

Self-Supervised Learning of Remote Sensing Scene Representations Using Contrastive Multiview Coding

Vladan Stojnić Vladimir Risojević

Faculty of Electrical Engineering, University of Banja Luka, Bosnia and Herzegovina

{vladan.stojnic, vladimir.risojevic}@etf.unibl.org

Abstract

In recent years self-supervised learning has emerged as a promising candidate for unsupervised representation learning. In the visual domain its applications are mostly studied in the context of images of natural scenes. However, its applicability is especially interesting in specific areas, like remote sensing and medicine, where it is hard to obtain huge amounts of labeled data. In this work, we conduct an extensive analysis of the applicability of self-supervised learning in remote sensing image classification. We analyze the influence of the number and domain of images used for self-supervised pre-training on the performance on downstream tasks. We show that, for the downstream task of remote sensing image classification, using self-supervised pre-training on remote sensing images can give better results than using supervised pre-training on images of natural scenes. Besides, we also show that self-supervised pre-training can be easily extended to multispectral images producing even better results on our downstream tasks.

1. Introduction

The advent of deep learning based methods enabled rapid progress in computer vision, as can be seen by the results on many benchmark datasets like ImageNet [29] and MS-COCO [23]. However, most of these improvements came from the usage of supervised learning based methods that need huge amounts of labeled training data. This need for large labeled datasets poses a significant problem as creating them is a highly time-consuming and error-prone task. This is especially visible in highly specialized areas, like remote sensing and medicine, where data labeling can only be done by domain experts.

Over the years transfer learning emerged as a standard technique for overcoming this need for big labeled datasets. It is based on the idea of using neural networks pre-trained on already available labeled datasets, like ImageNet, as either feature extractors or for finetuning on small datasets.

This allowed usage of these supervised deep learning based methods for small datasets, but some limitations remained. One major limitation in applying these techniques for remote sensing images is that most remote sensing images contain information outside of the visible spectrum, while the neural networks are usually pre-trained on RGB images as most of the big labeled datasets contain only RGB images. This forces us to drop all information outside of the visible spectrum if we want to use a neural network pre-trained on RGB images, which significantly limits the amount of information we use in the analysis. Even if remote sensing images only contain RGB color channels, a possible drawback of using models pre-trained on ImageNet is that images of natural scenes present in ImageNet look quite different from remote sensing images, especially if these are satellite images. Besides, different missions like Sentinel, Landsat, and others provide us with an ever-growing amount of unlabeled multispectral images which makes it even more interesting to study the possibilities of learning visual representations in an unsupervised manner.

Recently, self-supervised learning [20, 10, 17] emerged as a promising candidate for overcoming the need for large labeled datasets in representation learning. It is based on the idea of training a neural network to solve the pretext task for which labels can be obtained directly from the data. This pretext task should be complex enough to allow the neural network to learn meaningful image representations. A neural network trained in a self-supervised fashion can then be used for transfer learning in the same way as we would use neural networks pre-trained on supervised tasks. However, as this is an unsupervised technique of pre-training it allows us to take advantage of a large number of available unlabeled images. This possibility of using unlabeled images for representation learning is especially interesting in areas where we have multispectral images or images that simply look differently than images of natural scenes. In these areas using in-domain images for pre-training could allow us to overcome the issues that we have with classical transfer learning.

Applications of self-supervised learning, in working with images of natural scenes, have been studied quite well in the last few years and they show very promising results. In the last year, self-supervised learning provided representations that transfer to classification, segmentation, and detection tasks even better than representations learned by supervised learning [5, 4]. However, the application of self-supervised learning methods in remote sensing has not been studied a lot. Most of these applications either used small amounts, up to 50,000 images, of unlabeled training data [31, 30, 35] or tested learned representations on a small number of datasets [1, 19, 38].

For these reasons, in this work we give a detailed analysis of the applicability of self-supervised learning in remote sensing. First, we analyze the influence of the number and domain of images used for self-supervised pre-training on the performance on downstream tasks. Second, we compare representations learned by self-supervised learning to the representations learned by supervised learning. Third, we analyze different strategies for expanding a self-supervised algorithm proposed for RGB images to the case of multispectral images. Fourth, we analyze the possibility of finetuning these self-supervised models and the influence it has on results on downstream tasks. Finetuning is usually skipped in most self-supervised learning papers and these models are usually only used as feature extractors. Besides that, we compared these results to the results obtained by finetuning supervisedly trained models.

We show that, for the classification of remote sensing images, using self-supervised pre-training on remote sensing images can give better results than using supervised pre-training on images of natural scenes, even when we use significantly smaller number of images. We also apply self-supervised pre-training on multispectral images and show that it further improves the results on downstream tasks. Code and trained models used in this paper are available at <https://github.com/vladan-stojnic/CMC-RSSR>.

