




Learning Representations of Satellite Images From Metadata Supervision

Jules Bourcier^{1,2}, Gohar Dashyan¹, Karteek Alahari², and Jocelyn Chaussonot²

¹ Preligens, Paris, France

² Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, Grenoble, France

Abstract. Self-supervised learning is increasingly applied to Earth observation problems that leverage satellite and other remotely sensed data. Within satellite imagery, metadata such as time and location often hold significant semantic information that improves scene understanding. In this paper, we introduce Satellite Metadata-Image Pretraining (SatMIP), a new approach for harnessing metadata in the pretraining phase through a flexible and unified multimodal learning objective. SatMIP represents metadata as textual captions and aligns images with metadata in a shared embedding space by solving a metadata-image contrastive task. Our model learns a non-trivial image representation that can effectively handle recognition tasks. We further enhance this model by combining image self-supervision and metadata supervision, introducing SatMIPS. As a result, SatMIPS improves over its image-image pretraining baseline, SimCLR, and accelerates convergence. Comparison against four recent contrastive and masked autoencoding-based methods for remote sensing also highlight the efficacy of our approach. Furthermore, our framework enables multimodal classification with metadata to improve the performance of visual features, and yields more robust hierarchical pretraining. Code and pretrained models will be made available at: <https://github.com/preligens-lab/satmip>.

Keywords: Self-supervised and multimodal learning · Remote sensing

1 Introduction

In recent years, self-supervised learning (SSL) has become a staple pretraining paradigm in computer vision, and has received much attention in the domain of remote sensing and Earth observation (EO) [53]. This marked interest can be attributed to two broad reasons. First, for a wide variety of high-impact EO tasks, ranging from crop-yield prediction to urban planning [11, 13, 17, 29, 44, 46, 49], labels are scarce and difficult to obtain, while unlabeled satellite imagery is abundantly available. This makes SSL eminently practical. Second, the diversity of remote sensors yields unique challenges of specialization, context awareness, and multimodal fusion, with rich spatial, temporal, and spectral contexts. This calls for the development of tailored representation learning methods, in order to address limitations of existing generic vision models [43, 48].

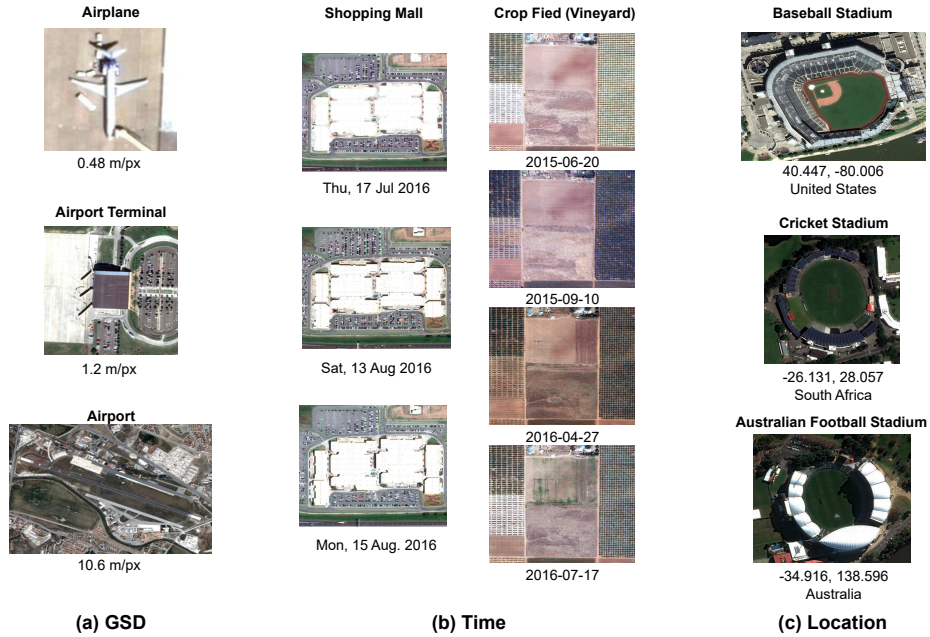


Fig. 1: Examples of satellite images where metadata can help recognizing objects. Ground sample distance (GSD) (a), which determines the area occupied by pixels, provides information on size and scale (e.g., airport/airport terminal/airplane); time (b) and location (c) can help understand the functional nature of man-made structures having different appearance depending on time (e.g., shopping mall’s parking lot are fuller on the weekend, crop fields undergo cycles) or place (e.g., some sports facilities correlated with regions on Earth), respectively.

EO satellites usually generate a rich set of metadata associated with the images they capture. They provide information about the size, scale, time of acquisition, as well as numerous other image properties. Some of these metadata, such as the ground sampling distance (GSD), timestamp, and geographic location, can be highly explanatory of semantic content present in images, as illustrated in Fig. 1.

Recent works have proposed SSL objectives using satellite metadata, such as predicting geolocation from images as a pretext task [2], using spatial neighbors or spatially-aligned images over time as positive views in contrastive learning [2, 24, 34, 39], or extending Vision Transformers (ViTs) [15] with positional encodings that integrate information from timestamps and spectral groups [12] or the GSD [41]. These methods have successfully improved contrastive/siamese learning algorithms [6, 7, 20] or masked autoencoders (MAEs) [4, 19]. However, these methods utilize different metadata fields via specialized model architectures or tasks. To our knowledge, there has not been a unified and flexible approach for incorporating heterogeneous metadata into a pretraining algorithm.

In this work, we propose a simple and effective model for learning visual representations from satellite metadata supervision. Our model, Satellite Metadata-Image Pretraining (SatMIP), encodes pairs of images and metadata as separate modalities and aligns them in a deep embedding space via a contrastive task, inspired by language-image pretraining [25, 40]. Through this task, we aim to learn a visual encoder that embeds metadata information, and their latent semantic characteristics, into image features. It requires metadata only during the pretraining phase, and not necessarily during transfer to downstream tasks. We pretrain ViTs backbones with SatMIP on the Functional Map of the World (fMoW) dataset [11], using GSD, timestamp, and geolocation, among other metadata fields. Through extensive experiments on various downstream classification datasets, we observe that the visual encoders pretrained with SatMIP generate non-trivial representations that generalize to downstream recognition tasks, showing that learning a joint embedding between images and metadata makes a meaningful pretext task.

