

SELF-SUPERVISED VISION TRANSFORMERS FOR JOINT SAR-OPTICAL REPRESENTATION LEARNING

Yi Wang^{1,2}, Conrad M Albrecht^{1,2}, Xiao Xiang Zhu^{1,2}

¹Remote Sensing Technology Institute, German Aerospace Center (DLR), Germany

²Data Science in Earth Observation, Technical University of Munich (TUM), Germany

ABSTRACT

Self-supervised learning (SSL) has attracted much interest in remote sensing and Earth observation due to its ability to learn task-agnostic representations without human annotation. While most of the existing SSL works in remote sensing utilize ConvNet backbones and focus on a single modality, we explore the potential of vision transformers (ViTs) for joint SAR-optical representation learning. Based on DINO, a state-of-the-art SSL algorithm that distills knowledge from two augmented views of an input image, we combine SAR and optical imagery by concatenating all channels to a unified input. Subsequently, we randomly mask out channels of one modality as a data augmentation strategy. While training, the model gets fed optical-only, SAR-only, and SAR-optical image pairs learning both inner- and intra-modality representations. Experimental results employing the BigEarthNet-MM dataset demonstrate the benefits of both, the ViT backbones and the proposed multimodal SSL algorithm DINO-MM.

Index Terms— Self-supervised learning, vision transformer, multimodal representation learning

1. INTRODUCTION

Recent advances in self-supervised learning (SSL) demonstrate great success in computer vision and remote sensing. Designed to generate task-agnostic representations from large-scale, unlabeled data, SSL is attracting significant attention due to an increasing amount of openly available earth observation data. A common SSL pipeline disassembles into two major steps: (1) based on a *self-supervision* objective, a model is trained to learn high-level representations from massive amounts of unlabeled input data; (2) the *pre-trained* model from (1) is then transferred to a supervised *downstream* task exploiting its ability to capture good representations.

The history of SSL yields various designs for self-supervision: (1) traditionally, generative methods like autoencoders [1] learn representations by reconstructing input data; (2) in the past few years, predictive methods work by solving

pretext tasks such as predicting the relative position of two cropped patches [2]; (3) recently, contrastive methods utilize Siamese architectures to contrast the similarity between two augmented views of the same input. A major benefit of contrastive learning is the freedom for the model to not depend on a specific pretext task, but rather learn general representations. Yet, this can also lead to a trivial solution—the identity mapping. Various methods have been designed to avoid this *model collapse*, namely: negative sampling [3], clustering [4], knowledge distillation [5], or redundancy reduction [6]. While most of recent SSL works in remote sensing are based on negative sampling [7, 8], our approach picks DINO (self distillation with **no** labels) [9], a state-of-the-art contrastive SSL method of category knowledge distillation.

Apart from the design of self-supervision, the choice of a backbone network is a key component to capture expressive representations. ConvNets, ResNets [10] in particular, are popular backbones in computer vision and remote sensing. The recent success of vision transformers (ViTs) [11] in computer vision transfers to ViTs as SSL model backbones [9, 12, 13]. Yet, the potential of self-supervised ViTs remains unlocked in remote sensing data analytics. Our work targets to bridge this gap.

Moreover, while a typical SSL algorithm considers a single modality, real-world problems usually require approaches to jointly analyze multiple modalities. This is particularly true in earth observation where a variety of sensors capture the physical properties of Earth’s surface. Two popular modalities in remote sensing are synthetic-aperture radar (SAR) and multispectral (optical) images. Although supervised SAR-optical fusion has been extensively studied, corresponding self-supervised techniques are currently in their early stage [14, 15, 16]. Meanwhile, it is also key to adjust those multimodal algorithms to situations where only a single modality is available. Concerning these two aspects, we propose a joint SAR-optical representation learning method, DINO-MM, simultaneously able to distill knowledge from a single or both modalities. With a newly proposed RandomSensorDrop module, the model randomly receives SAR, optical, or concatenated SAR-optical images as input for contrastive learning. Therefore, the pre-trained model is flexible to infer representations from any combination of available modalities.

This work is supported by Helmholtz Association’s Initiative and Networking Fund through Helmholtz AI. Codes and pre-trained weights are available at <https://github.com/zhu-xlab/DINO-MM>.




