

# Rethinking Transformers Pre-training for Multi-Spectral Satellite Imagery

Mubashir Noman<sup>1</sup> Muzammal Naseer<sup>1</sup> Hisham Cholakkal<sup>1</sup> Rao Muhammad Anwar<sup>1</sup>  
Salman Khan<sup>1,2</sup> Fahad Shahbaz Khan<sup>1,3</sup>

<sup>1</sup>Mohamed bin Zayed University of AI <sup>2</sup>Australian National University

<sup>3</sup>Linköping University

## Abstract

*Recent advances in unsupervised learning have demonstrated the ability of large vision models to achieve promising results on downstream tasks by pre-training on large amount of unlabelled data. Such pre-training techniques have also been explored recently in the remote sensing domain due to the availability of large amount of unlabelled data. Different from standard natural image datasets, remote sensing data is acquired from various sensor technologies and exhibit diverse range of scale variations as well as modalities. Existing satellite image pre-training methods either ignore the scale information present in the remote sensing imagery or restrict themselves to use only a single type of data modality. In this paper, we re-visit transformers pre-training and leverage multi-scale information that is effectively utilized with multiple modalities. Our proposed approach, named SatMAE++, performs multi-scale pre-training and utilizes convolution based upsampling blocks to reconstruct the image at higher scales making it extensible to include more scales. Compared to existing works, the proposed SatMAE++ with multi-scale pre-training is equally effective for both optical as well as multi-spectral imagery. Extensive experiments on six datasets reveal the merits of proposed contributions, leading to state-of-the-art performance on all datasets. SatMAE++ achieves mean average precision (mAP) gain of 2.5% for multi-label classification task on BigEarthNet dataset. Our code and pre-trained models are available at [https://github.com/techmn/satmae\\_pp](https://github.com/techmn/satmae_pp).*

## 1. Introduction

Remote sensing employs a wide range of sensor technologies to acquire data for earth observation and monitoring through satellites and aircrafts. The acquired data may differ in terms of Ground Sample Distance (GSD) based on sensor technology and altitude. GSD refers to the distance between the two adjacent pixels measured on ground i.e.,

GSD of 0.3 meter in an image means adjacent pixels in the image are 0.3 meter apart on the ground. Accordingly, an image with size of  $10 \times 10$  pixels may span over the city depending on the GSD of the image. Therefore, scale of the objects may vary considerably within a single image. In addition, multi-spectral (Sentinel-2) data captured from satellite utilize different sensors to acquire data. Consequently, multi-spectral data may possess different GSDs within the single image. Various bands combination in multi-spectral data can be used to highlight different information. For example, short wave and near infrared bands combination may be used for crop and agriculture monitoring. Therefore, it is desired to utilize the multi-scale information from various sensor's data for remote sensing tasks.

Even though the utilization of pre-trained models on large volume of data for remote sensing applications are increasingly popular, multi-scale information in remote sensing data is scarcely exploited by the vision community. Recently, self supervised learning on remote sensing data is widely explored [2, 16–18]. SatMAE [5] demonstrate the effectiveness of pre-training the transformers on large amount of data for various downstream remote sensing tasks. However, SatMAE [5] does not exploit the multi-scale information present in the remote sensing satellite imagery and is not able to generalize over domains having images with multiple scales. ScaleMAE [20] is a newly introduced framework that proposes a strategy to encode the multi-scale information present in optical remote sensing data. The authors propose a GSD based positional encoding (GSDPE) to inform the model about the position and scale of the patch tokens. However, the proposed GSDPE is constrained to be utilized with the RGB (optical) images only. As multi-spectral (Sentinel-2) images have different GSD resolutions for different channels (see Tab. 1 for more description) and ScaleMAE requires image channels to be stacked together [20], therefore, GSDPE cannot be utilized with multi-spectral data. Additionally, ScaleMAE introduced a sophisticated Laplacian pyramid based decoder enabling the model to learn multi-scale representations. Alternatively, SatMAE shows that the standard sinusoidal po-

sition encodings can be easily extended for multi-spectral data. We argue that the MAE can still learn better multi-scale representations without using GSDPE and sophisticated Laplacian based decoder. To this end, we propose a framework SatMAE++ that uses the standard position encoding and reconstruct the image at multiple scale levels by utilizing convolution based upsampling blocks. In general, our contributions are:

- We empirically demonstrate that the standard position encodings along with multi-scale reconstruction encourage the model to learn better feature representations.
- Based on our observations, we propose a simple yet effective multi-scale pre-training approach, named SatMAE++, that achieves impressive performance on both optical as well as multi-spectral satellite imagery.
- We perform extensive experimentation on six datasets to validate the efficacy of our multi-scale pre-training framework. Our SATMAE++ sets a new state-of-the-art on all six datasets. On the downstream task of land cover classification, our SatMAE++ achieves an absolute gain of 3.6% over the baseline [5].

