Improved skin lesion recognition by a Self-Supervised Curricular Deep Learning approach

Kirill Sirotkin, Marcos Escudero-Viñolo, Pablo Carballeira, Juan C. SanMiguel

Abstract—State-of-the-art deep learning approaches for skin lesion recognition often require pretraining on larger and more varied datasets, to overcome the generalization limitations derived from the reduced size of the skin lesion imaging datasets. ImageNet is often used as the pretraining dataset, but its transferring potential is hindered by the domain gap between the source dataset and the target dermatoscopic scenario. In this work, we introduce a novel pretraining approach that sequentially trains a series of Self-Supervised Learning pretext tasks and only requires the unlabeled skin lesion imaging data. We present a simple methodology to establish an ordering that defines a pretext task curriculum. For the multi-class skin lesion classification problem, and ISIC-2019 dataset, we provide experimental evidence showing that: i) a model pretrained by a curriculum of pretext tasks outperforms models pretrained by individual pretext tasks, and ii) a model pretrained by the optimal pretext task curriculum outperforms a model pretrained on ImageNet. We demonstrate that this performance gain is related to the fact that the curriculum of pretext tasks better focuses the attention of the final model on the skin lesion. Beyond performance improvement, this strategy allows for a large reduction in the training time with respect to ImageNet pretraining, which is especially advantageous for network architectures tailored for a specific problem.

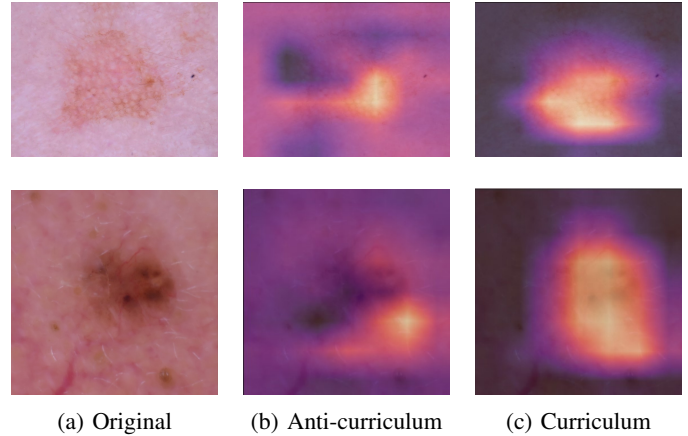(a) Original   (b) Anti-curriculum   (c) Curriculum

Fig. 1: Class activation maps (the brighter the more relevant the corresponding region in the image is for the model prediction) of classifiers pre-trained with pretext tasks ordered in an advantageous (curriculum) and disadvantageous (anti-curriculum) fashion. Each row corresponds to a different sample image.

I. INTRODUCTION

According to the American Institute for Cancer Research, skin cancer is one of the most commonly occurring cancers in humans, and affects more than one million people worldwide every year [1]. Timely cancer diagnosis significantly increases the patient's chances of successful recovery. Therefore, it is of paramount importance to recognize the developing cancer in its early stages. However, many malignant skin lesions might go unnoticed if primary care clinicians do not recognize them as such, due to their lack of expertise or limited diagnosis time. Recent developments in AI-based automated assessment tools led to the creation of systems with significant potential for aiding medical analysis, as acknowledged by the latest EU legislative acts [2]. In particular, AI-based systems can improve skin lesion diagnostics by signaling if qualified medical personnel is necessary for diagnosis. In this regard, methods based on dermatoscopic image analysis, that do not require biopsying the lesion, are the preferred choice for early signaling.

As of today, leading approaches for image-based skin lesion assessment [3] are based on Convolutional Neural Networks (CNNs), achieving up to 72.5% of accuracy on the multi-class classification problem of the popular benchmark dataset—ISIC-2019 [4]–[6], in which skin lesions are to be classified among 8 categories (i.e., melanoma, melanocytic nevus, vascular lesion, etc.). Traditionally, these solutions are trained in a supervised learning setting where the weights of the network are initialized randomly and automatically optimized to correctly predict the labels of the data in the training set by minimizing the value of the loss function that measures correctness of model's predictions [7]. This strategy poses three restrictions on the training dataset—it must be labeled, large in size and diverse [8], so that the features learned from the training data can accurately approximate those found in the validation set. For its application to dermatoscopic imaging analysis—and generally medical imaging, these conditions entail a significant obstacle, as the availability of images themselves is limited due to the privacy restrictions and rareness of certain medical conditions. Moreover, even if the data is available, annotating it requires the skills and expertise of qualified doctors, which further contributes to the difficulty of dataset creation. For these reasons, some of the largest available medical imaging datasets consist of just around 100 000 labeled images (NIH Chest X-Ray dataset [9]), while most others (PAD-UFES-20 [10], TCGA-

LUAD [11], Lumbar Spine MRI [12], etc.) are significantly smaller, sometimes containing only a few hundred images per dataset—tens to hundreds times less than ImageNet [13]. This makes supervised CNN training on such data largely ineffective.

To address the lack of data, it is beneficial to *pretrain* CNN models on (preferably) domain-similar datasets to obtain a better starting point for training on the *downstream task* (i.e., the target task, such as skin lesion classification) [14]. Usually, such starting points are obtained by supervised training of a CNN model on a large and widely accepted as representative dataset, such as ImageNet [13]. Importantly, even when the domains of datasets used for pretraining and downstream task training differ significantly, pretraining still yields better performance than using randomly initialized weights [14]. Nonetheless, this limits the model selection to architectures with publicly available ImageNet-pretrained models, implying a computationally expensive and time consuming training procedure for new model architectures, or architectures designed ad-hoc for specific tasks, such as skin lesion recognition.

In such situations, a common method to pretrain a CNN model is to use Self-Supervised Learning (SSL)—a subset of unsupervised learning methods that leverages automatically generated labels as training objectives. The SSL-driving tasks, known as *pretext tasks*, vary in their formulation and aim at solving problems that do not require manually annotated data: e.g., predicting randomly applied rotations [15], colorization [16] or solving jigsaw puzzles [17]. Previous works show the advantages of SSL-pretraining applied for object recognition [18], where SSL-pretrained models outperform models pretrained on ImageNet in a supervised regime, and skin lesion assessment [19], where SSL-pretraining makes models more robust to noise.

In this study we take a step further for SSL pretraining by showing that the consecutive use of properly ordered pretext tasks significantly improves the results on ISIC-2019, with respect to pretraining using individual pretext tasks. Curriculum learning strategies [20] propose to order samples during training according to their learning outcomes. Inspired by these techniques, we propose orderings of pretext tasks. Specifically, we hypothesize that, given a downstream task, an advantageous ordering of pretext tasks can be obtained by sorting the tasks according to the increasing order of their individual performances. We call orderings following this rule curriculum ones. Experimental results demonstrate that curriculum orderings tend to achieve higher performance on the downstream task than the rest of the possible orderings of the explored pretext tasks.

In addition to the analysis of their accuracies, we provide visual evidence that automatically obtained curriculum orderings of pretext tasks results in more effective internal representations that better focus on the skin lesion (see an example in Figure 1), which explain the increase in classification accuracy.

Overall, our analysis results in the following contributions:

- We show that if pretext tasks are applied sequentially, their ordering has a significant effect on the model's accuracy after fine-tuning.

