. Such behavior is desirable for deployment in a real-world clinical setting.

2. Related Work

Transfer Learning for Medical Image Analysis. Despite the differences in image statistics, scale, and task-relevant features, transfer learning from natural images is commonly used in medical image analysis [29, 31, 32, 46], and multiple empirical studies show that this improves performance [1, 15, 20]. However, detailed investigations from Raghu *et al.* [37] of this strategy indicate this does not always improve performance in medical imaging contexts. They, however, do show that transfer learning from ImageNet can speed up convergence, and is particularly helpful when the medical image training data is limited. Importantly, the study used relatively small architectures, and found pronounced improvements with small amounts of data especially when using their largest architecture of ResNet-50 (1×) [18]. Transfer learning from in-domain data can help alleviate the domain mismatch issue. For example, [7, 20, 26, 13] report performance improvements when pretraining on labeled data in the same domain. However, this approach is often infeasible for many medical tasks in which labeled data is expensive and time-consuming to obtain. Recent advances in self-supervised learning provide a promising alternative enabling the use of unlabeled medical data that is often easier to procure.

Self-supervised Learning. Initial works in self-supervised representation learning focused on the problem of learning embeddings without labels such that a low-capacity (commonly linear) classifier operating on these embeddings could achieve high classification accuracy [12, 14, 35, 49]. *Contrastive* self-supervised methods such as instance discrimination [45], CPC [21, 36], Deep InfoMax [22], Ye *et al.* [47], AMDIM [2], CMC [41], MoCo [10, 17], PIRL [33], and SimCLR [8, 9] were the first to achieve linear classification accuracy approaching that of end-to-end supervised training. Recently, these methods have been harnessed to achieve dramatic improvements in label efficiency for semi-supervised learning. Specifically, one can first pretrain in a task-agnostic, self-supervised fashion using all data, and then fine-tune on the labeled subset in a task-specific fashion with a standard supervised objective [8, 9, 21]. Chen *et al.* [9] show that this approach benefits from large (high-capacity) models for pretraining and fine-tuning, but after a large model is trained, it can be distilled to a much smaller model with little loss in accuracy.

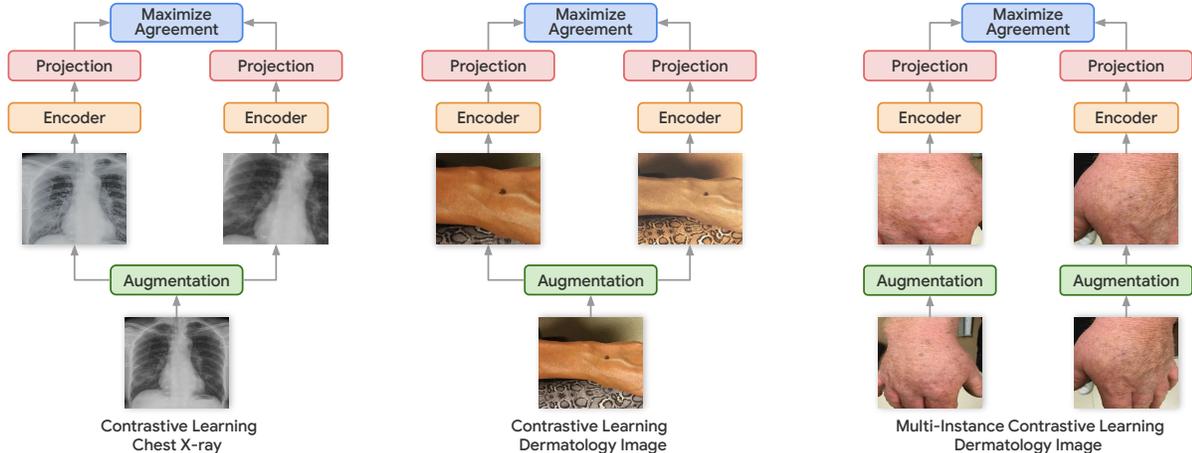


Figure 3: An illustrations of our self-supervised pretraining for medical image analysis. When a single image of a medical condition is available, we use standard data augmentation to generate two augmented views of the same image. When multiple images are available, we use two distinct images to directly create a positive pair of examples. We call the latter approach Multi-Instance Contrastive Learning (MICLe).

Our Multi-Instance Contrastive Learning approach is related to previous works in video processing where multiple views naturally arising due to temporal variation [38, 42]. These works have proposed to learn visual representations from video by maximizing agreement between representations of adjacent frames [42] or two views of the same action [38]. We generalize this idea to representation learning from image datasets, when sets of images containing the same desired class information are available and we show that benefits of MICLe can be combined with state-of-the-art self-supervised learning methods such as SimCLR.

Self-supervision for Medical Image Analysis. Although self-supervised learning has only recently become viable on standard image classification datasets, it has already seen some application within the medical domain. While some works have attempted to design domain-specific pretext tasks [3, 40, 53, 52], other works concentrate on tailoring contrastive learning to medical data [5, 19, 25, 27, 51]. Most closely related to our work, Sowrirajan *et al.* [39] explore the use of MoCo pretraining for classification of CheXpert dataset through linear evaluation.

Several recent publications investigate semi-supervised learning for medical imaging tasks (*e.g.*, [11, 28, 43, 50]). These methods are complementary to ours, and we believe combining self-training and self-supervised pretraining is an interesting avenue for future research (*e.g.*, [9]).

