

# ACPL: Anti-curriculum Pseudo-labelling for Semi-supervised Medical Image Classification

Fengbei Liu<sup>1\*</sup> Yu Tian<sup>1\*</sup> Yuanhong Chen<sup>1</sup> Yuyuan Liu<sup>1</sup>  
 Vasileios Belagiannis<sup>2</sup> Gustavo Carneiro<sup>1</sup>

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


<sup>2</sup> Universität Ulm, Germany

## Abstract


Effective semi-supervised learning (SSL) in medical image analysis (MIA) must address two challenges: 1) work effectively on both multi-class (e.g., lesion classification) and multi-label (e.g., multiple-disease diagnosis) problems, and 2) handle imbalanced learning (because of the high variance in disease prevalence). One strategy to explore in SSL MIA is based on the pseudo labelling strategy, but it has a few shortcomings. Pseudo-labelling has in general lower accuracy than consistency learning, it is not specifically design for both multi-class and multi-label problems, and it can be challenged by imbalanced learning. In this paper, unlike traditional methods that select confident pseudo label by threshold, we propose a new SSL algorithm, called anti-curriculum pseudo-labelling (ACPL), which introduces novel techniques to select informative unlabelled samples, improving training balance and allowing the model to work for both multi-label and multi-class problems, and to estimate pseudo labels by an accurate ensemble of classifiers (improving pseudo label accuracy). We run extensive experiments to evaluate ACPL on two public medical image classification benchmarks: Chest X-Ray14 for thorax disease multi-label classification and ISIC2018 for skin lesion multi-class classification. Our method outperforms previous SOTA SSL methods on both datasets<sup>12</sup>.

## 1. Introduction

Deep learning has shown outstanding results in medical image analysis (MIA) [24, 34, 35]. Compared to computer vision, the labelling of MIA training sets by medical experts is significantly more expensive, resulting in low availability of labelled images, but the high availability of unlabelled



(a) Diagram of our ACPL (top) and traditional pseudo-label SSL (bottom)



(b) Imbalanced distribution on multi-label Chest X-ray14 [39] (left) and multi-class ISIC2018 [36] (right)

Figure 1. In (a), we show diagrams of the proposed ACPL (top) and the traditional pseudo-label SSL (bottom) methods, and (b) displays histograms of images per label for the multi-label Chest X-ray14 [39] (left) and multi-class ISIC2018 [36] (right).

images from clinics and hospitals databases can be explored in the modelling of deep learning classifiers. Furthermore, differently from computer vision problems that tend to be mostly multi-class and balanced, MIA has a number of multi-class (e.g., a lesion image of a single class) and multi-label (e.g., an image from a patient can contain multiple diseases) problems, where both problems usually contain severe class imbalances because of the variable prevalence of diseases (see Fig. 1-(b)). Hence, MIA semi-supervised learning (SSL) methods need to be flexible enough to work with multi-label and multi-class problems, in addition to handle imbalanced learning.

State-of-the-art (SOTA) SSL approaches are usually based on the consistency learning of unlabelled data [5, 6, 32] and self-supervised pre-training [25]. Even though consistency-based methods show SOTA results on multi-class SSL problems, pseudo-labelling methods have shown

\*First two authors contributed equally to this work.

<sup>1</sup>Supported by Australian Research Council through grants DP180103232 and FT190100525.

<sup>2</sup>Code is available at <https://github.com/FBLADL/ACPL>

better results for multi-label SSL problems [29]. Pseudo-labelling methods provide labels to confidently classified unlabelled samples that are used to re-train the model [22]. One issue with pseudo-labelling SSL methods is that the confidently classified unlabelled samples represent the least informative ones [30] that, for imbalanced problems, are likely to belong to the majority classes. Hence, this will bias the classification toward the majority classes and most likely deteriorate the classification accuracy of the minority classes. Also, selecting confident pseudo-labelled samples is challenging in multi-class, but even more so in multi-label problems. Previous papers [2, 29] use a fixed threshold for all classes, but a class-wise threshold that addresses imbalanced learning and correlations between classes in multi-label problems would enable more accurate pseudo-label predictions. However, such class-wise threshold is hard to estimate without knowing the class distributions or if we are dealing with a multi-class or multi-label problem. Furthermore, using the model output for the pseudo-labelling process can also cause confirmation bias [1], whereby the assignment of incorrect pseudo-labels will increase the model confidence in those incorrect predictions, and consequently decrease the model accuracy.

In this paper, we propose the **anti-curriculum pseudo-labelling (ACPL)**, which addresses multi-class and multi-label imbalanced learning SSL MIA problems. First, we introduce a new approach to select the most informative unlabelled images to be pseudo-labelled. This is motivated by our argument that there exists a distribution shift between unlabelled and labelled samples for SSL. An effective learning curriculum must focus on informative unlabelled samples that are located as far as possible from the distribution of labelled samples. As a result, these informative samples are likely to belong to the minority classes in MIA imbalanced learning problems. Selecting these informative samples will naturally balance the training process and, given that they are selected before the pseudo-labelling process, we eliminate the need for estimating a class-wise classification threshold, facilitating our model to work well on multi-class and multi-label problems. The information content measure of an unlabelled sample is computed with our proposed cross-distribution sample informativeness that outputs how close an unlabelled sample is from the set of labelled anchor samples (anchor samples are highly informative labelled samples). Second, we introduce a new pseudo-labelling mechanism, called **informative mixup**, which combines the model classification with a K-nearest neighbor (KNN) classification guided by sample informativeness to improve prediction accuracy and mitigate confirmation bias. Third, we propose the **anchor set purification** method that selects the most informative pseudo-labelled samples to be included in the labelled anchor set to improve the pseudo-labelling accuracy of the KNN classi-

fier in later training stages.

To summarise, our ACPL approach selects highly informative samples for pseudo-labelling (addressing MIA imbalanced classification problems and allowing multi-label multi-class modelling) and uses an ensemble of classifiers to produce accurate pseudo labels (tackling confirmation bias to improve classification accuracy), where the main technical contributions are:

- A novel information content measure to select informative unlabelled samples named **cross-distribution sample informativeness**;
- A new pseudo-labelling mechanism, called **informative mixup**, which generates pseudo labels from an ensemble of deep learning and KNN classifiers; and
- A novel method, called **anchor set purification (ASP)**, to select informative pseudo-labelled samples to be included in the labelled anchor set to improve the pseudo-labelling accuracy of the KNN classifier.