## 2. Related Work

### 2.1. Representation Learning for Satellite Imagery

Recently, self-supervised learning has been widely utilized as pre-training step for various remote sensing tasks enabling the model to learn rich feature representation from unlabelled data. GASSL [2] and SeCo [16] employ the contrastive learning strategy to demonstrate the effectiveness of pre-training on different downstream tasks. Neumann et. al. [18] explore the idea of in-domain representation learning for remote sensing imagery. Mendieta et. al. [17] investigated the continual learning strategy and mask image modeling technique in their work. SatMAE [5] utilized the mask image modeling for multi-spectral, temporal and optical remote sensing data. ScaleMAE [20] introduces a ground sample distance (GSD) based positional encoding for optical satellite imagery to encode the scale information present in the remote sensing data. While ScaleMAE has shown promising results on optical satellite data, its complicated GSD positional encoding limits its ability to be utilized for multi-spectral data. In contrast, we simplify the idea of extracting multi-scale information in optical and multi-spectral satellite data.

### 2.2. Multi-scale Information

Visual images typically contain objects of various sizes and scales. It is therefore desired to utilize the multi-scale information to learn better semantic representations. Convolutional neural networks and transformers [1, 7, 9, 10, 14, 15, 19, 22] have employed the multi-scale information and show promising results on classification, detection, and

Table 1. Sentinel-2 bands description. The multi-spectral data contains three different GSD resolutions ranging from 10m to 60m. Different band combinations can be utilized to highlight specific information in the multi-spectral image.

Band	Description	GSD (m)	Central Wavelength (nm)
B1	Ultra Blue (Aerosol)	60	443
B2	Blue	10	490
B3	Green	10	560
B4	Red	10	665
B5	Red Edge 1	20	705
B6	Red Edge 2	20	740
B7	Red Edge 3	20	783
B8	Near Infra Red	10	842
B8A	Red Edge 4	20	865
B9	Water Vapor	60	940
B10	Cirrus	60	1375
B11	Short Wave Infra Red 1	20	1610
B12	Short Wave Infra Red 2	20	2190

segmentation tasks. ConvMAE [8] introduces a hierarchical masking strategy to learn multi-scale features for a hybrid convolution-transformer encoder. Point-M2AE [24] proposes a multi-scale auto encoder for 3D point clouds. ScaleMAE [20] propose a Laplacian based decoder to learn multi-scale information. We rationalize the method and introduce an extensible convolution based upsampling block to reconstruct the feature maps at multiple scales.

## 3. Background

**Masked Auto-Encoder (MAE)** takes an input  $I \in \mathbb{R}^{C \times H \times W}$ , and resizes the image into sequence  $S$  of non-overlapping patches having patch size  $P$ , where  $S \in \mathbb{R}^{N \times P^2 C}$  and  $N = (\frac{H}{P} \cdot \frac{W}{P})$  is the number of patches. A patch embedding layer  $f : \mathbb{R}^{P^2 C} \rightarrow \mathbb{R}^D$  is utilized to obtain the sequence  $\tilde{S} \in \mathbb{R}^{N \times D}$  of embedded tokens. A fraction  $m$  of the  $N$  tokens are masked and the remaining  $1 - m$  tokens are input to the transformer encoder that utilizes the position embedding to capture the spatial information of the patches in the image. In decoder, the visible tokens are placed back to their original sequence positions and learnable mask tokens are appended to obtain  $N$  tokens. The positional embeddings are added to the  $N$  tokens and fed to a series of transformer blocks. The decoder finally reconstruct the image by utilizing the mean squared error loss.

**Positional Encodings** Transformers utilize the positional encodings to learn the spatial locations of the patch tokens. In general, transformers use the sinusoidal positional encodings given as:

$$\begin{aligned}
 v_x(pos, 2i) &= \sin\left(\frac{k}{\Omega^{\frac{2i}{d}}}\right) \\
 v_y(pos, 2i + 1) &= \cos\left(\frac{k}{\Omega^{\frac{2i}{d}}}\right)
 \end{aligned} \tag{1}$$

where  $pos$  is the position,  $i$  is feature dimension index,  $d$  is the total possible positions, and  $\Omega$  is a large constant.