To go one step further, we combine SatMIP with the image SSL method SimCLR [7], introducing SatMIPS. By co-solving an image-image and a metadata-image contrastive task with an efficient “coupled” architecture, SatMIPS benefits from both sources of supervision, and improves over its SimCLR baseline, yielding better representations while converging faster. Moreover, on several downstream tasks, it outperforms multiple existing MAE and contrastive-based pretraining methods involving metadata for remote sensing. Furthermore, on downstream tasks with metadata, SatMIP allows deploying metadata features in tandem with visual features, which can further improve the classification performance. In addition, we also show that metadata supervision yields stronger results with hierarchical pretraining [42].

- We propose SatMIP, a novel self-supervised pretraining task and model for remote sensing inspired by CLIP [40], which aligns images with their metadata in a joint embedding space.
- We further propose SatMIPS, an evolution of the SLIP [36] architecture, which combines image-image and metadata-image contrastive learning.
- We conduct extensive experiments involving various downstream classification datasets, demonstrating the effectiveness and efficiency of our approach.

2 Related Work

Using geospatial metadata for visual representation learning. Satellite images systematically convey metadata that can be leveraged for free within SSL tasks similarly to pseudo-labels. One strategy involves employing metadata estimation as a pretext task: [56] proposes self-supervised time and location estimation tasks for learning geotemporal image features, while [2] solves a location classification as a subsidiary task to contrastive SSL [7, 20]. Contrary to predicting metadata information directly, our approach aligns image and metadata into a common embedding space.

Another avenue is the implicit incorporation of metadata information into existing pretext tasks, by using it to enrich the set of positive or negative instances and thereby learn the invariances driven by these augmentations. Building on this idea, some works use neighboring images in space as positives [24, 26, 39], while others use spatially-aligned images over time as positives [2, 33, 34]. Another line of recent works have extended masked autoencoders (MAE) [19] for remote sensing, and incorporated metadata information into positional encodings in ViTs. [12] proposes a spectral and temporal reconstruction task, and embeds timestamps and spectral bands into the positional encodings; [41] solves a super-resolution task and embeds GSD into positional encodings to incorporate scale information; [23] further extends this idea to multiple spectral bands, GSDs and sensors. In contrast to our model that is agnostic to network architectures or vision SSL frameworks, these approaches bake into ViTs and are specifically tailored to MAEs.

Embedding metadata. Recent studies aim to directly encode metadata along with images and perform a form of metadata-image pretraining. Close to our work, in [27, 32, 51], location encoders are learned through a contrastive image-location pretraining task, aiming to be deployed on downstream tasks involving location. In contrast to our approach, these works employ a two-step approach of training an image encoder, then a location encoder on frozen image features. Our objective being to train a visual encoder intended for downstream tasks without necessarily relying on metadata, we adopt the opposite approach by training a visual encoder with metadata supervision. Furthermore, instead of relying on location-specific encoder architectures as them, we employ a generic Transformer [50] that can be fed heterogeneous types of metadata beyond location.

Closer to our approach, [59] uses EXIF metadata and images in contrastive pretraining for learning to extract low-level camera properties of images for forensics tasks. Our approach differs in that we propose metadata-image pretraining for high-level representation learning. Furthermore, we enhance our metadata-image objective by concurrently solving an image-image contrastive task. Other work such as [18, 22] used medical images and biodata records as tabular features in contrastive learning for increasing performances on downstream visual diagnostic tasks. Such metadata is the result of a supervised and very expensive collection process that does not scale to large datasets.

Language vs. metadata supervision. Language-image pretraining (LIP) has emerged as a significant advance that bridges the gap between natural language and image representation learning. The works of CLIP [40] and ALIGN [25] showed that the straightforward pretext task of predicting which caption corresponds to a given image is an effective way to learn image representations on large-scale noisy (image, text) pairs. SLIP [36] extends CLIP via a multi-task learning framework combining it with image-image contrastive pretraining [7], showing that both objectives are synergistic. LIP has also recently garnered at-

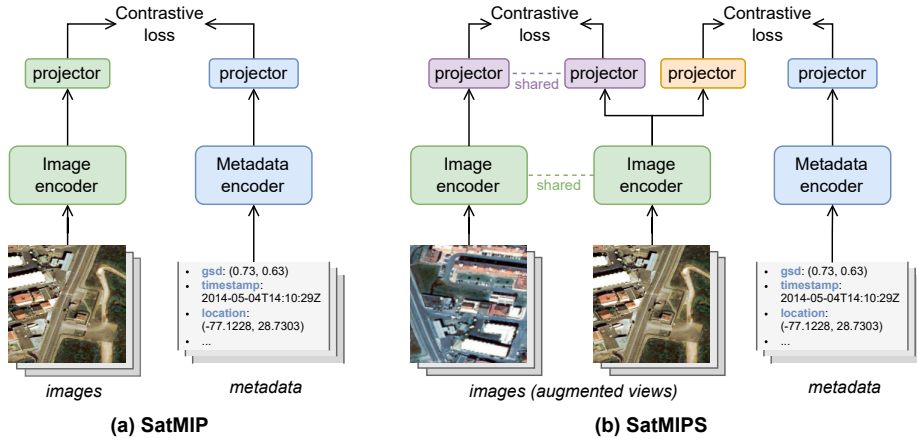


Fig. 2: (a) **Satellite Metadata-Image Pretraining (SatMIP)** learns a joint embedding between images and associated metadata. Batches of inputs are processed by a dual encoder and projection head, optimized through a contrastive loss, as in CLIP [40]. After pretraining, only the image encoder can be transferred to downstream tasks, or both can be used to perform bimodal (image, metadata) recognition. (b) **SatMIP with Self-Supervision (SatMIPS)** combines metadata with SimCLR-style [7] self-supervision: it learns both a joint embedding between augmented views of an image as well as an image and its metadata. The image encoder is shared between branches unlike the projection heads, and one image view is coupled between the two objectives for efficiency.

tention within the field of remote sensing [30, 61]. However, while image captions are available in huge web-crawled multimedia datasets such as YFCC100M [47] or LAION-5B [45], such human-provided captions are scarce for remotely sensed images. On the other hand, metadata are automatically produced by sensors and are therefore widely available, which means that the number of image-metadata pairs can follow the ever-growing number of satellite images. Seeking to harness metadata for complementing vision, our SatMIP (respectively SatMIPS) model is analogous to CLIP (respectively SLIP), but with metadata as an input modality instead of language.

