




# Multi-Feature Vision Transformer via Self-Supervised Representation Learning for Improvement of COVID-19 Diagnosis

Xiao Qi<sup>1</sup>(✉), David J. Foran<sup>4</sup>, John L. Noshier<sup>2</sup>, and Ilker Hacihaliloglu<sup>2,3</sup> 

<sup>1</sup> Department of Electrical and Computer Engineering, Rutgers University,  
New Brunswick, NJ, USA

[xiao.qi@rutgers.edu](mailto:xiao.qi@rutgers.edu)

<sup>2</sup> Department of Radiology, The University of British Columbia,  
Vancouver, BC, Canada

<sup>3</sup> Department of Medicine, The University of British Columbia,  
Vancouver, BC, Canada

<sup>4</sup> Rutgers Cancer Institute of New Jersey, New Brunswick, NJ, USA

**Abstract.** The role of chest X-ray (CXR) imaging, due to being more cost-effective, widely available, and having a faster acquisition time compared to CT, has evolved during the COVID-19 pandemic. To improve the diagnostic performance of CXR imaging a growing number of studies have investigated whether supervised deep learning methods can provide additional support. However, supervised methods rely on a large number of labeled radiology images, which is a time-consuming and complex procedure requiring expert clinician input. Due to the relative scarcity of COVID-19 patient data and the costly labeling process, self-supervised learning methods have gained momentum and has been proposed achieving comparable results to fully supervised learning approaches. In this work, we study the effectiveness of self-supervised learning in the context of diagnosing COVID-19 disease from CXR images. We propose a multi-feature Vision Transformer (ViT) guided architecture where we deploy a cross-attention mechanism to learn information from both original CXR images and corresponding enhanced local phase CXR images. By using 10% labeled CXR scans, the proposed model achieves 91.10% and 96.21% overall accuracy tested on total 35,483 CXR images of healthy (8,851), regular pneumonia (6,045), and COVID-19 (18,159) scans and shows significant improvement over state-of-the-art techniques. Code is available <https://github.com/endiqq/Multi-Feature-ViT>.

**Keywords:** Self-supervised learning · Vision transformer · Cross-attention · COVID-19 · Chest X-ray

## 1 Introduction

The rapid spread of COVID-19 outbreak caused a surge of patients to emergency departments and hospitalization. Compared to CT, chest X-ray (CXR) has

several advantages such as its wide availability, exposure to less radiation, and faster image acquisition times. Due to this CXR has become the primary diagnostic tool for improved management of COVID-19. However, the interpretation of CXR images, compared to CT, is more challenging due to low image resolution and COVID-19 image features being similar to regular pneumonia. Computer-aided diagnosis via deep learning has been investigated to help mitigate these problems and help clinicians during the decision-making process [20, 21, 24]. Most supervised deep learning methods rely on a large number of labeled radiology images. Medical image labeling is a time-consuming and complex procedure requiring expert clinician input.

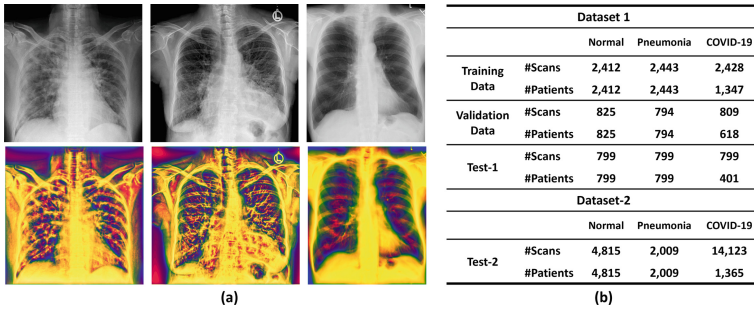
Semi-supervised learning methods have been proposed to provide a solution to scarcity of data in the context of COVID-19 diagnosis. [22] proposed a multi-feature guided teacher-student distillation approach. Most recently self-supervised learning methods, which utilize all the unlabeled data during learning, have been investigated for COVID-19 diagnosis [8, 9, 19]. [8] achieved 79.5% accuracy and 86.6% area under the receiver operating characteristic curve (AUC) on 426 COVID-19 scans while using 1% of the labeled data for the pretext task. [9] reported a mean average precision of 41.6% for 1,214 COVID-10 test data. [19] reported 99.5% accuracy using 607 COVID-19 scans. While these initial results on self-supervised learning are promising most of the prior work was evaluated on limited COVID-19 scans.