We evaluate ACPL on two publicly available medical image classification datasets, namely the Chest X-Ray14 for thorax disease multi-label classification [39] and the ISIC2018 for skin lesion multi-class classification [8, 36]. Our method outperforms the current SOTA methods in both datasets.

## 2. Related Work

We first review consistency-based and pseudo-labelling SSL methods. Then, we discuss the curriculum and anti-curriculum learning literature for fully and semi-supervised learning and present relevant SSL MIA methods.

**Consistency-based SSL** optimises the classification prediction of labelled images and minimises the prediction outputs of different views of unlabelled images, where these views are obtained from different types of image perturbations, such as spatial/temporal [21, 33], adversarial [27], or data augmentation [5, 6, 32]. The performance of the consistency-based methods can be further improved with self-supervised pre-training [25]. Even though consistency-based SSL methods show SOTA results in many benchmarks [32], they depend on a careful design of perturbation functions that requires domain knowledge and would need to be adapted to each new type of medical imaging. Furthermore, Rizve et al. [29] show that pseudo-labelling SSL methods are more accurate for multi-label problems.

**Pseudo-labelling SSL** methods [7, 29, 31, 41] train a model with the available labelled data, estimate the pseudo labels of unlabelled samples classified with high confidence [22], then take these pseudo-labelled samples to re-train the model. As mentioned above in Sec. 1 pseudo-label SSL approaches can bias classification toward the majority classes in imbalanced problems, is not seamlessly adaptable to multi-class and multi-label problems, and can also

lead to confirmation bias. We argue that the improvement of pseudo-labelling SSL methods depends on the selection of informative unlabelled samples to address the majority class bias and the adaptation to multi-class and multi-label problems, and an accurate pseudo-labelling mechanism to handle confirmation bias, which are two points that we target with this paper.

The selection of training samples based on their information content has been studied by fully supervised **curriculum and anti-curriculum learning methods** [40]. Curriculum learning focuses on the easy samples in the early training stages and gradually includes the hard samples in the later training stages, where easy samples [4, 15, 20] are usually defined as samples that have small losses during training, and hard samples tend to have large losses. On the other hand, anti-curriculum focuses on the hard samples first and transitions to the easy samples later in the training [14, 17]. The methods above have been designed to work in fully supervised learning. Cascante et al. [7] explored a pseudo labelling SSL method based on curriculum learning, but we are not aware of SSL methods that explore anti-curriculum learning. Since we target accurate SSL of imbalanced multi-class and multi-label methods, we follow anti-curriculum learning that pseudo-labels the most informative samples which are likely to belong to the minority classes (consequently, helping to balance the training) and enable the selection of samples without requiring the estimation of a class-wise classification threshold (enabling a seamless adaptation to multi-class and multi-label problems).

The main benchmarks for SSL in MIA study the multi-label classification of chest X-ray (CXR) images [13, 39] and multi-class classification of skin lesions [8, 36]. For **CXR SSL** classification, pseudo-labelling methods have been explored [2], but SOTA results are achieved with consistency learning approaches [9, 23, 25, 26, 37]. For **skin lesion SSL** classification, the current SOTA is also based on consistency learning [26], with pseudo-labelling approaches [3] not being competitive. We show that our proposed pseudo-labelling method ACPL can surpass the consistency-based SOTA on both benchmarks, demonstrating the value of selecting highly informative samples for pseudo labelling and of the accurate pseudo labels from the ensemble of classifiers. We also show that our ACPL improves the current computer vision SOTA [29] applied to MIA, demonstrating the limitation of computer vision methods in MIA and also the potential of our approach to be applied in more general SSL problems.

### 3. Methods

To introduce our SSL method ACPL, assume that we have a small labelled training set  $\mathcal{D}_L = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{D}_L|}$ , where  $\mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^{H \times W \times C}$  is the input image of size

---

#### Algorithm 1 Anti-curriculum Pseudo-labelling Algorithm

---

- 1: **require:** Labelled set  $\mathcal{D}_L$ , unlabelled set  $\mathcal{D}_U$ , and number of training stages  $T$
  - 2: **initialise**  $\mathcal{D}_A = \mathcal{D}_L$ , and  $t = 0$
  - 3: **warm-up train**  $p_{\theta_t}(\mathbf{x})$  **with**  
 $\theta_t = \arg \min_{\theta} \frac{1}{|\mathcal{D}_L|} \sum_{(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}_L} \ell(\mathbf{y}_i, p_{\theta}(\mathbf{x}_i))$
  - 4: **while**  $t < T$  **or**  $|\mathcal{D}_U| \neq 0$  **do**
  - 5:   **build pseudo-labelled dataset using CDSI from (2) and IM from (6):**  
 $\mathcal{D}_S = \{(\mathbf{x}, \tilde{\mathbf{y}}) | \mathbf{x} \in \mathcal{D}_U, h(f_{\theta_t}(\mathbf{x}), \mathcal{D}_A) = 1, \tilde{\mathbf{y}} = g(f_{\theta_t}(\mathbf{x}), \mathcal{D}_A)\}$
  - 6:   **update anchor set with ASP from (7):**  
 $\mathcal{D}_A = \mathcal{D}_A \cup \{(\mathbf{x}, \tilde{\mathbf{y}}), \text{ where } (\mathbf{x}, \tilde{\mathbf{y}}) \in \mathcal{D}_S, \text{ and } a(f_{\theta_t}(\mathbf{x}), \mathcal{D}_U, \mathcal{D}_A) = 1\}$
  - 7:    $t \leftarrow t + 1$
  - 8:   **optimise (1) using**  $\mathcal{D}_L, \mathcal{D}_S$  **to obtain**  $p_{\theta_t}(\mathbf{x})$
  - 9:   **update labelled and unlabelled sets:**  
 $\mathcal{D}_L \leftarrow \mathcal{D}_L \cup \mathcal{D}_S, \mathcal{D}_U \leftarrow \mathcal{D}_U \setminus \mathcal{D}_S$
  - 10: **end while**
  - 11: **return**  $p_{\theta_t}(\mathbf{x})$
- 