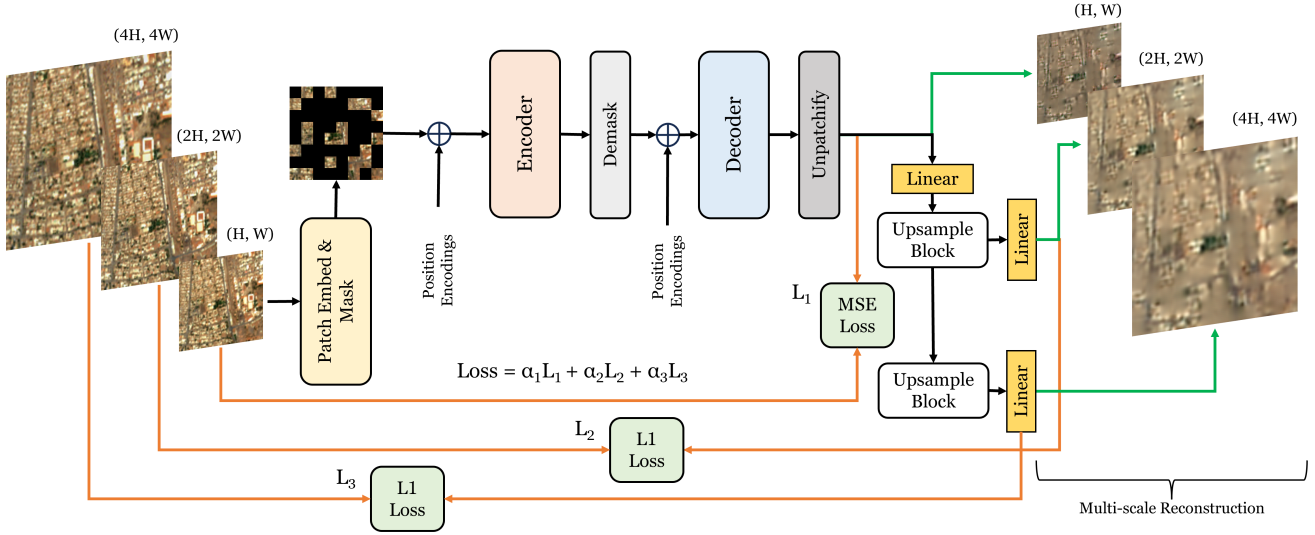


Figure 1. Illustration of Mask Autoencoder (MAE) framework for SatMAE++. The input image having spatial resolution of  $(4H, 4W)$  is downsampled twice to obtain the images of resolution  $(2H, 2W)$  and  $(H, W)$ , respectively. We then feed the image with resolution  $(H, W)$  to the MAE similar to the SatMAE [5] framework. The decoder reconstructs the image at the resolution  $(H, W)$  and apply MSE loss to measure reconstruction quality. The reconstructed output is projected back to the feature space and is upsampled through upsampling blocks to obtain features at  $(2H, 2W)$  and  $(4H, 4W)$  resolutions. The upsampled outputs are projected back to the image space and L1 loss is utilized to penalize the reconstructions at higher resolutions. The overall loss is the weighted mean of all the losses.

## 4. Method

### 4.1. Baseline Framework

We adapt the recent SatMAE [5] framework as our baseline. SatMAE follows the vanilla MAE based architecture for RGB data as discussed in section 3. To deal with multi-spectral data, SatMAE performs channel grouping based on the GSDs of the multi-spectral channels. To make the training stable and reduce the memory requirements of the self attention operation, spatial size of the input is reduced to  $96 \times 96$  instead of the standard input size of  $224 \times 224$  [5]. Accordingly, patch size is decreased to 8. Then it creates separate patch embeddings layers and concatenate them in the spatial dimension. SatMAE modifies the position encoding of MAE to incorporate the spectral information. To this end, separate encoding for each channel group is created and concatenated to the  $x_{pos,i}, y_{pos,i}$  positional encodings such that the final encoding dimension is  $D$ . Afterwards, it passes the features to the transformer encoder. Similar to the vanilla MAE, the decoder takes the output of transformer encoder, place the visible tokens to their original position and learnable mask tokens are appended to obtain the  $N$  tokens. Subsequently, the spectral encodings are added to the patch tokens and feed to the series of transformer blocks. Finally, mean squared error (MSE) loss is utilized to measure the reconstruction quality.

**Limitations:** Although the above baseline framework operates well on RGB and multi-spectral data, however it does not exploit the multi-scale information present in the re-

mote sensing data. In addition, the above baseline framework strives to generalize across domains having images at multiple scale levels. Furthermore, ScaleMAE [20] shows that the scale information can be encoded through a GSD based positional embeddings. However, the GSD based positional embeddings introduced by ScaleMAE can only be used with the RGB data that has same GSD resolution for each channel [20]. Therefore, it is desired to rethink the design of the SatMAE framework to learn the multi-scale information that is not constrained to single data modality.

### 4.2. Overall Architecture

Fig. 1 illustrates the overall framework of SatMAE++ that overcomes the limitations of baseline framework SatMAE [5] and other recent approaches for multi-scale pre-training on multi-spectral (fMoW-Sentinel) as well as RGB data (fMoW-RGB). We take input image at most of three scale levels and feed the image at lowest scale level to the SatMAE framework. The base framework takes the input image, apply the patch embedding and masking operations, and feed the visible tokens to the transformer encoder. Subsequently, decoder takes the output of the encoder and reconstructs the image having same spatial dimension as the lowest scale level input. The reconstructed output from the SatMAE model is utilized by the upsample blocks to perform fine grained reconstruction at higher scale levels. Reconstruction at higher scales encourage the model to learn multi-scale representations thereby improving the performance on various downstream tasks.

**Upsample Block:** UpsampleFig. 2 shows the architecture of the upsample block. It takes input features  $X \in \mathbb{R}^{C \times H \times W}$  and pass them through a transpose convolution layer to upsample the spatial resolution of the features. We normalize the upsampled features and apply the leaky relu activation operation afterwards. Then, a residual block, comprising of two convolution layers, is utilized for local feature enhancement. The enhanced features  $\tilde{X} \in \mathbb{R}^{C \times 2H \times 2W}$  are projected back to the spatial domain by using a linear projection layer and mean absolute error is utilized to compute the reconstruction error between the input and the reconstructed image. Next, we discuss in detail the multi-scale reconstruction procedure at two and three scale levels.

#### 4.2.1 Reconstruction at Two Scale Levels

Let  $I \in \mathbb{R}^{C \times H \times W}$  be an input image to the MAE. We take an image  $\hat{I}$  of resolution  $\mathbb{R}^{C \times 2H \times 2W}$  and down-scale it to obtain the image  $I$  of size  $\mathbb{R}^{C \times H \times W}$ . The image  $I$  is fed into the MAE which first utilizes the patch embedding layers to patchify the input image. In case of multi-spectral input, separate patch embedding layers is used for different groups of band channels. Afterwards, patch tokens of different groups are concatenated along the spatial dimension. We mask the 75% of the patch tokens similar to the other MAE methods. Then, positional encodings are added to the visible patch tokens.

