

# GEO-Bench-2: From Performance to Capability, Rethinking Evaluation in Geospatial AI

Naomi Simumba<sup>†</sup>  
IBM Research Europe

Nils Lehmann<sup>†</sup>  
Technical University Munich

Paolo Fraccaro<sup>†</sup>  
IBM Research Europe  
paolo.fraccaro@ibm.com

Hamed Alemohammad  
Clark University

Geeth De Mel  
IBM Research Europe

Salman Khan  
MBZUAI

Manil Maskey  
NASA Impact

Nicolas Longepe  
ESA Φ-lab

Xiao Xiang Zhu  
Technical University Munich

Hannah Kerner  
Arizona State University

Juan Bernabe-Moreno  
IBM Research Europe

Alexander Lacoste<sup>‡</sup>  
ServiceNow Research  
alexandre.lacoste@servicenow.com

## Abstract

*Geospatial Foundation Models (GeoFMs) are transforming Earth Observation (EO), but evaluation lacks standardized protocols. GEO-Bench-2 addresses this with a comprehensive framework spanning classification, segmentation, regression, object detection, and instance segmentation across 19 permissively-licensed datasets. We introduce "capability" groups to rank models on datasets that share common characteristics (e.g., resolution, bands, temporality). This enables users to identify which models excel in each capability and determine which areas need improvement in future work. To support both fair comparison and methodological innovation, we define a prescriptive yet flexible evaluation protocol. This not only ensures consistency in benchmarking but also facilitates research into model adaptation strategies, a key and open challenge in advancing GeoFMs for downstream tasks.*

*Our experiments show that no single model dominates across all tasks, confirming the specificity of the choices made during architecture design and pretraining. While models pretrained on natural images (ConvNext ImageNet, DINO V3) excel on high-resolution tasks, EO-specific models (TerraMind, Prithvi, and Clay) outperform them on multispectral applications such as agriculture and disaster response. These findings demonstrate that optimal model choice depends on task requirements,*

*data modalities, and constraints. This shows that the goal of a single GeoFM model that performs well across all tasks remains open for future research. GEO-Bench-2 enables informed, reproducible GeoFM evaluation tailored to specific use cases. Code, data, and leaderboard for GEO-Bench-2 are publicly released under a permissive license.*

## 1. Introduction

Deep learning and self-supervision have transformed Earth Observation (EO), enabling analysis of vast, heterogeneous geospatial data at unprecedented scales and accuracy. Geospatial Foundation Models (GeoFMs)—large-scale architectures pre-trained on diverse data—promise to generalize across tasks, sensors, and geographies while reducing reliance on task-specific supervision. However, their development remains hampered by data complexity and the absence of standardized evaluation protocols.

Benchmarking is critical in EO, where models must generalize across spatial, temporal, spectral, and sensor domains. While existing benchmarks like PANGAEA [42], Copernicus-Bench [59], and others [14, 19, 40] have made valuable contributions, they remain limited by narrow modality focus, restrictive licensing, incomplete task coverage, or limited guidelines for reproducible research. GEO-Bench [38] introduced structured evaluation protocols but had key gaps: limited task coverage (no detection or instance segmentation), lack of integration with fine-tuning tools, and inability to assess

<sup>†</sup> Equal Contribution.

<sup>‡</sup> Corresponding authors

Table 1. Model ranking across GEO-Bench-2 capabilities

|                           | Core | Pixel-wise | Classification | Detection | Multi Temporal | < 10m | >=10m | RGB/NIR | Multi-Spectral |
|---------------------------|------|------------|----------------|-----------|----------------|-------|-------|---------|----------------|
| Clay-V1 ViT-B             | 1    | 1          | 9              | 2         | 2              | 6     | 7     | 4       | 3              |
| ConvNext-XLarge ImageNet  | 2    | 3          | 4              | 1         | 8              | 3     | 4     | 2       | 5              |
| DinoV3-ConvNext Large-WEB | 3    | 6          | 1              | 5         | 12             | 2     | 5     | 3       | 4              |
| ConvNext-Large ImageNet   | 4    | 4          | 3              | 3         | 10             | 4     | 2     | 5       | 6              |
| Prithvi-EO-2.0-600M-TL    | 5    | 5          | 8              | 9         | 4              | 8     | 6     | 8       | 2              |
| TerraMind-V1 Large        | 6    | 2          | 6              | 8         | 1              | 7     | 1     | 7       | 1              |
| Satlas-SwinB Sentinel2    | 7    | 7          | 7              | 4         | 6              | 10    | 3     | 9       | 7              |
| DOFA-ViT-L                | 8    | 9          | 5              | 12        | 7              | 5     | 8     | 6       | 8              |
| Prithvi-EO-2.0-300M-TL    | 9    | 10         | 11             | 6         | 5              | 11    | 11    | 11      | 9              |
| DinoV3-ViT-L SAT          | 10   | 8          | 2              | 11        | 9              | 1     | 9     | 1       | 11             |
| TerraMind-V1 Base         | 11   | 12         | 12             | 10        | 3              | 13    | 10    | 12      | 10             |
| Satlas-SwinB Naip         | 12   | 11         | 10             | 7         | 11             | 9     | 12    | 10      | 12             |
| DeCUR-Resnet50            | 13   | 14         | 13             | 14        | 13             | 14    | 13    | 14      | 13             |
| Resnet50-ImageNet         | 14   | 13         | 14             | 13        | 14             | 12    | 14    | 13      | 14             |




Figure 1. Effect of pretraining factors on model performance on the Core capability.

specific model capabilities beyond coarse aggregation.

