

FoMo: Multi-Modal, Multi-Scale and Multi-Task Remote Sensing Foundation Models for Forest Monitoring

Nikolaos Ioannis Bountos^{1,2,3}, Arthur Ouaknine^{1,4}, Ioannis Papoutsis², David Rolnick^{1,4}

¹ Mila - Québec AI Institute

² Orion Lab, National Observatory of Athens & National Technical University of Athens

³ Harokopio University of Athens

⁴ McGill University

Abstract

Forests are vital to ecosystems, supporting biodiversity and essential services, but are rapidly changing due to land use and climate change. Understanding and mitigating negative effects requires parsing data on forests at global scale from a broad array of sensory modalities, and using them in diverse forest monitoring applications. Such diversity in data and applications can be effectively addressed through the development of a large, pre-trained foundation model that serves as a versatile base for various downstream tasks. However, remote sensing modalities, which are an excellent fit for several forest management tasks, are particularly challenging considering the variation in environmental conditions, object scales, image acquisition modes, spatio-temporal resolutions, *etc.* With that in mind, we present the first unified **Forest Monitoring Benchmark** (FoMo-Bench), carefully constructed to evaluate foundation models with such flexibility. FoMo-Bench consists of 15 diverse datasets encompassing satellite, aerial, and inventory data, covering a variety of geographical regions, and including multispectral, red-green-blue, synthetic aperture radar and LiDAR data with various temporal, spatial and spectral resolutions. FoMo-Bench includes multiple types of forest-monitoring tasks, spanning classification, segmentation, and object detection. To enhance task and geographic diversity in FoMo-Bench, we introduce **TalloS**, a global dataset combining satellite imagery with ground-based annotations for tree species classification across 1,000+ categories and hierarchical taxonomic levels. Finally, we propose **FoMo-Net**, a pre-training framework to develop foundation models with the capacity to process any combination of commonly used modalities and spectral bands in remote sensing. This work aims to inspire research collaborations between machine learning and forest biology researchers in exploring scalable multi-modal and multi-task models for forest monitoring and beyond.

Code and Data —

<https://github.com/RolnickLab/FoMo-Bench>

Extended version — <https://arxiv.org/abs/2312.10114>

1 Introduction

Forests are an essential part of Earth’s ecosystems and natural systems, providing critical services on which humanity

depends. They also play a vital role in biodiversity conservation, harboring countless species and maintaining ecological balance. However, forests are rapidly changing as a result of land use decisions, climate change and invasive species (Bonan 2008; Curtis et al. 2018; Hartmann et al. 2022) leading to significant biodiversity decline and economic losses. However, effective forest management and conservation are hampered by a lack of information, highlighting the need for scalable forest monitoring. This gap encompasses multiple specific problems, such as tree species identification, tree size assessment, land cover classification, *etc.*, which also vary considerably across geographies, climatic zones and landscapes.

Remote sensing (RS) has been explored for forest monitoring using various sensors and spatial scales (Guimarães et al. 2020; Michałowska and Rapiński 2021), while increasingly, deep learning methods have been used to address various specific forest-related problems (Kattenborn et al. 2021). Simultaneously, the field of computer vision has seen a growth in foundation models (Bommasani et al. 2022), *i.e.* models flexible to diverse modalities, scales and data regimes while performing well across downstream tasks (Kirillov et al. 2023; Jaegle et al. 2021, 2022). However, it has been noted that general-purpose foundation models in computer vision may be ineffective in specific problem domains (Rolf et al. 2024), presenting a need for foundation models explicitly designed for specific problem spaces, such as Earth observation (Lacoste et al. 2021) and weather or climate modeling (Mukkavilli et al. 2023). Foundation models offer the potential for a flexible backbone that can rapidly be fine-tuned on smaller annotated datasets. Given the diverse nature of RS data and applications, Earth observation foundation models should be able to operate on a wide array of Ground Sampling Distances (GSD), spectral bands, image acquisition modes, and environmental conditions. Assessing the capacity of such models is far more challenging than comparing performance on standard RS datasets. In this work, we propose a framework to evaluate general remote sensing foundation models utilizing the inherent diversity of critical forest-monitoring tasks. Our **Forest Monitoring Benchmark (FoMo-Bench)** brings together a wide range of publicly available datasets for forest monitoring, encompassing significant spatial, temporal, and spectral diversity as illustrated in Fig. 1. FoMo-Bench includes a mixture of

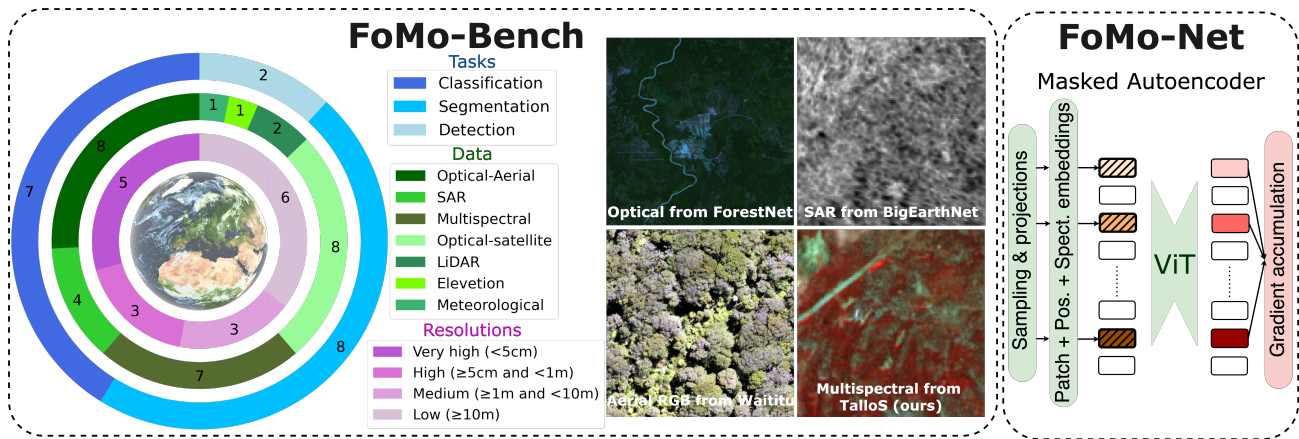


Figure 1: FoMo-Bench evaluation framework and the FoMo-Net pretraining framework for foundation models.

classification, segmentation and detection tasks, from land cover classification to tree crown segmentation. To improve the geographic diversity of FoMo-Bench, we also create **TalloS**, global, multi-modal dataset focused on tree species identification based on worldwide forest inventories (Jucker et al. 2022). Each sample included in TalloS combines data acquired from the Sentinel-2 and Sentinel-1 satellite missions, meteorological data from the ERA-5 (Hersbach et al. 2020) dataset, elevation information, and georeferenced tree species information for multi-label classification.

Finally, we develop **FoMo-Net**, a pretraining framework which can flexibly process most commonly used modalities and spectral bands in RS with a single backbone. **FoMo-Net serves as an initial response to an ambitious research question:** Can we unify foundation models trained on specific individual remote sensors and GSD into a singular, versatile foundation model applicable and adaptable to the requirements (*e.g.* combination of spectral bands, GSD, and geographic location) of a wide range of downstream tasks? To put FoMo-Net to the test, we utilize FoMo-Bench combined with unlabeled data archives to develop an initial multi-modal RS foundation model exploiting the FoMo-Net algorithm with a simple vision transformer (ViT) (Dosovitskiy et al. 2021) encoder. Following common practice in RS foundation models, we assess the capacity of our foundation model by finetuning to the FoMo-Bench tasks. While supervised baselines in FoMo-Bench set a high performance bar, FoMo-Net delivers strong predictive accuracy handling a wide array of input sensors, spectral band combinations, GSD and forest monitoring tasks. It does not require any change in the encoder, marking a crucial first step toward a multi-task, variable-modality model for forest monitoring, stimulating progress in this nascent field and laying the groundwork for future foundation models. We elaborate more on our findings and potential lines of research in Sec. 7. Overall, our contributions can be summarized as follows:

- We provide FoMo-Bench the first, curated benchmark for forest monitoring, with a global spatial coverage and a diverse set of evaluation tasks;
- We introduce TalloS, a novel multi-modal dataset for tree

species multi-label classification;

- We introduce FoMo-Net, a unique pre-training approach designed to train foundation models with great flexibility in terms of the input bands, their ground sampling distance, and geographic location.

