

An Evaluation of Self-Supervised Pre-Training for Skin-Lesion Analysis

Levy Chaves^{1,3}[0000-0002-7431-2440], Alceu Bissoto^{1,3}[0000-0003-2293-6160],
Eduardo Valle^{2,3}[0000-0001-5396-9868], and Sandra Avila^{1,3}[0000-0001-9068-938X]

¹ Institute of Computing (IC)

{levy.chaves, alceubissoto, sandra}@ic.unicamp.br

² School of Electrical and Computing Engineering (FEEC)

dovalle@dca.fee.unicamp.br

³ Recod.ai Lab., University of Campinas

Abstract. Self-supervised pre-training appears as an advantageous alternative to supervised pre-trained for transfer learning. By synthesizing annotations on pretext tasks, self-supervision allows pre-training models on large amounts of pseudo-labels before fine-tuning them on the target task. In this work, we assess self-supervision for diagnosing skin lesions, comparing three self-supervised pipelines to a challenging supervised baseline, on five test datasets comprising in- and out-of-distribution samples. Our results show that self-supervision is competitive both in improving accuracies and in reducing the variability of outcomes. Self-supervision proves particularly useful for low training data scenarios (<1500 and <150 samples), where its ability to stabilize the outcomes is essential to provide sound results.

Keywords: Self-supervision, Out-of-distribution, Skin lesions, Melanoma, Classification, Small Datasets

1 Introduction

Self-supervised learning bridges the gap between supervised learning, which leads to the most accurate models but requires human-annotated samples, and unsupervised learning, which can exploit non-annotated samples but often leads to disappointing accuracies. By using synthesized annotations on so-called *pretext tasks*, self-supervision is able to *pre-train* models on abundant pseudo-labels before tuning them for the downstream target task.

Applications for which annotated data is expensive or scarce — often the case for medical applications — especially benefit from self-supervision [2, 32, 43]. Training state-of-the-art Deep Learning models require extensive training datasets, which are seldom available for medical applications. We can mitigate the issue by applying transfer learning, i.e., pre-training the models (with classical supervised learning) on a large, unrelated dataset, and fine-tuning them on the target dataset, but there is a risk that the representations learned during pre-training will not fully adapt to the downstream task [31]. Self-supervised

pre-training has proved, thus, advantageous for transfer learning in many tasks, such as object localization [22], speech representation [26], and medical image classification [1, 40, 49].

In this work, we assess self-supervision pre-training for the automated diagnosis of skin lesions, an application for which traditionally transfer learning from models supervised on ImageNet is employed to mitigate the scarcity of data [31, 41]. Our work improves on existing self-supervised applications to medical applications [2, 21, 40, 42] by evaluating performances on out-of-distribution and low data (less than 150 samples in the stringiest case) regimens. We also evaluate adding an intermediate contrastive learning pre-training before performing the traditional fine-tuning protocol.

The main contributions of this work are:

- We assess five self-supervision learning candidates (BYOL, InfoMin, MoCo, SimCLR, and SwAV) against a competitive supervised baseline;
- We perform a systematic assessment of four transfer learning pipelines (the supervised baseline and three self-supervised contenders) in five publicly accessible test datasets, comprising in-distribution and out-distribution scenarios. Our results suggest that self-supervised models present superior performance in both in-distribution and out-of-distribution in almost all evaluated datasets;
- We assess the performance of our pipelines/datasets in a low-data training scenario (with as few as 148 samples). Again, we find performance improvement in favor of self-supervised pre-training.

We organized the remaining text as follows. We discuss the state-of-the-art on self-supervision in Sec. 2, comprising both general works and those dedicated to medical images, and skin lesions in particular. We detail our goals, datasets, pipelines, protocols, experimental design, and implementation details in Sec. 3. Experimental results and analyses appear in Sec. 4. Finally, we discuss our main findings, along with future research directions in Sec. 5.

2 Related Work

Self-supervised learning has attracted growing attention in the past decade, with hundreds of papers published in the past few years. For a comprehensive review, we refer the reader to the survey of Jing et al. [24]. We recommend the survey by Liu et al. [29] for a more fundamental/theoretical viewpoint on a broad scope of techniques. In this section, we will limit ourselves to the methods directly relevant to this work, and to a selection of methods used for medical images and, in particular, for skin-lesion analysis.

2.1 Self-supervised Learning for Visual Tasks

Self-supervised learning pre-trains models on auxiliary **pretext tasks** such as colorizing [48], predicting rotation angles [15], and in-painting [34], before fine-tuning them on the **downstream task**, i.e., the **target task**. This allows

pre-learning representations on unlabeled data and then refining those representations on labeled data. The base model in self-supervised learning, called the **encoder**, transforms the input image into the **(latent) representations**. ResNet-50 is often employed as a backbone, due to its ability to conciliating simplicity and accuracy [2, 7, 9, 16, 18, 47].

A critical breakthrough in self-supervised learning was the adoption of contrastive losses [24, 38], which explicitly organize the feature space by bringing together the representations for related (positive) pairs of samples, while pushing apart the representations for unrelated (negative) pairs.

InstDisc [47] is the seminal work on contrastive self-supervised image classification. InstDisc reframes class-level classification as instance-level discrimination: each training sample becomes one label, whose data-augmented views must be recognized against data-augmented views from all other training samples. The challenge is extending the loss for so many labels (millions, in ImageNet), which is conquered by reformulating the softmax loss. An ℓ_2 -normalization turns the dot products into cosine similarities. A temperature hyperparameter τ allows regulating the loss concentration. A memory bank caches the parameters for each label/instance. Finally, the softmax is approximated using the Noise-Contrastive Estimation [17], previously successful for training very large word embeddings. The technique creates very compact (128-d) representations, thus making storage and computation for the memory bank feasible, despite the large number of entries.