In order to break the challenges associated with scarcity of training data and to boost classification performance, we propose a new self-supervised learning approach where the proposed framework exploits local phase enhanced CXR image features to significantly improve the learning performance, specially when the data is limited. Our contributions and findings include the following: 1) We developed MoCo-COVID, a Vision Transformer (ViT) with modified Momentum Contrast (MoCo) pretraining on CXR images for self-supervised learning. This is the first study pretraining a ViT using MoCo for COVID-19 diagnosis from CXR images. 2) We demonstrated the performance of MoCo-COVID can be significantly improved by leveraging the local phase-based enhanced CXR scans specially in low data regime. 80.27% and 93.24% overall accuracy were achieved tested on 799 and 14,123 COVID-19 scans while using 1% of the labeled local phase-based enhanced data for the training. 3) A novel objective function was proposed using knowledge distillation to provide better generalization. 4) We developed a multi-feature ViT architecture based on cross-attention mechanism (MF-ViT CA) to further improve accuracy. The proposed MF-ViT CA achieves 95.03% and 97.35% mean accuracy on two large-scale test datasets including 14,922 COVID-19 scans and outperforms state-of-the-art semi-supervised learning and self-supervised learning methods.

## 2 Methods

**Datasets:** All images were collected from six public data repositories, which are BIMCV [12], COVIDx [30], COVID-19-AR [6], MIDRC-RICORD-1c [28],

COVID-19 Image Repository [31] and COVID-19-NY-SBU [5]. Our dataset consists of a total of 33,055 CXR scans from 18,252 patients with three classes: normal, pneumonia, and COVID-19. These images were split into two datasets: **1) Dataset-1:** A subset containing 12,108 CXR scans with a balanced distribution of classes (Fig. 1(b)). Dataset-1 was split into 60% training, 20% as validation, and 20% as testing dataset. No subject overlaps among train, validation, and test datasets. The test data in this dataset is referred to as *Test-1*. **2) Dataset-2:** Includes 20,947 CXR scans from 8,181 patients and has a larger number of scans in COVID-19 class. Dataset-2 only serves as an additional test dataset, referred to as *Test-2*, for evaluating the robustness of the proposed methods.



**Fig. 1.** (a) Top row original CXR images. Bottom row  $MF(x, y)$  images. The first two columns are from subjects who are diagnosed with COVID-19. The last column is from a healthy subject. (b) Class distribution of the evaluation datasets.

**Image Enhancement:** Local phase-based image analysis methods are more robust to intensity variations, usually arising from patient characteristics or image acquisition settings, and have been incorporated into various medical image processing tasks [1, 15, 32]. The enhanced local phase CXR image, denoted as  $MF(x, y)$ , is obtained by combining three different local phase image features: 1-Local weighted mean phase angle ( $LwPA(x, y)$ ), 2-Weighted local phase energy ( $LPE(x, y)$ ), and 3-Enhanced local energy attenuation image ( $ELEA(x, y)$ ).  $LPE(x, y)$  and  $LwPA(x, y)$  image features are extracted by filtering the CXR image in frequency domain using monogenic filter and  $\alpha$ -scale space derivative (ASSD) bandpass quadrature filters [21].  $ELEA(x, y)$  image is extracted, processing the  $LPE(x, y)$  image, by modeling the scattering and attenuation effects of lung tissue inside a local region using L1 norm-based contextual regularization method [21]. We have used the filter parameters reported in [21] for enhancing all the CXR images. Investigating Fig. 1(a) we can see that structural features inside the lung tissue are more dominant for COVID-19 CXR images compared to healthy lung (last column Fig. 1(a)). The enhanced local phase CXR images ( $MF(x, y)$ ) and the original CXR images are used to train proposed self-supervised learning methods explained in the next sections.

