



# Self-supervised Mean Teacher for Semi-supervised Chest X-Ray Classification

Fengbei Liu<sup>1(✉)</sup>, Yu Tian<sup>1,4</sup>, Filipe R. Cordeiro<sup>2</sup>, Vasileios Belagiannis<sup>3</sup>,  
Ian Reid<sup>1</sup>, and Gustavo Carneiro<sup>1</sup>

<sup>1</sup> Australian Institute for Machine Learning, University of Adelaide,  
Adelaide, Australia

`fengbei.liu@adelaide.edu.au`

<sup>2</sup> Universidade Federal Rural de Pernambuco, Recife, Brazil

<sup>3</sup> Universität Ulm, Ulm, Germany

<sup>4</sup> South Australian Health and Medical Research Institute, Adelaide, Australia

**Abstract.** The training of deep learning models generally requires a large amount of annotated data for effective convergence and generalisation. However, obtaining high-quality annotations is a labourious and expensive process due to the need of expert radiologists for the labelling task. The study of semi-supervised learning in medical image analysis is then of crucial importance given that it is much less expensive to obtain unlabelled images than to acquire images labelled by expert radiologists. Essentially, semi-supervised methods leverage large sets of unlabelled data to enable better training convergence and generalisation than using only the small set of labelled images. In this paper, we propose Self-supervised Mean Teacher for Semi-supervised ( $S^2MTS^2$ ) learning that combines self-supervised mean-teacher pre-training with semi-supervised fine-tuning. The main innovation of  $S^2MTS^2$  is the self-supervised mean-teacher pre-training based on the joint contrastive learning, which uses an infinite number of pairs of positive query and key features to improve the mean-teacher representation. The model is then fine-tuned using the exponential moving average teacher framework trained with semi-supervised learning. We validate  $S^2MTS^2$  on the multi-label classification problems from Chest X-ray14 and CheXpert, and the multi-class classification from ISIC2018, where we show that it outperforms the previous SOTA semi-supervised learning methods by a large margin. Our code will be available upon paper acceptance.

**Keywords:** Semi-supervised learning · Chest X-ray · Self-supervised learning · Multi-label classification

## 1 Introduction

Deep learning has shown outstanding results in medical image analysis problems [16, 19–21, 28, 29]. However, this performance usually depends on the avail-

F. Liu and Y. Tian—First two authors contributed equally to this work.

© Springer Nature Switzerland AG 2021

C. Lian et al. (Eds.): MLMI 2021, LNCS 12966, pp. 426–436, 2021.

[https://doi.org/10.1007/978-3-030-87589-3\\_44](https://doi.org/10.1007/978-3-030-87589-3_44)

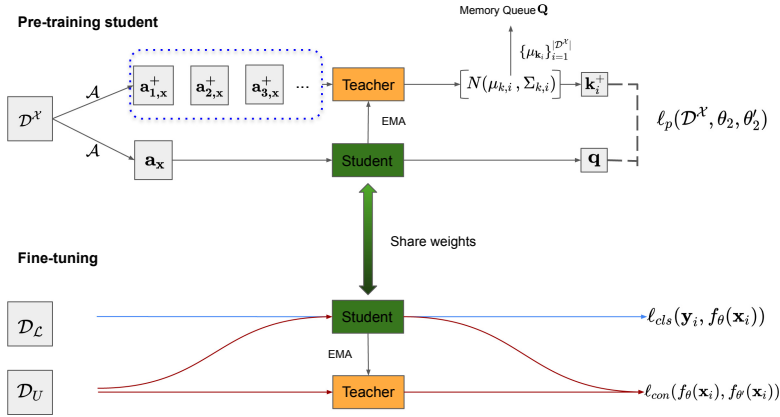
ability of labelled datasets, which is expensive to obtain given that the labelling process requires expert doctors. This limitation motivates the study of semi-supervised learning (SSL) methods that train models with a small set of labelled data and a large set of unlabelled data.