Instead of using a memory bank, SimCLR [9] employs end-to-end learning [18], adding an auxiliary dimension-reducing network (projection network) after the encoder and generating the representations on the fly for each batch. The pretext task and loss are very similar to InstDisc’s, but only the samples present in the batch are considered in the computation of the loss, without resorting to the memory bank. Thus, SimCLR requires very large batch sizes (4096 samples *vs.* InstDisc’ 256), and strong data-augmented views in order to be effective.

MoCo [18] proposes a dictionary of representations whose size is a hyperparameter that may be much larger than the batch size (which is limited to the GPU memory) while still being much smaller than the training set as in InstDisc. The entries in the dictionary are the past few batches, updated in a FIFO scheme. Pretext task and loss still work similar to InstDisc and SimCLR, but negative examples are now taken from the dictionary. The parameters for each label are updated using a “momentum” update, which prevents the representations from fluctuating too much. That is reminiscent of the proximal regularization of InstDisc, but the latter acts on the loss instead of the representations. Further, MoCo-V2 [11] added the projection network and strong data-augmented views as in SimCLR’s into the original MoCo formulation.

In BYOL [16], one slow network creates targets for a fast network. The parameters of the fast network are learned by backpropagation, and the parameters of the slow network are the exponential moving average of the parameters of the fast network. In that manner, BYOL bootstraps its own target representations.

BYOL still matches data-augmented views between positive pairs as pretext, but without resorting to negative pairs. Instead, it feeds one view to the fast, and the other to the slow network, and uses the cosine distance between the two outputs as loss.

SwAV [7] is an interesting technique that, instead of using instance-based pairwise positive/negative examples, creates pseudo-labels by clustering the representations online, batch by batch. The pretext task is assigning data-augmented views of the same training sample to the same cluster, with an equipartitioning constraint preventing the trivial solution of a single cluster.

In contrast to the techniques above, which use standard data augmentation techniques to create the views of the samples for contrastive learning, InfoMin [39] *learns* how to create the views, using a criterion of minimizing the mutual information between views. The motivation is creating a challenging but feasible pretext task for the model.

2.2 Self-supervised Learning on Medical Tasks

Currently, there are two paths to follow when the matter is using self-supervised learning in medical applications. One is simply to use the same pretext task designed for general purpose computer vision or propose a slightly adapted version of such tasks that fit best into the current medical application. Early medical applications leveraged self-supervised pretext tasks of reconstructing distorted or damaged inputs [6, 8, 20, 32], such as image reconstruction in retinal images [20], context restoration in fetal MRI [8], or depth estimation in monocular endoscopy [30]. Zhou et al. [49], working on X-ray images, employ a domain-general pretext task (the matching of data-augment views of instances of most methods of Sec. 2.1), and uses stronger baselines: both supervised pre-training on ImageNet, and self-supervised pre-training with MoCo [18], still showcasing improvements in downstream image classification. Their technique, *Comparing to Learn*, uses two networks, in knowledge-distillation teacher-student pair, for a momentum encoding scheme somewhat reminiscent of BYOL [16]. MoCo pre-training appears to be widely used in the medical field, bringing superior performance to other medical applications compared to their supervised counterpart for COVID diagnosis [37, 43], and pleural effusion classification [10].

On the other hand, the second way to explore self-supervised is to leverage knowledge about the medical domain – by experience or any domain expert involved – and computer vision to design a custom-built pretext task for the target medical application. Suitable pretext tasks are crucial for learning predictive representations, motivating some works to evaluate whether domain-specific might improve self-supervised learning for medical images. For instance, Jamaludin et al. [23] pre-train a Siamese Net with a contrastive loss in which the positive pairs are patches of spinal magnetic resonance images depicting the same vertebrae of a patient across exams, and the negative pairs are corresponding vertebrae in different patients. They found that the scheme improves the prediction of intervertebral disc degeneration. Wenjia et al. [3] use a pretext in which the model

has to predict the bounding boxes of anatomic features in heart magnetic resonance images, metadata ordinarily available in the DICOM files. They found improvements in the downstream task of segmenting the heart in the images. An issue with all those works is their choice of baseline, networks initialized with random weights, instead of stronger baselines such as models fine-tuned on ImageNet, or other schemes for self-supervision.

Most close to our work of performing a systematic evaluation, Truong et al. [40] assess four medical classifications and three distinct self-supervised pre-training in similar training regimens that ours but lacks evaluation regarding out-of-distribution performances. Hosseinzadeh et al. [21] evaluate the performance of fourteen self-supervised ImageNet models to a diverse set of tasks in medical image classification and segmentation. Again, out-of-distribution and low-data performance evaluations remained uncovered.

2.3 Self-supervised Learning on Skin Lesion Analysis

Wang et al. [44] employ a clustering pretext-task reminiscent of SwAV [7], but accumulating samples from several small batches, and employing different clustering and losses. Since the downstream task is *unsupervised* learning on the same clusters, although evaluated on the classes of ISIC 2018 Lesion Diagnostic Challenge, this work is in a gray zone between self-supervised and purely unsupervised learning. They found favorable results compared to other clustering techniques, but, not surprisingly, a large penalty compared to works that employ supervised fine-tuning. Segmentation tasks also benefits from self-supervision, such as in Li et al. [28], and Wang et al. [46], which both applied the self-supervised with color-based pretext tasks for segmenting skin lesions.

