







Co-Training with Active Contrastive Learning and Meta-Pseudo-Labeling on 2D Projections for Deep Semi-Supervised Learning

David Aparco-Cardenas   [University of Campinas | david.cardenas@ic.unicamp.br]

Jancarlo F. Gomes  [University of Campinas | jgomes@ic.unicamp.br]

Alexandre X. Falcão  [University of Campinas | afalcao@ic.unicamp.br]

Pedro J. de Rezende  [University of Campinas | rezende@ic.unicamp.br]

 Institute of Computing, University of Campinas, Av. Albert Einstein, 1251, Campinas, São Paulo, 13083-852, Brazil.

Received: 11 April 2025 • Accepted: 29 May 2026 • Published: 10 June 2026

Abstract A major challenge that prevents the training of deep learning models is the limited availability of large quantities of accurately labeled data. This shortcoming is particularly acute in areas such as medical and biological sciences where data annotation is an expert-demanding, time-consuming, and error-prone undertaking. In this regard, semi-supervised learning tackles this challenge by capitalizing on scarce labeled and abundant unlabeled data; however, state-of-the-art methods typically depend on pre-trained features and large validation sets to learn effective representations for classification tasks. In addition, the reduced set of labeled data is often randomly sampled, neglecting the selection of more informative samples. Here, we present *active Deep Feature Annotation* (active-DeepFA), a method that effectively combines contrastive learning, teacher-student-based meta-pseudo-labeling and active learning to train non-pre-trained CNN architectures for image classification in scenarios of scarcity of labeled and abundance of unlabeled data. It integrates deep feature annotation (DeepFA) into a co-training setup that implements two cooperative networks to mitigate confirmation bias arising from pseudo-labels. The method starts with a reduced set of labeled samples by warming up the networks with supervised contrastive learning. Afterward and at regular epoch intervals, label propagation is performed on the 2D projections of the networks' deep features. Next, the most reliable pseudo-labels are exchanged between networks in a cross-training fashion, while the most meaningful samples are annotated and added to the labeled set. The networks independently minimize an objective loss function comprising supervised contrastive, supervised, and semi-supervised loss components, enhancing the representations for image classification. Our approach is evaluated on seven challenging image datasets across three distinct domains with only 5% of labeled samples, surpassing baselines and outperforming seven established benchmarks. In addition, it reduces annotation effort by achieving comparable results to those of its counterparts with only 3% of labeled data.

Keywords: Semi-supervised learning, Contrastive learning, Pseudo labeling, Deep feature annotation

1 Introduction

Over the past few years, deep learning methodologies have garnered significant attention and acclaim due to their remarkable versatility in addressing a range of challenges across a wide spectrum of disciplines, including computer vision, natural language processing, and speech recognition. In addition, the predictive capability of deep learning, combined with its automatic feature extraction, makes it a favored technique in medical diagnosis [Dong *et al.*, 2021].

These methodologies have risen to prominence for achieving state-of-the-art results in various computer vision tasks, notably in image classification, segmentation, and object detection. This success is largely due to the vast availability of data and the simultaneous progress in both deep learning algorithms and computational power. Despite that, available data are frequently unlabeled in practical situations, presenting a significant hurdle for fully supervised methods that greatly depend on labeled data. This issue is particularly acute in fields such as medical and biological sciences, where data are very scarce and data labeling demands expertise and substantial

time investment.

In light of this scenario, *semi-supervised learning* (SSL) methods attempt to address the short supply of labeled data by utilizing unlabeled data to improve the generalization performance of predictive classifiers. In the field of deep neural networks, this approach is known as *deep semi-supervised learning* (DSSL). Over time, a wide array of DSSL methods has emerged, showcasing various strategies and approaches [Yang *et al.*, 2022]. Out of these, *pseudo-labeling* is renowned for its prominence due to its simple yet effective learning scheme. It essentially works by assigning pseudo-labels to unlabeled data and selecting the samples with the highest-confidence predictions, which in turn are used to train and improve the model.

Recently, Benato *et al.* [2021b, 2023] introduced *confidence Deep Feature Annotation* (conf-DeepFA), a teacher-student-based meta-pseudo-labeling method to train CNNs in semi-supervised settings. It builds on the *Deep Feature Annotation* (DeepFA) technique [Benato *et al.*, 2021a] by exploiting the non-linear 2D projections of the student pre-trained CNN model's deep features for downstream label propagation

within an iterative scheme. The 2D projection is achieved through t-SNE [Van Der Maaten, 2014] and labels are propagated from labeled to unlabeled samples in 2D by means of a teacher graph-based classifier called OPFSemi [Amorim et al., 2016]. In [Benato et al., 2021b], a threshold selection criterion is introduced to select the most confident pseudo-labeled samples, which are then utilized to train the student CNN model via cross-entropy loss. The backbone of the student CNN model comprises the VGG-16 encoder with pre-trained weights on the ImageNet dataset [Russakovsky et al., 2015].

While traditional pre-trained architectures can accelerate learning and reduce reliance on labeled data in annotation-scarce scenarios, they may not be viable in application domains that require more lightweight and specialized architectures. Moreover, they may introduce domain bias, inherit biases from the source dataset, and require careful fine-tuning, which is critical in certain applications [Steed and Caliskan, 2021].

The conf-DeepFA method has shown to achieve high classification accuracy on several datasets from a reduced set of labeled samples ($\sim 1\%$ of the dataset) and without the need for a validation set, further reducing the annotation effort. Nevertheless, its implementation using non-pre-trained custom CNN architectures remains an open problem since the required number of labeled samples to provide satisfactory results is rather high ($\sim 5\%$ of the dataset) [Aparco-Cardenas et al., 2024, 2025]. In addition, the occurrence of confirmation bias is a potential drawback for this method, where pseudo-labeling errors may have an adverse effect on the model's generalizability [Arazo et al., 2020].

More recently, Wang et al. [2021] presented a contrastive learning and pseudo-labeling-based iterative method to classify remote sensing images in semi-supervised scenarios where labeled data are scarce and unlabeled data are abundant. The method implements two independent yet complementary CNNs that are jointly trained in a co-training fashion. The intuition behind this setup is to create two views of the training data to improve the networks' ability to learn effective feature representations. It starts by independently training the encoder section of the networks by means of the minimization of a weighted contrastive loss. Next, a classification head is added to each network for subsequent image classification. In each iteration, the networks are trained in a cross-training manner with each network playing both the roles of teacher, by producing pseudo-labels, and student, by learning from the other network's pseudo-labels. The pseudo-labels with the highest 10% classification score are chosen to train the networks in the following iteration. An iterative categorical cross-entropy (ICE) loss function is implemented to reduce the impact of incorrect pseudo-labels during early iterations. Ultimately, an ensemble comprising the models trained across multiple iterations is used for inference on new data. This approach effectively combines contrastive learning with pseudo-labeling to achieve robust feature representation and to efficiently use unlabeled data, making it well-suited for scenarios with scarce labeled data, such as remote sensing imagery.

The work by Aparco-Cardenas et al. [2024], hereafter called *contrastive Deep Feature Annotation* (cont-DeepFA),

attempts to integrate contrastive learning and deep feature annotation in an iterative co-training fashion for the semi-supervised classification of biological images. It enables the use of non-pre-trained custom CNN architectures in the DeepFA framework by capitalizing on contrastive learning to improve the representation learning capability of two collaborative networks. Subsequently, during each iteration, the networks are trained in a teacher-student cross-training setting, where the teacher OPFSemi performs label propagation from labeled to unlabeled samples on the non-linear 2D projection of each student network's deep features; thereafter, the pseudo-labels with the highest 10% labeling confidence score, per class, are chosen to fine-tune the other network in a cross-training way. This has proven to effectively lessen confirmation bias and overfitting, while increasing the generalizability of the networks as iterations proceed. It was evaluated on three challenging unbalanced biological image datasets with only 5% of labeled samples.

Despite its efficacy, cont-DeepFA presents several drawbacks. It requires a larger proportion of labeled samples (5% vs 1% of the dataset) compared to approaches based on pre-trained CNN architectures. In addition, cont-DeepFA is storage-inefficient since it demands storing two models per iteration, resulting in $2T$ stored models after T iterations for training the final model. Moreover, its performance is highly dependent on the initial randomly selected labeled set, as this guides both supervised contrastive initialization and iterative pseudo-labeling. Finally, the absence of a mechanism for selecting representative samples limits its capacity to improve classification performance and generalization.

Active learning is a technique designed to identify the most informative samples from unlabeled data and present them to the oracle (e.g., a human annotator) for labeling, with the goal of minimizing labeling effort while preserving performance [Ren et al., 2021].

In this paper, we build on the work by Aparco-Cardenas et al. [2024] to introduce *active Deep Feature Annotation* (active-DeepFA), an active contrastive-based meta-pseudo-labeling method that follows the teacher-student paradigm. It adopts a co-training strategy and capitalizes on active learning to train non-pre-trained custom CNN under conditions of short supply of labeled and abundance of unlabeled data. It implements two cooperative CNNs enhancing pseudo-labeling accuracy and mitigating confirmation bias. Moreover, it exploits active learning through the periodic labeling of the most uncertain pseudo-labeled samples during training, which in turn improves the networks' generalizability and reduces the number of required labeled samples to achieve suitable results.

The active-DeepFA method starts with a reduced set of labeled samples ($\sim 1\%$ of the dataset) and works by initially pre-training the networks (students) through the minimization of a weighted supervised contrastive loss for a few epochs. Afterward, a loss equal to the sum of three losses: *i*) a weighted supervised contrastive loss, *ii*) a supervised loss, and *iii*) a semi-supervised loss, is minimized throughout the training process for each network. The set of labeled samples is used to minimize both (i) and (ii), while a set of reliable pseudo-labels is used to minimize (iii). Label propagation is performed in intervals of a few epochs through deep feature annotation,

so that the deep features of each network are projected onto 2D via t-SNE and OPFSemi (teacher) propagates labels from labeled to unlabeled samples. Next, the pseudo-labels with the highest 10% labeling confidence score, per class, are selected to fine-tune the other network, while a small number of samples, whose pseudo-labels have the lowest labeling confidence score, are chosen to be labeled by an oracle, and then added to the set of labeled samples to minimize (i) and (ii) during subsequent epochs. The aforesaid active learning procedure halts when the number of labeled samples reaches a preset value. Lastly, an ensemble of the two trained networks is used for inference on unseen data.

This method effectively combines contrastive learning, meta-pseudo-labeling and active learning to efficiently exploit unlabeled data, producing well-generalized models under conditions of scarcity of labeled and abundance of unlabeled data. Moreover, labeling the samples, identified as the most uncertain to classify, at regular epoch intervals results in improved generalization performance, since the model is continuously fed with meaningful samples during the training phase. The latter is in stark contrast to its counterparts, which randomly select the set of samples to be labeled and keep it fixed throughout the learning process.

Unlike the works by Wang *et al.* [2021] and Aparco-Cardenas *et al.* [2024], which use supervised contrastive learning only at the beginning of each iteration as a sort of pre-training, we incorporate it into the training loss during the whole training procedure to leverage the samples labeled by active learning. Moreover, the training phase is reduced to a single iteration, ending up with only two networks, thus optimizing storage efficiency.