2 Related Work

Large-scale RS data has been available for many years through publicly funded projects such as the Sentinel and Landsat missions. Building on such sources, various machine learning benchmarks have been developed, such as BigEarthNet-MM (Sumbul et al. 2021), Sen12MS (Schmitt et al. 2019), Sen4AgriNet (Sykas et al. 2022), and CropHarvest (Tseng et al. 2021). These initiatives have facilitated extensive research, particularly in land use and land cover (LULC) management and forest-related tasks such as tree species identification (Fassnacht et al. 2016).

Forest monitoring: High resolution image recordings from unmanned aerial vehicles (UAVs) are part of recent datasets for automated canopy mapping (Galuszynski et al. 2022), woody invasive species detection (Kattenborn et al. 2019) and tree species detection (Kattenborn et al. 2020). In particular, ReforestTree (Reiersen et al. 2022) combines UAV recordings with inventories for both detecting trees and estimating their aboveground biomass. The NeonTree dataset (Weinstein et al. 2021) is unique in grouping airborne recordings including light detection and ranging (LiDAR), red-blue-green (RGB) and hyperspectral measurements. Although inventories have been used alongside satellite data in Russia with less than a thousand samples (Brieger et al. 2019), we are not aware of any RS dataset providing fine annotations on a large volumetric and geographical scale.

RS benchmarks: Benchmarks across the space of computer vision have moved beyond single tasks to encompass a diversity of related tasks. The large-scale BigEarthNet-MM dataset, built from Sentinel-1 and Sentinel-2 data, is designed to assess both performance and training complexity of deep learning methods (Papoutsis et al. 2023). The SustainBench (Yeh et al. 2021) combines various satellite imagery and street-level images to monitor sustainable de-

velopment goals covering 15 different topics. Geo-Bench (Lacoste et al. 2024) proposes to group RS datasets – including SAR, multispectral, elevation and cloud probabilities data – to tackle six classification and six segmentation tasks focused on LULC. However, Geo-Bench’s limited spatial coverage reduces its potential for broad global applicability. There remains no benchmark focused on forest monitoring worldwide, combining the full suite of forest inventories, aerial and satellite data, and including the several relevant classification, segmentation and detection tasks.

RS foundation models: Foundation models have been investigated for RS datasets (Sun et al. 2022; Cong et al. 2022), grounded mainly on LULC related tasks. While these approaches have shown initial success, they are often limited to specific satellite products and lack diversity in spectral bands and GSD. To address the GSD limitation, ScaleMAE (Reed et al. 2023) was trained on the FMoW-RGB dataset (Christie et al. 2018), adopting a masked image modeling (MIM) approach. This model incorporated GSD positional encoding and a Laplacian-pyramid decoder to enforce multi-scale feature learning. Similarly, Cross-Scale MAE (Tang et al. 2023) leverages FMoW-RGB to learn scale-invariant representations by integrating contrastive loss in the MIM task and introducing a cross-prediction task, where lower-resolution images predict the higher-resolution ones. SpectralGPT (Hong et al. 2024), utilizes the spectral diversity of Sentinel-2 by employing 3D tokens, which presents scalability challenges when dealing with a large number of bands. SkySense (Guo et al. 2024) stands out as the largest RS foundation model to date with 2 billion parameters. It relies on three aligned modalities *i.e.* WorldView RGB, Sentinel-2 multispectral and Sentinel-1 radar data, though this limits its applicability to a broader range of RS sensors. SkySense pretraining utilizes contrastive learning and cross-modal alignment. In a similar fashion, OmniSAT (Astruc et al. 2025) and CROMA (Fuller, Millard, and Green 2023) also explore contrastive and reconstruction losses to efficiently fuse information from spatially aligned RS imagery. Both models use Sentinel-1 and Sentinel-2 data, with CROMA focusing on single timestamps and OmniSAT incorporating time series and high resolution aerial data. USatMAE (Ray, Pankow, and Basu 2016) employs paired NAIP and Sentinel-2 data, using dedicated patch embeddings for each spectral band and a spectral group pooling strategy that groups bands by spatial resolution. These tokens are then processed by a transformer trained via masked image modeling. Similarly, SpectralMAE (Zhu et al. 2023) focuses on paired hyperspectral data, using a transformer encoder with masking and attention mechanisms applied exclusively to the spectral dimension. The recently introduced Presto algorithm (Tseng et al. 2023) proposes a MIM framework for paired data, with a built-in functionality to process a variable number of spectral bands, but is restricted to single-pixel time series. Finally, DOFA (Xiong et al. 2024), aims to create a sensor agnostic transformer encoder through MIM, utilizing a wavelength-conditioned dynamic patch embedding which provides flexibility regarding input spectral bands.

Recently, there has been a surge of interest in flexible yet general RS foundation models that are capable of handling

diverse data sources and solving a wide range of downstream tasks. In this work, we aim to contribute to this line of research by: a) introducing an evaluation framework designed to assess the generalization ability of future models across various Earth observation modalities and conditions, with a specific focus on forest monitoring tasks that have high ecological, environmental, societal and economic impact, and b) proposing a novel pretraining paradigm that leverages any number of unaligned modalities to create sensor-agnostic foundation models. To the best of our knowledge, foundation models have yet to be trained on such a diverse range of inputs or applied to multiple downstream tasks at varying resolutions (from a few centimeters to tens of meters), while incorporating a wide array of spectral bands. This approach addresses a significant gap in the field, as highlighted by recent studies (Ouaknine et al. 2024).

3 FoMo-Bench

We introduce FoMo-Bench, which offers a robust testbed for multimodal RS foundation models, grounded on the diversity of forest-monitoring tasks. Unlike the envisioned multimodal models, conventional methods in FoMo-Bench are task-specific, delivering strong application-optimized baselines though incapable of cross-task adaptation. By carefully selecting datasets composing FoMo-Bench, we have ensured that it (1) encompasses a diversity of geographic locations and environmental conditions, (2) spans a range of task types, and (3) includes input data from a wide variety of sensors *e.g.* optical, SAR, multispectral, DEM and LiDAR, as well as a range of ground sampling distances, from a few centimeters to 60m per pixel. The diversity of the FoMo-Bench data and tasks is illustrated in Fig. 1 and described in more detail below. We have curated all the datasets in FoMo-Bench, spatially split them to train machine learning models and created a reproducible environment for future work with a unified set of PyTorch data loaders. To our knowledge, this is the first benchmark grouping such diverse forest monitoring tasks both in terms of problems they address as well as in regards to the input data.