Most related to our work, Azizi et al. [2] performed a well-designed, systematic evaluation of SimCLR [9], for two medical tasks: skin-lesion analysis on a private dataset of $> 450,000$ teledermatology clinical images, and X-rays on the publicly available CheXpert dataset. Contrasting SimCLR pre-training to two strong supervised pre-training baselines, they find it advantageous for the skin-lesion task, and similar for the X-rays task. Their study is complementary to ours, with two medical tasks, three encoder architectures, three pre-training datasets, and the evaluation of a novel pretext technique for exploiting multiple images of the same clinical case they call Multi-Instance Contrastive Learning. Our study, whose focal point is skin-lesion diagnosis, evaluates five test datasets with in- and out-of-distribution images, three pipelines for self-supervision pre-training (in contrast to a challenging supervised pipeline), and five candidate self-supervision schemes. Verdelho et al. [42] compare only two self-supervised learning approaches both quantitatively and qualitatively, but lacks low-data and out-of-distribution evaluation.

We stand out our work from the ones available in the literature in Table 1 by highlighting the contributions of our experimental design.

Table 1: Overview of related works that evaluate self-supervised *vs.* supervised pre-training.

| Work _{year} | #Evaluated Methods | Out-of-distribution Evaluation | Low-data Evaluation |
|--|--------------------|--------------------------------|---------------------|
| Azizi et al. [2] ₂₀₂₁ | 2 | No | Yes |
| Hosseinzadeh et al. [21] ₂₀₂₁ | 15 | No | No |
| Truong et al. [40] ₂₀₂₁ | 5 | No | Yes |
| Verdelho et al. [42] ₂₀₂₂ | 2 | No | No |
| Ours ₂₀₂₂ | 6 | Yes | Yes |

3 Materials and Methods

This section details the methodology, comprising the datasets and factors in our experimental design. We also discuss how we conduct the experimental evaluation of all pipelines.

3.1 Datasets

Following ISIC 2020 Challenge [35], our task is melanoma *vs.* benign lesions classification. We evaluate our experiments in five, high-quality, publicly available datasets (Table 2).

Table 2: Description of the datasets used in this work. Mel.: number of melanomas. †Split used for test if omitted

| Dataset (split†) | Size | Mel. Lesion Diagnoses | Other Information |
|---------------------|--------|--|---|
| isic19 [12] (train) | 14 805 | 3121 Melanoma <i>vs.</i> actinic keratosis, benign keratosis, dermatofibroma, melanocytic nevus, vascular lesion | Dermoscopic images. |
| isic19 (validation) | 1 931 | 224 Idem | Dermoscopic images, in-distribution. |
| isic19 (test) | 3 863 | 396 Idem | Idem. |
| isic20 [35] | 1 743 | 581 Melanoma <i>vs.</i> actinic keratosis, benign keratosis, lentigo, melanocytic nevus, unknown (benign) | Dermoscopic images, out-of-distribution, additional unknown diagnosis. |
| derm7pt-derm [25] | 872 | 252 Melanoma <i>vs.</i> melanocytic nevus, seborrheic keratosis | Dermoscopic images, out-of-distribution. |
| derm7pt-clinic [25] | 839 | 248 Idem | Clinical images, out-of-distribution. |
| pad-ufes-20 [33] | 1 261 | 52 Melanoma <i>vs.</i> actinic keratosis, Bowen’s disease, nevus, seborrheic keratosis | Clinical images, out-of-distribution, additional Bowen’s disease diagnosis. |

We performed all training and validation in splits of the isic19 dataset. We removed samples from isic20 present in the isic19 train/validation splits to avoid contaminating the former. We removed basal cell carcinomas and squamous cell carcinomas from all datasets, leaving melanoma as the only malignant class.

The diversity of test datasets aimed at mitigating bias in evaluation [4, 5, 14], providing both in-distribution (same dataset, same type of image, same classes) and out-of-distribution (cross-dataset, different types of image, different classes) scenarios.

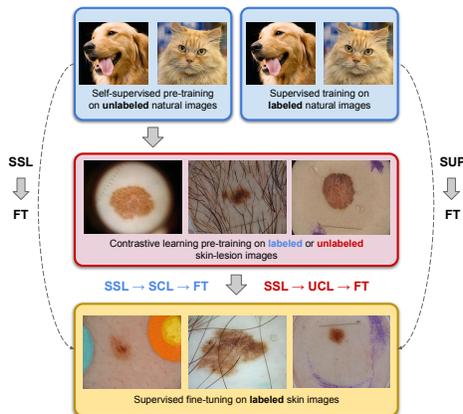


Fig. 1: An overview of our evaluated pipelines. In $\text{SSL} \rightarrow \text{FT}$ scheme we contrast the result of five fine-tuned SSL ImageNet pre-trained models on isic19 dataset (see Sec. 3.2) with the supervised counterpart. The $\text{SSL} \rightarrow \text{SCL} \rightarrow \text{FT}$ pipeline differs from $\text{SSL} \rightarrow \text{UCL} \rightarrow \text{FT}$ according to the employed contrastive loss. They both go through a pre-training stage (see Sec. 3.3) — which can be supervised (SCL) or unsupervised (UCL) — using the isic19 dataset and then performing a supervised fine-tuning. Figure inspired from Azizi et al. [2].

3.2 Experimental Design

We evaluate four alternative pipelines (Fig. 1), which vary in the pre-training and fine-tuning of the model. First, we wish to compare the baseline (supervised pre-training) pipeline ($\text{SUP} \rightarrow \text{FT}$) with the basic self-supervision pipeline ($\text{SSL} \rightarrow \text{FT}$) to establish whether self-supervision is advantageous. In addition, we wish to select a self-supervision scheme among five candidates (BYOL, InfoMin, MoCo-V2, SimCLR, and SwAV) to perform the remainder of the experiments. Selecting the most promising scheme at this stage is necessary for managing the number of experiments, as the next round of experiments will be extensive and, thus, expensive. The $\text{SSL} \rightarrow * \rightarrow \text{FT}$ pipelines have an additional, intermediate pre-training step on the isic19 train split using supervised (SCL) or unsupervised (UCL) contrastive loss (Sec. 3.3).