$H \times W$  with  $C$  colour channels, and  $\mathbf{y}_i \in \{0, 1\}^{|\mathcal{Y}|}$  is the label with the set of classes denoted by  $\mathcal{Y} = \{1, \dots, |\mathcal{Y}|\}$  (note that  $\mathbf{y}_i$  is a one-hot vector for multi-class problems and a binary vector in multi-label problems). A large unlabelled training set  $\mathcal{D}_U = \{\mathbf{x}_i\}_{i=1}^{|\mathcal{D}_U|}$  is also provided, with  $|\mathcal{D}_L| \ll |\mathcal{D}_U|$ . We assume the samples from both datasets are drawn from the same (latent) distribution. Our algorithm also relies on the pseudo-labelled set  $\mathcal{D}_S$  that is composed of pseudo-labelled samples classified as informative unlabelled samples, and an anchor set  $\mathcal{D}_A$  that contains informative pseudo-labelled samples. The goal of ACPL is to learn a model  $p_{\theta} : \mathcal{X} \rightarrow [0, 1]^{|\mathcal{Y}|}$  parameterised by  $\theta$  using the labelled, unlabelled, pseudo-labelled, and anchor datasets.

Below, in Sec. 3.1, we introduce our ACPL optimisation that produces accurate pseudo labels to unlabelled samples following an anti-curriculum strategy, where highly informative unlabelled samples are selected to be pseudo-labelled at each training stage. In Sec. 3.2, we present the information criterion of an unlabelled sample, referred to as *cross distribution sample informativeness (CDSI)*, based on the dissimilarity between the unlabelled sample and samples in the anchor set  $\mathcal{D}_A$ . The pseudo labels for the informative unlabelled samples are generated using the proposed *informative mixup (IM)* method (Sec. 3.3) that mixes up the results from the model  $p_{\theta}(\cdot)$  and a  $K$  nearest neighbor (KNN) classifier using the anchor set. At the end of each training stage, the anchor set is updated with the *anchor set purification (ASP)* method (Sec. 3.4) that only keeps

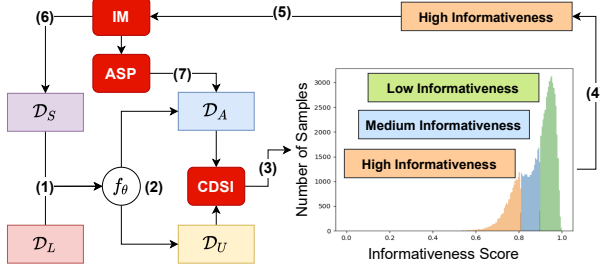


Figure 2. Anti-curriculum pseudo-labelling (ACPL) algorithm. The algorithm is divided into the following iterative steps: 1) train the model with  $\mathcal{D}_S$  and  $\mathcal{D}_L$ ; 2) extract the features from the anchor and unlabelled samples; 3) estimate information content of unlabelled samples with CDSI from (4) with anchor set  $\mathcal{D}_A$ ; 4) partition the unlabelled samples into high, medium and low information content using (2); 5) assign a pseudo label to high information content unlabelled samples with IM from (6); 6) update  $\mathcal{D}_S$  with new pseudo-labelled samples; and 7) update  $\mathcal{D}_A$  with ASP in (7).

the most informative subset of pseudo-labelled samples, according to the CDSI criterion.

### 3.1. ACPL Optimisation

Our ACPL optimisation, described in Alg. 1 and depicted by Fig. 2, starts with a warm-up supervised training of the parameters of the model  $p_\theta(\cdot)$  using only the labelled set  $\mathcal{D}_L$ . For the rest of the training, we use the sets of labelled and unlabelled samples,  $\mathcal{D}_L$  and  $\mathcal{D}_U$ , and update the pseudo-labelled set  $\mathcal{D}_S$  and the anchor set  $\mathcal{D}_A$  containing the informative unlabelled and pseudo-labelled samples, where  $\mathcal{D}_S$  start as an empty set and  $\mathcal{D}_A$  starts with the samples in  $\mathcal{D}_L$ . The optimisation iteratively minimises the following cost function:

$$\ell_{ACPL}(\theta, \mathcal{D}_L, \mathcal{D}_S) = \frac{1}{|\mathcal{D}_L|} \sum_{(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}_L} \ell(\mathbf{y}_i, p_\theta(\mathbf{x}_i)) + \frac{1}{|\mathcal{D}_S|} \sum_{(\mathbf{x}_i, \tilde{\mathbf{y}}_i) \in \mathcal{D}_S} \ell(\tilde{\mathbf{y}}_i, p_\theta(\mathbf{x}_i)), \quad (1)$$

where  $\ell(\cdot)$  denotes a classification loss (e.g., cross-entropy),  $\theta$  is the model parameter,  $\mathbf{y}_i$  is the ground truth, and  $\tilde{\mathbf{y}}_i$  is the estimated pseudo label. After optimising (1), the labelled and unlabelled sets are updated as  $\mathcal{D}_L = \mathcal{D}_L \cup \mathcal{D}_S$  and  $\mathcal{D}_U = \mathcal{D}_U \setminus \mathcal{D}_S$ , and a new iteration of optimisation takes place.

### 3.2. Cross Distribution Sample Informativeness (CDSI)

The function that estimates if an unlabelled sample has high information content is defined by

$$h(f_\theta(\mathbf{x}), \mathcal{D}_A) = \begin{cases} 1, & p_\gamma(\zeta = \text{high}|\mathbf{x}, \mathcal{D}_A) > \tau, \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

where  $\zeta \in \mathcal{Z} = \{\text{low}, \text{medium}, \text{high}\}$  represents the information content random variable,  $\gamma = \{\mu_\zeta, \Sigma_\zeta, \pi_\zeta\}_{\zeta \in \mathcal{Z}}$  denotes the parameters of the Gaussian Mixture Model (GMM)  $p_\gamma(\cdot)$ , and  $\tau = \max \{p_\gamma(\zeta = \text{low}|\mathbf{x}, \mathcal{D}_A), p_\gamma(\zeta = \text{medium}|\mathbf{x}, \mathcal{D}_A)\}$ . The function  $p_\gamma(\zeta|\mathbf{x}, \mathcal{D}_A)$  can be decomposed into  $p_\gamma(\mathbf{x}|\zeta, \mathcal{D}_A)p_\gamma(\zeta|\mathcal{D}_A)/p_\gamma(\mathbf{x}|\mathcal{D}_A)$ , where

