

SpectralGPT: Spectral Remote Sensing Foundation Model

Danfeng Hong, *Senior Member, IEEE*, Bing Zhang, *Fellow, IEEE*, Xuyang Li, Yuxuan Li, Chenyu Li, Jing Yao, *Member, IEEE*, Naoto Yokoya, *Member, IEEE*, Hao Li, *Member, IEEE*, Pedram Ghamisi, *Senior Member, IEEE*, Xiuping Jia, *Fellow, IEEE*, Antonio Plaza, *Fellow, IEEE*, Paolo Gamba, *Fellow, IEEE*, Jon Atli Benediktsson, *Fellow, IEEE*, Jocelyn Chanussot, *Fellow, IEEE*

Abstract—The foundation model has recently garnered significant attention due to its potential to revolutionize the field of visual representation learning in a self-supervised manner. While most foundation models are tailored to effectively process RGB images for various visual tasks, there is a noticeable gap in research focused on spectral data, which offers valuable information for scene understanding, especially in remote sensing (RS) applications. To fill this gap, we created for the first time a universal RS foundation model, named SpectralGPT, which is purpose-built to handle spectral RS images using a novel 3D generative pretrained transformer (GPT). Compared to existing foundation models, SpectralGPT 1) accommodates input images with varying sizes, resolutions, time series, and regions in a progressive training fashion, enabling full utilization of extensive RS big data; 2) leverages 3D token generation for spatial-spectral coupling; 3) captures spectrally sequential patterns via multi-target reconstruction; 4) trains on one million spectral RS images, yielding models with over 600 million parameters. Our evaluation highlights significant performance improvements with pretrained SpectralGPT models, signifying substantial potential in advancing spectral RS big data applications within the field of geoscience across four downstream tasks: single/multi-label scene classification, semantic segmentation, and change detection.

Index Terms—Artificial intelligence, deep learning, downstream, foundation model, tensor masked modeling, progressive, remote sensing, spectral data, transformer.

1 INTRODUCTION

SPECTRAL imaging is capable of capturing a vast array of spectral information, thereby enabling highly accurate

- D. Hong, X. Li, and Y. Li are with the Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing 100094, China, and also with the School of Electronic, Electrical and Communication Engineering, University of Chinese Academy of Sciences, Beijing 100049, China.
- B. Zhang is with the Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing 100094, China, and also with the College of Resources and Environment, University of Chinese Academy of Sciences, Beijing 100049, China.
- C. Li is with the Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing 100094, China, and also with the School of Mathematics, Southeast University, Nanjing 210096, China.
- J. Yao is with the Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing 100094, China.
- N. Yokoya is with the Graduate School of Frontier Sciences, the University of Tokyo, Chiba 277-8561, Japan.
- H. Li is with the Big Geospatial Data Management, Technical University of Munich, Munich 85521, Germany.
- P. Ghamisi is with the Helmholtz-Zentrum Dresden-Rossendorf, Helmholtz Institute Freiberg for Resource Technology, Freiberg 09599, Germany, and also with the Institute of Advanced Research in Artificial Intelligence, Vienna 1030, Austria.
- X. Jia is with the School of Engineering and Information Technology, University of New South Wales, Canberra, ACT 2612, Australia.
- A. Plaza is with the Hyperspectral Computing Laboratory, Department of Technology of Computers and Communications, Escuela Politécnica, University of Extremadura, Cáceres 10003, Spain.
- P. Gamba is with the Department of Electrical, Computer and Biomedical Engineering, University of Pavia, Pavia 27100, Italy.
- J. Benediktsson is with the Faculty of Electrical and Computer Engineering, University of Iceland, Reykjavík 102, Iceland.
- J. Chanussot is with Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, Grenoble 38000, France.
- B. Zhang is the corresponding author (zb@radi.ac.cn).

analysis and recognition of objects and scenes beyond what is possible with RGB data alone [1], [2], [3]. This has made multi/hyper-spectral (MS/HS) remote sensing (RS) data the preferred tool of choice and a key component in a wide range of Earth Observation (EO) applications [4], [5], [6], [7], including land use/land cover mapping, ecosystem monitoring, weather forecasting, energy resource development, biodiversity conservation, and geological exploration.

The rapid expansion in the availability and accessibility of spectral data from RS satellite missions, such as Landsat-8/9, Sentinel-2, Gaofen-1/2/6, and others, has further opened up opportunities for new discoveries and advancements in fields related to EO [8]. Nevertheless, this growth also gives rise to two challenging difficulties that require prompt attention and effective solutions.

- **Limited information extraction and mining capabilities from massive spectral data.** The existing expert-centric and data-driven models have reached their limits and are insufficient for effectively learning visual representations from such vast amounts of spectral RS data. There is an urgent need to create new-generation models to improve the intelligent processing and analysis capabilities of spectral RS big data to a level that matches its volume.
- **Limited prediction and interpretation abilities for downstream EO tasks on a few label cases.** Compared to the availability of spectral RS data, there is a scarcity of corresponding labels at both the pixel and image levels. This shortage of labeled data hinders the application of full supervision in deep learning

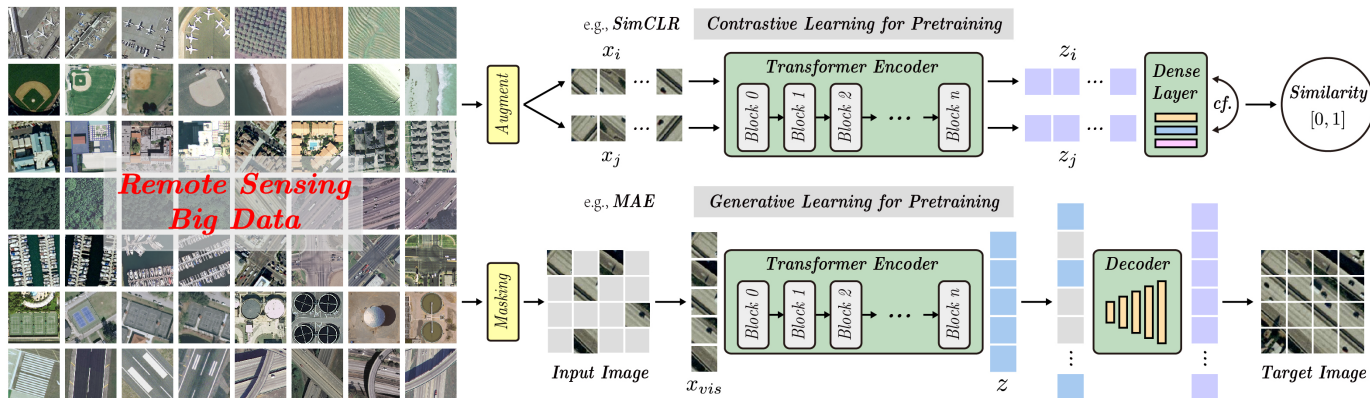


Fig. 1. An illustration to clarify the differences between contrastive learning and generative learning for pretraining in terms of the RS foundation model.

and AI models for practical EO tasks. Urgent action is needed to create RS foundation models embodied with spectral knowledge.

The development of foundation models [9] based on pretraining techniques is currently experiencing a remarkable surge, driven by advancements in self-supervised learning techniques [10] and the transformative capabilities of transformer-based methods [11]. Notably, this surge is particularly evident in the domains of natural language processing and computer vision. The generative pretrained transformer (GPT), deriving out the well-known ChatGPT, is an acknowledged and representative method in the research of foundation models. These pretraining agent tasks are usually divided into contrastive learning [12] and generative learning [13]. As the name suggests, the former aims to teach the model to differentiate between similar and dissimilar examples, while the latter focuses on training a model to generate new data or recover complete data from partial observations. Their differences are illustrated in Fig. 1. Two representative frameworks in contrastive learning are momentum contrast (MoCo) [14] and simple contrastive learning (SimCLR) [15]. MoCo introduces momentum updates to improve the contrastive learning process, while SimCLR leverages data augmentations to enhance the variety and complexity of the image pairs used for contrastive learning. There have been numerous variants of the MoCo and SimCLR frameworks developed since their initial proposals. These variants aim to address specific challenges or limitations of the original frameworks or to further improve their performance. For example, some variants of SimCLR have incorporated new types of data augmentations or improved the training objective [16], [17], [18], while some variants of MoCo have explored different momentum update strategies or used additional losses to improve contrastive learning [19], [20], [21]. Accompanied by the rise of vision transformers (ViT) [22], there has been significant progress in generative learning based on masked image modeling (MIM) for visual pretraining tasks. Bidirectional encoder representation from image transformers (BEiT), as presented in [23], is a prominent example of a MIM architecture built on top of the ViT. MIM allows for the input of all image patches, which provides flexibility for adapting to various network architectures. In [24], a novel

convolution and transformer joint network was proposed, which achieves high-accuracy spectral reconstruction in large-scale complex scenes. Yet the high computational cost associated with MIM can to some extent limit its practical use in certain applications. He *et al.* [25] proposed masked autoencoders (MAE) as a particular alternative to MIM. MAE has emerged as a crucial baseline architecture in the GPT family, which has shown great potential in pretraining tasks. In MAE, unmasked patches or pixels are used to reconstruct those that are masked. This approach is computationally more efficient and also enhances the inference ability of the pretrained models, thereby making it more practical for various applications.

However, these advanced models have been relatively underexplored in RS. Wang *et al.* [26] trained a plain vision transformer with 100 million parameters on RGB images towards RS task design and developed a new rotated varied-size window attention mechanism for fine-tuning the model on downstream tasks. Unlike MAE-based methods that rely on only a few visible image patches to infer the entire image, Sun *et al.* [27] considered all image patches, whether masked or unmasked, by implementing a MIM strategy in their RS pretraining models. Despite the increased computational cost and the reduction in inference efficiency, MIM allows for the flexible use of various deep architectures as network backbones, such as ViT and Swin transformers [28]. The success of these two initial studies indicates the significant potential of pretrained models for applications in RS. The rapid progress in imaging spectroscopy has solidified the significance of spectral RS in EO. This prominence arises from its unique ability to effectively utilize the wealth of spectral information available. However, existing RS foundation models encounter challenges when applied to spectral data due to their limited capacity to model multi-band data. The specific gaps between spectral data and existing foundation models can be summarized as follows.

- **Gap 1 (cf. existing RS foundation models):** They often struggle to capture spatial-spectral representations inherent in 3D tensor data. A majority of these models are predominantly designed for processing data resembling RGB imagery, which constrains their capability to fully capture and characterize spectral

information. Consequently, their applicability to handle such data types remains constrained.

- **Gap 2 (cf. foundation models for video data):** There have been foundation models designed for video data in computer vision [29], [30], yet there are significant differences between video data and spectral data. The main distinctions lie in the varying content between continuous frames in videos and the redundancy that often exists between all frames. As a result, the network designs pretrained for video data are often not well-suited for spectral data.
- **Gap 3 (cf. foundation models for spectral data):** Indeed, recent research about foundation models for spectral data has been relatively scarce. Only one conference paper, namely SatMAE [31], has delved into the utilization of pretrained transformers, e.g., MAE, for spectral satellite images. SatMAE's central design approach involves grouping adjacent spectral bands, akin to RGB bands. However, this practice inadvertently disrupts spectral continuity, leading to a suboptimal capture of 3D spatial-spectral coupling traits and spectrally sequential data. Moreover, constraints pertaining to the number of pretraining samples and effective training strategies have further impeded performance enhancements in this context.

To fill these gaps, we devise SpectralGPT, a groundbreaking RS foundation model meticulously tailored for spectral data. SpectralGPT features pioneering elements, such as a 3D masking strategy, an encoder for learning visual representations from spatial-spectral mixed tokens, and a decoder with multi-target reconstruction for preserving spectrally sequential characteristics. These innovations significantly enhance SpectralGPT's ability to learn intrinsic knowledge representations from spectral data, providing valuable insights for scene understanding in various RS downstream applications. Fig. 2 illustrates a visual overview of SpectralGPT's pretraining and its versatile application in diverse downstream tasks, underscoring its profound contributions.

- **Customized foundation model for spectral data:** SpectralGPT is the first purpose-built foundation model designed explicitly for spectral RS data. SpectralGPT considers unique characteristics of spectral data, i.e., spatial-spectral coupling and spectral sequentiality, in the MAE framework with a simple yet effective 3D GPT network.
- **Large-scale training data:** SpectralGPT is trained on an extensive dataset derived from the Sentinel-2 satellite with over one million spectral images. This effort culminates in the creation of three distinct model iterations—Base, Large, and Huge—comprising approximately 100 million, 300 million, and 600 million parameters, respectively.
- **Flexibility in pretraining:** SpectralGPT employs a progressive training strategy, enabling it to process input images with varying sizes, resolutions, time series, and geographical regions. SpectralGPT is a comprehensive framework and a collective term, referring to the model pretrained exclusively on a single dataset in our case. Correspondingly, the

model progressively pretrained on diverse datasets is denoted as SpectralGPT⁺. This innovative design exposes the model's encoder to a diverse array of information, ultimately enhancing its capability to represent a wide range of features effectively.

- **Advanced 3D masking and reconstruction:** SpectralGPT leverages a 3D tensor-shaped spatial-spectral mask with a masking rate of at least 90% on spectral RS data. Additionally, it employs a groundbreaking multi-target reconstruction strategy to comprehensively capture locally spatial-spectral characteristics and spectrally sequential information. These innovations substantially improve the model's learning capabilities through inference.
- **Superior performance across downstream tasks:** SpectralGPT's impact extends to downstream RS models, where it outperforms existing state-of-the-art (SOTA) competitors across various tasks, including single/multi-label scene classification, semantic segmentation, and change detection.
- **New benchmark dataset:** We curate a new benchmark dataset, named SegMunich, which focuses on urban areas and their adjacent neighborhoods within Munich City, Germany. This dataset is designed to cater to the requirements of semantic segmentation tasks with 13 classes to facilitate downstream analysis.