We present GEO-Bench-2, a comprehensive framework for evaluating GeoFMs across real-world EO applications. Our contributions include: (1) 19 curated datasets with permissive licenses organized into overlapping capability-specific subsets (pixel-wise, detection, multi-temporal, multi-spectral, etc.); (2) geographically balanced splits for scalable, representative evaluation; (3) a prescriptive, but flexible, protocol that ensures fair comparison while allowing for innovation in model fine-tuning and adaptation; (4) full integration with TerraTorch [27], lowering barriers to entry; and (5) an interactive leaderboard on Hugging Face for community engagement. GEO-Bench-2 enables targeted capability assessment and reproducible

comparison, accelerating progress toward general-purpose geospatial intelligence.

## 2. GEO-Bench-2

### 2.1. Dataset Selection and Transformation

There are hundreds of EO labeled datasets [49]. Through an in-depth investigation, we identified 35 datasets with the potential to meet the criteria below. After experimentation and thorough analysis, we selected 19 that we believe enable high-quality GeoFM evaluation (see Table 2 and full details in Appendix 6.1). Our selection criteria were:

**Challenging and Discriminative** Datasets must clearly






Figure 2. Global distribution of samples (500 random locations per dataset shown for visualization). The BioMassters and So2Sat datasets were originally released without geospatial information and are hence not represented in this figure.

distinguish strong GeoFMs from baseline models. We experimentally validated discriminative power across multiple models, discarding datasets with overlapping performance profiles. For example, most well-trained models can obtain above 98% accuracy on Eurosat [31] or saturate in multi-modal datasets like Sen1Floods11 [11].

**Open Licenses** We prioritized permissive licenses to enable academic and industry adoption, avoiding GPL and non-commercial licenses<sup>1</sup>.

**Diversity** The benchmark encompasses a wide range of tasks, modalities, and geographic regions, featuring samples from all seven continents (Figure 2). The comparatively higher representation of Europe stems from initiatives such as INSPIRE and Horizon Europe (see Figure 23 for relative continental coverage)<sup>2</sup>.

If needed, datasets were sub-sampled to reduce computational costs. Datasets were formatted following the TACO convention [9], ensuring FAIR compliance [62] with self-contained, ML-ready samples. If sub-sampling was applied,

<sup>1</sup>For some capabilities (e.g., detection), finding suitable open-license datasets proved challenging.

<sup>2</sup>To enhance global balance, we encourage the community to release open-source datasets from underrepresented regions, with an effort to avoid non-commercial restrictions. Additionally, to promote the creation of initiatives similar to INSPIRE and Horizon Europe worldwide.

a maximum of 20,000 was selected for classification and 4,000 for pixel-wise and detection tasks.

## 2.2. Capability Groups

Single-aggregate metrics obscure GeoFM’s strengths and constrain research directions. GEO-Bench-2 instead evaluates nine overlapping **capabilities** designed to highlight specific challenges in GeoFM development. These include architectural challenges, such as object detection and segmentation, as well as input-related challenges like sensor fusion and temporal understanding. (Table 3).

**Core** *Core* provides a balanced subset across all capabilities with the aim of reducing the compute needed for evaluating on the benchmark while evaluating on the more discriminative datasets.

**ML Task Type** *Classification* includes single- and multi-label image classification tasks. *Pixel-wise predictions* encompass semantic segmentation and regression, assigning classes or continuous values per pixel. *Detection* evaluates object detection and instance segmentation with bounding boxes and masks. Since the detection datasets can be challenging to work with from a computational and methodological perspective, they are included only in 2 subsets to make other capabilities more accessible.

**Temporality** *Multi-temporal* datasets include multiple timestamps for change detection and temporal modeling.

**Resolution** To study over-reliance on data coming from Sentinel or Landsat, we also create subsets to study performances under and above 10 m GSD.

**Spectral Diversity** Spectrally, *RGB/NIR* includes basic visible and near-infrared data, while *Multi-spectral-dependent* requires optical bands beyond RGB for task completion (validated via ablation, see Appendix 6.3.4). We do not explicitly define a SAR capability, but five datasets (KuroSiwo, PASTIS-R, CaFFe, BEN V2, CloudSEN12) could be used to study this axis independently.

## 3. Evaluation Protocol

The benchmark’s main objective is to advance GeoFMs’ development by measuring their performance on downstream tasks. However, since adapting a base model to a specific dataset remains an active research area with significant variation across models and task types, our second objective is to support this exploration. We therefore allow users flexibility in how they adapt GeoFMs to downstream tasks within our framework. Our experiments establish a **baseline adaptation protocol** and provide initial rankings for several GeoFMs, but we encourage users to submit their own leaderboard entries demonstrating improved adaptation strategies.