Satellite-based datasets: FoMo-Bench includes seven datasets involving satellite imagery, namely BigEarthNet-MM, Sen12MS, RapidAI4EO, ForestNet, FiveBillionPixels, TreeSatAI, and TalloS. The **BigEarthNet-MM** and **Sen12MS** (Schmitt et al. 2019) datasets are two of the most commonly used multi-modal datasets for LULC in RS. Both are based on the Copernicus Sentinel missions offering SAR and multispectral data from Sentinel-1 and Sentinel-2 respectively. For the BigEarthNet-MM dataset, we adapt the existing label structure, based on the Corine land cover database (EEA 2020), to encompass 19 distinct land cover classes including broad-leaved, coniferous and mixed forests. While BigEarthNet-MM covers 10 countries in Europe, the Sen12MS dataset has a global spatial coverage (see Fig. 2a), offering annotations for 17 classes based on MODIS land cover maps (Friedl et al. 2010) with 5 different classes for forest cover. FoMo-Bench also includes **RapidAI4EO** (Bhugra et al. 2022) as a third LULC dataset with a unique source of high resolution satellite imagery offering timeseries of Planet data, at 3m resolution, for 500,000

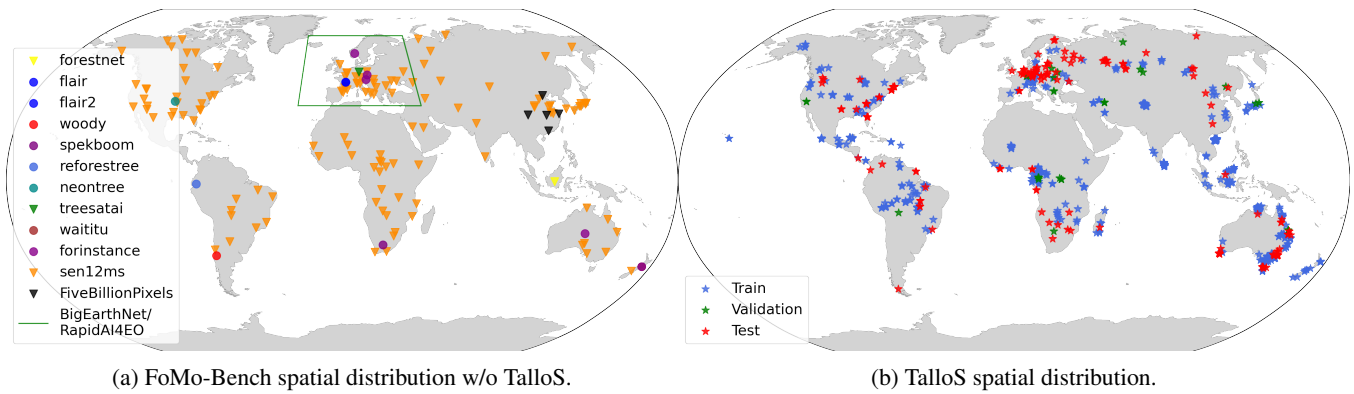


Figure 2: **FoMo-Bench spatial distribution.** (a) Spatial distribution of the datasets included in FoMo-Bench, without TalloS; satellite-based datasets are marked with arrows, aerial-based datasets with circles. (b) Distribution of the train, validation and test splits of the TalloS dataset worldwide.

locations across Europe. We adapt RapidAI4EO labels to encompass 24 classes, 12 of them relevant to forest monitoring. The **ForestNet** dataset (Irvin et al. 2020) aims to identify drivers leading to deforestation and uses Landsat-8 imagery. **FiveBillionPixels** dataset (Tong, Xia, and Zhu 2023) provides high resolution Gaofen-2 images (4m) for LULC of cities in China with 24 categories including 5 forest monitoring classes. The **TreeSatAI** dataset (Ahlsweide et al. 2023) offers multi-modal data (*e.g.* aerial images, SAR and multispectral) for tree species identification, including 20 tree species derived from forest administration data in Lower Saxony, Germany. Finally, we introduce a completely novel dataset, **TalloS**, detailed in Sec. 4.

Aerial-based datasets: The limited spatial resolution offered by free satellite data is not suited for certain RS tasks, such as tree canopy detection or segmentation. FoMo-Bench includes eight datasets with high resolution data from aerial sensor data, namely NeonTree, Woody, Reforestree, Spekboom, Waititu, FLAIR #1, FLAIR #2 and FORinstance. The **NeonTree** dataset (Weinstein et al. 2021) focuses on tree crown detection with RGB imagery, LiDAR and hyperspectral data recorded from airborne systems. FoMo-Bench includes the NeonTree high resolution RGB imagery (0.1m resolution) along with point cloud LiDAR data. The **Woody** dataset (Kattenborn et al. 2019) is composed of RGB images recorded from an UAV, at several cm resolution, for detection of several invasive tree species in Chile. Similarly, the **Spekboom** (Galuszynski et al. 2022) and **Waititu** (Kattenborn et al. 2020) datasets are composed of RGB images recorded from UAVs to map tree species in South Africa and New Zealand, respectively. In the same vein, the **Reforestree** dataset (Reiersen et al. 2022) is designed for tree crown detection with 6 classes using high-resolution RGB images recorded from an UAV, as well as ground measurements, such as tree height and diameter at breast height. The **FORinstance** dataset (Puliti et al. 2023) provides a unique source of data with LiDAR point clouds and semantic annotations collected in 5 countries (see Fig. 2a). Unlike the previous aerial datasets, **FLAIR #1** (Garioud et al. 2022, 2024) and **FLAIR #2** (Garioud et al.

2023, 2024) are large-scale datasets testing spatiotemporal generalization. Both datasets cover wide areas in France, providing high resolution aerial images (0.2m) and elevation maps (plus Sentinel-2 timeseries in the case of FLAIR #2). These datasets test LULC classification across 19 classes.

4 The TalloS Dataset

As shown in Fig. 2, the spatial coverage of pre-existing datasets in FoMo-Bench is limited. We improve the diversity and spatial coverage of FoMo-Bench by designing a new dataset, TalloS. This dataset couples manual tree inventories with multispectral and SAR timeseries imagery for a fine grained multi-label classification challenge. As described in the previous section, various forest monitoring tasks have relied on high-resolution UAV data, which, despite offering detailed insights, suffer from limited spatial coverage due to human and hardware constraints (Galuszynski et al. 2022; Reiersen et al. 2022; Kattenborn et al. 2019, 2020). TalloS is a large-scale multi-modal dataset of satellite-based imagery for tree species multi-label classification. It is based on the global Tallo database (Jucker et al. 2022), containing almost 500,000 manual georeferenced recordings aggregating information from forest inventories including the tree species as well as tree height, stem diameter, and crown radius. Due to the modest GSD of publicly accessible satellites ($\approx 10\text{m}$), we design TalloS to concentrate on areas containing a minimum of 50 trees, as regions with lower tree populations lack distinctive signatures in satellite imagery. This choice balances diversity and strong signal for model prediction. There are $\approx 335\text{k}$ locations including 50 or more trees, but only $\approx 251\text{k}$ with 100 or more, reducing the dataset size by 25.1% and the number of genera by more than 11%. TalloS data represent a time series from 2016 to 2023 with a temporal frequency of 15 days. Each timestep consists of a multispectral image from Sentinel-2, a SAR image from Sentinel-1, a digital elevation model (DEM) from Copernicus GLO-30 at 10m per pixel resolution, along with meteorological information from ERA-5. Sentinel-2 data is selected with cloud coverage below 30%, and Sentinel-1 captures are aligned as closely as possible to Sentinel-2 acquisi-

Dataset	Sensors	Spatial res.	Sampling Weight
SSL4EO-Landsat	Landsat 8-9	Low	0.2
RapidAI4EO	Planet, S-2	Medium, Low	0.2
TalloS	S-1, S-2, DEM	Low	0.2
FLAIR #1	Aerial	High	0.1
FiveBillionPixels	Gaofen-2	Medium	0.2
UAV-datasets	Aerial	Very high	0.1

Table 1: **FoMo-Net pre-training datasets.** Our framework has been trained with four satellite-based datasets and four aerial-based datasets. Their spatial resolution is defined as very high ($< 5\text{cm}$), high ($\geq 5\text{cm}$ and $< 1\text{m}$), medium ($\geq 1\text{m}$ and $< 10\text{m}$) or low ($\geq 10\text{m}$). The sampling weight is based on the frequency of each source in our pretraining datasets and is used for the modality sampling. S-1 and S-2 stand for Sentinel-1 and Sentinel-2 respectively.

tions. Spectral bands with lower GSD are resampled to 10m per pixel for consistency with conventional models. Each RS sample of TalloS is mapped to ground-validated and fine grained measurements of trees in its given area, including tree species, genus and family as well as their specific traits. For the purpose of this study, we focus on the genus level labels, as distinguishing between species may require higher resolution imagery. Throughout this paper, unless explicitly stated otherwise, we refer to the various genus categories of TalloS as its classes. TalloS represents a fine-grained classification problem involving $>1,000$ classes with a heavily skewed data distribution. The task is framed as multi-label classification, with the goal of predicting every tree genus that is depicted in the image timeseries. In Fig. 2b, we illustrate the wide spatial distribution of TalloS as well as the proposed training, validation and test splits. Our goal is to maximize the geographic spatial coverage of each split while ensuring that (a) regions used for training and testing do not overlap, (b) most classes present in the validation and test sets are also included in the training set. Accommodating both conditions proved challenging due to the limited geographic distribution of many tree genera. We present TalloS samples along with its class distribution in Appendix A.