The current state-of-the-art (SOTA) SSL is based on pseudo-labelling methods [18, 26], consistency-enforcing approaches [2, 17, 27], self-supervised and semi-supervised learning (S<sup>4</sup>L) [5, 34], and graph-based label propagation [1]. Pseudo-labelling is an intuitive SSL technique, where confident predictions from the model are transformed into pseudo-labels for the unlabelled data, which are then used to re-train the model [18]. Consistency-enforcing regularisation is based on training for a consistent output given model [22, 27] or input data [2, 17] perturbations. S<sup>4</sup>L methods are based on self-supervised pre-training [4, 12], followed by supervised fine-tuning using few labelled samples [5, 34]. Graph-based methods rely on label propagation on graphs [1]. Recently, Yang et al. [33] suggested that self-supervision pre-training provides better feature representations than consistency-enforcing approaches in SSL. However, previous S<sup>4</sup>L approaches use only the labelled data in the fine-tuning stage, missing useful training information present in the unlabelled data. Furthermore, self-supervised pre-training [4, 12] tends to use limited amount of samples to represent each class, but recently, Cai et al. [3] showed that better representation can be obtained with an infinite amount of samples. Also, recent research [26] suggests that the student-teacher framework, such as the mean-teacher [27], works better in multi-label semi-supervised tasks than other SSL methods. We speculate that this is because other methods are usually designed to work with softmax activation that only works in multi-class problems, while mean-teacher [27] does not have this constraint and can work in multi-label problems.

In this paper, we propose a self-supervised mean-teacher for semi-supervised (S<sup>2</sup>MTS<sup>2</sup>) learning approach that combines S<sup>4</sup>L [5, 34] with consistency-enforcing learning based on the mean-teacher algorithm [27]. The main contribution of our method is the self-supervised mean-teacher pre-training with the joint contrastive learning [3]. To the best of our knowledge, this is the first approach, in our field, to train the mean teacher model with self-supervised learning. This model is then fine-tuned with semi-supervised learning using the exponential moving average teacher framework [27]. We evaluate our proposed method on the thorax disease multi-label datasets ChestX-ray 14 [32] and CheXpert [15], and on the multi-class skin condition dataset ISIC2018 [8, 30]. We show that our method outperforms the SOTA on semi-supervised learning [1, 11, 22, 31]. Moreover, we investigate each component of our framework for their contribution to the overall model in the ablation study.

## 2 Related Works

SSL is a research topic that is gaining attention from the medical image analysis community due to the expensive image annotation process [7] and the growing number of large-scale datasets available in the field [32]. The current SOTA



**Fig. 1.** Description of the proposed self-supervised mean-teacher for semi-supervised ( $S^2MTS^2$ ) learning. The main contribution of the paper resides in the top part of the figure, with the self-supervised mean-teacher pre-training based on joint contrastive learning, which uses an infinite number of pairs of positive query and key features sampled from the unlabelled images to minimise  $\ell_p(\cdot)$  in (1). This model is then fine-tuned with the exponential moving average teacher in a semi-supervised learning framework that uses both labelled and unlabelled sets to minimise  $\ell_{cls}(\cdot)$  and  $\ell_{con}(\cdot)$  in (2).

SSL methods are based on consistency-enforcing approaches that leverage the unlabelled data to regularise the model prediction consistency [17,27]. Other related papers [9] extend the mean teacher [27] to encourage consistency between the prediction by the student and teacher models for atrium and brain lesion segmentation. The SOTA SSL method on Chest X-ray images [22] exploits the consistency in the relations between labelled and unlabelled data. None of these methods explores a self-supervised consistency-enforcing method to pre-train an SSL model, as we propose.

Self-supervised learning methods [4,12] are also being widely investigated in SSL because they can provide good representations [5,34]. However, these methods ignore the large amount of unlabelled data to be used during SSL, which may lead to unsatisfactory generalisation process. An important point in self-supervised learning is on how to define the classes to be learned. In general, each class is composed of a single pair of augmented images from the same image, and many pairs of augmentations from different images [3–5,12]. The use of a single pair of images to form a class has been criticised by Cai et al. [3], who propose the joint contrastive learning (JCL), which is an efficient way to form a class with an infinite number of augmented images from the same image to leverage the statistical dependency between different augmentations.

### 3 Method

In this section, we introduce our two-stage learning framework in detail (see Fig. 1). We assume that we have a small labelled dataset, denoted by  $\mathcal{D}_L = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{D}_L|}$ , where the image is represented by  $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{H \times W \times C}$ , and class  $\mathbf{y} \in \{0, 1\}^{|\mathcal{Y}|}$ , where  $\mathcal{Y}$  represents the label set. We consider a multi-label problem and thus  $\sum_{c=1}^{|\mathcal{Y}|} \mathbf{y}_i(c) \in [0, |\mathcal{Y}|]$ . The unlabelled dataset is defined by  $\mathcal{D}_U = \{\mathbf{x}_i\}_{i=1}^{|\mathcal{D}_U|}$  with  $|\mathcal{D}_L| \ll |\mathcal{D}_U|$ .