- Results indicate that curriculum orderings of the evaluated pretext tasks benefit the transferred performance over models not following these orderings.
- We demonstrate that—given a common architecture and using curriculum orderings, sequential SSL pretraining on ISIC-2019 dataset can outperform both supervised pretraining on ImageNet and the current winner of the ISIC-2019 challenge [3], reaching 75.44% of balanced accuracy after fine-tuning for classification, while requiring significantly less time and resources for training.

II. RELATED WORK

This section presents a brief overview of the current state-of-the-art in skin lesion recognition and discusses the use of Self-Supervised and Supervised Learning and in the presence of limited training data, specifically, medical imaging data.

A. Limitation of traditional approaches to image recognition

Most image recognition approaches are based on CNNs. A common strategy to improve a CNN's prediction accuracy on a given task is to increase the complexity of the model. This is done either by significantly changing the underlying network's architecture (i.e., AlexNet [21] vs. VGG [22] vs. ResNet [23]) or adjusting its depth [23] and/or width [24]. However, despite the growing performance on the ImageNet challenge, such approaches do not address the large limitation of any CNN, that is especially relevant to medical imaging datasets: their poor ability to generalize and learn efficient representations with limited training data.

Alternatively to changing the models themselves, one might resort to using ensembles of various networks or the same network trained multiple times (trials) with different initializations [25], [26]. However, this approach is very computationally expensive (as it requires to train the multiple models that comprise an ensemble, and later collect each model's prediction during inference). Furthermore, this paradigm is also not trivial to adapt to new tasks/domains, as it may imply the adaption of every model in the ensemble.

B. Current state of skin lesion recognition

Due to privacy restrictions, to the relatively low rate at which certain medical conditions develop, and to the uneven distribution of suitable capture devices worldwide, a very small amount of medical imaging data is publicly available. In the subfield of skin lesion analysis, one of the largest accessible collections of imaging data is the International Skin Imaging Collaboration (ISIC) dataset that aggregates the data from other sources, such as HAM_10000 dataset [4], MSK dataset [5] and BCN_20000 dataset [6]. Annual public challenges based on the ISIC dataset target multi-class (predicting the exact type of a skin lesion) or binary (malignant vs. benign) classification problems. Whereas the latter can be considered an almost solved problem (ISIC-2020—a binary melanoma recognition problem—challenge winners achieved 0.949 on ROC AUC metric), the multi-class problem proposed in ISIC-2019 is still an open one, where the best reported approach [3]

currently reaches only 72.5% ± 1.7% of *balanced accuracy* (average per-class accuracy).

Although ISIC-2019 is the largest publicly available skin lesion dataset with multi-class annotations, containing over 25000 labeled images, its size is small compared with those used for standard CNN training in well established tasks. Moreover, the number of samples per class ranges from 239 to 12875, making it highly imbalanced and further complicating CNN training—e.g., a vanilla ResNet-50 [23] does not reach 50% of balanced accuracy, as shown in Table I. Some works address this issue by designing new loss functions that account for severe class imbalances [27]. However, the general trend seems to be increasing the complexity of neural models and utilizing deeper architectures, such as DenseNets [28] or very deep ResNets [23]. Continuing in the same direction, the top three best performing approaches in ISIC-2019 skin lesion diagnosis challenge are based on ensembles of neural networks that leverage multiple models to infer predictions [3], [29], [30].

C. Transfer learning and Self-Supervised Learning

A different and popular approach to increase the classification accuracy of deep learning architectures is transfer learning: using a model trained for one task as a starting point for training for a different task. Recent studies introduce strategies for pretraining models without relying on labelled data [15], [18]. This is achieved by using SSL-based methods that solve a *pretext task* by formulating objectives that do not require manually labelled data. It is assumed that solving pretext tasks requires a high-level understanding of the data and, therefore, allows to learn efficient visual representations. Hereby, the features learned through solving *pretext tasks* are considered to be starting points suitable for transfer learning. For instance, a well-known *pretext task*—rotation prediction [15]—generates training labels that encode a series of rotations applied to the input image, anticipating that the information required to correctly rotate an image requires the *learning* of some fundamental object qualities.

D. Curriculum learning

Although the number of existing SSL strategies is constantly increasing, a categorization according to the underlying label generation strategy can be used to assign them to one of three categories: **geometric** (Rotation prediction [15], Relative patch location [31], Jigsaw puzzles [17]), **clustering** (Deep Clustering [32], Online Deep Clustering [33], ClusterFit [34]) and **contrastive** (Momentum Contrast [18], [35], Bootstrap Your Own Latent [36], Non-Parametric Image Discrimination [37], Simple Framework for Contrastive Learning of Visual Representations [38]).

Relative patch location prediction, being a typical **geometric** SSL model, splits an input image into patches, then samples two adjacent ones and trains to predict their relative location. On the other hand, **clustering** SSL models, such as Online Deep Clustering (ODC), use classical clustering algorithms to generate intermediate training labels. Finally, **contrastive** models rely on the use of the contrastive loss function to discriminate between positive and negative samples. Often the negative samples are generated through the data augmentations of positive samples, as it is done in the case of Momentum Contrast (MoCo) model.

Currently, regarding the accuracies achieved using ImageNet as the downstream classification task, **contrastive** models outperform **geometric** and **clustering** ones: MoCo v3 [39]—81.0%, SimCLR v2 [40]—79.8% and BYOL [36]—79.6%, while the closest **clustering** model—DeepCluster v2 [32]—yields only 75.2% of accuracy and a classic **geometric** model—Rotation prediction—reaches 55.4% of accuracy.

E. SSL in skin lesion recognition

The promise of learning useful representations without requiring labelled data led to the applications of SSL strategies to the medical imaging, where data labelling is challenging. In particular, a number of recent works target the skin lesion recognition problem leveraging SSL approaches for model pretraining. This narrows the gap between SSL and fully supervised pretraining on ImageNet but, in most cases, does not yet close it. For example, a recently proposed approach [41] reaches 80.6% of accuracy on a multi-class ISIC-2018 classification problem by employing self-supervision to obtain transformation-invariant features. Specifically, features extracted at each epoch from the image decoder module of a CycleGAN architecture [42] are assigned to N clusters without a prior knowledge of N, using the maximum modularity clustering algorithm [43]. The memberships of the samples are used as pseudo-labels to optimize the features. A different study compares the individual performances of five existing SSL models [44] (BYOL [36], MoCo [35], SimCLR [38], InfoMin [45], SwAV [46]) on the ISIC-2019 data, but only considers the binary classification problem, reaching 0.956 on the standard for binary problems metric—ROC AUC (area under the precision-recall curve). On the contrary, another recent work does consider the multi-class skin lesion classification problem and shows that SimCLR pretraining outperforms supervised pretraining [47]. However, this study relies on a private dermatoscopic dataset containing over 450 000 samples. Nevertheless, there is a gap in the current state-of-the-art as, to our knowledge, none of the existing approaches considers more than one SSL task in the pretraining stage, thus, limiting the resulting performance by not taking advantage of the pretext tasks of different nature. Moreover, with seldom exceptions, most works that use SSL pretraining in the skin lesion domain rely only on **contrastive** models (as these are generally the most accurate ones), thus, the potential contributions of **clustering** and **geometric** models are still barely explored. Finally, a multi-class skin lesion recognition problem is relatively unexplored by SSL methods, as the majority of works focus on the binary melanoma recognition task, and only a few studies target the multi-class problem of the older (and smaller) versions of the ISIC dataset.