5 FoMo-Net

Building on FoMo-Bench and inspired by the necessity for a generalized foundation model, we introduce FoMo-Net. FoMo-Net is a pretraining paradigm for learning to process all of the most common modalities in the RS domain with a single, sensor-agnostic foundation model. In this work we treat each spectral band of each sensor as a separate modality. Our approach is driven by three core considerations. First, it should maximize flexibility to process various input settings, without relying on specific sets of sensory modalities. Second, it should have the capacity to process information and generate meaningful representations for the whole globe. Third, it should be applicable to a wide range of downstream tasks.

Pre-training data: FoMo-Net pre-training scheme con-

tains rich multi-sensor information, both paired and unpaired, from most parts of the world. In particular, we use the RapidAI4EO dataset as a source of Planet and Sentinel-2 data; TalloS containing DEM, Sentinel-1, and Sentinel-2; SSL4EO-Landsat (Stewart et al. 2024) to acquire global information from Landsat 8 and 9; a combination of all available UAV datasets in FoMo-Bench; and the FiveBillionPixel dataset providing a unique source of Gaofen-2 high resolution satellite imagery. The datasets used in the FoMo-Net pre-training framework are detailed in Tab. 1. The proposed framework is able to process any combination of the 36 most common modalities in RS, stemming from Sentinel-1, Sentinel-2, Landsat 8-9, Planet, Gaofen-2 and UAV sensors with ground sampling distance spanning from a few centimeters to 60m per pixel. In this work, a distinct modality is defined as any band with a unique combination of wavelength and spatial resolution. Given its importance in Earth observation, digital elevation models (DEM) are also included in our definition as a specific band. The FoMo-Net pre-training pipeline is illustrated in Fig. 3 and will be detailed in the following sections. Additionally, the pseudocode of the pre-training pipeline is summarized in Appendix B. Point-cloud data are excluded from the pretraining set and initial FoMo-Net design due to the need for specialized processing, which is deferred to future work (Sec. 7).

Approach to variable spectral bands: Let $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_n\}$ be a set of n datasets, and $X = \{X_1, \dots, X_m\}$ be a set of m spectral bands, where $X_i \in \mathbb{R}^{H \times W}$ where H and W represent respectively the height and width dimensions. At each iteration, a training batch \mathcal{B} contains a variable number of spectral bands, each sampled with respective probabilities $\alpha_i \in \{\alpha_1, \dots, \alpha_m\}$. In our experiments, we set α_i proportional to the frequency with which the i -th band occurs in \mathcal{D} (see Tab. 1).

Each band X_i is tokenized into N patches of size P according to the transformation s , so that $s(X_i) \in \mathbb{R}^{N \times \lfloor \frac{H}{P} \rfloor \times \lfloor \frac{W}{P} \rfloor}$. The tokenized input $s(X_i)$ is then embedded using a linear transformation $t_i \in \{t_1, \dots, t_m\}$ learnt during the optimization process as $t_i: \mathbb{R}^{N \times \lfloor \frac{H}{P} \rfloor \times \lfloor \frac{W}{P} \rfloor} \rightarrow \mathbb{R}^{N \times d}$. We considered two projection setups: The first (noted FoMo-Net₁) uses a single linear projection for all spectral bands, while the second (denoted FoMo-Net_m) projects each spectral band with its own linear projection. Our initial experiments suggest that FoMo-Net₁ yields significantly better results. The input to our model, then, is provided as:

$$t_i(s(X_i)) \oplus \mathbf{S} \oplus \mathbf{P}, \quad (1)$$

where $\mathbf{S} \in \mathbb{R}^{N \times d}$, $\mathbf{P} \in \mathbb{R}^{N \times d}$ are respectively the learnt spectral and positional embeddings, and \oplus denotes the element-wise sum.

Spectral band MAE: The next step of the procedure follows a masked autoencoding (MAE) framework (He et al. 2022) that is shared by various high-performing foundation models for RS (see Sec. 2). However, our method differs from a typical MAE in two key aspects (see Fig. 3).

First, we introduce a random spectral band selection process by sampling in both \mathcal{D} and X , representing the datasets and spectral bands, respectively. The number of selected spectral bands is randomly defined at each iteration. This

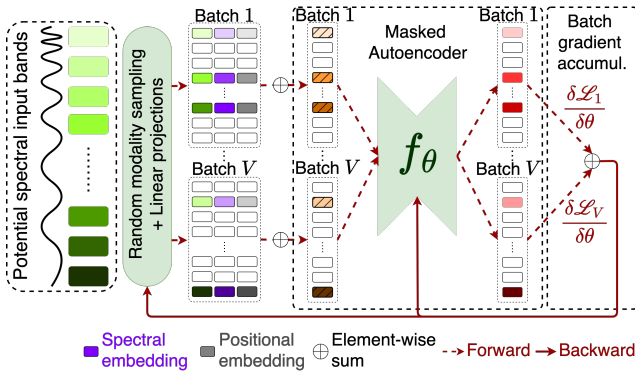


Figure 3: **FoMo-Net pre-training framework.** Considering a set of potential spectral input bands and a sub-training process of V batches, each batch includes a sub-set of bands w.r.t. to the sampled dataset and modalities it contains. For each batch, the input band embeddings, the spectral embeddings and the positional embeddings are summed element-wise to generate individual embeddings per band. The masked autoencoder, $f_\theta(\cdot)$ parameterized by θ , reconstructs the partially masked inputs, illustrated with dashed squared, of a given batch. The gradients of the loss $\mathcal{L}(\cdot)$ are accumulated and backpropagated through both $f_\theta(\cdot)$ and the linear projections. For the sake of clarity, we do not include the elevation modality which is not considered as a spectral band, but included in the pre-training scheme.

approach promotes modeling interactions between a wide range of band combinations, creating a highly flexible encoder able to process any set of spectral bands. Next, we incorporate both spectral and positional embeddings, allowing the model to capture information regarding the spatial position of the patch and its spectral band, as detailed in Eq. 1. Both types of embeddings are trainable and randomly initialized. Given the randomness in the input space, the spectral embedding is of utmost importance as no other information regarding the nature of the input is provided. Finally, as each band is tokenized independently the transformer is able to apply the attention mechanism to both spectral and spatial dimensions.

Batch gradient accumulation: The combination of the proposed modality sampling and token masking is particularly computationally efficient in contrast to the exhaustive generation and processing of tokens from the entire spectrum of information. However, randomly choosing datasets and spectral bands at each training iteration results in highly heterogeneous batches, potentially destabilizing the optimization process. To address this, the FoMo-Net pre-training algorithm exploits gradient accumulation (see Fig. 3). Letting f_θ be a neural network parameterized by θ , and \mathcal{L} be a loss function, the gradient is accumulated during V forward passes according to $\sum_{v=1}^V \delta \mathcal{L}_v(f_\theta(\mathcal{B}^v)) / \delta \theta$, where \mathcal{B}^v is the v -th batch of input data (see Eq. 1), and it is this sum that is backpropagated through f_θ .

FoMo-Net backbone: In this work we apply FoMo-Net on a ViT encoder with 12 layers, and 12 attention heads. We

limit the maximum number of sampled spectral bands at the pre-training stage to 4 and the pixel resolution to 64×64 for efficiency and train for 300 epochs.