2. Related Work

2.1. Self-supervised learning of visual representations

Self-supervised learning of visual representations is a fast-growing subfield of unsupervised learning and, over the years, many methods have been proposed and applied. All the proposed methods can be categorized into three groups depending on the type of pretext task they use.

In the first group, we have methods based on the pretext task of image reconstruction or solving some natural tasks. In these pretext tasks, self-supervised models are trained to remove a degradation artificially added to images or to solve some simple natural tasks in image understanding. For example, in [7] a pretext task is formulated in a

way that a model needs to predict how one block of the image is positioned compared to another block of the same image. A similar pretext task is presented in [24] where a model is trained to solve a jigsaw puzzle consisting of nine shuffled image blocks. In [27] the authors propose a pretext task based on solving the image inpainting problem. Another interesting pretext task is proposed in [9] where a model learns to predict the natural orientation of the image. A pretext task based on coloring an image is presented in [21, 43] while in [44] authors propose a pretext task where they train a model to predict both the colors of the image from its luminance as well as predict the luminance from its colors. The methods from this group perform generally worse than the methods from the other two groups. Moreover, some of these methods, such as image inpainting, as well as predicting natural orientations or relative positions of image blocks cannot be unambiguously defined for remote sensing images, which makes them inapplicable to this domain.

The second group of methods for construction of pretext tasks is based on the idea of predicting synthetically generated labels. In this case self-supervised models solve a classical classification problem, but instead of predicting human supplied labels they predict labels that are automatically generated. In [8] every image is its own class and other samples from that class are generated by applying different augmentations to the original image. Another way of generating labels is proposed in [3]. This time classes are determined by applying k-means clustering to the representations extracted from a self-supervised model. New classes are determined at the start of each epoch, while each image is assigned a label that corresponds to the index of the closest cluster.

The third group of self-supervised methods is based on the idea of contrastive learning. In the case of contrastive learning, the pretext task is defined in a way that representations of similar (positive) images should be close in the representation space, while the representations of dissimilar (negative) images should be as far as possible. Now the main task is how to find similar/dissimilar images. In [25, 15] the authors propose to use patches extracted from the same image as positive pairs and patches from different images as negative pairs. Another method for finding positive pairs is based on the idea of applying different augmentations to the same image and using them as positive pairs [5, 39, 42, 12], while the negative pairs again come from different images. A method based on using representations from different neural network layers as positive pairs is proposed in [16, 2]. In [36] the authors propose to use disjoint subsets of image channels as positive pairs while they use subsets of channels coming from different images as negative pairs. Another method is proposed in [4, 22]. This time authors propose to contrast image representations

Dataset	Number of images	Image sizes	Spatial resolution (m)
ImageNet	1,281,167	up to 4488x7056	-
NWPU-RESISC45	31,500	256x256	0.2 - 30
DOTA	2,806	800x800 - 4000x4000	not specified
NWPU VHR-10	800	381x601 - 1028x1728	0.08 - 2
BigEarthNet	269,695	up to 120x120	10 - 60

Table 1. Details of datasets used for self-supervised learning. (Number of images denotes number of images used for self-supervised learning and not the number of images in a dataset.)

with their cluster assignments. The process is quite similar to the previously mentioned methods that use clustering for generating synthetic image labels. They calculate representations of positive/negative images but instead of directly contrasting representations, they cluster the representations and contrast them with cluster assignments of their positive/negative pairs.

2.2. Self-supervised learning in remote sensing

Lately, some research has been done in the area of applying self-supervised learning to the analysis of remote sensing images. In [31, 30] the authors analyzed the possibilities of using split-brain autoencoder [44] in the analysis of aerial images. They analyzed the influence of the number of images used for self-supervised learning, as well as usage of different color channels, on the results obtained on the downstream task of aerial image classification. A method that is quite similar to contrastive learning with different image augmentations is proposed in [1]. However, instead of using different augmentations of the same image as a positive pair, they use satellite recordings of the same geospatial areas in different time frames. Also, besides the contrastive loss, they introduce a loss term that is based on the classification of the image geolocation. Classification of a geolocation is based on the idea of predicting artificial labels. Artificial labels are generated by k-means clustering of latitude and longitude coordinates of every image in a dataset. In [35] the authors analyze possibilities of using different self-supervised methods, specifically the ones based on image inpainting [27], context prediction [7], and contrastive learning using different augmentations. However, they only train these methods on small datasets of remote sensing images ranging from 20,000 to 30,000 images. A method proposed in [19] is based on contrastive learning with different augmentations. They train their self-supervised models on relatively big datasets with 100,000 images, but they only test them on two downstream tasks NAIP [18] and EuroSAT [14]. In [38] the authors propose a method similar to the image colorization, but they train a model to predict information from the visible part of the spectrum from the information outside of the visible spectrum. The authors trained the model using the BigEarthNet dataset [33, 34] that is quite big with 269,695 training im-

ages, but they again only tested learned representation on two downstream tasks.