$$p_\gamma(\mathbf{x}|\zeta, \mathcal{D}_A) = n(d(f_\theta(\mathbf{x}), \mathcal{D}_A)|\mu_\zeta, \Sigma_\zeta), \quad (3)$$

with  $n(\cdot; \mu_\zeta, \Sigma_\zeta)$  denoting a Gaussian function with mean  $\mu_\zeta$  and covariance  $\Sigma_\zeta$ ,  $p_\gamma(\zeta|\mathcal{D}_A) = \pi_\zeta$  representing the ownership probability of  $\zeta$  (i.e., the weight of mixture  $\zeta$ ), and  $p_\gamma(\mathbf{x}|\mathcal{D}_A)$  being a normalisation factor. The probability in (3) is computed with the density of the unlabelled sample  $\mathbf{x}$  with respect to the anchor set  $\mathcal{D}_A$ , as follows:

$$d(f_\theta(\mathbf{x}), \mathcal{D}_A) = \frac{1}{K} \sum_{\substack{(f_\theta(\mathbf{x}_A), \mathbf{y}_A) \in \mathcal{N}(f_\theta(\mathbf{x}), \mathcal{D}_A)}} \frac{f_\theta(\mathbf{x})^\top f_\theta(\mathbf{x}_A)}{\|f_\theta(\mathbf{x})\|_2 \|f_\theta(\mathbf{x}_A)\|_2}, \quad (4)$$

where  $\mathcal{N}(f_\theta(\mathbf{x}), \mathcal{D}_A)$  represents the set of K-nearest neighbors (KNN) from the anchor set  $\mathcal{D}_A$  to the input image feature  $f_\theta(\mathbf{x})$ , with each element in the set  $\mathcal{D}_A$  denoted by  $(f_\theta(\mathbf{x}_A), \mathbf{y}_A)$ . The  $F$ -dimensional input image feature is extracted with  $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^F$  from the model  $p_\theta(\cdot)$  with  $p_\theta(\mathbf{x}) = \sigma(f_\theta(\mathbf{x}))$ , where  $\sigma(\cdot)$  is the final activation function to produce an output in  $[0, 1]^{|Y|}$ . The parameters  $\gamma$  in (2) are estimated with the expectation-maximisation (EM) algorithm [10], every time after the anchor set is updated.

### 3.3. Informative Mixup (IM)

After selecting informative unlabelled samples with (2), we aim to produce reliable pseudo labels for them. We can provide two pseudo labels for each unlabelled sample  $\mathbf{x} \in \mathcal{D}_U$ : the model prediction from  $p_\theta(\mathbf{x})$ , and the K-nearest neighbor (KNN) prediction using the anchor set, as follows:

$$\tilde{\mathbf{y}}_{\text{model}}(\mathbf{x}) = p_\theta(\mathbf{x}), \quad \tilde{\mathbf{y}}_{\text{KNN}}(\mathbf{x}) = \frac{1}{K} \sum_{(f_\theta(\mathbf{x}_A), \mathbf{y}_A) \in \mathcal{N}(f_\theta(\mathbf{x}), \mathcal{D}_A)} \mathbf{y}_A. \quad (5)$$

$\mathbf{y}_A$  is the label of anchor set samples. However, using any of the pseudo labels from (5) can be problematic for model training. The pseudo label in  $\tilde{\mathbf{y}}_{\text{model}}(\mathbf{x})$  can cause confirmation bias, and the reliability of  $\tilde{\mathbf{y}}_{\text{KNN}}(\mathbf{x})$  depends on the size and representativeness of the initial labelled set to produce accurate classification. Inspired by MixUp [42], we propose the **informative mixup** method that constructs the pseudo-labelling function  $g(\cdot)$  in (1) with a linear combination of  $\tilde{\mathbf{y}}_{\text{model}}(\mathbf{x})$  and  $\tilde{\mathbf{y}}_{\text{KNN}}(\mathbf{x})$  weighted by the density



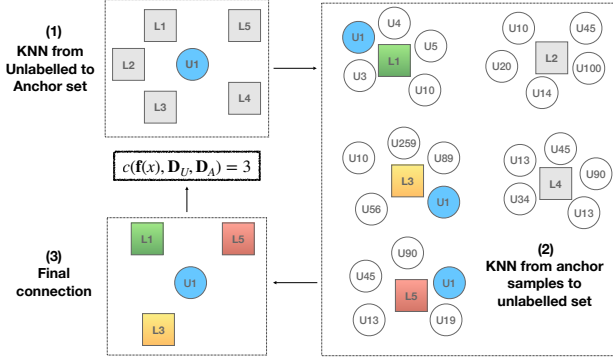


Figure 3. **ASP**: 1) find KNN samples from an informative unlabelled sample to the anchor set  $\mathcal{D}_A$ ; 2) find KNN samples from each anchor sample of (1) to the unlabelled set  $\mathcal{D}_U$ ; and 3) calculate the number of surviving nearest neighbours. Samples with the smallest values of  $c(\cdot)$  are selected to be inserted into  $\mathcal{D}_A$ .

score from (4), as follows:

$$\tilde{y} = g(f_\theta(\mathbf{x}), \mathcal{D}_A) = d(f_\theta(\mathbf{x}), \mathcal{D}_A) \times \tilde{y}_{\text{model}}(\mathbf{x}) + (1 - d(f_\theta(\mathbf{x}), \mathcal{D}_A)) \times \tilde{y}_{\text{KNN}}(\mathbf{x}). \quad (6)$$

The informative mixup in (6) is different from MixUp [42] because it combines the classification results of the same image from two models instead of the classification from the same model of two images. Furthermore, our informative mixup weights the classifiers with the density score to reflect the trade-off between  $\tilde{y}_{\text{model}}(\mathbf{x})$  and  $\tilde{y}_{\text{KNN}}(\mathbf{x})$ . Since informative samples are selected from a region of the anchor set with low feature density, the KNN prediction  $\tilde{y}_{\text{KNN}}(\mathbf{x})$  is less reliable than  $\tilde{y}_{\text{model}}(\mathbf{x})$ , so by default, we should trust more the model classification. The weighting between the two predictions in (6) reflects this observation, where  $\tilde{y}_{\text{model}}(\mathbf{x})$  will tend have a larger weight given that  $d(f_\theta(\mathbf{x}), \mathcal{D}_A)$  is usually larger than 0.5, as displayed in Fig. 2 (see the informativeness score histogram at the bottom-right corner). When the sample is located in a high-density region, we place most of the weight on the model prediction given that in such case, the model is highly reliable. On the other hand, when the sample is in a low-density region, we try to balance a bit more the contribution of both the model and KNN predictions, given the low reliability of the model.