In the first round of experiments, we attempt a few combinations of hyperparameters for each self-supervision scheme. We purposefully optimize the baseline pipeline more thoroughly to make it challenging. The exact search space appears in Sec. 3.4. We perform all searches on the isic19 validation split to avoid using privileged test information on this step [41]. To estimate the statistical variability of those experiments, we perform five replicates for every experiment, reflecting different random initializations for the training procedures (optimizer, scheduler, and augmentations).

The next round of experiments is a systematic evaluation of all pipelines (Fig. 1) under three data regimens: full training data with 100% of the samples,

and low training data with 10, and 1% of the samples. The latter intends to simulate the frequent scenario on medical images of insufficient training data.

For each combination of pipeline and hyperparameter (see Subsec. 3.4), we measure their performance on the isic19 validation split five times, reflecting different random initializations for the training procedures, and, on the low-data experiments, also different random training subsets. We pick the five non-unique best combinations of hyperparameters for each pipeline. For each combination, we perform five replicates on the isic19 test split, resulting in 25 measurements for each pipeline.

3.3 SSL \rightarrow UCL/SCL \rightarrow FT Pipelines

These two pipelines investigate the benefit of introducing an additional contrastive learning pre-training step before the traditional fine-tuning. Even though it adds an additional computational cost, works in the literature report some advantages [1,2,13], but their evaluation only consider domains with abundant data availability (compared to ours). To this end, we investigate if the same observed improvements also translate to domains with only a few hundred data available.

We evaluated two contrastive losses:

Unsupervised Contrastive Loss (UCL): we performed the pre-training on the isic19 training set using the self-supervised NT-Xent contrastive loss [9]:

$$\mathcal{L}_{UCL} = \frac{-1}{2N} \sum_{i=1}^{2N} \log \frac{\exp(z_i \cdot z_i^+)/\tau}{\sum_{k \neq i} \exp(z_i \cdot z_k)/\tau}, \quad (1)$$

where $z_* = f(x_*)$ is the representation for input x_i output by the encoder f , z_i^+ is the positive pair to z_i , and all vectors z_* are ℓ_2 -normalized. The scalar temperature τ hyperparameter regulates the concentration/spreading of the loss. Only the input data is exploited, and class labels are ignored.

Supervised Contrastive Loss (SCL): we performed pre-training using a straightforward extension of the loss above [27], which incorporates class labels by grouping as positive all examples in the same class (instead of just the augmented pair coming from the same instance):

$$\mathcal{L}_{SCL} = \frac{-1}{2N} \sum_{i=1}^{2N} \frac{1}{|Z_i^+|} \sum_{z^+ \in Z_i^+} \log \frac{\exp(z_i \cdot z^+)/\tau}{\sum_{k \neq i} \exp(z_i \cdot z_k)/\tau}, \quad (2)$$

where Z_i^+ is the set of all representations that are positive to z_i , and the other symbols are the same as in Eq. 1.

3.4 Implementation Details

We follow standard guidelines for self-supervised learning literature [36] and use ResNet-50 (1 \times) [19] as base encoder for all experiments. In SUP \rightarrow FT scheme, we strive to make the baseline challenging, by performing, on the isic19 validation split, a thorough a grid search comprising batch size (32, 128, 512), balanced batches (yes or no), starting learning rate (0.1, 0.05, 0.005, 0.009, 0.0001),

and learning rate scheduler (plateau, cosine). The optimizer is the SGD with a momentum of 0.9 and weight decay of 0.001. The plateau scheduler has patience of 10 epochs and a reduction factor of 10.

For the $\text{SSL} \rightarrow * \rightarrow \text{FT}$ pipelines, we employ two fully-connected layers to embed the ResNet-50 onto 128-dimensional representations, fed to the contrastive loss. We resized the input images to 224×224 and used SimCLR’s recommended heavy image augmentation pretexts — color jitter, horizontal and vertical flips, random resized crop, and grayscale. We omitted the Gaussian blur because, for skin lesions, it leads many images to be very similar, harming the results. We used a learning rate of 0.001 with a cosine decay on an Adam optimizer. We perform a grid search through pre-training batch size (80 or 512), balanced batches (yes or no), temperature scale (0.1, 0.5, 1.0), and pre-training epochs (50 or 200).

For all pipelines, the fine-tuning lasts for 100 epochs with early stopping with patience of 22 epochs, monitored on the validation loss. Both schedulers have a minimum learning rate of 10^{-5} . All experiments ran in a single RTX 5000 GPU, except for the $\text{SSL} \rightarrow \text{UCL/SCL} \rightarrow \text{FT}$ pipelines which required two Quadro RTX 8000 GPUs. The source code to reproduce our work in addition to detailed descriptions about the data is available on our public repository⁴.

4 Results

As explained in Sec. 3.2, we organized our extensive experimental design in two rounds, corresponding to the next two subsections. In a third subsection, we analyze the second round of experiments in a low training data scenario.

4.1 Self-supervision Schemes *vs.* Baseline Comparison

In this first round of experiments, we compared the baseline pipeline ($\text{SUP} \rightarrow \text{FT}$) to the basic self-supervision pipeline ($\text{SSL} \rightarrow \text{FT}$) with five self-supervision schemes (BYOL, InfoMin, MoCo-V2, SimCLR, and SwAV). We optimized the baseline and self-supervised pipelines as explained in Sec. 3.4.

The results (Table 3) show that, despite having no access to the labels during the pre-training, and being less thoroughly optimized during the final fine-tuning, the models with self-supervised pre-training are very competitive. Indeed, two of the pipelines (SimCLR and SwAV) had averages above the ones in the baseline.

This first round of experiments intended to validate the applicability of self-supervised learning, and to select one self-supervised scheme for the expensive round of systematic evaluations in the next round. Thus, it comes with the important caveat that both optimization and evaluation were conducted in the isic19 validation set. The second round of experiments will evaluate the ability of the pipelines to generalize performance in the rigorous setting of a held-out test set.