Fig. 1: DINO-MM: the proposed joint SAR-optical SSL algorithm. The concatenated SAR-optical image is taken as raw input. It is randomly transformed into two augmented views and fed into a DINO-based teacher-student network. The transformations $T_{1/2}$ include both regular data augmentations and a *RandomSensorDrop* module that randomly masks out SAR or optical channels at a given probability. This way, the model contrasts views of either a single modality or both, learning both inner- and inter-modality representations. The abbreviation *ema* denotes *exponential moving average*, and *sg* represents *stop gradient*.

2. METHODOLOGY

2.1. Vision transformers

The vision transformer (ViT) [11] is a neural network model for image analysis. It employs a Transformer-like architecture stacking attention blocks in order to draw global dependencies of input sequences [17] over patches of an image. Unlike in natural language processing (NLP) where a sentence is a sequence, an image needs splitting into fixed-size patches. These patches are then linearly embedded to build the input sequence for a ViT. After adding position embeddings, the resulting sequence of vectors is fed to a standard Transformer encoder. For classification, a learnable *classification token* gets added to the sequence. Similar to ConvNets, ViTs may work as model backbones for SSL with the output of the last layer’s classification token identified as encoded representation vectors.

2.2. DINO-MM

Our work is based on DINO [9]—a recent contrastive SSL algorithm with proven performance when applied in conjunction with ViT backbones. We extend DINO to multimodal SSL, and reference it by DINO-MM. It adds a simple data augmentation module, *RandomSensorDrop*, that randomly masks out SAR or optical channels of the concatenated input image for joint multimodal representation learning.

2.2.1. DINO

Self-distillation with **no** labels (DINO) is a knowledge distillation-based contrastive SSL algorithm maximizing the similarity of representations spawning from augmented views of the same input image. Figure 1 depicts the model structure. An initial step generates two distorted views¹ of an image,

¹DINO follows a multi-crop strategy that invokes more than two views [4]. To illustrate, we limit the discussion to two views.

x_1 and x_2 . These two views are sent to a Siamese teacher-student network, $E_{student}$ and $E_{teacher}$. Both networks share the same architecture g with separate sets of parameters θ_s and θ_t . The neural network g builds on a backbone f (e.g. a ViT encoder) and a multi-layer perceptron (MLP) projection head h : $g = h \circ f$. The output of the teacher network gets centered by a mean computed over the batch, which can be interpreted as adding a bias term c : $g_t(x) \leftarrow g_t(x) + c$. The center c is updated with an exponential moving average:

$$c \leftarrow mc + (1 - m) \frac{1}{B} \sum_{i=1}^B g_{\theta_t}(x_i), \quad (1)$$

where m is a rate parameter, and B denotes the batch size. Each network outputs a K -dimensional feature vector normalized by a temperature softmax, $\sim \exp(z_k/\tau)$, per feature dimension z_k with $k = 1 \dots K$. In a final step, the similarity of the features from both networks is quantified by a cross-entropy loss.

DINO avoids model collapse by *centering* the teacher network outputs through c , and by *sharpening* with a low teacher softmax temperature, i.e. it holds $\tau_t < \tau_s$ for the softmax normalizations. A stop-gradient operator applies to the teacher on training to make the gradients backpropagate through the student network, only. Teacher parameters θ_t receive corresponding updates through an exponential moving average (*ema*) of the student parameters θ_s —in analogy to eq. (1). Once self-supervised pre-training converged, the network backbone f serves as a representation generator to any downstream task.

2.2.2. RandomSensorDrop

We propose *RandomSensorDrop* as an additional data augmentation module to upgrade DINO for joint SAR-optical representation learning. Consider a batch of SAR-optical image pairs with dimensions $[B, C_{SAR/optical}, W, H]$ where B

Table 1: Possible combinations of the two augmented views.

$x_1 \backslash x_2$	SAR (S)	Optical (O)	SAR-optical (M)
SAR (S)	SS	SO	SM
Optical (O)	OS	OO	OM
SAR-optical (M)	MS	MO	MM

represents batch size, and $C_{SAR/optical}$ represents the number of SAR or optical channels, respectively. For DINO-MM, we concatenate the pair of images into a raw input x with dimension $[B, C_{optical} + C_{SAR}, W, H]$. Once established, in addition to regular data augmentations of DINO not touching channel information, we add the RandomSensorDrop module. This module randomly picks one of the *three* options for the *two* views x_1 and x_2 : replace SAR, optical, or none of the channel values by zero.