Modelling challenges: FoMo-Net must account for interactions across a wide range of spectral bands potentially useful for downstream tasks, unlike standard foundation models tailored to specific sensors. Additionally, varying GSD requires capturing features and patterns at multiple scales, some of which are a) underrepresented and b) lack sufficient data points with wide spatial coverage. Finally, given the diversity of information in the input and the large-scale pretraining dataset, scaling up the backbone could further improve performance. However, this would require significantly more computational resources.

6 Experiments

We train a set of standard computer vision models tailored to each specific task. Each baseline is trained in a supervised fashion on the data for that single task due to the heterogeneity of the tasks in both input and output. In the case of multi-label classification and object detection, we also evaluate a simple approach using supervised learning on multiple datasets *i.e.* PolyNet. PolyNet denotes a ResNet-50 model pretrained on several datasets simultaneously. Since PolyNet has been trained in a supervised fashion on FoMo-Bench tasks, we evaluate it via linear probing either in-domain (denoted PolyNet_{in}), *i.e.* on a dataset included in pre-training, or out-of-domain (denoted PolyNet_{out}), *i.e.* on a dataset not included in pre-training. For object detection and semantic segmentation we train a task-specific head to learn the downstream task for each dataset independently. For all experiments, we present the average performance along with the standard deviation over three independent runs using different random seeds. Detailed training settings and qualitative results are provided in Appendices C and D respectively.

Semantic segmentation of grid-based data: FoMo-Bench includes seven datasets suited to semantic segmentation. Our baseline models are UNet (Ronneberger, Fischer, and Brox 2015), DeepLabv3 (Chen et al. 2017) and UPerNet (Xiao et al. 2018).

For UNet and DeepLabv3 architectures, we use a ResNet50 (He et al. 2016) backbone pretrained on ImageNet. The UPerNet is based on a Swin-Base backbone from (MMSegmentation Contributors 2020). We assess FoMo-Net by extracting the encoded tokens, upsampling and feeding them to a small decoder similar to FloodViT (Bountos et al. 2024). We report the mean intersection over union (mIoU) and the F1-Score at micro level for all datasets. Results are presented in Tab. 3.

Classification: Due to the large GSD of sensors in RS, classification tasks are usually multi-label. The multi-label classification setting is thus used for BigEarthNet-MM, Sen12MS, RapidAI4EO, TreeSat AI and TalloS for either LULC or tree species recognition. Our baselines are ResNet50, ViT and ConvNext (Liu et al. 2022) models. We initialize all models with weights pretrained on ImageNet. We pretrain PolyNet on RapidAI4EO, TalloS, Sen12MS and BigEarthNet, with their common bands. In-domain evaluation is performed on these datasets, while TreeSatAI is used

Dataset	Modalities	#C	Method	#P (Mil.)	F1 micro	F1 macro
BigEarthNet (Sumbul et al. 2021)	Multispec. SAR	19	ResNet50	23.58	77.18 (± 0.32)	70.32 (± 1.32)
			ViT	87.97	78.09 (± 0.34)	72.29 (± 0.61)
			ConvNext	87.60	<u>77.25</u> (± 0.26)	<u>69.66</u> (± 0.73)
			PolyNet _{in}	0.038	73.51 (± 0.42)	66.07 (± 0.20)
			FoMo-Net ₁	66.39	<u>73.42</u> (± 0.10)	<u>61.36</u> (± 0.55)
			FoMo-Net _m	73.57	66.11 (± 0.07)	50.0 (± 0.25)
Sen12MS (Schmitt et al. 2019)	Multispec. SAR	17	ResNet50	23.57	51.34 (± 1.71)	<u>32.62</u> (± 0.33)
			ViT	87.97	<u>51.49</u> (± 1.24)	31.77 (± 1.72)
			ConvNext	87.60	50.86 (± 3.91)	30.02 (± 2.01)
			PolyNet _{in}	0.034	51.70 (± 0.70)	33.54 (± 0.74)
			FoMo-Net ₁	66.39	<u>42.87</u> (± 1.69)	<u>25.15</u> (± 0.45)
			FoMo-Net _m	73.57	38.07 (± 0.65)	19.44 (± 0.45)
TalloS (ours)	Multispec. SAR DEM	1160	ResNet50	26.34	27.31 (± 1.02)	1.35 (± 0.09)
			ViT	89.20	15.23 (± 3.97)	0.37 (± 0.12)
			ConvNext	88.98	9.08 (± 2.58)	0.43 (± 0.12)
			PolyNet _{in}	2.79	4.57 (± 0.88)	0.13 (± 0.03)
			FoMo-Net ₁	67.27	<u>18.50</u> (± 2.33)	<u>0.52</u> (± 0.15)
			FoMo-Net _m	74.45	14.25 (± 2.81)	0.27 (± 0.09)
RapidAI4EO (Bhugra et al. 2022)	Optical Multispec.	24	ResNet50	23.60	<u>62.84</u> (± 0.17)	<u>46.33</u> (± 0.81)
			ViT	88.37	51.04 (± 2.13)	25.59 (± 3.33)
			ConvNext	87.61	64.07 (± 0.23)	47.20 (± 5.24)
			PolyNet _{in}	0.053	59.21 (± 0.29)	41.11 (± 0.40)
			FoMo-Net ₁	66.40	<u>54.69</u> (± 0.06)	<u>33.47</u> (± 0.56)
			FoMo-Net _m	73.58	49.12 (± 5.59)	24.63 (± 8.46)
TreeSatAI (Ahlsweide et al. 2023)	Multispec.	15	ResNet50	23.56	53.01 (± 0.86)	30.19 (± 1.7)
			ViT	87.57	<u>55.48</u> (± 1.17)	<u>36.63</u> (± 1.99)
			ConvNext	87.60	64.15 (± 0.29)	43.99 (± 1.53)
			PolyNet _{out}	0.03	42.41 (± 0.11)	21.41 (± 0.28)
			FoMo-Net ₁	66.39	<u>45.02</u> (± 0.51)	<u>18.93</u> (± 0.24)
			FoMo-Net _m	73.57	27.96 (± 1.02)	12.08 (± 0.41)

Table 2: **Performance on classification tasks.** #C denotes the number of classes and #P the number of trainable parameters (in millions). We report the F1-score at the micro and macro levels. We **emphasize** the best results and underline the second best.

for out-of-domain evaluation. We report micro and macro F1 scores for all datasets in Tab. 2.

Object detection The object detection task is useful in forest monitoring to count and precisely locate trees, focusing in some cases on particular species. We use the NeonTree and ReforestTree datasets to test binary and multi-class object detection respectively, testing the Faster R-CNN (Ren et al. 2016), RetinaNet (Lin et al. 2020) and YOLOS (Fang et al. 2021) models (see Tab. 4). We measure performance by the mIoU, the mean average precision with an IoU = 0.5 threshold (mAP₅₀) and the mean average precision with a variable threshold of IoU $\in [0.5, 0.95]$ (mAP_{50:95}).

The models include ResNet-50 and DeiT-S (Touvron et al. 2021) backbones pre-trained on the COCO dataset. We also include a Faster R-CNN with a DINOv2 backbone (Oquab et al. 2024) fully finetuned. FoMo-Net is augmented with a randomly initialized Faster R-CNN head for finetuning. We pretrain PolyNet simultaneously on NeonTree and ReforestTree. In-domain evaluation (PolyNet_{in}) is then conducted on each dataset separately. For the out-of-domain evaluation (PolyNet_{out}), we pretrain the backbone on one dataset while

evaluating on the other.

Semantic segmentation of point clouds: Processing point cloud representations from LiDAR sensors aims at better understanding the geometry of trees at large scale. FoMo-Bench includes NeonTree and FORinstance datasets for binary and multi-class point cloud segmentation respectively. For these datasets, we use the baselines PointNet (Qi et al. 2017a), PointNet++ (Qi et al. 2017b) and Point Transformer (Zhao et al. 2021). These experiments are presented in Appendix E.