Following [5], we utilize the general positional encodings (as illustrated in Eq. 1) that does not depend on the GSD information. The visible patch tokens are fed to a series of transformer blocks that produce the encoded visible features. Similar to the SatMAE, decoder takes the encoded visible features from the transformer encoder and apply a linear projection to reduce the embedding dimensions. Then, visible features are placed back to their original index positions and learnable mask tokens are appended to the visible tokens. Afterwards, RGB or multi-spectral positional encodings are added to the patch tokens. Eventually, the patch tokens are fed to the decoder transformer and a final projection layer maps the decoded features back to the spatial domain. The decoded image  $F$  is reshaped to the original input dimensions and mean squared error is utilized to compute reconstruction quality.

$$L_1 = \frac{1}{n} \sum_{i=1}^n (F - I)^2 \quad (2)$$

After obtaining the reconstructed input from the decoder, we use linear projection to map the reconstructed image back to the feature space. We then utilize transpose convolution to upsample the feature maps at a resolution of  $(2H \times 2W)$ . The upsampled feature maps are passed from a residual block that is composed of two convolution layers. Finally, we project the features back to the image space to

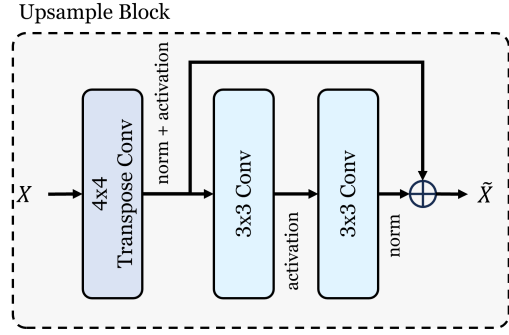


Figure 2. The illustration of upsample block used in SatMAE++ framework. Input features  $X$  are upsampled by utilizing the transpose convolution operation. Afterwards, a residual block which is composed of two convolution layers is employed to enhance the upsampled features given as  $\tilde{X}$ .

obtain the scaled reconstructed image  $\hat{F}$  and apply the L1 loss to analyze the reconstruction performance of model at a higher scale.

$$L_2 = \frac{1}{n} \sum_{i=1}^n |\hat{F} - \hat{I}| \quad (3)$$

Likewise super-resolution methods, we utilize the L1 loss at higher scales so that the model can learn to reconstruct the actual image values. Overall loss is given as:

$$Loss = \alpha_1 L_1 + \alpha_2 L_2 \quad (4)$$

#### 4.2.2 Reconstruction at Three Scale Levels

For multi-spectral data, we reconstruct the image at three scale levels as the resolution of the model input is smaller compared to the RGB data (RGB uses  $224 \times 224$  pixels). Here, we take image  $\bar{I}$  at a higher resolution of  $\mathbb{R}^{C \times 4H \times 4W}$ . We then down sample the image  $\bar{I}$  twice to obtain the images  $\hat{I} \in \mathbb{R}^{C \times 2H \times 2W}$  and  $I \in \mathbb{R}^{C \times H \times W}$ , respectively. As discussed in the previous section, we reconstruct the image  $F$  at spatial resolution of  $(H, W)$ . Afterwards, a linear projection layer is applied to project the data to feature space and input it to the upsample block. The upsample block uses the transpose convolution to increase the spatial resolution by two times which is then fed to the residual convolution block to obtain features  $\hat{F}$  having spatial resolution  $(2H, 2W)$ . The features  $\hat{F}$  are input to another upsample block to obtain features  $\bar{F}$  with dimensions  $(4H, 4W)$ . Both features  $\hat{F}$  and  $\bar{F}$  are projected back to the image space and L1 loss is applied to measure reconstructed features quality. The overall loss is the weighted average of the three losses given as:

