Boosting Few-Shot Learning with Disentangled Self-Supervised Learning and Meta-Learning for Medical Image Classification

Eva Pachetti[1,2] Sotirios A. Tsaftaris [3] Sara Colantonio [1]

[1] Institute of Information Science and Technologies, National Research Council, Pisa, Italy
{eva.pachetti, sara.colantonio}@isti.cnr.it
[2] Department of Information Engineering, University of Pisa, Pisa, Italy
[3] School of Engineering, University of Edinburgh, Edinburgh, UK
S.Tsaftaris@ed.ac.uk

ABSTRACT

Background and objective: Employing deep learning models in critical domains such as medical imaging poses challenges associated with the limited availability of training data. We present a strategy for improving the performance and generalization capabilities of models trained in low-data regimes.

Methods: The proposed method starts with a pre-training phase, where features learned in a self-supervised learning setting are disentangled to improve the robustness of the representations for downstream tasks. We then introduce a meta-fine-tuning step, leveraging related classes between meta-training and meta-testing phases but varying the granularity level. This approach aims to enhance the model's generalization capabilities by exposing it to more challenging classification tasks during meta-training and evaluating it on easier tasks but holding greater clinical relevance during meta-testing. We demonstrate the effectiveness of the proposed approach through a series of experiments exploring several backbones, as well as diverse pre-training and fine-tuning schemes, on two distinct medical tasks, i.e., classification of prostate cancer aggressiveness from MRI data and classification of breast cancer malignity from microscopic images.

Results: Our results indicate that the proposed approach consistently yields superior performance w.r.t. ablation experiments, maintaining competitiveness even when a distribution shift between training and evaluation data occurs.

Conclusion: Extensive experiments demonstrate the effectiveness and wide applicability of the proposed approach. We hope that this work will add another solution to the arsenal of addressing learning issues in data-scarce imaging domains.

Keywords Few-shot learning · Meta-learning · Disentangled representation learning · Self-supervised learning · Classification · Medical image analysis

1 Introduction

Developing deep learning models that are able to perform competitively on small datasets is essential, especially in domains where data are scarce. Few-shot learning (FSL) is a paradigm designed to train models by using a limited number of examples [1]. One of the most relevant applications of FSL is in the medical image domain [2, 3, 4, 5], where datasets are often small because of the difficulty and cost of acquiring such images and the need to protect patient privacy. The use of FSL has become increasingly popular since the advent of meta-learning, defined as the ability to 'learn to learn.' Indeed, in a meta-learning framework, models are trained on tasks rather than data, enabling the model to generalize better [6]. Several studies [7, 8] have shown that employing self-supervised learning (SSL) on

additional unlabeled data as a pre-training step can boost the effectiveness of a meta-trained model. However, the main challenge with SSL methods, particularly contrastive-learning methods, is that they tend to disentangle features related to image augmentations used to generate additional instances for the contrastive setting (such as rotation and color jitter) rather than class-specific ones [9]. This way, models rely on the wrong features to perform the downstream task, resulting in poorer generalization. To learn more disentangled features beyond augmentations during SSL, Wang *et al.* [9] proposed a new method called Iterative Partition-based Invariant Risk Minimization (IP-IRM). This iterative algorithm systematically disentangles one feature at a time by partitioning the training samples into two subsets based on an entangled feature (e.g., color) and minimizing a contrastive loss invariant w.r.t. the two subsets. The invariance between subsets ensures the disentanglement of the selected feature [9]. Our work builds on a simple idea: leveraging the power of disentangled SSL as a pre-training step to extract robust feature representations and coupling it with the generalization capabilities of a meta-learning framework to enhance the classification performance of models trained in the FSL regime. Specifically, we use IP-IRM to pre-train a convolutional backbone to generate robust features and then fine-tune it with a meta-learning approach based on Prototypical Networks called Meta Deep Brownian Distance Covariance (Meta DeepBDC) [10]. We also propose a novel meta-learning approach that uses different but related classes during the meta-training and meta-testing stages. Specifically, we use a dataset from the same domain with finer-grained classes for meta-training and coarser-grained classes for meta-testing, i.e., meta-training classes are subgroups of meta-testing classes. The aim is to train the model on more complex and diverse classification tasks, where complexity comes from higher granularity and diversity comes from a larger number of classes from which to choose in each episode. On the other hand, meta-testing episodes are easier to address since the classes are broader and the episodes less varied. This approach may be beneficial when coarse classes are still clinically relevant and there is insufficient data to expect good performance on more detailed classes.

To prove the versatility of our method, we evaluate our approach on two different medical imaging datasets and clinical tasks: the PI-CAI challenge dataset [11] of multiparametric MRI (mpMRI) prostate images and the Breast cancer Histopathological Image (BreakHis) challenge dataset [12] of microscopic breast images. Using the first dataset, our objective is to classify the prostate mpMRI images based on the tumour severity by predicting a prognostic value defined by the International Society of Urological Pathology (ISUP) [13]. Following the mpMRI investigation, a biopsy examination (an invasive and uncomfortable procedure) is usually required to determine the tumour severity in suspected lesions [14]. In this sense, a prediction model that performs accurate diagnosis directly from MRI images would help radiologists limit overdiagnosis, thus reducing patient discomfort.

Concerning the BreakHis dataset, we consider as a clinical task the classification of benign and malignant lesions, which affects the patient's treatment plan. Breast cancer diagnosis from microscopic images is a challenging and time-consuming work as it requires manually detecting cancer nuclei [15]. This leads to high workload for pathologists, long diagnosis time and potential misdiagnosis due to human factors such as eye fatigue, in addition to device-dependant influences [16]. For all these reason, developing an automated classification of breast cancer biopsy samples could support pathologists by improving diagnosis accuracy and breast cancer early detection.

Below, we outline the primary contributions of our work:

- We propose to exploit the strengths of disentangled SSL as a pre-training step for a meta-fine-tuned model to enhance the generalizability performance in the FSL domain.
- We propose a novel meta-learning approach that trains the model on more complex tasks than the ones used during meta-testing by employing related classes with different levels of granularity. Specifically, we use finer classes during meta-training and coarser ones during meta-testing, which are still clinically relevant.
- We examine the robustness of our approach across distribution shifts by using data from different vendors between the training and evaluation phases.

We structured this paper as follows: Section 2 provides a comprehensive review of related works; Section 3 yields a theoretical background of the methods employed in this study; in Section 4, we delineate our proposed approach and detail the experimental settings, whereas the results of our experiments are presented in Section 5. In Section 6, we discuss and explain our findings. Eventually, in Section 7, we summarize our contribution and the applicability range of our study, finally outlining potential future directions.

2 Related works

Our work synergistically integrates meta-learning, SSL, and feature disentanglement methods to enhance the performance of models trained in an FSL setting. This section provides an overview of the state-of-the-art research in the interplay of these techniques, drawing parallels and highlighting differences with our proposed approach. First, we discuss existing works that leverage meta-learning methods to perform medical image classification in an FSL regime,

aligning with our primary task. Next, similar to our approach, we examine studies that employ an SSL pre-training step to improve the performance of downstream tasks performed in an FSL setting. In the subsequent paragraph, we outline papers that propose different methodologies for improving FSL performance by leveraging feature disentanglement. Finally, we delve into the works that promote feature disentanglement in an SSL setting, culminating in a study that evaluates the efficacy of this approach in an FSL regime.

2.1 FSL in Medical Image Classification

Several studies have tackled the challenge of medical image classification within an FSL framework. Some of these works exclusively rely on meta-learning algorithms. For instance, Singh *et al.* [17] applied the Reptile algorithm by meta-training the model on more prevalent diseases and evaluating it on rare diseases. Other works augmented the meta-learning paradigm by incorporating additional modules and pre-training steps. Dai *et al.* [3], for example, enhanced a gradient-based meta-learning algorithm by integrating it with a prior-guided Variational Autoencoder (VAE) to improve target features. From another perspective, Jiang *et al.* [4] merged meta-learning with transfer learning, utilizing a multi-learner model (autoencoder, metric-learner, and task learner) trained based on either a transfer learning or a meta-learning criterion at different stages. In this work, we propose to boost the meta-learning framework capabilities according to two main approaches: embedding a pre-training step that leverages disentangled SSL and elevating the generalization capabilities of the episodic meta-learning by varying the class granularity between meta-training and meta-testing episodes.

2.2 SSL as a pre-training step in FSL

In line with our research, the concept of improving the performance of FSL-trained models through an SSL pre-training step has been explored in various studies. For instance, Chen *et al.* [18] promoted a more robust data representation in downstream tasks by incorporating a contrastive SSL pre-training step followed by episodic training on natural images. In the facial expression recognition domain, Chen *et al.* [19] integrated SSL and FSL by performing an SSL pre-training, followed by a classical fully-supervised fine-tuning phase of the feature extractor. The authors then used the fine-tuned model to build the prototype and perform a few-shot classification. In another approach Medina *et al.* [20] adopted a transfer learning strategy that constructs a metric embedding, closely clustering unlabeled prototypical samples and their augmentations through a self-supervised approach. The pre-trained embedding serves as a starting point for few-shot classification, achieved by prototypical fine-tuning of the final classification layer. Yang *et al.* [21] suggest incorporating contrastive learning into both the pre-training and meta-fine-tuning stages to enhance the performance of few-shot classification. In the pre-training stage, they introduced a self-supervised contrastive loss that leverages global and local information to learn effective initial representations. Concerning the meta-training stage, instead, they proposed a cross-view episodic training mechanism that involves performing nearest centroid classification on two different views of the same episode and employing a distance-scaled contrastive loss based on these views. Our work employs an enhanced SSL version for pre-training, which ensures the disentanglement of features, aiming at improving the quality of the features extracted during the subsequent downstream task.

2.3 Feature disentanglement for FSL

The concept of promoting feature disentanglement to improve performance in the FSL domain has been explored in the literature as well, particularly within the realm of natural images. Hu *et al.* [22] employed feature disentanglement to generate augmented features through hallucination, aiming to mitigate data sparsity in few-shot classification. Similarly, the utilization of feature disentanglement for hallucination was adopted by Lin *et al.* [23], where they extracted class-specific and appearance-specific features for base categories and then employed them to hallucinate image features for novel categories. In another study, Cheng *et al.* [24] presented a disentangled feature representation framework tailored for few-shot applications. This framework adaptively decouples discriminative features, modeled by the classification branch, from the class-irrelevant component of the variation branch, thereby enhancing performance on few-shot tasks.

Unlike the described approaches that employ feature disentanglement in a supervised setting to directly improve FSL classification, our work leverages the disentanglement of features at the unsupervised pre-training level. This unsupervised approach enhances the extraction of robust features, which subsequently benefits the FSL downstream task.