A promising method for improving the accuracy of deep learning architectures emerges from the careful selection and ordering of training samples—curriculum learning [20], [48],
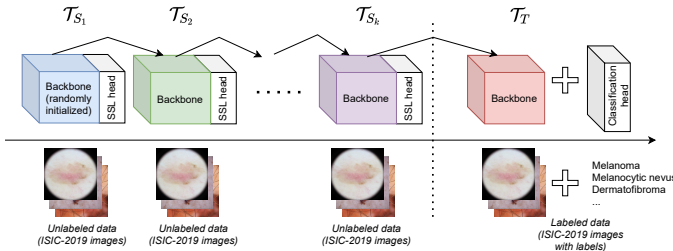
Fig. 2: Proposed pretraining scheme.

[49]. The idea behind it is that, inspired by human learning processes, the training of deep learning architectures may benefit from ordering of data samples based on some evidence of their *complexity*, e.g., their contribution in the training loss.

Generally, a curriculum of samples starts by using for training those samples that are learned faster and progressively incorporates the rest. For visual tasks, reported results indicate that curriculum training results in models that yield better performance than those obtained when samples are selected at random [20]. However, despite the benefits of curriculum learning at the samples level, to our knowledge no study investigates the effects of a different orderings of models pre-trained with multiple sequential tasks. In this work we define a curriculum order of pretext tasks that optimizes the performance of a model pre-trained sequentially on those tasks. Finally, we investigate applications of curriculum ordering of pretext tasks and their effects on the skin lesion recognition performance.

III. PROPOSED APPROACH

In this study we propose to increase the accuracy on the ISIC-2019 multi-class problem by improving the pretraining pipeline, frequently applied in transfer learning approaches to the skin lesion recognition problem. The basic idea of our method (see Figure 2) is to extend the recent works on SSL pretraining for skin lesion recognition and leverage multiple, instead of a single, pretext tasks in the pretraining stage.

A. Preliminaries

To clarify the proposed approach we first formalize the notation used in this work, following the nomenclature proposed in previous works [50]:

Domain: A domain $\mathcal{D} = \{\mathcal{X}, P(X)\}$ is comprised of the complete feature space $\mathcal{X}$ (i.e., all images) and a marginal distribution $P(X)$, where $X = \{x | x_i \in \mathcal{X}, i = 1, ..., n\}$ is the set of images defining the domain (i.e., skin lesion images).

Task: A task $\mathcal{T} = \{\mathcal{Y}, \mathcal{F}\}$ is comprised of the label space $\mathcal{Y}$ (i.e., melanoma, melanocytic nevus,...) and the implicit decision function $\mathcal{F}$ which is expected to be learned by the model from the input data (i.e., conditional probability of labels, given the samples).

Transfer learning: Given (a) data distribution(s) $\{\mathcal{D}_S, \mathcal{T}_{S_i} | i = 1, ..., m\}$ drawn from the *source* (S) domain(s)

TABLE I: Accuracies for the evaluated single- and multi-source transfer settings for the ISIC-19 skin lession recognition task. The right-most column indicates whether the pretraining strategy led to a higher classification accuracy than supervised pretraining on ImageNet. The column "δ" indicates how the performance of a combination of pretext tasks differs from an individual pretext task. The left-most column shows whether a combination follows Curriculum (C), Anti-Curriculum (AC) or Mixed Curriculum (MC) ordering.

	1st task	2nd task	3rd task	Balanced accuracy (%)	δ (%)	Better than ImageNet
-	Rel. loc.			69.52	-	No
(AC)	Rel. loc.	ODC		70.68	1.16	No
(MC)	Rel. loc.	ODC	MoCo v2	75.00	5.49	**Yes**
(C)	Rel. loc.	MoCo v2		74.10	4.58	**Yes**
(MC)	Rel. loc.	MoCo v2	ODC	74.38	4.86	**Yes**
-	MoCo v2			72.74	-	No
(AC)	MoCo v2	ODC		72.72	-0.02	No
(MC)	MoCo v2	ODC	Rel. loc.	67.00	-5.74	No
(AC)	MoCo v2	Rel. loc.		66.72	-6.02	No
(AC)	MoCo v2	Rel. loc.	ODC	69.80	-2.95	No
-	ODC			63.52	-	No
(C)	ODC	Rel. loc.		68.23	4.71	No
(C)	ODC	Rel. loc.	MoCo v2	73.36	9.84	No
(C)	ODC	MoCo v2		75.44	11.92	**Yes**
(MC)	ODC	MoCo v2	Rel. loc.	65.73	2.21	No
	ISIC-2019 challenge winner [3]			72.5 ± 1.7	-	-
	Supervised ImageNet			73.76	-	-
	No pretraining			49.27	-	-

and task(s), and a data distribution $\{\mathcal{D}_T, \mathcal{T}_T\}$ drawn from the *target* (T) domain and task, transfer learning uses the knowledge from the source domain(s) to improve the performance of the decision function $\mathcal{F}_\mathcal{T}$ on the target domain.

Single-source transfer learning: Transfer learning setup in which $m = 1$. In other words, the knowledge is transferred from a single source domain to the target domain. Most transfer learning approaches follow single-source scenario.

Multi-source transfer learning: Transfer learning setup in which $m \geq 2$. In other words, the knowledge is transferred from multiple source domains to the target domain.

B. Self-supervised Curricula

With these definitions, we formulate the proposed approach as follows: given a single source domain and k source tasks $\{\mathcal{D}, \mathcal{T}_{S_i} | i = 1, ..., k\}$, we define a multi-source transfer learning approach that uses the knowledge from $\{\mathcal{D}, \mathcal{T}_{S_i}\}$ to improve the performance of the learned target decision function $\mathcal{F}_T$. In this setting, the knowledge is sequentially transferred from $\{\mathcal{D}, \mathcal{T}_{S_i}\}$ to $\{\mathcal{D}, \mathcal{T}_{S_{i+1}}\}$ (the first task $\mathcal{T}_1$ is trained from scratch). At step k we transfer the knowledge from $\{\mathcal{D}, \mathcal{T}_{S_k}\}$ to the target task $\{\mathcal{D}, \mathcal{T}_T\}$.

Note, that in our case the source $\mathcal{D}_S$ and the target $\mathcal{D}_T$ domains are the same, therefore, the whole process can be

TABLE II: Comparison of SSL pretraining approaches on the ISIC-2017 binary classification challenge. Pretext-Invariant Representation Learning (PIRL) [51] with image transformations produced by the Jigsaw [17] and Rotation prediction [15] pretext tasks, along with individual Relative Location [31], ODC [33] and MoCo v2 [35] pretext tasks are outperformed by a combination of Relative Location, ODC and MoCo v2. ROC AUC stands for the area under the receiver operating characteristic curve, PR AUC stands for the area under the precision-recall curve.