At the time of submission, the trained model of SpectralGPT and SpectralGPT⁺ has been released at <https://doi.org/10.5281/zenodo.10533809>. Interested users can apply for the SpectralGPT/SpectralGPT⁺ encoded feature maps by emailing authors the selected image files with the required format along with a short description of the intended use. After the review, we will release an open-source implementation of the proposed model along with instructions for training/testing at the same link to further encourage public use.

2 THE PROPOSED SPECTRALGPT

2.1 A Brief Recall of MAE

MAE is a simple autoencoding method [32] that enables the reconstruction of the original signal. Like all autoencoders, MAE includes an encoder that maps the observed signal to a potential representation and a decoder that reconstructs the original signal from the potential representation. However, in contrast to classical autoencoders, MAE uses an asymmetric design that enables the encoder to operate only on partial and observed signals (without mask tokens). Additionally, MAE employs a lightweight decoder to reconstruct the complete signal from potential representations and mask tokens.

In detail, the implementation process of the MAE can be broken down into the following steps:

Step 1. Given the input image \mathbf{x} with $H \times W$ pixels by C dimensions, the strategy in ViT is adapted to divide it into regular, non-overlapping patches with the size of $p \times p \times C$, denoted as $\mathbf{x} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_i, \dots, \mathbf{x}_{\frac{H}{p} \times \frac{W}{p}}\}$.

Step 2. Next, a masking operation is performed on these patches to identify visible (or unmasked) and masked

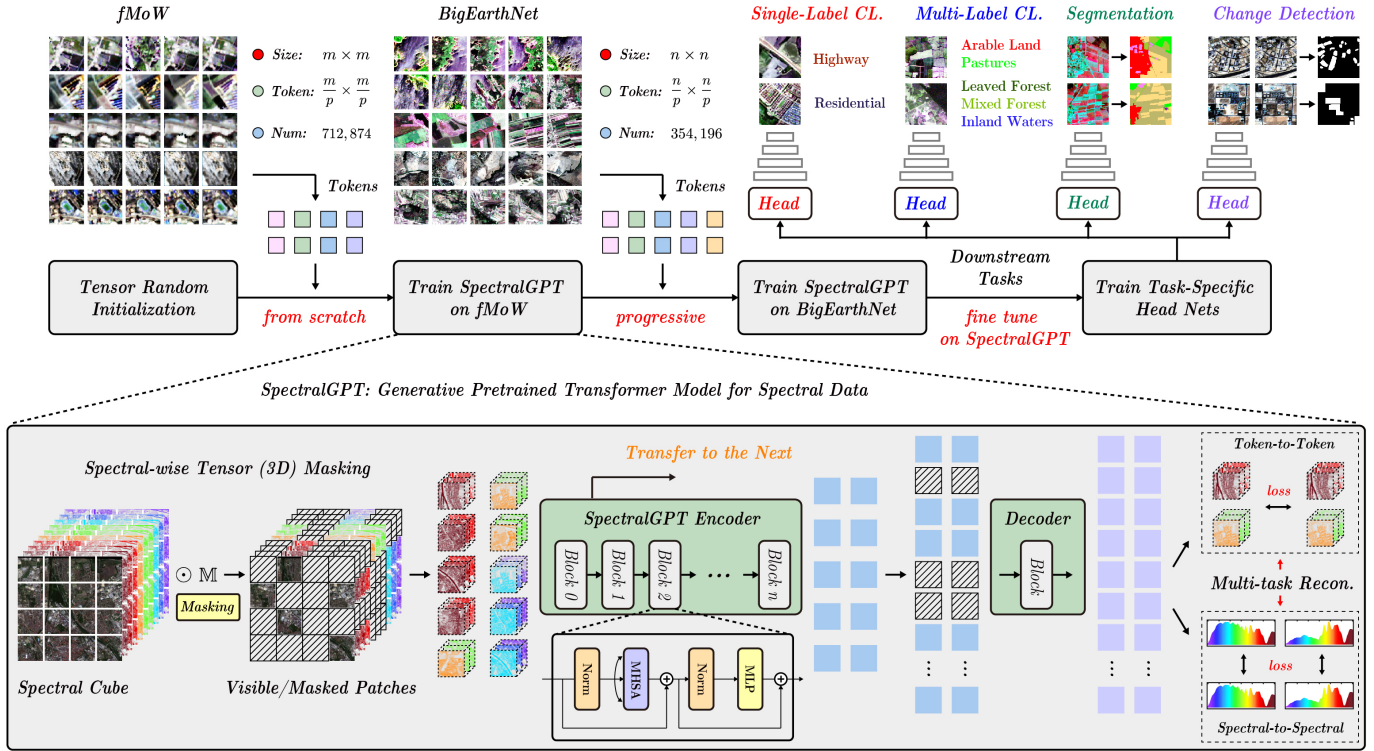


Fig. 2. An illustrative workflow of the proposed SpectralGPT foundation model and its adaptation to downstream tasks. In the pretraining phase, SpectralGPT starts to train the model from scratch on one dataset (e.g., fMoW-S2, with 712,874 images) with a (3D) tensor-based random weight initialization. Subsequently, the model undergoes progressive training on more datasets (e.g., BigEarthNet-S2, with 354,196 images) with varying image sizes, time series information, and geographic regions. SpectralGPT is constructed following the MAE architecture [25] and incorporates spectral-wise tensor (3D) masking, where 90% of the tokens are masked out. For downstream tasks, such as classification, segmentation, and change detection, the pretrained SpectralGPT is connected with task-specific Head networks to be trained and then performs fine-tuning.

patches, i.e., $\mathbf{x}_{vis} = \{\mathbf{x}_i | i \in vis\}$. Only the visible patches are sent into the to-be-learned encoder.

Step 3. The encoder f_{en} is implemented using ViT, where each visible patch is first linearly projected by a shared matrix \mathbf{E}_s , combined with positional embeddings \mathbf{E}_{pos} , and then processed through a series of transformer blocks. Thus, the encoder output in the i -th patch can be expressed as $\mathbf{z}_i = f_{en}(\mathbf{E}_s \mathbf{x}_i + \mathbf{E}_{pos})$.

Step 4. The input to the MAE decoder, denoted by g_{de} , is a complete set of tokens that includes the encoded visible patches and mask tokens (e.g., \mathbf{z}_m). The encoded features, which are the latent representations from the encoder, and the mask tokens are used as inputs and combined with positional embeddings to the lightweight ViT decoder. The final layer of the decoder is a linear projection (e.g., \mathbf{W}) that outputs several channels equal to the number of pixels in a patch. The output is then reshaped to reconstruct the image as $\hat{\mathbf{x}} = g_{de}(\mathbf{W}([\mathbf{z}_{vis}, \mathbf{z}_m]) + pos)$, where \mathbf{z}_{vis} is the encoded representations of visible patches.

Step 5. The loss function used in MAE is the mean squared error (MSE), and it is calculated for the visible and masked patches (similar to BERT [33]), i.e., $\mathcal{L} = \frac{1}{vis+mask} \sum_{i \in vis \cup mask} (\mathbf{x}_i - \hat{\mathbf{x}}_i)^2$.

It is worth noting that a normalization approach is performed, i.e., the mean and standard deviation of pixel values in each patch are calculated, and the patch is normalized accordingly, i.e., $\mathbf{x}_{norm} = \{\frac{\mathbf{x}_i - \mu_i}{\sigma_i} | i \in vis\}$. In this case, the encoder reconstruction task changes to reconstruct

normalized pixel values.

2.2 Methodological Overview of SpectralGPT

Our SpectralGPT model is structured with three key components: 3D masking for processing spectral data, an encoder to learn spectrally visual representations, and a decoder for multi-target reconstruction. What sets our approach apart is a progressive training manner, where the model is trained using diverse types of spectral data. This strategy enhances the proposed SpectralGPT foundation model, endowing it with greater flexibility, robustness, and generalization capabilities. Fig. 2 provides an illustrative workflow of the proposed SpectralGPT with various downstream tasks.

2.3 3D Masking on Spectral Data

Inspired by the spacetime-agnostic sampling in video-like data using the MAE-based framework [30], we model multi-band spectral images as 3D tensor data. To achieve this, we implement a 3D cube masking strategy that enables efficient processing of tensor-like spectral data. Our approach leverages a 90% masking rate to capture spatially and spectrally visual representations in an effective way, leading to more accurate and diverse knowledge extraction from input spectral data.

Given a 3D cube-shaped spectral image $\mathbf{x} \in \mathbb{R}^{H \times W \times D}$, we partition it into non-overlapping 3D tensor tokens along both the spatial and spectral dimensions. Each token has

a size of $p \times p \times k$, where p and k are the token sizes in spatial and spectral dimensions, respectively. Using these settings, we then have $\frac{H}{p} \times \frac{W}{p} \times \frac{D}{k}$ tokens, denoted as $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_{\frac{H}{p} \times \frac{W}{p} \times \frac{D}{k}}\}$. This results in the following visible and masked representations (e.g., \mathbf{x}_{vis} and \mathbf{x}_{mask}) on the spectral data.

$$[\mathbf{x}_{vis}, \mathbf{x}_{mask}] = \mathbb{M} \odot \mathbf{x}, \quad (1)$$

where $\mathbb{M} \in \{0, 1\}^{\frac{H}{p} \times \frac{W}{p} \times \frac{D}{k}}$ is a token-wise binary mask indicating which tokens should be masked, i.e., all pixels in the token are set to zero. Accordingly, \odot denotes the masking operation.

2.4 Encoder for Visible Tokens

Similar to the encoder in MAE, all visible tokens $\{\mathbf{x}_i | i \in vis\}$ spatial-spectral mixed representations are first transformed into feature embeddings using a shared linear projection \mathbf{E}_s . The learned representations via the encoder f_θ with respect to the variable θ are denoted as $f_\theta(\mathbf{E}_s \mathbf{x}_i + \mathbf{E}_{pos})$, where \mathbf{E}_{pos} represents the positional encoding. The encoder f_θ comprises several stacking self-attention (SA) transformer blocks. The SA module used in the encoder can be constructed as follows.

- The input embeddings \mathbf{z}_i are linearly transformed to *query* \mathbf{Q}_i , *key* \mathbf{K}_i , and *value* \mathbf{V}_i embeddings using learnable projection matrices \mathbf{W}_Q , \mathbf{W}_K , and \mathbf{W}_V , respectively.
- The attention scores \mathbf{S}_i between *query* and *key* embeddings are computed as a dot product scaled by $\frac{1}{\sqrt{d}}$ and passed through a softmax function. The resulting scores are then used to weight the *value* embeddings, which are summed to produce the final output embeddings, i.e., $\mathbf{z}_i = \mathbf{S}_i \mathbf{V}_i$. The formula of the original SA's complete process can be written as

$$\begin{aligned} \mathbf{Q}_i &= \mathbf{x}_i \mathbf{W}_Q, \quad \mathbf{K}_i = \mathbf{x}_i \mathbf{W}_K, \quad \mathbf{V}_i = \mathbf{x}_i \mathbf{W}_V, \\ \mathbf{S}_i &= \text{softmax}\left(\frac{\mathbf{Q}_i \mathbf{K}_i^\top}{\sqrt{d}}\right), \\ \mathbf{z}_i &= \text{Attention}(\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i) = \mathbf{S}_i \mathbf{V}_i, \end{aligned} \quad (2)$$

where d denotes the dimension of embeddings.

- Finally, the output features \mathbf{z}_i have the same dimension as \mathbf{x}_i and can be further processed by subsequent encoders.

2.5 Lightweight Decoder with Multi-Target Reconstruction

Given the encoder output features \mathbf{z} , we simultaneously train a lightweight decoder with a multi-target reconstruction strategy with respect to the variable ϕ on \mathbf{z} to recover the original image tokens from potential embeddings of visible and masked image tokens. Mathematically, the reconstructed image tokens $\hat{\mathbf{x}}$ can be formulated as $\hat{\mathbf{x}} = g_\phi(f_\theta(\mathbb{M} \odot \mathbf{x}))$. The decoder g_ϕ is typically narrower and shallower than the encoder, and usually consists of a few transformer blocks and a linear reconstruction layer. The proposed SpectralGPT trains the encoder f_θ and decoder g_ϕ in an end-to-end manner to minimize the reconstruction loss between the reconstructed image tokens $\hat{\mathbf{x}}$ and the original

image tokens \mathbf{x} . In our approach, the reconstruction loss consists of two components: token-to-token and spectral-to-spectral. This multi-target reconstruction allows the learned representations to capture spatial-spectral coupling characteristics and spectrally sequential information effectively. This overall loss \mathcal{L} is quantified mathematically using MSE in the pixel space, defined as follows:

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_{token} + \lambda \mathcal{L}_{spectral} \\ &= \frac{1}{m} \sum_{i \in vis} (\mathbf{x}_i - \hat{\mathbf{x}}_i)^2 + \frac{1}{n} \sum_{j=1}^n (\mathbf{x}_j - \hat{\mathbf{x}}_j)^2 \\ &= \frac{1}{m} \sum_{i \in vis} (\mathbf{x}_i - \hat{\mathbf{x}}_i)^2 \\ &\quad + \frac{1}{n} \sum_{j=1}^n ([\mathbf{x}_{r,c,1}, \dots, \mathbf{x}_{r,c,\frac{D}{k}}]_j - [\hat{\mathbf{x}}_{r,c,1}, \dots, \hat{\mathbf{x}}_{r,c,\frac{D}{k}}]_j)^2, \end{aligned} \quad (3)$$

where m and n represent the number of masked tokens ($\frac{H}{p} \times \frac{W}{p}$) and spectral tokens (each spectral token consists of $\frac{D}{k}$ standard tokens along with spectral dimension), and (r, c) denotes the standard token in the i -th row and the j -th column of the spectral data.