In this work, we train the self-supervised algorithm on three large datasets of: (a) images of natural scenes, (b) high-resolution remote sensing images, and (c) low-resolution remote sensing images. In addition, we apply it to both RGB and multispectral remote sensing images. Finally, we evaluate learned representations on four downstream tasks of both single-label and multi-label classification problems.

3. Datasets

To perform the experiments in this paper we use two groups of datasets. One group contains datasets used for training of the self-supervised algorithm and the second group contains datasets used for the evaluation of the representations obtained using self-supervised pre-training.

3.1. Datasets used for self-supervised pre-training

To train self-supervised models we use three different datasets. One dataset consists of images of natural scenes and will be used as a baseline, while the two remaining datasets consist of remote sensing images. One of these datasets contains high-resolution images, while the other one contains low-resolution images. It is important to note that even though most of these datasets contain labels we do not use them in any way for training of our self-supervised models. Details about these datasets can be seen in Table 1.

For the dataset of natural scenes, we chose ImageNet [29] as it is a standard dataset used for training both supervised and self-supervised models in the literature. In this work, we use only the training set of 1.2M images from ImageNet1000 dataset.

As the high-resolution remote sensing dataset, we use the dataset created by combining images from NWPU-RESISC45 [6], DOTA [40], and NWPU VHR-10 [32] datasets. Some of the images in DOTA dataset contain a black border, so we manually extracted regions without it from every image in the dataset. From these regions, we then extracted non-overlapping patches of 256x256 pixels that we used as samples in our dataset for self-supervised learning. For NWPU VHR-10 dataset, we only

Dataset	Annotation type	Number of categories	Number of images	Image sizes	Spatial resolution (m)
AID	single-label	30	10,000	600x600	0.5 - 8
MLRSNet	multi-label	60	109,161	256x256	0.1 - 10
BigEarthNet	multi-label	19	519,284	up to 120x120	10 - 60
So2Sat	single-label	17	376,485	32x32	10

Table 2. Details of datasets used for the evaluation of self-supervised representations.

extracted non-overlapping patches of 256x256 pixels from the original images and use these patches as samples in our dataset. Since NWPU-RESISC45 dataset originally consists of 256x256 pixel images we use these original images as samples in our dataset. Using this procedure we created a dataset of 196,215 images that we use for training of our self-supervised algorithm. We call this dataset NWPU-DOTA throughout this paper.

For the low-resolution dataset, we use images from BigEarthNet [33, 34] dataset. This dataset consists of 519,284 multispectral images collected by Sentinel-2 satellite, but we only use images from the training set of BigEarthNet, 269,695 images in total, as the dataset for training of the self-supervised algorithm.

3.2. Datasets used for the evaluation of learned representations

To evaluate the representations learned by the self-supervised algorithm we use four remote sensing datasets: two with high-resolution and two with low-resolution images. Details about these datasets can be seen in Table 2.

AID [41] is a high-resolution scene classification dataset of aerial images. Every image in this dataset is labeled with one category from the set of 30 predefined categories. On this dataset, we use 50% - 50% training/test split for training and evaluation of our models, as previously used in literature.

MLRSNet [28] is a dataset of high-resolution aerial images. This dataset was developed for the problem of scene classification in a multi-label setting. Every image in this dataset is assigned 1 to 13 labels from the set of 60 predefined categories. For training and evaluation of our models, we use a 20%-80% training/test split as suggested in [28].

BigEarthNet is a dataset of multispectral satellite images for multi-label land use classification. There are two versions of this dataset with 45 [33] and 19 [34] predefined categories, but in this work, we only consider a version of the dataset with 19 categories. For training/validation/test split we use the same split defined in [34] with 269,695, 123,723, and 125,866 images, respectively.

So2Sat [45] represents a dataset of multispectral satellite images collected by Sentinel-2 and SAR images collected by Sentinel-1 designed for local climate zones classification. Every image in this dataset is labeled with one category from the set of 17 predefined categories. In this

work, we use only multispectral images and the official training/validation split with 352,366 and 24,119 images, respectively.

4. Method

The method we use in this work consists of two steps: (1) self-supervised pre-training and (2) transfer learning and evaluation of learned representations.

4.1. Self-supervised pre-training

As the first step in our experiments, we use Contrastive Multiview Coding - CMC [36] for self-supervised pre-training. We chose CMC because it gave good results for images of natural scenes and could be readily adapted to multispectral images. We believe that similar results would be obtained by using the majority of the methods from the second and third groups in Section 2.