### 3.4. Anchor Set Purification (ASP)

After estimating the pseudo label for informative unlabelled samples, we aim to update the anchor set with informative pseudo-labelled samples to maintain density score from (4) accurate in later training stages. However, adding all pseudo-labelled samples will cause anchor set over-sized and increase hyper-parameter sensitivity. Thus, we propose

the Anchor Set Purification (ASP) module to select the least connected pseudo-labelled samples to be inserted in the anchor set, as in (see Fig. 3):

$$a(f_\theta(\mathbf{x}), \mathcal{D}_U, \mathcal{D}_A) = \begin{cases} 1, & c(f_\theta(\mathbf{x}), \mathcal{D}_U, \mathcal{D}_A) \leq \alpha, \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

where the pseudo-labelled samples with  $a(f_\theta(\mathbf{x}), \mathcal{D}_U, \mathcal{D}_A) = 1$  and  $\tilde{y} = g(f_\theta(\mathbf{x}), \mathcal{D}_A)$  from (6) are inserted into the anchor set. The information content  $c(f_\theta(\mathbf{x}), \mathcal{D}_U, \mathcal{D}_A)$  of a pseudo-labelled sample  $f_\theta(\mathbf{x})$  in (7) is computed in three steps (see Fig. 3): 1) find the KNN samples  $\mathcal{N}(f_\theta(\mathbf{x}), \mathcal{D}_A)$  from  $f_\theta(\mathbf{x})$  to the anchor set  $\mathcal{D}_A$ ; 2) for each of the  $K$  elements  $(\mathbf{x}_A, y_A) \in \mathcal{N}(f_\theta(\mathbf{x}), \mathcal{D}_A)$ , find the KNN set  $\mathcal{N}(f_\theta(\mathbf{x}_A), \mathcal{D}_U)$  from  $f_\theta(\mathbf{x}_A)$  to the unlabelled set  $\mathcal{D}_U$ ; and 3)  $c(f_\theta(\mathbf{x}), \mathcal{D}_U, \mathcal{D}_A)$  is calculated to be the number of times that the pseudo-labelled sample  $\mathbf{x}$  appears in the KNN sets  $\mathcal{N}(f_\theta(\mathbf{x}_A), \mathcal{D}_U)$  for the  $K$  elements of set  $\mathcal{N}(f_\theta(\mathbf{x}), \mathcal{D}_A)$ . The threshold  $\alpha$  in (7) is computed with  $\alpha = \min_{\mathbf{x} \in \mathcal{D}_S} c(f_\theta(\mathbf{x}), \mathcal{D}_U, \mathcal{D}_A)$ .

## 4. Experiments

For the experiments below, we use the Chest X-Ray14 [39] and ISIC2018 [8, 36] datasets.

**Chest X-Ray14** contains 112,120 CXR images from 30,805 different patients. There are 14 labels (each label is a disease) and No Finding class, where each patient can have multiple labels, forming a multi-label classification problem. To compare with previous papers [2, 26], we adopt the official train/test data split [39]. We report the classification result on the test set (26K samples) using area under the ROC curve (AUC), and the learning process uses training sets containing different proportions of the labelled data in {2%, 5%, 10%, 15%, 20%}.

**ISIC2018** is a skin lesion dataset that contains 10,015 images with seven labels. Each image is associated with one of the labels, forming a multi-class classification problem. We follow the train/test split from [26] for fair comparison, where the training set contains 20% of labelled samples and 80% of unlabelled samples. We report the AUC, Sensitivity, and F1 score results.

### 4.1. Implementation Details

For both datasets, we use DenseNet-121 [12] as our backbone model. For Chest X-Ray14, the dataset pre-processing consists of resizing the images to  $512 \times 512$  for faster processing. For the optimisation, we use Adam optimizer [19], batch size 16 and learning rate 0.05. During training, we use data augmentation based on random crop and resize, and random horizontal flip. We first train 20 epochs on the initial labelled subset to warm-up the model for feature extraction. Then we train for 50 epochs, where in every 10 epochs we update the anchor set with ASP from

Table 1. Mean AUC testing set results over the 14 disease classes of Chest X-Ray14 for different labelled set training percentages. \* indicates the methods that use DenseNet-169 as backbone architecture. **Bold** number means the best result per label percentage and underline shows previous best results.

Method Type	Label Percentage	2%	5%	10%	15%	20%
Consistency based	SRC-MT* [26]	66.95	72.29	75.28	77.76	79.23
	NoTeacher [37]	72.60	77.04	77.61	N/A	79.49
	S <sup>2</sup> MTS <sup>2</sup> [25]	74.69	78.96	79.90	<u>80.31</u>	81.06
Pseudo Label	Graph XNet* [2]	53.00	58.00	63.00	68.00	78.00
	UPS [29]	65.51	73.18	76.84	78.90	79.92
	Ours	<b>74.82</b>	<b>79.20</b>	<b>80.40</b>	<b>81.06</b>	<b>81.77</b>

Table 2. Class-level AUC testing set results comparison between our approach and other semi-supervised SOTA approaches trained with **20%** of labelled data on Chest Xray-14. \* denotes the models use DenseNet-169 as backbone. **Bold** number means the best result per class and underline shows second best results.