Our approach is assessed on seven challenging image datasets from multiple domains and benchmarked against a number of well-established DSSL methods to prove its effectiveness and efficiency. The datasets are unbalanced and present a high degree of resemblance among classes, making classification quite difficult.

The main contributions of this work can be summarized as:

1. We present a novel active contrastive-based meta-pseudo-labeling approach that advances previous work to train non-pre-trained custom CNN architectures for image classification in conditions of short supply of labeled and abundance of unlabeled data.
2. We integrate active learning with deep feature annotation within a co-training setting to efficiently harness the unlabeled data, which in turn improves the networks' generalizability and reduces the annotation effort by labeling the most difficult-to-classify samples.

The remainder of this paper is organized as follows. Section 2 briefly reviews the related work. Next, Section 3 presents the proposed methodology in detail. Further, Section 4 describes the experiments conducted and discusses the results. Lastly, Section 5 concludes the paper and provides future and ongoing work directions.

2 Related Work

Based on different choices of architectures and unsupervised loss functions or regularization terms, DSSL methods can be divided into several categories [Yang *et al.*, 2022], among which pseudo-labeling, consistency regularization and hybrid methods stand out due to their straightforwardness and effectiveness. Also, deep semi-supervised methods in the literature capitalize on active learning to improve image classification while alleviating annotation burden. These methods are described in the following sub-sections.

2.1 Pseudo-labeling methods

Pseudo-labeling methods assign pseudo-labels to unlabeled data using only a reduced set of labeled data. In general, these can be categorized into self-training [Lee *et al.*, 2013] and disagreement-based models [Zhou, 2008].

Self-training approaches use the confidence of the model's predictions to infer pseudo-labels for unlabeled data. However, recent works have shown that the model's performance can be further enhanced by using an auxiliary model (*e.g.*, graph-based methods) in a teacher-student setup [Pham *et al.*, 2021; Benato *et al.*, 2021a,b, 2023; Arazo *et al.*, 2020]. It is noteworthy that a number of issues may arise from naive approaches, such as confirmation bias [Arazo *et al.*, 2020], imbalance bias [Wang *et al.*, 2022a], or concept drift [Cascante-Bonilla *et al.*, 2021], which adversely impact the generalizability of the model. Some approaches to mitigate the above issues are: to consider the pseudo-labeling confidence during the regularization process [Arazo *et al.*, 2020], to reset the parameters of the model at the start of each learning iteration [Cascante-Bonilla *et al.*, 2021], and to include the selection of the most confident pseudo-labeled samples to ensure labeling correctness [Benato *et al.*, 2021b, 2023].

A well-known example in this category is *Pseudo-label* [Lee *et al.*, 2013], a self-training method that fine-tunes a pre-trained network simultaneously with both labeled and unlabeled data. The network is initially pre-trained on the labeled data in a usual supervised manner, while pseudo-labels for the unlabeled data are recomputed after each batch update and subsequently used to regularize the model.

2.2 Consistency regularization methods

Consistency regularization methods build on the manifold hypothesis (data points situated on the same low-dimensional manifold should share the same label) or smoothness hypothesis (if two samples are close within the input space, their corresponding labels should match) by integrating consistency constraints into the loss function. These constraints can be adopted under different modes of perturbation (*e.g.*, input perturbation, weights perturbation or layer perturbation) leading to a wide spectrum of methods [Yang *et al.*, 2022]. Well-known representatives in this category are Π -Model, Mean Teacher, VAT, and UDA, which are briefly described below.

Π -Model [Laine and Aila, 2017] forms a consensus prediction from the ensemble of predictions of instances of a single network across different training epochs, and under

different regularization and input augmentation conditions. *Mean Teacher* [Tarvainen and Valpola, 2017] works in a teacher-student setup. The teacher maintains an exponential moving average (EMA) of the student’s weights over training steps, while a consistency cost minimizes the distance of the predictions of both networks. *Virtual Adversarial Training* (VAT) [Miyato et al., 2018] uses a consistency constraint that ensures similarity between the model’s outputs of the original input sample and its corresponding adversarially transformed image. *Unsupervised Data Augmentation* (UDA) [Xie et al., 2020] examines the role of noise injection in consistency training by replacing simple noising operations with advanced data augmentation methods.

2.3 Hybrid methods

Hybrid methods capitalize on both consistency regularization and pseudo-labeling to create more accurate and effective models in label-scarce settings. A well-known representative is *FixMatch* [Sohn et al., 2020], which assigns pseudo-labels in a self-training manner to both weakly and strongly augmented views of an input sample, while a consistency constraint ensures their similarity.

A more recent method in this category is *SimMatch* [Zheng et al., 2022], which extends the hybrid paradigm by jointly considering semantic and instance similarity. Using a labeled memory buffer, it propagates information across both the semantic and instance levels, producing more reliable pseudo-labels and mitigating overconfidence.

2.4 Active learning methods

SSL methods leverage active learning to label the most informative and representative samples, improving model performance while minimizing labeling effort. The identification of these samples is carried out through various approaches [Ren et al., 2021]. Some works in this category are briefly presented as follows.

Gao et al. [2020] propose a consistency-based semi-supervised active learning framework that works in two stages. Firstly, labeled and unlabeled data are used to train the model with cross-entropy loss and consistency-based loss, respectively. Secondly, the model outputs of unlabeled data and their augmentations are measured by a consistency-based metric. Next, unlabeled samples with low-consistency scores are selected for labeling and added to the labeled pool. Yuan et al. [2022] propose ActiveMatch, an end-to-end SSL method that effectively combines unsupervised and supervised contrastive learning with self-training pseudo-labeling to actively select the most informative samples for labeling during the training cycle, leading to better representations for image classification. The unlabeled samples with the lowest classification score are submitted to the oracle for labeling and then included in the labeled set. Zhang et al. [2014] introduce an SSL algorithm that combines co-training and active learning to improve classification performance while spending the same human annotation effort. It exploits co-training to select the most reliable samples according to the criteria of high confidence and nearest neighbor to improve the classifier, while

the most informative samples are selected and submitted to a human annotator for labeling.

3 Methodology

In this section, we thoroughly present our method, outlining its constituent components and describing its execution steps.

3.1 Notations and Definitions

In a Semi-Supervised Learning (SSL) setting, a training set $\mathcal{Z} = \{\mathbf{x} | \mathbf{x} \in X\}$ comprises two disjoint subsets: a labeled set $\mathcal{Z}_L = \{(\mathbf{x}, y) | \mathbf{x} \in X, y \in Y\}$ and an unlabeled set $\mathcal{Z}_U = \{\mathbf{x} | \mathbf{x} \in X\}$, where $\mathcal{Z} = \mathcal{Z}_L \cup \mathcal{Z}_U$, $\mathcal{Z}_L \cap \mathcal{Z}_U = \emptyset$, and $|\mathcal{Z}_L| \ll |\mathcal{Z}_U|$. In this notation, \mathbf{x} represents a sample, and y denotes its corresponding class label. We proceed under the assumption that the samples in both \mathcal{Z}_L and \mathcal{Z}_U are drawn from the same data distribution. Let us denote $B_{\mathcal{X}}$ as a set of batches for a given set \mathcal{X} .

Let N_1 and N_2 be two independent and collaborative networks that share the same architecture and structure during the training phase. Each network comprises a backbone encoder E and two projection heads P_c , to perform contrastive learning, and P_x , for classification tasks. The output latent space of E is used for 2D projection and label propagation.

3.2 Network Architecture

The architecture of both N_1 and N_2 comprises an *encoder E* consisting of four convolutional layers of 64, 128, 256, and 512 filters each, followed by a dense (fully-connected) layer of 4096 neurons. Each of the convolutional layers consists of a filter bank with kernel size 3×3 , batch normalization, ReLU activation, max-pooling with kernel size 3×3 . P_c consists of a dense layer of 1024 neurons, while P_x encompasses a dense layer with a number of neurons equal to the number of classes. Figure 1 shows the networks’ architecture.




Figure 1. The networks’ architecture consisting of an encoder E and two projection heads P_x and P_c .

This architecture was empirically selected based on previous related works [Aparco-Cardenas et al., 2024, 2025] and preliminary experiments that confirmed its suitability for the datasets considered in this manuscript. Since the focus of this

work is on the training strategy rather than architectural exploration, we adopted a straightforward CNN design without further architectural ablations.

3.3 Supervised Contrastive Learning on \mathcal{Z}_L

In an attempt to circumvent the encoder’s pre-training prerequisite for DeepFA-based methods, we adopt supervised contrastive learning as a warm-up step during early epochs. In addition to that, it is a component of the loss function during the whole training procedure, which enhances the generalization ability of the networks.

In this regard, contrastive learning acts as a dimensional-reduction technique by mapping, in a contrastive manner, a set of high-dimensional data points (e.g., images) to representations in a lower-dimensional space, such that representations of semantically similar pairs are drawn closer, whereas those of dissimilar pairs are pushed away from each other in the lower-dimensional space. During the learning phase, each network adopts a Siamese architecture comprising two weight-sharing subnetworks as shown in Figure 2. Let $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{Z}_L$ be a pair of labeled input images, and let $\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2$ be their respective augmented versions. Over the course of training, $\tilde{\mathbf{x}}_1$ and $\tilde{\mathbf{x}}_2$ serve as input to the different branches of the Siamese network. The distance function \mathcal{D} between the lower-dimensional representations of $\tilde{\mathbf{x}}_1$ and $\tilde{\mathbf{x}}_2$ is defined as the Euclidean distance:

$$\mathcal{D}(\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2) = \|P_c(E(\tilde{\mathbf{x}}_1)) - P_c(E(\tilde{\mathbf{x}}_2))\| \quad (1)$$

Let y_t be a binary label assigned to the pair $\mathbf{x}_1, \mathbf{x}_2$, where $y_t = 0$ if they are deemed similar and $y_t = 1$ otherwise. In our setting, similarity labels are derived directly from the class labels in \mathcal{Z}_L : two samples are regarded as similar if they belong to the same class, and dissimilar otherwise. No additional pairwise annotation is required. To compensate for the large difference between the number of similar and dissimilar pairs, a weight factor τ is incorporated into the contrastive loss function. The weighted supervised contrastive loss function is defined as

$$\mathcal{L}_S^{cl} = \sum_{i=1}^{|S|} L(p^{(i)}) \quad (2)$$

$$L(p^{(i)}) = \tau(1 - y_t^{(i)}) \frac{1}{2} (\mathcal{D}^{(i)})^2 + y_t^{(i)} \frac{1}{2} \{\max(0, m - \mathcal{D}^{(i)})\}^2 \quad (3)$$

where $p^{(i)} = (y_t, \tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2)^{(i)}$ is the i -th labeled sample pair and S is the set of possible sample pairs for a given batch $B^l \in B_{\mathcal{Z}_L}$, such that $|S| = \binom{|B^l|}{2}$. The margin value is set to $m = 2$, consistent with previous related work [Aparco-Cardenas et al., 2024].