$$L_3 = \frac{1}{n} \sum_{i=1}^n |\bar{F} - \bar{I}| \quad (5)$$

$$Loss = \alpha_1 L_1 + \alpha_2 L_2 + \alpha_3 L_3$$

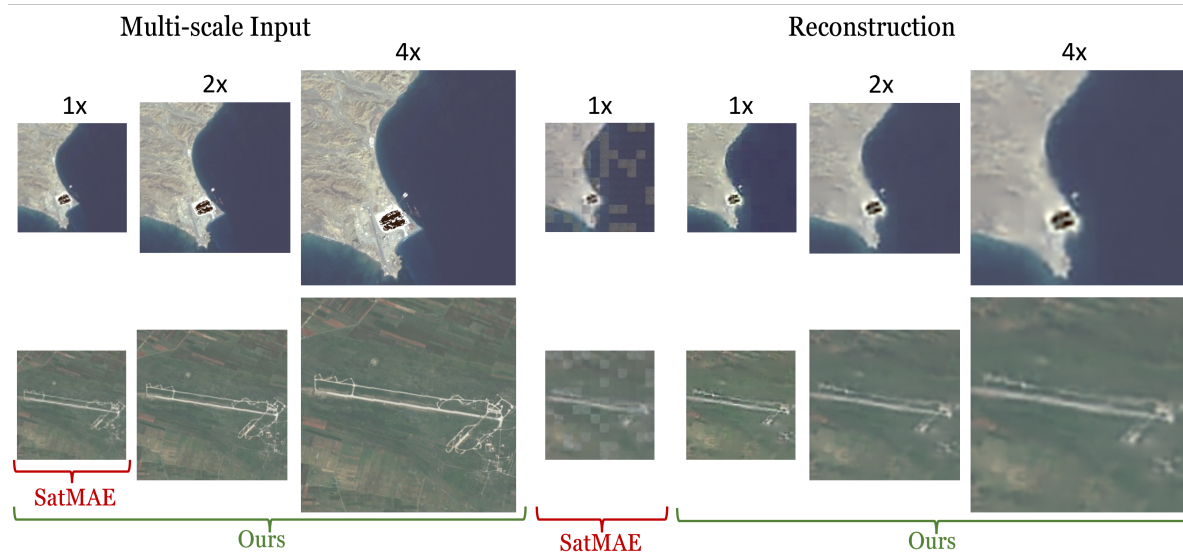


Figure 3. SatMAE++ reconstruction results at multi-scale level. Examples from fMoW-Sentinel dataset are shown here. For illustration, we show the RGB channels of the multi-spectral data here. The images are reconstructed at resolutions of  $(H, W)$ ,  $(2H, 2W)$ , and  $(4H, 4W)$ , respectively. We observe that the proposed model provide better reconstruction results compared to SatMAE at resolution of  $(H, W)$ .

## 5. Experiments

In this section, we first discuss about the mainstream datasets. Then, we explain the pre-training and finetuning procedures utilized for the mainstream benchmarks. Afterwards, we discuss about the datasets used for transfer learning experiments and its finetuning methodology. We further discuss the performance of learned representations through simplified multi-scale MAE pre-training for SatMAE++ and also present the fine tuning results on downstream tasks.

### 5.1. Pre-training Datasets

We utilized two publicly available large scale datasets for pre-training the vision transformer on multi-spectral and RGB satellite data.

**fMoW-RGB:** Functional Map of the World (fMoW) [4] is a large scale publicly available dataset of high resolution satellite images. The dataset is divided into 62 categories for classification task and comprises of about 363k training and 53k test images.

**fMoW-Sentinel:** SatMAE [5] refines and extended the fMoW-RGB for classification task and includes the Sentinel-2 data for the images. Similar to the fMoW-RGB, this dataset has 62 class categories. The dataset contains greater number of images and comprises of 712874 training, 84939 validation, and 84966 test images.

**Reconstruction Results:** We present the multi-scale reconstruction results on fMoW-Sentinel dataset in Fig. 3. We observe the improvement in reconstruction quality of the image by employing multi-scale pre-training. Furthermore,

Table 2. Comparison of state-of-the-art methods on the validation set of fMoW-RGB dataset. \* represents that we reproduce the results using the public codebase provided by the authors and pre-train the model for 800 epochs on fMoW-RGB, which are consistent to the performance reported in ScaleMAE [20]

Method	Backbone	Top1 Acc.
GASSL [2]	ResNet-50	71.55
MoCo-V2 [11]	ResNet-50	64.34
MAE [12]	ViT-Large	68.4
ConvMAE [8]	Conv ViT-Large	74.1
ScaleMAE [20]	ViT-Large	77.9
SatMAE [5]	ViT-Large	77.84
SatMAE* [5]	ViT-Large	73.87
<b>SatMAE++ (Ours)</b>	ViT-Large	<b>78.14</b>

we compare the reconstruction results of our approach with the baseline approach (SatMAE) in Fig. 4. The reconstruction quality of SatMAE on visible patches is much worse in comparison with the masked patches. However, our multi-scale approach improves the quality of reconstruction for both visible and masked patches. While achieving favorable reconstruction compared to SatMAE (see Fig. 3 and 4), we observe our approach to struggle in some cases to reconstruct fine-grained structural details.

### 5.2. Pre-training and Finetuning on fMoW-RGB

**Pre-training:** Similar to the SatMAE [5] and ScaleMAE [20], we pre-train the ViT-Large [6] model on fMoW-RGB dataset. The configuration of model is same as utilized in




Figure 4. Here, we compare the reconstruction performance of our framework with the baseline SatMAE. We observe that the reconstruction results of SatMAE on visible patches is worse compared to the masked patches. Whereas our framework provides much better results on all the patches including the visible patches. The above reported results demonstrate the effectiveness of multi-scale pre-training framework SatMAE++.

[5] i.e., input image of spatial resolution ( $224 \times 224$ ) and patch size of 16. During pre-training, we resize the shorter side of the image to 448 pixels and then perform random crop of  $448 \times 448$  pixels from the resized image. The image of size  $448 \times 448$  pixels is then downsampled to a lower resolution of  $224 \times 224$  pixels by using bilinear interpolation. We then input the image of size  $224 \times 224$  pixels to the model. We use AdamW optimizer and cosine learning rate scheduler in our pre-training experiments. The initial learning rate is set to  $7e-4$  and batch size of 64 is used for single GPU. Similar to [5] 75% of the patches are masked while pre-training the model. We utilize 8 NVIDIA V100 GPUs to train the model for 800 epochs.