Our model consists of a student and a teacher model [27], denoted by parameters  $\theta, \theta' \in \Theta$ , respectively, which parameterize the classifier  $f_\theta : \mathcal{X} \rightarrow [0, 1]^{|\mathcal{Y}|}$ . This classifier can be decomposed as  $f_\theta = h_{\theta_1} \circ g_{\theta_2}$ , with  $g_{\theta_2} : \mathcal{X} \rightarrow \mathcal{Z}$  and  $h_{\theta_1} : \mathcal{Z} \rightarrow [0, 1]^{|\mathcal{Y}|}$ . The first stage (top of Fig. 1) of the training consists of a self-supervised learning that uses the images from  $\mathcal{D}_L$  and  $\mathcal{D}_U$ , denoted by  $\mathcal{D}^\mathcal{X} = \{\mathbf{x}_i | \mathbf{x}_i \in \mathcal{D}_L^\mathcal{X} \cup \mathcal{D}_U\}_{i=1}^{|\mathcal{D}^\mathcal{X}|}$ , with  $\mathcal{D}_L^\mathcal{X}$  representing the images from the set  $\mathcal{D}_L$ , where our method minimises the joint contrastive learning loss [3], defined in (1). This means that during this first stage, we only learn the parameters for  $g_{\theta_2}$ . The second stage (bottom of Fig. 1) fine-tunes this pre-trained student-teacher model using the semi supervised consistency loss defined in (2). Below we provide details on the losses and training.

#### 3.1 Joint Contrastive Learning to Self-supervise the Mean-Teacher Pre-training

The self-supervised pre-training of the mean-teacher using joint contrastive learning (JCL) [3], presented in this section, is the main technical contribution of this paper. The teacher and student process an input image to return the keys  $\mathbf{k} \in \mathcal{Z}$  and the queries  $\mathbf{q} \in \mathcal{Z}$  with  $\mathbf{k} = g_{\theta_2'}(\mathbf{x})$  and  $\mathbf{q} = g_{\theta_2}(\mathbf{x})$ . We also assume that we have a set of augmentation functions, i.e., random crop and resize, rotation and Gaussian blur, denoted by  $\mathcal{A} = \{a_l : \mathcal{X} \rightarrow \mathcal{X}\}_{l=1}^{|\mathcal{A}|}$ . Then JCL minimises the following loss [3]:

$$\ell_p(\mathcal{D}^\mathcal{X}, \theta_2, \theta_2') = -\frac{1}{|\mathcal{D}^\mathcal{X}|} \frac{1}{M} \sum_{i=1}^{|\mathcal{D}^\mathcal{X}|} \sum_{m=1}^M \left[ \log \frac{\exp \left[ \frac{1}{\tau} \mathbf{q}_i^\top \mathbf{k}_{i,m}^+ \right]}{\exp \left[ \frac{1}{\tau} \mathbf{q}_i^\top \mathbf{k}_{i,m}^+ \right] + \sum_{j=1}^K \exp \left[ \frac{1}{\tau} \mathbf{q}_i^\top \mathbf{k}_{i,j}^- \right]} \right], \quad (1)$$

where  $\tau$  is the temperature hyper-parameter, the query  $\mathbf{q}_i = g_{\theta_2}(a(\mathbf{x}_i))$ , with  $a \in \mathcal{A}$ . the positive key  $\mathbf{k}_{i,m}^+ \sim p(\mathbf{k}_i^+)$ , with  $p(\mathbf{k}_i^+) = \mathcal{N}(\mu_{\mathbf{k}_i}, \Sigma_{\mathbf{k}_i})$  and  $\mathbf{k}_i = g_{\theta_2'}(a(\mathbf{x}_i))$  (i.e., a sample from the data augmentation distribution for  $\mathbf{x}$ ), and the negative keys  $\mathbf{k}_{i,j}^- \in \{\mu_{\mathbf{k}_j}\}_{i,j \in \{1, \dots, |\mathcal{D}^\mathcal{X}|\}, i \neq j}$  represents a negative key for query  $\mathbf{q}_i$ . In (1),  $M$  denotes the number of positive keys, and Cai et al. [3] describe a loss that minimises a bound to (1) for  $M \rightarrow \infty$  – below, the minimisation of  $\ell_p(\cdot)$  in (1) is realised by the minimisation of this bound. As defined above, the generative model  $p(\mathbf{k}_i^+)$  is denoted by the Gaussian  $\mathcal{N}(\mu_{\mathbf{k}_i}, \Sigma_{\mathbf{k}_i})$ , where the mean  $\mu_{\mathbf{k}_i}$  and covariance  $\Sigma_{\mathbf{k}_i}$  are estimated from a set of keys  $\{\mathbf{k}_{i,l}^+ = g_{\theta_2'}(a_l(\mathbf{x}_i))\}_{a_l \in \mathcal{A}}$  formed by different views of  $\mathbf{x}_i$ . The set of negative keys  $\{\mu_{\mathbf{k}_j}\}_{i,j \in \{1, \dots, |\mathcal{D}^\mathcal{X}|\}, i \neq j}$  is stored

in a memory queue [12] that is updated in a first-in-first-out way, where the mean of the keys in  $\{\mu_{\mathbf{k}_i}\}_{i=1}^{|\mathcal{D}^X|}$  are inserted to the memory queue to replace the oldest key means from previous training iterations. The memory queue has been designed to increase the number of negative samples without sacrificing computation efficiency.