2.6 Progressive Pretraining

The proposed SpectralGPT model has the advantage of being highly adaptable to different input image sizes, which is especially useful for processing large datasets with images of varying size, resolution, temporal variability, and geographical coverage. This is achieved by dividing the input image into fixed-sized 3D tokens (e.g., $8 \times 8 \times 3$), which are then independently processed through the encoder-decoder pipeline. The resulting tokens are then stitched back together to form the final output image. This approach ensures that the model can handle images of arbitrary dimensions in theory, without requiring any changes to the architecture or hyperparameters. With this characteristic, the proposed model allows for feeding varying-sized images into the encoder networks and also might enable the input images with different sensors, resolutions, time series, and modalities, as long as the 3D tokens are cropped to a fixed size.

It is worth emphasizing that the progressive feeding of different types of input images into the networks is not only useful for enabling greater flexibility in the type and size of input images but also for improving the model's ability to extract valuable knowledge from diverse data sources, thereby enhancing model generalization. This can be achieved, for instance, by first inputting images with a size of 96×96 pixels and then progressively feeding in the images with the size of 128×128 pixels, or by starting with Sentinel-2 data and then transitioning to Landsat-8 or Gaofen-2 data. More broadly, the ability to process varying types and sizes of input images can lead to more robust and generalizable features that are not limited to a specific input image type or size, thus improving model generalization and performance on previously unseen data. Moreover, this flexibility in input image size and type is particularly beneficial in practical applications, where input images may come from different sources or have varying

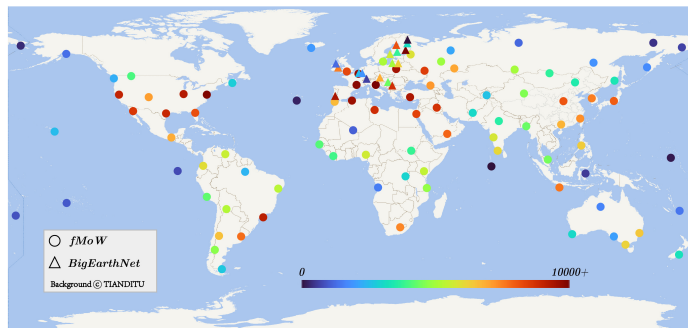


Fig. 3. Sample density distribution collected from Sentinel-2 sources (i.e., fMoW, BigEarthNet) over the Earth’s inhabited areas, amounting to a total of 1,473,105 images.

resolutions, allowing the model to be more adaptable to real-world scenarios where input images can be unpredictable.

The input order of the datasets (i.e., feeding fMoW-S2 prior to BigEarthNet-S2) is arranged in the progressive training strategy. The motivation for such design is motivated by two primary considerations. On one hand, the quality of spectral data in fMoW-S2 is relatively lower than that in BigEarthNet-S2, whereas its scale is much larger. As a result, the model benefits from initiating learning from data with comparatively lower quality but large scale, followed by fine-tuning with higher-quality data. On the other hand, the image size of fMoW-S2 is smaller than that of BigEarthNet-S2. We found that the model tends to initially learn from the small-size data and further extend to the larger one. Furthermore, the order setting not only expedites network convergence but also facilitates higher-efficiency training in networks toward a better solution.

2.7 Pretrained Dataset

Our foundation model is trained on a comprehensive dataset comprising over one million spectral images from the Sentinel-2 satellite. This dataset encompasses 12 spectral bands and draws from two primary sources: fMoW-S2 [31], a globally diverse collection labeled with 62 categories based on the Functional Map of the World (fMoW) [34], and BigEarthNet [35], a regional dataset originating from over ten European countries. To provide an overview of the dataset, Fig. 3 illustrates the distribution of image samples over Earth’s inhabited areas, amounting to a total of 1,473,105 images.

Researchers at Stanford University meticulously curated a dataset by leveraging geo-coordinates and timestamps from the fMoW dataset. This process aims to construct a time series of Sentinel-2 images. To ensure data quality, locations with exclusively pre-Sentinel-2 fMoW images were excluded. For locations with partially preceding fMoW images, selective curation is performed, involving the exclusion of these specific images and the introduction of supplementary captures at 6-month intervals to enrich the temporal sequence. This approach culminates in the creation of the fMoW Sentinel-2 dataset, denoted as fMoW-S2. This dataset predominantly covers fMoW locations and preserves labels mirroring those in the original fMoW dataset.

The fMoW-S2 dataset comprises Sentinel-2 spectral images (B1-12 and B8A) and is partitioned into three sub-

sets: 712,874 training images, 84,939 validation images, and 84,966 test images, totaling 882,779 images. Each image has an average dimension of approximately 45 pixels in height and 60 pixels in width. For additional details about the fMoW-S2 dataset, interested parties can refer to the dedicated website¹.

Furthermore, the study incorporates the BigEarthNet dataset, specifically the BigEarthNet-S2 variant², comprising 590,326 distinct, non-overlapping Sentinel-2 spectral image tokens. The pretraining of different variant models using the proposed method involves the use of 712,874 fMoW-S2 images and 354,196 BigEarthNet-S2 images. Notably, only 10% of the BigEarthNet-S2 images with labels, amounting to 35,420 images, are utilized for fine-tuning in downstream tasks.

2.8 Implementation Details and Experimental Setup

Following established conventions, we acknowledge that Sentinel-2 images comprise 13 spectral bands. However, to harmonize datasets across pretrained and downstream tasks in terms of channel composition, we have opted to retain the 12 dominant bands, excluding band B10, in all fMoW dataset images. To ensure data consistency, we normalize the spectral images band by band, scaling their values to a standardized range of 0 to 1. Subsequently, we follow established methodologies [31] for preprocessing. This involves random image cropping within the range of 0.2x to 1.0x of the original size, resizing them to 96×96 pixels, and applying horizontal flips. These meticulous steps collectively contribute to the robustness and compatibility of our spectral foundation model.

We employ the vanilla ViT-Base architecture as the network backbone. To adapt the model to spectral data, we employ a token size of $8 \times 8 \times 3$ pixels, effectively partitioning the images. For instance, an image with a size of $96 \times 96 \times 12$ pixels is segmented into $12 \times 12 \times 4$ tokens. Drawing inspiration from a prior work [30], our approach incorporates two learnable positional embeddings. One of these embeddings is dedicated to spatial information, while the other is tailored to capture variations across spectral channels. This augmentation further refines the model’s ability to extract meaningful features from the spectral input.

Our pretraining closely adheres to the approach outlined in a prior study [31]. Utilizing the computational power of 8 NVIDIA GeForce RTX 4090 GPUs and AMD EPYC 7Y83 CPU, we implement the AdamW optimizer [36] with a foundational learning rate of 10^{-4} , coupled with a half-cycle cosine decay schedule. To ensure robustness, we adopt a 3D masking ratio of 90%, facilitating effective training. The model undergoes a comprehensive pretraining regimen spanning 200 epochs on the fMoW-S2 dataset. After this phase, the model’s training continues on the BigEarthNet-S2 dataset for 100 epochs. While this phase necessitates a modification in input dimensions to $128 \times 128 \times 12$, the other settings remain consistent. This meticulous strategy effectively enhances the model’s adaptability and performance across diverse datasets.

1. <https://purl.stanford.edu/vg497cb6002>
2. <https://bigearth.net>

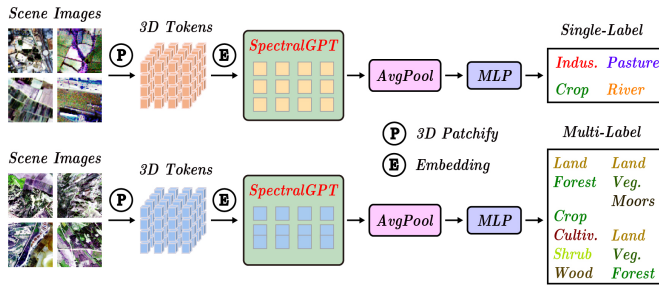


Fig. 4. Network architecture for downstream tasks in terms of single-label (Top) RS scene classification and multi-label (Bottom) RS scene classification by leveraging our pretrained SpectralGPT model. The Avg-Pool and MLP denote the average pooling operation and the multilayer perception, respectively.

3 EXPERIMENTS

In this section, we rigorously evaluate the performance of our SpectralGPT model by benchmarking it against several SOTA foundation models: ResNet50 [37], SeCo [38], ViT [22], and SatMAE [31]. Further, we assess its capabilities across four downstream EO tasks, including single-label scene classification, multi-label scene classification, semantic segmentation, and change detection, as well as extensive ablation studies.

We quantitatively assess the performance of pretrained foundation models across four downstream tasks in terms of recognition accuracy for the single-label RS scene classification task, macro and micro mean average precision (mAP), i.e., macro-mAP (micro-mAP), for the multi-label RS scene classification task, overall accuracy (OA) and mean intersection over union (mIoU) for the semantic segmentation task, and precision, recall, and F1 score for the change detection. Additionally, we conduct insightful ablation studies, exploring critical factors such as masking ratio, decoder depth, model size, patch size, and training epochs. Utilizing the computational power of 4 NVIDIA GeForce RTX 4090 GPUs, we meticulously fine-tune pretrained foundation models for both downstream tasks and ablation studies, thereby offering comprehensive insights into SpectralGPT’s capabilities and adaptability within the RS domain.

3.1 Single-Label RS Scene Classification on EuroSAT

For the downstream single-label RS scene classification task, we employ the EuroSAT dataset [39]. This dataset consists of 27,000 Sentinel-2 satellite images collected from 34 European countries. These images are classified into 10 land use classes, each containing between 2,000 to 3,000 labeled images. Each image in this dataset has a resolution of 64×64 pixels and encompasses 13 spectral bands. It’s worth noting that, for consistency with prior data processing, band B10 has been excluded from all images. Additionally, we follow the train/validation splits as recommended in [40].

On the EuroSAT dataset, these pretrained models undergo fine-tuning, spanning 150 epochs with a batch size of 512. This fine-tuning process employs the AdamW optimizer with a base learning rate of 2×10^{-4} , and it incorporates data augmentations in line with prior work [25], including weight decay (0.05), drop path (0.1), reprob (0.25), mixup (0.8), and cutmix (1.0). The foundational encoder of

TABLE 1

Quantitative results of SOTA pretrained foundation models for the downstream single-label RS scene classification task in terms of recognition accuracy on the EuroSAT dataset. The best result is shown in bold.

Method	Pretrained Dataset	Acc. (%)
ResNet50 [37]	ImageNet-1k	96.72
SeCo [38]	SeCo	97.23
ViT [22]	Random Init.	98.73
ViT-22k [22]	ImageNet-22k	98.91
SatMAE [31]	fMoW-S2	99.09
SpectralGPT	fMoW-S2	99.15
SpectralGPT+	fMoW-S2+BigEarthNet	99.21

the pretrained model is utilized, and its output is passed through an average pooling layer to generate predictions. The training objective is to minimize the cross-entropy loss. Fig. 4 illustrates the network architecture for the downstream single-label scene classification task.

The pretrained model’s encoder serves as the foundational backbone, and its output is subject to an average pooling layer to generate predictions. The training objective involves minimizing the cross-entropy loss. In Table 1, we present a comparative analysis of our proposed method against alternative pretraining models, reporting the highest Top1 accuracy on the validation set. The obtained results highlight the efficacy of the proposed approach, achieving an impressive accuracy of 99.15%. Furthermore, when the model undergoes pretraining on both the fMoW-S2 and BigEarthNet datasets, a noteworthy performance boost is observed, culminating in a remarkable accuracy of 99.21%. This underscores the advantage of leveraging diverse data sources for improved model performance.

3.2 Multi-Label RS Scene Classification on BigEarth-Net

For the multi-label RS scene classification task, we utilize the BigEarthNet-S2 dataset [35]. This extensive dataset consists of 125 Sentinel-2 tiles and comprises 590,326 12-band images that span 19 classes for multi-label classification. The images encompass resolutions ranging from 10 to 60 meters, with 12% of low-quality images being excluded. The training and validation sets align with prior research [40], with 354,196 training samples and 118,065 validation samples. To prepare the images for model training, those of varying resolutions are standardized to a uniform size of 128×128 pixels using bilinear interpolation.

On the BigEarthNet-S2 dataset, these foundation models are fine-tuned using a 10% subset of the training data, following similar settings to those applied in the EuroSAT fine-tuning experiments, except for an increased learning rate of 2×10^{-4} , which is in line with findings from previous research [38], [31]. Most existing methods, including those that use pretrained foundation models, typically utilize all available images for training in the BigEarthNet-S2 dataset. In contrast, our proposed SpectralGPT achieves higher classification performance, even when utilizing only 10% of the training samples. Given the multi-label classification nature of this task, our training objective involves the multi-label soft margin loss, and performance evaluation is based on the

TABLE 2

Quantitative results of SOTA pretrained foundation models for the downstream multi-label RS scene classification task in terms of mean average precision (mAP) on the BigEarthNet dataset. The best result is shown in bold.

Method	Pretrained Dataset	macro-mAP	micro-mAP
ResNet50 [37]	ImageNet-1k	80.76	80.06
SeCo [38]	SeCo	83.16	82.82
ViT [22]	Random Init.	81.57	80.15
ViT-22k [22]	ImageNet-22k	85.08	84.67
SatMAE [31]	fMoW-S2	85.21	84.93
SpectralGPT	fMoW-S2	86.03	85.61
SpectralGPT ⁺	fMoW-S2+BigEarthNet	88.22	87.50

mAP metric. Notably, we calculate mAP using both macro and micro mAP measurements. This approach is particularly relevant for the BigEarthNet-S2 dataset, which exhibits class imbalance. The multi-label classification framework is shown in Fig. 4.