CMC is based on the idea of contrasting representations of different image views. More specifically, given a set of images $\{x_1, x_2, \dots, x_N\}$, each image x_i is split into two views, v_i^1 and v_i^2 , containing different subsets of channels. Each of these views is then processed using encoding, $z_i^1 = E_1(v_i^1)$, $z_i^2 = E_2(v_i^2)$, and projection, $h_i^1 = P_1(z_i^1)$, $h_i^2 = P_2(z_i^2)$, networks that output data and contrastive representations, respectively. The model is then trained to pull together contrastive representations of views coming from the same image, h_i^1 and h_i^2 , while pushing apart contrastive representations of views from different images, h_i^1 and h_j^2 , as shown in Figure 1. This is achieved by training the encoding and projection networks to minimize the loss

$$L^{V^1, V^2} = - \mathbb{E}_{\{v_1^1, v_1^2, \dots, v_{k+1}^2\}} \left[\log \frac{s(h_1^1, h_1^2)}{\sum_{j=1}^{k+1} s(h_1^1, h_j^2)} \right], \quad (1)$$

where view V^1 is an anchor, and one positive and k negative images are sampled from view V^2 . Discriminating function $s(\cdot, \cdot)$ is cosine similarity of contrastive representations with temperature hyperparameter τ

$$s(h^1, h^2) = \exp \left(\frac{h^1 \cdot h^2}{\|h^1\| \|h^2\| \tau} \right). \quad (2)$$

Finally, the objective function for CMC is obtained by computing the losses (1) using both views as anchors and summing them

$$L = L^{V_1, V_2} + L^{V_2, V_1}. \quad (3)$$

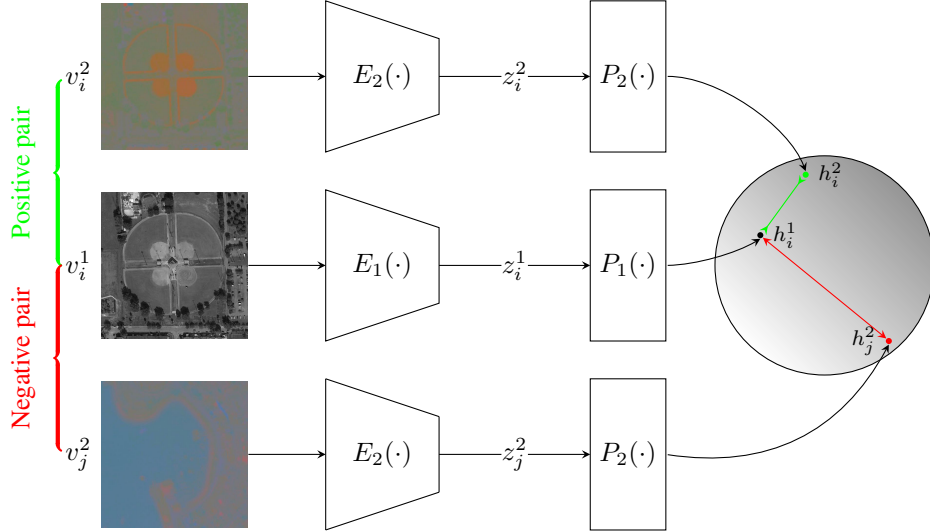


Figure 1. Schematic representation of the CMC algorithm.

Since k in (1) can be very large, Noise Contrastive Estimation [11] is used to approximately compute the loss and reduce computational requirements.

In this work, we use a ResNet50 [13] based architecture for the encoding network of the CMC model, and for the projection network we use a fully-connected linear layer followed by L2 normalization. We train our CMC models for 400 epochs using SGD with batch size 100, learning rate 0.03, momentum 0.9, and weight decay 10^{-4} . Learning rate is further reduced by a factor of 10 after the 250th, 300th and 350th epoch. We also employ data augmentation using random resized cropping and horizontal flipping. For calculating NCE based loss we use $k = 4096$ negative samples and the temperature parameter of $\tau = 0.07$.

4.2. Transfer learning

As our main goal in this paper is to evaluate representations learned using self-supervised learning, in the next step we use previously trained CMC models for transfer learning. To conduct transfer learning, the trained CMC models are used both as feature extractors and as starting points for finetuning.

For feature extraction, we pass 224x224 pixel images through the CMC model and use concatenated outputs of both encoder networks as image representation. To evaluate these representations we train a linear classifier on top of them for our testing datasets. For single-label datasets, we use softmax activation function with categorical cross-entropy loss, while for multi-label datasets we use sigmoid activation function with binary cross-entropy loss. A linear classifier is trained for 50 epochs using Adam optimizer with the batch size of 256 and weight decay 10^{-2} . The learning rate is set to 10^{-2} for MLRSNet dataset and 10^{-3}

for other datasets. This learning rate is further reduced by a factor of 5 after the 30th, 35th, 40th, and 45th epoch.