3. Self-Supervised Pretraining

Our approach comprises the following steps. First, we perform self-supervised pretraining on unlabeled images using contrastive learning to learn visual representations. For contrastive learning we use a combination of unlabeled ImageNet dataset and task specific medical images. Then, if multiple images of each medical condition are available the

Multi-Instance Contrastive Learning (MICLe) is used for additional self-supervised pretraining. Finally, we perform supervised fine-tuning on labeled medical images. Figure 1 shows the summary of our proposed method.

3.1. A Simple Framework for Contrastive Learning

To learn visual representations effectively with unlabeled images, we adopt SimCLR [8, 9], a recently proposed approach based on contrastive learning. SimCLR learns representations by maximizing agreement [4] between differently augmented views of the same data example via a contrastive loss in a hidden representation of neural nets.

Given a randomly sampled mini-batch of images, each image x_i is augmented twice using random crop, color distortion and Gaussian blur, creating two views of the same example x_{2k-1} and x_{2k} . The two images are encoded via an encoder network $f(\cdot)$ (a ResNet [18]) to generate representations h_{2k-1} and h_{2k} . The representations are then transformed again with a non-linear transformation network $g(\cdot)$ (a MLP projection head), yielding z_{2k-1} and z_{2k} that are used for the contrastive loss.

With a mini-batch of encoded examples, the contrastive loss between a pair of positive example i, j (augmented from the same image) is given as follows:

$$\ell_{i,j}^{\text{NT-Xent}} = -\log \frac{\exp(\text{sim}(z_i, z_j)/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(\text{sim}(z_i, z_k)/\tau)}, \quad (1)$$

Where $\text{sim}(\cdot, \cdot)$ is cosine similarity between two vectors, and τ is a temperature scalar.

3.2. Multi-Instance Contrastive Learning (MICLe)

In medical image analysis, it is common to utilize multiple images per patient to improve classification accuracy

and robustness. Such images may be taken from different viewpoints or under different lighting conditions, providing complementary information for medical diagnosis. When multiple images of a medical condition are available as part of the training dataset, we propose to learn representations that are invariant not only to different augmentations of the same image, but also to different images of the same medical pathology. Accordingly, we can conduct a multi-instance contrastive learning (MICLe) stage where positive pairs are constructed by drawing two crops from the images of the same patient as demonstrated in Fig. 3.

In MICLe, in contrast to standard SimCLR, to construct a mini-batch of $2N$ representation, we randomly sample a mini-batch of N bags of instances and define the contrastive prediction task on positive pairs retrieved from the bag of images instead of augmented views of the same image. Each bag, $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M\}$ contains images from a same patient (*i.e.*, same pathology) captured from different views and we assume that M could vary for different bags. When there is two or more instances in a bag ($M = |X| \geq 2$), we construct positive pairs by drawing two crops from two randomly selected images in the bag. In this case, the objective still takes the form of Eq. (1), but images contributing to each positive pair are distinct. Algorithm 1 summarizes the proposed method.

Leveraging multiple images of the same condition using the contrastive loss helps the model learn representations that are more robust to the change of viewpoint, lighting conditions, and other confounding factors. We find that multi-instance contrastive learning significantly improves the accuracy and helps us achieve the state-of-the-art result on the dermatology condition classification task.

4. Experiment Setup

4.1. Tasks and datasets

We consider two popular medical imaging tasks. The first task is in the dermatology domain and involves identifying skin conditions from digital camera images. The second task involves multi-label classification of chest X-rays among five pathologies. We chose these tasks as they embody many common characteristics of medical imaging tasks like imbalanced data and pathologies of interest restricted to small local patches. At the same time, they are also quite diverse in terms of the type of images, label space and task setup. For example, dermatology images are visually similar to natural images whereas the chest X-rays are gray-scale and have standardized views. This, in turn, helps us probe the generality of our proposed methods.

Dermatology. For the dermatology task, we follow the experiment setup and dataset of [29]. The dataset was collected and de-identified by a US based tele-dermatology service with images of skin conditions taken using consumer grade digital cameras. The images are heteroge-

Algorithm 1: Multi-Instance Contrastive Learning.

Input: batch size N , constant τ , $g(\cdot)$, $f(\cdot)$, \mathcal{T}
while stopping criteria not met **do**
 Sample mini-batch of $\{X\}_{k=1}^N$ **for** $k \leftarrow 1$ **to** $k = N$
 do
 Draw augmentation functions t and $t' \sim \mathcal{T}$;
 if $|X_k| \geq 2$ **then**
 Randomly select \mathbf{x}_k and $\mathbf{x}'_k \in X_k$;
 else
 $\mathbf{x}_k = \mathbf{x}'_k \leftarrow$ the only element of X_k ;
 end
 $\tilde{\mathbf{x}}_{2k-1} = t(\mathbf{x}_k)$; $\tilde{\mathbf{x}}_{2k} = t'(\mathbf{x}'_k)$;
 $\mathbf{z}_{2k-1} = g(f(\tilde{\mathbf{x}}_{2k-1}))$; $\mathbf{z}_{2k} = g(f(\tilde{\mathbf{x}}_{2k}))$;
 end
 for $i \in \{1, \dots, 2N\}$ and $j \in \{1, \dots, 2N\}$ **do**
 $s_{i,j} = \mathbf{z}_i^\top \mathbf{z}_j / (\|\mathbf{z}_i\| \|\mathbf{z}_j\|)$;
 end
 $\ell(i, j) \leftarrow \ell_{i,j}^{\text{NT-Xent}}$ in Eq. (1);
 $\mathcal{L} = \frac{1}{2N} \sum_{k=1}^N [\ell(2k-1, 2k) + \ell(2k, 2k-1)]$;
end
return Trained encoder network $f(\cdot)$

neous in nature and exhibit significant variations in terms of the pose, lighting, blur, and body parts. The background also contains various noise artifacts like clothing and walls which adds to the challenge. The ground truth labels were aggregated from a panel of several US-board certified dermatologists who provided differential diagnosis of skin conditions in each case.