Let K be the number of classes of the dataset and let p be the size of each class for a given balanced dataset. The ratio between the number of similar and dissimilar pairs can be computed as

$$\frac{K \binom{p}{2}}{\binom{Kp}{2} - K \binom{p}{2}} = \frac{p - 1}{p(K - 1)} \approx \frac{1}{K - 1}. \quad (4)$$




Figure 2. The weight-sharing Siamese architecture adopted by each network during contrastive learning.

Therefore, the τ value is set to $K - 1$ to compensate for the contribution of similar pairs in \mathcal{L}_S^{cl} as suggested in Wang et al. [2021].

Unlike other contrastive learning approaches that employ cosine similarity as the distance metric, we adopt the Euclidean distance to remain consistent with the prior work on which our method builds [Aparco-Cardenas et al., 2024] and to ensure methodological coherence with OPFSemi, which also computes edge weights in the graph using the Euclidean distance.

3.4 Supervised Learning on \mathcal{Z}_L

For a given batch $B^l \in B_{\mathcal{Z}_L}$, the supervised loss function is defined as

$$\mathcal{L}_S = \frac{1}{|B^l|} \sum_{i=1}^{|B^l|} CE(y_i, \mathbf{q}_i) \quad (5)$$

where y_i and \mathbf{q}_i are the label and the logit prediction for the i -th sample, and $CE(\cdot, \cdot)$ represents the cross-entropy loss.

3.5 Semi-Supervised Learning on \mathcal{Z}_U

The pseudo-labeling procedure is carried out via label propagation by means of OPFSemi [Amorim et al., 2016], a parameter-free semi-supervised graph-based technique inspired by the Optimum-Path Forest (OPF) methodology, which has been adopted in other recent works for the same purpose [Benato et al., 2021a,b, 2023; Aparco-Cardenas et al., 2024, 2025]. OPFSemi was selected because it efficiently exploits both labeled and unlabeled data without requiring parameter tuning, while also providing robustness in modeling complex decision boundaries, which makes it well-suited for scenarios with limited labeled samples.

Let $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^2$ be a function that non-linearly maps the latent space output of the encoder E into 2D. For each network, $\phi(\mathcal{Z})$ is computed by forward passing \mathcal{Z} through E and projecting the output deep features into 2D via t-SNE. The OPFSemi technique is executed on the graph induced by the samples in the projection set $\phi(\mathcal{Z})$. It starts by modeling $\phi(\mathcal{Z})$ as a complete graph, where samples are encoded as nodes and edges are weighted by the Euclidean distance between their endpoints. The projection subset of labeled samples $\phi(\mathcal{Z}_L)$ acts as the set of prototypes and their labels are propagated to their most closely connected unlabeled samples in $\phi(\mathcal{Z}_U)$,

resulting in a graph partitioned into optimum-path trees rooted at $\phi(\mathcal{Z}_L)$.

Let $\lambda'(u)$ be the pseudo-label assigned to $u \in \phi(\mathcal{Z}_U)$ via OPFSemi's label propagation, and let $\lambda(u)$ be the ground-truth label of $s \in \phi(\mathcal{Z}_L)$, such that s is the root of the optimum-path tree to which u is connected – i.e., $\lambda'(u) = \lambda(s)$. Let $c(u)$ be the cost of the optimum-path offered from s to u , and let c' be the cost of the second smallest optimum-path offered to u from $t \in \phi(\mathcal{Z}_L)$, where $s \neq t$ and $c' > c(u)$. At the end of the execution, a labeling confidence value defined as

$$v(u) = \frac{c'}{c(u) + c'} \in [0, 1] \quad (6)$$

is calculated and assigned to each $u \in \phi(\mathcal{Z}_L)$.

Afterward, we pick the pseudo-labels with the highest 10% labeling confidence value, per class. This procedure guarantees the selection of reliable pseudo-labels that encompass all classes in the dataset. In contrast, using a fixed labeling confidence threshold may fail to select pseudo-labels from classes with a low mean labeling confidence value. The 10% cutoff was adopted in accordance to previous works [Wang et al., 2021; Aparco-Cardenas et al., 2024].

Let Y' be the set of pseudo-labels with the highest 10% labeling confidence value, per class. We define the set $\mathcal{Z}_U^{pl} = \{(\mathbf{x}, y') \mid \mathbf{x} \in \mathcal{Z}_U, y' \in Y'\}$, where y' is the pseudo-label assigned to sample \mathbf{x} . For a given batch $B^{pl} \in B_{\mathcal{Z}_U^{pl}}$, the semi-supervised loss function is defined as

$$\mathcal{L}_U^{ssl} = \frac{1}{|B^{pl}|} \sum_{i=1}^{|B^{pl}|} CE(y'_i, \mathbf{q}_i) \quad (7)$$

where y'_i and \mathbf{q}_i are the pseudo-label and the logit prediction for the i -th sample. During training, we apply strong augmentations to the samples of \mathcal{Z}_U^{pl} , as defined in Sohn et al. [2020], to enhance the learning performance and the generalizability of the network.

3.6 Active Learning and Cross-Training

Once label propagation is concluded, we identify the most uncertain samples as the ones whose pseudo-labels have the lowest labeling confidence value, which are the ones located at the decision boundary of the optimum-path trees. These samples are regarded as the most difficult to classify and thus assumed to contain meaningful information to improve generalization performance. Let $\mathcal{Z}_A^{(i)}$ be the set of samples chosen by active learning for N_i , $i \in \{1, 2\}$. The set $\mathcal{Z}_A^{(i)}$ comprises the k_{active} samples whose pseudo-labels have the lowest labeling confidence values, for a fixed value of k_{active} . These sets are then joined and the k_{active} samples with the lowest labeling confidence values are sent to the oracle to be labeled and included in the labeled set \mathcal{Z}_L to minimize both \mathcal{L}_S^{cl} and \mathcal{L}_S during the following epochs. This procedure is halted when a preset maximum number of samples labeled by active learning n_{active} is achieved.

The conventional approach for the pseudo-labeling training stage is to re-train the network with the pseudo-labels generated by itself. Nevertheless, this may inadvertently provoke

overfitting and confirmation bias, thus hampering network learning. In this respect, the idea of using the pseudo-labels generated by another model stems from the notion that the generalization capability of the network is enhanced when exposed to new information. Let $\mathcal{Z}_U^{pl(i)}$ be the set of picked pseudo-labeled samples for N_i , $i \in \{1, 2\}$. Therefore, after each pseudo-labeling stage, the set $\mathcal{Z}_U^{pl(2)}$ is used to minimize \mathcal{L}_U^{ssl} for network N_1 and vice versa for N_2 , performing this in a cross-training manner. Ultimately, the objective loss function that each network minimizes during the training phase is defined as

$$\mathcal{L} = \mathcal{L}_S^{cl} + \mathcal{L}_S + \mathcal{L}_U^{ssl}. \quad (8)$$

After the training process has been completed, we use both N_1 and N_2 in an ensemble fashion to predict new data. Let $\vec{p}^i(u)$ be the predicted probability vector of N_i for sample \mathbf{x} . The label of \mathbf{x} , denoted as $\lambda(\mathbf{x})$, is computed as the index of the highest value of the average of the predicted probability vectors of N_1 and N_2 in the following way

$$\lambda(\mathbf{x}) = \operatorname{argmax} \left\{ \frac{\vec{p}^1(\mathbf{x}) + \vec{p}^2(\mathbf{x})}{2} \right\}. \quad (9)$$

3.7 Overview of the Proposed Method

An overview of the procedural steps of active-DeepFA is outlined in Algorithm 1. The method starts with two networks N_1, N_2 (students) sharing the same architecture and structure. Since we desire that each network produce a unique representation for the same input image, we begin by independently initializing the weights of each network using Kaiming He initialization [He et al., 2015] (Line 1). We also keep a counter c_{active} of the number of samples labeled by active learning, which is initially set to zero (Line 2). The training process is executed for a preset number of epochs, n_{epochs} . As a form of pre-training, we perform contrastive learning by minimizing \mathcal{L}_S^{cl} on \mathcal{Z}_L for a fixed number of warm-up epochs, w_{epochs} , where $w_{\text{epochs}} < n_{\text{epochs}}$ (Lines 4-9).

Afterward, for the epoch immediately after the warm-up step – i.e., at epoch $w_{\text{epochs}} + 1$ – and at intervals of e_{int} epochs – i.e., at every epoch divisible by e_{int} , we project the output deep features of $E(\mathcal{Z})$ onto 2D using t-SNE and execute label propagation from $\phi(\mathcal{Z}_L)$ to $\phi(\mathcal{Z}_U)$ on the 2D space via OPFSemi (teacher). Next, the sets \mathcal{Z}_U^{pl} and \mathcal{Z}_A are constructed exploiting the labeling confidence values assigned to the pseudo-labels (Lines 11-15). Let $\mathcal{Z}_U^{pl(1)}, \mathcal{Z}_A^{(1)}$ and $\mathcal{Z}_U^{pl(2)}, \mathcal{Z}_A^{(2)}$ be the sets obtained from N_1 and N_2 , respectively. If c_{active} is less than a fixed maximum number n_{active} , we proceed to construct the set \mathcal{Z}'_A by joining the sets $\mathcal{Z}_A^{(1)}, \mathcal{Z}_A^{(2)}$ and selecting the k_{active} samples with the lowest labeling confidence value. The samples of \mathcal{Z}'_A are submitted to the oracle to be labeled and subsequently added to \mathcal{Z}_L . Next, the value of c_{active} is updated accordingly (Lines 16-20).