Table 2 presents a comparative analysis of our pretrained model against other proposed pretrained models and models trained from scratch, showcasing the exceptional performance of the proposed approach. In particular, when compared to ViT pretrained on ImageNet-22k and SatMAE, our SpectralGPT model outperforms them by 0.84% (0.82%) and 0.71% (0.68%) in terms of macro-mAP (micro-mAP), respectively. Notably, the introduction of additional pre-training data (BigEarthNet), i.e., SpectralGPT⁺, leads to a significant performance boost, with the model achieving an impressive 88.22% (87.50%) macro-mAP (micro-mAP), surpassing the model solely trained on fMoW-S2 by 2.19% (1.86%). This substantial improvement can be attributed to two key factors. Firstly, the model's initial pretraining on BigEarthNet (even without labels) equips it with a strong grasp of the dataset's distribution, accelerating convergence during fine-tuning and enhancing mAP. Secondly, the adoption of the MIM method as a pretraining pretext task, coupled with a substantial data scale, necessitates alignment with the training strategy, emphasizing the significance of the random masking framework and a 90% masking ratio to facilitate more robust representation learning. Furthermore, as our evaluation focuses on a multi-label classification task and employs only 10% of the training data, the results underscore the superior generalization and few-shot learning capabilities of our proposed model in tackling challenging downstream tasks.

3.3 RS Semantic Segmentation on SegMunich

For the semantic segmentation task, we create a new SegMunich dataset, which is derived from the Sentinel-2 spectral satellite [42]. This dataset consists of a 10-band best-pixel composite with dimensions of $3,847 \times 2,958$ pixels and a spatial resolution of 10 meters. It captures Munich's urban landscape over a span of three years up to April 2020 and includes a segmentation mask that meticulously delineates 13 Land Use and Land Cover (LULC) classes. The data for this mask is sourced from various places, including OpenStreetMap for street network data and the OSMLULC

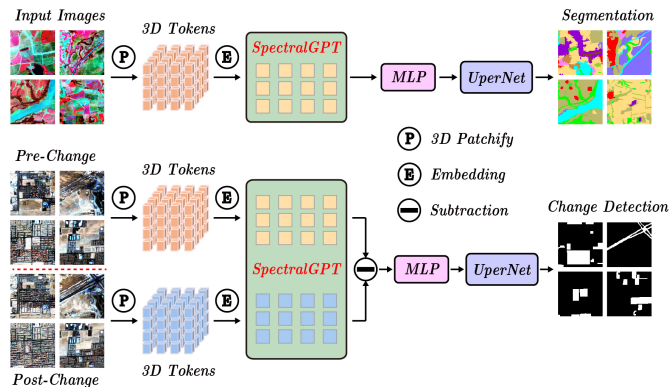


Fig. 5. Network architecture for downstream tasks in terms of semantic segmentation (Top) and change detection (Bottom) by leveraging our pretrained SpectralGPT model and training a follow-up UperNet [41] Head. The MLP denotes the multilayer perception.

platform³ for the remaining 12 classes, all obtained at the same 10-meter spatial resolution. To create a comprehensive feature representation for semantic segmentation, the dataset combines the 10-meter spectral bands (B1, B2, B3, and B4) with resampled 20-meter spectral bands (B5, B6, B7, B8A, B11, B12), which have been upsampled to match the 10-meter resolution. This amalgamation of spectral bands ensures that the dataset provides rich and informative data for the semantic segmentation task.

On the SegMunich dataset, we employ the UperNet framework [41] in conjunction with the pretrained foundation models, initially consolidating the four tokens per pixel from the encoder's final layer into a single token. Image data is divided into 128×128 -pixel tokens with a 50% overlap. The dataset is then split into a training-validation ratio of 8:2 and subjected to data augmentation techniques, including random flips and rotations. During fine-tuning on this dataset, we use a batch size of 96 and set the base learning rate to 5×10^{-4} . The optimization and loss functions remain consistent with those employed in the EuroSAT experiment, ensuring a coherent and uniform approach to model training and evaluation. The segmentation architecture is detailed in Fig. 5.

Table 3 lists quantitative results in terms of OA and mIoU for the semantic segmentation task. Our SpectralGPT (SpectralGPT⁺) outperforms all others, exhibiting a significant lead with a 1.1% (2.3%) higher mIoU than the second-best result (i.e., SatMAE). Fig. 6(a) offers a visual depiction of the Munich area under study for the segmentation task, along with the proportions of the 13 classes. The qualitative comparison in several ROIs, as shown in Fig. 6(b), highlights our model's superior ability to recognize a wider range of land use categories well compared to the competing models in most instances. Furthermore, when considering ViT-22k as the baseline for performance comparison, our model consistently excels across all segmentation classes, as evident in Fig. 6(c), particularly for categories, such as *crops*, *pastures*, *open spaces*, *vegetations*, and others. By amalgamating category statistics with class-wise IoU outcomes, it becomes evident that our SpectralGPT model excels in mitigating

3. <https://osmlanduse.org/>

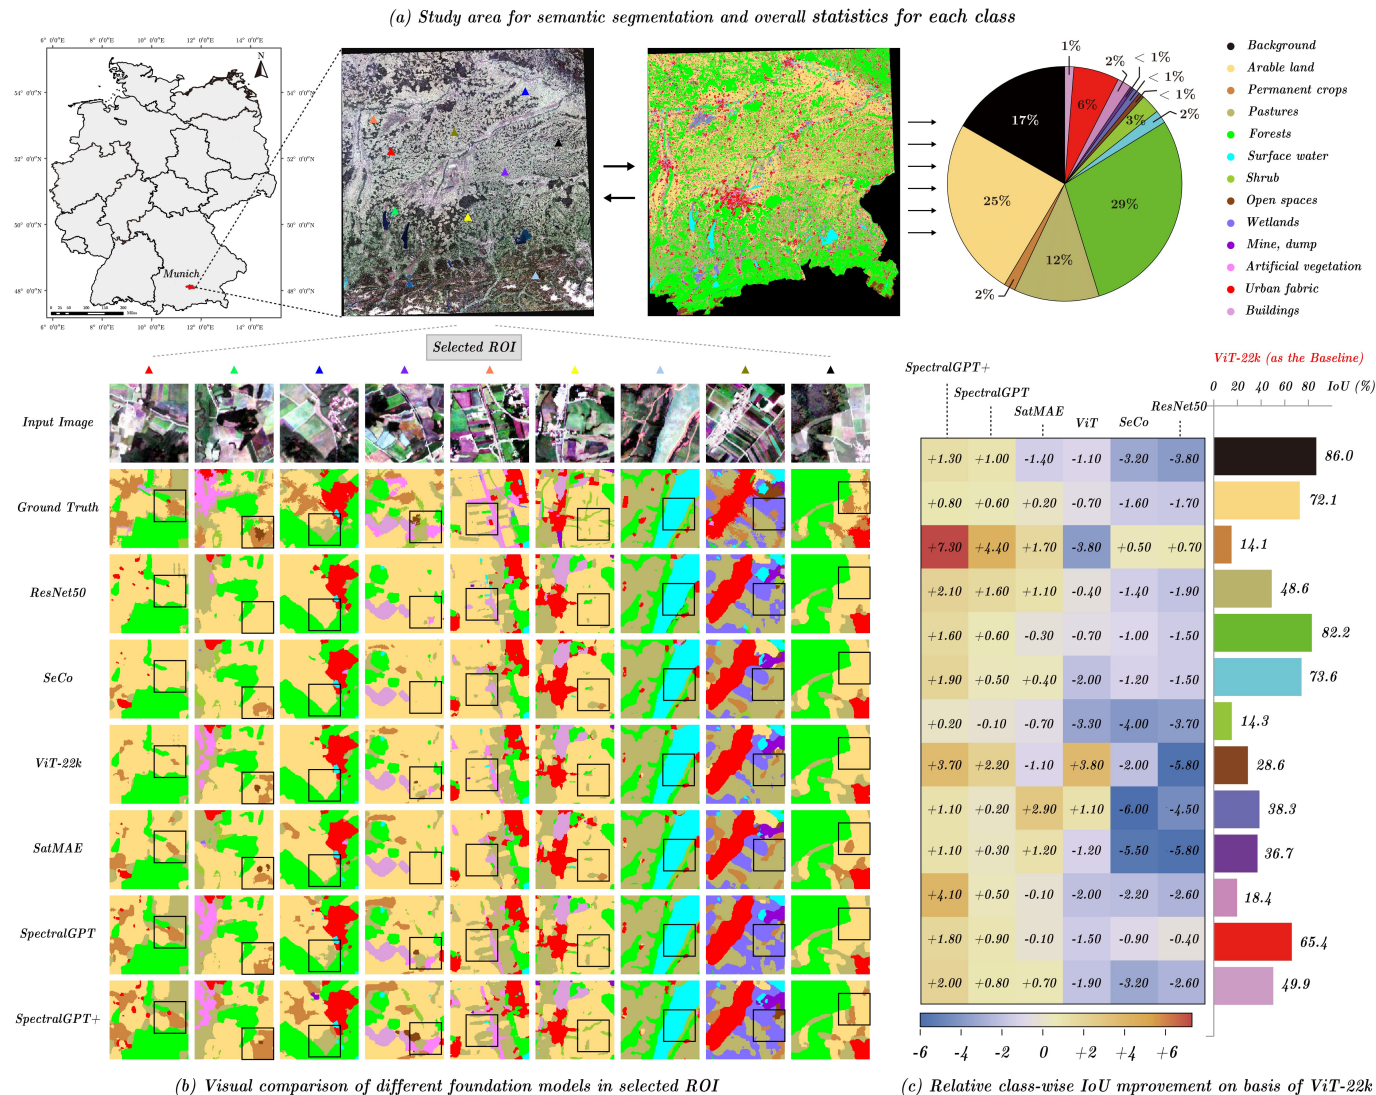


Fig. 6. Qualitative and quantitative semantic segmentation results of different pretrained foundation models on the SegMunich dataset. (a) Study area and overall statistics for each class in the semantic segmentation task. (b) Segmentation visualization maps of different foundation models in selected ROIs. (c) Relative IoU performance improvement for each class on the basis of ViT-22k.

the challenges posed by category-imbalanced classification. This results in a substantial enhancement in performance when compared to other foundation models.

3.4 RS Change Detection on OSCD

For the change detection task, we utilize the OSCD dataset [43]. Fig. 7(a) shows several examples at the city scale. This dataset comprises 24 cities of Sentinel-2 images, with 14 images used for training and 10 images for evaluation. These images were captured between 2015 and 2018 and encompassed 13 spectral bands with resolutions of 10m, 20m, and 60m. The dataset is annotated at the pixel level to indicate changes, specifically focusing on urban developments.

On the OSCD dataset, we perform image cropping to create patches of size 128×128 pixels with a 50% overlap rate, and we apply random flips and rotations as data augmentation techniques. For each pair of images, both are simultaneously processed through a shared encoder, and

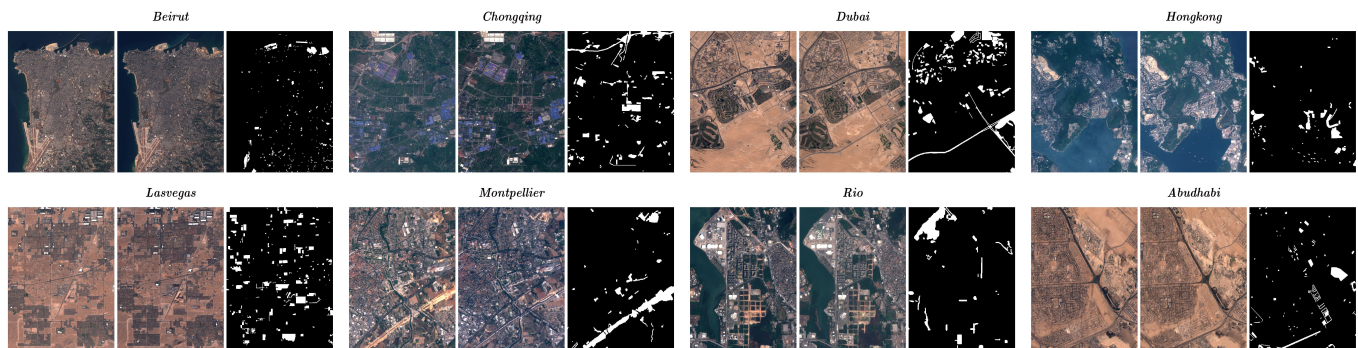
the difference between their features is computed and then passed to a UperNet. Each feature pixel consists of 4 tokens, similar to the segmentation approach, and we use a linear layer to consolidate these 4 tokens into 1 token. The model is trained for 60 epochs with a batch size of 64 and a learning rate set to 1×10^{-3} , using negative log-likelihood loss as the training objective. The entire framework for leveraging the pretrained SpectralGPT model in the change detection task is depicted in Fig. 5.