Method Type	Supervised	Consistency based			Pseudo-labelling		
Method	Densenet-121	MT [33] *	SRC-MT [26] *	S <sup>2</sup> MTS <sup>2</sup> [25]	GraphXNet [2]	UPS [29]	Ours
Atelectasis	75.75	75.12	75.38	<u>78.57</u>	71.89	77.09	<b>79.53</b>
Cardiomegaly	80.71	87.37	87.7	<u>88.08</u>	87.99	85.73	<b>89.03</b>
Effusion	79.87	80.81	81.58	<u>82.87</u>	79.2	81.35	<b>83.56</b>
Infiltration	69.16	70.67	70.4	70.68	<b>72.05</b>	70.82	<u>71.40</u>
Mass	78.40	77.72	78.03	<b>82.57</b>	80.9	81.82	<u>82.49</u>
Nodule	74.49	73.27	73.64	<u>76.60</u>	71.13	76.34	<b>77.73</b>
Pneumonia	69.55	69.17	69.27	72.25	<b>76.64</b>	70.96	<u>73.86</u>
Pneumothorax	84.70	85.63	86.12	<u>86.55</u>	83.7	85.86	<b>86.95</b>
Consolidation	71.85	72.51	73.11	<u>75.47</u>	73.36	74.35	<b>75.50</b>
Edema	81.61	82.72	82.94	<u>84.83</u>	80.2	83.56	<b>84.95</b>
Emphysema	89.75	88.16	88.98	<u>91.88</u>	84.07	91.00	<b>93.36</b>
Fibrosis	79.30	78.24	79.22	<u>81.73</u>	80.34	80.87	<b>81.86</b>
Pleural Thicken	73.46	74.43	75.63	<u>76.86</u>	75.7	75.55	<b>77.60</b>
Hernia	86.05	<b>87.74</b>	<u>87.27</u>	85.98	87.22	85.62	85.89
Mean	78.19	78.83	79.23	<u>81.06</u>	78.00	79.92	<b>81.77</b>

Sec. 3.4. For the KNN classifier in (2), we set K to be 200 for 2% and 5% (of labelled data) and 50 for remaining label proportions. These values are set based on validation results, but our approach is robust to a large range K values – we show an ablation study that compares the performance of our method for different values of K. For ISIC2018, we resize the image to  $224 \times 224$  for fair comparison with baselines. For the optimisation, we use Adam optimizer [19], batch size 32 and learning rate 0.001. During training, data augmentation is also based on random crop and resize, and random horizontal flip. We warm-up the model for 40 epochs and then we train for 100 epochs, where in every 20 epochs, we update the anchor set with ASP. For the KNN classifier, K is set to 100 based on validation set. The code is written in Pytorch [28] and we use two RTX 2080ti Gpus for all experiments. KNN computation takes 5 sec for Chest X-ray14 unlabelled samples with Faiss [16] library for faster processing. We follow [25, 26, 33] to maintain an exponential moving average (EMA) version of the trained model, which is only used for evaluation not for training.

## 4.2. Thorax Disease Classification Result

For the results on Chest X-Ray14 in Table 1, our method, NoTeacher [37], UPS [29], and S<sup>2</sup>MTS<sup>2</sup> [25] use the DenseNet-121 backbone, while SRC-MT [26] and GraphXNet [2] use DenseNet-169 [12]. SRC-MT [26] is a consistency-based SSL; NoTeacher [37] extends MT by replacing the EMA process with two networks combined with a probabilistic graph model; S<sup>2</sup>MTS<sup>2</sup> [25] combines self-supervised pre-training with MT fine-tuning; and GraphXNet [2] constructs a graph from dataset samples and assigns pseudo labels to unlabelled samples through label propagation; and UPS [29] applies probability and uncertainty thresholds to enable the pseudo labelling of unlabelled samples. All methods use the official test set [39]. Our approach achieves the SOTA results for all percentages of training labels. Compared to the pseudo-labelling approaches UPS and GraphXNet, our approach outperforms them by a margin between 3% to 20%. Compared to the consistency-based approaches SRC-MT and NoTeacher, our method consistently achieves 2% improvement for all cases, even though we use a backbone archi-

Table 3. AUC, Sensitivity and F1 testing results on ISIC2018, where 20% of the training set is labelled. **Bold** shows the best result per measure, and underline shows second best results.

Method	AUC	Sensitivity	F1
Supervised	90.15	65.50	52.03
SS-DCGAN [11]	91.28	67.72	54.10
TCSE [23]	92.24	68.17	58.44
TE [21]	92.70	69.81	59.33
MT [33]	92.96	69.75	59.10
SRC-MT [26]	93.58	71.47	60.68
Self-training [3]	90.58	67.63	54.51
Ours	<b>94.36</b>	<b>72.14</b>	<b>62.23</b>

texture of lower capacity (i.e., DenseNet-121 instead of DenseNet-169). Compared with the previous SOTA, our method outperforms  $S^2MTS^2$  [25] by 1% AUC in all cases, which is remarkable because our method is initialised with an ImageNet pre-trained model instead of an expensive self-supervised pre-training approach.

The class-level performances using 20% of the labelled data of SSL methods are shown in Table 2, which demonstrates that our method achieves the best result in 10 out of the 14 classes. Our method surpasses the previous pseudo-labelling method GraphXNet by 3.7% and threshold based pseudo-labelling method [29] by 1.8%. Our method also outperforms consistency-based methods MT [33] and SRC-MT [26] by more than 2%. For method  $S^2MTS^2$  [25] with self-supervised learning, our method can outperform it using an ImageNet pre-trained model, alleviating the need of a computationally expensive self-supervised pre-training.

### 4.3. Skin Lesion Classification Result

We show the results on ISIC2018 in Table 3, where competing methods are based on self-training [3], generative adversarial network (GAN) to augment the labelled set [11], temporal ensembling [21], MT [33] and its extension [26], and a DenseNet-121 [12] baseline trained with 20% of the training set. Compared with consistency-based approaches [23, 26, 33], our method improves between 0.7% and 3% in AUC and around 1% in F1 score. Our method also outperforms previous self-training approach [3] by a large margin in all measures.

### 4.4. Ablation Study

For the ablation study, we test each of our three contributions and visualize the data distribution of selected subset with high and low informative samples on the Chest X-Ray14 [39] with 2% labelled training set, where for CDSI and ASP, we run each experiment three times and show the mean and standard deviation of the AUC results.

**Cross-distribution Sample Informativeness (CDSI).** We first study in Table 4 how performance is affected by




Figure 4. (Left) Mean AUC testing results for different values for K in the KNN (for CDSI in (4) and pseudo-labelling in (5)), where the green region uses ASP and blue region does not use ASP. (Right) Mean size of  $\mathcal{D}_L$  at every training stage when adding unlabelled samples of high, medium and low information content according to (2). Model is trained on Chest X-Ray14, where 2% of the training is labelled.