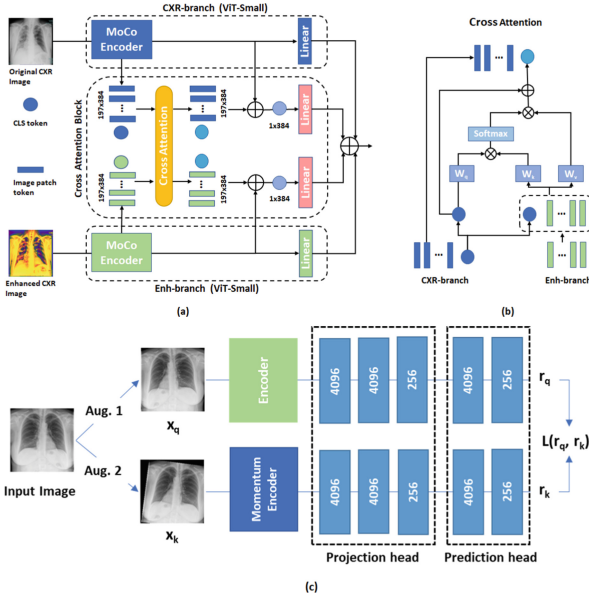
**MoCo-COVID-Self-Supervised Pretraining Using MoCo:** We introduce a MoCo-COVID framework (Fig. 2(c)). During the self-supervised training, a CXR image is first transformed via two random augmentations (Aug.1 and Aug.2) into images  $x_q$  and  $x_k$ .  $x_q$  is passed through an encoder network, while  $x_k$  is passed through a momentum encoder network. We choose ViT-Small (ViT-S) [7] with 6 heads as the backbone, instead of 12 heads used by MoCo v3, to have a lower parameter count, and a faster throughput [3]. Sine-cosine variant [29] is added to the sequence as positional embedding. The representations generated by each network are then passed into the projection head followed by a prediction head. The projection head has three layers and the prediction head has two layers. Each layer follows with batch normalization (BN) and a rectified linear unit (ReLU) except the last layer of both projection head and prediction head. Then the InfoNCE contrastive loss function [18] is adopted to promote the similarity between the representations  $r_q$  and  $r_k$ :

$$L(r_q, r_k) = -\log \frac{\exp(r_q \cdot r_k + 1/\tau)}{\sum_{i=1}^K \exp(r_q \cdot r_k, i/\tau)}, \quad (1)$$

where  $\tau$  is a temperature hyperparameter and  $K$  is the number of currently stored representations. The momentum coefficient follows a cosine schedule changing from 0.9 to 0.999 during the MoCo-COVID pretraining.

**Multi-Feature Fusion Vision Transformer via Cross Attention:** Figure 2(a) illustrates the architecture of our proposed Multi-Feature Vision Transformer with cross-attention (CA) block (MF-ViT CA), which consists of two branches and a CA block. *CXR-branch* is used for processing the original CXR image, *Enh-branch* is used for processing the enhanced local phase CXR image ( $MF(x, y)$ ), and CA block for extracting information from both branches. During the forward pass, an original CXR image and corresponding  $MF(x, y)$  image is first passed to the pretrained MoCo encoder in parallel to obtain a tensor of dimension  $197 \times 384$  for each image. These two tensors are then fed as inputs to our CA block. In the CA block, the CLS token of one branch fuses with patch tokens of the other branch using CA mechanism, and after fusion the CLS token concatenates with its own patch tokens again to produce an output in the same dimension of input tensor for each branch, described more details in below. The outputs of cross-attention and MoCo encoder are then combined via element-wise summation for each branch. The CLS token of dimension  $1 \times 384$  from each branch as a compact representation, which encodes the information from both the original CXR image and  $MF(x, y)$  image, is fed into a linear projection layer with three units. After that, the outputs of the two projection functions are fused using an element-wise summation.

**Loss Function:** We used hard-label distillation, which is a variation of distillation introduced by [27], for training CA block. In our study, *CXR-branch* and *Enh-branch*, which are obtained with the MoCo-COVID-pretrained ViT-S by using original and enhanced CXR scans (denoted as *cxr* and *enh* in Eq. 2), are considered as the teacher models. The CA block is considered as the student model. The objective function with this hard-label distillation is:



**Fig. 2.** (a) Proposed multi-feature fusion ViT architecture using the cross attention block which is shown in (b). (c) The proposed MoCo-COVID architecture.

$$L_{\text{global}}^{\text{hardDistill}} = \frac{1}{3}L_{\text{CE}}(\psi(Z_s), y) + \frac{1}{3}L_{\text{CE}}(\psi(Z_s), y_{t_{cxr}}) + \frac{1}{3}L_{\text{CE}}(\psi(Z_s), y_{t_{enh}}). \quad (2)$$

where  $Z_s$  is the logits of the student model and  $\psi$  is the softmax function. The idea is to use both the real target  $y$  and the target generated by the teachers  $y_t = \text{argmax}_c Z_t(c)$ , where  $Z_t$  is the logits of the teacher models. For a given image, the predicted label associated with the teacher may change due to data augmentation during the forward pass. Thus, the teacher model is aiming at producing predicted labels that are similar but not identical to true label.