To ensure fair comparisons across GeoFMs, we require users to follow [mandatory guidelines highlighted in blue](#)

| Task                  | Dataset                  | Domain           | Modalities      | GSD  | Train/Val/Test  | Classes        | License             |
|-----------------------|--------------------------|------------------|-----------------|------|-----------------|----------------|---------------------|
| Classification        | BEN V2 [16]              | Land cover       | S1+S2           | 10m  | 20000/4000/4000 | 19             | CDLA-Permissive-1.0 |
|                       | TreeSatAI [5]            | Tree Species     | S2 TS           | 10m  | 20000/4000/4000 | 13             | CC-BY-4.0           |
|                       | So2sat [38, 65]          | Climate Zones    | S2              | 10m  | 19992/986/986   | 17             | CC-BY-4.0           |
|                       | Forestnet [33, 38]       | Tree Species     | L8              | 15m  | 6464/989/989    | 12             | CC-BY-4.0           |
| Pixel Regression      | BioMassters [44]         | Biomass estim.   | S1+S2 TS        | 10m  | 4000/1000/2000  | $\mathbb{R}_+$ | CC-BY-4.0           |
| Semantic Segmentation | CaFFe [28]               | Glacier zones    | S1 SAR          | 10m  | 4000/1000/2000  | 4              | CC-BY-4.0           |
|                       | CloudSen12 [8]           | Cloud/shadow     | S1+S2           | 10m  | 4000/1000/2000  | 4              | CC0                 |
|                       | NASA Burn Scars [52]     | Burn scars       | HLS             | 30m  | 524/160/120     | 2              | CC-BY-4.0           |
|                       | Dynamic EarthNet [53]    | LULC (temporal)  | Planet          | 3m   | 4000/1000/2000  | 7              | CC-BY-4.0           |
|                       | FLAIR 2 [22]             | LULC             | Aerial RGBN+DEM | 0.2m | 4000/1000/2000  | 13             | Open License 2.0    |
|                       | FTW [36]                 | Field boundaries | S2              | 10m  | 4000/1000/2000  | 2              | CC-BY-SA            |
|                       | KuroSiwo [13]            | Flood extent     | S1+DEM+Slope    | 10m  | 4000/1000/2000  | 4              | MIT                 |
|                       | PASTIS (R) [48]          | Crop type map.   | S1+S2           | 10m  | 1455/482/496    | 19             | CC-BY-4.0           |
|                       | SpaceNet2 [56]           | Building         | Worldview       | 0.3m | 4000/1000/2000  | 2              | CC-BY-SA-4.0        |
|                       | SpaceNet7 [56]           | Building         | Planet          | 3m   | 3888/652/1152   | 2              | CC-BY-SA-4.0        |
|                       | EverWatch [23]           | Bird species     | Aerial RGB      | 0.1m | 4429/500/196    | 9              | CC0                 |
| Object Detection      | m-nzcattle [2]           | Cattle           | Aerial RGB      | 0.1m | 524/66/65       | 2              | CC-BY-4.0           |
|                       |                          |                  |                 |      |                 |                |                     |
| Instance Segmentation | PASTIS (R) panoptic [48] | Crop type map.   | S1+S2           | 10m  | 1455/482/496    | 19             | CC-BY-4.0           |
|                       | Substations [35]         | Substations      | S2              | 10m  | 4000/500/500    | 2              | CC-BY-4.0           |

Table 2. Anonymized-Bench dataset overview grouped by task category. S1: Sentinel-1; S2: Sentinel-2; DEM: Digital Elevation Model; TS: Time series; LULC: Land Use Land Cover.

| Dataset             | Core | Pixel wise | Classi-<br>fication | Detection | Multi<br>Temporal | < 10m<br>Res | ≥ 10m<br>Res | RGB/<br>NIR | Multi<br>Spectral |
|---------------------|------|------------|---------------------|-----------|-------------------|--------------|--------------|-------------|-------------------|
| BEN V2              | ✓    |            | ✓                   |           |                   |              | ✓            |             | ✓                 |
| TreeSatAI           | ✓    |            | ✓                   |           |                   | ✓            |              | ✓           |                   |
| So2Sat              |      |            | ✓                   |           |                   |              | ✓            |             | ✓                 |
| ForestNet           |      |            | ✓                   |           |                   |              | ✓            |             |                   |
| BioMassters         | ✓    | ✓          |                     |           | ✓                 |              | ✓            |             | ✓                 |
| CaFFe               |      | ✓          |                     |           |                   |              | ✓            |             |                   |
| CloudSEN12          | ✓    | ✓          |                     |           |                   |              | ✓            |             | ✓                 |
| NASA Burn Scars     | ✓    | ✓          |                     |           |                   |              | ✓            |             | ✓                 |
| Dynamic Earth Net   |      | ✓          |                     |           | ✓                 | ✓            |              | ✓           |                   |
| FLAIR 2             | ✓    | ✓          |                     |           |                   | ✓            |              | ✓           |                   |
| FTW                 | ✓    | ✓          |                     |           |                   |              | ✓            | ✓           | ✓                 |
| KuroSiwo            | ✓    | ✓          |                     |           | ✓                 |              | ✓            |             |                   |
| PASTIS              | ✓    | ✓          |                     |           | ✓                 |              | ✓            |             | ✓                 |
| SpaceNet 2          |      | ✓          |                     |           |                   | ✓            |              |             | ✓                 |
| SpaceNet 7          | ✓    | ✓          |                     |           |                   | ✓            |              | ✓           |                   |
| EverWatch           | ✓    |            |                     | ✓         |                   |              |              |             |                   |
| NZCattle            |      |            |                     | ✓         |                   |              |              |             |                   |
| PASTIS (R) panoptic |      |            |                     | ✓         |                   |              |              |             |                   |
| Substations         | ✓    |            |                     | ✓         |                   |              |              |             |                   |