2.4 SSL and Feature disentanglement

Disentangled representation learning is a well-studied problem in supervised learning [25, 26, 27, 28, 29, 30]. On the other hand, forcing the feature disentanglement in unsupervised learning is considered more challenging [31]. Most of

the existing methods for unsupervised disentanglement use generative models such as Generative Adversarial Networks (GAN) [32, 33, 34] and VAE [35, 36, 37]. Wang *et al.* [9] proposed the first method for unsupervised disentanglement called IP-IRM, which leverages the IRM algorithm [38] to implement a disentangled SSL. Here, the authors assess the efficacy of their method in an FSL setting but differently from the present work, without fine-tuning the feature extractor on the downstream task. Indeed, they directly classify the feature embeddings generated by the feature extractor using both a k-nearest neighbors algorithm and a standalone linear classifier.

On the contrary, in this work, we consider the use of the IP-IRM algorithm as a pre-training step, intending to improve the learning of robust data representation for the downstream task, which is performed by fine-tuning the feature extractor through a meta-learning approach.

3 Background and Motivation

Based on the state-of-the-art methods, we defined a new approach aimed at further improving the capabilities of models trained in an FSL manner. In this section, we offer a theoretical insight into the algorithms utilized for this purpose, motivating their relevance and adoption in our approach.

3.1 SimCLR

Zhang *et al.* [39] emphasized that for downstream classification tasks, predictive and contrastive SSL methods yield superior results compared to generative SSL approaches due to their ability to focus on high-level anatomical structures rather than pixel-wise information. Specifically, they demonstrated the robust performance of SimCLR, BYOL [40], and RPL [41] across various tasks, including segmentation and classification, highlighting their effectiveness in learning robust representations. Following these guidelines, we considered a contrastive method, namely SimCLR [42], as our SSL algorithm.

As demonstrated by [9], a significant portion of the disentangled features learned by SimCLR are related to augmentations, whereas other meaningful features remain entangled. This results in a decreased generalization performance when evaluating the model on downstream tasks. To overcome this limitation, we employed the IP-IRM algorithm [9] to enforce the disentanglement of features beyond augmentation-related ones during SSL. In the subsequent paragraph, we provide a more detailed description of the IRM algorithm [38] and its derivative IP-IRM algorithm [9].

3.2 Iterative Partition-based Invariant Risk Minimization

It is often the case that during training the model environment-related features instead of class-specific ones. To identify and eliminate such features, we can leverage the assumption that if data are collected in different environments, spurious correlations are unstable, as they vary with changes in the environment itself [43]. Based on this principle, the IRM algorithm disentangles features by seeking a data representation where the optimal classifier remains consistent across all environments. In simpler terms, the same classifier performs effectively regardless of the environment of the training data, enabling robust generalization over out-of-distribution data. In summary, we aim for a data representation that is effective in producing accurate predictions and capable of generating an invariant predictor across all environments. Mathematically speaking, given a set of different environments e, a risk under the environment R_e, and a fixed and dummy classifier w represented as a scalar, we want to find a representation of the data ϕ such that:

$$\min_{\phi} \sum_{e} R_e(\phi) + \lambda \cdot \|\nabla_{w|w=1} R_e(w \cdot \phi)\|^2. \tag{1}$$

The first term corresponds to Empirical Risk Minimization (ERM) [44], wherein the goal is to minimize the average risk across environments (predictive power) using a predictor ϕ. The second term encourages the predictor ϕ to exhibit invariance across different environments. The parameter λ represents a regularization term that balances these two components.

The authors in [1] leveraged this idea to force the disentanglement of features when the training is performed in an unsupervised, specifically in a self-supervised way, by iterating the IRM algorithm on different subsets of an unlabeled dataset. Specifically, the idea is to subdivide the unlabeled dataset into two partitions based on a partition matrix **P**. For the i-th image, $\mathcal{P}_{i,k} = 1$ if it belongs to the k-th subset and 0 otherwise. Each partition is then utilized to train the feature extractor ϕ in a self-supervised manner. The pretext task loss ($\mathcal{L}$) is defined by the contrastive loss, which takes the following form:

$$\mathcal{L} = \sum_{x \in \mathcal{X}_k} -\log\left(\frac{\exp(x^T x^* \cdot \theta)}{\sum_{x' \in \mathcal{X}_k \cup \mathcal{X}^* \backslash x} \exp(x^T x^* \cdot \theta)}\right), \tag{2}$$

where $\theta = 1$ is a parameter to evaluate the invariance of the SSL loss across the subsets, $\mathcal{X}_k$ represents the features of the k-th subset, and $\mathcal{X}^*$ the features of their augmented version.

In the first step, the parameters of ϕ are updated by:

$$\min_{\Phi} \sum_{\mathbf{P} \in \mathcal{P}} \sum_{k=1}^{2} [\mathcal{L} + \lambda_1 \|\nabla_{\theta=1} \mathcal{L}\|^2], \tag{3}$$

where λ_1 is a hyperparameter that regulates the IRM loss term, with its minimization promoting invariance across subsets.

In the second step, the parameters of ϕ are held fixed, and the objective is to identify a new partition to be disentangled, denoted by $\mathbf{P}^*$. The desired partition is the one that maximizes the invariance between the losses, i.e.:

$$\mathbf{P}^* = \arg \max_{\mathbf{P}} \sum_{k=1}^{2} \left[\mathcal{L} + \lambda_2 \|\nabla_{\theta=1} \mathcal{L}\|^2 \right]. \tag{4}$$

The $\mathcal{P}$ matrix is updated as $\mathcal{P} \leftarrow \mathcal{P} \cup \mathbf{P}^*$. These two steps are iterated until convergence is achieved. For further details on the IP-IRM algorithm, refer to [9].

3.3 Meta Deep Brownian Distance Covariance

Meta DeepBDC is a meta-learning method based on Prototypical Networks [45]. In traditional Prototypical Networks, data are represented by their first moment (mean), and similarity between class prototypes and query embeddings is assessed using metrics like Euclidean or cosine distance. However, studies have shown that incorporating richer statistics, including second moments, while adopting the Frobenius norm or Kullback-Leiberler divergence as similarity measures results in improved performance [46, 47]. It should be noted that these studies often neglect joint distributions, limiting the learning of relationships between the two embeddings. To overcome this limitation, Xie *et al.* [10] proposed using the BDC metric, defined as the Euclidean distance between the joint characteristic function and the product of the marginals of two random variables, X and Y. BDC is mathematically formulated as:

$$\rho(X, Y) = \int_{\mathbb{R}^p} \int_{\mathbb{R}^q} \frac{|\Phi_{XY}(t,s) - \Phi_X(t)\Phi_Y(s)|^2}{c_p c_q \|t\|^{1+p} \|s\|^{1+q}} dt ds, \tag{5}$$

where $\Phi_X(t)$ and $\Phi_Y(s)$ are the marginal distributions of X and Y, respectively, $\Phi_{XY}(t, s)$ is the joint characteristic function of the two random variables and c_p is defined as $c_p = \pi^{(1+p)/2} / \Gamma((1+p)/2)$, where Γ is the complete gamma function.

In the image domain, the similarity between two images can be computed as the inner product between their respective BDC matrices. In the context of Meta DeepBDC, the algorithm calculates a BDC matrix for each support and query embedding. Subsequently, the prototype for each class is determined as the average of the BDC matrices belonging to that class. Classification is then executed by computing a distance distribution between the query BDC matrix and each class prototype. See [10] for more details.

4 Material and Methods

This section delves into the proposed approach, outlining the core concept and the experiments undertaken to validate its effectiveness.

4.1 The main idea

Our work aims to improve the capabilities of FSL-trained models by synergistically combining the strengths of SSL, feature disentanglement, and meta-learning methods. Specifically, we propose to perform a pre-training step exploiting an SSL setting where feature disentanglement is promoted, addressing the limitations of classical contrastive SSL. This step aims to equip the model with robust features that enhance its performance during the downstream task. To effectively disentangle features, we employ the IP-IRM algorithm. This algorithm iteratively partitions an unlabeled dataset into two partitions based on a feature to be disentangled. The disentanglement is performed by minimizing the invariance of the SSL loss across the two partitions.

Subsequently, we employ meta-learning to fine-tune the pre-trained model on the downstream task using a labeled dataset. Meta-learning further enhances the model's generalization abilities, enabling effective performance in low-data settings. Specifically, we divide the labeled dataset into three subsets to perform meta-training, meta-validation, and meta-test. To further improve the meta-learning framework's effectiveness, we propose utilizing related classes at different granularity levels between the meta-training and meta-testing episodes, i.e., the meta-training subset is composed of images labeled ad a finer-grained level, whereas meta-validation and meta-test images are labeled with coarser-grained ones. We illustrate our proposed approach in Fig. 1.

4.2 Problem definition

Given a set of datasets $D = \{D^1, D^2, \ldots, D^N\}$, each one comprises both labeled and unlabeled images, denoted as $D^i = \{(x, y)_1, \ldots, (x, y)_m, (x)_{m+1}, \ldots, (x)_n\}$, where $(x, y)_k$ represents an image-label pair, and $(x)_h$ is the image alone. We split each dataset as $D^i_{lab} = \{(x, y)_1, (x, y)_2, \ldots, (x, y)_m\}$ and $D^i_{unlab} = \{(x)_{m+1}, (x)_{m+2}, \ldots, (x)_n\}$. We further separate the labeled dataset into two distinct sets of classes, namely C^i_{fine} and C^i_{coarse}, where $|C^i_{coarse}| < |C^i_{fine}|$. Specifically, C^i_{fine} contains classes that are subgroups of classes in C^i_{coarse}. In line with this, we further divide D^i_{lab} into three subsets for training, validation, and testing, labeling the images in each subset as follows: $D^i_{meta-train} = \{(x, y)_k, y_k \in C^i_{fine}\}$, $D^i_{meta-val} = \{(x, y)_k, y_k \in C^i_{coarse}\}$, and $D^i_{meta-test} = \{(x, y)_k, y_k \in C^i_{coarse}\}$.

Our objective is to utilize D^i_{unlab} to pre-train a feature extractor ϕ following an SSL approach, where feature disentanglement is enforced. Subsequently, we conduct fine-tuning on ϕ using a meta-learning approach, leveraging $D^i_{meta-train}$, $D^i_{meta-val}$, and $D^i_{meta-test}$.

4.3 Pre-training phase

4.3.1 Proposed pre-training step

The initial step in our approach involves a disentangled SSL pre-training phase, which leverages the SimCLR and IP-IRM algorithms. As with the IP-IRM algorithm, our pre-training phase is conducted iteratively. Specifically, at each iteration, the D^i_{unlab} is subdivided into two subsets w.r.t. a feature to be yet disentangled, and the SSL loss invariance is minimized between the two subsets.