The training of the student-teacher model [12, 27, 35] is achieved by updating the student parameter using the loss in (1), as in  $\theta_2(t) = \theta_2(t-1) - \nabla_{\theta_2} \ell_p(\mathcal{D}^X, \theta_2, \theta'_2)$ , where  $t$  is the training iteration. The teacher model parameter is updated with exponential moving average (EMA) with  $\theta'_2(t) = \alpha\theta'_2(t-1) + (1-\alpha)\theta_2(t)$ , with  $\alpha \in [0, 1]$ . For this pre-training stage, we notice that training for more epochs always improve the model regularisation given that it is difficult to overfit the training set with the loss in (1). Hence, we select the last epoch student model  $g_{\theta_2}(\cdot)$  to initialise the fine-tuning stage, defined below in Sec. 3.2.

### 3.2 Fine-Tuning the Mean Teacher

To fine tune the mean teacher, we follow the approach in [12, 27] using the following loss to train the student model:

$$\ell_t(\mathcal{D}_L, \mathcal{D}_U, \theta, \theta') = \frac{1}{|\mathcal{D}_L|} \sum_{(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}_L} \ell_{cls}(\mathbf{y}_i, f_\theta(\mathbf{x}_i)) + \frac{1}{|\mathcal{D}|} \sum_{\mathbf{x}_i \in \mathcal{D}} \ell_{con}(f_\theta(\mathbf{x}_i), f_{\theta'}(\mathbf{x}_i)), \quad (2)$$

where  $\ell_{cls}(\mathbf{y}_i, f_\theta(\mathbf{x}_i)) = -\mathbf{y}_i^\top \log(f_\theta(\mathbf{x}_i))$ ,  $\ell_{con}(f_\theta(\mathbf{x}_i), f_{\theta'}(\mathbf{x}_i)) = \|f_\theta(\mathbf{x}_i) - f_{\theta'}(\mathbf{x}_i)\|^2$ , and  $\mathcal{D} = \mathcal{D}_U \cup \mathcal{D}_L^X$ . The training of the student-teacher model [12, 27, 35] is achieved by updating the student parameter using the loss in (2), as in  $\theta(t) = \theta(t-1) - \nabla_{\theta} \ell_t(\mathcal{D}_L, \mathcal{D}_U, \theta, \theta')$ , where  $t$  is the training iteration. The teacher model parameter is updated with exponential moving average (EMA) with  $\theta'(t) = \alpha\theta'(t-1) + (1-\alpha)\theta(t)$ , with  $\alpha \in [0, 1]$ . After finishing the fine-tuning stage, we select the teacher model  $f_{\theta'}(\cdot)$  to estimate the multi-label classification for test images.

## 4 Experiment

### 4.1 Dataset Setup

We use Chest X-ray14 [32], CheXpert [15] and ISIC2018 [8, 30] datasets to evaluate our method.

**Chest X-ray14** contains 112,120 chest x-ray images from 30,805 different patients. There are 14 different labels (each label represents a disease) in the dataset, where each patient can have multiple diseases at the same time, forming a multi-label classification problem. To compare with previous papers [1, 22], we adopt the official train/test data split. For the self-supervised pre-training of the mean teacher, we used all the unlabelled images (86k samples) from the training set. For the semi-supervised fine-tuning of the mean teacher, we follow

the papers [1, 22] and experiment with training sets containing different proportions of labelled data (2%, 5%, 10%, 15%, 20%). We report the classification result on the official test set (26,000 samples) using area under the ROC curve (AUC).

**CheXpert** contains around 220,000 images with 14 different diseases, and similarly to Chest X-ray14, each patient can have multiple diseases at the same time. For pre-processing, we remove all lateral view images and treat uncertain label as negative labels. We follow the semi-supervised setup from [11], and experiment with 100/200/300/400/500 labelled samples per class. We report results on the official test set using AUC.

**ISIC2018** is a multi-class skin condition dataset that contains 10,015 images with seven different labels. Each image is associated with one of the seven labels, forming a multi-class classification problem. We follow [22] train/test split for fair comparison, where the training contains 20% of the samples labelled, and the remaining 80% unlabelled. We report the AUC, Sensitivity, and F1 score results to compare with baselines.