The procedure outlined enables the model to digest any possible combinations of the two modalities as summarized by Table 1. Consequently, the model is expected to learn both inner- and inter-modality representations:

- x_1 and x_2 contain SAR images, only, i.e. the model learns SAR-only representations;
- x_1 and x_2 contain optical images, only, i.e. the model distills optical-only representations;
- x_1 and x_2 contain both modalities, i.e. the model learns joint SAR-optical representations.

As a result, the pre-trained model is able to infer data representations with either a single modality or with both modalities available.

3. EXPERIMENTS

3.1. Dataset and implementation details

Dataset. We evaluate DINO-MM on BigEarthNet-MM [18]—a large-scale, multi-label scene classification dataset that contains 590,326 pairs of Sentinel-1 and Sentinel-2 image patches. We split the dataset into 311,667 training pairs, 103,944 validation pairs, and 118,065 testing pairs. Patches got dropped where the scene is fully covered in seasonal snow, or those including artifacts of clouds. All patches are re-sampled to 10m spatial resolution such that each image has the size of 120x120 pixels. We perform self-supervised pre-training on the training split without labels.

Model. We use ViT-S/8 as backbone where S indicates a small patch embedding dimension of 384, and /8 signals input images get split into patches of size 8x8. We design the projection head to a 3-layer MLP with hidden dimension 2048. It is followed by l_2 normalization and a weight-normalized fully connected layer of dimension 65,536.

Data augmentation. We add the proposed RandomSensorDrop augmentation to the end of the existing DINO aug-

Table 2: Linear classification results on the BigEarthNet-MM dataset [18]. We report random initialization, self-supervised pre-training with a single modality (DINO-S1/2), joint SAR-optical pre-training (DINO-MM), and fully supervised learning. We also report the performance under label-limited regimes where only 1% of the labels are available to train the linear classifier.

	100%			1%		
	S1	S2	S1+S2	S1	S2	S1+S2
Random	54.6	62.0	64.5	52.7	59.0	62.4
DINO-S1/2	76.2	86.0	–	68.7	82.0	–
DINO-MM	79.5	87.1	87.1	75.3	82.9	82.8
Supervised	77.1	86.7	88.6	63.7	73.6	75.0

mentations including: RandomResizedCrop, RandomHorizontalFlip, RandomColorJitter, RandomGrayscale, RandomGaussianBlur, and RandomSolarize. We follow the multi-crop strategy with 2 global and 8 local crops.

Self-supervised pre-training. We train the model for 100 epochs applying the AdamW optimizer with batch size 256. The training takes about 8 hours on 4 NVIDIA A100 GPUs. The learning rate is linearly ramped up to $5 \cdot 10^{-4}$ over the first 10 epochs and then decayed with a cosine schedule. The temperature τ_s is set to 0.1 and a linear warm-up increases τ_t from 0.04 to 0.07 over the first 30 epochs.

Performance evaluation. We quantify performance of the proposed method by *linear classification*, i.e., we train a linear classifier on top of the parameter-frozen encoder f . We define the teacher backbone as the pre-trained model. We train the linear classifier for 100 epochs with a batch size of 256, SGD optimizer, and a cosine-decayed learning rate starting from 0.01. As reference, we compare the resulting performance to fully supervised learning running on the AdamW optimizer and a cosine-decayed learning rate starting at 10^{-3} . The `torch.nn.MultiLabelSoftMarginLoss` is applied for supervised training. We report the performance score in terms of mean average precision.

3.2. Results

Table 2 verifies the success of self-supervised ViTs and highlights the benefits of DINO-MM. While table rows represent different training models, corresponding columns list downstream tasks with full labels (100%) and a fraction of labels (1%), respectively.

Self-supervised ViTs. In contrast to random initialization (row 1), self-supervised pre-training significantly boosts performance by a margin of $\geq 20\%$ points (rows 2, 3). This result underlines the capacity of ViTs to generate comprehensive representations for remote sensing data—an observation verified when compared to supervised learning (row 4): training a linear classifier with 100% labels, ViT-based SSL reaches performance close to fully supervised learning.