In all, the dataset has cases from a total of 12,306 unique patients. Each case includes between one to six images. This further split into development and test sets ensuring no patient overlap between the two. Then, cases with the occurrence of multiple skin conditions or poor quality images were filtered out. The final $\mathcal{D}_{\text{Derm}}^{\text{Train}}$, $\mathcal{D}_{\text{Derm}}^{\text{Validation}}$, and $\mathcal{D}_{\text{Derm}}^{\text{Test}}$ include a total of 15,340 cases, 1,190 cases, and 4,146 cases, respectively. There are 419 unique condition labels in the dataset. For the purpose of model development, we identified and use the most common 26 skin conditions and group the rest in an additional ‘Other’ class leading to a final label space of 27 classes for the model. We refer to this as $\mathcal{D}_{\text{Derm}}$ in the subsequent sections. We also use an additional de-identified $\mathcal{D}_{\text{Derm}}^{\text{External}}$ set to evaluate the generalization performance of our proposed method under distribution shift. Unlike $\mathcal{D}_{\text{Derm}}$, this dataset is primarily focused on skin cancers and the ground truth labels are obtained from biopsies. The distribution shift in the labels make this a particular challenging data to evaluate the zero-shot (*i.e.* without any additional fine-tuning) transfer performance of models.

For SimCLR pretraining, we combine the images from $\mathcal{D}_{\text{Derm}}^{\text{Train}}$ and additional unlabeled images from the same source leading to a total of 454,295 images for self-supervised pretraining. We refer to this as the $\mathcal{D}_{\text{Derm}}^{\text{Unlabeled}}$. For MICLe pretraining, we only use images coming from

the 15,340 cases of the $\mathcal{D}_{\text{Derm}}^{\text{Train}}$. Additional details are provided in the Appendix A.1.

Chest X-rays. CheXpert [23] is a large open source dataset of de-identified chest radiograph (X-ray) images. The dataset consists of a set of 224,316 chest radiographs coming from 65,240 unique patients. The ground truth labels were automatically extracted from radiology reports and correspond to a label space of 14 radiological observations. The validation set consists of 234 manually annotated chest X-rays. Given the small size of the validation dataset and following [34, 37] suggestion, for the downstream task evaluations we randomly re-split the training set into 67,429 training images, 22,240 validation images, and 33,745 test images. We train the model to predict the five pathologies used by Irvin and Rajpurkar *et al.* [23] in a multi-label classification task setting. For SimCLR pretraining for the chest X-ray domain, we only consider images coming from the train set of the CheXpert dataset discarding the labels. We refer to this as the $\mathcal{D}_{\text{CheXpert}}^{\text{Unlabeled}}$. In addition, we also use the NIH chest X-ray dataset, \mathcal{D}_{NIH} , to evaluate the zero-shot transfer performance which consist of 112,120 de-identified X-rays from 30,805 unique patients. Additional details on the dataset can be found here [44] and also are provided in the Appendix A.2.

4.2. Pretraining protocol

To assess the effectiveness of self-supervised pretraining using big neural nets, as suggested in [8], we investigate ResNet-50 (1 \times), ResNet-50 (4 \times), and ResNet-152 (2 \times) architectures as our base encoder networks. Following SimCLR [8], two fully connected layers are used to map the output of ResNets to a 128-dimensional embedding, which is used for contrastive learning. We also use LARS optimizer [48] to stabilize training during pretraining. We perform SimCLR pretraining on $\mathcal{D}_{\text{Derm}}^{\text{Unlabeled}}$ and $\mathcal{D}_{\text{CheXpert}}^{\text{Unlabeled}}$, both with and without initialization from ImageNet self-supervised pretrained weights. We indicate pretraining initialized using self-supervised ImageNet weights, as ImageNet \rightarrow Derm, and ImageNet \rightarrow CheXpert in the following sections.

Unless otherwise specified, for the dermatology pretraining task, due to similarity of dermatology images to natural images, we use the same data augmentation used to generate positive pairs in SimCLR. This includes random color augmentation (strength=1.0), crops with resize, Gaussian blur, and random flips. We find that the batch size of 512 and learning rate of 0.3 works well in this setting. Using this protocol, all of models were pretrained up to 150,000 steps using $\mathcal{D}_{\text{Derm}}^{\text{Unlabeled}}$.