Similar to the feature extraction, for finetuning we put a linear layer on top of the concatenated outputs of both encoder networks and use the same activation and loss functions. This time we train the whole resulting model for 100 epochs, without warm-up, using Adam optimizer with the batch size of 100, learning rate of 10^{-4} and weight decay 10^{-4} . Learning rate is further reduced by a factor of 5 after the 60th, 70th, 80th and 90th epoch. In the process of training these finetuned models, we first rescale training images to 256x256 pixels and then extract 224x224 pixel blocks from them using random resized cropping. These blocks are also horizontally flipped with the probability of 0.5. In the process of evaluation of the finetuned models, we pass them input images that are rescaled to 224x224 pixels.

As evaluation metrics, we use accuracy for single-label datasets and macro-averaged mean average precision (mAP) for multi-label datasets.

5. Experimental Setup

To assess the possibilities of CMC for learning representations of remote sensing images, we perform three experiments. First, we want to analyze the influence of the number and domain of the images used for training of the CMC model on the results obtained by training a linear classifier for our downstream tasks. To do this we train CMC models using subsets of 10,000 (10k), 25,000 (25k), 50,000 (50k), 100,000 (100k), and 196,215 (200k) images from ImageNet, NWPU-DOTA, and BigEarthNet datasets. It is important to note that in these experiments we use only RGB color channels even for datasets with multispectral images. To be precise, we train our CMC models as proposed

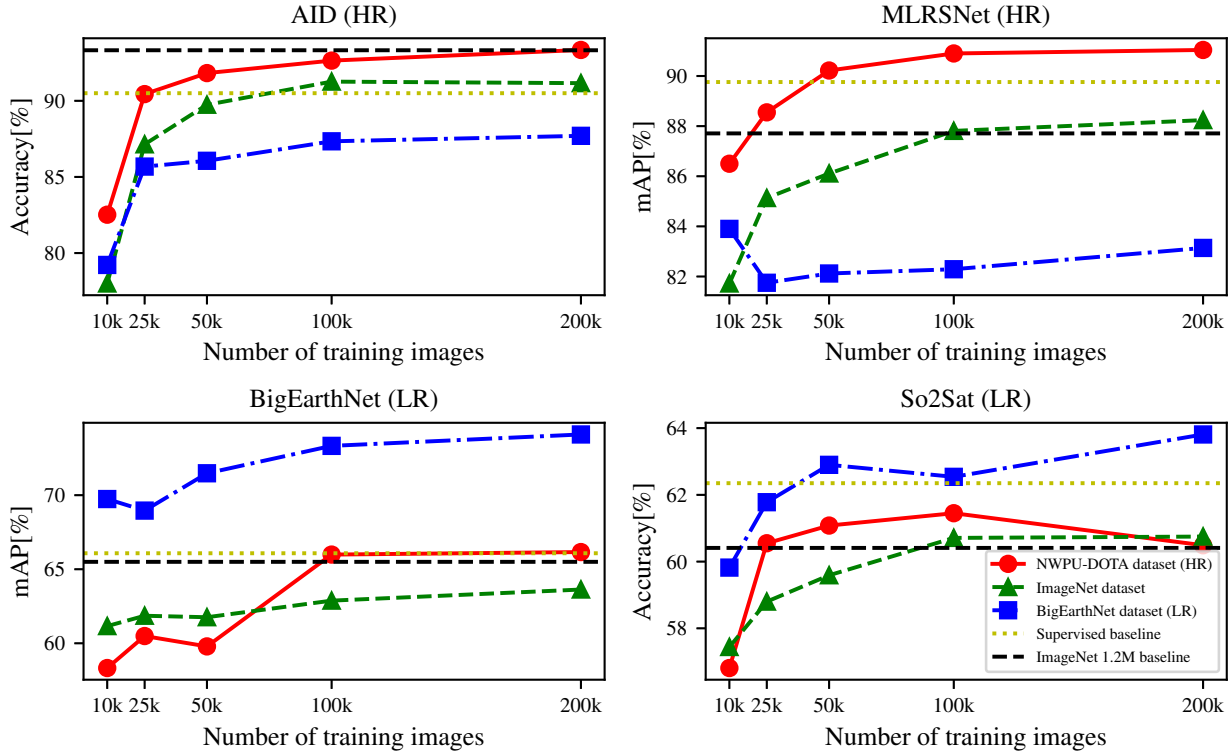


Figure 2. Results obtained using a linear classifier on top of the representations learned by supervised and self-supervised learning. Abbreviations HR and LR represent high-resolution and low-resolution remote sensing datasets, respectively.

in [36], RGB images are converted to $L^*a^*b^*$ color space and L^* component is used as one view of the image while a^*b^* components are used as the second view.