Table 3. Overview of dataset capabilities.

throughout this section. The remaining details describe our baseline adaptation protocol. As a first guideline, **We request users to document precisely their adaptation protocol, following the provided template.** Also, to improve comparison, **if a new GeoFM is evaluated, we suggest users to provide at least one entry according to the baseline protocol.**

**Why fine-tuning?** It is possible to use the benchmark with frozen backbones, and we provide a second leaderboard for users who would like to operate in this fashion. However,

in Section 6.3.2, we demonstrate an important performance degradation and significant change in ranking of models compared to a fine-tuning approach. For this reason, we recommend using the benchmark with a fine-tuning and hyperparameter selection approach.

### 3.1. Hyperparameter Optimization

**Search Space:** The search space of hyperparameters should be mostly uniform across the benchmark to prevent users

from crafting a search space for each dataset. Users may design multiple search spaces for different task types (e.g., separate spaces for classification vs. segmentation), [but must document these search spaces and the rationale behind them](#). Leaderboard submissions will be reviewed and may be flagged as *over-engineered* if search spaces are excessively customized. As a rule of thumb, we recommend a search space cover at least 4 datasets to ensure statistical validity of the HPO process.

**Search Budget:** To balance exploration with compute fairness, we impose a [maximum search budget of 16 trials per dataset](#), and the selection of the best configuration can only be done on the provided validation set. We allow Bayesian optimizers such as Optuna.

**Repeated Test:** Once the best configuration is obtained on the validation set, it is requested to [repeat the fine-tuning of the best configuration 5 times](#) and evaluate on the test set with the required metric (see section 3.6) to account for training stochasticity.

**Baseline protocol:** Detailed configuration of our experiments is provided in Appendix 6.2. At a glance, we use Optuna [6] in TerraTorch Iterate to automate the process, and we encourage users to do so when possible. The optimization is conducted over a fixed hyperparameter space across all models and all datasets, tuning Learning rate and Batch size while keeping other hyperparameters constant. Each HPO and repeated experiment iteration lasted for 50 epochs.

### 3.2. Pre-Processing and Data Augmentation

[Pre-processing or data augmentations cannot be informed by the test set, and the methodology needs to be uniform across all datasets.](#)

**Baseline protocol:** We apply per-band Z-score normalization to each model<sup>3</sup>, and parameters are estimated over the training split. The same normalization strategy is used for target data in regression datasets. During training, standard augmentations include horizontal and vertical flipping. For pixel-wise tasks, random cropping is applied to images larger than  $224 \times 224$  pixels during training, while tiled inference is utilized for validation and testing to process the full image extent.

### 3.3. Base Model Adaptation

Users can employ different methodologies to adapt a base model for different categories of tasks, but [a methodology should be generic and valid for at least 4 different datasets](#).

**Baseline protocol:**

**Classification:** We used a single linear layer with softmax on the encoder output.

---

<sup>3</sup>Note, this might penalize some base models expecting a different input distribution. However, we believe that the fine-tuning procedure mitigates this potential discrepancy.

**Pixel-wise Tasks:** (Segmentation and Regression) We used a UNet decoder for both semantic segmentation and regression tasks, where we fed equally spaced features from the encoder’s output into the UNet [43]. For transformer-based models, *LearnedFeatureInterpolation*, available in TerraTorch, is applied to hierarchically structure the encoder output before passing it to the UNet [47]. This is a technique demonstrated to work well for ViTs compared to other choices (i.e., UPerNet) [43].

**Object Detection and Instance Segmentation:** We used Faster R-CNN [26] and Mask R-CNN [30], respectively. For both tasks, a *FeaturePyramidNetwork* [41] was applied, following the initial transformation of ViT features as described for pixel-wise tasks.

### 3.4. Multi-Spectral Bands, SAR and Multi-Modal Datasets

We leave the choice to users on how to process the available data sources in each dataset. Below, a description of the approach we used.

**Baseline Protocol:**

**Multi-Spectral Bands:** All bands contained in a given dataset are utilized, provided they are compatible with the model. Where an exact match did not exist, model bands were matched to the dataset bands with the closest wavelength (i.e. matching S2 bands to WorldView).

**Synthetic Aperture Radar:** In case a model could not handle Synthetic Aperture Radar (SAR) natively, VV and VH polarization bands were loaded as the model’s RGB channels in the order VV, VH, and VV.

**Multi-Modal Datasets:** It remains unclear how much S1 helps when S2 is available [42]. To investigate this, we ran an ablation (Appendix 6.3.6) on three multi-modal models (Terramind, DOFA, Clay), comparing S2-only performance to S1+S2 for multi-modal datasets. For the main results, we report the best outcome per model–dataset pair, which led to using both modalities only for Terramind on BEN V2 and BioMassters (Appendix 6.3.6); all other models use S2 only.

### 3.5. Multi-Temporality

**Baseline Protocol:** For datasets with multi-temporal inputs, each timestamp was passed through the encoder separately. The resulting encoder outputs were then averaged along the embedding dimensions before being passed to the decoder.