3 Method

We aim to learn a visual representation of remotely sensed images that embeds the semantic information contained within metadata that is obtained directly from the imaging sensor. To this end, we introduce SatMIP, a pretraining strategy that learns a joint embedding between an image and its metadata. Then, we introduce SatMIPS, which leverages both image self-supervision and metadata supervision. The architectures of SatMIP and SatMIPS are presented in Fig. 2.

3.1 SatMIP: Contrastive pretraining of metadata and image embeddings

Contrastive pretext task. We assume that we have access to an unlabeled dataset as $\mathbb{X} = \{(\mathbf{v}_i, \mathbf{m}_i)\}_{i=1}^N$, where $(\mathbf{v}_i, \mathbf{m}_i) \sim p(\mathcal{I}, \mathcal{M})$ are associated images and metadata pairs, sampled from their respective spaces \mathcal{I}, \mathcal{M} . The metadata space \mathcal{M} can be composed of a set of numerical variables (*e.g.*, GSD, location coordinates, or look angle) and categorical variables (*e.g.*, sensor name). We define two neural network encoders, one for images $e^{\mathcal{I}} : \mathcal{I} \rightarrow \mathbb{R}^d$ and one for metadata $e^{\mathcal{M}} : \mathcal{M} \rightarrow \mathbb{R}^d$. Each encoder $e^{\mathcal{I}/\mathcal{M}}$ is composed of a backbone $f^{\mathcal{I}/\mathcal{M}}$ and a projection head $g^{\mathcal{I}/\mathcal{M}}$. Given a sampled batch of K image and metadata pairs, we compute embeddings of images $\mathbf{z}_i^{\mathcal{I}} = e^{\mathcal{I}}(\mathbf{v}_i)$ and metadata $\mathbf{z}_i^{\mathcal{M}} = e^{\mathcal{M}}(\mathbf{m}_i)$. Following CLIP [40], we use a contrastive loss by considering matching images and metadata as positives and non-matching images and metadata across the batch as negatives. Let us define the generic contrastive loss function [38]:

$$\mathcal{L}^{\text{clr}}(a_i, b_i) = -\log \frac{\exp(s(a_i, b_i)/\tau)}{\sum_{j=1}^K \exp(s(a_i, b_j)/\tau)}, \quad (1)$$

where (a_i, b_i) are two vectors of equal dimension, s is the cosine similarity and τ is a parameter that adjusts the dynamic range. We define the loss of SatMIP as the symmetrized contrastive loss between images and metadata embeddings:

$$\mathcal{L}_i^{\text{MI}}(\mathbf{z}_i^{\mathcal{I}}, \mathbf{z}_i^{\mathcal{M}}) = \frac{1}{2} (\mathcal{L}^{\text{clr}}(\mathbf{z}_i^{\mathcal{I}}, \mathbf{z}_i^{\mathcal{M}}) + \mathcal{L}^{\text{clr}}(\mathbf{z}_i^{\mathcal{M}}, \mathbf{z}_i^{\mathcal{I}})). \quad (2)$$

Through this objective, the weights of encoders $e^{\mathcal{I}}$ and $e^{\mathcal{M}}$ are optimized simultaneously to embed feature vectors of the matching image and metadata nearby in a common latent space. The objective can be naturally interpreted as one of classifying the correct metadata from the image, and correct image from the metadata. Intuitively, the image encoder will learn implicit neural features from the metadata, and vice versa. Note that there does not exist a simple 1:1 mapping between images and metadata, because metadata can match many image variations and vice versa (*e.g.*, due to the non-deterministic nature of weather). This prevents the model from solely overfitting the pretext task. In addition, we apply data augmentation to the images which further regularizes the task.

Transfer to downstream tasks. Our primary goal with SatMIP is to learn a visual representation from the image backbone $f^{\mathcal{I}}$, to be transferred to downstream tasks. In this case, the metadata encoder is used only as a proxy for pretraining and is then discarded. Alternatively, provided that a downstream task provides a subset of metadata fields used to train the metadata encoder, we can encode metadata alongside images with the dual backbone $(f^{\mathcal{I}}, f^{\mathcal{M}})$ and fuse their embeddings as input to a supervised model. Several techniques exist to fuse multimodal embeddings [1, 3]. We use the simplest strategy of concatenating both vectors and fitting a parametric classifier on top to learn an optimal combination of features for a given task.

3.2 SatMIPS: Combining self- with metadata supervision

We further introduce SatMIPS which combines the previously described SatMIP with the SSL method SimCLR [7].

SimCLR learns a joint embedding between two augmented views of the same image with a contrastive loss, which makes it very similar to CLIP. Let the image encoder be $e = g \circ f : \mathcal{I} \rightarrow \mathbb{R}^c$. Given a batch of positive views $\{(\mathbf{v}_i, \mathbf{v}'_i)\}_{k=1}^K$, it computes embeddings $\mathbf{z}_i = e(\mathbf{v}_i)$ and $\mathbf{z}'_i = e(\mathbf{v}'_i)$ and employs the following symmetrized contrastive loss to align the embeddings of matching views:

$$\mathcal{L}_i^{\text{Sim}}(\mathbf{z}_i, \mathbf{z}'_i) = \frac{1}{2} (\mathcal{L}^{\text{clr}}(\mathbf{z}_i, \mathbf{z}'_i) + \mathcal{L}^{\text{clr}}(\mathbf{z}'_i, \mathbf{z}_i)) \quad (3)$$

Multi-task framework. [36] proposed SLIP as a combination of CLIP and SimCLR through multi-tasking. Following their framework, we express the SatMIPS task as a linear combination of the two SimCLR and SatMIP objectives. We share the visual backbone $f = f^{\mathcal{I}}$ between both models and optimize the sum of their loss:

$$\mathcal{L}_i^{\text{MI+Sim}}(\mathbf{z}_i^{\mathcal{I}}, \mathbf{z}_i^{\mathcal{M}}, \mathbf{z}_i, \mathbf{z}'_i) = \mathcal{L}_i^{\text{MI}} + \lambda \mathcal{L}_i^{\text{Sim}} \quad (4)$$

where λ is a hyperparameter that weights the promicence of the SimCLR loss relative to the SatMIP loss. We find that $\lambda = 1$, *i.e.*, equal weighting works well (we show this in additional ablations in the supplementary material).