When reducing the fraction of labels to 1%, self-supervised

pre-training outperforms supervised learning on any combination of Sentinel-1/2 imagery. This result advertises SSL as an efficient approach in situations with a limited amount of labels available. In addition, while the performance score for supervised learning significantly drops with number of labels reduced ($\geq 13\%$ points), the performance of linear classification based on pre-trained models is affected little ($\leq 5\%$ points). This demonstrates the benefits of SSL to help minimize human annotation efforts.

Joint multimodal SSL. A comparison of SSL based on a single modality, and SSL with both modalities is presented in rows 2 and 3 of Table 2, respectively. With both modalities (DINO-MM on S1+S2), the performance is improved compared to a single modality (DINO-S1 on S1 or DINO-S2 on S2)—a result expected since an additional modality adds information. Further, the results suggest optical imagery capture more relevant information than SAR raster data for the specific scene classification task at hand. Most significantly, Table 2 supports the benefits of joint multimodal SSL: a model pre-trained with both modalities transfers well to single-modality tasks, and it can even boost performance. That is, DINO-MM on S1 or S2 outperforms DINO-S1 on S1 or DINO-S2 on S2, and it even outperforms supervised learning with all the labels. This result is critical to remote sensing applications such as early change detection for disaster monitoring where either one or the other satellite captures a scene first.

4. CONCLUSION

In this paper, we explored the potential of self-supervised vision transformers (ViTs) for multimodal remote sensing image understanding. Based on a state-of-the-art self-supervised learning (SSL) algorithm, DINO, we proposed and evaluated a model extension, DINO-MM, to jointly learn representations of SAR-optical data. We introduced a simple data augmentation module, RandomSensorDrop, that randomly masks out channels of one modality while training. Experimental results verify the success of ViT backbones for SSL in remote sensing and prove the benefits of the proposed multimodal SSL method.

5. REFERENCES

- [1] D. H. Ballard, “Modular learning in neural networks.” in *AAAI*, vol. 647, 1987, pp. 279–284.
- [2] 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.
- [3] A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018.
- [4] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, “Unsupervised learning of visual features by contrasting cluster assignments,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 9912–9924, 2020.
- [5] X. Chen and K. He, “Exploring simple siamese representation learning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 750–15 758.
- [6] J. Zbontar, L. Jing, I. Misra, Y. LeCun, and S. Deny, “Barlow twins: Self-supervised learning via redundancy reduction,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 12 310–12 320.
- [7] N. Jean, S. Wang, A. Samar, G. Azzari, D. Lobell, and S. Ermon, “Tile2vec: Unsupervised representation learning for spatially distributed data,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, pp. 3967–3974.
- [8] K. Ayush, B. Uzcent, C. Meng, K. Tanmay, M. Burke, D. Lobell, and S. Ermon, “Geography-aware self-supervised learning,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 181–10 190.
- [9] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, “Emerging properties in self-supervised vision transformers,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9650–9660.
- [10] 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.
- [11] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *International Conference on Learning Representations*, 2020.
- [12] X. Chen, S. Xie, and K. He, “An empirical study of training self-supervised vision transformers,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9640–9649.
- [13] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, “Masked autoencoders are scalable vision learners,” *arXiv preprint arXiv:2111.06377*, 2021.
- [14] K. Cha, J. Seo, and Y. Choi, “Contrastive multiview coding with electro-optics for sar semantic segmentation,” *IEEE Geoscience and Remote Sensing Letters*, 2021.
- [15] A. Montanaro, D. Valsesia, G. Fracastoro, and E. Magli, “Self-supervised learning for joint sar and multispectral land cover classification,” *arXiv preprint arXiv:2108.09075*, 2021.
- [16] Y. Chen and L. Bruzzone, “Self-supervised sar-optical data fusion of sentinel-1/-2 images,” *IEEE Transactions on Geoscience and Remote Sensing*, 2021.
- [17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [18] G. Sumbul, A. De Wall, T. Kreuziger, F. Marcelino, H. Costa, P. Benevides, M. Caetano, B. Demir, and V. Markl, “Bigearthnet-mm: A large-scale, multimodal, multilabel benchmark archive for remote sensing image classification and retrieval [software and data sets],” *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 3, pp. 174–180, 2021.