For the CheXpert dataset, we pretrain with learning rate in {0.5, 1.0, 1.5}, temperature in {0.1, 0.5, 1.0}, and batch size in {512, 1024}, and we select the model with best performance on the down-stream validation set. We also tested

a range of possible augmentations and observe that the augmentations that lead to the best performance on the validation set for this task are random cropping, random color jittering (strength = 0.5), rotation (upto 45 degrees) and horizontal flipping. Unlike the original set of proposed augmentation in SimCLR, we do not use the Gaussian blur, because we think it can make it impossible to distinguish local texture variations and other areas of interest thereby changing the underlying disease interpretation the X-ray image. We leave comprehensive investigation of the optimal augmentations to future work. Our best model on CheXpert was pretrained with batch size 1024, and learning rate of 0.5 and we pretrain the models up to 100,000 steps.

We perform MICLE pretraining only on the dermatology dataset as we did not have enough cases with the presence of multiple views in the CheXpert dataset to allow comprehensive training and evaluation of this approach. For MICLE pretraining we initialize our model using SimCLR pretrained weights, and then incorporate the multi-instance procedure as explained in Section 3.2 to further learn a more comprehensive representation using multi-instance data of $\mathcal{D}_{\text{Derm}}^{\text{Train}}$. Due to memory limits caused by stacking up to 6 images per patient case, we train with a smaller batch size of 128 and learning rate of 0.1 for 100,000 steps to stabilize the training. Decreasing the learning rate for smaller batch size has been suggested in [8]. The rest of the settings, including optimizer, weight decay, and warmup step are the same as our previous pretraining protocol.

In all of our pretraining experiments, images are resized to 224 \times 224. We use 16 to 64 Cloud TPU cores depending on the batch size for pretraining. With 64 TPU cores, it takes \sim 12 hours to pretrain a ResNet-50 (1 \times) with batch size 512 and for 100 epochs. Additional details about the selection of batch size and learning rate, and augmentations are provided in the Appendix B.

4.3. Fine-tuning protocol

We train the model end-to-end during fine-tuning using the weights of the pretrained network as initialization for the downstream supervised task dataset following the approach described by Chen *et al.* [8, 9] for all our experiments. We trained for 30,000 steps with a batch size of 256 using SGD with a momentum parameter of 0.9. For data augmentation during fine-tuning, we performed random color augmentation, crops with resize, blurring, rotation, and flips for the images in both tasks. We observe that this set of augmentations is critical for achieving the best performance during fine-tuning. We resize the Derm dataset images to 448 \times 448 pixels and CheXpert images to 224 \times 224 during this fine-tuning stage.

For every combination of pretraining strategy and downstream fine-tuning task, we perform an extensive hyperparameter search. We selected the learning rate and weight decay after a grid search of seven logarithmically spaced

Table 1: Performance of dermatology skin condition and Chest X-ray classification model measured by top-1 accuracy (%) and area under the curve (AUC) across different architectures. Each model is fine-tuned using transfer learning from pretrained model on ImageNet, only unlabeled medical data, or pretrained using medical data initialized from ImageNet pretrained model (e.g. ImageNet → Derm). Bigger models yield better performance. pretraining on ImageNet is complementary to pretraining on unlabeled medical images.

Architecture	Dermatology Classification			Chest X-ray Classification	
	Pretraining Dataset	Top-1 Accuracy(%)	AUC	Pretraining Dataset	Mean AUC
ResNet-50 (1×)	ImageNet	62.58 ± 0.84	0.9480 ± 0.0014	ImageNet	0.7630 ± 0.0013
	Derm	63.66 ± 0.24	0.9490 ± 0.0011	CheXpert	0.7647 ± 0.0007
	ImageNet→Derm	63.44 ± 0.13	0.9511 ± 0.0037	ImageNet→CheXpert	0.7670 ± 0.0007
ResNet-50 (4×)	ImageNet	64.62 ± 0.76	0.9545 ± 0.0007	ImageNet	0.7681 ± 0.0008
	Derm	66.93 ± 0.92	0.9576 ± 0.0015	CheXpert	0.7668 ± 0.0011
	ImageNet→Derm	67.63 ± 0.32	0.9592 ± 0.0004	ImageNet→CheXpert	0.7687 ± 0.0016
ResNet-152 (2×)	ImageNet	66.38 ± 0.03	0.9573 ± 0.0023	ImageNet	0.7671 ± 0.0008
	Derm	66.43 ± 0.62	0.9558 ± 0.0007	CheXpert	0.7683 ± 0.0009
	ImageNet→Derm	68.30 ± 0.19	0.9620 ± 0.0007	ImageNet→CheXpert	0.7689 ± 0.0010

learning rates between $10^{-3.5}$ and $10^{-0.5}$ and three logarithmically spaced values of weight decay between 10^{-5} and 10^{-3} , as well as no weight decay. For training from the supervised pretraining baseline we follow the same protocol and observe that for all fine-tuning setups, 30,000 steps is sufficient to achieve optimal performance. For supervised baselines we compare against the identical publicly available ResNet models¹ pretrained on ImageNet with standard cross-entropy loss. These models are trained with the same data augmentation as self-supervised models (crops, strong color augmentation, and blur).

4.4. Evaluation methodology

After identifying the best hyperparameters for fine-tuning a given dataset, we proceed to select the model based on validation set performance and evaluate the chosen model multiple times (10 times for chest X-ray task and 5 times for the dermatology task) on the test set to report task performance. Our primary metrics for the dermatology task are top-1 accuracy and Area Under the Curve (AUC) following [29]. For the chest X-ray task, given the multi-label setup, we report mean AUC averaged between the predictions for the five target pathologies following [23]. We also use the non-parametric bootstrap to estimate the variability around the model performance and investigating any statistically significant improvement. Additional details are provided in Appendix B.1.1.