Then, for each network, we compute the losses $\mathcal{L}_S^{cl}, \mathcal{L}_S$ on \mathcal{Z}_L , and \mathcal{L}_U^{ssl} on the strong augmented samples of $\mathcal{Z}_U^{pl(2)}$ for N_1 and vice versa for N_2 . The losses are subsequently added and the network is updated accordingly to minimize this sum

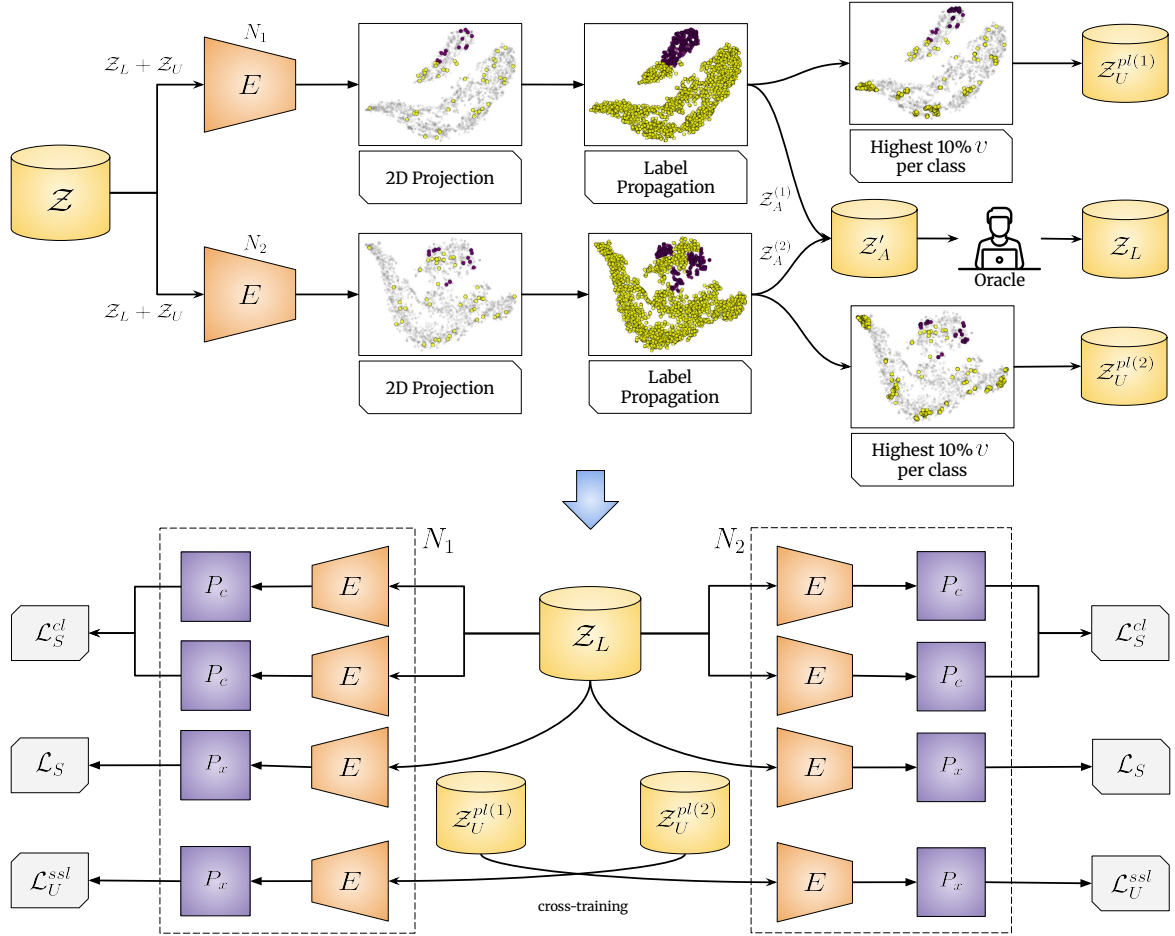


Figure 3. Overview of the active-DeepFA method. For each network, the output deep features of $E(\mathcal{Z})$ are projected into 2D using t-SNE and labeling propagation is executed by OPFSemi on the 2D space. Next, the sets $\mathcal{Z}_U^{pl(i)}$, $\mathcal{Z}_A^{(i)}$ are constructed for N_i , $i \in \{1, 2\}$. The set \mathcal{Z}'_A is created by joining the sets $\mathcal{Z}_A^{(1)}$, $\mathcal{Z}_A^{(2)}$ and picking the k_{active} samples whose pseudo-labels have the lowest v value. Then, the elements of \mathcal{Z}'_A are sent to the oracle to be labeled and added to \mathcal{Z}_L . Afterward, the losses \mathcal{L}_S^{cl} , \mathcal{L}_S are computed on \mathcal{Z}_L , while \mathcal{L}_U^{ssl} is calculated on $\mathcal{Z}_U^{pl(2)}$ for N_1 and vice versa for N_2 in a cross-training manner.

(Lines 21-26). Lastly, the trained networks N_1 , N_2 are used in an ensemble manner to predict on new data.

4 Experiments and Results

This section evaluates the performance of active-DeepFA on seven datasets from multiple domains, compares it to baselines and established DSSL methods, and provides an analysis and discussion of the experimental results.

4.1 Datasets

We evaluate the active-DeepFA method on seven challenging datasets spanning three distinct domains: biological imaging, remote sensing and medical imaging. The first three datasets are from a parasite biological image collection [Suzuki *et al.*, 2013]. The collection consists of optical microscopy images of the most common human intestinal parasites in Brazil. These parasites are widespread in countries with tropical, subtropical, and equatorial climates, posing substantial public health challenges. The images are rescaled to $200 \times 200 \times 3$ pixels. The datasets comprising this collection are listed as


follows: *i) Helminth larvae*; *ii) Helminth eggs*; and *iii) Protozoan cysts*. These datasets are characterized by their class imbalance and the presence of similar features among distinct classes, which increases the difficulty of network learning. Table 1 shows a brief description of the datasets, while some images of parasites and impurities for (i)–(iii) are displayed in Figure 4. The fourth dataset is the 10-class Moving and Stationary Target Acquisition and Recognition (MSTAR) [Song *et al.*, 2016], a standard Synthetic Aperture Radar (SAR) benchmark in remote sensing. The images are grayscale images of 128×128 pixels. Table 2 summarizes the categories and class sizes, while Figure 5 provides one image example per class. Lastly, the final three datasets – PneumoniaMNIST, PathMNIST and TissueMNIST – are derived from the MedMNIST medical image collection [Yang *et al.*, 2023]. PneumoniaMNIST consists of grayscale chest X-ray images of size 128×128 , annotated for binary classification between pneumonia and normal cases. PathMNIST is constructed from colon pathology slides, providing RGB image patches of size $128 \times 128 \times 3$ categorized into nine distinct tissue types. TissueMNIST, in turn, contains grayscale images of size 128×128 representing human kidney cortex cells, with labels corresponding to eight different cell types. Table 2

Algorithm 1 active-DeepFA algorithm

Input: $\mathcal{Z}, N_1, N_2, n_{\text{epochs}}, w_{\text{epochs}}, e_{\text{int}}, n_{\text{active}}, k_{\text{active}}$
Output: Trained networks N_1, N_2

- 1: Initialize N_1, N_2 with Kaiming He initialization
- 2: $c_{\text{active}} \leftarrow 0$
- 3: **for** $e \leftarrow 1$ to n_{epochs} **do**
- 4: **if** $e \leq w_{\text{epochs}}$ **then**
- 5: **for** $B^l \in B_{\mathcal{Z}_L}$ **do**
- 6: **for** $N \in \{N_1, N_2\}$ **do**
- 7: Compute \mathcal{L}_S^{cl} using B^l
- 8: $\mathcal{L} \leftarrow \mathcal{L}_S^{cl}$
- 9: Update N to minimize \mathcal{L}
- 10: **else**
- 11: **if** $e \% e_{\text{int}} = 0$ **or** $e = w_{\text{epochs}} + 1$ **then**
- 12: **for** $N \in \{N_1, N_2\}$ **do**
- 13: Project $E(\mathcal{Z})$ onto 2D
- 14: Propagate labels from $\phi(\mathcal{Z}_L)$ to $\phi(\mathcal{Z}_U)$
- 15: Construct the sets \mathcal{Z}_U^{pl} and \mathcal{Z}_A
- 16: **if** $c_{\text{active}} < n_{\text{active}}$ **then**
- 17: Construct \mathcal{Z}'_A from $\mathcal{Z}_A^{(1)}$ and $\mathcal{Z}_A^{(2)}$
- 18: Send \mathcal{Z}'_A to the oracle for labeling
- 19: $\mathcal{Z}_L \leftarrow \mathcal{Z}_L \cup \mathcal{Z}'_A$
- 20: $c_{\text{active}} \leftarrow c_{\text{active}} + k_{\text{active}}$
- 21: **for** $B^l \in B_{\mathcal{Z}_L}, B^{pl} \in B_{\mathcal{Z}_U^{pl}}$ **do**
- 22: **for** $N \in \{N_1, N_2\}$ **do**
- 23: Compute $\mathcal{L}_S^{cl}, \mathcal{L}_S$ using B^l
- 24: Compute \mathcal{L}_U^{ssl} using B^{pl} ▷ cross-training
- 25: $\mathcal{L} \leftarrow \mathcal{L}_S^{cl} + \mathcal{L}_S + \mathcal{L}_U^{ssl}$
- 26: Update N to minimize \mathcal{L}
- 27: **return** N_1, N_2

summarize the datasets' characteristics, while Figures 6, 7, and 8 illustrate representative examples from each dataset.



(a) Helminth larvae (b) Helminth eggs (c) Protozoan cysts

Figure 4. Images from the Helminth larvae, Helminth eggs and Protozoan cysts datasets. The first row shows images of parasites, while the second row exhibits images of impurities.





Figure 5. Images from the MSTAR dataset.



(a) Normal (b) Pneumonia
Figure 6. Images from the PneumoniaMNIST dataset.




Figure 7. Images from the PathMNIST dataset.




Figure 8. Images from the TissueMNIST dataset.

4.2 Dataset Preparation

We split the first three datasets \mathcal{Z}_D into $\frac{2}{3}$ for the training set \mathcal{Z} and the remaining $\frac{1}{3}$ for the test set \mathcal{Z}_T ($\mathcal{Z}_D = \mathcal{Z} \cup \mathcal{Z}_T$) by means of three-fold cross-validation – creating 3 train/test splits. A validation set is omitted, as its inclusion would contradict the purpose of this work by increasing the annotation effort. We emulate the conditions of short supply of labeled data by partitioning the training set for each split into labeled \mathcal{Z}_L and unlabeled \mathcal{Z}_U sets ($\mathcal{Z} = \mathcal{Z}_L \cup \mathcal{Z}_U$), such that $|\mathcal{Z}_L| = 1\% \times |\mathcal{Z}_D|$ and $|\mathcal{Z}_U| = 65.66\% \times |\mathcal{Z}_D|$. The 1% of labeled samples are randomly picked from \mathcal{Z} for each split in a stratified manner. For the remaining four datasets – MSTAR, PneumoniaMNIST, PathMNIST, and TissueMNIST – we retain the original train/test split and perform three independent random samplings of 1% of the dataset for labeling. Tables 3 and 4 show the number of labeled and unlabeled samples for each dataset.

Table 1. Description of the intestinal parasites datasets.

Dataset	Category	Size	Class ID
Helminth larvae	<i>Strongyloides stercoralis</i>	446	1
	Impurities	3068	2
	Total	3514	
Helminth eggs	<i>Hymenolepis nana</i>	348	1
	<i>Hymenolepis diminuta</i>	80	2
	<i>Ancylostomatidae</i>	148	3
	<i>Enterobius vermicularis</i>	122	4
	<i>Ascaris lumbricoides</i>	337	5
	<i>Trichuris trichiura</i>	375	6
	<i>Schistosoma mansoni</i>	122	7
	<i>Taenia spp.</i>	236	8
	Impurities	3344	9
Total	5112		
Protozoan cysts	<i>Entamoeba coli</i>	719	1
	<i>Entamoeba histolytica / E. dispar</i>	78	2
	<i>Endolimax nana</i>	724	3
	<i>Giardia duodenalis</i>	641	4
	<i>Iodamoeba bütschlii</i>	1501	5
	<i>Blastocystis hominis</i>	189	6
	Impurities	5716	7
Total	9568		

Table 2. Description of the MSTAR and MedMNIST datasets.