**Finetuning:** We finetune the ViT-Large model by loading the pre-training weights in an end-to-end manner. The initial learning rate is set to  $1e-3$  and batch size of 8 is used for single GPU. We use the AdamW optimizer with cosine scheduler and keep the configurations of augmentations and weight decay similar to the [5]. We finetune the model for 50 epochs on 8 NVIDIA V100 GPUs.

**Discussion:** We present the finetuning results of our approach in Tab. 2. The state-of-the-art ScaleMAE [20] shows an improved performance of 77.9% over the SatMAE [5] by utilizing sophisticated GSD based positional encodings and




Figure 5. Illustration of finetuning convergence on validation set of fMoW-Sentinel dataset. We observe that the model pre-trained with multi-scales achieves faster convergence as compared to the model pre-trained with single or less scales. The model trained with single scale achieves highest score of 61.61 at 20th epoch whereas the model that utilised three scales in pre-training converges earlier and achieves highest score at 12th epoch.

Laplacian decoder. However, our approach outperforms the state-of-the-art ScaleMAE [20] by providing the score of 78.14% without using the sophisticated GSD based positional encodings. Compared to the baseline method [5], our approach provides an absolute gain of 4.27% over the reproduced results and 0.3% over the numbers reported in the paper.

### 5.3. Pre-training and Finetuning on fMoW-Sentinel

**Pre-training:** Following [5], we pre-train the ViT-Large model on fMoW-Sentinel dataset having image size of  $96 \times 96$  and patch size of 8. We employ the SatMAE+Group+IM strategy of grouping the channels and independent masking for the multi-spectral data. All 13 bands of the Sentinel-2 images are not utilized and we discard the channels B1, B9, and B10 during pre-training and finetuning. We follow the channel grouping approach of [5] and create (i) group containing B2, B3, B4, B8 channels (ii) group comprising of B5, B6, B7, B8A channels, and (iii) group of B11, B12 channels, respectively. The groups are selected to have same GSD resolution. During pre-training stage, the shorter side of the image is resized to 384 pixels and random cropping is used to obtain image of size  $384 \times 384$  pixels. We then interpolate  $384 \times 384$  sized image to obtain two down sampled images of sizes  $192 \times 192$  and  $96 \times 96$  pixels, respectively. The resized image having spatial resolution of  $96 \times 96$  is input to the model. The model is pre-trained with a base learning rate of  $1e-4$  and batch size of 8 for 50 epochs on 8 NVIDIA V100 GPUs. The remaining settings of opti-

Table 3. Finetuning results on the validation set of fMoW-Sentinel dataset. \* represents that we reproduce the results using the publicly available codebase provided by the authors and pre-train the model on fMoW-Sentinel for 50 epochs.

Method	Backbone	Top1 Acc.
MoCo-V3	ViT-Base	50.45
MoCo-V3+Group	ViT-Base	51.33
SatMAE* [5]	ViT-Base	60.01
<b>SatMAE++ (Ours)</b>	ViT-Base	<b>62.69</b>
ScaleMAE* [20]	ViT-Large	42.21
SatMAE [5]	ViT-Large	61.48
SatMAE* [5]	ViT-Large	61.61
<b>SatMAE++ (Ours)</b>	ViT-Large	<b>63.23</b>

mizer, learning rate scheduler, weight decay, masking ratio, are kept same as listed in [5].

**Finetuning:** In the finetuning stage, we load the pre-training weights of the ViT-Large model and finetune it for 30 epochs. We used the base learning rate of  $2e-4$ , input size of  $96 \times 96$  pixels and patch size of 8. Following [5], we utilize AdamW optimizer, cosine scheduler and same data augmentations. We employ 8 NVIDIA V100 GPUs for finetuning the model.

**Discussion:** We report the state-of-the-art comparison on fMoW-Sentinel dataset in Tab. 3. Our approach that employs multi-scale pre-training provides significant improvement over the SatMAE [5]. In comparison to the state-of-the-art method SatMAE that utilizes ViT-Large backbone, our approach provides an improved accuracy score of 1.75%. In case of ViT-Base backbone, SatMAE++ achieves a gain of 2.68% in the terms of top1 accuracy compared to the SatMAE method which is even better than the score of SatMAE method obtained by utilizing the ViT-Large backbone. We observe that the multi-scale pre-training encourage the model to learn better feature representations especially when the image resolution varies considerably in the given dataset.

**Ablation Studies:** To demonstrate the efficacy of multi-scale pre-training, we pre-train the ViT-Large model with one, two and three scales. Afterwards, we finetune the model by loading the respective pre-trained weights on fMoW-Sentinel dataset. Tab. 4 shows the results of finetuning by loading the pre-train weights of one, two and three scales. We observe consistent improvement in the performance of the model as pre-training scales are increased. Compared to the single scale level, the improvement of the model is significant when two scales are utilized i.e., a gain of 1.25% is achieved. However, the performance gain is 0.37% when we shift from the two scale to three scale levels. We hypothesize that after a certain number of scale levels, the effect of multi-scale pre-training may be insignificant.

Table 4. Ablation study demonstrating the effect of using multiple scales in pre-training. Our SatMAE++ with three scales provides a superior performance with top-1 accuracy of 63.23%.