Table 4. Ablation study on Chest X-ray14 (2% labelled). Starting with a baseline classifier (DenseNet-121), we test the selection of unlabelled samples (to be provided with a pseudo-label) with different information content, according to (2) (i.e., low, medium, high), and the use of the anchor set purification (ASP) module.

Information Content	ASP	AUC $\pm$ std
Baseline		65.84 $\pm$ 0.14
Low	✗	67.18 $\pm$ 2.40
	✓	67.76 $\pm$ 1.05
Medium	✗	70.83 $\pm$ 1.49
	✓	71.16 $\pm$ 0.51
High	✗	73.81 $\pm$ 0.75
	✓	<b>74.44 <math>\pm</math> 0.38</b>

pseudo-labelling unlabelled samples with different degrees of informativeness (low, medium and high) using our CDSI. Starting from the baseline classifier DenseNet-121 that reaches an AUC of 65%, we observe that pseudo-labelling low-information content unlabelled samples yields the worst result (around 67% AUC) and selecting high-information content unlabelled samples produces the best result (around 73% AUC). Figure 4 (right) plots how the size of the labelled set  $\mathcal{D}_L$  during training depends on the degree of informativeness of the unlabelled samples to be pseudo-labelled. These results show that: 1) unlabelled samples of high-information content enables the construction of a smaller labelled set (compared with unlabelled samples of low- or medium-information content), allowing a more efficient training process that produces a more accurate KNN classifier; and 2) the standard deviation of the results in Table 4 are smaller when selecting the unlabelled samples of high-information content, compared with the low- or medium-information content. This second point can be explained by the class imbalance issue in Chest X-Ray14, where the selection of low-information content samples will enable the training of majority classes, possibly producing an ineffective training for the minority classes that can increase the variance in the results.

Table 5. AUC testing set results on Chest X-ray14 (2% labelled) for different pseudo labelling strategies ( $\alpha$  denotes the linear coefficient combining the model and KNN predictions).

Pseudo-label Strategies	Methods	AUC
Baseline	-	65.84
Single Prediction	Model prediction	72.63
	KNN prediction	72.45
Mixup	random sampled $\alpha$	73.23
	MixUp [42]	69.28
Ours	Informative Mixup	<b>74.44</b>

**Anchor Set Purification (ASP).** Also in Table 4, we compare ASP with an alternative method that selects all pseudo-labelled samples to be included into the anchor set for the low-, medium- and high-information content unlabelled samples. Results show that the ASP module improves AUC between 0.3% and 1.0% and reduces standard deviation between 0.4% and 1.4%. This demonstrates empirically that the ASP module enables the formation of a more informative anchor set that improves the pseudo-labelling accuracy, and consequently the final AUC results. Furthermore, in Figure 4 (left), ASP is shown to stabilise the performance of the method with respect to  $K \in \{50, 100, 150, 200, 250, 300\}$  for the KNN classifier of (4). In particular, with ASP, the difference between the best and worst AUC results is around 1%, while without ASP, the difference grows to 2%. This can be explained by the fact that without ASP, the anchor set grows quickly with relatively less informative pseudo-labelled samples, which reduces the stability of the method.

**Informative Mixup (IM)** In Table 5, we show that our proposed IM in (6) produces a more accurate pseudo-label, where we compare it with alternative pseudo-label methods, such as with only the model prediction, only the KNN prediction, random sample  $\alpha$  from beta distribution to replace the density score in (6), and regular MixUp [42]. It is clear that the use of model or KNN predictions as pseudo labels does not work well most likely because of confirmation bias (former case) or the inaccuracy of the KNN classifier (latter case). MixUp [42] does not show good accuracy either, as also observed in [38] and [18], when MixUp is performed in multi-label images or multiple single-object images. The random sampling of  $\alpha$  for replacing density score shows a better result than MixUp, but the lack of an image-based weight to balance the two predictions, like in (6), damages performance. Our proposed IM shows a result that is at least 1.5% better than any of the other pseudo-labelling approaches, showing the importance of using the density of the unlabelled sample in the anchor set to weight the contribution of the model and KNN classifiers.

The **imbalanced learning mitigation** is studied in Figure 5, which shows the histogram of label distribution in percentage (for a subset of four disease minority classes and the No




Figure 5. The selection of highly informative unlabelled samples (blue) promote a more balanced learning process, where the difference in the number of samples belonging to the minority or majority classes is smaller than if we selected unlabelled samples with low informativeness (yellow). Green shows the original data distribution]. Full 14-class distributions are shown in the supplementary material.

Finding majority class) by selecting unlabelled samples of high (blue) and low (yellow) information content. We also show the original label distribution in green for reference.

Notice that the selection of highly informative samples significantly increases the percentage of disease minority classes (from between 5% and 10% to almost 30%) and decreases the percentage of the No Finding majority class (from 60% to 30%), creating a more balanced distribution of these five classes. This indicates that our informative sample selection can help to mitigate the issue of imbalanced learning. We include the full 14-classes histograms in the supplementary material.

## 5. Discussion and Conclusion

In this work, we introduced the anti-curriculum pseudo-labelling (ACPL) SSL method. Unlike traditional pseudo-labelling methods that use a threshold to select confidently classified samples, ACPL uses a new mechanism to select highly informative unlabelled samples for pseudo-labelling and an ensemble of classifiers to produce accurate pseudo-labels. This enables ACPL to address MIA multi-class and multi-label imbalanced classification problems. We show in the experiments that ACPL outperforms previous consistency-based, pseudo-label based and self-supervised SSL methods in multi-label Chest X-ray14 and multi-class ISIC2018 benchmarks. We demonstrate in the ablation study the influence of each of our contributions and we also show how our new selection of informative samples addresses MIA imbalanced classification problems. For future work, it is conceivable that ACPL can be applied can be applied to more general computer vision problems, so we plan to test ACPL in traditional computer vision benchmarks. We would also explore semi-supervised classification with out-of-distribution (OOD) data in the initial labelled and unlabelled sets as our method currently assume all samples are in-distribution.