The major benefit of the proposed hard label distillation is being parameter-free compared to soft distillation in [11]. In our study, the probability distribution has the correct class at a very high probability, with all other class probabilities close to 0. The problem with these weak probabilities is they do not capture desirable information for the student model to learn effectively. To address this issue, we proposed this hard label distillation loss function by adding two distillation losses, which are losses between the predicted label of teacher model and logits of student model. Distillation loss would soften the distributions predicted by the teacher model so that the student model can learn more information, and this is especially useful when dealing with small datasets [11]. And the proposed objective function ensures the student model inherits better quality from the teacher model and mitigates the over-confidence issue of neural networks by improving the generalization.

**Cross-Attention (CA) Mechanism:** The proposed CA mechanism is inspired by [2], which is designed for fusing multi-scale features from two branches. In our study, we intend to extract information from two types of images via cross-attention, and thus we removed projection function used in [2] for changing the dimension of feature maps. Figure 2(b) presents the structure of our proposed cross-attention. Similar to [2], we utilize the CLS token at each branch to fuse with patch tokens from the other branch and then project it back to its own branch in order to exchange information among the patch tokens from the other branch. Since the CLS token already learns the extracted information from all patch tokens in its own branch, interacting with the patch tokens from the other branch helps to include information from different feature inputs. In the following section, we provide a detailed explanation about the CA mechanism for the *CXR-branch*. *Enh-branch* follows the same process.

**CXR-Branch-CA:** An illustration of CA for *CXR-branch* is presented in the Fig. 2(b). Firstly, the CLS tokens from *CXR-branch*, denoted as  $x_{cls}^{cxr}$ , concatenates the patch tokens from the *Enh-branch* to form  $x'^{cxr}$  shown as Eq. 3.

$$x'^{cxr} = [ x_{cls}^{cxr} \parallel x_{patch}^{enh} ] \quad (3)$$

Then, the CA is performed between  $x_{cls}^{cxr}$  and  $x'^{cxr}$  using linear projections to computer queries, keys and values ( $\mathbf{Q}$ ,  $\mathbf{K}$  and  $\mathbf{V}$ ), where CLS token is the only information used in query. And it uses the scaled dot product for calculating the attention weights between  $\mathbf{Q}$  and  $\mathbf{K}$  and then aggregates  $\mathbf{V}$ . The CA can be expressed as below:

$$\begin{aligned} \mathbf{Q} &= x_{cls}^{cxr} \mathbf{W}_q, \mathbf{K} = x'^{cxr} \mathbf{W}_k, \mathbf{V} = x'^{cxr} \mathbf{W}_v, \\ \mathbf{CA} &= \text{softmax}(\mathbf{Q}\mathbf{K}^T / \sqrt{D/h})\mathbf{V}, \end{aligned} \quad (4)$$

In the Eq. 4,  $\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v \in \mathbb{R}^{C \times (C/h)}$  are learnable parameters, where  $C$  and  $h$  are the embedding dimension and number of heads. We use the three heads in the cross attention in this study. In the end, the new CLS token of *CXR-branch*, which is obtained by cross attention and residual shortcut, concatenates with patch tokens from *CXR-branch* as the output of *CXR-branch* shown as below:

$$y_{cls}^{cxr} = x_{cls}^{cxr} + \mathbf{CA}, z^{cxr} = [ y_{cls}^{cxr} \parallel x_{patch}^{cxr} ] \quad (5)$$

### 3 Experiments and Results

**MoCo-COVID Pretraining:** We pretrained our MoCo-COVID end-to-end for both original CXR image (denoted as CXR-ViT-S) and  $MF(x, y)$  (denoted as Enh-ViT-S), on all training dataset without label information. MoCo-COVID pretraining initialization was performed using the weights obtained on ImageNet-initialized models for a faster convergence [23]. Data augmentation included resizing to a  $224 \times 224$  grid, random rotation ( $10^\circ$ ), and horizontal flipping similar to [25]. We maintained hyperparameters related to momentum, weight

decay, and feature dimension from MoCo v3 [3]. To be specific, the model was optimized by AdamW [17], a weight decay parameter of 0.1 and batch size of 16 on 2 NVIDIA GTX 1080 using PyTorch’s DistributedDataParallel framework [14]. We pretrained for 300 epochs (40 warm-up epochs) with a cosine linear-rate scheduler and set the initial learning rate as  $lr \times \text{BatchSize}/4$ , where  $lr$  is  $1.5e^{-4}$ .