### 3.6. Evaluation Metrics

[All submissions should report results using the following standardized metrics:](#)

- Semantic Segmentation: Multiclass Jaccard Index
- Single label Classification: Accuracy
- Multilabel classification: F1-score
- Pixel Wise Regression: Root Mean Squared Error (RMSE)

- Instance segmentation and Object detection: Mean Average Precision (mAP)

### 3.7. Aggregated performance

To obtain a valid aggregation and account for uncertainty introduced by random seeds, we follow the methodology introduced in [38]. It is recommended to use the leaderboard or the provided code for computing these scores based on raw results to minimize the chance of mistakes.

**Renormalization:** To report scores for a capability, we aggregate results across multiple datasets. Simple averaging can distort comparisons, as strong performance on one dataset may overshadow improvements elsewhere [18]. Therefore, we normalize each dataset’s scores to [0,1] using a linear transformation with the worst and best result for a given dataset as the reference points<sup>4</sup>. These normalization factors, computed from our fixed set of models, are published in our Anonymous-Bench-Repo along with tools for calculating final aggregated scores. All users should reuse these factors for consistency.

**Bootstrap:** To obtain the uncertainty over the final aggregated score, we use a stratified bootstrap over the collection of 5 normalized repeats for each dataset with 100 resamples.

**IQM aggregation:** On each normalized bootstrapped iteration, results from each dataset in a capability are averaged across using the interquartile mean (IQM) to discard outliers and obtain a more statistically stable mean [3].

**Final Aggregation:** Results across the 100 normalized IQM bootstrap iterations are averaged to obtain a final score and standard deviation.

### 3.8. Leaderboard

To foster a community effort, we provide a leaderboard and encourage a wide range of users to submit their experiments, the good ones and the bad ones. We also encourage users to reproduce existing experiments. To lower the barrier of entry, users can submit an entry for a single capability without having to evaluate on all datasets of the benchmark. Finally, we reserve the right to reject submissions based on the following:

- Does not follow the required guidelines
- Cannot be reproduced, e.g., closed source model
- Per-dataset over-engineering
- Any other reasons that would undermine the main objectives of this benchmark

## 4. Results

We present normalized bootstrapped IQM results obtained by applying the GEO-Bench-2 protocol across all benchmarking capabilities in Figure 3. Table 1 provides model rankings sorted by Core capability performance. Raw performance metrics for individual datasets are available in Appendix 6.3.1.

<sup>4</sup>Similar to z-score normalization, but [0,1] improves interpretability.

### 4.1. Models used in our Experiments

To explore the value of our proposed benchmark, we have chosen models that reflect a diverse range of paradigms, model sizes, and pretraining datasets. They also include foundational general-purpose computer vision models like the latest DINOv3 [50] to explore their generalizability to EO tasks. The characteristics of the benchmarked models are detailed in Table 4. All experiments in the main text are conducted with end-to-end finetuning, meaning the pretrained backbone was unfrozen during the training process.

### 4.2. General Trends

Despite some models that perform more consistently across a number of capabilities (i.e., Clay-V1 ViT-B, or ConvNeXt architectures), our results show that there is no model that dominates across all datasets. On the contrary, looking at Figure 3, we can see that different models have a clear lead in different capabilities. Some examples include TerraMind-V1-Large for its multi-spectral and Multi-Temporal capabilities, or DINOv3-ViT-L-SAT for the < 10 m GSD and RGB/NIR options. Although not ideal, given the goal of GeoFMs to perform well across all possible tasks, this is not surprising. In fact, based on their pretraining task, pretraining data distribution, target data, architecture, and size, some models will be more suitable to tackle specific use cases.

### 4.3. Impact of Architecture, Size, and Pretraining Dataset

Figure 1 shows how architecture, model size, and pretraining dataset size relate to Core capability performance. Larger models consistently outperform their smaller counterparts, with ResNet-50 models showing notably poor performance regardless of pretraining data. However, Clay-V1 ViT-B achieves top Core performance with only 86M parameters. This is likely due to its diverse pretraining on 70M heterogeneous EO samples (1-30m GSD) and smaller patch size of 8, which captures finer details at 4× computational cost compared to standard ViT models with patch size 16. Notably, ConvNeXt architectures pretrained on natural images adapt effectively to EO tasks through full finetuning, independent of model size or pretraining domain.

### 4.4. Importance of Multi-Spectral Bands

DINOv3-ViT-L-SAT and the ConvNeXt models are top performers in the RGB/NIR or the below 10 m GSD capability, which mostly includes high-resolution RGB data to perform the task. Looking at Figure 3, we can see that despite TerraMind-V1-Large, Prithvi-EO-2.0-600-TL, and Clay-V1 ViT-B taking the podium in the Multi-Spectral-Dependent capability, the ConvNeXt models are still close in normalized performance. This is somewhat surprising because this capability only includes datasets that reported a statistically significant drop in performance once multi-spectral bands were removed to complete the task (see Appendix 6.3.4). However, when we look