Starting point	ROC AUC	PR AUC	Accuracy
Random initialization	0.598	0.257	79.67
ImageNet	0.791	0.525	83.03
PIRL (Jigsaw) [52]	0.664	0.32	79.8
PIRL (Rotation) [52]	0.732	0.37	80.3
Relative Location	0.663	0.33	78.5
ODC	0.696	0.37	80.33
MoCo v2	0.731	0.4	80.16
RL $\rightarrow$ ODC $\rightarrow$ MoCo v2 (ours)	**0.754**	**0.48**	**82.3**

viewed as the sequential training of $k+1$ models on the source domain.

Currently, optimal ordering of k pretext tasks remains an open question. We define the best performing ordering as the curriculum one and make a hypothesis that it coincides with the ordering $\mathcal{A}_1, \mathcal{A}_2, ...\mathcal{A}_k$ of accuracies obtained with each pretext task individually, where $\mathcal{A}_1 \leq \mathcal{A}_2 \leq ... \leq \mathcal{A}_k$. Further, in Section IV, we present experimental findings that support this hypothesis.

IV. EXPERIMENTS

A. Setup

We follow the method described in Section III to assess the effects of sequential pretraining of multiple pretext tasks on the ISIC-2019 multi-class classification challenge (as the latest version of ISIC that addresses the multi-class problem). As consecutive learning of pretext tasks changes the learned representations, it can be expected that same-category pretext tasks learned consecutively will not change the learned representations significantly as they are optimized on the basis of similar objectives (i.e., the representations learned by the pretext tasks of the same category—**geometric**, **clustering** or **contrastive**—could be expected to be similar). Therefore, we pick three representative examples of each group of pretext tasks (Relative Location, MoCo v2, ODC) to ensure that the representations learned by them differ between each other, define curriculum orderings and compare them against all other possible combinations of them in the training pipeline. The following steps are used to define the order of pretext tasks:

1) Given three pretext tasks, we perform single-source transfer learning with the same target and source domains and record the achieved accuracies: $\mathcal{A}_{RL}, \mathcal{A}_{ODC}, \mathcal{A}_{MoCo\_v2}$. All three pretext tasks use the same backbone architecture: ResNet-50 [23].
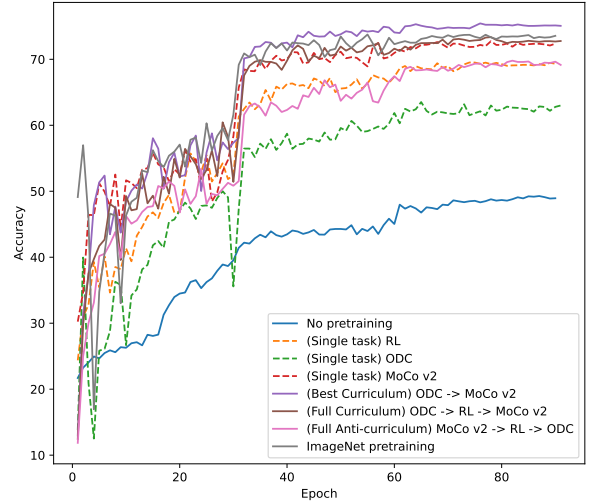

Fig. 3: Validation accuracies of classifiers (ResNet-50 models) trained on the features learned with different pretext tasks (or their combinations).

2) Following the hypothesis outlined in Section III we define as curriculum orderings those coinciding with the increasing single-source transfer learning accuracies (accuracy hereafter). We group all orderings as follows:
   - Curriculum: ordered by increasing accuracy.
   - Anti-curriculum: ordered by decreasing accuracy.
   - Full (anti-)curriculum: (anti-)curriculum comprised by all tasks.

   Any other ordering is considered a Mixed-curriculum.

Given the individual accuracies of pretext tasks, we define all the curriculum orderings of two and three pretext tasks for the multi-source transfer learning and train corresponding models. Further, in Section IV-B, we show that these curriculum orderings are the best performing ones (see Tables I and II), thereby, proving our hypothesis.

Training regimes and hyperparameters for the pretext tasks are presented in Table III. When transferring the knowledge to the target task, we select the checkpoint from the last pretext task in the pretraining pipeline at the epoch at which the best accuracy is achieved on the validation set (Relative Location) or we simply load the weights from the last epoch when validation set is not available due to the nature of the pretext task (ODC, MoCo v2).

As baselines for comparison we use a model trained from scratch on the target—skin-lesion recognition—task (No pretraining) and a model fine-tuned end-to-end on the target task but pre-trained on the ImageNet dataset (Supervised ImageNet). Training hyperparameters for these two models are compiled in the last two rows of Table III

B. Results for skin-lesion classification

Table I presents the classification accuracies on the ISIC-2019 dataset achieved with different single-source and multi-

source transfer learning setups. Given the classification accuracies obtained with the three single-source setups (first row of each block in Table I): $\mathcal{A}_{ODC} = 63.5$, $\mathcal{A}_{RL} = 69.5$, $\mathcal{A}_{MoCo\_v2} = 72.7$, we define the full curriculum in the multi-source approach as ODC → Relative Location → MoCo v2 and the full anti-curriculum as MoCo v2 → Relative Location → ODC.

The results shown in Table I also demonstrate the effect of multi-source pretext task pretraining on the classification accuracies on ISIC-2019. Specifically, when the pretraining begins with Relative Location or ODC, thereby avoiding the supposed anti-curriculum ordering, the resulting accuracy is always superior to single-source pretext task pretraining. On the contrary, orderings that are expected to be anti-curriculum (or mixed curriculum) starting with MoCo v2 (e.g., MoCo v2 → ODC; MoCo v2 → ODC → RL; MoCo v2 → RL; MoCo v2 → RL → ODC) always harm the performance of the individual model (MoCo v2).

Moreover, five pretraining setups lead to state-of-the-art results after transfer learning, including one where only a single SSL model was trained (MoCo v2); two pretraining setups outperform the current ISIC-2019 winner [3] (Relative Location → ODC → MoCo v2 and ODC → MoCo v2); and four pretraining setups outperform the results achieved with fully supervised ImageNet pretraining, with the best result of 75.44% achieved with a pretraining pipeline of two pretext tasks following a curriculum ordering (ODC → MoCo v2). Overall, combinations of pretext tasks following a curriculum ordering outperform the individual models in all cases and reach the highest performances. Hence, the results support our definition of curriculum.

Figure 3 also demonstrates the accuracy curves on the target task for models pretrained with different setups and reveals that more accurate models do not only reach higher performances but also do it faster than less accurate ones.

To our knowledge, no work targeting a multi-class problem on ISIC-2019 dataset used SSL methods. Therefore, we additionally evaluate our approach on an older version of the ISIC dataset—ISIC-2017 [53]—where SSL approaches were tested before. Table II summarizes the results and demonstrates that while none of the three SSL models tested by us (Relative Location, ODC and MoCo v2) outperforms previously published results, their curriculum combination does.

Finally, the bar chart shown on Figure 4 indicates that the accuracy gains achieved by different pretraining setups are evenly distributed among all categories of ISIC-2019 and do not come from a large increase in accuracy in only one or few categories: combinations following the curriculum ordering (ODC → MoCo v2 and ODC → RL → MoCo v2) consistently improve the results of the two individual pretext tasks (ODC and MoCo v2).