## References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020. [2](#)
- [2] Angelica I Aviles-Rivero, Nicolas Papadakis, Ruoteng Li, Philip Sellars, Qingnan Fan, Robby T Tan, and Carola-Bibiane Schönlieb. Graphxnet chest x-ray classification under extreme minimal supervision. *arXiv preprint arXiv:1907.10085*, 2019. [2](#), [3](#), [5](#), [6](#)
- [3] Wenjia Bai, Ozan Oktay, Matthew Sinclair, Hideaki Suzuki, Martin Rajchl, Giacomo Tarroni, Ben Glocker, Andrew King, Paul M Matthews, and Daniel Rueckert. Semi-supervised learning for network-based cardiac mr image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 253–260. Springer, 2017. [3](#), [7](#)
- [4] Yoshua Bengio, Jerome Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009. [3](#)
- [5] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*, 2019. [1](#), [2](#)
- [6] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems*, 32:5049–5059, 2019. [1](#), [2](#)
- [7] Paola Cascante-Bonilla et al. Curriculum labeling: Revisiting pseudo-labeling for semi-supervised learning. *arXiv preprint arXiv:2001.06001*, 2020. [2](#), [3](#)
- [8] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kalloo, Konstantinos Liopyris, Michael Marchetti, et al. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368*, 2019. [2](#), [3](#), [5](#)
- [9] Wenhui Cui, Yanlin Liu, Yuxing Li, Menghao Guo, Yiming Li, Xiuli Li, Tianle Wang, Xiangzhu Zeng, and Chuyang Ye. Semi-supervised brain lesion segmentation with an adapted mean teacher model. In *International Conference on Information Processing in Medical Imaging*, pages 554–565. Springer, 2019. [3](#)
- [10] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977. [4](#)
- [11] Andres Diaz-Pinto, Adrián Colomer, Valery Naranjo, Sandra Morales, Yanwu Xu, and Alejandro F Frangi. Retinal image synthesis and semi-supervised learning for glaucoma assessment. *IEEE transactions on medical imaging*, 38(9):2211–2218, 2019. [7](#)
- [12] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. [5](#), [6](#), [7](#)
- [13] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 590–597, 2019. [3](#)
- [14] Angela H Jiang, Daniel L-K Wong, Giulio Zhou, David G Andersen, Jeffrey Dean, Gregory R Ganger, Gauri Joshi, Michael Kaminsky, Michael Kozuch, Zachary C Lipton, et al. Accelerating deep learning by focusing on the biggest losers. *arXiv preprint arXiv:1910.00762*, 2019. [3](#)
- [15] Lu Jiang, Deyu Meng, Qian Zhao, Shiguang Shan, and Alexander G Hauptmann. Self-paced curriculum learning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015. [3](#)
- [16] Jeff Johnson et al. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017. [6](#)
- [17] Kenji Kawaguchi and Haihao Lu. Ordered sgd: A new stochastic optimization framework for empirical risk minimization. In *International Conference on Artificial Intelligence and Statistics*, pages 669–679. PMLR, 2020. [3](#)
- [18] Jang-Hyun Kim, Wonho Choo, Hosan Jeong, and Hyun Oh Song. Co-mixup: Saliency guided joint mixup with super-modular diversity. *arXiv preprint arXiv:2102.03065*, 2021. [8](#)
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#), [6](#)
- [20] M Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. *Advances in neural information processing systems*, 23:1189–1197, 2010. [3](#)
- [21] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. [2](#), [7](#)
- [22] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. 2013. [2](#)
- [23] Xiaomeng Li, Lequan Yu, Hao Chen, Chi-Wing Fu, and Pheng-Ann Heng. Semi-supervised skin lesion segmentation via transformation consistent self-ensembling model. *arXiv preprint arXiv:1808.03887*, 2018. [3](#), [7](#)
- [24] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen AWM Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. A survey on deep learning in medical image analysis. *Medical image analysis*, 42:60–88, 2017. [1](#)
- [25] Fengbei Liu, Yu Tian, Filipe R Cordeiro, Vasileios Belagianis, Ian Reid, and Gustavo Carneiro. Self-supervised mean teacher for semi-supervised chest x-ray classification. *arXiv preprint arXiv:2103.03629*, 2021. [1](#), [2](#), [3](#), [6](#), [7](#)

- [26] Quande Liu et al. Semi-supervised medical image classification with relation-driven self-ensembling model. *IEEE transactions on medical imaging*, 39(11):3429–3440, 2020. 3, 5, 6, 7
- [27] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018. 2
- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 6
- [29] Mamshad Nayeem Rizve et al. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *International Conference on Learning Representations*, 2020. 2, 3, 6, 7
- [30] Burr Settles. Active learning literature survey. 2009. 2
- [31] Weiwei Shi, Yihong Gong, Chris Ding, Zhiheng MaXiaoyu Tao, and Nanning Zheng. Transductive semi-supervised deep learning using min-max features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 299–315, 2018. 2
- [32] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in Neural Information Processing Systems*, 33, 2020. 1, 2
- [33] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in Neural Information Processing Systems*, 30, 2017. 2, 6, 7
- [34] Yu Tian, Gabriel Maicas, Leonardo Zorron Cheng Tao Pu, Rajvinder Singh, Johan W Verjans, and Gustavo Carneiro. Few-shot anomaly detection for polyp frames from colonoscopy. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 274–284. Springer, 2020. 1
- [35] Yu Tian, Guansong Pang, Fengbei Liu, Seon Ho Shin, Johan W Verjans, Rajvinder Singh, Gustavo Carneiro, et al. Constrained contrastive distribution learning for unsupervised anomaly detection and localisation in medical images. *MICCAI2021*, 2021. 1
- [36] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1):1–9, 2018. 1, 2, 3, 5
- [37] Balagopal Unnikrishnan, Cuong Manh Nguyen, Shafa Balaram, Chuan Sheng Foo, and Pavitra Krishnaswamy. Semi-supervised classification of diagnostic radiographs with noteacher: A teacher that is not mean. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 624–634. Springer, 2020. 3, 6
- [38] Qian Wang, Ning Jia, and Toby P Breckon. A baseline for multi-label image classification using an ensemble of deep convolutional neural networks. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 644–648. IEEE, 2019. 8
- [39] 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. 1, 2, 3, 5, 6, 7
- [40] Xiaoxia Wu, Ethan Dyer, and Behnam Neyshabur. When do curricula work? In *International Conference on Learning Representations*, 2021. 3
- [41] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10687–10698, 2020. 2
- [42] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 4, 5, 8