4.3.2 Pre-training ablations

To evaluate the efficacy of the initial stage of our proposed approach, we compared it with two alternative pre-training methodologies. The first approach involves the use of the conventional SimCLR algorithm, corresponding to the initial iteration of the IP-IRM algorithm, where all samples belong to the first subset (i.e., the first column of **P** is set to 1). In this scenario, pre-training was conducted on the same dataset used by the IP-IRM algorithm, specifically D^i_{unlab}. The second approach entails fully-supervised pre-training on the Imagenet dataset, utilizing pre-trained parameters provided by the PyTorch library [48].

4.4 Fine-tuning phase

4.4.1 Proposed Meta-learning framework

After pre-training the feature extractor ϕ, we meta-fine-tuned it on the downstream task using the Meta DeepBDC algorithm [10], which employs an episodic training approach. During the episodic training, the model is presented with multiple classification tasks, each comprising a support set and a query set. The support set represents the training data for that specific task, while the query set is used to evaluate the model's performance on that task. Following the *N-way K-shot* paradigm, each support set consists of *N* classes, each represented by *K* image examples. In Meta DeepBDC, a BDC matrix is calculated for each support and query sample during each meta-training episode. The average of the support BDC matrices belonging to the same class forms the class prototype, and classification within each episode is performed by computing a similarity distribution of the query BDC matrix against all the class prototypes. Once the meta-fine-tuning is finalized, we evaluated its performance on various meta-testing tasks employing the same approach as in the meta-training phase.

In traditional meta-learning experiments, meta-training and meta-testing tasks typically involve unrelated sets of classes [45, 49, 50], e.g., distinguishing dog breeds in meta-training and cat breeds in meta-testing. However, some studies adopted a different approach by using distinct classes for meta-training and meta-testing that still belong to the same underlying data distribution. For instance, the authors in [4, 17] used the same image modalities and anatomical
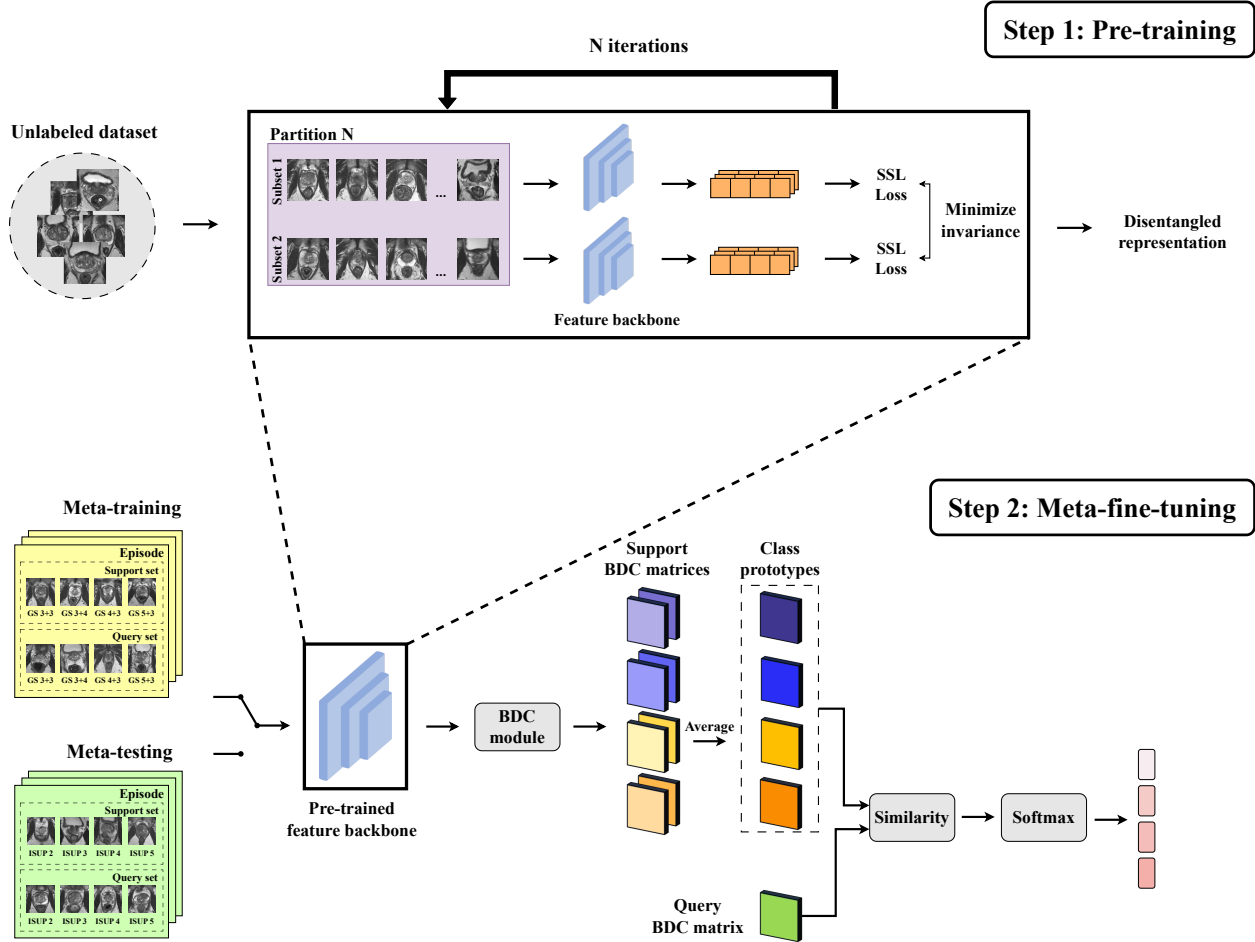
Figure 1: Illustration of the proposed approach. In the pre-training step, the feature backbone undergoes pre-training using the IP-IRM algorithm. At each iteration, the unlabeled dataset is divided into two subsets to maximize invariance between SSL losses. Subsequently, these subsets are employed to update the feature backbone parameters by minimizing the invariance between the SSL losses. The pre-trained backbone then undergoes meta-fine-tuning using the Meta DeepBDC algorithm. Meta-training episodes contain finer-grained classes while meta-testing coarser-grained ones belonging to the same source dataset. In both the meta-training and meta-testing phases a BDC matrix is computed for each support and query sample. Class prototypes are derived by averaging the BDC matrices of all support samples for that class. Classification is achieved by computing a similarity distribution of the query BDC matrix w.r.t. the class prototypes.

structure but explored different types of diseases between meta-training and meta-testing. We propose a novel approach wherein the model is tasked with similar classification tasks during meta-training and meta-testing but at different levels of granularity. Specifically, the meta-training classes (C_{fine}) constitute subgroups of the meta-testing ones (C_{coarse}), i.e., $|C_{fine}| > |C_{coarse}|$. This design choice serves two purposes. First, during meta-training, the model is challenged to distinguish images at a finer granularity level. Second, since the number of *ways* (distinct classes) remains constant between meta-training and meta-testing, and the total number of meta-training classes is higher, the episodes in meta-training become more diversely composed. These two facets amplify the complexity of the training phase, consequently strengthening the model's generalization capabilities. Meanwhile, since the model is tasked with similar but simpler tasks during meta-testing, we expect that its classification capabilities will improve accordingly. Our reasoning behind this approach holds relevance in contexts where the meta-testing tasks carry higher clinical significance despite being easier.

4.4.2 Fine-tuning ablations

To evaluate the effectiveness of the described approach, we conducted three ablation fine-tuning experiments. To assess whether using finer-grained classes during meta-training compared to meta-testing contributes to higher model generalizability, we performed two additional meta-learning experiments by generating the meta-training set from a different source dataset than the one used for meta-testing. Formally, we considered an additional dataset D^j where $j \neq i$ and new sets of fine-grained and coarse-grained classes, namely C^j_{fine} and C^j_{coarse} respectively. In the first ablation experiment, the meta-training classes belonged to C^j_{fine}, i.e., $D^j_{meta-train} = \{(x, y)_k, y_k \in C^j_{fine}\}$. In the second one, we considered the coarser-grained set of classes for meta-learning as well, $D^j_{meta-train} = \{(x, y)_k, y_k \in C^j_{coarse}\}$. Meta-validation and meta-testing sets remained unchanged.

Furthermore, as a third ablation experiment, we fine-tuned the pre-trained backbones using the classical fully-supervised approach. For this experiment, we utilized the same dataset splitting as in the proposed meta-learning approach but with a consistent class set across all subsets, denoted as C^i_{coarse}. Specifically, the datasets considered were as follows: $D^i_{train} = \{(x, y)_k, y_k \in C^i_{coarse}\}$, $D^i_{val} = \{(x, y)_k, y_k \in C^i_{coarse}\}$, and $D^i_{test} = \{(x, y)_k, y_k \in C^i_{coarse}\}$.

4.5 Datasets

We employed two datasets: the PI-CAI challenge dataset [dataset] [11], consisting of 1500 mpMRI acquisitions of prostate cancer and the BreakHis dataset [dataset] [12], specifically its processed version curated by Pereira [dataset] [51]. Both datasets were confirmed to have been collected with institutional/ethical review board approval.

4.5.1 PI-CAI dataset

For our experiments, we focused exclusively on T2-weighted images of patients with both cancerous and benign lesions. All acquisitions of benign lesions, specifically those without biopsy or with a negative biopsy (for which we lacked ground truth), were used as D_{unlab}. This dataset comprised 11202 images from 849 patients. Conversely, we utilized the acquisitions of cancerous lesions as D_{lab} for the fine-tuning phase. The labeled dataset contains 2049 images from 382 patients, which we divided into 1611 for training, 200 for validation, and 238 for testing. During the splitting process, we ensured patient stratification, i.e., all images from the same patient were grouped in the same subset, avoiding any data leakage.

Each mpMRI acquisition was provided together with a biopsy report indicating the severity of each lesion, assessed by the Gleason Score (GS), which can assume values ranging from $1 + 1$ to $5 + 5$ based on the severity of the two most common patterns in the biopsy specimen. Related to the GS, each lesion was assigned a prognostic score according to the ISUP guidelines [13], which ranges from 1 to 5, according to the tumor severity. Fig. 2a visually represents the relationship between ISUP and GS. For this study, our focus was exclusively on lesions with ISUP scores ≥ 2, as annotations for ISUP-1 lesions, necessary to perform the image pre-processing steps, were unavailable. Consequently, we considered the eight GS classes (from $3 + 4$ to $5 + 5$) as the C_{fine} class set and the four ISUP classes (from 2 to 5) as the C_{coarse} set. For clarification purposes, we provided the labeling criteria for the training, validation, and test subsets in Table 1.