## 4.2 Implementation Details

For all datasets, we use DenseNet121 [14] as our backbone model. For self-supervised pre-training, we follow [5] and replace the two-layer multi-layer perceptron (MLP) projection head by a three-layer MLP. For dataset pre-processing, we resized Chest X-ray14 images to  $512 \times 512$  for faster processing and CheXpert and ISIC2018 to  $224 \times 224$  for fair comparison with baselines. We use the data augmentation proposed in [4], consisting of random resize and crop, random rotation, random horizontal flipping, except for random grayscale because X-ray images are originally in grayscale. The batch size is 128 for Chest X-ray14 and 256 for CheXpert and ISIC2018, and learning rate is 0.05. For the fine-tuning stage, we use batch size 32 with 16 labelled and 16 unlabelled. The fine-tuning takes 30 epochs with learning rate decayed by 0.1 at 15 and 25 epochs for all datasets. We use the SGD optimiser with 0.9 momentum for the pre-training stage, and Adam optimiser in fine-tuning stage. The code is written in Pytorch [24]. We use 4 Nvidia Volta-100 for the self-supervised stage and 1 Nvidia RTX 2080ti for fine-tuning.

## 4.3 Experimental Results

We evaluate our approach on the official test set of ChestX-ray14 using different percentage of labelled training data (i.e., 2%, 5%, 10%, 15%, 20%), as shown in Table 1. The set of labelled data used for each percentage above follows the same strategy of previous works [1, 22]. Our S<sup>4</sup>L achieves the SOTA AUC results on all different percentages of labels. Our model surpasses the previous SOTA SRC-MT [22] by a large margin of 8.7% and 6.8% AUC for the 2% and 5% labelled set cases, respectively, where we use a backbone architecture of lower complexity (Densenet121 instead of the DenseNet169 of [22]). Using the same Densenet121 backbone, GraphXnet [1] fails to classify precisely for the 2% and 5% labelled set cases. Our method surpasses GraphXnet by more than 20% AUC in both cases.

Furthermore, we achieve the SOTA results of the field for the 10%, 15% and 20% labelled set cases, outperforming all previous semi-supervised methods [1, 22]. It is worth noting that our model trained with 5% of the labelled set achieves better results than SRC-MT with 15% of labelled. We also compare with a recently proposed self-supervised pre-training methods, MoCo V2 [6], adapted to our semi-supervised task, followed by the fine-tuning stage using different percentages of labelled data. Our method outperforms MoCo V2 by almost 10% AUC when using 2% of labelled set, and almost 3% AUC for 10% of labelled set. Our result for 20% labelled set achieves comparable 81.06% AUC performance as the supervised learning approaches – 81.20% from MoCo V2 (Densenet 121) and 81.75% from SRC-MT (Densenet 169) using 100% of the labelled samples. Such result indicates the effectiveness of our proposed S<sup>2</sup>MTS<sup>2</sup> in SSL benchmark problems.

We also show the class-level performance using 20% of the labelled data and compare with other SOTA methods in Table 2. We compare with the previous baselines, namely original mean teacher (MT) with Densenet169, SRC-MT with Densenet169, MoCo V2, and GraphXNet with Densenet121. We also train a baseline Densenet121 model with 20% labelled data using Imagenet pre-trained model. Our method achieves the best results on nine classes, surpassing the original MT [27] and its extension SRC-MT [22] by a large margin, demonstrating the effectiveness of our self-supervised learning.

Furthermore, we compare our approach on the fully-supervised Chest X-ray14 benchmark in Table 3. To the best of our knowledge, Hermoza et al. [13] has the SOTA supervised classification method containing a complex structure (relying on the weakly-supervised localisation of lesions) with a mean AUC of 82.1% (over the 14 classes), while ours reports a mean AUC of 82.5%. Hence, our model, using the whole labelled set, achieves the SOTA performance on 8 classes and an average that surpasses the previous supervised methods by a minimum of 0.4% and a maximum of 8% AUC.

The results on CheXpert and ISIC2018 datasets are shown in Tables 4 and 5, respectively. In particular, for CheXpert in Table 4, we compare our method with LatentMixing [11] and our result is better in all cases. For ISIC2018 on Table 5, using the test set from SRC-MT [22], our method outperforms all baselines (Supervised, MT, and SRC-MT) for all measures.

#### 4.4 Ablation Study

We study the impact of different components of our proposed S<sup>2</sup>MTS<sup>2</sup> in Table 6 using Chest X-Ray14. Using the proposed self-supervised learning with just the student model, our model achieves at least 72.95% mean AUC on various percentages of labelled training data. Adding the JCL component improves the baseline by around 1% mean AUC on each training percentage. Adding the mean teacher boosts the result by 1.5% to 2% mean AUC on each training percentage. The combination of all our proposed three components achieves SOTA performance on semi-supervised task.