Efficient view coupling. The main issue of SLIP is that it increases the number of images processed from 1 to 3, resulting in approximately $3\times$ more activations [36] that increase training time and memory footprint significantly. To alleviate this cost, in SatMIPS, we couple one of the image view \mathbf{v}_i between SimCLR and SatMIP. The output of the backbone f on this shared view is directed through the specific projection heads of SimCLR (g) and SatMIP ($g^{\mathcal{I}}$). This design is driven by our tests which showed that SatMIP works well with the same strong augmentation policy as SimCLR (*cf.* Sec. 4.7). Thanks to view coupling, we can largely reduce the overhead without an impact on downstream performance (*cf.* Sec. 4.7).

We have selected SimCLR as the image SSL method in SatMIPS for its simplicity and conceptual similarity with SatMIP. However, the general design of SatMIPS is agnostic to this choice, and the metadata-image objective could be blended with another SSL method.

4 Experiments and results

We evaluate our SatMIP and SatMIPS models by studying the performance of their learned representations on a set of remote sensing downstream classification tasks. We conduct experiments to benchmark the quality of representations

under k-nearest neighbors (kNN) and linear probing classification, the rate of convergence of pretraining, and the application of hierarchical pretraining. We then perform an ablation study of important components of our models.

4.1 Datasets

Pretraining. To pretrain our models we use the training set of the Functional Map of the World (fMoW) dataset [11], similar to previous work [2, 12, 41]. It consists of 363k global, very high-resolution images and associated metadata obtained by MAXAR optical satellites. We use the fMoW-RGB product, composed of the RGB pansharpened images. The metadata is composed of a diverse set of metadata fields, including, among others, GSD, timestamp, location, location-derived information such as UTM zone and country, cloud cover, and various imaging angles; we describe the full metadata considered in the supplementary material. We exclude any field that is obtained through manual annotation such as areas of interests and land use categories. Unless otherwise specified, in our SatMIP experiments, we used a combination of three fields: the GSD, timestamp, and location, described in Tab. 1 and visualized in the supplementary material. As we can see, the fields span an extensive range of values. We preprocess each source image by cropping over the annotated area of interest and resizing to 224×224 pixels, and we transform the GSD and location fields to match the preprocessed images.

Evaluation. To evaluate the performance of pretraining methods, we use a diverse set of 7 remote sensing RGB image classification datasets: (1) The labeled version of fMoW with 62 classes of functional land use (sharing the same training data as for pretraining); (2) RESISC45 [10] for land use/land cover classification of multi-sensor (satellite & aerial) images; (3) Optimal31 [52] for land use/land cover classification of multi-sensor images; (4) UC Merced [55] for land use classification of very-high resolution aerial images; (5) FGSC23 [58] for fine-grained ship classification in high-resolution multi-sensor images; (6) EuroSAT [21] for land use/land cover classification of Sentinel-2 images; (7) So2Sat [62] for local climate zone classification of Sentinel-2 images. We report macro-averaged F1 score on fMoW and FGSC23, and top-1 accuracy on the other datasets.

4.2 Setup

Baseline and state-of-the-art. We adopt SimCLR [7] as the natural baseline to compare SatMIP and SatMIPS to. We also compare to existing SSL models for remote sensing that are pretrained on fMoW-RGB: Geo, TP, and Geo-TP from [2], which are originally based on MoCo [8, 20]; SatMAE [12], Scale-MAE [41], and SatMAE++ [37], which are based on MAE [19]. We reproduce contrastive methods on top of SimCLR for an even comparison with our models, while for MAE-based models, we take pretrained weights available on their official repositories.

Table 1: The three fields of metadata we use to train our models, and descriptive statistics of values in the fMoW training set.

Field	Description	Min	Median	Max
Ground sample distance	Physical distance between pixel centers, in x and y directions (m)	0.08, 0.06	0.76, 0.60	23.13, 22.35
Timestamp	Date and time of acquisition (UTC)	2002-01-28 07:04:18	2015-08-06 10:08:02	2017-07-12 08:25:25
Location	Latitude and longitude of the image centroid (degrees)	-54.9320, -179.8810	37.9951, 7.0395	71.6118, 179.0439

Implementation details. *Visual encoders:* Unless otherwise noted, we use the MoCoV3 [9] version of a ViT-Small with patch size 16, consisting of 21.7M parameters.

Metadata encoders: To encode metadata, we experimented with two different types of encoders: (a) a textual encoder, which first converts metadata to text and then tokenizes it into as a sequence, processed like the language modality in CLIP [40] (inspired by [59] which applied this approach to EXIF metadata); (b) a tabular encoder, which natively supports numerical fields. For experiments targeting image-only classification, we use a textual encoder, while we use a tabular encoder for experiments involving bimodal (visual and metadata) classification. Both approaches work within our models, but perform differently depending on the downstream use case: we present a comparative study in the supplementary material. We use Transformer models [50] composed of 3 layers with 8 attention heads and a width of 512. For the textual approach, we use a BERT-style Transformer [14], while for the tabular approach, we use the FT-Transformer model [16].

Projection heads: In SatMIP, we use a linear layer for the visual and textual encoders. In SimCLR and SatMIPS, the projection head for the image-image objective is a 3-layer MLP.

Data augmentation: For all models, the inputs image views are generated with the augmentation policy of [7], with the addition of vertical flips and rotations [5, 60]. We do not apply augmentation to the metadata.

Pretraining: We use a global batch size of 1024 for all models and train with the AdamW [31] optimizer for 200 epochs, unless otherwise noted. All other hyperparameters of SimCLR and SatMIP (S) are provided in the supplementary material.

Evaluation schemes. The keyword we apply for evaluating the models is *practicality*, of both the learned representation and the pretraining algorithm. To measure the achievement of these criteria, we evaluate the representation quality by fitting kNN or linear classifiers on frozen features extracted from the training

Table 2: kNN classification performance on various downstream datasets of SSL methods pretrained on fMoW. We compare **SatMIP** and **SatMIPS** against the baseline SimCLR and existing contrastive and MAE-based methods, under a consistent evaluation. For reference, Random indicates a feature extractor with random weights. We **bold** the best accuracies per dataset. R45: RESISC45, O31: Optimal31, UCM: UC Merced, F23: FGSC-23, Euro: EuroSAT, So2: So2Sat, Acc.: Accuracy (Top-1).