Furthermore, to emulate a realistic scenario with distinct distributions in training and testing data, we considered only Siemens vendor data for training and only Philips vendor data for both validation and testing. We opted for the analogous distributions among validation and test datasets since, as emphasized by Setlur *et al.* [52], using validation samples that are not independent and identically distributed with the test samples can result in unreliable results when determining the optimal model, especially the one that maximizes performance on the test set. The vendor used for the training dataset was also employed for the pre-training dataset, ensuring a complete separation in data distributions between the pre-training/training and evaluation phases.

As for the data pre-processing, we utilized the provided whole prostate segmentation to extract the mask centroid for each slice. We standardized the field of view (FOV) at 100 mm to ensure consistency across all acquisitions and cropped each image based on this value around the found centroid. We determined the number of rows (N_{rows}) and columns (N_{cols}) for the fixed FOV, leveraging the pixel spacing in millimeters along the x-axis (px) and the y-axis (py) according to the following formulae: $N_{cols} = \frac{FOV_x}{px}$ and $N_{rows} = \frac{FOV_y}{py}$. Furthermore, we resized all images to a uniform matrix size of 128×128 pixels to maintain a consistent pixel count. Finally, we performed a z-score normalization on all images using an in-volume approach.
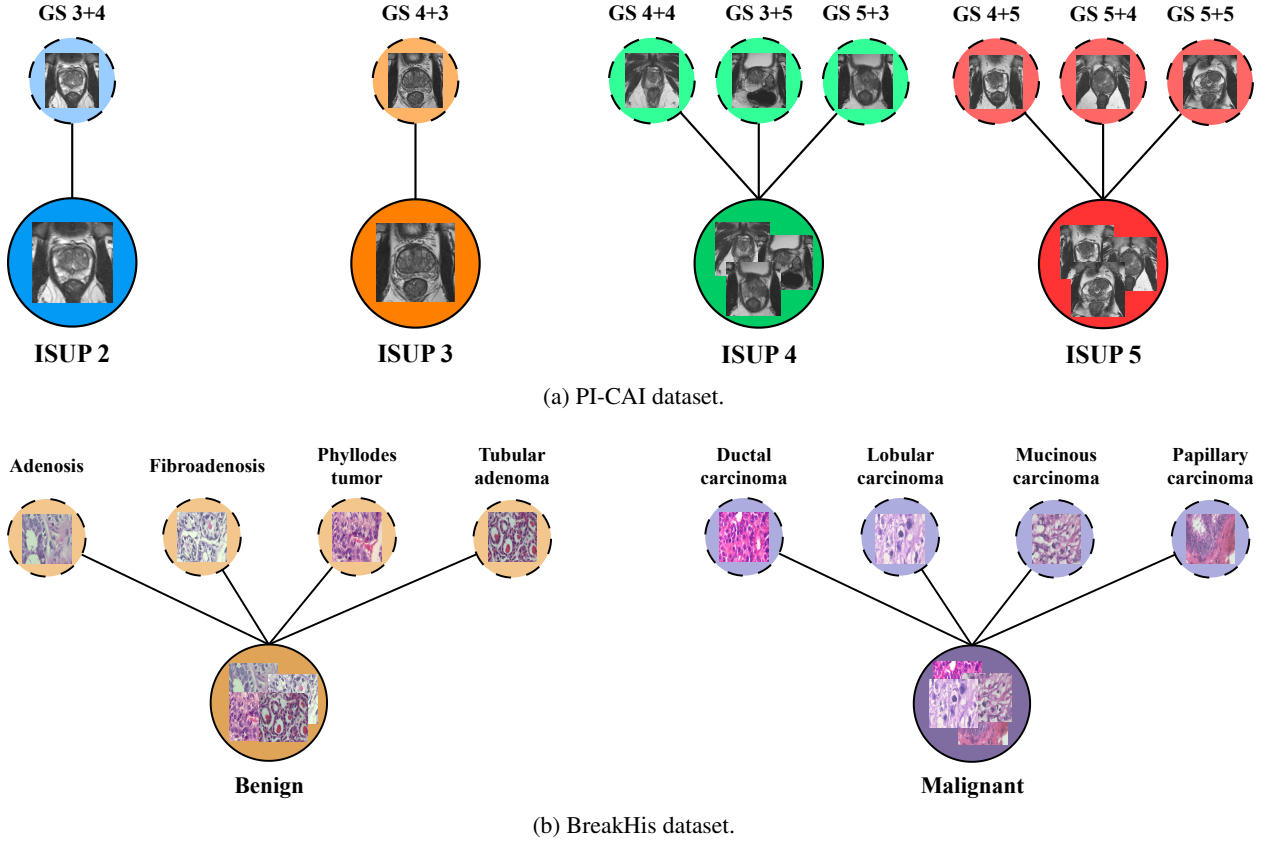
(a) PI-CAI dataset.

(b) BreakHis dataset.

Figure 2: Relationship between fine and coarse labels in (a) PI-CAI dataset and (b) BreakHis dataset.

4.5.2 BreakHis dataset

The dataset consists of 7909 microscopic images of breast tumor tissues from 82 patients, taken at four different magnification levels: 40X, 100X, 200X, and 400X. Each image has a size of 224 by 224 pixels. We used the images at 40X, 100X, and 200X magnification levels for a total of 6090 images as D_{unlab} by ignoring their labels. On the other hand, we used the images at 400X magnification level, namely 1819 images, as D_{lab} to perform the fine-tuning step. We split this dataset into training, validation, and test sets, with 1475, 165, and 183 images, respectively. The dataset has eight classes of lesions which we considered as C_{fine} set: *adenosis*, *tubular adenoma*, *fibroadenoma*, *phyllodes tumor*, *papillary carcinoma*, *lobular carcinoma*, *mucinous carcinoma*, and *ductal carcinoma*. The dataset also provides a binary classification of the lesions into benign and malignant, which we considered as C_{coarse} set. Specifically, the first four fine classes are benign, while the last four are malignant. We provided a visual representation of the relatioship between fine and coarse classes in Fig. 2b. For clarification purposes, we also provided the labeling criteria for training, validation, and test sets of the BreakHis dataset in Table 1.

Table 1: A summary of the labeling procedure for the two examined datasets. ISUP = International Society of Urological Pathology, GS = Gleason Score, AD = Adenosis, TA = Tubular Adenoma, FA = Fibroadenoma, PT = Phyllodes Tumor, PC = Papillary Carcinoma, LC = Lobular Carcinoma, MC = Mucinous Carcinoma, DC = Ductal Carcinoma.

Dataset	Subset	Meta-fine-tuning labels	Fully-supervised fine-tuning labels
PI-CAI	(Meta) Training	GS: $3+4, 4+3, 4+4, 3+5, 5+3, 4+5, 5+4, 5+5$	ISUP: 2, 3, 4, 5
	(Meta) Validation	ISUP: 2, 3, 4, 5	ISUP: 2, 3, 4, 5
	(Meta) Test	ISUP: 2, 3, 4, 5	ISUP: 2, 3, 4, 5
BreakHis	(Meta) Training	AD, TA, FA, PT, PC, LC, MC, DC	Benign, Malignant
	(Meta) Validation	Benign, Malignant	Benign, Malignant
	(Meta) Test	Benign, Malignant	Benign, Malignant

4.6 Experiments

We performed all our experiments using four popular CNNs as backbones, namely, ResNet-18, ResNet-50, VGG-16 and DenseNet-121.

4.6.1 Pre-training

In both the vanilla SSL and SSL+IP-IRM pre-training phases, we conducted a hyperparameter optimization to determine the best configuration for our experiments. This involved evaluating various values for weight decay (WD), batch size (BS), and the number of epochs, as outlined in Table 2. Following the approach established in the original SimCLR paper [42], we set the learning rate (LR) as a function of the BS, as follows:

$$LR = \frac{0.3 \cdot BS}{256}. \tag{6}$$

We opted for large BS values, aligning with the observation made by Chen *et al.* [42] that larger BS generally yield improved performance. Consistent with the original implementation, we conducted our experiments using the SGD optimizer and employed the cosine decay schedule to adjust the learning rate throughout epochs. To mitigate the risk of performance degradation associated with inconsistent input sizes [39], we maintained the same input size during both the pre-training and fine-tuning phases.

As highlighted in [39], performance on proxy tasks does not consistently correlate positively with performance on downstream tasks. Therefore, we selected the optimal pre-training hyperparameters as the ones that maximize the performance of the fine-tuned model. Finally, regarding the composition of augmentations, we utilized the combination that yielded the best performance according to the original paper, i.e., random cropping and random color distortion augmentations.

Table 2: A summary of the hyperparameters optimized during pre-training and meta-fine-tuning phases. The provided learning rate values must be considered as initial values before decay. The learning rate is decreased by a factor of 10 with each decay epoch.

Hyperparameter	Values
Pre-training phase	
Weight decay	$[10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}]$
Batch size	[128,256,512]
Epochs	[100,400]
Meta-fine-tuning phase	
Learning rate	$[10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}]$
Weight decay	$[10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}]$
Epochs	[100,400]
Decay epochs for 100 epochs	[20,50]
Decay epochs for 400 epochs	[80,200]

4.6.2 Fine-tuning

To establish a fair comparison across all experiments, we defined a baseline model, namely the fully-supervised-pretrained model on the ImageNet dataset, fine-tuned using the traditional fully-supervised approach. We optimized its hyperparameters, such as LR, WD, and the number of epochs, and we employed these optimized hyperparameters for all subsequent fine-tuning experiments. To enhance the training process, particularly in terms of optimization and generalization, we exploited a learning rate decay strategy. This strategy involves reducing the LR by a factor of 10 after a predetermined number of warm-up epochs. We reported the hyperparameters evaluated during the baseline optimization in Table 2.

As for our meta-fine-tuning experiments, we performed a *4-way K-shot* training for the PI-CAI dataset and a *2-way K-shot* for the BreakHis dataset. For both datasets, we evaluated both *1-shot* and *5-shot* configurations, resulting in support sets containing either one or five examples per class, respectively. In contrast, the query set always consisted of 10 examples for each class. We conducted 600 episodes throughout both the meta-training and meta-validation phases, as well as the meta-testing phase. Additionally, we replicated the meta-testing evaluation five times.

In all fine-tuning experiments, we refrained from using any data augmentation. This decision was based on evidence suggesting that data augmentation can diminish the positive impact of SSL pre-training and may even negatively impact model performance, as discussed in [39].