⁴ <https://github.com/VirtualSpaceman/ssl-skin-lesions>

Table 3: Results for the first round of experiments, comparing the supervised SUP \rightarrow FT baseline to the basic SSL \rightarrow FT pipeline with five SSL schemes. The metric is the AUC on the isic19 validation split. Despite the baseline using label information on pre-training, and being more thoroughly optimized, self-supervision pre-training is still very competitive with it.

| Method | AUC (%) | Hyperparameters | | | |
|-------------------|----------------------------------|-----------------|------------|------------|-----------|
| | | learning rate | batch size | batches | scheduler |
| Sup. baseline | 94.8 ± 0.6 | 0.009 | 128 | balanced | plateau |
| SimCLR [9] | 95.6 ± 0.3 | 0.01 | 32 | unbalanced | plateau |
| SwAV [7] | 95.3 ± 0.6 | 0.01 | 32 | unbalanced | plateau |
| BYOL [16] | 94.6 ± 0.5 | 0.01 | 32 | unbalanced | plateau |
| InfoMin [39] | 94.4 ± 0.5 | 0.001 | 32 | unbalanced | plateau |
| MoCo-V2 [11] | 93.9 ± 0.7 | 0.001 | 32 | unbalanced | plateau |

4.2 Systematic Evaluation of Pipelines

In the second round of experiments, we performed a systematic evaluation of the baseline pipeline, pre-trained with supervision (SUP \rightarrow FT) against the three pipelines pre-trained with self-supervision (SSL \rightarrow FT, SSL \rightarrow UCL \rightarrow FT, and SSL \rightarrow SCL \rightarrow FT). In this round, we only evaluated SimCLR as the self-supervision scheme for several reasons: it showed the best performance in the preliminary experiments (Sec. 4.1), it allows introducing annotation information easily with a supervised contrastive loss, it had one hyperparameter less than SwAV to optimize (number of clusters), and the ablation studies in the original papers helped to decide on a range of reasonable values for the temperature value.

As explained in Sec. 3.2, this round of experiments simulates a realistic machine-learning protocol, in which first we optimize the hyperparameters for each pipeline on the isic19 validation split, then evaluate the performance on a held-out test set. The test sets considered as in-distribution is isic19 test split; and isic20, derm7pt-derm, derm7pt-clinic, and pad-ufes20 as out-of-distribution. Those cross-dataset evaluations are critical to evaluate how well the pipelines generalize to different classes, image acquisition techniques, or even to subtle dataset variations across institutions.

The results appear in the topmost plot of Fig. 2, where each boxplot shows the distribution of 25 individual measurements (small black dots), corresponding to the best five non-unique hyperparameterizations, with five replicates for each of them. The boxplots show, as usual, the three quartiles (box), and the range of the data (whiskers) up to $1.5\times$ the interquartile range (samples outside that range are plotted individually as “outliers”). The large red dots show the means for each experiment. The metric is the AUC on the test datasets labeled on the right vertical axis. To make the horizontal axis comparable across its domain, we linearize the AUC using the logit (i.e., the logarithm of the odds) in base 2, shown on the bottom axis. The original AUC values appear on the top axis.

The plots reveal two advantages for the self-supervised pipelines: first, performances (means and medians) tend to be higher; second, the variability (width of the boxes) tended to be smaller. That shows the ability of the self-supervised pre-training in improving the results and in making them more stable.

No consistent advantage in terms of trend improvement (mean, median) is evident among the different self-supervised pipelines, but in terms of variability reduction, the double-pre-trained pipelines (SSL \rightarrow SCL/UCL \rightarrow FT) appear to have a slight advantage.

4.3 Low Training Data Scenario

These results follow the same protocol as those in the previous section but with drastically reduced train datasets. The results appear in the middle and bottom-most plots in Fig. 2, for 10% (1480 samples) and 1% (148 samples), respectively, of the original train dataset. Other than for this restriction, the interpretation of the plots is the same as in the previous section.

The results are much noisier than the full-data experiments: in part, this is intrinsic to the smaller training sets, but the random choice of training subsets also contributes to increased variability.

Again, the self-supervised pipelines appear advantageous, both in terms of trend improvement (mean, median) and in terms of variability reduction, but here the advantage of the double-pre-trained pipelines (SSL \rightarrow SCL/UCL \rightarrow FT) seems more decisive, especially for the lowest data regimen, where it brings a clear improvement both in trend and variability. As we will discuss in the conclusions, such variability reduction is critical for the soundness of the deployment of low-data models.

4.4 Qualitative Analysis

We performed a qualitative analysis, since we would like to have clues of what different pre-training methods learned to make the decision. Therefore, to analyze the differences between the learned presentations for each pre-training, we performed the Score-CAM [45], method for visualizing the features learned by a neural network and the regions that activate a certain label. We preferred a gradient-free method which overcome both saturation, and false confidence issues [45] from gradient-based techniques.

Fig. 3 shows the Score-CAM results for the top-3 self-supervised models (according to Table 3) *vs.* the supervised baseline for each training regime (100%, 10%, and 1% of training data). We randomly sampled three images from the isic19-test to show true positive, true negative, false positive, and false negatives cases. Apart from having different confidences about the prediction of each sample (bottom left rectangle) each model appears to highlight distinct regions. In general, both confidences and activation maps in true positive, and true negative seems to focus on similar regions in the 100% training regime, but the attention appears sparser in the 10 and 1% training data regime.