Second, we want to analyze the possibilities of using multispectral images for the training of CMC models. To accomplish this we train our CMC models on subsets of 10,000 (10k), 25,000 (25k), 50,000 (50k), 100,000 (100k), and 196,215 (200k) images from BigEarthNet dataset and use them as feature extractors on BigEarthNet and So2Sat downstream tasks. We propose two possible ways to create multispectral image views for the CMC model. Our first method is based on the idea that the first view should contain the spectral bands with short wavelengths, while the second view should contain the ones with long wavelengths. In addition, we also split visible bands into different views so that the blue color channel ends up in the view with long-wavelength bands, while red and green end up in the other view. Views created in this way are shown in Table 3 where we used only 10 spectral bands available from Sentinel-2 as proposed in [34]. The second method for defining CMC views is based on the idea that it is better to use decorrelated channels when training CMC models [36, 37]. To decorrelate previously used spectral bands we employ Principal Component Analysis - PCA. To determine the PCA decomposition we randomly sampled 144 pixel values from

View	Spectral bands
1	2, 8, 8A, 11, 12
2	3, 4, 5, 6, 7

Table 3. One possible method for defining CMC views of multi-spectral images.

each of 200k training images from BigEarthNet. After determining the PCA decomposition, original image channels are projected into the PCA basis and new CMC views are defined so that they contain approximately the same variance. To be precise, the first view contains the channel that explains the most data variance and four channels that explain the least amount of variance, while the second view contains the remaining channels.

Finally, we evaluated the possibility of using previously trained CMC models for finetuning on our downstream tasks. For this, we used only CMC models trained using 200k images from different datasets and finetuned them using the procedure presented in Section 4.2.

All experiments are implemented on a computer with two NVIDIA GeForce RTX 2080Ti GPUs and using a PyTorch [26] library. Depending on the number and type (RGB or multispectral) of images used for training of the CMC




Figure 3. Results obtained using a linear classifier on top of the representations learned by self-supervised learning on multispectral images. Supervised baseline represents a baseline obtained using supervisedly pre-trained network on ImageNet, and CMC baseline with RGB images represents the best result obtained in previous experiments.

models, training a model takes from few hours to few days. As for the transfer learning, all feature extraction models can be trained in a couple of minutes, while finetuning can take up to a few days in the case of multispectral images.

6. Experimental Results

6.1. Feature extraction with RGB images

In Figure 2 the results obtained using a linear classifier on top of the representations learned by CMC models trained on different datasets with varying sizes are shown. In the same Figure, we also show the baseline results obtained by using the representations from the CMC model trained using 1.2M ImageNet images and ResNet50 model trained in supervised fashion on ImageNet dataset, available in PyTorch [26]. As we can see, using bigger datasets for the training of CMC models, in general, gives better representations for our downstream tasks, but this can be expected from most deep learning models. More interesting is to look at the influence of the domains of images used for CMC training on the results on downstream tasks. For downstream tasks on AID and MLRSNet datasets, we can see that using NWPU-DOTA dataset for self-supervised pre-training gives better results than using ImageNet or BigEarthNet for pre-training while using ImageNet gives better results than using BigEarthNet. Interestingly, using even as low as 25k or 50k images from NWPU-DOTA gives better results than using 200k images from ImageNet and in the case of MLRSNet even better than using 1.2M images from ImageNet. Besides that, using NWPU-DOTA dataset for self-supervised pre-training can give results that are better than the results obtained using representations learned by supervised learning on ImageNet. In the case of testing on BigEarthNet and So2Sat datasets, it is visible that using BigEarthNet dataset for self-supervised pre-training gives better results than using NWPU-DOTA or ImageNet for

pre-training while using NWPU-DOTA gives better results than using ImageNet. For the case of testing on BigEarthNet, we can see that using even 10k images from BigEarthNet for self-supervised pre-training gives results that are better than any other result obtained either using supervised or self-supervised learning. As for the So2Sat, we can see that using 25k or more images from BigEarthNet, for self-supervised learning, gives results that are better than the results obtained using self-supervised learning on other datasets. Besides that, using 50k images from BigEarthNet for self-supervised learning exceeds the results obtained by using the supervisedly pre-trained network.

The aforementioned results show us that the choice of images used for the training of the CMC model is crucial to the performance on the downstream tasks if we use CMC models as feature extractors. We can see that on the downstream task of remote sensing image classification it is much better to train the CMC model on remote sensing images than on images of natural scenes. This is even more pronounced when the spatial resolutions of the images used in the pretext and downstream tasks are similar. Interestingly, if the spatial resolution between pre-training and downstream images significantly differs we get some opposed results. For the case of high-resolution downstream tasks, we get better results by pre-training on images of natural scenes than on low-resolution remote sensing images. On the other hand, when we do testing on low-resolution tasks it is better to pre-train on high-resolution remote sensing images than on images of natural scenes. We think that this discrepancy might be due to the inability of learning sharp edge detection filters when we pre-train on low-resolution images. On the other hand, if we pre-train on high-resolution remote sensing images we can learn both filters for sharp edge detection and filters that would detect features interesting for the special case of remote sensing images.