4.6.3 Evaluation metrics and loss function

To evaluate the performance of our models, we employed both binary and multi-class Area Under the ROC Curve (AUROC), which proves more stable in the context of accuracy when handling the imbalance often found in medical imaging datasets. For the computation of multi-class AUROC, we adopted the *One-vs-rest* approach, averaging the binary AUROC between each class and all the others. In the meta-learning experiments, AUROC was assessed in each episode, and the final AUROC performance was determined by averaging the AUROC of each episode during the meta-testing phase.

Given our objective of maximizing the AUROC metric, we adopted the AUC margin loss (AUC-M) as our loss function. Introduced in [53] and defined in the LibAUC library [54], AUC-M loss is a surrogate loss function designed to maximize the ROC curve. Specifically, AUC-M loss is a min-max loss function that encourages the model to learn a decision boundary that separates positive and negative examples with a large margin. This property makes the AUC-M loss more robust to noisy data and not adversely affected by simple data. As an optimizer algorithm, following [54], we employed the Proximal Epoch Stochastic method (PESG) [55], which is a stochastic method designed to solve smooth non-convex strongly-concave min-max problems, such as deep AUC maximization [55].

5 Results

The hyperparameter optimization of the baseline model led to the following values for both ResNet-18 and ResNet-50: 10^{-2} LR, 10^{-2} WD, and 100 epochs for the PI-CAI dataset, and 10^{-1} LR, 10^{-2} WD, and 100 epochs for the BreakHis dataset. For the VGG-16, we found that the optimal hyperparameters are 10^{-3} LR, 10^{-5} WD and 100 epoch for the PI-CAI dataset, 10^{-3} LR, 10^{-4} WD and 100 epochs for the BreakHis dataset. Finally, concerning the DenseNet-121 backbone, the hyperparameter optimization provided the following optimal hyperparameters: 10^{-2} LR, 10^{-4} WD and 100 epochs for the PI-CAI dataset and 10^{-2} LR, 10^{-5} WD and 100 epochs for the BreakHis dataset.

Concerning model performance, we presented the results for the PI-CAI dataset in Table 3. There are two main perspectives from which to analyze the results. The first perspective pertains to the pre-training type, which varies according to the Table 3 rows. Specifically, for each backbone, the last row (**SimCLR+IP-IRM**) indicates the pre-training step of the proposed approach, while the other two rows (**Fully-supervised** and **SimCLR**) represent the ablation pre-training experiments. The second perspective relates to the fine-tuning method, which varies with the columns in Table 3. We presented the results of the proposed meta-fine-tuning approach in the grouped column **Meta-train on PI-CAI** *(fine)*. In contrast, we detailed the ablation fine-tuning experiments in the remaining three columns, namely **Meta-train on BreakHis** *(fine)*, **Meta-train on BreakHis** *(coarse)*, and **Fully-supervised**. The **Meta-train on BreakHis** *(fine)* and **Meta-train on BreakHis** *(coarse)* columns showcase the results wherein we utilized different sources for the meta-train and meta-testing datasets, considering both fine-grained and coarse-grained classes during meta-training, respectively. On the other hand, the **Fully-supervised** column presents the results of a classical fully-supervised fine-tuning on the entire dataset. To offer a more immediate visualization, we depicted the results for the PI-CAI dataset in a bar chart presented in Fig. 3a. In this representation, for each backbone and each k-shot setting (1-shot or 5-shot), we delineated the outcomes of various pre-training experiments using bars of different colors, and we portrayed the distinct fine-tuning experiments through separate groups of bars.

Similarly, we presented the results on the BreakHis dataset in Table 4. As in the approach taken for the PI-CAI datasets, we organized the results of different pre-training experiments as distinct rows and the fine-tuning experiments as separate columns in Table 4. In this context, we detailed the outcomes of the proposed fine-tuning approach in column **Meta-train of BreakHis** *(fine)*. Conversely, we delineated the ablation experiment in columns **Meta-train on PI-CAI** *(fine)*, **Meta-train on PI-CAI** *(coarse)*, where a diverse source dataset for meta-training and meta-testing is utilized, and in column **Fully-supervised**, representing the conventional fully-supervised fine-tuning. In line with the PI-CAI dataset, we illustrated the results in Fig. 3b. In the same way as for the PI-CAI dataset, for each backbone and k-shot setting we represented the diverse pre-training experiments through bars of different colors while differentiating the fine-tuning experiments by distinct bar groups.

Pre-training results. For each fine-tuning experiment, in both Table 3 and Table 4, we highlighted the pre-training approach that performs best in bold. Specifically, for the PI-CAI dataset, SimCLR+IP-IRM pre-training produces the most favorable results for all the 5-shot and fully-supervised fine-tuning experiments. In contrast, for three out of six 1-shot experiments (two for the ResNet-18 and one for the ResNet-50), the optimal performance is achieved through SimCLR alone. A similar trend is evident in the BreakHis dataset, where SimCLR alone serves as the best pre-training

(a) PI-CAI dataset.
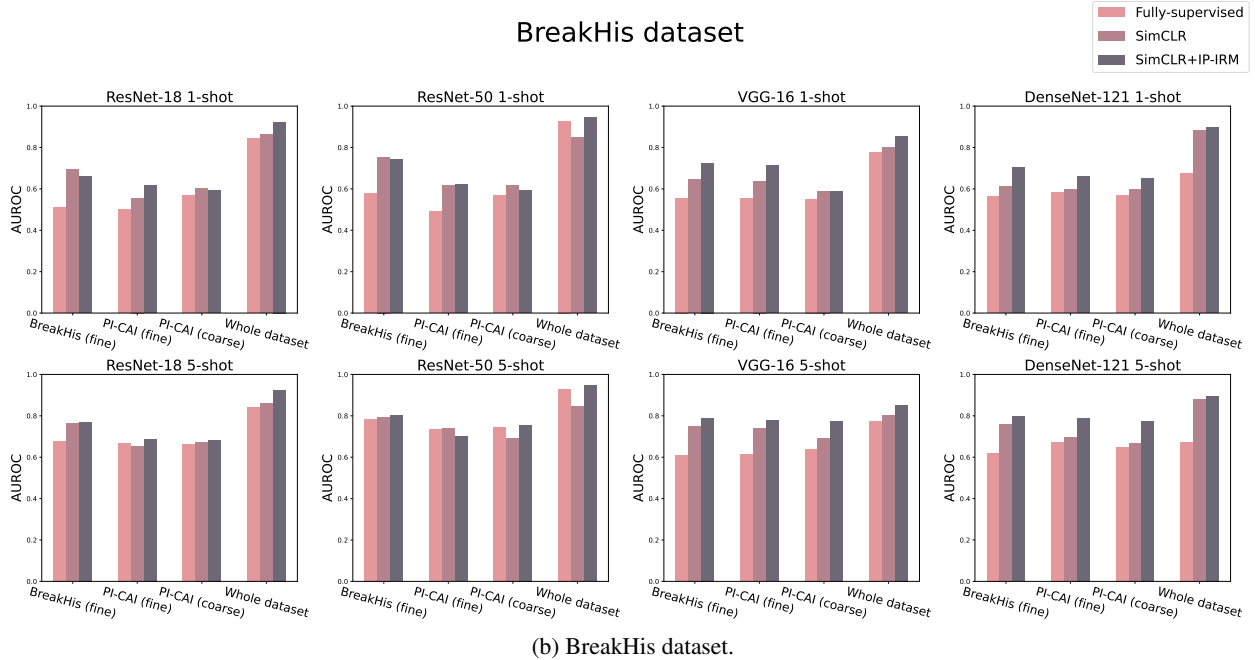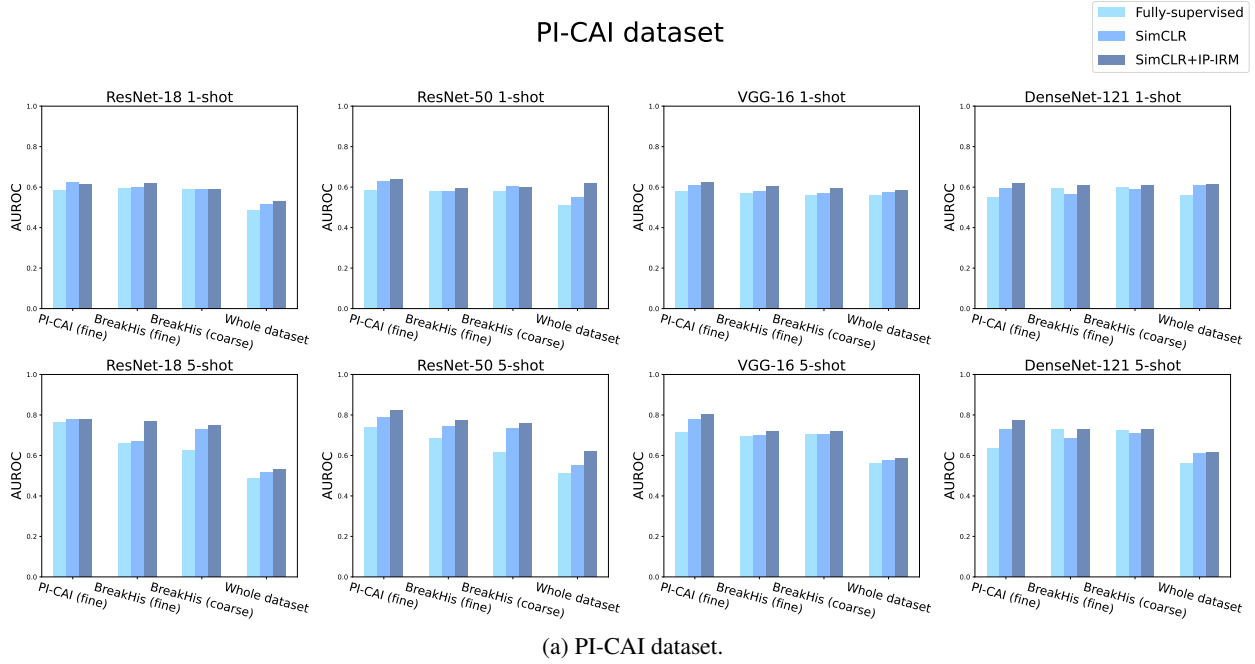

(b) BreakHis dataset.

Figure 3: Results visual representation for (a) PI-CAI dataset and (b) BreakHis dataset. Each plot represents the results of a backbone in a 1-shot or 5-shot setting. We represented each fine-tuning scheme with a three-column group. In each group, the three colours indicate the three pre-training approaches.