Model performance is assessed through precision, recall, and F1 score, with quantitative outcomes shown in Table 4 on the OSCD dataset, where our proposed model achieves the highest F1 score, surpassing the second-best model (i.e., SatMAE) by a substantial margin of 0.75% (1.53%). However, it is worth noting that our model excels in F1 score and recall but exhibits relatively lower precision compared to other models. This phenomenon can be attributed to two main factors. Firstly, extreme imbalance of the inherent data within the change detection task (see Fig.7(b)), where

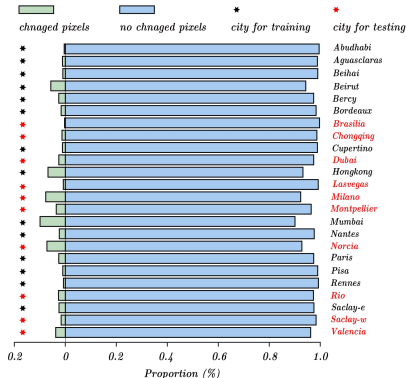
TABLE 3

Quantitative results of SOTA pretrained foundation models that are finely tuned using UperNet for the downstream RS semantic segmentation task in terms of overall accuracy (OA) and mean intersection over union (mIoU) on the SegMunich dataset. The best result is shown in bold.


Method	Pretrained Dataset	OA	Background	Arable land	Perm. Crops	Pastures	Forests	Surface water	Shrub	Open spaces	Wetlands	Mine, dump	Artificial veg.	Urban fabric	Buildings	mIoU
ResNet50 [37]	ImageNet-1k	80.1	82.2	70.4	14.8	46.7	80.7	72.1	10.6	22.8	33.8	30.9	15.8	65.0	47.3	45.6
SeCo [38]	SeCo	80.3	82.8	70.5	14.6	47.2	81.2	72.4	10.3	26.6	32.3	31.2	16.2	64.5	46.7	45.9
ViT [22]	Random Init	81.0	84.9	71.4	10.3	48.2	81.5	71.6	11.0	32.4	39.4	35.5	16.4	63.9	48.0	47.3
ViT-22k [22]	ImageNet-22k	81.7	86.0	72.1	14.1	48.6	82.2	73.6	14.3	28.6	38.3	36.7	18.4	65.4	49.9	48.3
SatMAE [31]	fMoW-S2	81.5	84.6	72.3	15.8	49.7	81.9	74.0	13.6	27.5	41.2	37.9	18.3	65.3	50.6	48.7
SpectralGPT	fMoW-S2	82.5	87.6	73.1	16.3	50.6	83.6	74.7	14.2	32.5	39.2	37.7	19.4	67.0	51.7	49.8
SpectralGPT+	fMoW-S2+BigEarthNet	82.7	88.0	73.1	22.5	51.0	84.1	76.0	14.5	33.7	39.6	38.7	22.5	67.4	52.0	51.0



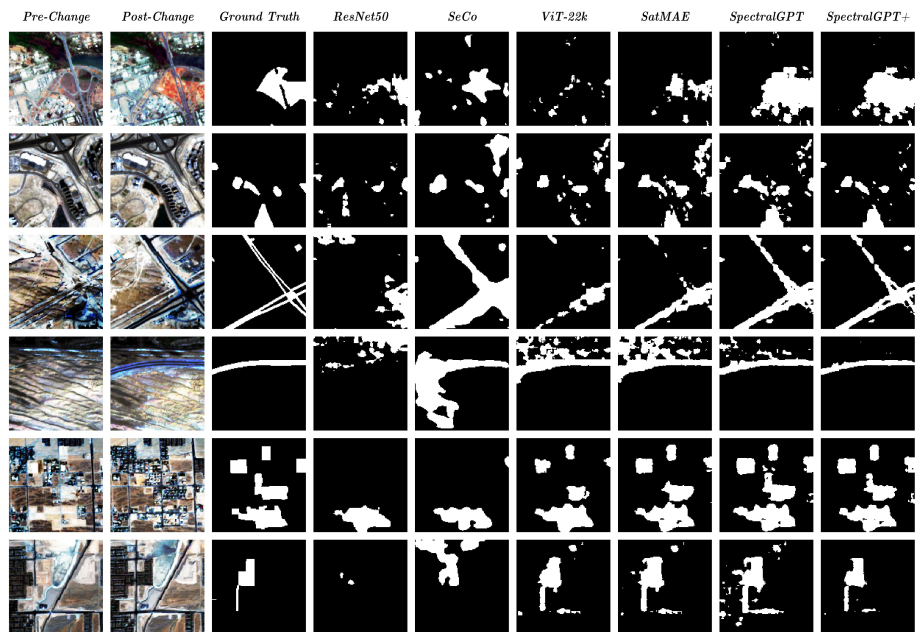
(a) Some examples of change detection at city scale



(b) Statistics of changed and unchanged pixels in 24 cities



(c) F1 scores of 7 foundation models over 10 testing cities



(d) Visual comparison of different foundation models in selected ROI

Fig. 7. Qualitative and quantitative change detection results of different pretrained foundation models on the OSCD dataset. (a) Some studied examples of change detection at the city scale. (b) Statistics of changed and unchanged pixels in 24 studied cities, including training and testing scenes. (c) F1 scores of different foundation models over 10 testing cities. (d) Change detection visualization maps of different foundation models in selected ROIs.

the number of positive and negative samples varies significantly, may lead the model to classify negative cases as positive to improve recall at the cost of precision. Secondly, the complexity of the ViT architecture demands a substantial amount of data to mitigate overfitting. The model may struggle with overfitting and become less adaptable

to out-of-domain data. Addressing this challenge could involve providing additional fine-tuning data or reducing the model's rank. In terms of qualitative results, our model excels in predicting change pixels with fewer false negatives in selected ROIs of Fig. 7(d). Significantly, Fig. 7(c) accentuates the exceptional performance of SpectralGPT, with our

TABLE 4

Quantitative results of SOTA pretrained foundation models that are finely tuned using UperNet for the downstream RS change detection task in terms of precision, recall, and F1 score on the OSCD dataset. The best result is shown in bold.

Method	Pretrained Dataset	Precision	Recall	F1
ResNet50 [37]	ImageNet-1k	65.42	38.86	48.10
SeCo [38]	SeCo	57.71	49.23	49.82
ViT [22]	Random Init.	56.71	47.52	51.71
ViT-22k [22]	ImageNet-22k	52.09	52.37	52.23
SatMAE [31]	fMoW-S2	55.18	50.54	52.76
SpectralGPT	fMoW-S2	51.65	56.15	53.51
SpectralGPT+	fMoW-S2+BigEarthNet	52.39	57.20	54.29

model achieving the top results in half of the testing cities. In addition, there is a consistent performance trend among the compared foundation models across 10 different testing cities, with *Lasvegas* and *Montpellier* consistently achieving the highest and second-highest F1 scores, respectively.

3.5 Ablation Studies

During the pretraining stage, we conduct a comprehensive study of various factors that may impact downstream task performance. These factors encompass masking ratio, ViT patch size, data scale, reconstruction target, decoder depth, and model size. To provide a more rigorous assessment of pretrained models, we subject all ablation models to fine-tuning on the BigEarthNet multi-label classification dataset with only a 10% subset of the train set, which poses a more formidable challenge, evaluated using the mAP measurement. Our choice of ViT-B as the backbone model ensures consistency across experiments. Except for ablations involving data scale and training schedule length, all models undergo pretraining on the fMoW-S2 dataset for a duration of 200 epochs. This comprehensive evaluation framework enables us to gain deeper insights into the impact of these factors on model performance.

3.5.1 Token Size

Table 5(a) Fig. 8(a) offers crucial insights into the impact of token size on model performance, consistently demonstrating that larger patch sizes lead to reduced model performance, aligning with prior research findings [31]. This phenomenon can be attributed to the intrinsic characteristics of ViT architectures. With larger token sizes, such as 16×16 , each image contains fewer tokens, resulting in a diminished level of fine-grained spatial information as the model progresses through its deeper layers. Consequently, this reduction in spatial detail negatively affects the model's overall performance. However, it is noteworthy that the pretrained model consistently enhances mAP regardless of the token size settings, emphasizing its capacity to improve performance across various configurations. Significantly, the recognition performance with a token size of 8×8 is notably superior to that with 16×16 , despite the input images being 96×96 or 128×128 in size, underscoring the versatility and efficacy of the pretrained model.

3.5.2 Data Scale

Table 5(b) and Fig. 8(b) present a comprehensive analysis focusing on the impact of pretraining data in our research.

We conducted pretraining using two datasets (i.e., fMoW-S2, BigEarthNet) while maintaining a standardized input image size of 96×96 . To delve deeper into this comparison, we initially pretrained models exclusively on fMoW-S2, followed by a seamless continuation of pretraining on BigEarthNet without any intermediate fine-tuning steps. Our pretraining datasets consisted of the extensive train set of fMoW-S2, which encompasses an impressive 712,874 images from around the world, and the BigEarthNet training set, which comprises 351,496 images within the European region after excluding those affected by snow, clouds, or cloud shadows.

This analysis in Table 5(b) underscores the substantial influence of both data scale and distribution on model pretraining. Models pretrained on the same dataset as the downstream task consistently demonstrate superior performance, highlighting the crucial role of dataset coherence in effective transfer learning. Furthermore, fMoW-S2 outperforms BigEarthNet in pretraining, primarily due to its larger dataset and broader geographic coverage. Interestingly, the concept of continual pretraining, which combines both datasets, results in models with higher mAP scores. This improvement can be attributed in part to the transition from 96×96 images during fMoW-S2 pretraining to 128×128 images during BigEarthNet pretraining, underscoring the beneficial impact of increasing image size on overall model efficacy.

3.5.3 Masking Ratio

Table 5(c) and Fig. 8(c) shed light on the impact of the masking ratio, revealing a noteworthy trend where higher masking ratios correspond to improved model performance. Unlike the conventional masking ratio of 75% often applied to natural RGB images, we find that the optimal masking ratio for multi-spectral images is 90%. This observation aligns with the hypothesis presented in [30] that the masking ratio in MIM methods is intricately linked to the information redundancy within the data. Multi-spectral images inherently exhibit greater information redundancy, with strong correlations among their spectral bands. Consequently, a higher masking ratio is essential for the model to effectively learn meaningful representations from these images. Moreover, a 90% masking ratio significantly enhances the efficiency of the pretraining stage, reducing memory complexity and expediting training times, offering a practical advantage in model development.

3.5.4 Reconstruction Target

Table 5(d) and Fig. 8(d) conduct an insightful analysis of the influence of reconstruction targets on normalized, standardized data and raw data without normalization or standardization in the context of multi-spectral images. Normalization, which scales all data to the $[0, 1]$ range, and standardization, which transforms data to have a mean of 0 and a standard deviation of 1, are the two examined targets. Remarkably, the results show minimal disparity in model performance between normalization and standardization reconstruction targets, primarily because both targets pertain to pixel-level data transformations. However, the model pretrained on raw data performs much worse than the models with normalized reconstruction targets. We attribute this phenomenon to the characteristics of multi-spectral images.

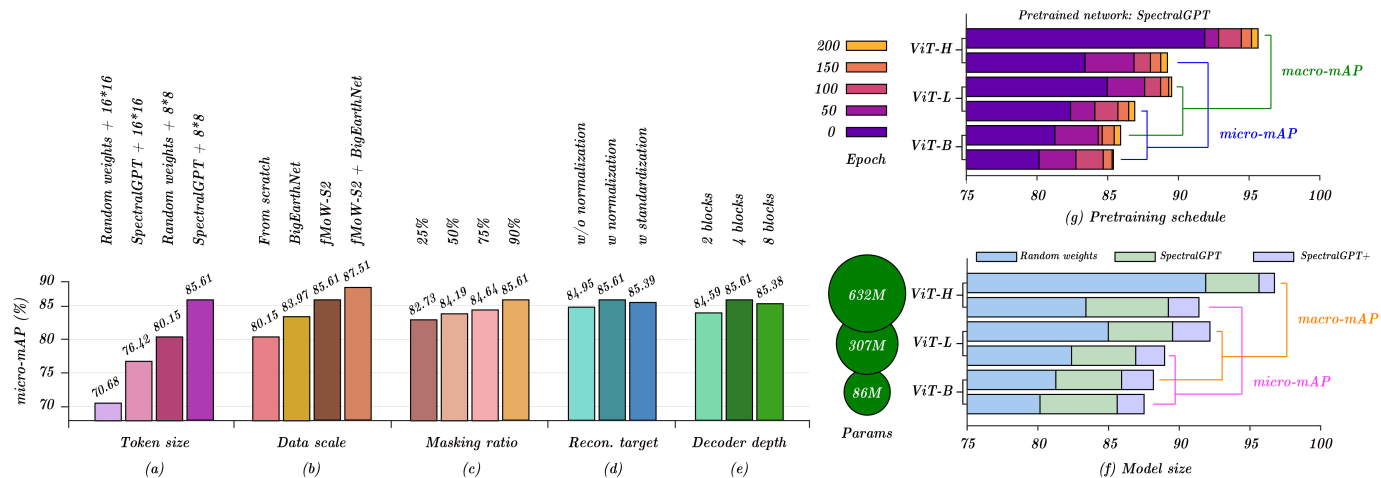


Fig. 8. An illustration for ablation analysis of the proposed SpectralGPT foundation model on BigEarthNet-S2 dataset in terms of (a) token size, (b) data scale, (c) masking ratio, (d) reconstruction target, and (e) decoder depth, as well as (f) model size (i.e., ViT-Base, ViT-Large, and ViT-Huge) and their (g) pretraining length/epoch.

TABLE 5

Ablation Analysis of the proposed SpectralGPT foundation model in terms of masking ratio, patch size, data scale, reconstruction target, and decoder depth, respectively. The best result is shown in bold.