We fine-tuned models with different fractions of labeled data. Label fraction represents the percentage of labeled data retained during fine-tuning. For example, a model fine-tuned with 1% label fraction meaning the model will only have access to 1% of the training dataset as labeled dataset, and the remaining 99% are hidden from the model as unseen data. The label fractions of training dataset are 0.25% (18 scans), 1% (72 scans), 10% (728 scans), 30% (2184 scans), and 100%. Fine-tuning was repeated five times for each label fraction. Label fractions less than 100% are random samples from the training dataset.

We conducted two fine-tuning ablations, which are linear probing (LP) and end-to-end fine-tuning (FT). LP means the pretrained weight values of the MoCo-COVID encoder were frozen and, after removing the projection and prediction heads, a new linear classifier with randomly reinitialized weights was added and fine-tuned using labeled data. FT allowed the entire model including MoCo-COVID encoder to fine-tune not just the newly added classifier. The models were fine-tuned using 90 epochs. All fine-tunings used the cosine annealing learning rate decay [16] and SGD [13] optimizer.

**Ablations of Multi-Feature ViT Using Cross-Attention Mechanism (MF-ViT CA):** The weights of the *CXR-branch* and *Enh-branch* are initialized to the MoCo-COVID pretrained weight values of CXR-ViT-S FT and Enh-ViT-S FT respectively, and the weights of the CA block are randomly reinitialized with a uniform distribution [10]. The CA block was fine-tuned using labeled data with hard-label distillation. We compared the proposed MF-ViT CA with a model, where it also has two MoCo-COVID pretrained branches (*CXR branch* and *Enh branch*) without using cross-attention block, denoted as MF-ViT LP. During the fine-tuning, we only fine-tuned the linear layers of *CXR-branch* and *Enh-branch*.

**Baselines:** As baseline comparison we report results from MoCo-CXR [25], which uses MoCo v2 pretrained Dense121, when trained using original CXR and enhanced local phase CXR images ( $MF(x, y)$ ). We report end-to-end fine-tuning protocol (FT) results for MoCo-CXR [25] and the network architecture was optimized to achieve the best results using our dataset to provide a fair comparison. Finally, we also compare our results against [22] semi-supervised, and fully supervised methods where  $MF(x, y)$  were used.

**Quantitative Results:** Quantitative results are displayed in Table 1. The proposed MF-ViT CA architecture achieved the best accuracy compared to the rest of the self-supervised learning models when label fractions were more than 1%. MF-ViT CA performed significantly better (paired t-test  $p < 0.05$ ) when tested on *Test-2* data compared to *Test-1* data proving the robustness of the method

when tested on large scale COVID-19 data (14,123 images). The proposed Enh-ViT-S FT model had the highest accuracy for label fraction less than 10%. We can also observe that end-to-end fine-tuning protocol (FT) results are significant improvements over linear probing (LP) protocol in the proposed MoCo-COVID model. The accuracy of the baseline architecture MoCo-CXR [25] and all the proposed architectures improve when enhanced local phase CXR images ( $MF(x, y)$ ) were used as an input. Finally, we have also observed that the CA mechanism significantly improves the results of the proposed MF-ViT LP model (paired t-text  $p < 0.05$ ) (Table 1). From the results, we observe that the MF-ViT CA yields a statistically significant gain (paired t-text  $p < 0.05$ ) compared to MF-TS [22] at 0.25% (18 scans) and 1% label fractions (72 scans). This indicates that the proposed models provide high-quality representations, better generalization capability, and transferable initialization for COVID-19 interpretation for minimal label fractions and when evaluated on large test data.

**Table 1.** Accuracy results obtained from *Test-1* and *Test-2* data. Green shaded region corresponds to the highest scores obtained. \* indicates statistical improvement compared with second best self-supervised learning method ( $p < 0.05$ ) using paired t-test.