Model	ViT Epochs		fMoW	R45	O31	UCM	F23	Euro	So2
	size		F1	Acc.	Acc.	Acc.	Acc.	Acc.	Acc.
Random	S	–	5.1 \pm 0.1	32.2 \pm 0.4	28.5 \pm 0.6	44.0 \pm 0.9	26.7 \pm 1.4	70.0 \pm 0.1	33.3 \pm 0.2
<i>Contrastive-based</i>									
SimCLR [7]	S	200	61.1 \pm 0.6	88.5 \pm 0.4	86.0 \pm 1.2	95.0 \pm 0.4	57.3 \pm 2.1	94.3 \pm 0.6	56.9\pm0.5
SimCLR-Geo [2]	S	200	59.0 \pm 0.3	88.9 \pm 0.5	87.0 \pm 1.0	95.2 \pm 0.6	60.4 \pm 0.9	95.5 \pm 0.2	57.5\pm0.4
SimCLR-TP [2]	S	200	65.2 \pm 0.3	90.4 \pm 0.5	87.6 \pm 1.1	97.6 \pm 0.4	61.2 \pm 2.3	95.00 \pm 0.2	57.3\pm0.7
SimCLR-Geo-TP [2]	S	200	65.8\pm0.6	91.3\pm0.5	89.7\pm1.0	97.1\pm0.4	63.0\pm0.4	95.6\pm0.2	57.2\pm0.6
<i>MAE-based</i>									
SatMAE [12]	L	800	46.3	75.2	69.6	86.2	44.7	91.3	53.7
Scale-MAE [41]	L	800	51.4	85.9	81.6	89.0	48.2	96.1	56.7
SatMAE++ [37]	L	800	38.0	77.1	67.8	84.9	46.9	93.1	51.5
SatMIP	S	200	55.2 \pm 0.2	87.5 \pm 0.1	84.8 \pm 0.6	95.2 \pm 0.8	56.4 \pm 0.2	95.7\pm0.5	55.9 \pm 0.2
SatMIPS	S	200	62.3 \pm 0.0	89.7 \pm 0.2	87.9 \pm 0.2	94.9 \pm 0.7	60.8 \pm 0.6	95.1 \pm 0.1	57.1\pm0.5

set of downstream tasks, following regular protocols [6, 54]. We measure resource efficiency by comparing performance across amounts of pretraining epochs and total training time.

4.3 Quality of visual representations

We present the results of our kNN classification experiments on fMoW and the four downstream datasets in Tab. 2. First, we observe that SatMIP learns non-trivial representations: it clearly outperforms random features, and even all MAE-based methods on all datasets except Scale-MAE on Euro and So2. It also competes with SimCLR on UCM and outperforms it on Euro. This validates that using metadata as supervision is effective for learning high-level semantic representations. Second, we see that SatMIPS outperforms SimCLR on all datasets, with the exception of UCM and So2 where they tie. This shows that image and metadata self-supervision interact constructively in SatMIPS to improve the quality of the shared features. Moreover, SatMIPS outperforms SimCLR-Geo on fMoW, R45, O31, and F23, and is comparable on other datasets. This tends to indicate that integrating metadata into a joint embedding objective is more effective than directly predicting metadata, as is done by Geo. Still, SatMIPS is mostly outperformed by SimCLR-TP and Geo-TP: although adding temporal positives makes a stronger extension than adding metadata supervision, we note that both methods could potentially be combined. However, this is beyond the scope of this work.

4.4 Classification on image and metadata features

Table 3: Classification with varying modalities on fMoW and EuroSAT: performance of image and combined image and metadata features learned via our models, using linear probing.

Model	fMoW F1			EuroSAT Acc.		
	Image	Image+Meta.	Δ	Image	Image+Meta.	Δ
SatMIP	59.3 \pm 0.3	63.1 \pm 0.1	+3.8 \pm 0.2	94.5 \pm 0.6	95.5 \pm 0.1	+1.0 \pm 0.4
SatMIPS	65.8 \pm 0.1	68.6 \pm 0.2	+2.8 \pm 0.2	95.8 \pm 0.3	96.4 \pm 0.2	+0.6 \pm 0.3

We investigate how the metadata modality can provide further benefits when used in downstream classification tasks. We perform bimodal classification using the combined representations of image and metadata encoders, on fMoW (using GSDs, timestamps and locations) and also on EuroSAT, for which we use the supplied GSDs and locations. We concatenate the features from each modality, and fit a linear classifier on the combined features. Results are presented in Tab. 3. On fMoW, we observe that for both SatMIP and SatMIPS, bimodal features provide a substantial performance improvement compared to features from images alone. On EuroSAT, bimodal features also achieve modest improvements. These results show that, in addition to forming useful supervision for pretraining, the learned metadata features are complementary to visual features for downstream tasks and can further improve performance.

4.5 Convergence speed analysis




Fig. 3: kNN classification performance of various datasets as a function of total pre-training time for SatMIP, SimCLR and SatMIPS. Error bars indicate mean \pm std. dev. over 3 runs. Data points correspond to 25, 50, 100 and 200 epochs in order of increasing time. Training times are relative to SimCLR at 25 epochs. Performance metric is accuracy except for fMoW and FGSC23 where it is F1 score.

To investigate the resource efficiency of pretraining procedures, we compare the performance obtained for different amounts of pretraining epochs and the resulting training time of SimCLR, SatMIP and SatMIPS. Results are presented

in Fig. 3. First, we note that SatMIP trains faster than SimCLR (by about 44%), and while it generally underperforms the baseline at 200 epochs, it generally does better when comparing at equivalent training time. Second, we observe that SatMIPS converges faster than SimCLR: the performance of the former at 100 epochs equates or surpasses the one of the latter at 200 epochs, on all the datasets. Even though SatMIPS is marginally slower to train than SimCLR for an equal amount of iterations/epochs (by about 5%), the higher convergence rate makes it a more efficient method. This shows that our metadata-image objective makes pretraining efficient, in addition to being effective.

4.6 Hierarchical pretraining

Table 4: kNN classification performance employing hierarchical pretraining (HPT) [42], with base pretraining on YFCC [47] with SLIP [36], and further pre-training each on fMoW with each model. Abbreviations share those of Tab. 2.