5. Experiments & Results

In this section we investigate whether self-supervised pretraining with contrastive learning translates to a better performance in models fine-tuned end-to-end across the selected medical image classification tasks. To this end, first, we explore the choice of the pretraining dataset for medical imaging tasks. Then, we evaluate the benefits of our proposed multi-instance contrastive learning (MICLE) for

dermatology condition classification task, and compare and contrast the proposed method against the baselines and state of the art methods for supervised pretraining. Finally, we explore label efficiency and transferability (under distribution shift) of self-supervised trained models in the medical image classification setting.

5.1. Dataset for pretraining

One important aspect of transfer learning via self-supervised pretraining is the choice of a proper unlabeled dataset. For this study, we use architectures of varying capacities (i.e ResNet-50 (1×), ResNet-50 (4×) and ResNet-152 (2×) as our base network, and carefully investigate three possible scenario for self-supervised pretraining in the medical context: (1) using ImageNet dataset only, (2) using the task specific unlabeled medical dataset (i.e. Derm and CheXpert), and (3) initializing the pretraining from ImageNet self-supervised model but using task specific unlabeled dataset for pretraining, here indicated as ImageNet → CheXpert and ImageNet → CheXpert. Table 1 shows the performance of dermatology skin condition and chest X-ray classification model measured by top-1 accuracy (%) and area under the curve (AUC) across different architectures and pretraining scenarios. Our results suggest that, best performance are achieved when both ImageNet and task specific unlabeled data are used. Combining ImageNet and Derm unlabeled data for pretraining, translates to $(1.92 \pm 0.16)\%$ increase in top-1 accuracy for dermatology classification over only using ImageNet dataset for self-supervised transfer learning. This results suggests that pretraining on ImageNet is likely complementary to pretraining on unlabeled medical images. Moreover, we observe that larger models are able to benefit much more from self-supervised pretraining underscoring the importance of model capacity in this setting.

As shown in Table 1, on CheXpert, we once again observe that self-supervised pretraining with both ImageNet

¹<https://github.com/google-research/simclr>

Table 2: Evaluation of multi instance contrastive learning (MICLe) on Dermatology condition classification. Our results suggest that MICLe consistently improves the accuracy of skin condition classification over SimCLR on different datasets and architectures.

Model	Dataset	MICLe	Top-1 Accuracy
ResNet-50 (4×)	Derm	No	66.93±0.92
	Derm	Yes	67.55±0.52
	ImageNet→Derm	No	67.63±0.32
	ImageNet→Derm	Yes	68.81±0.41
ResNet-152 (2×)	Derm	No	66.43±0.62
	Derm	Yes	67.16±0.35
	ImageNet→Derm	No	68.30±0.19
	ImageNet→Derm	Yes	68.43±0.32

and in-domain CheXpert data is beneficial, outperforming self-supervised pretraining on ImageNet or CheXpert alone.

5.2. Performance of MICLe

Next, we evaluate whether utilizing multi-instance contrastive learning (MICLe) and leveraging the potential availability of multiple images per patient for a given pathology, is beneficial for self-supervised pretraining. Table 2 compares the performance of dermatology condition classification models fine-tuned on representations learned with and without MICLe pretraining. We observe that MICLe consistently improves the performance of dermatology classification over the original SimCLR method under different pretraining dataset and base network architecture choices. Using MICLe for pretraining, translates to $(1.18 \pm 0.09)\%$ increase in top-1 accuracy for dermatology classification over using only original SimCLR.

5.3. Comparison with supervised transfer learning

We further improves the performance by providing more negative examples with training longer for 1000 epochs and a larger batch size of 1024. We achieve the best-performing top-1 accuracy of $(70.02 \pm 0.22)\%$ using the ResNet-152 (2×) architecture and MICLe pretraining by incorporating both ImageNet and Derm dataset in dermatology condition classification. Tables 3 and 4 show the comparison of transfer learning performance of SimCLR and MICLe models with supervised baselines for the dermatology and the chest X-ray classification. This result shows that after fine-tuning, our self-supervised model significantly outperforms the supervised baseline when ImageNet pretraining is used ($p < 0.05$). We specifically observe an improvement of over 6.7% in top-1 accuracy in the dermatology task when using MICLe. On the chest X-ray task, the improvement is 1.1% in mean AUC without using MICLe.

Though using ImageNet pretrained models is still the norm, recent advances have been made by supervised pretraining on large scale (often noisy) natural datasets [24, 30] improving transfer performance on downstream tasks.

Table 3: Comparison of best self-supervised models vs. supervised pretraining baselines on dermatology classification.

Architecture	Method	Pretraining Dataset	Top-1 Accuracy
ResNet-152 (2×)	Supervised	ImageNet	63.36 ± 0.12
ResNet-101 (3×)	BiT [24]	ImageNet-21k	68.45 ± 0.29
ResNet-152 (2×)	SimCLR	ImageNet	66.38 ± 0.03
ResNet-152 (2×)	SimCLR	ImageNet→Derm	69.43 ± 0.43
ResNet-152 (2×)	MICLe	ImageNet→Derm	70.02 ± 0.22

Table 4: Comparison of best self-supervised models vs. supervised pretraining baselines on chest X-ray classification.