Downstream task	Finetuning initialization						
	RGB images					Multispectral images	
	NWPU-DOTA HR 200k	BigEarthNet LR 200k	ImageNet 200k	ImageNet 1.2M	ImageNet supervised	Without PCA	With PCA
AID (HR)	94.76	94.60	94.60	94.16	95.58	-	-
MLRSNet (HR)	95.49	95.02	95.31	92.59	96.43	-	-
BigEarthNet (LR)	78.34	79.56	78.63	76.21	81.62	82.00	82.90
So2Sat (LR)	58.43	58.50	58.51	57.06	60.84	60.66	60.68

Table 4. Finetuning results on our downstream tasks. Abbreviations HR and LR represent high-resolution and low-resolution remote sensing datasets, respectively.

6.2. Feature extraction with multispectral images

In Figure 3, we can see the results obtained using a linear classifier on top of the representations learned by self-supervised learning on multispectral images. As we would have expected, using bigger datasets for self-supervised pre-training gives better results on downstream tasks. Besides, we can see that on the BigEarthNet downstream task using PCA-based views turns out to be better by around 2% than using non-decorrelated views. On the other hand, testing on the So2Sat dataset shows that better results are obtained if we use non-decorrelated views, but differences in results are in general much smaller. We think that this difference might come from the fact that the PCA decomposition was learned using images from the BigEarthNet and the learned basis may be too specialized to that dataset.

If we compare the results obtained using multispectral images with the ones obtained using only RGB images we can see that using multispectral information boosted the results by 2% to 4%.

6.3. Finetuning

Results obtained using finetuning are shown in Table 4. We can see that in the case of RGB images the best results are obtained using neural networks pre-trained in a supervised fashion. These models show around 1% to 2% better results than the models pre-trained in a self-supervised manner. It is also interesting that all self-supervised models give almost the same results with only a slight favor to the models trained with images that have similar resolution as the images in the downstream task. When we use multispectral images, finetuning a self-supervised network will give around 1% better results on the BigEarthNet dataset compared to the supervisedly pre-trained network and almost the same results on the So2Sat. However, it is important to note that on the So2Sat finetuned models give worse results than linear classifiers. We have observed a significant overfitting problem on the So2Sat using both linear and finetuned models that we were not able to solve using architectures and hyperparameters used in this work.

7. Conclusion

In this paper, we evaluate the applicability of self-supervised learning in remote sensing image classification. We show that if we use pre-trained models as feature extractors, self-supervised models trained on RGB remote sensing images give better results than supervised models trained on ImageNet. This can be seen even if the datasets used for self-supervised pre-training are 6 to 24 times smaller than the ImageNet dataset. However, if finetuning is used, the models pre-trained in a supervised fashion on ImageNet give up to 2% better results than the models trained using self-supervised learning. Nevertheless, self-supervised pre-training still has merit because finetuning of the network is time and resource consuming. Furthermore, models pre-trained on multispectral images are rare and we showed that self-supervised pre-training on multispectral images boosts the performance on downstream tasks up to 4% and allows us to match or surpass the results obtained by finetuning a supervised model pre-trained on RGB images.

There are still a lot of directions for future research in this area. In this paper we focused on scene classification, but it would be important to investigate whether similar results would be obtained for different downstream tasks, such as segmentation or object detection. Next, new methods for creating views for multispectral images are needed, since PCA-based views fail to generalize to different downstream tasks. It would also be interesting to train self-supervised methods on huge amounts of unlabeled data to find the limits of these methods. Finally, further investigation why supervisedly trained models give better results with finetuning is necessary.

Acknowledgements This research was funded in part by the Ministry of Scientific and Technological Development, Higher Education and Information Society of the Republic of Srpska under contract 19.032/961-102/19, “Service for Classification of Remote Sensing Images”.