Model	HPT from SLIP-YFCC	fMoW F1	R45 Acc.	O31 Acc.	UCM Acc.	F23 Acc.	Euro Acc.	So2 Acc.	Avg. Δ
SimCLR		61.1 \pm 0.6	88.5 \pm 0.4	86.0 \pm 1.2	95.0 \pm 0.4	57.3 \pm 2.1	94.3 \pm 0.6	56.9 \pm 0.5	
	✓	63.7 \pm 0.1	90.0 \pm 0.2	88.5 \pm 0.8	94.8 \pm 0.0	55.1 \pm 1.3	94.2 \pm 0.3	57.1 \pm 0.4	+0.6
SatMIP		55.2 \pm 0.2	87.5 \pm 0.1	84.8 \pm 0.6	95.2 \pm 0.8	56.4 \pm 0.2	95.7 \pm 0.5	55.9 \pm 0.2	
	✓	61.2 \pm 0.4	90.5 \pm 0.2	88.2 \pm 0.1	96.5 \pm 0.2	58.2 \pm 0.4	96.0 \pm 0.4	57.1 \pm 0.7	+2.4
SatMIPS		62.3 \pm 0.04	89.7 \pm 0.2	87.9 \pm 0.2	94.9 \pm 0.7	60.8 \pm 0.6	95.1 \pm 0.1	57.1 \pm 0.5	
	✓	66.3 \pm 0.2	91.4 \pm 0.3	89.7 \pm 0.1	96.2 \pm 0.3	59.9 \pm 0.3	95.9 \pm 0.1	57.9 \pm 0.5	+1.4

We consider the compatibility of models with hierarchical pretraining (HPT) as it is known to be a practical way to increase the performance of SSL models, especially in remote sensing [35, 42, 57]. To do so, we initialize the ViT backbones with openly-available base weights pretrained with SLIP, on the large generalist YFCC15M [40, 47] dataset. Results are presented in Tab. 4. We see that, on average, HPT provides greater performance improvements with SatMIP and SatMIPS. With SatMIP, the gains are also more steady across datasets, while HPT gives negative results on some datasets with SimCLR (UCM and F23) and SatMIPS (F23). Thus, SatMIPS advantage over SimCLR is reinforced, but more surprisingly, SatMIP outperforms SimCLR on all datasets except fMoW and O31, showing that the performance gap between metadata supervision and image self-supervision is largely closed when leveraging HPT.

4.7 Ablation study

We ablate key components of SatMIP and SatMIPS. For these experiments, we pretrain only for 25 epochs on fMoW. Additional ablations are provided in the supplementary material.

Table 5: Influence of the choice of metadata fields in SatMIP. We report downstream linear probing performance on fMoW and RESISC45, with “(I)” meaning classification on image features, while “(I,M)” means bimodal classification with combined (image, metadata) features. Time.: Timestamp, Loc.: Location. We highlight the defaults in **blue** and **bold** the best accuracies.

Metadata fields				Performance		
GSD	Time.	Loc.	12 other fields	fMoW _(I) F1	fMoW _(I,M) F1	R45 _(I) Acc.
✓				37.4 \pm 0.6	40.5 \pm 0.9	71.3 \pm 0.4
	✓			43.7 \pm 0.1	46.5 \pm 0.1	74.3 \pm 0.3
		✓		47.1 \pm 0.2	53.3 \pm 0.2	76.1 \pm 0.7
	✓	✓		46.3 \pm 0.1	51.1 \pm 0.1	76.0 \pm 0.2
✓			✓	51.3\pm0.7	58.3\pm0.7	77.9 \pm 0.3
✓	✓			49.4 \pm 0.2	54.2 \pm 0.3	76.9 \pm 0.4
✓	✓	✓		50.7\pm0.4	57.7\pm0.5	78.5\pm0.2
✓	✓	✓	✓	50.7\pm0.1	58.0\pm0.2	78.6\pm0.1

Choice of metadata in SatMIP. To study how each metadata field impacts the quality of representations learned by SatMIP, we ablate the set of fields used for pretraining and measure the performance of visual and bimodal classification on downstream tasks. We compare different combinations of the three GSD, timestamp and location fields, and using an “extended” set of fields composed of these three, and 12 other fields available in fMoW, detailed in the supplementary material. We present these results in Table 5. Regarding sets of single fields, we see that GSD performs the worst, followed by timestamp and location. This suggests that location contains the most useful semantic information for supervision. fMoW being a globally distributed dataset with a high variability of locations (*cf.* supplementary material), it likely contributes to the importance of location in SatMIP. Nevertheless, all single fields yield non-trivial performances, and are useful in bimodal classification on fMoW, demonstrating SatMIP’s ability to incorporate metadata flexibly. Overall, combining multiple fields together tends to improve performance over using single fields. On fMoW image and bimodal classification, SatMIP reaches best performance with GSD and location-only for both visual and bimodal classification, and enlarging to timestamp and other fields does not improve performance further. On R45 however, we observe best performance with GSD, timestamp and location. These results suggests that the fields of GSD, timestamp and location contain most of the useful information present in satellite metadata, and that SatMIP can constructively combine the complementary information present in these heterogeneous fields.

Coupling in SatMIPS. We justify the SatMIPS design choice of coupling image views between SimCLR and SatMIP. First, we ablate the image augmentation policy in SatMIP, comparing the strong policy used for the SimCLR

Table 6: Impact of image augmentation in SatMIP and view coupling in SatMIPS. Training time and GPU memory consumption are reported relative to SimCLR. We highlight the defaults in **green**/**blue** and **bold** the best numbers per model.

Model	Augmentation	Coupling?	Training time (rel.)	Memory /GPU (rel.)	fMoW F1	R45 Acc.
SimCLR	SimCLR-Sat	–	1.	1.	46.1\pm0.0	83.0\pm0.3
SimCLR	crop	–	1.	1.	33.4 \pm 0.2	77.5 \pm 0.2
SatMIP	SimCLR-Sat	–	0.56	0.62	45.6\pm0.4	81.9\pm0.2
SatMIP	crop	–	0.56	0.62	44.6 \pm 0.4	80.0 \pm 0.4
SatMIPS	SimCLR-Sat	no	1.53	1.58	54.5\pm0.1	87.5\pm0.1
SatMIPS	SimCLR-Sat	yes	1.05	1.11	53.9 \pm 0.1	87.1 \pm 0.0

baseline (SimCLR-Sat) with the light random resized cropping policy commonly used for LIP [36, 40]. We observe that a strong policy performs better, showing that it provides useful regularization while not altering the correctness of the metadata enough to perturb the contrastive task. Consequently, we adopt a common augmentation policy between SimCLR and SatMIP, and evaluate the impact on view coupling in SatMIPS on resource efficiency and performance. We observe that coupling reduces the training time and memory usage by about 30% relatively to decoupling, and has very little impact on the representation performance. Thanks to coupling, SatMIPS’ training time is only 5% higher than that of SimCLR, and memory usage is also contained.