Method	Backbone	Pre-training Scales	Top1 Acc.
SatMAE	ViT-Large	1	61.61
SatMAE++	ViT-Large	2	62.86
SatMAE++	ViT-Large	3	<b>63.23</b>

**Convergence Rate:** Fig. 5 shows the convergence of the model pre-trained with multiple scales on fMoW-Sentinel dataset. We observe that the model trained with single scale reaches its highest performance score of 61.61 at 20th epoch. Whereas when multi-scale pre-training is utilized, model converges earlier and achieves its highest performance score of 63.23 on the 12th epoch. From this, we infer that the multi-scale pre-training can uplift the performance of the models and provide faster convergence especially when data has diverse range of scale variations.

#### 5.4. Downstream Datasets

To demonstrate the effectiveness of our pre-training approach, we utilize following datasets corresponding to the land cover and multi-label classification tasks.

**EuroSAT [13]** is a publicly available remote sensing dataset for land use and land cover (LULC) classification. It is categorized into 10 classes and comprises of 27000 images. The dataset is available in both RGB and multi-spectral (Sentinel-2) format. Following [5], we use the data split provided by [18].

**RESISC-45 [3]** is another public remote sensing scene classification dataset comprising of 31500 images having 45 scene classes. We follow the [20] for train / val splits of the dataset.

**UC-Merced [23]** is a public land use remote sensing image dataset which contains 21 scene classes. Each class comprises of 100 images that are manually selected from the US regions. We follow the data split provided by [18] in our experiments.

**BigEarthNet [21]** is a multi-label land cover classification dataset and is publicly available for research purpose. The dataset is composed of 590326 Sentinel-2 images and categorized into 19 classes. Following [5], we use 10% of the train data in our experiments and utilize the data splits available at [18].

#### 5.5. Transfer Learning on Downstream Datasets

We finetune the pre-trained ViT-Large model (SatMAE++) on various downstream remote sensing tasks such as land cover and multi-label classification. SatMAE++ provides consistent improvement on the downstream tasks compared to the other state-of-the-art methods existing in literature.

Table 5. Land cover classification on EuroSAT dataset. † denotes that the model uses multi-spectral data.

Method	Backbone	Top1 Acc.
GASSL [2]	ResNet-18	89.51
SeCo [16]	ResNet-18	93.14
SatMAE [5]	ViT-Large	95.74
SatMAE† [5]	ViT-Large	98.98
<b>SatMAE++ (Ours)</b>	ViT-Large	<b>99.04</b>

Table 6. Finetuning results of land cover classification on RESISC-45 dataset.

Method	Backbone	Top1 Acc.
MAE [12]	ViT-Large	93.3
ConvMAE [8]	ViT-Large	95.0
SatMAE [5]	ViT-Large	94.8
ScaleMAE [20]	ViT-Large	95.7
<b>SatMAE++ (Ours)</b>	ViT-Large	<b>97.48</b>

### 5.5.1 Land Cover Classification

We present the transfer learning experiments for land cover classification task on three publicly available remote sensing datasets including EuroSAT [13], RESISC-45 [3], and UC-Merced [23]. We utilize the configuration settings used in fMoW-RGB finetuning experiment for transfer learning on these datasets.

**Finetuning on EuroSAT:** We present the finetuning results on EuroSAT dataset in Tab. 5. We observe that the multi-scale pre-training approach provides reasonable improvement over other approaches. It is notable that SatMAE++ surpasses the performance score of the state-of-the-art SatMAE without using the multi-spectral information and achieves an accuracy score of 99.01%.

**Finetuning on RESISC-45:** Tab. 6 shows the scene classification performance on RESISC-45 dataset. Among recent methods, ScaleMAE [20] achieves superior performance, however, ViT-Large when finetuned by using the pre-trained weights of our approach, surpasses the ScaleMAE score and achieves an accuracy of 97.48%.

**Finetuning on UC-Merced:** We further report the effectiveness of multi-scale MAE pre-training on another popular LULC dataset (Tab. 7). Here, we report the results of finetuning the ViT-Large model on UC-Merced dataset by loading the pre-trained weights of SatMAE and our approach, respectively. The ViT-Large model provides the accuracy score of 94.05% when utilizes the pre-trained weights of SatMAE. However, the performance of ViT-Large improves by 3.6% approx when finetuning is performed by loading the pre-trained weights of our approach.

Table 7. Finetuning results of land cover classification on UC-Merced dataset.

Method	Backbone	Top1 Acc.
SatMAE [5]	ViT-Large	94.05
<b>SatMAE++ (Ours)</b>	ViT-Large	<b>97.62</b>

Table 8. Multi-Label classification results on BigEarthNet [21] dataset. Following [5], we use mean Average Precision (mAP) metric and newer set of class labels. The reported results utilize 10% of training data. ‡ denotes that model uses RGB bands only.

Method	Backbone	mAP
GASSL [2]	Resnet-50	80.20
SeCo‡ [16]	Resnet-50	82.62
SatMAE [5]	ViT-Large	82.13
<b>SatMAE++ (Ours)</b>	ViT-Large	<b>85.11</b>

### 5.5.2 Multi-Label Classification

Finally, we report the performance of our approach on multi-label classification task. We finetune the ViT-Large model on BigEarthNet [21] dataset by loading the pre-train weights of fMoW-Sentinel dataset. We keep the same configurations for finetuning as utilized in fMoW-Sentinel experiment. As the task is multi-label classification, therefore *soft target cross entropy* loss is replaced with the *multi-label soft margin* loss. Following [5], we use the average precision metric as a performance measure of the model.