**Table 1.** Mean AUC result over the 14 disease classes of Chest X-Ray14 for different label set training percentages. \* indicates the methods that use Densenet169 as backbone architecture.

Label Percentage	2%	5%	10%	15%	20%	100%
Graph XNet* [1]	53.00	58.00	63.00	68.00	78.00	N/A
SRC-MT* [22]	66.95	72.29	75.28	77.76	79.23	81.75
NoTeacher [31]	72.60	77.04	77.61	N/A	79.49	N/A
MOCO V2 [6]	65.97	73.84	77.07	79.37	80.17	81.20
Ours	<b>74.69</b>	<b>78.96</b>	<b>79.90</b>	<b>80.31</b>	<b>81.06</b>	<b>82.51</b>

**Table 2.** Class-level AUC comparison between our S<sup>2</sup>MTS<sup>2</sup> and other semi-supervised SOTA approaches trained with **20% of labelled data** on Chest X-Ray14. \* denotes the methods that use Densenet-169 as backbone.

Method	Densenet-121	GraphXNet [1]	MOCO V2 [6]	MT [22] *	SRC-MT [22] *	Ours
Atelectasis	75.75	71.89	77.21	75.12	75.38	<b>78.57</b>
Cardiomegaly	80.71	87.99	85.84	87.37	87.7	<b>88.08</b>
Effusion	79.87	79.2	81.62	80.81	81.58	<b>82.87</b>
Infiltration	69.16	<b>72.05</b>	70.91	70.67	70.4	70.68
Mass	78.40	80.9	81.71	77.72	78.03	<b>82.57</b>
Nodule	74.49	71.13	<b>76.72</b>	73.27	73.64	76.60
Pneumonia	69.55	<b>76.64</b>	71.08	69.17	69.27	72.25
Pneumothorax	84.70	83.7	85.92	85.63	86.12	<b>86.55</b>
Consolidation	71.85	73.36	74.47	72.51	73.11	<b>75.47</b>
Edema	81.61	80.2	83.57	82.72	82.94	<b>84.83</b>
Emphysema	89.75	84.07	91.10	88.16	88.98	<b>91.88</b>
Fibrosis	79.30	80.34	80.96	78.24	79.22	<b>81.73</b>
Pleural Thicken	73.46	75.7	75.65	74.43	75.63	<b>76.86</b>
Hernia	86.05	87.22	85.62	<b>87.74</b>	87.27	85.98
Mean	78.19	78.88	80.17	78.83	79.23	<b>81.06</b>

**Table 3.** Class-level AUC comparison between our S<sup>2</sup>MTS<sup>2</sup> and other supervised SOTA approaches trained with **100% of labelled data** on Chest X-Ray14.

Method	Wang et al. [32]	Li et al. [20]	CheXNet [25]	CRAL [10]	Ma et al. [23]	Herzoza et al. [13]	Ours
Atelectasis	70	72.9	75.5	78.1	77.7	77.5	<b>78.7</b>
Cardiomegaly	81	84.6	86.7	88.3	<b>89.4</b>	88.1	87.4
Effusion	75.9	78.1	81.5	83.1	82.9	83.1	<b>83.8</b>
Infiltration	66.1	67.3	69.4	69.7	69.6	69.5	<b>70.9</b>
Mass	69.3	74.3	80.2	83	<b>83.8</b>	82.6	83.3
Nodule	66.9	75.8	73.5	76.4	77.1	78.9	<b>79.9</b>
Pneumonia	65.8	63.3	69.8	72.5	72.2	<b>74.1</b>	73.9
Pneumothorax	79.9	79.3	82.8	86.6	86.2	<b>87.9</b>	87.1
Consolidation	70.3	72	72.2	75.8	75	74.7	<b>75.9</b>
Edema	80.5	71	83.5	85.3	84.6	<b>84.6</b>	84.5
Emphysema	83.3	75.1	85.6	91.1	90.8	93.6	<b>93.7</b>
Fibrosis	78.6	76.1	80.3	82.6	82.7	83.3	<b>83.4</b>
Pleural Thicken	68.4	73	74.9	78	77.9	79.3	<b>79.3</b>
Hernia	87.2	66.8	89.4	91.8	<b>93.4</b>	91.7	93.3
Mean	74.5	73.9	78.9	81.6	81.7	82.1	<b>82.5</b>



**Table 4.** Mean AUC result (over the 14 disease classes) on CheXpert for different number of training samples per class.