(a) Token Size			(b) Data Scale		(c) Masking Ratio		(d) Reconstruction Target		(e) Decoder Depth	
Init. Weights	Patch Size	mAP	Pretrained Dataset	mAP	Ratio	mAP	Case	mAP	Blocks	mAP
Random	16	70.68	From scratch	80.15	25%	82.73	Without norm	84.95	2	84.59
SpectralGPT	16	76.42	BigEarthNet	83.97	50%	84.19	Normalization	85.61	4	85.61
Random	8	80.15	fMoW-S2	85.61	75%	84.64	Standardization	85.39	8	85.38
SpectralGPT	8	85.61	fMoW-S2+BigEarthNet	87.51	90%	85.61				

The spectral values are usually numerically large and vary from band to band, thereby the model pretraining on the raw data may need a longer pretraining schedule to converge and show the same performance compared with those models pretrained on normalized and standardized data. Our perspective suggests that employing a more semantically meaningful target in a specific representation space could potentially yield improved model performance.

3.5.5 Decoder Depth

Table 5(e) and Fig. 8(e) examine the impact of decoder depth on model performance, following the principles of MIM methods where the pretrained encoder serves as the backbone for downstream tasks while discarding the decoder component. Notably, the results reveal that a shallow decoder configuration is ill-suited for spectral model pretraining. This observation aligns with the hypothesis that spectral images, characterized by high dimensionality and complexity, require a decoder with enhanced capacity, consistent with prior findings in the field [30].

3.5.6 Model Size

Table 6 and Fig. 8(f) give a comparative analysis between fine-tuning results of ViT-B and ViT-L quantitatively and qualitatively, revealing compelling insights. The macro average precision and micro average precision are listed to comprehensively evaluate the performance of models. ViT-B, equipped with 12 transformer layers and 86 million

parameters, exhibits promising performance gains when employing the proposed method, achieving a mAP(micro) of 85.41, surpassing the ViT-B trained from scratch by 5.26. On the other hand, ViT-L, featuring 24 layers and 307 million parameters, notably outperforms ViT-B, with a mAP(micro) of 86.92, surpassing the model trained from scratch by a significant margin of 4.44. Besides, ViT-H, concluding 32 layers and 632 million parameters, highly enhances the performance of the neural network on BigEarthNet with the mAP(micro) of 89.23. Notably, though our models are only fine-tuned with 10% downstream training data, the ViT-H model employing SpectralGPT⁺ pretrained weights beats all the models even trained with the whole train set, with a SOTA mAP(micro) of 91.39. These results underscore the pivotal role of an appropriate pretraining strategy and indicate that larger ViT models are capable of learning more intricate image representations, rendering them highly suitable for tasks demanding superior accuracy.

3.5.7 Pretraining Schedule

In Fig. 8(g), we present the fine-tuning results for models trained with varying pre-training epochs, evaluated using the macro-mAP and micro-mAP metrics, respectively. Notably, the models pretrained for just 50 epochs exhibit significant performance gains compared to those trained from scratch. The observed trend in the figure indicates that the models continue to benefit from longer pretraining epochs, suggesting that extended training can further enhance per-

TABLE 6

Performance comparison using different pretrained models across three ViT-based network scales (i.e., base, large, huge) on the BigEarthNet dataset. The best result is shown in bold.

Network Scale	Params	Pretrained Network	macro-mAP	micro-mAP
ViT-Base	86M	Random Init.	81.27	80.15
		SpectralGPT	85.92	85.61
		SpectralGPT ⁺	88.17	87.50
ViT-Large	307M	Random Init.	84.98	82.38
		SpectralGPT	89.53	86.92
		SpectralGPT ⁺	92.17	88.96
ViT-Huge	632M	Random Init.	91.87	83.40
		SpectralGPT	95.64	89.23
		SpectralGPT ⁺	96.73	91.39

formance. Moreover, the results in Table 6 reinforce this finding, as ViT-L and ViT-H consistently achieve higher mAP compared to ViT-B, highlighting the effectiveness of both extended pretraining and larger model architectures.

3.6 Visual Comparison and Geo-characteristic Recoverability

With varying masking ratios (i.e., 50%, 75%, 90%, and 95%) as the input, Fig. 9 visually illustrates the image reconstruction results obtained using SatMAE and our SpectralGPT. Not unexpectedly, as the masking ratio increases, the reconstructed images deviate more from the originals. It is worth highlighting, however, that the proposed SpectralGPT outperforms SatMAE significantly in terms of spectral image reconstruction performance, particularly in preserving visual structures and textural details. To be specific, when utilizing 50% visible patches, the reconstructed results with SatMAE are comparable to those with SpectralGPT, albeit with some slight blurriness in certain fine details in the SatMAE results. As the percentage of masked patches increases (e.g., from a 75% mask to 90% and further to 95%), the reconstruction performance of SatMAE experiences a substantial decline. In contrast, our SpectralGPT exhibits a superior reconstruction capability (cf. SatMAE). Even with a masking rate exceeding 90%, critical structures, and shape components remain preserved in the visuals, demonstrating our model's robust learning, reasoning, and generalization capabilities.

In addition to the in-depth discussion and sensitivity analysis concerning the masking ratio, we have undertaken more extensive investigations into spectral-wise reconstruction capabilities by using only 10% of visible patches, with the remainder masked out. These investigations prioritize the representation of geographical characteristics, utilizing various spectral band combinations. As illustrated in Fig. 10, we present visualizations of eight different band combinations. These visualizations distinctly highlight the remarkable superiority of our proposed SpectralGPT (closer to that generated by original images) when compared to SatMAE, particularly in terms of band-wise spectral reconstruction capabilities and its application value in the context of EO tasks. In our study, eight geo-characteristics have been identified using various band combinations, corresponding to observation targets in practical applications⁴, as detailed

4. <https://www.sentinel-hub.com/>

TABLE 7

List of geographical characteristics and visualized band combinations corresponding to observation targets in practical geoscience applications.

Geo-characteristic	Band Combination	Observation Target
Agriculture Condition	B11, B8, B2	Crop Health
Bathymetric Survey	B4, B3, B1	Coast
Vegetation Health	B8, B4, B3	Chlorophyll
Geological Structure	B12, B11, B2	Faults, Lithology
Moisture Content	B11, B8A	Plant Water Pressure
Vegetation Density	B12, B8A, B4	Vegetation Cover, Soil, Building
Vegetation Index (NDVI)	B8, B4	Tree Crown, Urban, Waterscape
Atmospheric Penetration	B12, B11, B8A	Particle, Smoke, Haze, Thin Cloud

in Table 7. Furthermore, notable visual differences are evident in the geo-characteristics obtained using SatMAE and SpectralGPT. These pronounced visual disparities can be attributed to spectral degradation stemming from the comparatively limited reconstructive and inferential capabilities of SatMAE when compared to our more powerful SpectralGPT.

4 CONCLUSION

The explosive development of foundation models represents a significant technological revolution following the advent of deep learning. Currently, various industries are witnessing significant leaps in technology and application advancements, largely driven by the emergence of foundation models. The RS field is no exception, with numerous EO applications, reaping significant benefits.

Spectral imaging has gained recognition in EO for its ability to provide rich insights into the composition of observed objects and materials, making it a transformative technology with vast potential to address global challenges and reshape various industries. However, the ever-expanding availability of spectral data from various RS platforms undeniably presents formidable challenges. There is a pressing demand for the development of foundation models specifically designed for spectral RS data. To fully unlock and leverage the potential of spectral RS data, it is imperative to overcome and resolve several challenging obstacles. These include efficiently processing and utilizing diverse RS spectral big data from various sources, extracting meaningful knowledge representations from complex spatial-spectral mixed information, and tackling the spectral degradation of neighboring spectral relevance modeling.

In response to these challenges, we propose SpectralGPT, a customized spectral RS foundation model, featuring a novel 3D GPT architecture. With its innovative 3D GPT architecture, trained on over one million spectral images and over 600 million parameters, SpectralGPT empowers intelligent processing of spectral RS big data. SpectralGPT can flexibly handle diverse inputs in terms of size, resolution, temporal variability, and geographical coverage. This 3D masking strategy enables effective information extraction from spatial-spectral coupling tokens. Moreover, the innovative multi-target reconstruction is capable of capturing sequentially preserving spectral characteristics and meanwhile mitigating spectral degradation. Notably, our progressive training paradigm empowers the foundation

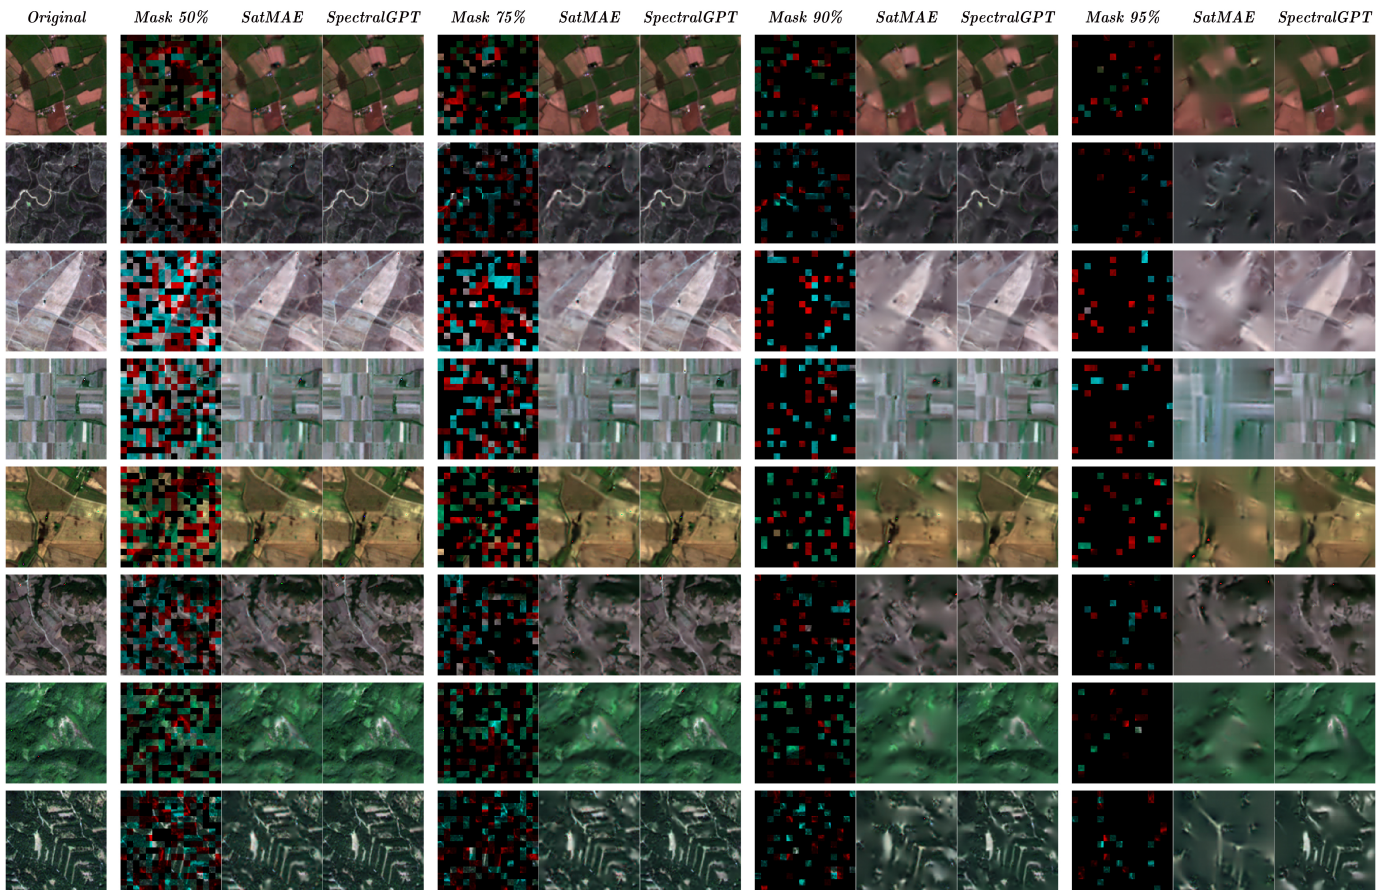


Fig. 9. Visual comparison from the nature-color (R: 4, G: 3, B: 2) image reconstruction perspective between SatMAE and SpectralGPT with varied masking ratios of 50%, 75%, 90%, and 95%, respectively. By masking out a greater number of patches, the reconstructed images exhibit noticeable differences from the originals (e.g., 50% vs. 95%), which is expected. It is worth noting, however, that SpectralGPT holds stronger reconstruction capability (*cf. SatMAE*), even if the masking rate has reached over 90%, showing its powerful learning, reasoning, and generalizing performance.

model, surpassing transitional points in performance. These breakthroughs achieved by SpectralGPT democratize access to spectral RS big data, rendering it more accessible and cost-effective for large-scale EO applications.

Our study also includes a comprehensive assessment of MAE-based pretrained foundation models, with a focus on spectral reconstruction capabilities. We systematically evaluated model performance with inputs ranging from 50% to as low as 5% visible tokens. This extensive analysis allows us to gauge their proficiency in spectral-wise reconstruction and inference, especially significant in Geo-fields, such as agriculture, oceanography, geology, and vegetation. Visualizing the band combinations of reconstructed spectral images using both SatMAE and SpectralGPT demonstrates the latter's potential in practical EO tasks and Geo-field applications.

Looking ahead, our research will pursue several objectives. We plan to expand the volume and diversity of RS data used for training, encompassing various modalities, resolutions, time series, and image sizes. This enrichment will enhance the robustness of the RS foundation model. Furthermore, we aim to extend SpectralGPT's capabilities by incorporating a wider range of downstream tasks. This will transform SpectralGPT into a versatile AI model with improved generalization, well-suited for diverse EO and geoscience applications.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China under Grant 42030111, 42271350, 42241109, and 62201553, and the MIAI@Grenoble Alpes (ANR-19-P3IA-0003).