Method	0.25%		1%		10%		30%		100%	
	Test1	Test2	Test1	Test2	Test1	Test2	Test1	Test2	Test1	Test2
<b>MoCo-CXR[25]-pretrained end-to-end Dense121</b>										
FT CXR	65.16	73.85	78.85	77.52	90.23	84.57	92.86	90.62	94.74	93.70
FT Enh	73.60	81.85	83.12	90.87	90.40	95.84	91.57	95.95	93.73	96.50
<b>MoCo-COVID Pretrained (ours)</b>										
CXR-ViT-S LP	72.64	74.18	77.35	78.33	83.76	81.71	85.10	83.88	86.93	86.15
Enh-ViT-S LP	79.97	91.71	84.06	94.00	87.81	94.87	89.10	95.85	90.51	94.14
CXR-ViT-S FT	73.00	73.37	78.23	81.13	88.59	85.94	91.66	89.48	93.26	92.19
Enh-ViT-S FT	<b>80.27</b>	<b>93.24*</b>	<b>84.10</b>	<b>94.00</b>	89.09	95.33	91.43	96.23	92.62	96.57
<b>Multi-Feature Model (ours)</b>										
MF-ViT LP	69.01	65.22	67.20	69.72	86.72	85.20	91.06	92.02	92.53	95.01
MF-ViT CA	79.88	89.91	82.72	92.57	<b>91.10*</b>	<b>96.21</b>	<b>93.27</b>	<b>96.84*</b>	<b>95.03</b>	<b>97.35*</b>
<b>Semi-Supervised Learning</b>										
MF-TS[22]	77.13	80.93	82.27	86.57	90.73	95.65	92.68	96.35	-	-
<b>Fully-Supervised Learning</b>										
XNet[4]	-	-	-	-	-	-	-	-	94.38	89.20
InceptionV4[26]	-	-	-	-	-	-	-	-	93.98	88.92



## 4 Conclusion

Our large quantitative evaluation results obtained using the largest COVID-19 data collected from different sites, show the significant improvements achieved using the local phase image features for self-supervised learning. Although we did not have access to the CXR machine type and non-image patient information (BMI, age, sex) we believe the large data used in this work represents images with varying image acquisition settings and intensity variations. Our quantitative results show significantly improved accuracy values over the investigated baselines proving the robustness of our proposed methods. Future work will include the extension of the method for diagnosing different lung diseases from CXR images.

## References

1. Alessandrini, M., Basarab, A., Liebgott, H., Bernard, O.: Myocardial motion estimation from medical images using the monogenic signal. *IEEE Trans. Image Process.* **22**(3), 1084–1095 (2012)
2. Chen, C.F.R., Fan, Q., Panda, R.: CrossViT: cross-attention multi-scale vision transformer for image classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 357–366 (2021)
3. Chen, X., Xie, S., He, K.: An empirical study of training self-supervised vision transformers. *CoRR abs/2104.02057* (2021). <https://arxiv.org/abs/2104.02057>
4. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1251–1258 (2017)
5. Clark, K., et al.: The cancer imaging archive (TCIA): maintaining and operating a public information repository. *J. Digital Imaging* **26**(6), 1045–1057 (2013). <https://doi.org/10.1007/s10278-013-9622-7>
6. Desai, S., et al.: Chest imaging representing a COVID-19 positive rural U.S. population. *Sci. Data* **7**, 1–6 (2020). <https://doi.org/10.1038/s41597-020-00741-6>
7. Dosovitskiy, A., et al.: An image is worth 16 × 16 words: transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
8. Gazda, M., Plavka, J., Gazda, J., Drotar, P.: Self-supervised deep convolutional neural network for chest x-ray classification. *IEEE Access* **9**, 151972–151982 (2021)
9. Hao, Y., Wang, Y., Wang, X.: Self-supervised pretraining for COVID-19 and other pneumonia detection from chest X-ray images. In: Xie, Q., Zhao, L., Li, K., Yadav, A., Wang, L. (eds.) *ICNC-FSKD 2021. LNDECT*, vol. 89, pp. 1000–1007. Springer, Cham (2022). [https://doi.org/10.1007/978-3-030-89698-0\\_102](https://doi.org/10.1007/978-3-030-89698-0_102)
10. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. *CoRR abs/1502.01852* (2015). <http://arxiv.org/abs/1502.01852>
11. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network (2015). <https://doi.org/10.48550/ARXIV.1503.02531>, <https://arxiv.org/abs/1503.02531>
12. de la Iglesia Vayá, M., et al.: BIMCV COVID-19+: a large annotated dataset of RX and CT images from COVID-19 patients. *arXiv preprint arXiv:2006.01174* (2020)
13. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization (2017)
14. Li, S., et al.: Pytorch distributed: experiences on accelerating data parallel training. *CoRR abs/2006.15704* (2020). <https://arxiv.org/abs/2006.15704>