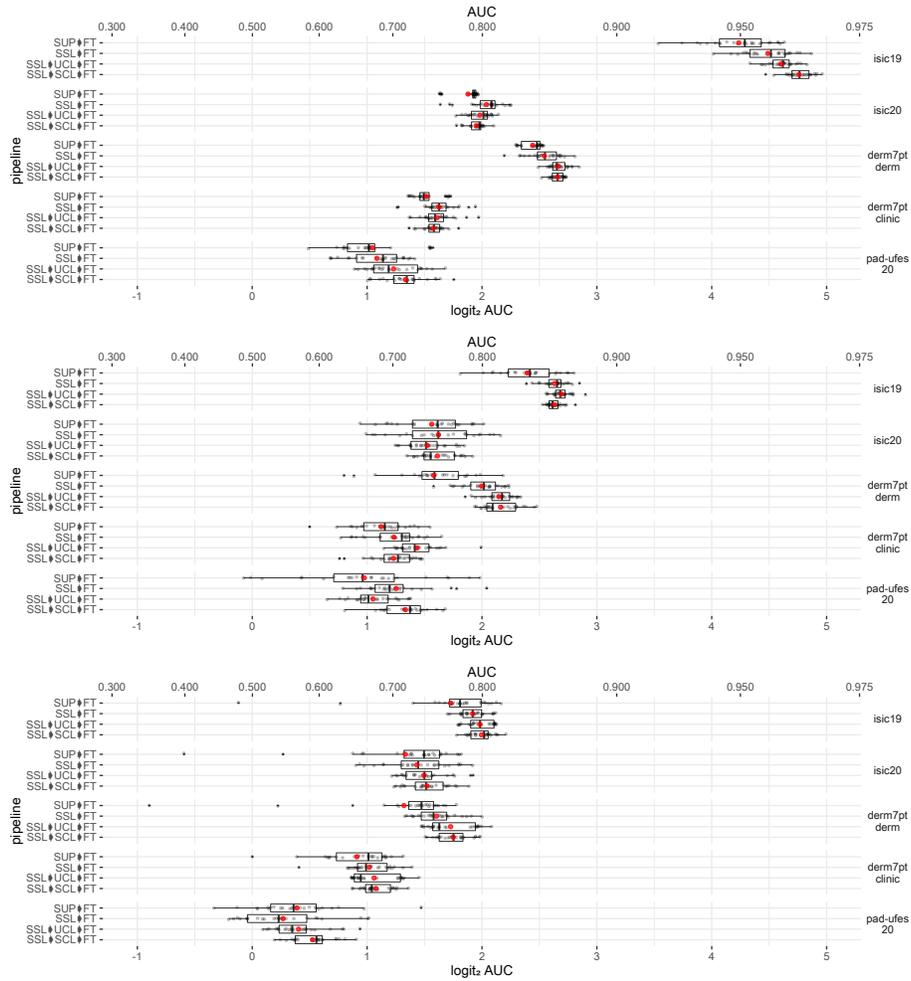


Fig. 2: Results for the second round of experiments, with a systematic comparison of the pipelines labeled on the left vertical axis at the datasets labeled on the right vertical axis. The top, middle, and bottom plots show results for 100, 10 and, 1% of the training data, respectively. Individual measurements represented by each boxplot appear as small black dots, whose means appear as larger red dots. In general, self-supervised pre-trained improves trends (medians, means) and reduces variability, both in the full-data and the low-data scenarios.

5 Conclusions

Our experiments show that self-supervised pre-training makes models easier to deploy than classical supervised pre-training. Even with a less thorough hyperoptimization, the former outperformed the latter in general trends and — especially — in variability.

It is hard to quantify this impression, but the models pre-trained with self-supervised also “felt” easier, faster, and more “ready-to-use” than the baseline models during training.

The advantage of the self-supervised pipelines was particularly prominent in the low-data scenarios, where their ability to stabilize the results, reducing variability, was even more noticeable. In those scenarios, especially the very-low data one, the double-pre-trained models appeared advantageous. We conjecture that self-supervised pre-training improved the mean performance due to the task granularity [1, 13], in which self-supervised models showed to perform particularly better when the pre-training is empowered by contrastive learning. However, those findings are shown experimentally, without any theoretical formulation. Understanding what circumstances make self-supervised competitive (or even superior) from a theoretical perspective is a promising research area.

Very low-data scenarios are not, unfortunately, rare in medical applications. Models trained on such regimens will experience large variances in performance in contrast to models trained with adequate samples, but model designers will often be unaware of such variance (since they cannot run a simulation such as ours, comparing their model to others trained in different datasets of the same size). Our results suggest that self-supervised pre-training may reduce that variability, leading to saner models. Of course, models trained on small samples may suffer from severe biases, and extensive exploration is necessary to evaluate whether self-supervised might reinforce those biases [4]. In addition, we performed qualitative analysis using Score-CAM to visualize the network’s attention.

Self-supervised learning is a thriving research area, and the possibility of creating domain-specific or at least domain-aware pretext tasks for skin-lesion analysis is an exciting avenue of continuation for this work. Domain-aware sampling methods for selecting the positive and negative pairs in contrastive learning — even while using commodity pretext tasks — also instigate the possibility of incorporating domain knowledge into self-supervision learning.

Acknowledgements

L. Chaves is partially funded by Santander and Google LARA 2021. A. Bisoto is funded by FAPESP 2019/19619-7. E. Valle is partially funded by CNPq 315168/2020-0. S. Avila is partially funded by CNPq 315231/2020-3, FAPESP 2013/08293-7, 2020/09838-0, and Google LARA 2021. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001. The Recod.ai lab is supported by projects from FAPESP, CNPq, and CAPES.