REFERENCES

- [1] A. F. Goetz, G. Vane, J. E. Solomon, and B. N. Rock, "Imaging spectrometry for earth remote sensing," *Science*, vol. 228, no. 4704, pp. 1147–1153, 1985. 1
- [2] S. Mei, C. Song, M. Ma, and F. Xu, "Hyperspectral image classification using group-aware hierarchical transformer," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2022. 1
- [3] M. Zhou, J. Huang, D. Hong, F. Zhao, C. Li, and J. Chanussot, "Rethinking pan-sharpening in closed-loop regularization," *IEEE Transactions on Neural Networks and Learning Systems*, 2023. 1
- [4] M. Reichstein, G. Camps-Valls, B. Stevens, M. Jung, J. Denzler, N. Carvalhais, and f. Prabhat, "Deep learning and process understanding for data-driven earth system science," *Nature*, vol. 566, no. 7743, pp. 195–204, 2019. 1
- [5] W. He, Q. Yao, C. Li, N. Yokoya, Q. Zhao, H. Zhang, and L. Zhang, "Non-local meets global: An iterative paradigm for hyperspectral image restoration," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 4, pp. 2089–2107, 2020. 1
- [6] D. Hong, J. Yao, C. Li, D. Meng, N. Yokoya, and J. Chanussot, "Decoupled-and-coupled networks: Self-supervised hyperspectral image super-resolution with subpixel fusion," *IEEE Transactions on Geoscience and Remote Sensing*, 2023. 1

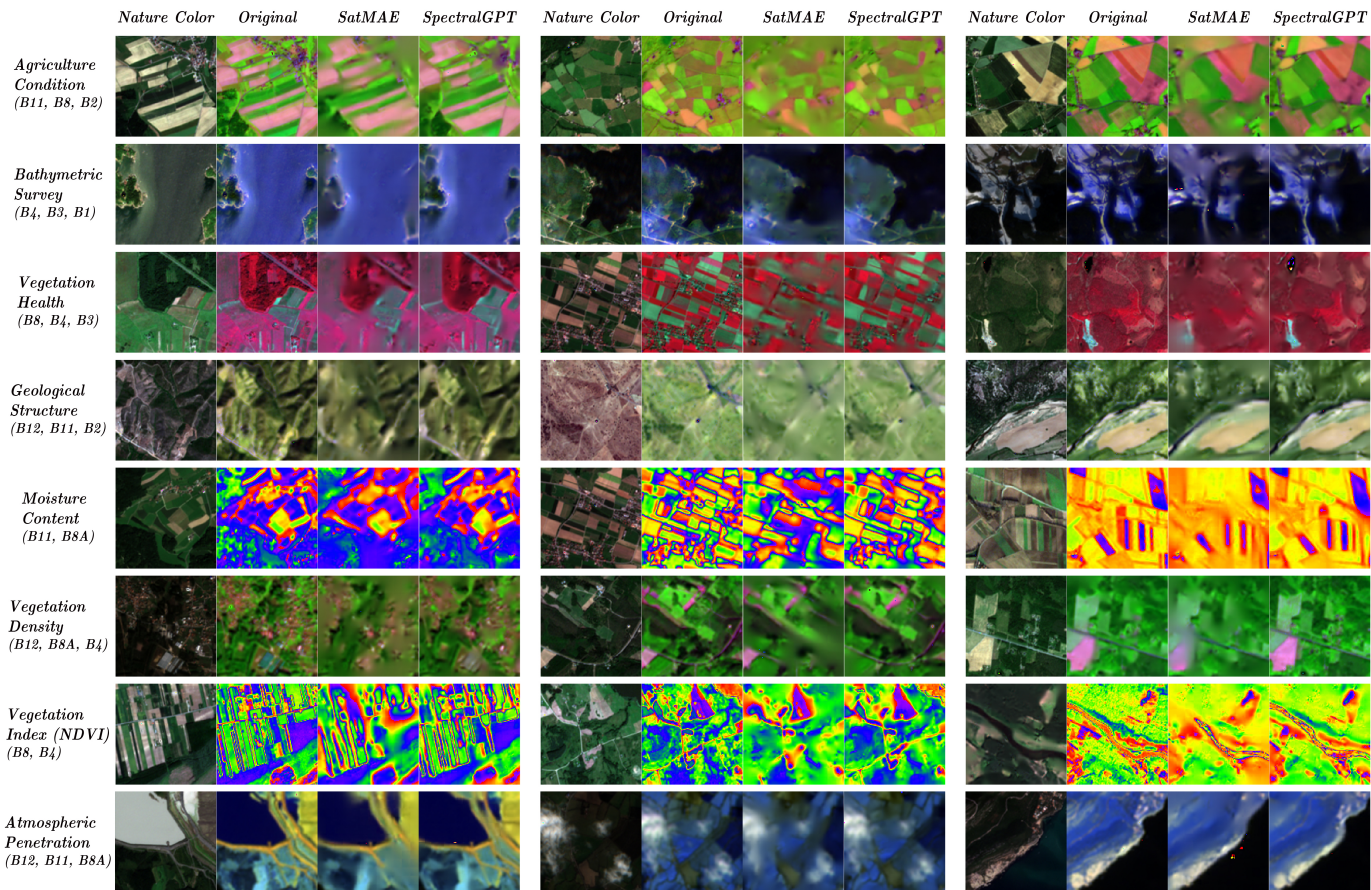


Fig. 10. Band combination visualization for geographical characteristic representations of reconstructed spectral images. These images are reconstructed and inferred by SatMAE and SpectralGPT models, respectively, only using 10% visible patches (masking 90% patches). Eight band combinations are given to highlight the characteristics, encompassing applications in agriculture, oceanography, geology, vegetation, atmosphere, and more, further serving a variety of EO applications.

- [7] S. Mei, R. Jiang, M. Ma, and C. Song, "Rotation-invariant feature learning via convolutional neural network with cyclic polar coordinates convolutional layer," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–13, 2023. **1**
- [8] D. Hong, B. Zhang, H. Li, Y. Li, J. Yao, C. Li, M. Werner, J. Chanussot, A. Zipf, and X. X. Zhu, "Cross-city matters: A multimodal remote sensing benchmark dataset for cross-city semantic segmentation using high-resolution domain adaptation networks," *Remote Sensing of Environment*, vol. 299, p. 113856, 2023. **1**
- [9] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, *et al.*, "On the opportunities and risks of foundation models," *arXiv preprint arXiv:2108.07258*, 2021. **2**
- [10] J. Tian, X. Sun, Y. Du, S. Zhao, Q. Liu, K. Zhang, W. Yi, W. Huang, C. Wang, X. Wu, *et al.*, "Recent advances for quantum neural networks in generative learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. **2**
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *NeurIPS*, vol. 30, 2017. **2**
- [12] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, "Learning transferable visual models from natural language supervision," in *ICML*, pp. 8748–8763, PMLR, 2021. **2**
- [13] X. Liu, D. Hong, J. Chanussot, B. Zhao, and P. Ghamisi, "Modality translation in remote sensing time series," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2021. **2**
- [14] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *CVPR*, pp. 9729–9738, 2020. **2**
- [15] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *ICML*, pp. 1597–1607, 2020. **2**
- [16] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. E. Hinton, "Big self-supervised models are strong semi-supervised learners," *NeurIPS*, vol. 33, pp. 22243–22255, 2020. **2**
- [17] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *NeurIPS*, vol. 33, pp. 9912–9924, 2020. **2**
- [18] X. Chen and K. He, "Exploring simple siamese representation learning," in *CVPR*, pp. 15750–15758, 2021. **2**
- [19] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," *arXiv preprint arXiv:2003.04297*, 2020. **2**
- [20] Y. Xiong, M. Ren, and R. Urtasun, "Loco: Local contrastive representation learning," *NeurIPS*, vol. 33, pp. 11142–11153, 2020. **2**
- [21] Z. Xie, Y. Lin, Z. Zhang, Y. Cao, S. Lin, and H. Hu, "Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning," in *CVPR*, pp. 16684–16693, 2021. **2**
- [22] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020. **2, 7, 8, 10, 11**
- [23] H. Bao, L. Dong, S. Piao, and F. Wei, "Beit: Bert pre-training of image transformers," in *ICLR*, 2022. **2**
- [24] D. Du, Y. Gu, T. Liu, and X. Li, "Spectral reconstruction from satellite multispectral imagery using convolution and transformer joint network," *IEEE Transactions on Geoscience and Remote Sensing*, 2023. **2**
- [25] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *CVPR*, pp. 16000–16009, 2022. **2, 4, 7**
- [26] D. Wang, Q. Zhang, Y. Xu, J. Zhang, B. Du, D. Tao, and L. Zhang, "Advancing plain vision transformer towards remote sensing

foundation model," *IEEE Transactions on Geoscience and Remote Sensing*, 2022. **2**

[27] X. Sun, P. Wang, W. Lu, Z. Zhu, X. Lu, Q. He, J. Li, X. Rong, Z. Yang, H. Chang, *et al.*, "Ringmo: A remote sensing foundation model with masked image modeling," *IEEE Transactions on Geoscience and Remote Sensing*, 2022. **2**

[28] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *ICCV*, pp. 10012–10022, 2021. **2**

[29] Z. Tong, Y. Song, J. Wang, and L. Wang, "Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training," *NeurIPS*, vol. 35, pp. 10078–10093, 2022. **3**

[30] C. Feichtenhofer, Y. Li, K. He, *et al.*, "Masked autoencoders as spatiotemporal learners," *NeurIPS*, vol. 35, pp. 35946–35958, 2022. **3, 4, 6, 11, 12**

[31] Y. Cong, S. Khanna, C. Meng, P. Liu, E. Rozi, Y. He, M. Burke, D. Lobell, and S. Ermon, "Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery," *NeurIPS*, vol. 35, pp. 197–211, 2022. **3, 6, 7, 8, 10, 11**

[32] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, P.-A. Manzagol, and L. Bottou, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *Journal of Machine Learning Research*, vol. 11, no. 12, 2010. **3**

[33] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018. **4**

[34] G. Christie, N. Fendley, J. Wilson, and R. Mukherjee, "Functional map of the world," in *CVPR*, pp. 6172–6180, 2018. **6**

[35] G. Sumbul, M. Charfuelan, B. Demir, and V. Markl, "Bigearthnet: A large-scale benchmark archive for remote sensing image understanding," in *IGARSS*, pp. 5901–5904, IEEE, 2019. **6, 7**

[36] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019. **6**

[37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, pp. 770–778, 2016. **7, 8, 10, 11**

[38] O. Manas, A. Lacoste, X. Giró-i Nieto, D. Vazquez, and P. Rodriguez, "Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data," in *ICCV*, pp. 9414–9423, 2021. **7, 8, 10, 11**

[39] P. Helber, B. Bischke, A. Dengel, and D. 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*, vol. 12, no. 7, pp. 2217–2226, 2019. **7**

[40] M. Neumann, A. S. Pinto, X. Zhai, and N. Houlsby, "In-domain representation learning for remote sensing," in *ICLR-AI for Earth Sciences Workshop*, 2020. **7**

[41] T. Xiao, Y. Liu, B. Zhou, Y. Jiang, and J. Sun, "Unified perceptual parsing for scene understanding," in *ECCV*, pp. 418–434, 2018. **8**

[42] D. Hong, B. Zhang, X. Li, Y. Li, C. Li, J. Yao, N. Yokoya, H. Li, X. Jia, A. Plaza, P. Gamba, J. A. Benediktsson, and J. Chanussot, "SpectralGPT: The first remote sensing foundation model customized for spectral data," Oct. 2023. **8**

[43] R. C. Daudt, B. Le Saux, A. Boulch, and Y. Gousseau, "Urban change detection for multispectral earth observation using convolutional neural networks," in *IGARSS*, pp. 2115–2118, Ieee, 2018. **9**



Danfeng Hong (IEEE Senior Member) received the Dr. -Ing degree (summa cum laude) from the Signal Processing in Earth Observation (SiPEO), Technical University of Munich (TUM), Munich, Germany, in 2019.

Since 2022, he has been a Full Professor with the Aerospace Information Research Institute, Chinese Academy of Sciences. Before joining CAS, he was a Research Scientist and led a Spectral Vision Working Group at the Remote Sensing Technology Institute (IMF), German Aerospace Center (DLR), Oberpfaffenhofen, Germany. He was also an Adjunct Scientist at GIPSA-lab, Grenoble INP, CNRS, Univ. Grenoble Alpes, Grenoble, France. His research interests include artificial intelligence, multimodal remote sensing, large foundation models, hyperspectral imaging, and large-scale Earth observation.

Dr. Hong is an Associate Editor for the IEEE Transactions on Geoscience and Remote Sensing (TGRS) and the Editorial Board Member of Information Fusion and ISPRS Journal of Photogrammetry and Remote Sensing. He was a recipient of the Best Reviewer Award of the IEEE TGRS in 2021 and 2022, the Best Reviewer Award of the IEEE JSTARS in 2022, the Jose Bioucas Dias Award for recognizing the outstanding paper at WHISPERS in 2021, the Remote Sensing Young Investigator Award in 2022, the IEEE GRSS Early Career Award in 2022, and a Highly Cited Researcher (Clarivate Analytics) in 2022 and 2023.



Bing Zhang (IEEE Fellow) received the B.S. degree in geography from Peking University, Beijing, China, in 1991, and the M.S. and Ph.D. degrees in remote sensing from the Institute of Remote Sensing Applications, Chinese Academy of Sciences, Beijing, in 1994 and 2003, respectively.