15. Li, Z., van Vliet, L.J., Stoker, J., Vos, F.M.: A hybrid optimization strategy for registering images with large local deformations and intensity variations. *Int. J. Comput. Assist. Radiol. Surg.* **13**(3), 343–351 (2017). <https://doi.org/10.1007/s11548-017-1697-z>
16. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with restarts. *CoRR abs/1608.03983* (2016). <http://arxiv.org/abs/1608.03983>
17. Loshchilov, I., Hutter, F.: Fixing weight decay regularization in Adam. *CoRR abs/1711.05101* (2017). <http://arxiv.org/abs/1711.05101>
18. Van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv e-prints. arXiv-1807* (2018)
19. Park, J., Kwak, I.Y., Lim, C.: A deep learning model with self-supervised learning and attention mechanism for COVID-19 diagnosis using chest X-ray images. *Electronics* **10**(16), 1996 (2021)
20. Park, S., et al.: Multi-task vision transformer using low-level chest X-ray feature corpus for COVID-19 diagnosis and severity quantification. *Med. Image Anal.* **75**, 102299 (2022)
21. Qi, X., Brown, L.G., Foran, D.J., Noshier, J., Hacihaliloglu, I.: Chest X-ray image phase features for improved diagnosis of COVID-19 using convolutional neural network. *Int. J. Comput. Assist. Radiol. Surg.* **19**, 1–10 (2020)
22. Qi, X., Foran, D.J., Noshier, J.L., Hacihaliloglu, I.: Multi-feature semi-supervised learning for COVID-19 diagnosis from chest X-ray images. In: Lian, C., Cao, X., Rekić, I., Xu, X., Yan, P. (eds.) *MLMI 2021. LNCS*, vol. 12966, pp. 151–160. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-87589-3\\_16](https://doi.org/10.1007/978-3-030-87589-3_16)
23. Raghu, M., Zhang, C., Kleinberg, J.M., Bengio, S.: Transfusion: understanding transfer learning with applications to medical imaging. *CoRR abs/1902.07208* (2019). <http://arxiv.org/abs/1902.07208>
24. Serena Low, W.C., et al.: An overview of deep learning techniques on chest X-ray and CT scan identification of COVID-19. In: *Computational and Mathematical Methods in Medicine 2021* (2021)
25. Sowrirajan, H., Yang, J., Ng, A.Y., Rajpurkar, P.: MOCO pretraining improves representation and transferability of chest X-ray models. *CoRR abs/2010.05352* (2020). <https://arxiv.org/abs/2010.05352>
26. Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.: Inception-v4, inception-resnet and the impact of residual connections on learning (2016)
27. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. *CoRR abs/2012.12877* (2020). <https://arxiv.org/abs/2012.12877>
28. Tsai, E.B., et al.: The RSNA international COVID-19 open annotated radiology database (RICORD). *Radiology.* **299**, E204 (2021). <https://doi.org/10.1148/radiol.2021203957>, PMID: 33399506
29. Vaswani, A., et al.: Attention is all you need. *CoRR abs/1706.03762* (2017). <http://arxiv.org/abs/1706.03762>
30. Wang, L., Lin, Z.Q., Wong, A.: COVID-net: a tailored deep convolutional neural network design for detection of COVID-19 cases from chest X-ray images. *Sci. Rep.* **10**(1), 19549 (2020). <https://doi.org/10.1038/s41598-020-76550-z>
31. Winther, H.B., et al.: COVID-19 image repository (2020). <https://doi.org/10.6084/m9.figshare.12275009.v1>
32. Zhao, Y., Liu, Y., Wu, X., Harding, S.P., Zheng, Y.: Retinal vessel segmentation: an efficient graph cut approach with Retinex and local phase. *PLoS ONE* **10**(4), e0122332 (2015)