Table 3: Test set results in terms of multi-class AUROC for the PI-CAI dataset. In all the experiments, the downstream task consists of classifying the images into four ISUP classes. We experimented with performing meta-training on both finer-grained classes (eight GS classes and eight breast tumor classes, denoted as *fine*) and coarser ones (four malignant breast tumor classes, denoted as *coarse*). For each experiment, we highlighted the proposed and ablation approaches accordingly.

Backbone	Pre-training type	Fine-tuning type						
		Meta-train on PI-CAI (*fine*) (Proposed)		Meta-train on BreakHis (*fine*) (Ablation)		Meta-train on BreakHis (*coarse*) (Ablation)		Fully-supervised (Ablation)
		4-way 1-shot	4-way 5-shot	4-way 1-shot	4-way 5-shot	4-way 1-shot	4-way 5-shot	Whole dataset
ResNet-18	Fully-supervised (Ablation)	0.585	0.763	0.594	0.659	0.588	0.626	0.487
	SimCLR (Ablation)	**0.624**	0.779	0.600	0.667	**0.592**	0.727	0.515
	SimCLR + IP-IRM (Proposed)	0.615	**0.780**	0.618	0.767	0.591	0.749	0.533
ResNet-50	Fully-supervised (Ablation)	0.587	0.740	0.582	0.683	0.578	0.616	0.509
	SimCLR (Ablation)	0.628	0.790	0.581	0.744	**0.605**	0.732	0.553
	SimCLR + IP-IRM (Proposed)	**0.640**	**0.821**	0.595	0.773	0.602	0.756	0.620
VGG-16	Fully-supervised (Ablation)	0.579	0.712	0.569	0.693	0.563	0.704	0.562
	SimCLR (Ablation)	0.608	0.776	0.578	0.701	0.571	0.706	0.576
	SimCLR + IP-IRM (Proposed)	**0.626**	**0.803**	0.603	0.720	**0.593**	0.717	0.585
DenseNet-121	Fully-supervised (Ablation)	0.549	0.635	0.597	0.727	0.599	0.726	0.560
	SimCLR (Ablation)	0.597	0.730	0.566	0.685	0.589	0.709	0.611
	SimCLR + IP-IRM (Proposed)	**0.621**	**0.771**	**0.612**	**0.729**	**0.610**	**0.727**	**0.617**

Table 4: Test set results in terms of binary AUROC for the BreakHis dataset. In all the experiments, the downstream task consists of classifying the images into two classes (Benign vs. Malignant). We experimented with performing meta-training on both finer-grained classes (eight breast tumor classes and eight GS classes, denoted as *fine*) and coarser ones (four ISUP classes, denoted as *coarse*). For each experiment, we highlighted the proposed and ablation approaches accordingly.

Backbone	Pre-training type	Fine-tuning type						
		Meta-train on BreakHis (*fine*) (Proposed)		Meta-train on PI-CAI (*fine*) (Ablation)		Meta-train on PI-CAI (*coarse*) (Ablation)		Fully-supervised (Ablation)
		2-way 1-shot	2-way 5-shot	2-way 1-shot	2-way 5-shot	2-way 1-shot	2-way 5-shot	Whole dataset
ResNet-18	Fully-supervised (Ablation)	0.510	0.678	0.499	0.669	0.570	0.661	0.843
	SimCLR (Ablation)	**0.694**	0.763	0.556	0.654	**0.601**	0.675	0.862
	SimCLR + IP-IRM (Proposed)	0.658	**0.772**	0.618	0.688	0.594	**0.685**	0.923
ResNet-50	Fully-supervised (Ablation)	0.576	0.783	0.490	0.735	0.568	0.746	0.928
	SimCLR (Ablation)	**0.754**	0.793	0.615	**0.739**	**0.615**	0.692	0.847
	SimCLR + IP-IRM (Proposed)	0.742	**0.802**	0.620	0.700	0.591	**0.754**	0.946
VGG-16	Fully-supervised (Ablation)	0.556	0.608	0.554	0.613	0.551	0.639	0.775
	SimCLR (Ablation)	0.644	0.750	0.638	0.739	**0.590**	0.692	0.802
	SimCLR + IP-IRM (Proposed)	**0.725**	**0.791**	**0.712**	**0.777**	0.588	**0.775**	0.853
DenseNet-121	Fully-supervised (Ablation)	0.565	0.620	0.584	0.675	0.568	0.651	0.673
	SimCLR (Ablation)	0.613	0.761	0.597	0.699	0.596	0.669	0.883
	SimCLR + IP-IRM (Proposed)	**0.702**	**0.797**	**0.658**	**0.790**	**0.649**	**0.774**	**0.896**

step for five 1-shot configurations (two for the ResNet-18, two for the ResNet-50 and one for the VGG-16). However, for all 5-shot configurations except one, as well as the fully-supervised approach, consistent enhancements are observed with SimCLR+IP-IRM pre-training.

Meta-fine-tuning results. In describing the meta-fine-tuning results, among the pre-training approaches (three for each meta-fine-tuning strategy), we refer to the one that provided the best AUROC value. Concerning the PI-CAI dataset, according to Table 3, the optimal performance is obtained by the proposed approach (first fine-tuning type column of Table 3), yielding 0.780 AUROC for ResNet-18, 0.821 for ResNet-50, 0.803 for VGG-16 and 0.771 for DenseNet-121. In terms of the ablation experiments, employing the meta-learning approach with granularity and dataset shift (second fine-tuning type column in Table 3) provides the second-best performance, achieving the highest AUROC of 0.767 for ResNet-18, 0.773 for ResNet-50, 0.720 for VGG-16 and 0.729 for DenseNet-121. Lastly, when meta-training is conducted on coarser classes (third fine-tuning type column in Table 3), classification performance is further degraded, with an AUROC of 0.749 for ResNet-18, 0.756 for ResNet-50, 0.717 for VGG-16 and of 0.727 for DenseNet-121.

Exploring the BreakHis dataset, the proposed meta-fine-tuning configuration (first fine-tuning type column of Table 4) yields the most favorable results for all the backbones, with AUROCs of 0.772, 0.802, 0.791 and 0.797 for ResNet-18, ResNet-50, VGG-16 and DenseNet-121, respectively. Considering the ablation experiments, for ResNet-18, VGG-16 and DenseNet-121, the second-best performance is observed when both granularity and dataset source shift are present (second fine-tuning type column of Table 4), resulting in 0.688, 0.777 and 0.790 AUROC, respectively. On the other hand, for the same three backbones, the least favorable outcome is obtained when coarser-grained classes are employed for meta-training (third fine-tuning type column of Table 4), yielding an AUROC of 0.685, 0.775 and 0.774, respectively. Conversely, for ResNet-50, using coarser classes for meta-training produces superior results, achieving a 0.754 AUROC compared to 0.739 AUROC obtained with finer-grained classes.

Fully-supervised fine-tuning results. As the last ablation experiment, we considered performing a fully-supervised fine-tuning on the entire dataset. For the PI-CAI dataset, this fine-tuning strategy (results provided in the fourth fine-tuning type column of Table 3) exhibits the worst performance w.r.t. all the 5-shot settings and even w.r.t. several 1-shot, especially when the proposed meta-fine-tuning approach is employed. Indeed, for the four backbones the higher AUROC is 0.533, 0.620, 0.585 and 0.617, respectively. This underscores that the meta-learning approach significantly enhances the generalization ability of both backbones, no matter the low-data regime, and that the proposed meta-fine-tuning approach surpasses the fully-supervised one even in the 1-shot regime. Conversely, for the BreakHis dataset, the fully-supervised fine-tuning approach (fourth fine-tuning type column of Table 4) exhibits considerable strength, overcoming all the results provided by the meta-fine-tuning strategies. Indeed, this approach yields 0.923, 0.946, 0.853 and 0.896 AUROC for ResNet-18, ResNet-50, VGG-16 and DenseNet-121, respectively.

6 Discussion

In this work, we introduced an approach to enhance the performance of models trained in an FSL paradigm. The proposed strategy comprises a pre-training step that utilizes SSL coupled with the IP-IRM algorithm, which forces the feature disentanglement, ultimately guiding the model in learning more robust feature representation for the downstream task and making it distinguish between invariant and task-specific features. In addition, we proposed an enhanced meta-learning framework that leverages classes between the meta-training and meta-testing phases at different granularity levels. This approach aims to improve the model's generalization ability by exposing it to a more diverse and structured meta-training phase.

To evaluate the efficacy of the proposed pre-training strategy, we conducted two additional ablation experiments, i.e., a vanilla SSL method and a classical fully-supervised approach. As for the pre-training phase, to assess the effectiveness of the proposed meta-learning framework, we explored three additional ablation fine-tuning experiments as well, namely two meta-fine-tuning and a fully-supervised fine-tuning experiment. We delve into the meta-fine-tuning experiments in the following:

- **Same source dataset with granularity shift (Proposed)**: Both meta-training and meta-testing sets derive from the same source dataset, but the granularity of classes differs. We employed finer-grained classes for meta-training to promote detailed feature learning and coarser-grained ones for meta-testing to assess the model's ability to generalize to broader categories.

- **Different source datasets with granularity shift (Ablation)**: Meta-training and meta-testing sets are drawn from distinct source datasets, but the granularity relationship between the classes remains consistent as in the previous case. This approach aims to highlight the effectiveness of the proposed granularity shift even when meta-training and meta-testing have different source datasets.

- **Different Source Datasets with granularity consistency (Ablation)**: Meta-training and meta-testing sets stem from different source datasets, but the granularity between classes is maintained consistently. We designed this approach to evaluate whether reducing class granularity during the meta-training impacts the model's generalizability in the subsequent meta-testing phase.

In terms of pre-training, concerning the PI-CAI dataset, the results presented in Table 3 demonstrate that using the IP-IRM algorithm consistently provides the best-performing outcomes across all fully-supervised and 5-shot fine-tuning experiments. However, in the 1-shot experiments, the SimCLR alone outperforms the SimCLR+IP-IRM in three out of twelve configurations. Regarding the BreakHis dataset, the proposed pre-training approach exhibits consistent improvement in most of the experiments. However, similar to the PI-CAI dataset, some 1-shot experiments reveal that SimCLR alone yields better results w.r.t. the IP-IRM algorithm. One possible explanation for this behaviour is that during the pre-training phase, the IP-IRM algorithm may lead to an over-disentangled data representation, resulting in excessively specialized representations and overly sensitive to variations in the training set. Consequently, in a 1-shot setting where the model must learn from only one example and may not have the opportunity to grasp all relevant features, it might struggle to generalize effectively, thus providing worse performance.