C. Qualitative analysis of the results

Here, we use Class Activation Mapping (CAM) [54] to demonstrate the advantages of a curriculum ordering of pretext tasks in terms of how accurately the model focuses on areas of the input image that are relevant to the target task. The use
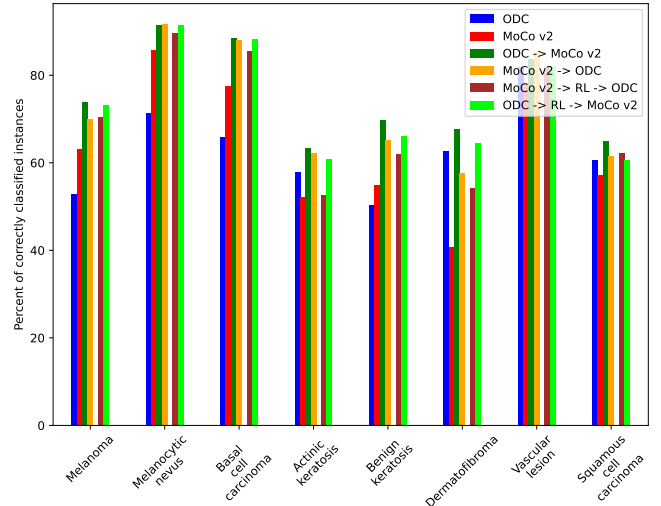

Fig. 4: Percentage of correctly classified images per class by models with different pretraining strategies.

of CAM allows to visualizr the discriminative image areas that have a higher impact in the class prediction of a trained network.

Figure 5 compares CAMs of models without pretraining, with those using supervised ImageNet or SSL pretraining, Figures 6 and 7 present CAMs of models pretrained with two individual pretext tasks and their combinations arranged according to curriculum or anti-curriculum orderings: Figure 6 compares the effects of ODC and MoCo v2, and Figure 7—Relative Location and MoCo v2 (we compare ODC and Relative Location against MoCo v2 as it results in the best individual performance). Figure 1 extends the analysis to the full curriculum and anti-curriculum orderings of pretext tasks.

Hereinafter, our visual analysis is based on the simple assumption that—for lesion assessment, the lesion and its borders are more important than other skin tissues (i.e. image regions that are not part of the lesion and its borders).

Figure 5 presents evidence that for the same input sample image, a non-pretrained model (Figure 5b) does not focus attention on the skin lesion as much as the pretrained models do (Figures 5c-5e)—CAM highlights indicate that a non-pretrained model focuses more attention on the parts of the images that do not contain lesions and gives more importance to the black areas surrounding the image that carry no useful information for image classification. Moreover, Figure 5e demonstrates that the best combination of pretext tasks leads to a model that focuses on the lesions better than the ImageNet-pretrained model: CAM highlights cover the lesion and are absent in other areas of the images (in the presented examples, only the SSL-pretrained models led to the correct classification of skin lesions). Figures 6 and 7 further compare different SSL pretraining setups and show that curriculum ordering of pretext tasks significantly better focuses on the skin lesions than anti-curriculum ordering and individual tasks—in the CAM images the lesions are highlighted in brighter colors while the rest of the image is dark. In examples presented on

(a) Original   (b) No pretraining   (c) Supervised pretraining (ImageNet)   (d) MoCo v2 (best single)   (e) ODC → MoCo v2 (best combination)
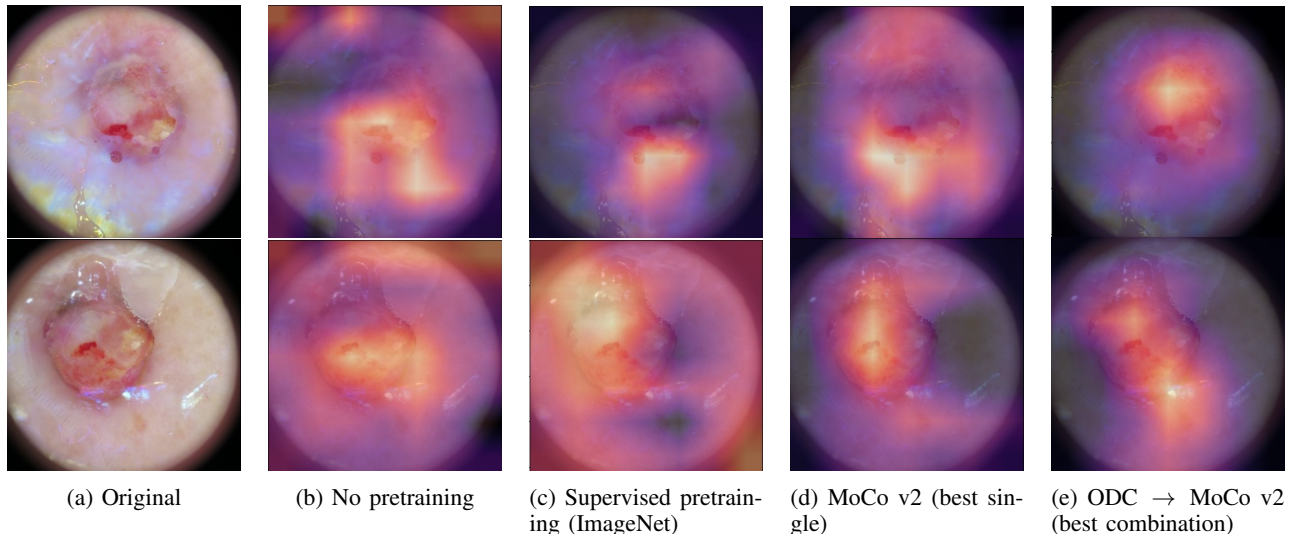
Fig. 5: Class activation maps of classifiers with and without pretraining on pretext tasks. In the shown samples, classifiers that had no pretraining tend to focus on irrelevant parts of the images (black surrounding areas) and incorrectly classify skin lesions.

(a) Original   (b) ODC   (c) MoCo v2   (d) MoCo v2 → ODC (anti-curriculum)   (e) ODC → MoCo v2 (curriculum)
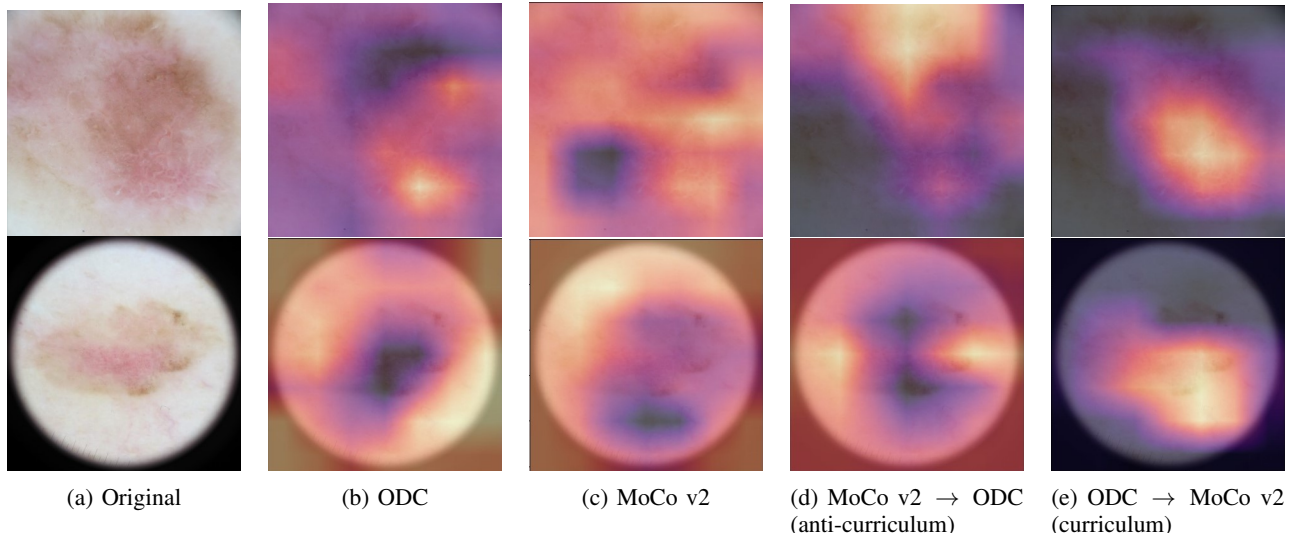
Fig. 6: Class activation maps of classifiers pre-trained with ODC and/or MoCo v2. In the presented samples only the curriculum combination of pretext tasks correctly recognizes type of the lesion.