Dataset	Category	Train	Test	Class ID
MSTAR	2S1	299	274	1
	BMP2	233	195	2
	BRDM2	298	274	3
	BTR60	256	195	4
	BTR70	233	196	5
	D7	299	274	6
	T62	299	273	7
	T72	232	196	8
	ZIL131	299	274	9
	ZSU234	299	274	10
	Total	2757	2425	
PneumoniaMNIST	Normal	1214	234	1
	Pneumonia	3494	390	2
	Total	4708	624	
PathMNIST	Adipose	9366	1338	1
	Background	9509	847	2
	Debris	10360	339	3
	Lymphocytes	10401	634	4
	Mucus	8006	1035	5
	Smooth muscle	12182	592	6
	Normal colon mucosa	7886	741	7
	Cancer-associated stroma	9401	421	8
	Colorectal AdCa epith.	12885	1233	9
Total	89996	8180		
TissueMNIST	Collecting duct/tubule	53075	15942	1
	Distal convoluted tubule	7814	3010	2
	Glomerular endothel. cells	5866	2454	3
	Interstitial endothel. cells	15406	5179	4
	Leukocytes	11789	4146	5
	Podocytes	8482	2979	6
	Proximal tubule segments	39980	11978	7
	Thick ascending limb	25385	7808	8
Total	167797	53196		

Table 3. Number of labeled and unlabeled samples for the parasites and MSTAR datasets.

	Helminth larvae	Helminth eggs	Protozoan cysts	MSTAR
$ \mathcal{Z}_L $	36	56	100	56
$ \mathcal{Z}_U $	2308	3356	6281	2691

4.3 Experimental Setup

All implementations were carried out in PyTorch 1.8.0. For t-SNE, we used the implementation from the Multicore t-SNE library [Ulyanov, 2016] with a perplexity value of 50. The

Table 4. Number of labeled and unlabeled samples for the MedMNIST datasets.

	PneumoniaMNIST	PathMNIST	TissueMNIST
$ \mathcal{Z}_L $	54	977	2133
$ \mathcal{Z}_U $	4654	89019	163333

OPFSemi algorithm is parameter-free, so no tuning was required. As for the optimization hyperparameters, we used stochastic gradient descent (SGD) as optimizer with batch size of 32, learning rate of 3×10^{-4} , momentum value of 0.9, weight decay value of 5×10^{-4} and Nesterov momentum. The learning rate adopts a cosine learning rate decay [Loshchilov and Hutter, 2017], whose value at step t is $lr = lr_0 \times \cos(\frac{7\pi t}{16T})$, where lr_0 is the initial learning rate value and T is the total number of training steps. The values of t-SNE perplexity and network optimization parameters follow prior works on the same first three datasets [Aparco-Cardenas et al., 2024, 2025].

For the sake of comparison, we use the results from Aparco-Cardenas et al. [2024] for the first three datasets, since the same three-fold cross-validation splits are employed and the networks have an almost identical architecture. The only difference lies on the projection head, where the architecture presented in Aparco-Cardenas et al. [2024] has a single projection head consisting of two hidden dense (fully-connected) layers of 512 and 256 neurons each. The architecture used for our method is presented in Section 3.2.

The active-DeepFA method starts with $|\mathcal{Z}_L| = 1\% \times |\mathcal{Z}_D|$. We set the number of contrastive warm-up epochs equal to $w_{\text{epochs}} = 15$ to ensure meaningful representations for both subsequent pseudo-labeling and active learning tasks. To ensure frequent pseudo-labeling cycles and prioritize the active labeling of the most uncertain samples, we empirically set $e_{\text{int}} = 5$ for all datasets, while $k_{\text{active}} = 2$ for the three biological image datasets, $k_{\text{active}} = 5$ for MSTAR, $k_{\text{active}} = 10$ for PneumoniaMNIST, $k_{\text{active}} = 162$ for PathMNIST, and $k_{\text{active}} = 355$ for TissueMNIST, aiming to enhance generalizability while maintaining efficiency.

Given the comparatively larger size of the PathMNIST and TissueMNIST datasets, 20% of the unlabeled data is randomly sampled during each pseudo-labeling cycle and projected alongside the labeled data, rather than using the entire unlabeled data. This strategy reduces computational overhead and accelerates processing while maintaining a representative sampling of the dataset.

The number of epochs n_{epochs} varies in function of the maximum number of labeled samples by active learning n_{active} . Since our aim is to make a comparison with the results in Aparco-Cardenas et al. [2024], we continue labeling samples via active learning until we reach $|\mathcal{Z}_L| = 5\% \times |\mathcal{Z}_D|$. In this regard, the values of n_{epochs} and n_{active} for each dataset are shown in Tables 5 and 6.

Table 5. Values of n_{epochs} and n_{active} for the parasites and MSTAR datasets.

	Helminth larvae	Helminth eggs	Protozoan cysts	MSTAR
n_{epochs}	375	535	983	230
n_{active}	140	200	379	259

It is worth noting that when applying the approach to other

Table 6. Values of n_{epochs} and n_{active} for the MedMNIST datasets.

	PneumoniaMNIST	PathMNIST	TissueMNIST
n_{epochs}	135	150	140
n_{active}	214	3886	8508

datasets, these hyperparameters should be adjusted to reflect differences in data distribution and class balance, ensuring effectiveness and efficiency across diverse scenarios.

We use accuracy and Cohen’s κ , hereafter denoted as κ , to evaluate the performance of the method. The κ metric, where $\kappa \in [-1, 1]$, provides a more reliable measure than accuracy for unbalanced datasets. It quantifies the degree of disagreement between the classifier’s prediction and the ground-truth, where $\kappa = 1$ and $\kappa = -1$ indicate full agreement and disagreement, respectively. For each method, the mean and standard deviation of both accuracy and κ across splits are presented.

4.4 Comparison with baselines

In this experiment, we intend to benchmark our method against its direct baselines, from which this work derives. The main baselines of our approach are the method by Wang *et al.* [2021], conf-DeepFA [Benato *et al.*, 2023] and cont-DeepFA [Aparco-Cardenas *et al.*, 2024]. It is worth noting that the active-DeepFA method and its baselines do not utilize a validation set. The experiment configuration and hyperparameter values for all baselines can be found in Aparco-Cardenas *et al.* [2024]. It is important to highlight that all the baselines start with $|\mathcal{Z}_L| = 5\% \times |\mathcal{Z}_D|$, while active-DeepFA starts with only $|\mathcal{Z}_L| = 1\% \times |\mathcal{Z}_D|$ and expands \mathcal{Z}_L through active learning to reach $|\mathcal{Z}_L| = 5\% \times |\mathcal{Z}_D|$, thus they all end up using the same number of labeled data.

The mean and standard deviation of both accuracy and κ across splits are shown in Table 7. It can be observed that active-DeepFA achieves better generalization on the test set for all seven datasets, showcasing the effectiveness of combining supervised contrastive learning, active learning and co-training.

4.4.1 Influence of supervised contrastive learning

The results suggest that supervised contrastive learning serves as an effective weight initializer for DeepFA-based methods. In this regard, we circumvent the requirement for pre-trained encoders while reducing annotation costs by learning from a reduced set of labeled samples, thus enabling the use of non-pre-trained custom CNN architectures for DeepFA-based methods as conf-DeepFA.

In stark contrast to the baselines that uniquely use supervised contrastive learning as a sort of pre-training step, we further exploit it by incorporating it into the objective loss function throughout the whole training process. This procedure is enhanced by the dynamic labeled set, which is periodically expanded by active learning at regular epoch intervals. These informative samples incrementally improve the low-dimensional feature representations learned by the network from the weighted supervised contrastive loss. Moreover, the implementation of an additional projection head allows the simultaneous minimization of the three losses – super-

vised, semi-supervised and contrastive – thereby, boosting the representation ability of the network as demonstrated in the results. In contrast, the utilization of a single projection head for both contrastive and classification tasks may interfere with each other’s learning by overwriting the previously updated weights.

The work by Wang *et al.* [2021] performs contrastive learning directly on the output representations of the last convolutional layer, bypassing the use of a projection head. In this scenario, the Euclidean distance is applied to a high-dimensional space, making it susceptible to the curse of dimensionality, which adversely impacts the quality of learned representations. On the other hand, the approach by Aparco-Cardenas *et al.* [2024] employ a projection head to perform contrastive learning, afterward a decision layer is attached on top of it for downstream classification tasks. It should be noted that this strategy may not be suitable for the concomitant learning of both contrastive and classification tasks due to weight overwriting issues.

4.4.2 Influence of active learning

Most semi-supervised methods rely on a fixed reduced set of labeled samples to train a network in conditions of labeled data shortage. Previous DeepFA-based approaches also followed this strategy [Benato *et al.*, 2023; Aparco-Cardenas *et al.*, 2024, 2025]. This reduced set is often sampled randomly and then submitted to an expert for labeling. However, this sampling scheme overlooks the selection of informative samples that could enhance the learning process while lessening the annotation workload. The active-DeepFA method builds on this insight by systematically retrieving the samples whose pseudo-labels have the lowest labeling confidence value, derived after label propagation by OPFSemi. These samples exhibit the characteristic of being the most difficult to classify. Figure 9 shows the heatmap of the labeling confidence values following OPFSemi’s label propagation on the 2D projection of the *Helminth egg*’s output deep features. It can be seen that samples whose pseudo-labels have the lowest confidence values lie in decision boundary regions, where samples of different classes are either in proximity or overlapping. Therefore, it is reasonable to infer that these samples contain meaningful information, which can be leveraged to improve the generalization performance of the networks. This assumption is validated by the results achieving better classification scores than its counterparts with the same amount of labeled samples $|\mathcal{Z}_L| = 5\% \times |\mathcal{Z}_D|$.

Table 8 and Figure 10 display the test results for different percentages of labeled data after the active learning expansion phase. It is observed that comparable results to the baselines are attained with $|\mathcal{Z}_L| = 2\% \times |\mathcal{Z}_D|$ for the MSTAR and PneumoniaMNIST datasets, $|\mathcal{Z}_L| = 3\% \times |\mathcal{Z}_D|$ for the *Helminth larvae* and *Protozoan cysts* datasets, and $|\mathcal{Z}_L| = 4\% \times |\mathcal{Z}_D|$ for the *Helminth eggs*, PathMNIST and TissueMNIST datasets. Therefore, we show that the active learning procedure effectively reduces the annotation effort by requiring a smaller number of labeled samples to obtain similar results to those of the baselines.

In addition, the recurrent expansion of the labeled set through active learning has a critical impact on various com-

Table 7. Test results of accuracy and κ for baselines with $|\mathcal{Z}_L| = 5\% \times |\mathcal{Z}_D|$.