Regarding the fine-tuning phase on the PI-CAI dataset, it is clear that, for all backbones, the fully-supervised approach establishes a lower bound, suggesting that the traditional training methodology alone is inadequate for this task. Indeed, in all the experiments, the proposed meta-fine-tuning scheme exceeds the fully-supervised performance even in a 1-shot setting. In particular, the proposed meta-learning scheme delivers the best-performing results in both 1-shot and 5-shot scenarios, achieving a peak classification performance of 0.780, 0.821, 0.803 and 0.771 multiclass AUROC for ResNet-18, ResNet-50, VGG-16 and DenseNet-121 models, respectively. For each backbone, we present the ROC curves for a randomly selected meta-testing episode of the best-performing combination of pre-training and meta-fine-tuning schemes in Fig.4a, Fig.4b, Fig.4c, and Fig.4d.

In contrast, for the BreakHis dataset, the fully-supervised model demonstrates strong generalization capabilities. This robust performance may be attributed to the lower complexity of the task compared to the classification of prostate cancer aggressiveness, stemming from several reasons. Firstly, the breast malignity classification task involves distinguishing only two classes in contrast to the four classes of the prostate case. Additionally, while microscopic images of breast cancer typically exhibit the presence of the disease throughout the image with cell nuclei indicating the presence of the tumor [56] highlighted compared to background structures [57], prostate MRI images depict lesions typically occupying only a few pixels within the image [58]. Furthermore, the lesion signal in prostate T2-weighted images may be isointense w.r.t. to the surrounding tissues [59], making its detection and categorization more challenging for the model. Lastly, in our PI-CAI experiments, we deliberately introduced an additional layer of complexity by explicitly considering different vendors between training and evaluation data. This aspect provides an additional level of complexity not present in the BreakHis dataset. Consequently, for the BreakHis dataset, the restricted data availability during meta-learning training tends to negatively impact the overall performance despite the enhanced generalization capabilities of the meta-learning framework.

Among the meta-learning experiments, the proposed meta-learning strategy performs best in both 1-shot and 5-shot settings, achieving peak performance with an AUROC of 0.772, 0.802, 0.791 and 0.797 for ResNet-18, ResNet-50, VGG-16 and DenseNet-121, respectively. We illustrate the ROC curves for a randomly chosen meta-testing episode of the best-performing pre-training and meta-finetuning combination for each backbone in Fig. 4e, Fig. 4f, Fig. 4g and Fig. 4h. Concerning the two ablation meta-learning experiments, for ResNet-18, VGG-16 and DenseNet-121, the performance trend is generally coherent, i.e., the performance degrades when considering a different meta-training source and especially when using a coarser class set for meta-training. The only exception is provided by the ResNet-50 backbone where, in the 5-shot setting, meta-training on coarser classes yields a higher AUROC (0.754) compared to meta-training on finer ones (0.739 AUROC).

(a) ResNet-18, PI-CAI dataset   (b) ResNet-50, PI-CAI dataset   (c) VGG-16, PI-CAI dataset   (d) DenseNet-121, PI-CAI dataset

(e) ResNet-18, BreakHis dataset   (f) ResNet-50, BreakHis dataset   (g) VGG-16, BreakHis dataset   (h) DenseNet-121, BreakHis dataset
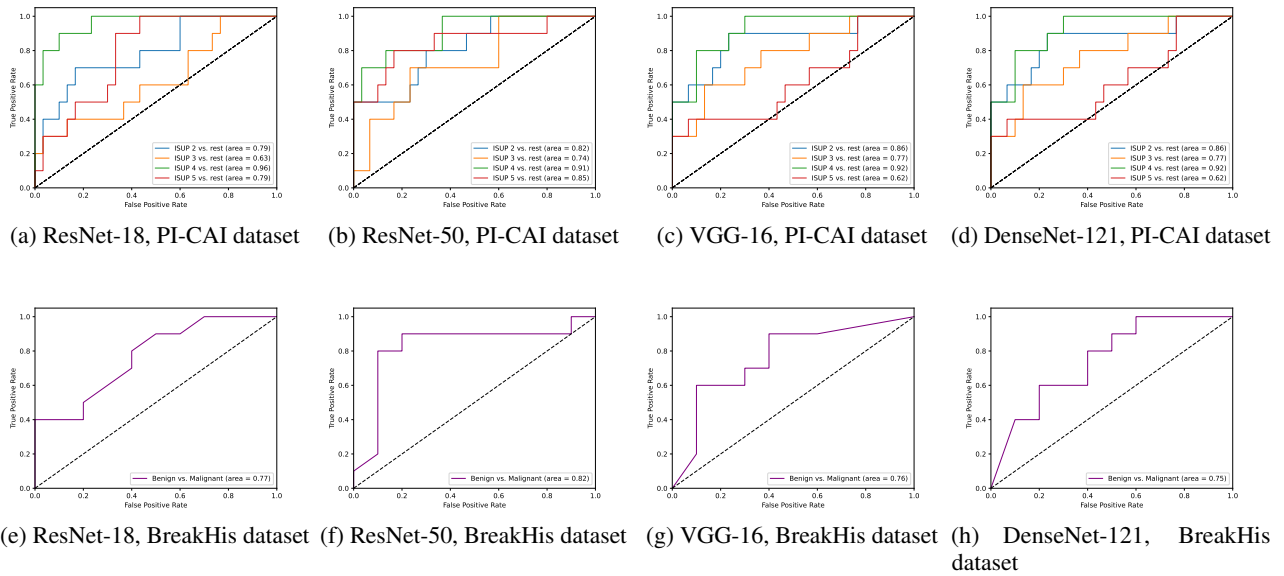
Figure 4: ROC curves for one random episode of meta-testing of each backbone pre-trained with SimCLR+IP-IRM and meta-fine-tuned in a 5-shot setting. Meta-training and meta-testing classes have the same source dataset (PI-CAI or BreakHis).

Overall, our experiments demonstrate the effectiveness of our proposed methods, which boost the FSL capabilities by embedding a disentangled SSL pre-training step and a novel meta-training scheme, outperforming most of the ablation experiments in both clinical challenges. Specifically, the disentangled SSL pre-training step proves to be particularly effective as the training data increases, whereas, in a 1-shot setting, using vanilla SimCLR may sometimes provide better performance. Additionally, when the classification task is particularly challenging, as in the PI-CAI dataset, the proposed strategy significantly improves the classification capabilities of the model compared to a classical fully-supervised approach on the entire dataset. Furthermore, our results on the PI-CAI dataset demonstrate the strong generalization capabilities of the proposed approach even when training and evaluation data come from different vendors. Either way, when the task is easier, as for the BreakHis dataset, and the fully-supervised baseline is stronger, our approach allows for achieving competitive classification capabilities in a few-data regime. Finally, our findings demonstrate that even

when meta-training and meta-testing data originate from different sources (unrelated meta-training and meta-testing classes), generating meta-training episodes from finer class sets contributes to improved generalization capabilities.

7 Conclusions

In this study, we proposed a novel method for boosting the capabilities of FSL-trained models by leveraging disentangled SSL and an enhanced meta-learning framework, and we demonstrated its effectiveness through extensive experiments and ablation studies on relevant medical imaging tasks. Furthermore, the diversity of clinical scenarios experimented demonstrated the versatility and potential broad application of our approach. Future work may delve into enriching the meta-learning algorithm's intrinsic capabilities by generating more informative prototypes leveraging support embeddings hallucination.

CRediT authorship contribution statement
Eva Pachetti: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data Curation, Writing - Original Draft, Writing - Review & Editing, Visualization. **Sotirios A. Tsaftaris**: Conceptualization, Methodology, Writing - Review & Editing, Supervision. **Sara Colantonio**: Conceptualization, Resources, Writing - Review & Editing, Supervision, Project administration, Funding acquisition.

Declaration of Competing Interest
The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3):1–34, 2020. doi:https://doi.org/10.1145/3386252.

[2] Gianluca Carloni, Eva Pachetti, and Sara Colantonio. Causality-driven one-shot learning for prostate cancer grading from mri. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2616–2624, 2023. doi:10.1109/ICCVW60793.2023.00276.

[3] Zhiyong Dai, Jianjun Yi, Lei Yan, Qingwen Xu, Liang Hu, Qi Zhang, Jiahui Li, and Guoqiang Wang. Pfemed: Few-shot medical image classification using prior guided feature enhancement. *Pattern Recognition*, 134:109108, 2023. doi:https://doi.org/10.1016/j.patcog.2022.109108.

[4] Hongyang Jiang, Mengdi Gao, Heng Li, Richu Jin, Hanpei Miao, and Jiang Liu. Multi-Learner Based Deep Meta-Learning for Few-Shot Medical Image Classification. *IEEE Journal of Biomedical and Health Informatics*, 27(1):17–28, 2022. doi:10.1109/JBHI.2022.3215147. Publisher: IEEE.

[5] Eva Pachetti and Sara Colantonio. A systematic review of few-shot learning in medical imaging. *arXiv*, 2023. doi:https://doi.org/10.48550/arXiv.2309.11433.

[6] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. *arXiv*, 2019. doi:https://doi.org/10.48550/arXiv.1904.04232.

[7] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. Boosting few-shot visual learning with self-supervision. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8059–8068, 2019. doi:10.1109/ICCV.2019.00815.

[8] Jong-Chyi Su, Subhransu Maji, and Bharath Hariharan. When does self-supervision improve few-shot learning? In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, pages 645–666. Springer, 2020. doi:https://doi.org/10.1007/978-3-030-58571-6_38.

[9] Tan Wang, Zhongqi Yue, Jianqiang Huang, Qianru Sun, and Hanwang Zhang. Self-supervised learning disentangled group representation as feature. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman

Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 18225–18240. Curran Associates, Inc., 2021. URL `https://proceedings.neurips.cc/paper_files/paper/2021/file/97416ac0f58056947e2eb5d5d253d4f2-Paper.pdf`.

[10] Jiangtao Xie, Fei Long, Jiaming Lv, Qilong Wang, and Peihua Li. Joint distribution matters: Deep brownian distance covariance for few-shot classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7972–7981, 2022. doi:10.1109/CVPR52688.2022.00781.

[11] Anindo Saha, Joeran Bosma, Jasper Twilt, Bram van Ginneken, Derya Yakar, Mattijs Elschot, Jeroen Veltman, Jurgen Fütterer, Maarten de Rooij, et al. Artificial intelligence and radiologists at prostate cancer detection in mri—the pi-cai challenge. In *Medical Imaging with Deep Learning, short paper track*, 2023.

[12] Fabio A Spanhol, Luiz S Oliveira, Caroline Petitjean, and Laurent Heutte. A dataset for breast cancer histopathological image classification. *Ieee transactions on biomedical engineering*, 63(7):1455–1462, 2015. doi:10.1109/TBME.2015.2496264.