Architecture	Method	Pretraining Dataset	Mean AUC
ResNet-152 (2×)	Supervised	ImageNet	0.7625 ± 0.001
ResNet-101 (3×)	BiT [24]	ImageNet-21k	0.7720 ± 0.002
ResNet-152 (2×)	SimCLR	ImageNet	0.7671 ± 0.008
ResNet-152 (2×)	SimCLR	CheXpert	0.7702 ± 0.001
ResNet-152 (2×)	SimCLR	ImageNet→CheXpert	0.7729 ± 0.001

We therefore also evaluate a supervised baseline from Kolesnikov *et al.* [24], a ResNet-101 (3×) pretrained on ImageNet21-k called Big Transfer (BiT). This model contains additional architectural tweaks included to boost transfer performance, and was trained on a significantly larger dataset (14M images labelled with one or more of 21k classes, vs. the 1M images in ImageNet) which provides us with a strong supervised baseline². ResNet-101 (3×) has 382M trainable parameters, thus comparable to ResNet-152 (2×) with 233M trainable parameters. We observe that the MICLe model is better than this BiT model for the dermatology classification task improving by 1.6% in top-1 accuracy. For the chest X-ray task, self supervised model is better by about 0.1% mean AUC. We surmise that with additional in-domain unlabeled data (we only use the CheXpert dataset for pretraining), self-supervised pretraining can surpass the BiT baseline by a larger margin. At the same time, these two approaches are complementary but we leave further explorations in this direction to future work.

5.4. Self-supervised models generalize better

We conduct further experiments to evaluate the robustness of self-supervised pretrained models to distribution shifts. For this purpose, we use the model post pretraining and end-to-end fine-tuning (i.e. CheXpert and Derm) to make predictions on an additional shifted dataset without any further fine-tuning (zero-shot transfer learning). We use the $\mathcal{D}_{Derm}^{External}$ and \mathcal{D}_{NIH} as our target shifted datasets. Our results generally suggest that self-supervised pretrained models can generalize better to distribution shifts.

For the chest X-ray task, we note that self-supervised pretraining with either ImageNet or CheXpert data im-

²This model is also available publicly at https://github.com/google-research/big_transfer

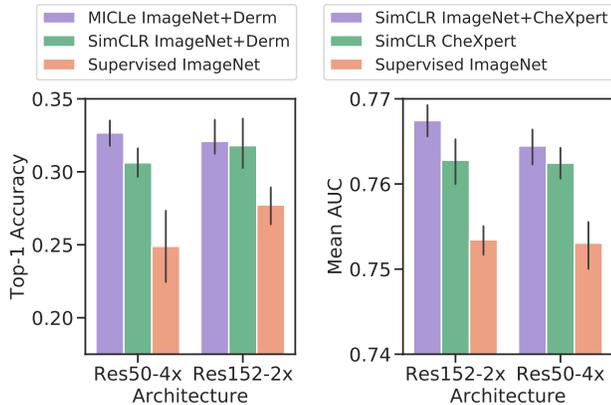


Figure 4: Evaluation of models on distribution-shifted datasets (left: $\mathcal{D}_{\text{Derm}}^{\text{Unlabeled}} \rightarrow \mathcal{D}_{\text{Derm}}^{\text{External}}$, right: $\mathcal{D}_{\text{CheXpert}}^{\text{Unlabeled}} \rightarrow \mathcal{D}_{\text{NIH}}$) shows that self-supervised training using both ImageNet and the target domain significantly improves robustness to distribution shift.

proves generalisation, but stacking them both yields further gains. We also note that when only using ImageNet for self-supervised pretraining, the model performs worse in this setting compared to using in-domain data for pretraining.

Further we find that the performance improvement in the distribution-shifted dataset due to self-supervised pretraining (both using ImageNet and CheXpert data) is more pronounced than the original improvement on the CheXpert dataset. This is a very valuable finding, as generalisation under distribution shift is of paramount importance to clinical applications. On the dermatology task, we observe similar trends suggesting the robustness of the self-supervised representations is consistent across tasks.

5.5. Self-supervised models are more label-efficient

To investigate label-efficiency of the selected self-supervised models, following the previously explained fine-tuning protocol, we fine-tune our models on different fractions of labeled training data. We also conduct baseline fine-tuning experiments with supervised ImageNet pretrained models. We use the label fractions ranging from 10% to 90% for both Derm and CheXpert training datasets. Fine-tuning experiments on label fractions are repeated multiple times using the best parameters and averaged. Figure 4 shows how the performance varies using the different available label fractions for the dermatology task. First, we observe that pretraining using self-supervised models can significantly help with label efficiency for medical image classification, and in all of the fractions, self-supervised models outperform the supervised baseline. Moreover, these results suggest that MICLe yields proportionally larger gains when fine-tuning with fewer labeled examples. In fact, MICLe is able to match baselines using only 20% of the training data for ResNet-50 (4x) and 30% of the training data for ResNet-152 (2x). Results on the CheXpert dataset are included in Appendix B.2, where we observe a similar trend.