7 Discussion

FoMo-Bench: Interestingly, traditional CNNs consistently rank among the top-performing models in both classification and segmentation tasks (Tab. 3, 2). ResNet50 exceptional performance on TalloS is striking, with FoMo-Net₁ coming close. TalloS high number of classes along with their skewed distribution make it an extreme multilabel classification problem, reflecting the complexities of global-scale forest monitoring. Tree species are distributed unequally around the globe (see Fig. 2b), with many classes present

Dataset	Modalities	#C	Method	#P (Mil.)	F1 micro	mIoU
Waititu (Kattenborn et al. 2020)	Aerial RGB	3	UNet	32.52	<u>77.96</u> (± 3.36)	<u>64.0</u> (± 4.55)
			DeepLabv3	26.67	80.81 (± 2.81)	67.90 (± 3.96)
			UPerNet	121.30	75.36 (± 3.04)	60.56 (± 3.94)
			FoMo-Net ₁	79.31	74.26 (± 2.00)	59.10 (± 2.51)
			FoMo-Net _m	86.49	71.52 (± 0.0)	55.66 (± 0.0)
FLAIR #1 (Garioud et al. 2022, 2024)	Aerial multisp.	19	UNet	32.52	<u>72.80</u> (± 0.17)	<u>57.24</u> (± 0.22)
			DeepLabv3	26.68	74.01 (± 0.3)	58.75 (± 0.39)
			UPerNet	121.30	63.89 (± 1.03)	46.94 (± 1.12)
			FoMo-Net ₁	86.39	66.16 (± 0.05)	49.43 (± 0.06)
			FoMo-Net _m	93.57	60.22 (± 0.39)	43.08 (± 0.39)
FLAIR #2 (Garioud et al. 2023, 2024)	Aerial multisp.	19	UNet	32.56	72.99 (± 1.80)	57.50 (± 2.22)
			DeepLabv3	26.7	83.97 (± 0.27)	72.37 (± 0.40)
			UPerNet	121.33	<u>76.93</u> (± 0.28)	<u>62.52</u> (± 0.37)
			FoMo-Net ₁	121.78	64.69 (± 0.07)	47.81 (± 0.08)
			FoMo-Net _m	128.96	57.95 (± 0.17)	40.80 (± 0.16)
Spekboom (Galuszynski et al. 2022)	Aerial RGB	2	UNet	32.52	97.70 (± 0.02)	<u>95.50</u> (± 0.03)
			DeepLabv3	26.67	97.57 (± 0.17)	95.27 (± 0.34)
			UPerNet	121.30	97.81 (± 0.01)	95.72 (± 0.03)
			FoMo-Net ₁	79.31	95.59 (± 0.06)	91.56 (± 0.11)
			FoMo-Net _m	86.49	91.83 (± 0.05)	70.97 (± 1.15)
FiveBillionPixels (Tong, Xia, and Zhu 2023)	Optical	22	UNet	32.52	74.40 (± 0.29)	59.24 (± 0.36)
			DeepLabv3	26.68	73.75 (± 0.32)	58.42 (± 0.40)
			UPerNet	121.31	67.43 (± 4.19)	51.02 (± 4.88)
			FoMo-Net ₁	82.85	<u>72.75</u> (± 0.19)	<u>57.18</u> (± 0.24)
			FoMo-Net _m	90.03	67.39 (± 0.30)	50.82 (± 0.34)
ForestNet (Irvin et al. 2020)	Multispec.	5	UNet	32.52	<u>95.26</u> (± 0.02)	<u>90.96</u> (± 0.04)
			DeepLabv3	26.68	95.27 (± 0.03)	90.97 (± 0.05)
			UPerNet	121.31	95.23 (± 0.05)	90.90 (± 0.10)
			FoMo-Net ₁	79.31	<u>95.22</u> (± 0.02)	<u>90.89</u> (± 0.04)
			FoMo-Net _m	86.49	95.24 (± 0.0)	90.92 (± 0.0)
Woody (Kattenborn et al. 2019)	Aerial RGB	4	UNet	32.52	<u>93.37</u> (± 0.23)	<u>87.57</u> (± 0.42)
			DeepLabv3	26.67	93.30 (± 0.40)	87.44 (± 0.71)
			UPerNet	121.30	93.87 (± 0.17)	88.46 (± 0.31)
			FoMo-Net ₁	79.31	89.7 (± 0.04)	81.32 (± 0.06)
			FoMo-Net _m	86.49	84.33 (± 0.84)	72.93 (± 1.27)

Table 3: **Semantic segmentation tasks for grid-based data.** #C denotes the number of classes and #P the number of classes and number of trainable parameters (in millions). We report the F1-score and the mean intersection over union (mIoU) at the micro levels. We **emphasize** the best results and underline the second best.

only in specific locations. Incorporating tree species occurrence maps as biological prior knowledge could prove helpful. PolyNet’s performance raises interesting research questions on the utilization of the semantic information of forest monitoring datasets for the creation of foundation models. With the exception of TalloS, where it is severely outperformed by FoMo-Net, PolyNet shows good in-domain predictive skill, including attaining the best performance for Sen12MS. FoMo-Net pretraining yielded promising results, even though it is still outperformed by specialized models. Notably, in some cases it surpasses the original ViT, particularly in TalloS and RapidAI4EO. As discussed in Sec. 5, digesting such diverse information simultaneously might benefit from a larger encoder. Despite the lack of built-in functionality in ViT for processing hierarchical features in dense

prediction, FoMo-Net’s performance in segmentation tasks (Table 3) is highly competitive to supervised specialized architectures *e.g.* UPerNet. Although the supervised baselines in FoMo-Bench set a high bar for comparison, they are not directly comparable to models like FoMo-Net. While supervised-learning models can excel when trained on large datasets, they struggle when applied to different tasks and are fundamentally limited to the specific tasks they were trained on. FoMo-Net pushes the boundaries of this emerging field, establishing a basis for the development of future foundation models. Another key finding is the superiority of FoMo-Net₁ over FoMo-Net_m. Contrary to intuition, a single shared projection to the embedding space outperforms dedicated projections for each spectral band. This enhances FoMo-Net’s flexibility as new modalities can rely on the

Dataset	Modality	#C	Method	Backbone	#P (Mil.)	mIoU	mAP ₅₀	mAP _{50:95}
NeonTree (Weinstein et al. 2021)	Airborne	1	F. R-CNN	ResNet-50	41.08	61.49 (± 0.087)	38.37 (± 1.68)	13.72 (± 0.68)
			RetinaNet	ResNet-50	36.13	<u>56.27</u> (± 1.45)	6.42 (± 4.37)	2.10 (± 1.36)
			YOLOS	DeiT-S	30.65	53.36 (± 0.18)	28.34 (± 0.48)	<u>12.64</u> (± 0.31)
			F. R-CNN	PolyNet _{in}	14.50	58.42 (± 0.04)	<u>35.93</u> (± 0.10)	11.95 (± 0.03)
			F. R-CNN	PolyNet _{out}	14.50	56.05 (± 0.10)	27.77 (± 0.21)	8.70 (± 0.07)
			F. R-CNN	DINOv2	94.36	<u>44.65</u> (± 0.65)	1.08 (± 0.73)	0.22 (± 0.17)
			F. R-CNN	FoMo-Net ₁	81.37	51.95 (± 0.46)	<u>5.67</u> (± 0.82)	<u>1.52</u> (± 0.25)
			F. R-CNN	FoMo-Net _m	88.55	45.12 (± 0.42)	1.1 (± 0.56)	0.24 (± 0.13)
ReforesTree (Reiersen et al. 2022)	Aerial	6	F. R-CNN	ResNet-50	41.10	64.67 (± 0.78)	6.83 (± 2.33)	<u>2.91</u> (± 1.11)
			RetinaNet	ResNet-50	36.21	65.49 (± 0.45)	4.69 (± 1.02)	2.13 (± 0.53)
			YOLOS	DeiT-S	30.65	62.85 (± 0.60)	2.30 (± 0.49)	1.33 (± 0.23)
			F. R-CNN	PolyNet _{in}	14.52	62.87 (± 0.25)	8.14 (± 0.11)	3.30 (± 0.03)
			F. R-CNN	PolyNet _{out}	14.52	59.47 (± 0.48)	<u>7.19</u> (± 0.22)	2.73 (± 0.09)
			F. R-CNN	DINOv2	94.38	46.90 (± 0.36)	0.08 (± 0.02)	0.02 (± 0.01)
			F. R-CNN	FoMo-Net ₁	81.39	<u>51.63</u> (± 0.25)	<u>1.32</u> (± 0.50)	<u>0.46</u> (± 0.16)
			F. R-CNN	FoMo-Net _m	88.57	46.38 (± 0.24)	0.20 (± 0.08)	0.04 (± 0.02)

Table 4: **Object detection evaluation.** #C denotes the number of classes and #P number of trainable parameters (in millions). We report the mIoU and the mean average precision (mAP) for both IoU = 0.5 and IoU \in [0.5, 0.95] thresholds. We **emphasize** the best results and underline the second best.

shared embedding. We compare FoMo-Net with other foundation models in Appendix F.