[13] Lars Egevad, Brett Delahunt, John R. Srigley, and Hemamali Samaratunga. International Society of Urological Pathology (ISUP) grading of prostate cancer - An ISUP consensus on contemporary grading. *APMIS: acta pathologica, microbiologica, et immunologica Scandinavica*, 124(6):433–435, June 2016. ISSN 1600-0463. doi:10.1111/apm.12533.

[14] Hendrik Van Poppel, Renée Hogenhout, Peter Albers, Roderick C.N. van den Bergh, Jelle O. Barentsz, and Monique J. Roobol. Early detection of prostate cancer in 2020 and beyond: Facts and recommendations for the european union and the european commission. *European Urology*, 79(3):327–329, 2021. ISSN 0302-2838. doi:https://doi.org/10.1016/j.eururo.2020.12.010. URL `https://www.sciencedirect.com/science/article/pii/S0302283820309581`.

[15] Kalpana George, Praveen Sankaran, and Paul Joseph K. Computer assisted recognition of breast cancer in biopsy images via fusion of nucleus-guided deep convolutional features. *Computer Methods and Programs in Biomedicine*, 194:105531, 2020. ISSN 0169-2607. doi:https://doi.org/10.1016/j.cmpb.2020.105531. URL `https://www.sciencedirect.com/science/article/pii/S0169260720302911`.

[16] Yassir Benhammou, Boujemâa Achchab, Francisco Herrera, and Siham Tabik. Breakhis based breast cancer automatic diagnosis using deep learning: Taxonomy, survey and insights. *Neurocomputing*, 375:9–24, 2020. ISSN 0925-2312. doi:https://doi.org/10.1016/j.neucom.2019.09.044. URL `https://www.sciencedirect.com/science/article/pii/S0925231219313128`.

[17] Rishav Singh, Vandana Bharti, Vishal Purohit, Abhinav Kumar, Amit Kumar Singh, and Sanjay Kumar Singh. MetaMed: Few-shot medical image classification using gradient-based meta-learning. *Pattern Recognition*, 120:108111, 2021. doi:https://doi.org/10.1016/j.patcog.2021.108111. Publisher: Elsevier.

[18] Da Chen, Yuefeng Chen, Yuhong Li, Feng Mao, Yuan He, and Hui Xue. Self-supervised learning for few-shot image classification. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1745–1749, 2021. doi:10.1109/ICASSP39728.2021.9413783.

[19] Xuanchi Chen, Xiangwei Zheng, Kai Sun, Weilong Liu, and Yuang Zhang. Self-supervised vision transformer-based few-shot learning for facial expression recognition. *Information Sciences*, 634:206–226, 2023. doi:https://doi.org/10.1016/j.ins.2023.03.105.

[20] Carlos Medina, Arnout Devos, and Matthias Grossglauser. Self-supervised prototypical transfer learning for few-shot classification. *arXiv*, 2020. doi:https://doi.org/10.48550/arXiv.2006.11325.

[21] Zhanyuan Yang, Jinghua Wang, and Yingying Zhu. Few-shot classification with contrastive learning. In *European Conference on Computer Vision*, pages 293–309. Springer, 2022. doi:https://doi.org/10.1007/978-3-031-20044-1_17.

[22] Zixuan Hu, Li Shen, Shenqi Lai, and Chun Yuan. Task-adaptive feature disentanglement and hallucination for few-shot classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 2023. doi:10.1109/TCSVT.2023.3238804.

[23] Chia-Ching Lin, Hsin-Li Chu, Yu-Chiang Frank Wang, and Chin-Laung Lei. Joint feature disentanglement and hallucination for few-shot image classification. *IEEE Transactions on Image Processing*, 30:9245–9258, 2021. doi:10.1109/TIP.2021.3124322.

[24] Hao Cheng, Yufei Wang, Haoliang Li, Alex C Kot, and Bihan Wen. Disentangled feature representation for few-shot image classification. *IEEE Transactions on Neural Networks and Learning Systems*, 2023. doi:10.1109/TNNLS.2023.3241919.

[25] Zhenyao Zhu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Multi-view perceptron: a deep model for learning face identity and view representations. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL https://proceedings.neurips.cc/paper_files/paper/2014/file/140f6969d5213fd0ece03148e62e461e-Paper.pdf.

[26] Jun-Ting Hsieh, Bingbin Liu, De-An Huang, Li Fei-Fei, and Juan Carlos Niebles. Learning to decompose and disentangle representations for video prediction. *CoRR*, abs/1806.04166, 2018. doi:https://doi.org/10.48550/arXiv.1806.04166.

[27] Ruichu Cai, Zijian Li, Pengfei Wei, Jie Qiao, Kun Zhang, and Zhifeng Hao. Learning disentangled semantic representation for domain adaptation. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 2060–2066. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi:10.24963/ijcai.2019/285. URL https://doi.org/10.24963/ijcai.2019/285.

[28] Scott Reed, Kihyuk Sohn, Yuting Zhang, and Honglak Lee. Learning to disentangle factors of variation with manifold interaction. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1431–1439, Bejing, China, 22–24 Jun 2014. PMLR. URL https://proceedings.mlr.press/v32/reed14.html.

[29] Theofanis Karaletsos, Serge J. Belongie, and Gunnar Rätsch. Bayesian representation learning with oracle constraints. *arXiv*, 2015. doi:https://doi.org/10.48550/arXiv.1506.05011.

[30] Xiao Liu, Pedro Sanchez, Spyridon Thermos, Alison Q O'Neil, and Sotirios A Tsaftaris. Learning disentangled representations in the imaging domain. *Medical Image Analysis*, 80:102516, 2022. doi:https://doi.org/10.1016/j.media.2022.102516.

[31] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4114–4124. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/locatello19a.html.

[32] Xi Chen, Yan Duan, Rein Houthooft, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/file/7c9d0b1f96aebd7b5eca8c3edaa19ebb-Paper.pdf.

[33] Zinan Lin, Kiran Thekumparampil, Giulia Fanti, and Sewoong Oh. InfoGAN-CR and ModelCentrality: Self-supervised model training and selection for disentangling GANs. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6127–6139. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/lin20e.html.

[34] Utkarsh Ojha, Krishna Kumar Singh, Cho-Jui Hsieh, and Yong Jae Lee. Elastic-infogan: Unsupervised disentangled representation learning in class-imbalanced data. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 18063–18075. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/d1e39c9bda5c80ac3d8ea9d658163967-Paper.pdf.

[35] Ricky T. Q. Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/1ee3dfcd8a0645a25a35977997223d22-Paper.pdf.

[36] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2649–2658. PMLR, 10–15 Jul 2018. URL https://proceedings.mlr.press/v80/kim18b.html.

[37] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*, 2016.

[38] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv*, 2019. doi:10.48550/arXiv.1907.02893.

[39] Chuyan Zhang, Hao Zheng, and Yun Gu. Dive into the details of self-supervised learning for medical image analysis. *Medical Image Analysis*, 89:102879, 2023. doi:https://doi.org/10.1016/j.media.2023.102879.

[40] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.

[41] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015. doi:https://doi.org/10.1109/ICCV.2015.167.

[42] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/chen20j.html.

[43] James Woodward. *Making things happen: A theory of causal explanation*. Oxford university press, 2005.

[44] V. Vapnik. Principles of risk minimization for learning theory. In J. Moody, S. Hanson, and R.P. Lippmann, editors, *Advances in Neural Information Processing Systems*, volume 4. Morgan-Kaufmann, 1991. URL https://proceedings.neurips.cc/paper_files/paper/1991/file/ff4d5fbbafdf976cfdc032e3bde78de5-Paper.pdf.

[45] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/cb8da6767461f2812ae4290eac7cbc42-Paper.pdf.

[46] Davis Wertheimer and Bharath Hariharan. Few-shot learning with localization in realistic settings. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6558–6567, 2019. doi:10.1109/CVPR.2019.00672.

[47] Wenbin Li, Lei Wang, Jing Huo, Yinghuan Shi, Yang Gao, and Jiebo Luo. Asymmetric distribution measure for few-shot learning. *arXiv*, 2020. doi:https://doi.org/10.48550/arXiv.2002.00153.

[48] 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*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.

[49] Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, et al. Meta-dataset: A dataset of datasets for learning to learn from few examples. *arXiv*, 2019. doi:https://doi.org/10.48550/arXiv.1903.03096.

[50] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/file/90e1357833654983612fb05e3ec9148c-Paper.pdf.

[51] Mayke Pereira. Breakhis - breast cancer histopathological database, 2023.

[52] Amrith Setlur, Oscar Li, and Virginia Smith. Two sides of meta-learning evaluation: In vs. out of distribution. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 3770–3783. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/1e932f24dc0aa4e7a6ac2beec387416d-Paper.pdf.

[53] Zhuoning Yuan, Yan Yan, Milan Sonka, and Tianbao Yang. Large-scale robust deep auc maximization: A new surrogate loss and empirical studies on medical image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3040–3049, 2021. doi:10.1109/ICCV48922.2021.00303.

[54] Zhuoning Yuan, Dixian Zhu, Zi-Hao Qiu, Gang Li, Xuanhui Wang, and Tianbao Yang. Libauc: A deep learning library for x-risk optimization. In *29th SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023.

[55] Zhishuai Guo, Yan Yan, Zhuoning Yuan, and Tianbao Yang. Fast objective & duality gap convergence for non-convex strongly-concave min-max problems with pl condition. *Journal of Machine Learning Research*, 24: 1–63, 2023.

[56] Barbara Young, Geraldine O'Dowd, and Phillip Woodford. *Wheater's functional histology E-Book: a text and colour atlas*. Elsevier Health Sciences, 2013.

[57] Lei He, L Rodney Long, Sameer Antani, and George R Thoma. Histology image analysis for carcinoma detection and grading. *Computer methods and programs in biomedicine*, 107(3):538–556, 2012. doi:https://doi.org/10.1016/j.cmpb.2011.12.007.

[58] Indrani Bhattacharya, David S Lim, Han Lin Aung, Xingchen Liu, Arun Seetharaman, Christian A Kunder, Wei Shao, Simon JC Soerensen, Richard E Fan, Pejman Ghanouni, et al. Bridging the gap between prostate radiology and pathology through machine learning. *Medical Physics*, 49(8):5160–5181, 2022. doi:https://doi.org/10.1002/mp.15777.

[59] Catherine Elizabeth Lovegrove, Mudit Matanhelia, Jagpal Randeva, David Eldred-Evans, Henry Tam, Saiful Miah, Mathias Winkler, Hashim U Ahmed, and Taimur T Shah. Prostate imaging features that indicate benign or malignant pathology on biopsy. *Translational andrology and urology*, 7(Suppl 4):S420, 2018. doi:10.21037/tau.2018.07.06.