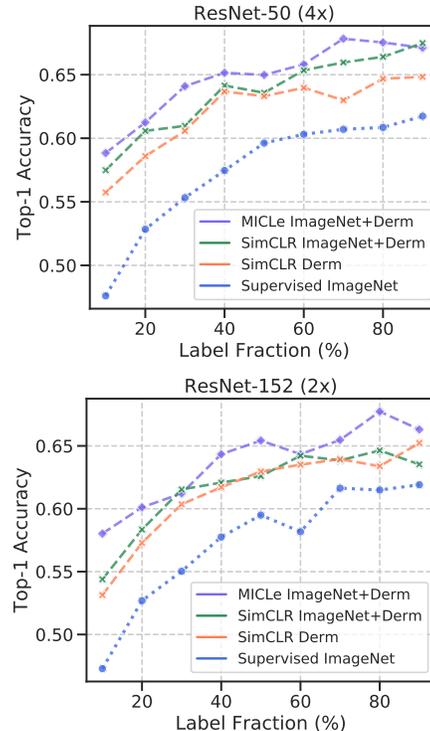


Figure 5: Top-1 accuracy for dermatology condition classification for MICLe, SimCLR, and supervised models under different unlabeled pretraining dataset and varied sizes of label fractions.

6. Conclusion

Supervised pretraining on natural image datasets such as ImageNet is commonly used to improve medical image classification. This paper investigates an alternative strategy based on self-supervised pretraining on unlabeled natural and medical images and finds that self-supervised pretraining significantly outperforms supervised pretraining. The paper proposes the use of multiple images per medical case to enhance data augmentation for self-supervised learning, which boosts the performance of image classifiers even further. Self-supervised pretraining is much more scalable than supervised pretraining since class label annotation is not required. A natural next step for this line of research is to investigate the limit of self-supervised pretraining by considering massive unlabeled medical image datasets. Another research direction concerns the transfer of self-supervised learning from one imaging modality and task to another. We hope this paper will help popularize the use of self-supervised approaches in medical image analysis yielding improvements in label efficiency across the medical field.

Acknowledgement

We would like to thank Yuan Liu for valuable feedback on the manuscript. We are also grateful to Jim Winkens, Megan Wilson, Umesh Telang, Patricia Macwilliams, Greg Corrado, Dale Webster, and our collaborators at DermPath AI for their support of this work.

References

- [1] Laith Alzubaidi, Mohammed A Fadhel, Omran Al-Shamma, Jinglan Zhang, J Santamaría, Ye Duan, and Sameer R Oleiwi. Towards a better understanding of transfer learning for medical imaging: a case study. *Applied Sciences*, 10(13):4523, 2020. [2](#)
- [2] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pages 15535–15545, 2019. [2](#)
- [3] Wenjia Bai, Chen Chen, Giacomo Tarroni, Jinming Duan, Florian Guitton, Steffen E Petersen, Yike Guo, Paul M Matthews, and Daniel Rueckert. Self-supervised learning for cardiac MR image segmentation by anatomical position prediction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 541–549. Springer, 2019. [3](#)
- [4] Suzanna Becker and Geoffrey E Hinton. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 355(6356):161–163, 1992. [3](#)
- [5] Krishna Chaitanya, Ertunc Erdil, Neerav Karani, and Ender Konukoglu. Contrastive learning of global and local features for medical image segmentation with limited annotations. *arXiv preprint arXiv:2006.10511*, 2020. [3](#)
- [6] 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. [14](#), [16](#)
- [7] Sihong Chen, Kai Ma, and Yefeng Zheng. Med3d: Transfer learning for 3D medical image analysis, 2019. [2](#)
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. [1](#), [2](#), [3](#), [5](#), [14](#)
- [9] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020. [1](#), [2](#), [3](#), [5](#), [14](#), [16](#)
- [10] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [2](#)
- [11] Veronika Cheplygina, Marleen de Bruijne, and Josien PW Pluim. Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *Medical image analysis*, 54:280–296, 2019. [3](#)
- [12] 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](#)
- [13] Robin Geyer, Luca Corinzia, and Viktor Wegmayr. Transfer learning by adaptive merging of multiple models. In M. Jorge Cardoso, Aasa Feragen, Ben Glocker, Ender Konukoglu, Ipek Oguz, Gozde Unal, and Tom Vercauteren, editors, *Proceedings of Machine Learning Research*. PMLR, 2019. [2](#)
- [14] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. [2](#)
- [15] Mara Graziani, Vincent Andrearczyk, and Henning Müller. Visualizing and interpreting feature reuse of pretrained cnns for histopathology. In *MVIP 2019: Irish Machine Vision and Image Processing Conference Proceedings*. Irish Pattern Recognition and Classification Society, 2019. [1](#), [2](#)
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019. [1](#)
- [17] 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*, pages 9729–9738, 2020. [2](#)
- [18] 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. [2](#), [3](#)
- [19] Xuehai He, Xingyi Yang, Shanghang Zhang, Jinyu Zhao, Yichen Zhang, Eric Xing, and Pengtao Xie. Sample-efficient deep learning for COVID-19 diagnosis based on CT scans. *medRxiv*, 2020. [3](#)
- [20] Michal Heker and Hayit Greenspan. Joint liver lesion segmentation and classification via transfer learning. *arXiv preprint arXiv:2004.12352*, 2020. [1](#), [2](#)
- [21] Olivier J Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272*, 2019. [2](#)
- [22] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. 2019. [2](#)
- [23] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 590–597, 2019. [1](#), [5](#), [6](#), [12](#), [13](#)
- [24] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (BiT): General visual representation learning. *arXiv preprint arXiv:1912.11370*, 6, 2019. [7](#), [14](#)
- [25] Hongwei Li, Fei-Fei Xue, Krishna Chaitanya, Shengda Liu, Ivan Ezhov, Benedikt Wiestler, Jianguo Zhang, and Bjoern Menze. Imbalance-aware self-supervised learning for 3d radiomic representations. *arXiv preprint arXiv:2103.04167*, 2021. [3](#)
- [26] Gaobo Liang and Lixin Zheng. A transfer learning method with deep residual network for pediatric pneumonia diagnosis. *Computer methods and programs in biomedicine*, 187:104964, 2020. [2](#)
- [27] Jingyu Liu, Gangming Zhao, Yu Fei, Ming Zhang, Yizhou Wang, and Yizhou Yu. Align, attend and locate: Chest x-ray