Figures 6 and 7 only the curriculum ordering resulted in the correct classification of skin lesions.

D. Computational complexity of multiple pretext task pretraining

In this subsection we demonstrate that our proposed SSL-pretraining approach is computationally efficient and requires much less computational effort than typical supervised pretraining on ImageNet. We compare the computational complexity of different pretraining setups by comparing the number of iterations each setup requires to achieve the results indicated in Table I, assuming that the impact on complexity associated with the differences in model architecture for different tasks (SSL or ImageNet-supervised heads) is negligible.

TABLE III: Hyperparameters used in the training of SSL and supervised models.

	Training epochs	Batch size	Learning rate (LR)	LR policy
Relative location	70	64	0.2	step [30; 50]
MoCo v2	200	32	0.03	Cosine Annealing
ODC	200	100	0.06	step [170]
ImageNet pretraining	90	256	0.1	step [30; 60; 90]
Classifier	90	100	0.1	step [30; 60; 80], warm-up: 10 epochs warm-up ratio: 0.0001

The first three rows of the table present the hyperparameters used during pretraining by the SSL models. The last two rows show the hyperparameters used during the ImageNet-pretraining and the hyperparameters used during transferring to the target task.

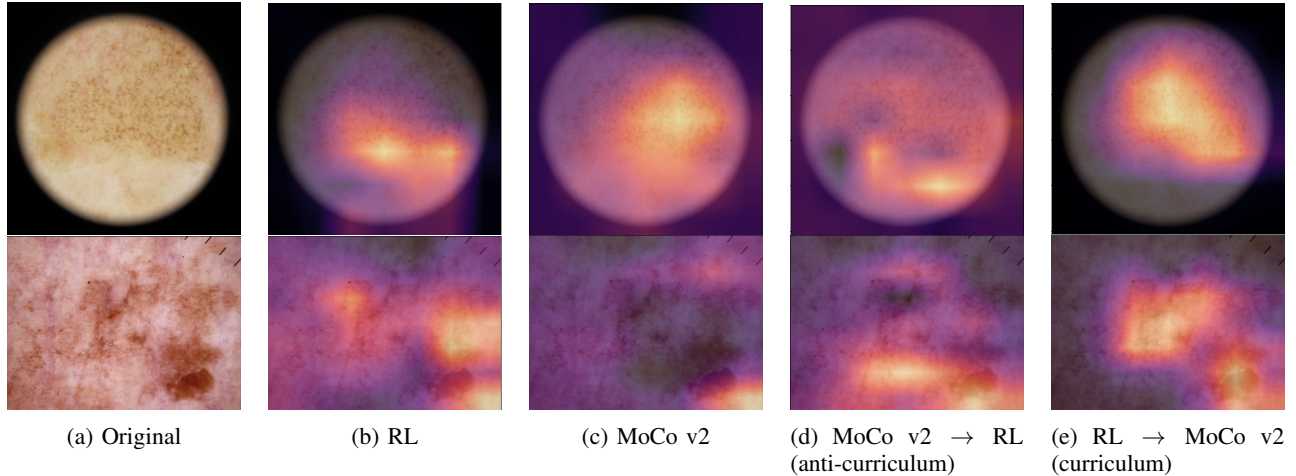| (a) Original | (b) RL | (c) MoCo v2 | (d) MoCo v2 → RL (anti-curriculum) | (e) RL → MoCo v2 (curriculum) |

Fig. 7: Class activation maps of classifiers pre-trained with Relative Location and/or MoCo v2. In the presented samples only the curriculum combination of pretext tasks correctly recognizes type of the lesion. RL stands for Relative Location.

The number of iterations required to complete the training is expressed as:

$$\mathcal{I} = N \cdot E / B, \qquad (1)$$

where N is the number of training instances, E is the number of training epochs and B is the batch size. Therefore, for equal batch sizes, Relative Location requires $\mathcal{I}_{RL} = 7125$ iterations to train on 20264 instances (80% of ISIC-2019), $\mathcal{I}_{ODC} = 15832$ iterations and $\mathcal{I}_{MoCo\_v2} = 15832$ iterations. On the other hand, supervised pretraining on ImageNet with the same batch size takes $\mathcal{I}_{ImageNet} = 421875$ iterations. Finally, transferring from an SSL model to the skin-lesion recognition task requires $\mathcal{I}_{transfer} = 7125$ iterations.

Pretraining a model on a sequence of curriculum ordered pretext tasks requires to:

1) Train all k individual pretext tasks and perform single-source transfer learning from them to establish the curriculum.
2) Train $k-1$ pretext tasks in the curriculum (the first one was already trained from scratch at the step one).

Hence, the number of required iterations required:

$$\mathcal{I}_{curriculum} = \mathcal{I}_{task_1} + 2 \cdot \sum_{k=2}^{K} \mathcal{I}_{task_k} + K \cdot (\mathcal{I}_{transfer}) \quad (2)$$

In our case, $k = 3$ (full curriculum). Thus, $\mathcal{I}_{curriculum} = \mathcal{I}_{ODC} + 2 \cdot \mathcal{I}_{RL} + 2 \cdot \mathcal{I}_{MoCo\_v2} + 3 \cdot \mathcal{I}_{transfer} = 83121$, at most.

Therefore, the complexity of our full-curriculum approach $\mathcal{I}_{curriculum}$ is almost five times lower than time complexity of supervised ImageNet pretraining $I_{ImageNet} = 421875$.

This complexity reduction is especially important for novel architectures that have not yet been pretrained on ImageNet, such as architectures that are tailor-designed, or automatically learned via Neural Architecture Search [55], to target a specific problem.

V. DISCUSSION

It is common ground that transfer learing from a pretrained model usually leads to faster convergence and better performance than training from scratch on natural images [56] or skin lesion images [52]. This hypothesis is also confirmed for the models pretrained with SSL pretext tasks, as shown in Section IV.

To our knowledge, we are the first studying the impact of the ordering of multiple sequential SSL tasks on overall performance gains: the full curriculum setup outperforms all three individual pretext tasks, while the full anti-curriculum setup sits between the two worst-performing single tasks. Interestingly, the highest performance results are obtained by a combination of just two pretext tasks (ODC and MoCo v2) in a curriculum ordering, instead of all three.