Method	Metric	Dataset						
		<i>Helminth larvae</i>	<i>Helminth eggs</i>	<i>Protozoan cysts</i>	MSTAR	PneumoniaMNIST	PathMNIST	TissueMNIST
Wang et al. [2021]	accuracy	0.970 ± 0.022	0.906 ± 0.004	0.886 ± 0.013	0.849 ± 0.012	0.821 ± 0.036	0.891 ± 0.018	0.594 ± 0.003
	κ	0.878 ± 0.078	0.822 ± 0.006	0.806 ± 0.021	0.832 ± 0.014	0.581 ± 0.091	0.874 ± 0.021	0.481 ± 0.006
conf-DeepFA [2021b]	accuracy	0.966 ± 0.016	0.884 ± 0.010	0.799 ± 0.060	0.815 ± 0.030	0.804 ± 0.011	0.871 ± 0.005	0.498 ± 0.006
	κ	0.867 ± 0.064	0.785 ± 0.024	0.663 ± 0.086	0.794 ± 0.034	0.538 ± 0.032	0.852 ± 0.006	0.368 ± 0.008
cont-DeepFA [2024]	accuracy	0.976 ± 0.015	0.946 ± 0.005	0.922 ± 0.009	0.898 ± 0.024	0.830 ± 0.027	0.917 ± 0.007	0.594 ± 0.003
	κ	0.889 ± 0.064	0.901 ± 0.009	0.870 ± 0.014	0.886 ± 0.027	0.603 ± 0.068	0.905 ± 0.008	0.482 ± 0.003
active-DeepFA (Ours)	accuracy	0.992 ± 0.003	0.965 ± 0.004	0.962 ± 0.009	0.966 ± 0.004	0.910 ± 0.004	0.918 ± 0.007	0.600 ± 0.004
	κ	0.959 ± 0.013	0.937 ± 0.007	0.937 ± 0.015	0.962 ± 0.004	0.799 ± 0.010	0.906 ± 0.008	0.490 ± 0.007




Figure 9. Heatmap of the labeling confidence values after labeling propagation on the 2D projection of the *Helminth eggs*'s deep features via OPFSemi. Left: the prototypes from $\phi(\mathcal{Z}_L)$ are non-gray-colored with a different color per class, while samples from $\phi(\mathcal{Z}_U)$ are gray-colored. Middle: result of label propagation from $\phi(\mathcal{Z}_L)$ to $\phi(\mathcal{Z}_U)$. Right: heatmap of the labeling confidence values, where blue-colored samples with the lowest labeling confidence values lie close to decision boundary regions.




Figure 10. Plots of the test results showing the mean and standard deviation of accuracy and κ for active-DeepFA across different percentage values of labeled data during active learning expansion.

Table 8. Test results of accuracy and κ for active-DeepFA over different percentages values of labeled data during active learning expansion.

$ \mathcal{Z}_L $	Metric	Dataset						
		<i>Helminth larvae</i>	<i>Helminth eggs</i>	<i>Protozoan cysts</i>	MSTAR	PneumoniaMNIST	PathMNIST	TissueMNIST
$2\% \times \mathcal{Z}_D $	accuracy	0.953 ± 0.019	0.890 ± 0.027	0.897 ± 0.015	0.866 ± 0.039	0.861 ± 0.044	0.868 ± 0.006	0.531 ± 0.008
	κ	0.802 ± 0.076	0.799 ± 0.048	0.829 ± 0.023	0.851 ± 0.044	0.680 ± 0.111	0.849 ± 0.007	0.391 ± 0.022
$3\% \times \mathcal{Z}_D $	accuracy	0.975 ± 0.009	0.935 ± 0.016	0.935 ± 0.009	0.921 ± 0.019	0.871 ± 0.017	0.891 ± 0.016	0.528 ± 0.025
	κ	0.889 ± 0.042	0.880 ± 0.032	0.881 ± 0.015	0.912 ± 0.021	0.706 ± 0.041	0.874 ± 0.017	0.399 ± 0.023
$4\% \times \mathcal{Z}_D $	accuracy	0.983 ± 0.005	0.954 ± 0.004	0.949 ± 0.004	0.946 ± 0.011	0.881 ± 0.003	0.906 ± 0.005	0.576 ± 0.024
	κ	0.922 ± 0.026	0.915 ± 0.007	0.915 ± 0.007	0.940 ± 0.012	0.732 ± 0.006	0.892 ± 0.006	0.451 ± 0.039
$5\% \times \mathcal{Z}_D $	accuracy	0.992 ± 0.003	0.965 ± 0.004	0.962 ± 0.009	0.966 ± 0.004	0.910 ± 0.004	0.918 ± 0.007	0.600 ± 0.004
	κ	0.959 ± 0.013	0.937 ± 0.007	0.937 ± 0.015	0.962 ± 0.004	0.799 ± 0.010	0.906 ± 0.008	0.490 ± 0.007

ponents of the active-DeepFA method. The inclusion of these meaningful samples incrementally enhance the low-dimensional space learned by the contrastive projection head P_c via the minimization of \mathcal{L}_S^c , which consequently improves the representation ability of the encoder E . Moreover, these

meaningful samples gradually boost the feature space learned by the classification projection head P_x , which minimizes both \mathcal{L}_S and \mathcal{L}_U^{ssl} . Furthermore, during the label propagation phase, this set of samples compose the set of prototypes from which OPFSemi propagates the labels. This progressively

improves the feature space of the encoder E as epochs evolve, which is demonstrated by the results in Table 8.

4.4.3 Influence of cross-training

The integration of two collaborative networks within a co-training scheme into the DeepFA framework plays a pivotal role in mitigating confirmation bias and overfitting. This is achieved by periodically exposing the networks to new information through the mutual exchange of reliable pseudo-labels. By doing this, the networks prevent overfitting the pseudo-labels generated from the 2D projection of the output feature space of their own encoders, which can be detrimental to the learning process if the pseudo-labels are inaccurate. This is revealed by the results in Table 7, where the methods that implement this strategy – Wang *et al.* [2021], cont-DeepFA and active-DeepFA – achieve better generalization performance on the test set than conf-DeepFA, which trains the network on the pseudo-labels generated from the 2D projection of its own feature space.

During the active learning procedure, each network provides its own set of least-confident pseudo-labels. We capitalize on these two different views of the data by joining the sets and selecting the most uncertain pseudo-labels, thus retrieving more informative samples and enhancing the learning process. Also, active-DeepFA presents an improvement in terms of storage efficiency over its counterparts due to being run in a single iteration and thus producing only two networks. In contrast, Wang *et al.* [2021] and cont-DeepFA generate a number of networks equal to twice the number of iterations.

4.4.4 Influence of label propagation by OPFSemi

The performance of OPFSemi is boosted by conducting label propagation on two distinct data representations derived from the output feature space of the networks’ encoders. The periodic exchange of the most reliable pseudo-labeled samples between the networks improves the feature space as epochs proceed. The 2D projection of this feature space progressively exhibits more defined clusters comprising samples of the same class (see Table 9). This provides suitable projection conditions for OPFSemi, enhancing the label propagation procedure and yielding more accurate pseudo-labels. Table 9 shows the evolution of the encoder’s output feature space for increasing values of percentages of labeled data as the labeled set expands through active learning. It can be seen that as epochs advance, samples of the same class (color) form more separated and compact clusters, which improves the classification metrics, as evidenced by the results in Table 8. The aforementioned effect can also be attributed in part to the samples periodically labeled by active learning that act as prototypes. These samples are regarded as the most difficult to classify and their 2D projections are in close proximity to decision boundaries. Therefore, their inclusion in the set of prototypes during label propagation encourages the production of better-defined and more separated clusters.

In addition, the integration of the OPFSemi classifier in the co-training setup, as implemented in both cont-DeepFA and active-DeepFA, shows an improvement over the generalization performance of Wang *et al.* [2021] (see Table 7),

which conducts pseudo-labeling in a self-training way using the same networks to assign pseudo-labels to the unlabeled samples. This result reveals the effectiveness of OPFSemi as an effective tool for label propagation, providing reliable pseudo-labels in appropriate projection scenarios. Moreover, the application of strong augmentations on the images from the sets $\mathcal{Z}_U^{pl(*)}$, $* \in \{1, 2\}$, during the training of the networks by minimizing \mathcal{L}_U^{ssl} , acts as a key factor in achieving better generalization ability. This is validated by the results shown in Table 7.