Limitations and next steps: While FoMo-Bench offers strong baselines across a broad range of forest monitoring tasks, assessing the true capability of flexible foundation models requires more than just performance metrics. Given the diversity in input modalities we need methods to evaluate how well these flexible models recognize the unique characteristics of each modality, their inter-modal relationships, and the features of the Earth’s surface across multiple scales. Also, as FoMo-Net tokenizes spectral bands independently the amount of tokens per sample increases proportionally to the number of spectral bands, creating a computational bottleneck due to the memory requirements and the quadratic complexity of attention posing challenges for scalability and making optimization a priority. Given PolyNet’s potential, combining FoMo-Net with supervised pre-training might result in a flexible pre-training method usable under variable computational constraints and with variable amounts of labeled data. Finally, extending FoMo-Net to non-grid data, *e.g.* point clouds, could prove valuable.

8 Conclusion

This work presents FoMo-Bench, a global benchmark, providing a robust evaluation framework for future forest monitoring flexible foundation models. To further expand the tasks and spatial coverage of FoMo-Bench, we introduce TalloS, a global, multi-temporal and multi-modal dataset for tree species classification. Building on these contributions, we craft a pre-training framework to construct flexible foundation models *i.e.* FoMo-Net, and provide an initial assessment on FoMo-Bench identifying current limitations and potential opportunities for this promising field.

Acknowledgments

This work has received funding from the IVADO program on “AI, Biodiversity and Climate Change”, the Canada CIFAR AI Chairs program and the European Union’s Horizon Europe research and innovation program under grant agreement No. 101130544 (ThinkingEarth).

References

- Ahlswede, S.; Schulz, C.; Gava, C.; et al. 2023. *TreeSatAI Benchmark Archive: a multi-sensor, multi-label dataset for tree species classification in remote sensing.* *Earth System Science Data*.
- Astruc, G.; Gonthier, N.; Mallet, C.; et al. 2025. OmniSat: Self-supervised Modality Fusion for Earth Observation. In *ECCV 2024*, 409–427. Springer Nature Switzerland.
- Bhugra, P.; Bischke, B.; Werner, C.; et al. 2022. Rapidai4Eo: Mono-and Multi-Temporal Deep Learning Models for Updating the Corine land Cover Product. In *2022 IEEE IGARSS*, 2247–2250.
- Bommasani, R.; Hudson, D.; Adeli, E.; et al. 2022. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Bonan, G. B. 2008. Forests and Climate Change: Forcings, Feedbacks, and the Climate Benefits of Forests. *Science*, 320(5882): 1444–1449.
- Bountos, N. I.; Sdraka, M.; Zavras, A.; Karasante, I.; Karavias, A.; Herekakis, T.; Thanasou, A.; Michail, D.; and Pappoutsis, I. 2024. Kuro Siwo: 33 billion m^2 under the water. A global multi-temporal satellite dataset for rapid flood mapping. *arXiv:2311.12056*.
- Brieger, F.; Herzsuh, U.; Pestryakova, L.; et al. 2019. Advances in the Derivation of Northeast Siberian Forest Met-

- rics Using High-Resolution UAV-Based Photogrammetric Point Clouds. *Remote Sensing*, 11(12): 1447.
- Chen, L.-C.; Papandreou, G.; Schroff, F.; and Adam, H. 2017. Rethinking atrous convolution for semantic image segmentation. arXiv:1706.05587.
- Christie, G.; Fendley, N.; Wilson, J.; and Mukherjee, R. 2018. Functional map of the world. In *Proceedings of the IEEE CVPR*, 6172–6180.
- Cong, Y.; Khanna, S.; Meng, C.; et al. 2022. Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery. *NeurIPS*, 35: 197–211.
- Curtis, P. G.; Slay, C. M.; Harris, N. L.; Tyukavina, A.; and Hansen, M. C. 2018. Classifying drivers of global forest loss. *Science*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; et al. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv:2010.11929.
- EEA. 2020. CORINE Land Cover 2018 (vector), Europe, 6-yearly. European Environment Agency.
- Fang, Y.; Liao, B.; Wang, X.; et al. 2021. You Only Look at One Sequence: Rethinking Transformer in Vision through Object Detection. In *NeurIPS*, volume 34, 26183–26197.
- Fassnacht, F.; Latifi, H.; Stereńczak, K.; et al. 2016. Review of studies on tree species classification from remotely sensed data. *Remote Sensing of Environment*, 186: 64–87.
- Friedl, M.; Sulla-Menashe, D.; Tan, B.; et al. 2010. MODIS Collection 5 global land cover: Algorithm refinements and characterization of new datasets. *Remote Sensing of Environment*, 114(1): 168–182.
- Fuller, A.; Millard, K.; and Green, J. R. 2023. CROMA: Remote Sensing Representations with Contrastive Radar-Optical Masked Autoencoders. In *NeurIPS*, volume 36.
- Galuszynski, N.; Duker, R.; Potts, A.; and Kattenborn, T. 2022. Automated mapping of *Portulacaria afra* canopies for restoration monitoring with convolutional neural networks and heterogeneous unmanned aerial vehicle imagery. *PeerJ*, 10: e14219.
- Garioud, A.; De Wit, A.; Poupée, M.; et al. 2023. FLAIR #2: textural and temporal information for semantic segmentation from multi-source optical imagery. arXiv:2305.14467.
- Garioud, A.; Gonthier, N.; Landrieu, L.; et al. 2024. FLAIR: a Country-Scale Land Cover Semantic Segmentation Dataset From Multi-Source Optical Imagery. In *NeurIPS*, volume 36.
- Garioud, A.; Peillet, S.; Bookjans, E.; Giordano, S.; and Wattrelos, B. 2022. FLAIR #1: semantic segmentation and domain adaptation dataset. arXiv:2211.12979.
- Guimarães, N.; Pádua, L.; Marques, P.; et al. 2020. Forestry remote sensing from unmanned aerial vehicles: A review focusing on the data, processing and potentialities. *Remote Sensing*, 12(6): 1046.
- Guo, X.; Lao, J.; Dang, B.; et al. 2024. Skysense: A multi-modal remote sensing foundation model towards universal interpretation for earth observation imagery. In *CVPR*, 27672–27683.
- Hartmann, H.; Bastos, A.; Das, A.; et al. 2022. Climate change risks to global forest health: emergence of unexpected events of elevated tree mortality worldwide. *Annual Review of Plant Biology*, 73(1): 673–702.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *CVPR*, 16000–16009.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hersbach, H.; Bell, B.; Berrisford, P.; et al. 2020. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730): 1999–2049.
- Hong, D.; Zhang, B.; Li, X.; et al. 2024. SpectralGPT: Spectral remote sensing foundation model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8): 5227–5244.
- Irvin, J.; Sheng, H.; Ramachandran, N.; et al. 2020. Forestnet: Classifying drivers of deforestation in indonesia using deep learning on satellite imagery. Tackling Climate Change with Machine Learning NeurIPS Workshop, arXiv:2011.05479.
- Jaegle, A.; Borgeaud, S.; Alayrac, J.-B.; et al. 2022. Perceiver IO: A General Architecture for Structured Inputs & Outputs. In *ICLR*.
- Jaegle, A.; Gimeno, F.; Brock, A.; et al. 2021. Perceiver: General Perception with Iterative Attention. In *ICML*, 4651–4664.
- Jucker, T.; Fischer, F.; Chave, J.; et al. 2022. Tallo: A global tree allometry and crown architecture database. *Global change biology*, 28(17): 5254–5268.
- Kattenborn, T.; Eichel, J.; Wisler, S.; et al. 2020. Convolutional Neural Networks accurately predict cover fractions of plant species and communities in Unmanned Aerial Vehicle imagery. *Remote Sensing in Ecology and Conservation*, 6(4): 472–486.
- Kattenborn, T.; Leitloff, J.; Schiefer, F.; and Hinz, S. 2021. Review on Convolutional Neural Networks (CNN) in vegetation remote sensing. *ISPRS journal of photogrammetry and remote sensing*, 173: 24–49.
- Kattenborn, T.; Lopatin, J.; Förster, M.; et al. 2019. UAV data as alternative to field sampling to map woody invasive species based on combined Sentinel-1 and Sentinel-2 data. *Remote sensing of environment*, 227: 61–73.
- Kirillov, A.; Mintun, E.; Ravi, N.; et al. 2023. Segment anything. In *IEEE/CVF ICCV*, 4015–4026.
- Lacoste, A.; Lehmann, N.; Rodriguez, P.; et al. 2024. GEO-Bench: Toward Foundation Models for Earth Monitoring. *NeurIPS*, 36.
- Lacoste, A.; Sherwin, E.; Kerner, H.; et al. 2021. Toward foundation models for earth monitoring: Proposal for a climate change benchmark. In *NeurIPS Workshops*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; et al. 2020. Focal Loss for Dense Object Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2): 318–327.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A convnet for the 2020s. In *CVPR*.