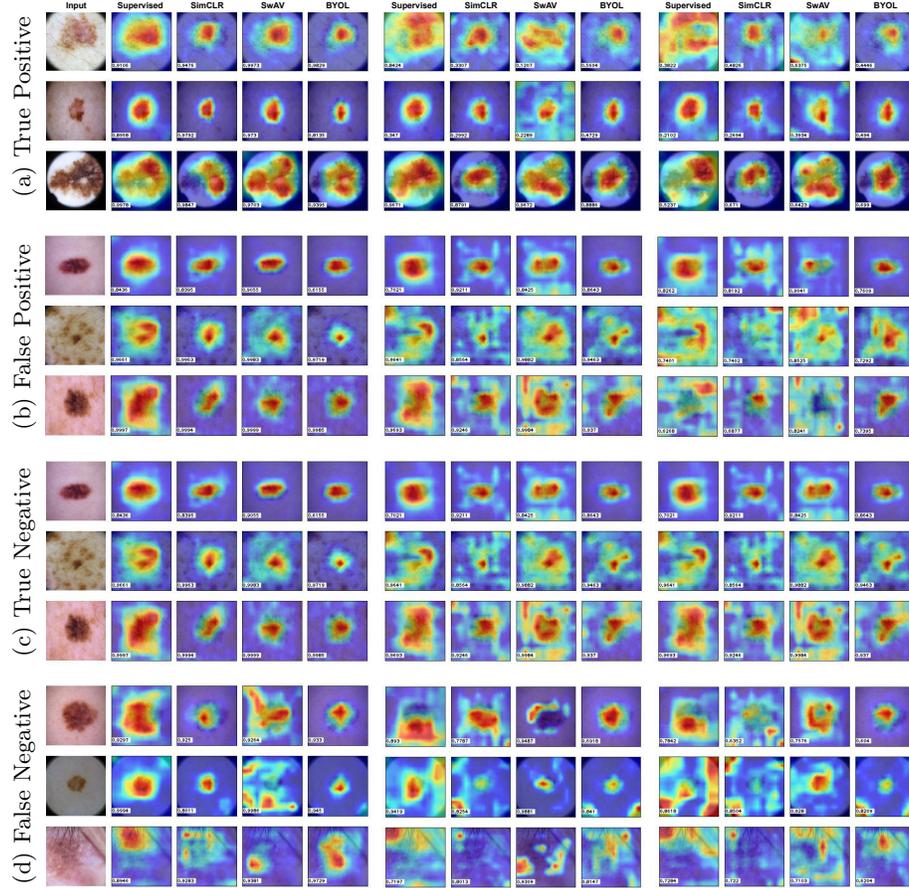


Fig. 3: Results of Score-CAM visualization for the top-3 self-supervised models *vs.* the supervised baseline for each training data regime. The brighter the more relevant the corresponding region in the image is for the model prediction. We randomly sampled three images from the *isic19* test set to show true positive, false positive, true negative, and false negative scenarios; considering the full-data training case. For example, in a true positive scenario, we take three random images in which all images are classified as malignant correctly.

References

1. Azizi, S., Culp, L., Freyberg, J., Mustafa, B., Baur, S., Kornblith, S., Chen, T., MacWilliams, P., Mahdavi, S.S., Wulczyn, E., et al.: Robust and efficient medical imaging with self-supervision. arXiv preprint arXiv:2205.09723 (2022)
2. Azizi, S., Mustafa, B., Ryan, F., Beaver, Z., Freyberg, J., Deaton, J., Loh, A., Karthikesalingam, A., Kornblith, S., Chen, T., et al.: Big self-supervised models advance medical image classification. In: International Conference on Computer Vision (ICCV) (2021)
3. Bai, W., Chen, C., Tarroni, G., Duan, J., Guitton, F., Petersen, S.E., Guo, Y., Matthews, P.M., Rueckert, D.: Self-supervised learning for cardiac mr image segmentation by anatomical position prediction. In: Medical Image Computing and Computer Assisted Intervention (MICCAI) (2019)
4. Bissoto, A., Fornaciali, M., Valle, E., Avila, S.: (De)Constructing bias on skin lesion datasets. In: Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2019)
5. Bissoto, A., Valle, E., Avila, S.: Debiasing skin lesion datasets and models? Not so fast. In: Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2020)
6. Boyd, J., Liashuha, M., Deutsch, E., Paragios, N., Christodoulidis, S., Vakalopoulou, M.: Self-supervised representation learning using visual field expansion on digital pathology. In: International Conference on Computer Vision (ICCV) (2021)
7. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: Advances in Neural Information Processing Systems (NeurIPS) (2020)
8. Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., Rueckert, D.: Self-supervised learning for medical image analysis using image context restoration. *Medical Image Analysis* **58**, 101539 (2019)
9. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International Conference on Machine Learning (ICML) (2020)
10. Chen, X., Yao, L., Zhou, T., Dong, J., Zhang, Y.: Momentum contrastive learning for few-shot covid-19 diagnosis from chest ct images. *Pattern recognition* **113**, 107826 (2021)
11. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
12. Codella, N., Gutman, D., Celebi, M.E., Helba, B., Marchetti, M.A., et al.: Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (ISBI), hosted by the international skin imaging collaboration (ISIC). In: International Symposium on Biomedical Imaging (ISBI) (2018)
13. Cole, E., Yang, X., Wilber, K., Mac Aodha, O., Belongie, S.: When does contrastive visual representation learning work? In: Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
14. Geirhos, R., Jacobsen, J.H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., Wichmann, F.A.: Shortcut learning in deep neural networks. *Nature Machine Intelligence* **2**(11), 665–673 (2020)
15. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: International Conference on Learning Representations (ICML) (2018)