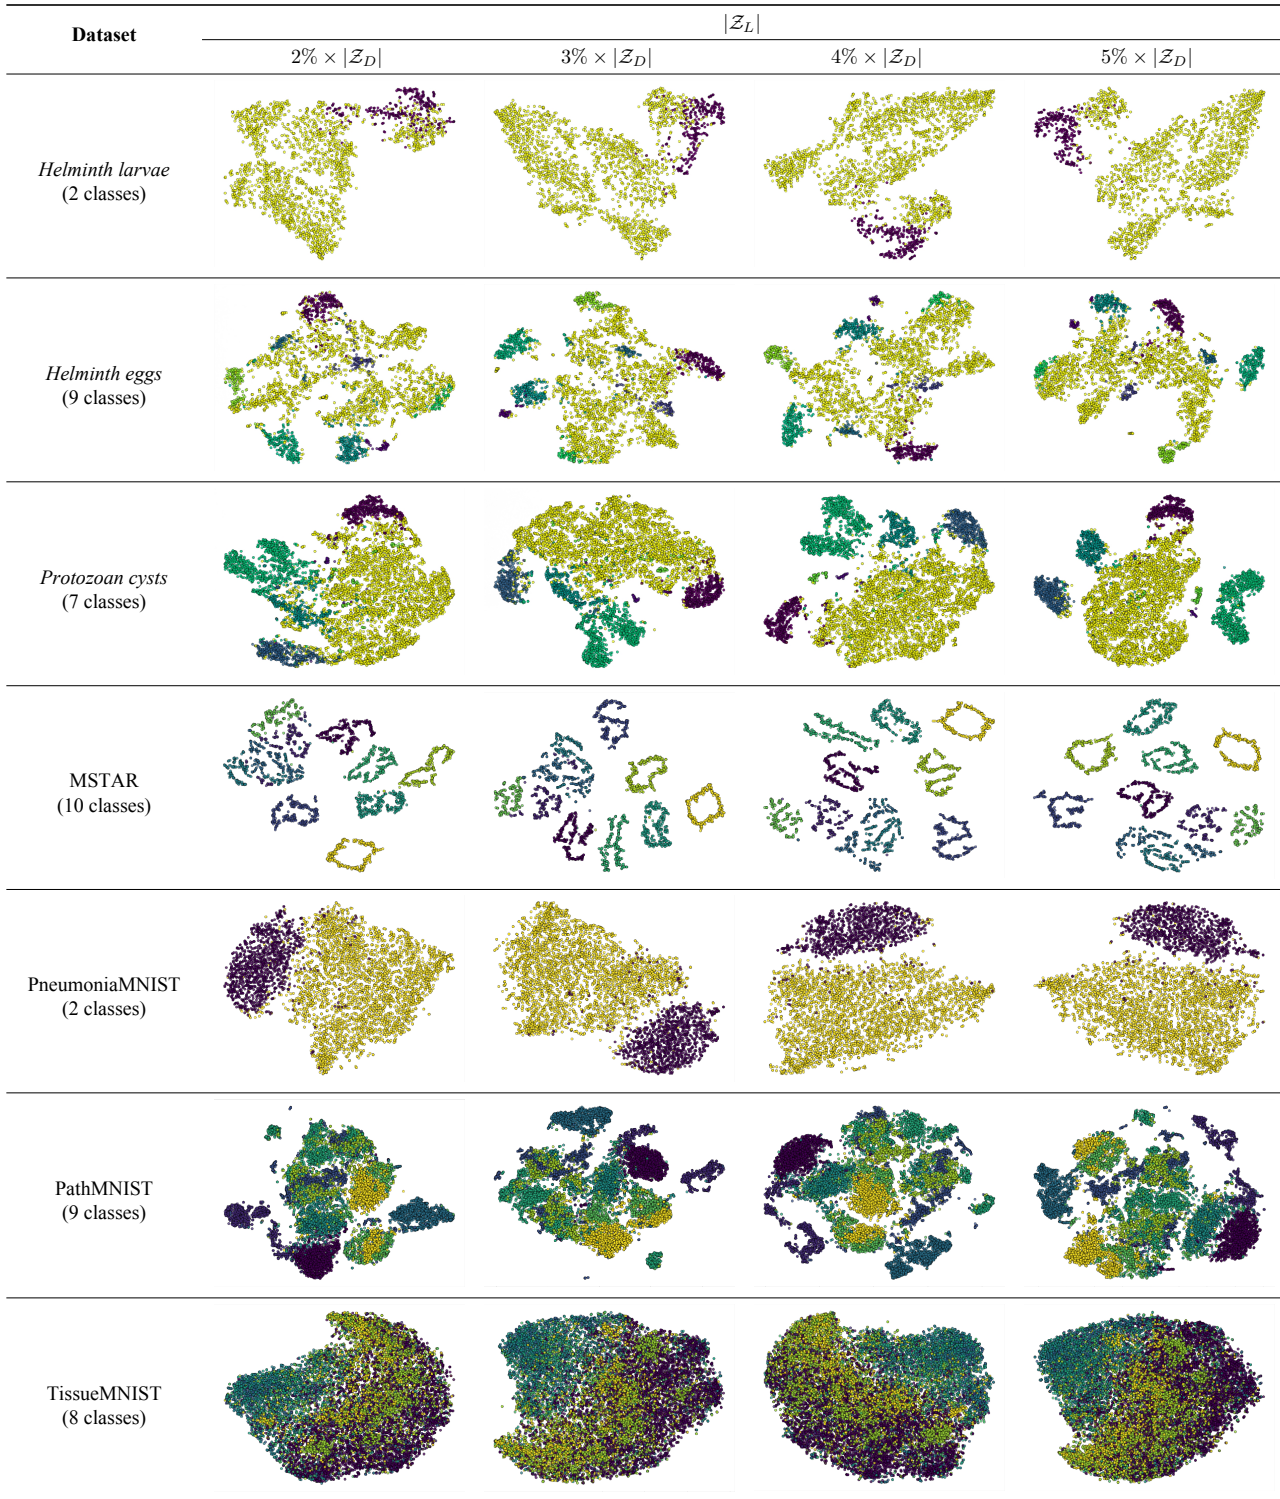
4.5 Comparison with established DSSL benchmarks

This section focuses on benchmarking active-DeepFA with other established DSSL methods from the literature. We selected seven widely adopted DSSL methods categorized into three distinct groups based on their underlying solution strategy: *Pseudo-labeling*: Pseudo-label [Lee *et al.*, 2013], *Consistency Regularization*: Π -Model [Laine and Aila, 2017], Mean Teacher [Tarvainen and Valpola, 2017], VAT [Miyato *et al.*, 2018] and UDA [Xie *et al.*, 2020], and *Hybrid*: FixMatch [Sohn *et al.*, 2020] and SimMatch [Zheng *et al.*, 2022]. We employ the Unified SSL Benchmark (USB) library [Wang *et al.*, 2022b] that contains a PyTorch implementation of the aforesaid DSSL methods. All seven DSSL methods use $|\mathcal{Z}_L| = 5\% \times |\mathcal{Z}_D|$ from the outset. In addition, the labeled set \mathcal{Z}_L is used as validation set since all methods rely on it to select the best-performing model.

Table 10 shows the mean and standard deviation of both accuracy and κ for all seven methods. It can be observed that active-DeepFA excels by consistently surpassing all its counterparts in both metrics across the seven datasets. From Tables 8 and 10, it can be seen that with only $|\mathcal{Z}_L| = 3\% \times |\mathcal{Z}_D|$, active-DeepFA – as with the baselines – achieves results comparable to those of the DSSL methods (which use 5%) for most of the datasets, therefore further validating its efficacy in reducing the annotation burden while attaining satisfactory results.

In particular, active-DeepFA demonstrates evident superiority over Pseudo-label, pointing that the meta-pseudo-labeling approach within a teacher-student setting enhances the conventional self-training pseudo-labeling scheme where the network itself assigns pseudo-labels to unlabeled data. Also, this result supports the role of the OPFSemi classifier as teacher within the teacher-student paradigm. On the other hand, consistency regularization-based methods improve the results achieved by Pseudo-label across all seven datasets, suggesting that incorporating data augmentation techniques and consistency constraints into the objective loss function could further improve the effectiveness of our method. It can also be noted that FixMatch, a hybrid approach that leverages on both self-training pseudo-labeling and consistency regularization, achieves consistent results across all datasets, thus suggesting that incorporating a teacher (*e.g.*, OPFSemi) within a teacher-student setup could enhance its overall performance. Similarly, SimMatch achieves the most competitive performance among all evaluated DSSL methods, underscoring the strength of such frameworks. The comparison

Table 9. Evolution of the 2D projection of the output deep features of encoder E across incremental values of percentages of labeled data as the labeled set expands through active learning.



between active-DeepFA and Mean Teacher, another teacher-student-based approach, highlights the superiority of teacher-student-based pseudo-labeling within a co-training setting. The active-DeepFA method invariably demonstrates superior performance over its consistency regularization counterparts, further substantiating its effectiveness and validating the efficacy of the proposed active learning strategy. An additional point to emphasize is that the unbalanced composition of the datasets coupled with the pronounced similarity among cer-

tain classes – particularly between impurities and parasites for the first three datasets – can account for the divergence observed across the results produced by different methods. In this context, active-DeepFA proves its robustness and flexibility in addressing new data yielding consistent performance across all seven datasets.

Table 10. Test results of accuracy and κ for established DSSL methods with $|\mathcal{Z}_L| = 5\% \times |\mathcal{Z}_D|$.

Method	Metric	Dataset						
		<i>Helminth larvae</i>	<i>Helminth eggs</i>	<i>Protozoan cysts</i>	MSTAR	PneumoniaMNIST	PathMNIST	TissueMNIST
Pseudo-label [2013]	accuracy	0.960 ± 0.010	0.916 ± 0.008	0.903 ± 0.011	0.758 ± 0.015	0.838 ± 0.012	0.854 ± 0.004	0.563 ± 0.020
	κ	0.811 ± 0.057	0.850 ± 0.008	0.836 ± 0.023	0.730 ± 0.017	0.613 ± 0.030	0.833 ± 0.005	0.453 ± 0.011
PI-Model [2017]	accuracy	0.962 ± 0.020	0.922 ± 0.003	0.912 ± 0.009	0.763 ± 0.021	0.839 ± 0.014	0.833 ± 0.015	0.568 ± 0.013
	κ	0.836 ± 0.081	0.855 ± 0.004	0.853 ± 0.016	0.736 ± 0.023	0.618 ± 0.035	0.808 ± 0.017	0.455 ± 0.015
Mean Teacher [2017]	accuracy	0.969 ± 0.009	0.928 ± 0.017	0.906 ± 0.003	0.758 ± 0.020	0.844 ± 0.016	0.847 ± 0.011	0.562 ± 0.017
	κ	0.860 ± 0.037	0.870 ± 0.033	0.843 ± 0.005	0.730 ± 0.023	0.630 ± 0.043	0.824 ± 0.013	0.450 ± 0.021
VAT [2018]	accuracy	0.965 ± 0.006	0.933 ± 0.013	0.912 ± 0.010	0.749 ± 0.040	0.862 ± 0.016	0.850 ± 0.024	0.569 ± 0.012
	κ	0.844 ± 0.031	0.880 ± 0.021	0.854 ± 0.018	0.721 ± 0.044	0.679 ± 0.044	0.829 ± 0.028	0.456 ± 0.012
UDA [2020]	accuracy	0.966 ± 0.009	0.933 ± 0.002	0.914 ± 0.004	0.799 ± 0.031	0.872 ± 0.031	0.858 ± 0.007	0.597 ± 0.006
	κ	0.851 ± 0.033	0.880 ± 0.003	0.857 ± 0.007	0.777 ± 0.034	0.705 ± 0.078	0.837 ± 0.008	0.483 ± 0.003
FixMatch [2020]	accuracy	0.968 ± 0.015	0.925 ± 0.011	0.914 ± 0.006	0.793 ± 0.017	0.889 ± 0.016	0.884 ± 0.020	0.579 ± 0.019
	κ	0.857 ± 0.060	0.862 ± 0.020	0.856 ± 0.011	0.769 ± 0.019	0.750 ± 0.040	0.867 ± 0.023	0.475 ± 0.012
SimMatch [2022]	accuracy	0.988 ± 0.001	0.916 ± 0.018	0.927 ± 0.005	0.843 ± 0.039	0.900 ± 0.009	0.888 ± 0.014	0.596 ± 0.002
	κ	0.945 ± 0.004	0.850 ± 0.033	0.881 ± 0.004	0.823 ± 0.048	0.779 ± 0.020	0.889 ± 0.008	0.488 ± 0.010
active-DeepFA (Ours)	accuracy	0.992 ± 0.003	0.965 ± 0.004	0.962 ± 0.009	0.966 ± 0.004	0.910 ± 0.004	0.918 ± 0.007	0.600 ± 0.004
	κ	0.959 ± 0.013	0.937 ± 0.007	0.937 ± 0.015	0.962 ± 0.004	0.799 ± 0.010	0.906 ± 0.008	0.490 ± 0.007

5 Final Remarks

In the present work, we introduced active-DeepFA, an active contrastive-based meta-pseudo-labeling approach that follows a teacher-student paradigm within a co-training setup. It enables the training of non-pre-trained custom CNN architectures for image classification in conditions of short supply of labeled and abundance of unlabeled data. It achieves this by incorporating deep feature annotation (DeepFA) into a cross-training procedure that implements two collaborative networks, which in turn minimize an objective loss function comprising three components: *i*) a weighted supervised contrastive loss, *ii*) a supervised loss, and *iii*) a semi-supervised loss. Label propagation via OPFSemi and labeled set expansion by active learning are performed at regular epoch intervals. Then, the most reliable pseudo-labels are submitted to the networks for retraining in a cross-training manner, mitigating confirmation bias and overfitting while improving the networks’ generalization performance as epochs proceed. The active-DeepFA method capitalizes on supervised contrastive learning to enhance the networks’ representation ability and to eliminate the requirement for pre-trained CNN architectures. It leverages active learning by periodically retrieving the most uncertain/informative samples identified through label propagation and submitting them to an oracle for annotation. This labeled set expansion has a critical impact on the three components of the objective loss function. The expanded labeled set improves the training set for both supervised contrastive and supervised learning. Also, it enhances the set of prototypes for label propagation via OPFSemi generating more reliable pseudo-labels and thus boosting semi-supervised learning. We assessed active-DeepFA on seven challenging image datasets from three different domains. The labeled data scarcity scenario is emulated by randomly labeling only 1% of the samples of each dataset in a stratified fashion. The labeled set is incrementally expanded at regular epoch intervals through active learning until 5% of the dataset is labeled. We also evaluated active-DeepFA with intermediate percentage values of expanded labeled data during the active learning procedure. The active-DeepFA method is compared to three direct baselines and to seven other well-

established DSSL methods that fall into three different categories according to their solution approach. Our method notably improves its baselines and outperforms most of the DSSL methods across all seven datasets with 5% of labeled data, thus demonstrating its effectiveness and robustness. On the *Helminth larvae*, PathMNIST, and TissueMNIST datasets, active-DeepFA shows results comparable to SimMatch, indicating that its performance gain narrows in cases where the dataset is substantially larger or complex. Furthermore, it achieves comparable results to those of its counterparts with a lower percentage of labeled data, proving their efficacy in reducing the annotation effort while achieving satisfactory results.