5 Conclusion

In this paper, we proposed a new self-supervised model for harnessing semantic information specific to satellite imagery metadata. We considered metadata as a complimentary modality to images, and demonstrated that SatMIP successfully learns useful visual and metadata representations. Our results have shown that metadata supervision is a strong competitor to traditional image-based SSL objectives, and that, within a multi-task framework, they are highly synergistic.

Our work on metadata supervision is focused on experiments with RGB images, but it could be applied to other remote sensing modalities, such as multispectral or radar images, since it does not make assumptions about the visual encoder. Also, we expect such representations to benefit from combining diverse sensors, thanks to the increased visual and metadata diversity. In addition, explicit information about the spectral bands could be included into the metadata encoding (*e.g.*, wavelengths or calibration parameters [28]) for learning spectrally-aware representations.

Additionally, our evaluation is focused on classification tasks on frozen features. One more avenue for future work is to study the behavior of models when finetuned, and exploring the transferrability of our models to dense prediction tasks.

Acknowledgments

We thank the anonymous reviewers for their constructive comments. This work has been supported in part by MIAI@Grenoble Alpes, (ANR-19-P3IA-0003), and has also benefited from access to the HPC resources of IDRIS under the allocation 2024-AD011013097R2 made by GENCI.

References

1. Audebert, N., Herold, C., Slimani, K., Vidal, C.: Multimodal deep networks for text and image-based document classification. In: Machine Learning and Knowledge Discovery in Databases: International Workshops of ECML PKDD 2019, Proceedings, Part I. pp. 427–443. Springer (2020)
2. Ayush, K., Uz Kent, B., Meng, C., Tanmay, K., Burke, M., Lobell, D., Ermon, S.: Geography-aware self-supervised learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10181–10190 (2021)
3. Bakkali, S., Ming, Z., Coustaty, M., Rusiñol, M.: Visual and textual deep feature fusion for document image classification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 562–563 (2020)
4. Bao, H., Dong, L., Piao, S., Wei, F.: BEit: BERT pre-training of image transformers. In: International Conference on Learning Representations (2022)
5. Bourcier, J., Dashyan, G., Chanussot, J., Alahari, K.: Evaluating the label efficiency of contrastive self-supervised learning for multi-resolution satellite imagery. In: Image and Signal Processing for Remote Sensing XXVIII. vol. 12267, pp. 152–161. SPIE (2022)
6. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. 2021 IEEE/CVF International Conference on Computer Vision (ICCV) pp. 9630–9640 (2021)
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
8. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
9. Chen, X., Xie, S., He, K.: An empirical study of training self-supervised vision transformers. 2021 IEEE/CVF International Conference on Computer Vision (ICCV) pp. 9620–9629 (2021)
10. Cheng, G., Han, J., Lu, X.: Remote sensing image scene classification: Benchmark and state of the art. Proceedings of the IEEE **105**(10), 1865–1883 (2017)
11. Christie, G., Fendley, N., Wilson, J., Mukherjee, R.: Functional map of the world. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6172–6180 (2018)
12. Cong, Y., Khanna, S., Meng, C., Liu, P., Rozi, E., He, Y., Burke, M., Lobell, D., Ermon, S.: SatMAE: Pre-training transformers for temporal and multi-spectral satellite imagery. Advances in Neural Information Processing Systems **35**, 197–211 (2022)
13. Demir, I., Koperski, K., Lindenbaum, D., Pang, G., Huang, J., Basu, S., Hughes, F., Tuia, D., Raskar, R.: Deepglobe 2018: A challenge to parse the earth through satellite images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 172–181 (2018)

14. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
15. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
16. Gorishniy, Y., Rubachev, I., Khrulkov, V., Babenko, A.: Revisiting deep learning models for tabular data. *Advances in Neural Information Processing Systems* **34**, 18932–18943 (2021)
17. Gupta, R., Goodman, B., Patel, N., Hosfelt, R., Sajeev, S., Heim, E., Doshi, J., Lucas, K., Choset, H., Gaston, M.: Creating xBD: A dataset for assessing building damage from satellite imagery. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. pp. 10–17 (2019)
18. Hager, P., Menten, M.J., Rueckert, D.: Best of both worlds: Multimodal contrastive learning with tabular and imaging data. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 23924–23935 (2023)
19. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16000–16009 (2022)
20. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9729–9738 (2020)
21. Helber, P., Bischke, B., Dengel, A., Borth, D.: EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **12**(7), 2217–2226 (2019)
22. Huang, W.: Multimodal contrastive learning and tabular attention for automated alzheimer’s disease prediction. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2473–2482 (2023)
23. Irvin, J., Tao, L., Zhou, J., Ma, Y., Nashold, L., Liu, B., Ng, A.Y.: USat: A unified self-supervised encoder for multi-sensor satellite imagery. *arXiv preprint arXiv:2312.02199* (2023)
24. Jean, N., Wang, S., Samar, A., Azzari, G., Lobell, D., Ermon, S.: Tile2vec: Unsupervised representation learning for spatially distributed data. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 3967–3974 (2019)
25. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: *International conference on machine learning*. pp. 4904–4916. PMLR (2021)
26. Kang, J., Fernandez-Beltran, R., Duan, P., Liu, S., Plaza, A.J.: Deep unsupervised embedding for remotely sensed images based on spatially augmented momentum contrast. *IEEE Transactions on Geoscience and Remote Sensing* **59**(3), 2598–2610 (2020)
27. Klemmer, K., Rolf, E., Robinson, C., Mackey, L., Rußwurm, M.: Towards global, general-purpose pretrained geographic location encoders. In: *NeurIPS 2023 Workshop on Distribution Shifts: New Frontiers with Foundation Models* (2023)
28. Kuester, M., Ochoa, T.: Absolute radiometric calibration is an essential tool to imagery science, but what is it? <https://blog.maxar.com/tech-and-tradecraft/2020/absolute-radiometric-calibration-is-an-essential-tool-to-imagery-science-but-what-is-it> (Feb 2020)