16. Grill, J.B., Strub, F., Alché, F., Tallec, C., Richemond, P., Buchatskaya, E., Dörsch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent – A new approach to self-supervised learning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2020)
17. Gutmann, M., Hyvärinen, A.: Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)* (2010)
18. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
20. Hervella, Á.S., Rouco, J., Novo, J., Ortega, M.: Retinal image understanding emerges from self-supervised multimodal reconstruction. In: *Medical Image Computing and Computer Assisted Intervention (MICCAI)* (2018)
21. Hosseinzadeh Taher, M.R., Haghighi, F., Feng, R., Gotway, M.B., Liang, J.: A systematic benchmarking analysis of transfer learning for medical image analysis. In: *Domain Adaptation and Representation Transfer, and Affordable Healthcare and AI for Resource Diverse Global Health*, pp. 3–13. Springer (2021)
22. Hu, D., Qian, R., Jiang, M., Tan, X., Wen, S., Ding, E., Lin, W., Dou, D.: Discriminative sounding objects localization via self-supervised audiovisual matching. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 33 (2020)
23. Jamaludin, A., Kadir, T., Zisserman, A.: Self-supervised learning for spinal mris. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pp. 294–302. Springer (2017)
24. Jing, L., Tian, Y.: Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
25. Kawahara, J., Daneshvar, S., Argenziano, G., Hamarneh, G.: Seven-point checklist and skin lesion classification using multitask multimodal neural nets. *IEEE Journal of Biomedical and Health Informatics* **23**(2), 538–546 (2019)
26. Kawakami, K., Wang, L., Dyer, C., Blunsom, P., van den Oord, A.: Learning robust and multilingual speech representations. In: *Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2020)
27. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2020)
28. Li, Y., Chen, J., Zheng, Y.: A multi-task self-supervised learning framework forscopy images. In: *International Symposium on Biomedical Imaging (ISBI)* (2020)
29. Liu, X., Zhang, F., Hou, Z., Mian, L., Wang, Z., Zhang, J., Tang, J.: Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering* (2021)
30. Liu, X., Sinha, A., Unberath, M., Ishii, M., Hager, G.D., Taylor, R.H., Reiter, A.: Self-supervised learning for dense depth estimation in monocular endoscopy. In: *OR 2.0 Context-Aware Operating Theaters, Computer Assisted Robotic Endoscopy, Clinical Image-Based Procedures, and Skin Image Analysis*, pp. 128–138. Springer (2018)
31. Menegola, A., Fornaciali, M., Pires, R., Bittencourt, F.V., Avila, S., Valle, E.: Knowledge transfer for melanoma screening with deep learning. In: *International Symposium on Biomedical Imaging (ISBI)* (2017)

32. Morís, D.I., Hervella, Á.S., Rouco, J., Novo, J., Ortega, M.: Context encoder self-supervised approaches for eye fundus analysis. In: International Joint Conference on Neural Networks (IJCNN) (2021)
33. Pacheco, A.G., Lima, G.R., Salomão, A.S., Krohling, B., Biral, I.P., de Angelo, G.G., Alves Jr, F.C., Esgario, J.G., Simora, A.C., Castro, P.B., et al.: Pad-ufes-20: A skin lesion dataset composed of patient data and clinical images collected from smartphones. *Data in brief* **32**, 106221 (2020)
34. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
35. Rotemberg, V., Kurtansky, N., Betz-Stablein, B., Caffery, L., Chousakos, E., Codella, N., Combalia, M., Dusza, S., Guitera, P., Gutman, D., et al.: A patient-centric dataset of images and metadata for identifying melanomas using clinical context. *Scientific data* **8**(1), 1–8 (2021)
36. Srikar Appalaraju, Yi Zhu, Y.X., Fehervari, I.: Towards good practices in self-supervised representation learning. In: Advances in Neural Information Processing Systems Workshops (NeurIPS-W) (2020)
37. Sriram, A., Muckley, M., Sinha, K., Shamout, F., Pineau, J., Geras, K.J., Azour, L., Aphinyanaphongs, Y., Yakubova, N., Moore, W.: Covid-19 prognosis via self-supervised representation learning and multi-image prediction. *arXiv preprint arXiv:2101.04909* (2021)
38. Tack, J., Mo, S., Jeong, J., Shin, J.: Csi: Novelty detection via contrastive learning on distributionally shifted instances. In: Advances in Neural Information Processing Systems (NeurIPS) (2020)
39. Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., Isola, P.: What makes for good views for contrastive learning? In: Advances in Neural Information Processing Systems (NeurIPS) (2020)
40. Truong, T., Mohammadi, S., Lenga, M.: How transferable are self-supervised features in medical image classification tasks? In: Machine Learning for Health. pp. 54–74. PMLR (2021)
41. Valle, E., Fornaciali, M., Menegola, A., Tavares, J., Bittencourt, F.V., Li, L.T., Avila, S.: Data, depth, and design: Learning reliable models for skin lesion analysis. *Neurocomputing* **383**, 303–313 (2020)
42. Verdelho, M.R., Barata, C.: On the impact of self-supervised learning in skin cancer diagnosis. In: International Symposium on Biomedical Imaging (2022)
43. Vu, Y.N.T., Wang, R., Balachandar, N., Liu, C., Ng, A.Y., Rajpurkar, P.: Medaug: Contrastive learning leveraging patient metadata improves representations for chest x-ray interpretation. In: Machine Learning for Healthcare Conference. pp. 755–769 (2021)
44. Wang, D., Pang, N., Wang, Y., Zhao, H.: Unlabeled skin lesion classification by self-supervised topology clustering network. *Biomedical Signal Processing and Control* **66**, 102428 (2021)
45. Wang, H., Wang, Z., Du, M., Yang, F., Zhang, Z., Ding, S., Mardziel, P., Hu, X.: Score-cam: Score-weighted visual explanations for convolutional neural networks. In: Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2020)
46. Wang, Z., Lyu, J., Luo, W., Tang, X.: Superpixel inpainting for self-supervised skin lesion segmentation from dermoscopic images. In: International Symposium on Biomedical Imaging (ISBI) (2022)

47. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
48. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: European Conference on Computer Vision (ECCV) (2016)
49. Zhou, H.Y., Yu, S., Bian, C., Hu, Y., Ma, K., Zheng, Y.: Comparing to learn: Surpassing imagenet pretraining on radiographs by comparing image representations. In: Medical Image Computing and Computer Assisted Intervention (MICCAI) (2020)