- diagnosis via contrast induced attention network with limited supervision. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10632–10641, 2019. [3](#)
- [28] Quande Liu, Lequan Yu, Luyang Luo, Qi Dou, and Pheng Ann Heng. Semi-supervised medical image classification with relation-driven self-ensembling model. *IEEE Transactions on Medical Imaging*, 2020. [3](#)
- [29] Yuan Liu, Ayush Jain, Clara Eng, David H Way, Kang Lee, Peggy Bui, Kimberly Kanada, Guilherme de Oliveira Marinho, Jessica Gallegos, Sara Gabriele, et al. A deep learning system for differential diagnosis of skin diseases. *Nature Medicine*, pages 1–9, 2020. [1](#), [2](#), [4](#), [6](#), [12](#)
- [30] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 181–196, 2018. [7](#)
- [31] Scott Mayer McKinney, Marcin Sieniek, Varun Godbole, Jonathan Godwin, Natasha Antropova, Hutan Ashrafian, Trevor Back, Mary Chesus, Greg C Corrado, Ara Darzi, et al. International evaluation of an AI system for breast cancer screening. *Nature*, 577(7788):89–94, 2020. [1](#), [2](#)
- [32] Afonso Menegola, Michel Fornaciali, Ramon Pires, Flávia Vasques Bittencourt, Sandra Avila, and Eduardo Valle. Knowledge transfer for melanoma screening with deep learning. In *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, pages 297–300. IEEE, 2017. [1](#), [2](#)
- [33] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6707–6717, 2020. [2](#)
- [34] Behnam Neyshabur, Hanie Sedghi, and Chiyuan Zhang. What is being transferred in transfer learning? *Advances in Neural Information Processing Systems*, 33, 2020. [5](#), [12](#), [13](#)
- [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](#)
- [36] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [2](#)
- [37] Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. In *Advances in neural information processing systems*, pages 3347–3357, 2019. [2](#), [5](#), [12](#), [13](#), [18](#)
- [38] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, and Sergey Levine. Time-contrastive networks: Self-supervised learning from video. In *IEEE International Conf. on Robotics and Automation (ICRA)*, pages 1134–1141. IEEE, 2018. [3](#)
- [39] Hari Sowrirajan, Jingbo Yang, Andrew Y Ng, and Pranav Rajpurkar. Moco pretraining improves representation and transferability of chest X-ray models. *arXiv:2010.05352*, 2020. [3](#)
- [40] Hannah Spitzer, Kai Kiwitz, Katrin Amunts, Stefan Harmeling, and Timo Dickscheid. Improving cytoarchitectonic segmentation of human brain areas with self-supervised siamese networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 663–671. Springer, 2018. [3](#)
- [41] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019. [2](#)
- [42] Michael Tschannen, Josip Djolonga, Marvin Ritter, Aravindh Mahendran, Neil Houlsby, Sylvain Gelly, and Mario Lucic. Self-supervised learning of video-induced visual invariances. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, 2020. [3](#)
- [43] Dong Wang, Yuan Zhang, Kexin Zhang, and Liwei Wang. Focalmix: Semi-supervised learning for 3d medical image detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3951–3960, 2020. [3](#)
- [44] 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. [5](#)
- [45] 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. [2](#)
- [46] Huidong Xie, Hongming Shan, Wenxiang Cong, Xiaohua Zhang, Shaohua Liu, Ruola Ning, and Ge Wang. Dual network architecture for few-view CT-trained on imagenet data and transferred for medical imaging. In *Developments in X-Ray Tomography XII*, volume 11113, page 111130V. International Society for Optics and Photonics, 2019. [1](#), [2](#)
- [47] 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 Conference on computer vision and pattern recognition*, pages 6210–6219, 2019. [2](#)
- [48] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017. [5](#)
- [49] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016. [2](#)
- [50] Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. Contrastive learning of medical visual representations from paired images and text. *arXiv preprint arXiv:2010.00747*, 2020. [3](#)
- [51] 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 *MICCAI*, pages 398–407. Springer, 2020. [3](#)

- [52] Jiuwen Zhu, Yuexiang Li, Yifan Hu, Kai Ma, S Kevin Zhou, and Yefeng Zheng. Rubik's cube+: A self-supervised feature learning framework for 3D medical image analysis. *Medical Image Analysis*, page 101746, 2020. 3
- [53] 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 *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 420–428. Springer, 2019. 3