| Model                          | Type      | # Backbone Params | Learning Technique | Data                                                  | Res    | N     | T | License    |
|--------------------------------|-----------|-------------------|--------------------|-------------------------------------------------------|--------|-------|---|------------|
| Resnet50-ImageNet              | ResNet-50 | 25M               | Supervised         | ImageNet-22k                                          | NA     | 14M   | 1 | Apache 2.0 |
| ConvNext-Large-ImageNet [61]   | ConvNext  | 230M              | Supervised         | ImageNet-22k                                          | NA     | 14M   | 1 | Apache 2.0 |
| ConvNext-XLarge-ImageNet [61]  | ConvNext  | 390M              | Supervised         | ImageNet-22k                                          | NA     | 14M   | 1 | Apache 2.0 |
| DINOv3-ViT-L-SAT [50]          | ViT       | 300M              | Distillation       | Maxar RGB                                             | 0.6 m  | 493M  | 1 | DINO V3    |
| DINOv3-ConvNext-Large-WEB [50] | ConvNext  | 230M              | Distillation       | LVD-1689M                                             | NA     | 1689M | 1 | DINO V3    |
| Resnet50-DeCUR [58]            | ResNet-50 | 25M               | Contrastive        | Sentinel-2                                            | 10 m   | 1M    | 1 | Apache 2.0 |
| DOFA-ViT-300M [63]             | ViT       | 300M              | MAE                | Sentinel-1 and -2, EnMap, Gaofen, Landsat             | 1-30 m | 8M    | 1 | CC-BY-4.0  |
| Clay-V1 ViT-B [1]              | ViT       | 86M               | MAE                | Landsat 8 and 9, Sentinel-1 and -2, NAIP, LINZ, MODIS | 1-30 m | 70M   | 1 | Apache 2.0 |
| Satlas-SwinB-Sentinel2 [10]    | Swin      | 88M               | Supervised         | Sentinel-2                                            | 10 m   | NA    | 1 | ODC-BY     |
| Satlas-NAIP [10]               | Swin      | 88M               | Supervised         | NAIP                                                  | 1 m    | NA    | 1 | ODC-BY     |
| Prithvi-EO-V2-300M-TL [52]     | ViT       | 300M              | MAE                | HLS                                                   | 30 m   | 4.2M  | 4 | Apache 2.0 |
| Prithvi-EO-V2-600M-TL [52]     | ViT       | 600M              | MAE                | HLS                                                   | 30 m   | 4.2M  | 4 | Apache 2.0 |
| TerraMind-V1-Base [34]         | ViT       | 86M               | Correlation        | Sentinel-1 and -2, LULC, DEM, NDVI                    | 10 m   | 9M    | 1 | Apache 2.0 |
| TerraMind-V1-Large [34]        | ViT       | 300M              | Correlation        | Sentinel-1 and -2, LULC, DEM, NDVI                    | 10 m   | 9M    | 1 | Apache 2.0 |

The pretraining dataset size ( $N$ ) is estimated where not clearly reported.

Checkpoints source: Torchgeo (DOFA, DeCUR, and Satlas); TerraTorch (Prithvi and TerraMind); Timm (Resnet50-ImageNet and ConvNext-ImageNet); Meta (DINO V3).

Table 4. Characteristics of the models compared: backbone type (Type), number of parameters (# Param.), pretraining technique (Technique), pretraining data (Data), spatial resolution (Res.), number of samples ( $N$ ), and number of timestamps per sample ( $T$ ). Characteristics for Satlas refer to the model pretrained on Sentinel-2 downstream tasks. The pretraining sample size ( $N$ ) is estimated when not explicitly reported.




Figure 3. Normalized bootstrapped IQM performance by capability. Models are ordered based on core capability

more in detail at the individual datasets’ raw results in Figure 24, we see interesting insights emerging from datasets like NASA Burn Scars and PASTIS crop classification. Particularly, here there is a more marked difference (e.g. up to 10%) between the models using all multi-spectral bands available and the

ones using only RGB data. This confirms the importance of EO-specific models for particular applications in which going beyond RGB is the only way to achieve maximum performance.



#### 4.5. Model Performance Across Task Types

Looking at the different task-focused capabilities, some architectural choices seem to bring clear advantages. For example, Clay-V1-ViT-B tops the Pixel-wise capability, which is consistent with its architecture. As already mentioned, it is the only GeoFM that adopts a smaller patch size, granting it an inherent advantage in retaining fine-grained spatial details crucial for these tasks. In the classification and detection capabilities, the DINOv3 and Imagenet ConvNeXt models achieve the highest rankings. In the detection tasks, however, most GeoFMs based on a Vision Transformer architecture show comparatively lower performance. This may be related to the choice of frameworks used for this task (i.e., Faster R-CNN and Mask R-CNN), which are known to favor models with convolutional inductive biases.

#### 4.6. Ablations

Beyond the RGB versus Multi-Spectral ablation, we ran additional studies to assess the impact of some of our choices. We summarize these by reporting how rankings change, measured via Kendall tau distance—the percentage of model pairs with altered order (Figure 4). Full details on all ablations are in Appendices 6.3.2–6.3.9.

**Frozen versus Fine-tuned Backbones.** We first investigated whether freezing the backbone during training would cause performance degradation and ranking changes relative to full fine-tuning. Contrary to findings in [42], we observed that full fine-tuning consistently outperformed the frozen backbone approach, though the magnitude of performance drop varied across models. This resulted in ranking changes for over 20% of model pairs, as shown in Figure 4.

**Decoder Complexity.** Next, we examined whether a simpler linear decoder could match the performance of a more complex UNet architecture while offering computational advantages. As detailed in Appendix 6.3.3, the linear decoder consistently underperformed across most datasets, leading to notable ranking changes.