He is currently a Full Professor and the Deputy Director of the Aerospace Information Research Institute, China Academy of Sciences. He is an IEEE Fellow and a Fellow of the Geographical Society of China. He has long been engaged in research on Hyperspectral remote sensing technology and applications. He has published more than 400 journal papers. He is also a Highly Cited Researcher (Clarivate Analytics). His creative achievements were rewarded more than 10 important prizes from international institutions and the Chinese government, including the IEEE Geoscience and Remote Sensing Society (GRSS) Regional Leader Award, the National Science and Technology Advance Award of China, the Outstanding Scientific Achievement Award of Chinese Academy of Sciences, etc.



Xuyang Li received the B.S. degree from the School of Optoelectronic Science and Engineering, University of Electronic Science and Technology of China (UESTC), Chengdu, China, in 2023. He is currently pursuing his Ph.D. degree at the Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing, China. His research interests include machine learning, remote sensing foundation models, and multimodal interpretation.



Yuxuan Li received the B.E. degree from the Institute of Geophysics and Geomatics, China University of Geosciences, Wuhan, China, in 2023. He is currently pursuing his Ph.D. degree at the Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing, China. His research interests include deep learning and remote sensing image processing.




Chenyu Li received the B.S. and M.S. degrees from the School of Transportation, Southeast University, Nanjing, China, in 2018 and 2021, respectively.

She is currently pursuing her Ph.D. degree from the Department of Mathematics at Southeast University, Nanjing, China, and is also a joint Ph.D. student at the Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing, China. Her research interests include interpretable artificial intelligence, hyperspectral remote sensing, and target detection.



Jing Yao (IEEE Member) received the Ph.D. degree in mathematics from Xi'an Jiaotong University, Xi'an, China, in 2021.

Since 2021, he has been an Assistant Professor at the Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing, China. From 2019 to 2020, he was a visiting student at the Signal Processing in Earth Observation (SiPEO), Technical University of Munich (TUM), Munich, Germany, and at the Remote Sensing Technology Institute (IMF), German Aerospace Center (DLR), Oberpfaffenhofen, Germany. His research interests include hyperspectral and multimodal remote sensing image analysis, mainly including optimization and deep learning-based methods for image processing and interpretation tasks. He was the recipient of the Jose Bioucas Dias Award for recognizing an outstanding paper at WHISPERS in 2021.



Naoto Yokoya (IEEE Member) received the M.Eng. and Ph.D. degrees from the Department of Aeronautics and Astronautics, The University of Tokyo, Tokyo, Japan, in 2010 and 2013, respectively.

From 2013 to 2017, he served as an Assistant Professor at The University of Tokyo. During the period from 2015 to 2017, he held the position of an Alexander von Humboldt Fellow at the German Aerospace Center (DLR) in Oberpfaffenhofen, Germany, and the Technical University of Munich (TUM) in Munich, Germany. Since 2018, he has been with the RIKEN Center for Advanced Intelligence Project, Tokyo, where he leads the Geoinformatics Team. Since 2020, he has been with The University of Tokyo, where he is currently an Associate Professor. His research is focused on the development of image processing, data fusion, and machine learning algorithms for understanding remote sensing images, with applications to disaster management and environmental assessment.

Dr. Yokoya won first place in the 2017 IEEE GRSS Data Fusion Contest organized by the IADF TC. He was the Chair from 2019 to 2021, a Co-Chair of the IEEE GRSS IADF TC from 2017 to 2019, and also the Secretary of the IEEE GRSS All Japan Joint Chapter from 2018 to 2021. He was an Associate Editor for the IEEE JSTARS from 2018 to 2021. He has been an Associate Editor for the IEEE TGRS since 2021, as well as the ISPRS Journal of Photogrammetry and Remote Sensing since 2024. He has been designated a Clarivate Highly Cited Researcher since 2022.



Hao Li (Member, IEEE) received the double B.Sc. degrees in Cartography and Computer Science from Wuhan University, China in 2015, the M.Sc. degree in Geomatic Engineering from University of Stuttgart, Germany in 2018, and the Dr. rer. nat degree from the GIScience research group, Heidelberg University, Germany in 2022. Since 2022, he has been a Postdoc researcher at the Professorship of Big Geospatial Data Management at the Technical University of Munich (TUM), where he is pursuing his Habilitation degree.

Dr. Li's research interests include Volunteered Geographic Information (VGI), geo-semantics, remote sensing, and geospatial big data management. He was a recipient of the ISPRS Best Poster Award (TC IV) at the XXIV ISPRS Congress in Nice, in 2022.



Pedram Ghamisi (IEEE Senior Member) graduated with a Ph.D. in electrical and computer engineering from the University of Iceland in 2015. He currently serves as (1) the head of the machine learning group at Helmholtz-Zentrum Dresden-Rossendorf (HZDR), Germany, and (2) a visiting professor at Lancaster University, UK. His research interests encompass deep learning, with a specific focus on remote sensing applications. For detailed information, please visit <http://www.ai4rs.com>.



Xiuping Jia (IEEE Fellow) received a B.Eng. degree from the Beijing University of Posts and Telecommunications, Beijing, China, in January 1982. Later she received a PhD degree in Electrical Engineering, and a Graduate Certificate in Higher Education, from The University of New South Wales, Australia, in 1996 and 2005, respectively, via part-time study. She has a life-long academic career in higher education which she is always passionate about. She is currently an Associate Professor at the School of Engineering Technology, The University of New South Wales at Canberra, Australia. Her research interests include remote sensing, hyperspectral image processing, and spatial data analysis. Dr Jia has authored or co-authored more than 280 referred papers, including over 170 journal papers, addressing various topics including data correction, feature reduction, and image classification using machine learning techniques. She has co-authored the remote sensing textbook titled Remote Sensing Digital Image Analysis [Springer-Verlag, 3rd (1999) and 4th eds. (2006)]. Dr. Jia is the author of Field Guide to Hyperspectral / Multispectral Image Processing, SPIE, 2022. These publications are highly cited in the remote sensing and image processing community. She is the Editor-in-Chief of the IEEE TGRS. She is an IEEE Fellow.



Antonio Plaza (IEEE Fellow) is a Full Professor and the Head of the Hyperspectral Computing Laboratory at the Department of Technology of Computers and Communications, University of Extremadura, where he received the M.Sc. degree in 1999 and the PhD degree in 2002, both in Computer Engineering. His main research interests comprise hyperspectral data processing and parallel computing of remote sensing data. He has authored 850 publications in this field, including 439 JCR journal papers, 25 book

chapters, and 330 peer-reviewed conference proceeding papers. He has guest-edited 18 special issues for different journals. Prof. Plaza is a Fellow of IEEE, AAIA, and Academia Europaea, The Academy of Europe. He is a recipient of the recognition of Best Reviewers of the IEEE Geoscience and Remote Sensing Letters (in 2009) and a recipient of the recognition of Best Reviewers of the IEEE TGRS in 2010, for which he served as Associate Editor in 2007-2012. He was also an Associate Editor for IEEE Access (receiving the recognition of Outstanding Associate Editor for the journal in 2017). He is a recipient of the Best Column Award of the IEEE Signal Processing Magazine in 2015, the 2013 Best Paper Award of the JSTARS journal, and the most highly cited paper (2005-2010) in the Journal of Parallel and Distributed Computing. He served as the Director of Education Activities for the IEEE GRSS in 2011-2012, and President of the Spanish Chapter of IEEE GRSS in 2012-2016. He is currently serving as Chair of the Publications Awards Committee of IEEE GRSS and as Chair of the Fellow Evaluations Committee of IEEE GRSS. He served as the Editor-in-Chief of the IEEE Transactions on Geoscience and Remote Sensing journal for five years (2013-2017) and as the Editor-in-Chief of the IEEE Journal on Miniaturization for Air and Space Systems (2019-2020). He has been included in the 2018, 2019, 2020, 2021, 2022 and 2023 Highly Cited Researchers List (Clarivate Analytics). Additional information: <http://sites.google.com/view/antonioplaza>.



Paolo Gamba (IEEE Fellow) is a Professor at the University of Pavia, Italy in the Telecommunications and Remote Sensing Laboratory. He received the Laurea degree in Electronic Engineering "cum laude" from the University of Pavia, Italy, in 1989, and the Ph.D. in Electronic Engineering from the same University in 1993.

He served as Editor-in-Chief of the IEEE Geoscience and Remote Sensing Letters from 2009 to 2013, and as Chair of the Data Fusion Committee of the IEEE Geoscience and Remote Sensing Society (GRSS) from October 2005 to May 2009. He has been elected in the GRSS AdCom from 2014 to 2022, served as GRSS President from 2019 to 2020, and is currently the Editor-in-Chief of the IEEE Geoscience and Remote Sensing Magazine. He also served as Technical Co-Chair of the 2010, 2015, and 2020 IGARSS conferences, in Honolulu (Hawaii), Milan (Italy), and online, respectively.

He is a Fellow of IEEE, IAPR, AIAA, and the Academia Europaea. He has been invited to give keynote lectures and tutorials on several occasions about urban remote sensing, data fusion, EO data for physical exposure, and risk management. He published more than 190 papers in international peer-reviewed journals.



Jon Atli Benediktsson (IEEE Fellow) received the Cand.Sci. degree in electrical engineering from the University of Iceland, Reykjavik, Iceland, in 1984, and the M.S.E.E. and Ph.D. degrees in electrical engineering from Purdue University, West Lafayette, IN, USA, in 1987 and 1990, respectively.

Since July 2015, he has been the President and Rector of the University of Iceland. From 2009 to 2015, he was the Pro-Rector of Science and Academic Affairs and a Professor of Electrical and Computer Engineering at the University of Iceland. He is a Co-Founder of the biomedical startup company Oxymap, Reykjavik. His research interests are in remote sensing, biomedical analysis of signals, pattern recognition, image processing, and signal processing, and he has published extensively in those fields.

Prof. Benediktsson was a member of the 2014 IEEE Fellow Committee. He is a fellow of IEEE, SPIE, Academia Europea, VFI, and Tau Beta Pi. He received the Stevan J. Kristof Award from Purdue University in 1991 as an outstanding graduate student in remote sensing. In 1997, he was a recipient of the Icelandic Research Council's Outstanding Young Researcher Award. In 2000, he was granted the IEEE Third Millennium Medal. In 2004, he was a co-recipient of the University of Iceland's Technology Innovation Award. He received the Yearly Research Award from the Engineering Research Institute of the University of Iceland in 2006, the Outstanding Service Award from the IEEE GRSS in 2007, the IEEE/VFI Electrical Engineer of the Year Award in 2013, the IEEE GRSS Education Award in 2020, the IEEE GRSS David Landgrebe Award in 2018, and the OECE Award from the School of ECE, Purdue University, in 2016. He was a co-recipient of the 2012 IEEE Transactions on Geoscience and Remote Sensing Paper Award, the IEEE GRSS Highest Impact Paper Award in 2013, and the International Journal of Image and Data Fusion Best Paper Award in 2016 and 2018. In 2021, he was honored as a recipient of the Order of the Falcon from the President of Iceland. He was a Highly Cited Researcher (Clarivate Analysis) from 2018 to 2020. He was the 2011–2012 President of the IEEE GRSS and has been on the GRSS AdCom since 2000. He was the Editor-in-Chief of IEEE TGRS from 2003 to 2008 and has served as an Associate Editor for TGRS since 1999, the IEEE Geoscience and Remote Sensing Letters since 2003, and IEEE Access since 2013. He was on the Editorial Board of the Proceedings of the IEEE from 2015 to 2020, is on the International Editorial Board of the International Journal of Image and Data Fusion and the Editorial Board of Remote Sensing, and was the Chairperson of the Steering Committee of IEEE JSTARS from 2007 to 2010.



Jocelyn Chanussot (IEEE Fellow) received the M.Sc. degree in electrical engineering from the Grenoble Institute of Technology (Grenoble INP), Grenoble, France, in 1995, and the Ph.D. degree from the Université de Savoie, Annecy, France, in 1998. Since 1999, he has been with Grenoble INP, where he is currently a Professor of signal and image processing. His research interests include image analysis, hyperspectral remote sensing, data fusion, machine learning, and artificial intelligence. He has been a visiting scholar at Stanford University (USA), KTH (Sweden), and NUS (Singapore). Since 2013, he has been an Adjunct Professor at the University of Iceland. In 2015–2017, he was a visiting professor at the University of California, Los Angeles (UCLA). He holds the AXA chair in remote sensing and is an Adjunct professor at the Chinese Academy of Sciences, Aerospace Information Research Institute, Beijing.

Dr. Chanussot is the founding President of the IEEE Geoscience and Remote Sensing French chapter (2007–2010) which received the 2010 IEEE GRSS Chapter Excellence Award. He has received multiple outstanding paper awards. He was the Vice-President of the IEEE Geoscience and Remote Sensing Society, in charge of meetings and symposia (2017–2019). He was the General Chair of the first IEEE GRSS Workshop on Hyperspectral Image and Signal Processing, Evolution in Remote Sensing (WHISPERS). He was the Chair (2009–2011) and Cochair of the GRS Data Fusion Technical Committee (2005–2008). He was a member of the Machine Learning for Signal Processing Technical Committee of the IEEE Signal Processing Society (2006–2008) and the Program Chair of the IEEE International Workshop on Machine Learning for Signal Processing (2009). He is an Associate Editor for the IEEE Transactions on Geoscience and Remote Sensing and the Proceedings of the IEEE. He was the Editor-in-Chief of the IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing (2011–2015) and an Associate Editor for IEEE Transactions on Image Processing. In 2014 he served as a Guest Editor for the IEEE Signal Processing Magazine. He is a Fellow of the IEEE, a member of the Institut Universitaire de France (2012–2017), and a Highly Cited Researcher (Clarivate Analytics/Thomson Reuters).