- Michałowska, M.; and Rapiński, J. 2021. A review of tree species classification based on airborne LiDAR data and applied classifiers. *Remote Sensing*, 13(3): 353.
- MMSegmentation Contributors. 2020. MMSegmentation: OpenMMLab Semantic Segmentation Toolbox and Benchmark. <https://github.com/open-mmlab/mms Segmentation>.
- Mukkavilli, S.; Civitarese, D.; Schmude, J.; et al. 2023. AI Foundation Models for Weather and Climate: Applications, Design, and Implementation. arXiv:2309.10808.
- Oquab, m.; Darcet, T.; Moutakanni, T.; et al. 2024. DINOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research*.
- Ouaknine, A.; Kattenborn, T.; Laliberté, E.; and Rolnick, D. 2024. OpenForest: A data catalogue for machine learning in forest monitoring. arXiv:2311.00277.
- Papoutsis, I.; Bountos, N.; Zavras, A.; et al. 2023. Benchmarking and scaling of deep learning models for land cover image classification. *ISPRS Journal of Photogrammetry and Remote Sensing*, 195: 250–268.
- Puliti, S.; Pearse, G.; Surový, P.; et al. 2023. FOR-instance: a UAV laser scanning benchmark dataset for semantic and instance segmentation of individual trees. arXiv:2309.01279.
- Qi, C.; Su, H.; Mo, K.; and Guibas, L. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. arXiv:1706.02413.
- Ray, D.; Pankow, J. S.; and Basu, S. 2016. USAT: A unified score-based association test for multiple phenotype-genotype analysis. *Genetic epidemiology*, 40(1): 20–34.
- Reed, C. J.; Gupta, R.; Li, S.; Brockman, S.; Funk, C.; Clipp, B.; Keutzer, K.; Candido, S.; Uyttendaele, M.; and Darrell, T. 2023. Scale-mae: A scale-aware masked autoencoder for multiscale geospatial representation learning. In *ICCV*.
- Reiersen, G.; Dao, D.; Lütjens, B.; et al. 2022. ReforestTree: A dataset for estimating tropical forest carbon stock with deep learning and aerial imagery. In *AAAI Conference on Artificial Intelligence*, volume 36, 12119–12125.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2016. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6): 1137–1149.
- Rolf, E.; Klemmer, K.; Robinson, C.; and Kerner, H. 2024. Mission Critical – Satellite Data is a Distinct Modality in Machine Learning. arXiv:2402.01444.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference*, 234–241. Springer.
- Schmitt, M.; Hughes, L.; Qiu, C.; and Zhu, X. 2019. SEN12MS—A curated dataset of georeferenced multispectral sentinel-1/2 imagery for deep learning and data fusion. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, 4.
- Stewart, A.; Lehmann, N.; Corley, I.; Wang, Y.; Chang, Y.-C.; Ait Ali Braham, N. A.; Sehgal, S.; Robinson, C.; and Banerjee, A. 2024. Ssl4eo-1: Datasets and foundation models for landsat imagery. In *NeurIPS*, volume 36.
- Sumbul, G.; De Wall, A.; Kreuziger, T.; et al. 2021. BigEarthNet-MM: A large-scale, multimodal, multilabel benchmark archive for remote sensing image classification and retrieval [software and data sets]. *IEEE Geoscience and Remote Sensing Magazine*, 9(3): 174–180.
- Sun, X.; Wang, P.; Lu, W.; et al. 2022. RingMo: A remote sensing foundation model with masked image modeling. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–22.
- Sykas, D.; Sdraka, M.; Zografakis, D.; and Papoutsis, I. 2022. A sentinel-2 multiyear, multicountry benchmark dataset for crop classification and segmentation with deep learning. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15: 3323–3339.
- Tang, M.; Cozma, A. L.; Georgiou, K.; and Qi, H. 2023. Cross-Scale MAE: A Tale of Multiscale Exploitation in Remote Sensing. In *NeurIPS*, volume 36.
- Tong, X.-Y.; Xia, G.-S.; and Zhu, X.-X. 2023. Enabling country-scale land cover mapping with meter-resolution satellite imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 196: 178–196.
- Touvron, H.; Cord, M.; Douze, M.; et al. 2021. Training data-efficient image transformers & distillation through attention. In *ICML*, 10347–10357. PMLR.
- Tseng, G.; Zvonkov, I.; Nakalembe, C.; and Kerner, H. 2021. Cropharvest: A global dataset for crop-type classification. In *NeurIPS*.
- Tseng, G.; Zvonkov, I.; Purohit, M.; Rolnick, D.; and Kerner, H. 2023. Lightweight, Pre-trained Transformers for Remote Sensing Timeseries. *ArXiv*.
- Weinstein, B.; Graves, S.; Marconi, S.; et al. 2021. A benchmark dataset for canopy crown detection and delineation in co-registered airborne RGB, LiDAR and hyperspectral imagery from the National Ecological Observation Network. *PLOS Computational Biology*, 17(7): 1–18.
- Xiao, T.; Liu, Y.; Zhou, B.; Jiang, Y.; and Sun, J. 2018. Unified perceptual parsing for scene understanding. In *ECCV*, 418–434.
- Xiong, Z.; Wang, Y.; Zhang, F.; et al. 2024. Neural Plasticity-Inspired Multimodal Foundation Model for Earth Observation. arXiv:2403.15356.
- Yeh, C.; Meng, C.; Wang, S.; et al. 2021. Sustainbench: Benchmarks for monitoring the sustainable development goals with machine learning. In *NeurIPS*.
- Zhao, H.; Jiang, L.; Jia, J.; et al. 2021. Point Transformer. In *ICCV*, 16259–16268.
- Zhu, L.; Wu, J.; Biao, W.; Liao, Y.; and Gu, D. 2023. Spectralmae: Spectral masked autoencoder for hyperspectral remote sensing image reconstruction. *Sensors*, 23(7): 3728.