We hypothesize on the reasons behind the effects of different pretext task orderings in the following way; the fact that a pretext task leads to relatively low accuracy after transfer learning, may indicate that only basic pretext task specific features and representations have been learned. Differently, a pretext task resulting in higher downstream performances may be benefiting from the learning of generic features and representations. For this reason, if one of these last *generically*-trained pretext tasks is followed by a *specifically*-trained pretext task, the representations learned at the first stage may be harmed, as they will be countermanded to adapt for solving the specific pretext objective of the second task, loosing its generalization ability. Otherwise, if two pretext tasks result in generic features, their consecutive training is expected to either maintain or improve their individual performance. This reasoning allows to explain why curriculum (and also mixed curriculum) orderings perform better than anti-curriculum ones.

A pair of pretext tasks ("Best curriculum" on Figure 3) outperforms a triplet of pretext tasks ("Full curriculum" on Figure 3), despite containing the same pretext tasks, ordered in the same way. Explicitly, the use of the previous specific-generic hypothesis comparing the best and full curriculum

sequences would imply that the first (ODC) and the last (MoCo v2) tasks produce more effective and generic representations than the second task (Relative Location). Under this hypothesis, the presence of Relative Location might harm the overall performance, as MoCO v2 will not improve on the representations of ODC, but will try to *recover* from the representations of Relative Location. However, Relative Location yields a higher downstream performance than ODC. Our hypothesis is that the training directions of ODC and MoCo v2 features are aligned, i.e, they both follow similar training paths, whereas intermediate Relative Location features in the full curriculum somehow *shift* the feature space.

In order to test this hypothesis we conduct the following experiment: using the Central Kernel Alignment (CKA) [57] to measure the similarity of representations at different network layers, we compare features learned by the pairs of individual pretext tasks as shown on Figure 8. In accordance with our hypothesis: ODC and MoCo v2 produce very similar representations (the CKA values are close to one for all layers except the final average pooling layer), while the representations of Relative Location are very dissimilar from those of ODC and MoCo v2. This is also a strong evidence of the multi-path training idea, i.e. the fact that two models reaching similar downstream accuracy (as ODC and Relative Location) are relying on dissimilar representations.

This discussion also sheds light on the limitations of our approach: while the full curriculum ordering (where all pretext tasks are used) is better than any anti-curriculum and mixed curriculum orderings using all three tasks, it is not always the best option (which also hints that a simple increase of the number pretext tasks in the pretraining pipeline might not lead to better performance). However, we demonstrate that while it is not necessarily the best option, it shows competitive results and in our experiments outperforms the baselines. In this regard, as a direction for the future work, we suggest to incorporate the CKA analysis (alongside individual pretext tasks accuracy) to take into account the relationships between the learned representations of the individual tasks when establishing the curriculum ordering.

VI. CONCLUSIONS

We explore the use of a sequential self-supervised pretraining pipeline to improve the accuracy of skin lesion classification from dermatoscopic images. We show that a model pretrained using a sequence of self-supervised pretext tasks on the target skin lesion dataset outperform models pretrained with individual pretext tasks, and a model pretrained on the ImageNet dataset in a supervised way. This approach also outperforms other top-performing solutions to the multi-class skin lesion recognition problem on the ISIC-2019 dataset, such as model ensembles. The qualitative analysis of the results shows that this strategy helps to better tune the focus of the model on the relevant parts of the image—skin lesion and its borders.

The experimental results also show that the order of the pretext tasks in the pretraining stage has a crucial effect on the accuracy on the downstream task, and that effective curriculum orderings in the pretext tasks correlate with increasing
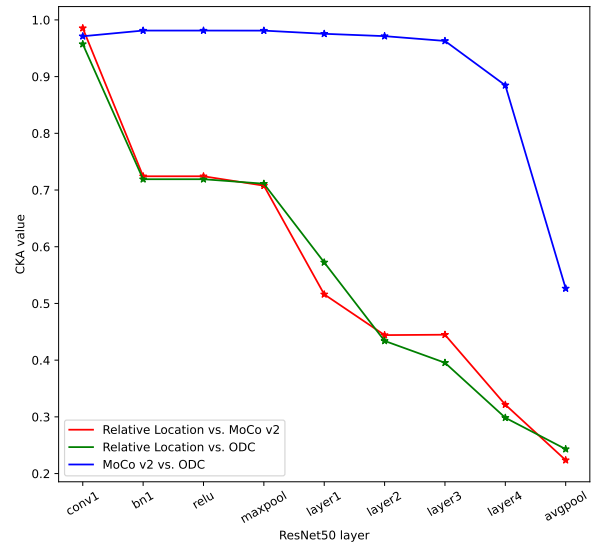

Fig. 8: Similarities of feature maps learned by pretext tasks per layer. The similarity of feature maps is measured with Central Kernel Alignment (CKA): values of CKA closer to 1 indicate strong similarity of feature maps, while values closer to 0 indicate that they are dissimilar.

downstream accuracies obtained for individual pretext task pretraining. However, results also suggest that the optimal ordering are not solely related to the individual pretext task accuracy, but also the similarity in the representations obtained from different pretext tasks. Further work will explore the use of representation similarity metrics to establish optimal curriculum orderings in an automatic manner.

Furthermore, our approach presents strong benefits for the computational complexity of the pretraining step, with respect to the ImageNet counterpart. It requires only a fraction of the ImageNet training time for relatively small datasets, such as ISIC-2019. This is especially relevant for ad-hoc network architectures for which a pretrained ImageNet model may be unavailable.

REFERENCES

[1] S. K. Clinton, E. L. Giovannucci, and S. D. Hursting, "The world cancer research fund/american institute for cancer research third expert report on diet, nutrition, physical activity, and cancer: impact and future directions," *The Journal of nutrition*, vol. 150, no. 4, pp. 663–671, 2020.

[2] High-Level Expert Group on Artificial Intelligence, "Coordinated Plan on Artificial Intelligence 2021 Review," https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419, accessed: 2021-10-26.

[3] N. Gessert, M. Nielsen, M. Shaikh, R. Werner, and A. Schlaefer, "Skin lesion classification using ensembles of multi-resolution efficientnets with meta data," *MethodsX*, vol. 7, p. 100864, 2020.

[4] P. Tschandl, C. Rosendahl, and H. Kittler, "The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. 2018; sci data (5): 180161," 2018.

[5] N. C. Codella, D. Gutman, M. E. Celebi, B. Helba, M. A. Marchetti, S. W. Dusza, A. Kalloo, K. Liopyris, N. Mishra, H. Kittler *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 *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*. IEEE, 2018, pp. 168–172.

[6] M. Combalia, N. C. Codella, V. Rotemberg, B. Helba, V. Vila-plana, O. Reiter, C. Carrera, A. Barreiro, A. C. Halpern, S. Puig *et al.*, "Bcn20000: Dermoscopic lesions in the wild," *arXiv preprint arXiv:1908.02288*, 2019.

[7] N. Ketkar, "Stochastic gradient descent," in *Deep learning with Python*. Springer, 2017, pp. 113–132.

[8] J. G. A. Barbedo, "Impact of dataset size and variety on the effectiveness of deep learning and transfer learning for plant disease classification," *Computers and Electronics in Agriculture*, vol. 153, pp. 46–53, 2018.

[9] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2097–2106.