**Multi-Temporal Processing.** For datasets in the Multi-Temporal capability, we confirmed that leveraging multiple timestamps is crucial. Using only a single timestamp resulted in performance drops (Appendix 6.3.7) and produced the largest ranking perturbations among all ablations (Figure 4).

### 5. Discussion and Limitations

Foundation models for EO are advancing rapidly [66], with a strong appeal for domain-specific applications. However, the field has yet to experience a paradigm shift comparable to that of generative modeling in text or computer vision. While our analysis confirms that larger models tend to perform better, EO has not exhibited the scaling laws observed in LLMs or image generation [32, 39].

The LLM community benefits from established benchmarks




Figure 4. Normalized Kendall tau distance showing the fraction of pairs where ranking is changing between 2 methods. Uncertainty was calculated through bootstrapping 2,000 times.

and centralized infrastructure for tracking progress [60]. With GEO-Bench-2, we provide such a framework for the EO domain. A key advantage of the GEO-Bench-2 leaderboard is enabling users to rank models across various capabilities and settings, allowing the community to identify models suited to specific use cases and resource constraints. Our prescriptive protocol is centered on computational budget, repeated experiments, and generalizability, providing consistency while allowing flexibility in how the community adapts foundation models for different capabilities.

Recent work has shown that models pretrained on natural images remain competitive on EO tasks [17], and well-tuned supervised models can match specialized FMs [64]. PANGAEA similarly found that most GeoFMs struggled to beat simpler models under frozen encoder settings [42], with TerraMind being the first exception [34].

Our experiments tell a more nuanced story. Under full fine-tuning settings, GeoFMs consistently outperformed simpler models (e.g., ResNet-50) but not larger models trained on natural images like ConvNeXt, particularly for high-resolution RGB-only tasks. However, GeoFMs demonstrated clear advantages when multi-spectral information is critical. This is precisely where GEO-Bench-2 can guide the field forward. By systematically evaluating models across diverse spectral bands, spatial resolutions, and temporal scales, we can identify where EO-specific architectures and pretraining truly matter. Our results align with DINOv3 findings [50], where even the 7B RGB-only model could not surpass the 10× smaller Prithvi-EO-2.0 in crop classification. This highlights the value of multi-spectral capabilities that GEO-Bench-2 explicitly tests.

Rather than discouraging, these findings reveal opportunities. GEO-Bench-2 pinpoints where current GeoFMs underperform and provides a testbed for developing models that leverage EO’s unique characteristics: multi-spectral data, temporal dynamics, and global coverage. These capabilities extend beyond what natural image pretraining can offer. This can guide the community toward further developing pretraining strategies optimized



for multi-spectral and temporal data. For instance, our analysis revealed competitive results from ConvNeXt architecture, rarely used in EO, suggesting unexplored architectural directions. By systematically testing models across GEO-Bench-2’s diverse tasks, the community can discover which innovations translate to real gains in the tasks that matter most: disaster response, environmental monitoring, and agriculture at scale.

## 5.1. Limitations

As with any benchmark, some arbitrary choices were necessary, but we believe they are reasonable and do not diminish the framework’s value. Our benchmark includes many leading models, though others remain untested, and some baseline decisions may not suit all model–dataset pairs. We encourage community contributions via the leaderboard. A SAR-specific capability is missing despite five SAR-inclusive datasets—future work should address this gap. Following [46], benchmarks should evolve toward real-world conditions with task-specific metrics and deployment-aware data, requiring cross-sector collaboration. Despite efforts toward global diversity, dataset coverage remains biased toward Europe and North America. Finally, predictive uncertainty is not assessed, though it is critical for real-world applications and should be prioritized.