Tab. 8 shows the finetuning results on BigEarthNet dataset. Similar to the other downstream tasks, our framework performs favorably in the case of multi-label classification task. Compared to the state-of-the-art frameworks, SeCo [16] performs fairly well and achieves the state-of-the-art score of 82.62%. SatMAE [5] has a slightly lower score of 82.13% as compared to the SeCo [16]. However, our framework provides significant improvement and achieves an average precision score of 85.11%.

## 6. Conclusion

Remote sensing imagery offers a wide range of resolutions and spectral bands having multi-scale information incorporated within it. Existing state-of-the-art methods struggle to effectively utilize the multi-scale information along with the multi-spectral data. We propose a framework, named SatMAE++, to incorporate multi-scale information thereby improving the model performance and achieving faster convergence during finetuning. Our SatMAE++ is easily extensible to multiple scale levels and is not restricted to single type of data modality. Extensive experimentation on several downstream and mainstream datasets reveal the efficacy of our approach. Future work includes extending the proposed multi-scale pre-training for dense prediction tasks.



## References

- [1] Abdulaziz Amer Aleissae, Amandeep Kumar, Rao Muhammad Anwer, Salman Khan, Hisham Cholakkal, Gui-Song Xia, and Fahad Shahbaz Khan. Transformers in remote sensing: A survey. *Remote Sensing*, 15(7), 2023. [2](#)
- [2] Kumar Ayush, Burak Uzcent, Chenlin Meng, Kumar Tanmay, Marshall Burke, David Lobell, and Stefano Ermon. Geography-aware self-supervised learning, 2022. [1](#), [2](#), [5](#), [8](#)
- [3] 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. [7](#), [8](#)
- [4] Gordon Christie, Neil Fendley, James Wilson, and Ryan Mukherjee. Functional map of the world. In *CVPR*, 2018. [5](#)
- [5] Yezhen Cong, Samar Khanna, Chenlin Meng, Patrick Liu, Erik Rozi, Yutong He, Marshall Burke, David B. Lobell, and Stefano Ermon. Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery. In *NeurIPS*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. [5](#)
- [7] H. Fan, B. Xiong, K. Mangalam, Y. Li, Z. Yan, J. Malik, and C. Feichtenhofer. Multiscale vision transformers. In *ICCV*, 2021. [2](#)
- [8] Peng Gao, Teli Ma, Hongsheng Li, Ziyi Lin, Jifeng Dai, and Yu Qiao. Convmoe: Masked convolution meets masked autoencoders, 2022. [2](#), [5](#), [8](#)
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [2](#)
- [10] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask r-cnn. In *ICCV*, 2017. [2](#)
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. [5](#)
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. [5](#), [8](#)
- [13] 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 Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. [7](#), [8](#)
- [14] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. [2](#)
- [15] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. [2](#)
- [16] Oscar Mañas, Alexandre Lacoste, Xavier Giró-i Nieto, David Vazquez, and Pau Rodríguez. Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data. In *ICCV*, 2021. [1](#), [2](#), [8](#)
- [17] Matias Mendieta, Boran Han, Xingjian Shi, Yi Zhu, and Chen Chen. Towards geospatial foundation models via continual pretraining. In *ICCV*, 2023. [2](#)
- [18] Maxim Neumann, Andre Susano Pinto, Xiaohua Zhai, and Neil Houlsby. In-domain representation learning for remote sensing. In *ICLR*, 2020. [1](#), [2](#), [7](#)
- [19] Mubashir Noman, Mustansar Fiaz, Hisham Cholakkal, Salman Khan, and Fahad Shahbaz Khan. Elgc-net: Efficient local-global context aggregation for remote sensing change detection. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–11, 2024. [2](#)
- [20] Colorado J Reed, Ritwik Gupta, Shufan Li, Sarah Brockman, Christopher Funk, Brian Clipp, Kurt Keutzer, Salvatore Candido, Matt Uyttendaele, and Trevor Darrell. Scale-mae: A scale-aware masked autoencoder for multiscale geospatial representation learning. In *ICCV*, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [21] Gencer Sumbul, Marcela Charfuelan, Begum Demir, and Volker Markl. Bigearthnet: A large-scale benchmark archive for remote sensing image understanding. In *IGARSS*, 2019. [7](#), [8](#)
- [22] Jianwei Yang, Chunyuan Li, Xiyang Dai, and Jianfeng Gao. Focal modulation networks. In *NeurIPS*, 2022. [2](#)
- [23] Yi Yang and S. Newsam. Bag-of-visual-words and spatial extensions for land-use classification. In *ACM SIGSPATIAL International Workshop on Advances in Geographic Information Systems*, 2010. [7](#), [8](#)
- [24] Renrui Zhang, Ziyu Guo, Peng Gao, Rongyao Fang, Bin Zhao, Dong Wang, Yu Qiao, and Hongsheng Li. Point-m2ae: Multi-scale masked autoencoders for hierarchical point cloud pre-training. *arXiv preprint arXiv:2205.14401*, 2022. [2](#)