[10] A. G. Pacheco, G. R. Lima, A. S. Salomão, B. Krohling, I. P. Biral, G. G. de Angelo, F. C. Alves Jr, J. G. Esgario, A. C. Simora, P. B. Castro *et al.*, "Pad-ufes-20: A skin lesion dataset composed of patient data and clinical images collected from smartphones," *Data in brief*, vol. 32, p. 106221, 2020.

[11] B. Albertina, M. Watson, C. Holback, R. Jarosz, S. Kirk, Y. Lee, and J. Lemmerman, "Radiology data from the cancer genome atlas lung adenocarcinoma [tcga-luad] collection," *The Cancer Imaging Archive*, 2016.

[12] S. Sudirman, A. Al Kafri, F. Natalia, H. Meidia, N. Afriliana, W. Al-Rashdan, M. Bashtawi, and M. Al-Jumaily, "Label image ground truth data for lumbar spine MRI dataset," *Mendeley Data*, 2019.

[13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, pp. 1097–1105, 2012.

[14] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Advances in Neural Information Processing Systems*, vol. 27, 2014.

[15] N. Komodakis and S. Gidaris, "Unsupervised representation learning by predicting image rotations," in *International Conference on Learning Representations (ICLR)*, 2018.

[16] R. Zhang, P. Isola, and A. A. Efros, "Colorful image colorization," in *European conference on computer vision*, 2016.

[17] M. Noroozi and P. Favaro, "Unsupervised learning of visual representations by solving jigsaw puzzles," in *European conference on computer vision*, 2016.

[18] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9729–9738.

[19] F. Navarro, C. Watanabe, S. Shit, A. Sekuboyina, J. C. Peeken, S. E. Combs, and B. H. Menze, "Evaluating the robustness of self-supervised learning in medical imaging," *arXiv preprint arXiv:2105.06986*, 2021.

[20] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 41–48.

[21] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, pp. 1097–1105, 2012.

[22] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.

[23] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[24] S. Zagoruyko and N. Komodakis, "Wide residual networks," *arXiv preprint arXiv:1605.07146*, 2016.

[25] T. G. Dietterich, "Ensemble methods in machine learning," in *Multiple Classifier Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 1–15.

[26] Y. Cao, T. A. Geddes, J. Y. H. Yang, and P. Yang, "Ensemble deep learning in bioinformatics," *Nature Machine Intelligence*, vol. 2, no. 9, pp. 500–508, 2020.

[27] T. D. Tô, D. T. Lan, T. T. H. Nguyen, T. T. N. Nguyen, H.-P. Nguyen, L. Phuong, and T. Z. Nguyen, "Ensembled skin cancer classification (ISIC-2019 challenge submission)," *ISIC Challenge*, 2019.

[28] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.

[29] S. Zhou, Y. Zhuang, and R. Meng, "Multi-category skin lesion diagnosis using dermoscopy images and deep cnn ensembles," *ISIC Challenge*, 2019.

[30] F. Pollastri, J. Maroñas *et al.*, "Aimagelab-prhlt at isic challenge 2019," *AImageLab, Tech. Rep*, 2019.

[31] C. Doersch, A. Gupta, and A. A. Efros, "Unsupervised visual representation learning by context prediction," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1422–1430.

[32] M. Caron, P. Bojanowski, A. Joulin, and M. Douze, "Deep clustering for unsupervised learning of visual features," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 132–149.

[33] X. Zhan, J. Xie, Z. Liu, Y.-S. Ong, and C. C. Loy, "Online deep clustering for unsupervised representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6688–6697.

[34] X. Yan, I. Misra, A. Gupta, D. Ghadiyaram, and D. Mahajan, "Clusterfit: Improving generalization of visual representations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6509–6518.

[35] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," *arXiv preprint arXiv:2003.04297*, 2020.

[36] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar, B. Piot, k. kavukcuoglu, R. Munos, and M. Valko, "Bootstrap your own latent - a new approach to self-supervised learning," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 21 271–21 284.

[37] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, "Unsupervised feature learning via non-parametric instance discrimination," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3733–3742.

[38] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.

[39] X. Chen*, S. Xie*, and K. He, "An empirical study of training self-supervised vision transformers," *arXiv preprint arXiv:2104.02057*, 2021.

[40] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. E. Hinton, "Big self-supervised models are strong semi-supervised learners," in *Advances in Neural Information Processing Systems*, vol. 33, 2020.

[41] D. Wang, N. Pang, Y. Wang, and H. Zhao, "Unlabeled skin lesion classification by self-supervised topology clustering network," *Biomedical Signal Processing and Control*, vol. 66, p. 102428, 2021.

[42] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.

[43] U. Brandes, D. Delling, M. Gaertler, R. Gorke, M. Hoefer, Z. Nikoloski, and D. Wagner, "On modularity clustering," *IEEE transactions on knowledge and data engineering*, vol. 20, no. 2, pp. 172–188, 2007.

[44] L. Chaves, A. Bissoto, E. Valle, and S. Avila, "An evaluation of self-supervised pre-training for skin-lesion analysis," *arXiv preprint arXiv:2106.09229*, 2021.

[45] Y. Tian, C. Sun, B. Poole, D. Krishnan, C. Schmid, and P. Isola, "What makes for good views for contrastive learning?" in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 6827–6839.

[46] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," in *Advances in Neural Information Processing Systems*, vol. 33, 2020.

[47] S. Azizi, B. Mustafa, F. Ryan, Z. Beaver, J. Freyberg, J. Deaton, A. Loh, A. Karthikesalingam, S. Kornblith, T. Chen *et al.*, "Big self-supervised models advance medical image classification," *arXiv preprint arXiv:2101.05224*, 2021.

[48] J. L. Elman, "Learning and development in neural networks: The importance of starting small," *Cognition*, vol. 48, no. 1, pp. 71–99, 1993.

[49] K. A. Krueger and P. Dayan, "Flexible shaping: How learning in small steps helps," *Cognition*, vol. 110, no. 3, pp. 380–394, 2009.

[50] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2020.

[51] I. Misra and L. v. d. Maaten, "Self-supervised learning of pretext-invariant representations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6707–6717.

[52] A. Kwasigroch, M. Grochowski, and A. Mikołajczyk, "Self-supervised learning to increase the performance of skin lesion classification," *Electronics*, vol. 9, no. 11, p. 1930, 2020.

[53] N. C. Codella, D. Gutman, M. E. Celebi, B. Helba, M. A. Marchetti, S. W. Dusza, A. Kalloo, K. Liopyris, N. Mishra, H. Kittler *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 *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*. IEEE, 2018, pp. 168–172.

[54] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2921–2929.

[55] T. Elsken, J. H. Metzen, and F. Hutter, "Neural architecture search: A survey," *The Journal of Machine Learning Research*, vol. 20, no. 1, pp. 1997–2017, 2019.

[56] E. S. Olivas, J. D. M. Guerrero, M. Martinez-Sober, J. R. Magdalena-Benedito, L. Serrano *et al.*, *Handbook of research on machine learning applications and trends: Algorithms, methods, and techniques: Algorithms, methods, and techniques*. IGI Global, 2009.

[57] S. Kornblith, M. Norouzi, H. Lee, and G. Hinton, "Similarity of neural network representations revisited," in *International Conference on Machine Learning*. PMLR, 2019, pp. 3519–3529.