Labelled	100	200	300	400	500
LatentMixing [11]	65.12	66.41	67.39	67.96	68.47
Ours	<b>66.15</b>	<b>67.85</b>	<b>70.83</b>	<b>71.37</b>	<b>71.58</b>

**Table 5.** AUC, Sensitivity and F1 result on ISIC2018 using 20% of labelled training samples.

Method	AUC	Sensitivity	F1
Supervised	90.15	65.50	52.03
MT	92.96	69.75	59.10
SRC-MT [22]	93.58	71.47	60.68
Ours	<b>94.71</b>	<b>72.14</b>	<b>62.67</b>

**Table 6.** Ablation studies of our method with different components on Chest X-Ray14. “Self-supervised” indicates the traditional self-supervised learning with contrastive loss [12]. “JCL” replaces contrastive loss with (1), “MT” stands for fine-tuned with student-teacher learning instead only fine-tuned on only labelled samples.

Self-supervised	JCL	MT	AUC (2%)	AUC (5%)	AUC (10%)	AUC (15%)	AUC (20%)
✓			72.95	76.82	78.54	79.28	80.14
✓	✓		73.60	77.46	79.18	79.83	80.62
✓		✓	74.80	77.66	79.08	79.70	80.57
✓	✓	✓	75.69	78.96	79.90	80.31	81.06

## 5 Conclusion

In this paper, we presented a novel semi-supervised framework, the Self-supervised Mean Teacher for Semi-supervised ( $S^2MTS^2$ ) learning. The main contribution of  $S^2MTS^2$  is the self-supervised mean teacher pre-trained based on joint contrastive learning [3], using an infinite number of pairs of positive query and key features. This model is then fine-tuned with the exponential moving average teacher framework.  $S^2MTS^2$  is validated on the thorax disease multi-label classification problem from the datasets Chest X-ray14 [32] and CheXpert [15], and the multi-class classification from the skin condition dataset ISIC2018 [8, 30]. The experiments show that our method outperforms the previous SOTA semi-supervised learning methods by a large margin in all benchmarks containing a varying percentage of labelled data. We also show that the method holds the SOTA results on Chest X-ray14 [32] even for the fully-supervised problem. The ablation study shows the importance of three main components of the method, namely self-supervised learning, JCL, and the mean-teacher model. We will investigate the performance of our method on other semi-supervised medical imaging benchmarks in the future.

**Acknowledgement.** This work was supported by Australian Research Council through grants DP180103232 and FT190100525.

## References

1. Aviles-Rivero, A.I., et al.: GraphXNET - chest X-ray classification under extreme minimal supervision. arXiv preprint [arXiv:1907.10085](https://arxiv.org/abs/1907.10085) (2019)
2. Berthelot, D., et al.: ReMixMatch: semi-supervised learning with distribution alignment and augmentation anchoring. arXiv preprint [arXiv:1911.09785](https://arxiv.org/abs/1911.09785) (2019)
3. Cai, Q., Wang, Y., Pan, Y., Yao, T., Mei, T.: Joint contrastive learning with infinite possibilities. arXiv preprint [arXiv:2009.14776](https://arxiv.org/abs/2009.14776) (2020)
4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International Conference on Machine Learning, pp. 1597–1607. PMLR (2020)
5. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.: Big self-supervised models are strong semi-supervised learners. arXiv preprint [arXiv:2006.10029](https://arxiv.org/abs/2006.10029) (2020)
6. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint [arXiv:2003.04297](https://arxiv.org/abs/2003.04297) (2020)
7. Cheplygina, V., de Bruijne, M., Pluim, J.P.: Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *Med. Image Anal.* **54**, 280–296 (2019)
8. Codella, N., et al.: Skin lesion analysis toward melanoma detection 2018: a challenge hosted by the international skin imaging collaboration (ISIC). arXiv preprint [arXiv:1902.03368](https://arxiv.org/abs/1902.03368) (2019)
9. Cui, W., et al.: Semi-supervised brain lesion segmentation with an adapted mean teacher model. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) IPMI 2019. LNCS, vol. 11492, pp. 554–565. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-20351-1\\_43](https://doi.org/10.1007/978-3-030-20351-1_43)
10. Guan, Q., Huang, Y.: Multi-label chest X-ray image classification via category-wise residual attention learning. *Pattern Recogn. Lett.* **130**, 259–266 (2020)
11. Gyawali, P.K., Ghimire, S., Bajracharya, P., Li, Z., Wang, L.: Semi-supervised medical image classification with global latent mixing. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12261, pp. 604–613. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-59710-8\\_59](https://doi.org/10.1007/978-3-030-59710-8_59)
12. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9729–9738 (2020)
13. Hermoza, R., Maicas, G., Nascimento, J.C., Carneiro, G.: Region proposals for saliency map refinement for weakly-supervised disease localisation and classification. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12266, pp. 539–549. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-59725-2\\_52](https://doi.org/10.1007/978-3-030-59725-2_52)
14. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2017)
15. Irvin, J., et al.: CheXpert: a large chest radiograph dataset with uncertainty labels and expert comparison. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, pp. 590–597 (2019)
16. Jonmohamadi, Y., et al.: Automatic segmentation of multiple structures in knee arthroscopy using deep learning. *IEEE Access* **8**, 51853–51861 (2020)
17. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. arXiv preprint [arXiv:1610.02242](https://arxiv.org/abs/1610.02242) (2016)
18. Lee, D.H., et al.: Pseudo-label: the simple and efficient semi-supervised learning method for deep neural networks. In: Workshop on Challenges in Representation Learning, ICML, vol. 3 (2013)