29. Lam, D., Kuzma, R., McGee, K., Dooley, S., Laielli, M., Klaric, M., Bulatov, Y., McCord, B.: xview: Objects in context in overhead imagery. arXiv preprint arXiv:1802.07856 (2018)
30. Li, X., Wen, C., Hu, Y., Yuan, Z., Zhu, X.X.: Vision-language models in remote sensing: Current progress and future trends. *IEEE Geoscience and Remote Sensing Magazine* (2024)
31. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: *International Conference on Learning Representations* (2019)
32. Mai, G., Lao, N., He, Y., Song, J., Ermon, S.: Csp: Self-supervised contrastive spatial pre-training for geospatial-visual representations. In: *International Conference on Machine Learning*. PMLR (2023)
33. Mall, U., Hariharan, B., Bala, K.: Change-aware sampling and contrastive learning for satellite images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5261–5270 (2023)
34. Manas, O., Lacoste, A., Giró-i Nieto, X., Vazquez, D., Rodriguez, P.: Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9414–9423 (2021)
35. Mendieta, M., Han, B., Shi, X., Zhu, Y., Chen, C.: Towards geospatial foundation models via continual pretraining. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 16806–16816 (2023)
36. Mu, N., Kirillov, A., Wagner, D., Xie, S.: Slip: Self-supervision meets language-image pre-training. In: *European Conference on Computer Vision*. pp. 529–544. Springer (2022)
37. Noman, M., Naseer, M., Cholakkal, H., Anwer, R.M., Khan, S., Khan, F.S.: Rethinking transformers pre-training for multi-spectral satellite imagery. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 27811–27819 (2024)
38. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
39. Pantazis, O., Brostow, G.J., Jones, K.E., Mac Aodha, O.: Focus on the positives: Self-supervised learning for biodiversity monitoring. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10583–10592 (2021)
40. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International Conference on Machine Learning*. pp. 8748–8763. PMLR (2021)
41. Reed, C.J., Gupta, R., Li, S., Brockman, S., Funk, C., Clipp, B., Keutzer, K., Candido, S., Uyttendaele, M., Darrell, T.: Scale-mae: A scale-aware masked autoencoder for multiscale geospatial representation learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4088–4099 (2023)
42. Reed, C.J., Yue, X., Nrusimha, A., Ebrahimi, S., Vijaykumar, V., Mao, R., Li, B., Zhang, S., Guillory, D., Metzger, S., et al.: Self-supervised pretraining improves self-supervised pretraining. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 2584–2594 (2022)
43. Rolf, E., Klemmer, K., Robinson, C., Kerner, H.: Mission critical – satellite data is a distinct modality in machine learning. arXiv preprint arXiv:2402.01444 (2024)
44. Schmitt, M., Hughes, L.H., Qiu, C., Zhu, X.X.: Sen12ms—a curated dataset of georeferenced multi-spectral sentinel-1/2 imagery for deep learning and data fusion. arXiv preprint arXiv:1906.07789 (2019)

45. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C.W., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S.R., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: LAION-5b: An open large-scale dataset for training next generation image-text models. In: Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (2022)
46. Sumbul, G., Charfuelan, M., Demir, B., Markl, V.: Bigearthnet: A large-scale benchmark archive for remote sensing image understanding. In: IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium. pp. 5901–5904. IEEE (2019)
47. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J.: Yfcc100m: The new data in multimedia research. *Communications of the ACM* **59**(2), 64–73 (2016)
48. Tuia, D., Roscher, R., Wegner, J.D., Jacobs, N., Zhu, X.X., Camps-Valls, G.: Toward a collective agenda on ai for earth science data analysis. *IEEE Geoscience and Remote Sensing Magazine* **9**(2), 88–104 (2021)
49. Van Etten, A., Lindenbaum, D., Bacastow, T.M.: Spacenet: A remote sensing dataset and challenge series. *arXiv preprint arXiv:1807.01232* (2018)
50. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
51. Vivanco Cepeda, V., Nayak, G.K., Shah, M.: Geoclip: Clip-inspired alignment between locations and images for effective worldwide geo-localization. *Advances in Neural Information Processing Systems* **36** (2024)
52. Wang, Q., Liu, S., Chanussot, J., Li, X.: Scene classification with recurrent attention of vhr remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing* **57**(2), 1155–1167 (2018)
53. Wang, Y., Albrecht, C.M., Braham, N.A.A., Mou, L., Zhu, X.X.: Self-supervised learning in remote sensing: A review. *IEEE Geoscience and Remote Sensing Magazine* **10**(4), 213–247 (2022)
54. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3733–3742 (2018)
55. Yang, Y., Newsam, S.: Bag-of-visual-words and spatial extensions for land-use classification. In: *Proceedings of the 18th SIGSPATIAL international conference on advances in geographic information systems*. pp. 270–279 (2010)
56. Zhai, M., Salem, T., Greenwell, C., Workman, S., Pless, R., Jacobs, N.: Learning geo-temporal image features. In: *British Machine Vision Conference* (2019)
57. Zhang, T., Gao, P., Dong, H., Zhuang, Y., Wang, G., Zhang, W., Chen, H.: Consecutive pre-training: A knowledge transfer learning strategy with relevant unlabeled data for remote sensing domain. *Remote Sensing* **14**(22), 5675 (2022)
58. Zhang, X., Lv, Y., Yao, L., Xiong, W., Fu, C.: A new benchmark and an attribute-guided multilevel feature representation network for fine-grained ship classification in optical remote sensing images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **13**, 1271–1285 (2020)
59. Zheng, C., Shrivastava, A., Owens, A.: Exif as language: Learning cross-modal associations between images and camera metadata. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6945–6956 (2023)
60. Zheng, X., Kellenberger, B., Gong, R., Hajnsek, I., Tuia, D.: Self-supervised pre-training and controlled augmentation improve rare wildlife recognition in uav im-

- ages. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 732–741 (2021)
61. Zhou, Y., Feng, L., Ke, Y., Jiang, X., Yan, J., Yang, X., Zhang, W.: Towards vision-language geo-foundation model: A survey. arXiv preprint arXiv:2406.09385 (2024)
 62. Zhu, X.X., Hu, J., Qiu, C., Shi, Y., Kang, J., Mou, L., Bagheri, H., Haberle, M., Hua, Y., Huang, R., et al.: So2sat lcz42: A benchmark data set for the classification of global local climate zones [software and data sets]. IEEE Geoscience and Remote Sensing Magazine **8**(3), 76–89 (2020)