As future work, we intend to improve the results here presented by incorporating data augmentation strategies and consistency constraints into the objective loss function to enhance the generalization performance of the models and further reduce the quantity of labeled data to yield satisfactory results. Moreover, we plan to replace t-SNE’s 2D projection with the low-dimensional output of a projection head to make the method faster.

Since the proposed method relies on 2D projections of the data, it naturally enables the integration of interactive active learning. By visualizing the projected samples, practitioners can directly inspect the distribution of labeled and unlabeled data, identify uncertain or informative instances, and guide the annotation process in a more intuitive manner. This interactive component has the potential to further reduce annotation effort, as experts can leverage visual cues to select samples that maximize the impact on model performance.

Another promising direction is to explore the integration of self-supervised approaches, which have recently shown strong performance in reducing the gap with supervised methods. While these models typically rely on large collections of unlabeled data and substantial computational resources, their adaptation to domains with limited and specialized datasets, such as medical and biomedical imaging, represents an exciting avenue for future research.

Declarations

Authors' Contributions

DA, AXF, and PJR contributed to the conception of this study. DA performed the experiments. JFG contributed to the data curation. AXF and PJR contributed to the supervision. DA wrote the original draft. DA, AXF, and PJR contributed to the review and editing. DA is the main contributor and writer of this manuscript. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Funding

The authors acknowledge grants from *Santander Bank Brazil*; CAPES/STIC-AMSUD, #88887.878869/2023-00; CNPq, #140930/2021-3, #314293/2023-0, #313329/2020-6, #442950/2023-3, #304711/2023-3; FAPESP, #2023/14427-8, #2023/04318-7, #2014/12236-1, #2013/07375-0.

Availability of data and materials

The datasets analyzed during the current study are available at <https://github.com/LIDS-UNICAMP/intestinal-parasites-datasets>.

References

- Amorim, W. P., Falcão, A. X., Papa, J. P., and Carvalho, M. H. (2016). Improving semi-supervised learning through optimum connectivity. *Pattern Recognition*, 60:72–85. DOI: 10.1016/j.patcog.2016.04.020.
- Aparco-Cardenas, D., Gomes, J. F., Falcão, A. X., and De Rezende, P. J. (2024). Contrastive learning and iterative meta-pseudo-labeling on 2d projections for deep semi-supervised learning. In *2024 37th SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pages 1–6. IEEE. DOI: 10.1109/SIBGRAPI62404.2024.10716313.
- Aparco-Cardenas, D., Gomes, J. F., Falcão, A. X., and de Rezende, P. J. (2025). Consensus-based iterative meta-pseudo-labeling for deep semi-supervised learning. *Information Sciences*, 692:121671. DOI: 10.1016/j.ins.2024.121671.
- Arazo, E., Ortego, D., Albert, P., O'Connor, N. E., and McGuinness, K. (2020). Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE. DOI: 10.1109/IJCNN48605.2020.9207304.
- Benato, B. C., Gomes, J. F., Telea, A. C., and Falcão, A. X. (2021a). Semi-automatic data annotation guided by feature space projection. *Pattern Recognition*, 109:107612. DOI: 10.1016/j.patcog.2020.107612.
- Benato, B. C., Telea, A. C., and Falcão, A. X. (2021b). Iterative pseudo-labeling with deep feature annotation and confidence-based sampling. In *2021 34th SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pages 192–198. IEEE. DOI: 10.1109/SIBGRAPI54419.2021.00034.
- Benato, B. C., Telea, A. C., and Falcão, A. X. (2023). Deep feature annotation by iterative meta-pseudo-labeling on 2d projections. *Pattern Recognition*, 141:109649. DOI: 10.1016/j.patcog.2023.109649.
- Cascante-Bonilla, P., Tan, F., Qi, Y., and Ordonez, V. (2021). Curriculum labeling: Revisiting pseudo-labeling for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6912–6920. DOI: 10.1609/AAAI.V35I8.16852.
- Dong, S., Wang, P., and Abbas, K. (2021). A survey on deep learning and its applications. *Computer Science Review*, 40:100379. DOI: 10.1016/j.cosrev.2021.100379.
- Gao, M., Zhang, Z., Yu, G., Arık, S. Ö., Davis, L. S., and Pfister, T. (2020). Consistency-based semi-supervised active learning: Towards minimizing labeling cost. In *European Conference on Computer Vision*, pages 510–526. Springer. DOI: 10.1007/978-3-030-58607-2_30.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1026–1034. DOI: 10.1109/ICCV.2015.123.
- Laine, S. and Aila, T. (2017). Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations*. DOI: 10.48550/arxiv.1610.02242.
- Lee, D.-H. et al. (2013). Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, ICML*, volume 3, page 896. Atlanta. Available at: https://www.researchgate.net/publication/280581078_Pseudo-Label_The_Simple_and_Efficient_Semi-Supervised_Learning_Method_for_Deep_Neural_Networks.
- Loshchilov, I. and Hutter, F. (2017). SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*. Available at: <https://openreview.net/forum?id=Skq89Scxx>.
- Miyato, T., Maeda, S.-i., Koyama, M., and Ishii, S. (2018). Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):1979–1993. DOI: 10.1109/TPAMI.2018.2858821.
- Pham, H., Dai, Z., Xie, Q., and Le, Q. V. (2021). Meta pseudo labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11557–11568. DOI: 10.1109/CVPR46437.2021.01139.
- Ren, P., Xiao, Y., Chang, X., Huang, P.-Y., Li, Z., Gupta, B. B., Chen, X., and Wang, X. (2021). A survey of deep active learning. *ACM Computing Surveys (CSUR)*, 54(9):1–40. DOI: 10.1145/3472291.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. (2015). Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115:211–252. DOI: 10.1007/s11263-015-0816-y.
- Sohn, K., Berthelot, D., Li, C.-L., Zhang, Z., Carlini, N., Cubuk, E. D., Kurakin, A., Zhang, H., and Raffel, C.

- (2020). Fixmatch: simplifying semi-supervised learning with consistency and confidence. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*. Curran Associates Inc.. DOI: 10.5555/3495724.3495775.
- Song, H., Ji, K., Zhang, Y., Xing, X., and Zou, H. (2016). Sparse representation-based sar image target classification on the 10-class mstar data set. *Applied Sciences*, 6(1):26. DOI: 10.3390/app6010026.
- Steed, R. and Caliskan, A. (2021). Image representations learned with unsupervised pre-training contain human-like biases. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 701–713. DOI: 10.1145/3442188.3445932.
- Suzuki, C. T. N., Gomes, J. F., Falcão, A. X., Shimizu, S. H., and Papa, J. P. (2013). Automated diagnosis of human intestinal parasites using optical microscopy images. In *IEEE 10th International Symposium on Biomedical Imaging*, pages 460–463. DOI: 10.1109/ISBI.2013.6556511.
- Tarvainen, A. and Valpola, H. (2017). Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in Neural Information Processing Systems*, 30. DOI: 10.5555/3294771.3294885.
- Ulyanov, D. (2016). Multicore-tsne. Available at: <https://github.com/DmitryUlyanov/Multicore-TSNE>.
- Van Der Maaten, L. (2014). Accelerating t-sne using tree-based algorithms. *The Journal of Machine Learning Research*, 15(1):3221–3245. DOI: 10.5555/2627435.2697068.
- Wang, C., Gu, H., and Su, W. (2021). Sar image classification using contrastive learning and pseudo-labels with limited data. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5. DOI: 10.1109/LGRS.2021.3069224.
- Wang, X., Wu, Z., Lian, L., and Yu, S. X. (2022a). Debiased learning from naturally imbalanced pseudo-labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14647–14657. DOI: 10.1109/CVPR52688.2022.01424.
- Wang, Y., Chen, H., Fan, Y., Sun, W., Tao, R., Hou, W., Wang, R., Yang, L., Zhou, Z., Guo, L.-Z., Qi, H., Wu, Z., Li, Y.-F., Nakamura, S., Ye, W., Savvides, M., Raj, B., Shinozaki, T., Schiele, B., Wang, J., Xie, X., and Zhang, Y. (2022b). Usb: A unified semi-supervised learning benchmark for classification. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. DOI: 10.5555/3600270.3600555.
- Xie, Q., Dai, Z., Hovy, E., Luong, T., and Le, Q. (2020). Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33:6256–6268. DOI: 10.5555/3495724.3496249.
- Yang, J., Shi, R., Wei, D., Liu, Z., Zhao, L., Ke, B., Pfister, H., and Ni, B. (2023). Medmnist v2-a large-scale lightweight benchmark for 2d and 3d biomedical image classification. *Scientific Data*, 10(1):41. DOI: 10.1038/s41597-022-01721-8.
- Yang, X., Song, Z., King, I., and Xu, Z. (2022). A survey on deep semi-supervised learning. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–20. DOI: 10.1109/TKDE.2022.3220219.
- Yuan, X., Li, Z., and Wang, G. (2022). Active-match: End-to-end semi-supervised active representation learning. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 1136–1140. DOI: 10.1109/ICIP46576.2022.9898008.
- Zhang, Y., Wen, J., Wang, X., and Jiang, Z. (2014). Semi-supervised learning combining co-training with active learning. *Expert Systems with Applications*, 41(5):2372–2378. DOI: 10.1016/j.eswa.2013.09.035.
- Zheng, M., You, S., Huang, L., Wang, F., Qian, C., and Xu, C. (2022). Simmatch: Semi-supervised learning with similarity matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14471–14481. DOI: 10.1109/CVPR52688.2022.01407.
- Zhou, Z.-H. (2008). Semi-supervised learning by disagreement. In *IEEE International Conference on Granular Computing*, pages 93–93. DOI: 10.1109/GRC.2008.4664785.