19. Lee, J.G., et al.: Deep learning in medical imaging: general overview. *Korean J. Radiol.* **18**(4), 570 (2017)
20. Li, Z., et al.: Thoracic disease identification and localization with limited supervision. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8290–8299 (2018)
21. Liu, F., Jonmohamadi, Y., Maicas, G., Pandey, A.K., Carneiro, G.: Self-supervised depth estimation to regularise semantic segmentation in knee arthroscopy. In: Martel, A.L., et al. (eds.) *MICCAI 2020. LNCS*, vol. 12261, pp. 594–603. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-59710-8\\_58](https://doi.org/10.1007/978-3-030-59710-8_58)
22. Liu, Q., et al.: Semi-supervised medical image classification with relation-driven self-ensembling model. *IEEE Trans. Med. Imaging* **39**(11), 3429–3440 (2020)
23. Ma, C., Wang, H., Hoi, S.C.H.: Multi-label thoracic disease image classification with cross-attention networks. In: Shen, D., et al. (eds.) *MICCAI 2019. LNCS*, vol. 11769, pp. 730–738. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-32226-7\\_81](https://doi.org/10.1007/978-3-030-32226-7_81)
24. Paszke, A., et al.: Pytorch: an imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703* (2019)
25. Rajpurkar, P., et al.: ChexNet: radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225* (2017)
26. Rizve, M.N., Duarte, K., Rawat, Y.S., Shah, M.: In defense of pseudo-labeling: an uncertainty-aware pseudo-label selection framework for semi-supervised learning. *arXiv preprint arXiv:2101.06329* (2021)
27. Tarvainen, A., Valpola, H.: Mean teachers are better role models: weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780* (2017)
28. Tian, Yu., Maicas, G., Pu, L.Z.C.T., Singh, R., Verjans, J.W., Carneiro, G.: Few-shot anomaly detection for polyp frames from colonoscopy. In: Martel, A.L., et al. (eds.) *MICCAI 2020. LNCS*, vol. 12266, pp. 274–284. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-59725-2\\_27](https://doi.org/10.1007/978-3-030-59725-2_27)
29. Tian, Y., Pu, L.Z., Singh, R., Burt, A.D., Carneiro, G.: One-stage five-class polyp detection and classification. In: *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pp. 70–73. IEEE (2019)
30. Tschandl, P., Rosendahl, C., Kittler, H.: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Sci. Data* **5**(1), 1–9 (2018)
31. Unnikrishnan, B., Nguyen, C.M., Balaram, S., Foo, C.S., Krishnaswamy, P.: Semi-supervised classification of diagnostic radiographs with NoTeacher: a teacher that is not mean. In: Martel, A.L., et al. (eds.) *MICCAI 2020. LNCS*, vol. 12261, pp. 624–634. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-59710-8\\_61](https://doi.org/10.1007/978-3-030-59710-8_61)
32. Wang, X., Peng, Y.a.: ChestX-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: *CVPR* (2017)
33. Yang, Y., Xu, Z.: Rethinking the value of labels for improving class-imbalanced learning. *arXiv preprint arXiv:2006.07529* (2020)
34. Zhai, X., Oliver, A., Kolesnikov, A., Beyler, L.: S4L: self-supervised semi-supervised learning. In: *ICCV*, pp. 1476–1485 (2019)
35. Zhou, H.-Y., Yu, S., Bian, C., Hu, Y., Ma, K., Zheng, Y.: Comparing to learn: surpassing ImageNet pretraining on radiographs by comparing image representations. In: Martel, A.L., et al. (eds.) *MICCAI 2020. LNCS*, vol. 12261, pp. 398–407. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-59710-8\\_39](https://doi.org/10.1007/978-3-030-59710-8_39)