References

- [1] Kumar Ayush, Burak Uzcent, Chenlin Meng, Marshall Burke, David Lobell, and Stefano Ermon. Geography-aware self-supervised learning. *arXiv preprint arXiv:2011.09980*, 2020. 2, 3
- [2] Philip Bachman, R. Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *NeurIPS*, 2019. 2
- [3] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. 2
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 2
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2
- [6] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017. 3
- [7] Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 2, 3
- [8] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin A. Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(9):1734–1747, 2016. 2
- [9] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018. 2
- [10] Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. In *ICCV*, 2019. 1
- [11] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, 2010. 5
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 2
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [14] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.*, 12(7):2217–2226, 2019. 3
- [15] Olivier J. Héroux. Data-efficient image recognition with contrastive predictive coding. In *ICML*, 2020. 2
- [16] R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *ICLR*, 2019. 2
- [17] Ashish Jainwal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1), 2021. 1
- [18] Neal Jean, Sherrie Wang, Anshul Samar, George Azzari, David B. Lobell, and Stefano Ermon. Tile2vec: Unsupervised representation learning for spatially distributed data. In *AAAI*, 2019. 3
- [19] Jian Kang, Rubén Fernández-Beltrán, Puhong Duan, Sicong Liu, and Antonio J. Plaza. Deep unsupervised embedding for remotely sensed images based on spatially augmented momentum contrast. *IEEE Trans. Geosci. Remote. Sens.*, 59(3):2598–2610, 2021. 2, 3
- [20] Alexander Kolesnikov, Xiaohua Zhai, and Lucas Beyer. Revisiting self-supervised visual representation learning. In *CVPR*, 2019. 1
- [21] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a proxy task for visual understanding. In *CVPR*, 2017. 2
- [22] Junnan Li, Pan Zhou, Caoming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *ICLR*, 2021. 2
- [23] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, 2014. 1
- [24] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 2
- [25] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2
- [26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019. 6, 7
- [27] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 2, 3
- [28] Xiaoman Qi, Panpan Zhu, Yuebin Wang, Liqiang Zhang, Junhuan Peng, Mengfan Wu, Jialong Chen, Xudong Zhao, Ning Zang, and P. Takis Mathiopoulos. Mlrsnet: A multi-label high spatial resolution remote sensing dataset for semantic scene understanding. *ISPRS J. Photogramm. and Remote. Sens.*, 169:337–350, 2020. 4
- [29] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, 2015. 1, 3
- [30] Vladan Stojnić and Vladimir Risojević. Analysis of color space quantization in split-brain autoencoder for remote sensing image classification. In *NEUREL*, 2018. 2, 3
- [31] Vladan Stojnić and Vladimir Risojević. Evaluation of split-brain autoencoders for high-resolution remote sensing scene classification. In *ELMAR*, 2018. 2, 3

- [32] Hao Su, Shunjun Wei, Min Yan, Chen Wang, Jun Shi, and Xiaoling Zhang. Object detection and instance segmentation in remote sensing imagery based on precise mask R-CNN. In *IGARSS*, 2019. 3
- [33] Gencer Sumbul, Marcela Charfuelan, Begüm Demir, and Volker Markl. Bigearthnet: A large-scale benchmark archive for remote sensing image understanding. In *IGARSS*, 2019. 3, 4
- [34] Gencer Sumbul, Jian Kang, Tristan Kreuziger, Filipe Marcelino, Hugo Costa, Pedro Benevides, Mario Caetano, and Begüm Demir. Bigearthnet dataset with a new class-nomenclature for remote sensing image understanding. *arXiv preprint arXiv:2001.06372*, 2020. 3, 4, 6
- [35] Chao Tao, Ji Qi, Weipeng Lu, Hao Wang, and Haifeng Li. Remote sensing image scene classification with self-supervised paradigm under limited labeled samples. *IEEE Geosci. Remote. Sens. Lett.*, 2020. 2, 3
- [36] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, 2020. 2, 4, 6
- [37] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? In *NeurIPS*, 2020. 6
- [38] Stefano Vincenzi, Angelo Porrello, Pietro Buzzega, Marco Cipriano, Pietro Fronte, Roberto Cucu, Carla Ippoliti, Annamaria Conte, and Simone Calderara. The color out of space: learning self-supervised representations for earth observation imagery. In *ICPR*, 2020. 2, 3
- [39] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 2
- [40] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge J. Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. DOTA: A large-scale dataset for object detection in aerial images. In *CVPR*, 2018. 3
- [41] Gui-Song Xia, Jingwen Hu, Fan Hu, Baoguang Shi, Xiang Bai, Yanfei Zhong, Liangpei Zhang, and Xiaoqiang Lu. AID: A benchmark data set for performance evaluation of aerial scene classification. *IEEE Trans. Geosci. Remote. Sens.*, 55(7):3965–3981, 2017. 4
- [42] Mang Ye, Xu Zhang, Pong C. Yuen, and Shih-Fu Chang. Unsupervised embedding learning via invariant and spreading instance feature. In *CVPR*, 2019. 2
- [43] Richard Zhang, Phillip Isola, and Alexei A. Efros. Colorful image colorization. In *ECCV*, 2016. 2
- [44] Richard Zhang, Phillip Isola, and Alexei A. Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, 2017. 2, 3
- [45] Xiao Xiang Zhu, Jingliang Hu, Chunping Qiu, Yilei Shi, Jian Kang, Lichao Mou, Hossein Bagheri, Matthias Häberle, Yuansheng Hua, Rong Huang, et al. So2sat lcz42: A benchmark dataset for global local climate zones classification. *arXiv preprint arXiv:1912.12171*, 2019. 4