## References

- [1] Clay foundation model. <https://clay-foundation.github.io/model/#clay-foundation-model>, 2025. Accessed on January 09, 2025. 7
- [2] Diab Abuaiadah and Alexander Switzer. Remote sensing dataset for detecting cows from high resolution aerial images. 2022. 4
- [3] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in neural information processing systems*, 34:29304–29320, 2021. 6
- [4] Steve Ahlswede, Christian Schulz, Christiano Gava, Patrick Helber, Benjamin Bischke, Michael Förster, Florencia Arias, Jörn Hees, Begüm Demir, and Birgit Kleinschmit. Treesatai benchmark archive: A multi-sensor, multi-label dataset for tree species classification in remote sensing. *Earth System Science Data Discussions*, 2022:1–22, 2022. 15
- [5] Steve Ahlswede, Christian Schulz, Christiano Gava, Patrick Helber, Benjamin Bischke, Michael Förster, Florencia Arias, Jörn Hees, Begüm Demir, and Birgit Kleinschmit. Treesatai benchmark archive: A multi-sensor, multi-label dataset for tree species classification in remote sensing. *Earth System Science Data Discussions*, 2022:1–22, 2022. 4
- [6] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631, 2019. 5, 17
- [7] Guillaume Astruc, Nicolas Gonthier, Clement Mallet, and Loic Landrieu. Omnisat: Self-supervised modality fusion for earth observation. In *European Conference on Computer Vision*, pages 409–427. Springer, 2024. 15
- [8] Cesar Aybar, Lesly Bautista, David Montero, Julio Contreras, Daryl Ayala, Fernando Prudencio, Jhomira Loja, Luis Ysuhualas, Fernando Herrera, Karen Gonzales, et al. Cloudsen12+: The largest dataset of expert-labeled pixels for cloud and cloud shadow detection in sentinel-2. *Data in Brief*, 56:110852, 2024. 4, 12
- [9] Cesar Aybar, Julio Contreras, Chen Ma, Oscar J Pellicer-Valero, Gonzalo Mateo-García, Luis Gómez-Chova, Gustau Camps-Valls, Nils Lehmann, Mikolaj Czerkawski, David Montero, et al. The missing piece: Standardising for ai-ready earth observation datasets. In *TerraBytes-ICML 2025 workshop*, 2025. 3
- [10] Favien Bastani, Piper Wolters, Ritwik Gupta, Joe Ferdinando, and Aniruddha Kembhavi. Satlaspretrain: A large-scale dataset for remote sensing image understanding, 2023. 7
- [11] Derrick Bonafilia, Beth Tellman, Tyler Anderson, and Erica Issenberg. Sen1floods11: A georeferenced dataset to train and test deep learning flood algorithms for sentinel-1. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 210–211, 2020. 3
- [12] Nikolaos Ioannis Bountos, Maria Sdraka, Angelos Zavras, Andreas Karavias, Ilektra Karasante, Themistocles Herekakis, Angeliki Thanasou, Dimitrios Michail, and Ioannis Papoutsis. Kuro siwo: 33 billion m<sup>2</sup> under the water. a global multi-temporal satellite dataset for rapid flood mapping. *Advances in Neural Information Processing Systems*, 37:38105–38121, 2024. 14
- [13] Nikolaos Ioannis Bountos, Maria Sdraka, Angelos Zavras, Andreas Karavias, Ilektra Karasante, Themistocles Herekakis, Angeliki Thanasou, Dimitrios Michail, and Ioannis Papoutsis. Kuro siwo: 33 billion m<sup>2</sup> under the water. a global multi-temporal satellite dataset for rapid flood mapping. In *Advances in Neural Information Processing Systems*, pages 38105–38121. Curran Associates, Inc., 2024. 4
- [14] Nikolaos Ioannis Bountos, Arthur Ouaknine, Ioannis Papoutsis, and David Rolnick. Fomo: Multi-modal, multi-scale and multi-task remote sensing foundation models for forest monitoring. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 27858–27868, 2025. 1
- [15] Kai Norman Clasen, Leonard Hackel, Tom Burgert, Gencer Sumbul, Begüm Demir, and Volker Markl. reben: Refined bigearthnet dataset for remote sensing image analysis. *arXiv preprint arXiv:2407.03653*, 2024. 12
- [16] Kai Norman Clasen, Leonard Hackel, Tom Burgert, Gencer Sumbul, Begüm Demir, and Volker Markl. reben: Refined bigearthnet dataset for remote sensing image analysis. *arXiv preprint arXiv:2407.03653*, 2024. 4
- [17] Isaac Corley, Caleb Robinson, Rahul Dodhia, Juan M Lavista Ferres, and Peyman Najafirad. Revisiting pre-trained remote sensing model benchmarks: resizing and normalization matters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3162–3172, 2024. 8
- [18] Janez Demšar. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine learning research*, 7(Jan):1–30, 2006. 6
- [19] Casper Fibaek, Luke Camilleri, Andreas Luyts, Nikolaos Dionelis, and Bertrand Le Saux. Phileo bench: Evaluating geo-spatial foundation models, 2024. URL <https://arxiv.org/abs/2401.04464>. 1
- [20] Anatol Garioud, Stéphane Peillet, Eva Bookjans, Sébastien Giordano, and Boris Wattrelos. Flair# 1: semantic segmentation

- and domain adaptation dataset. *arXiv preprint arXiv:2211.12979*, 2022. 14
- [21] Anatol Garioud, Apolline De Wit, Marc Poupée, Marion Valette, Sébastien Giordano, and Boris Wattrélos. Flair# 2: textural and temporal information for semantic segmentation from multi-source optical imagery. *arXiv preprint arXiv:2305.14467*, 2023. 14
- [22] Anatol Garioud, Nicolas Gonthier, Loic Landrieu, Apolline De Wit, Marion Valette, Marc Poupée, Sébastien Giordano, and Boris Wattrélos. Flair: a country-scale land cover semantic segmentation dataset from multi-source optical imagery. In *Advances in Neural Information Processing Systems (NeurIPS) 2023*, 2023. 4
- [23] Lindsey Garner, Ben Weinstein, Michael Rickershauser, Melissa Baldino, Holly Coates, Mary Commins, Tracey Faber, Jonah Gula, Shayly Van Ert, Peter Frederick, Ethan White, and S.K. Morgan Ernest. Everwatch benchmark: training and evaluation data for detection and species classification of everglades wading birds from airborne imagery, 2024. 4
- [24] L Garner, B Weinstein, M Rickershauser, M Baldino, H Coates, M Commins, T Faber, J Gula, S Van Ert, E White, et al. Everwatch benchmark: training and evaluation data for detection and species classification of everglades wading birds from airborne imagery. *Zenodo*, 2024. 13
- [25] Vivien Sainte Fare Garnot, Loic Landrieu, and Nesrine Chehata. Multi-modal temporal attention models for crop mapping from satellite time series. *ISPRS Journal of Photogrammetry and Remote Sensing*, 187:294–305, 2022. 14
- [26] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 5
- [27] Carlos Gomes, Benedikt Blumenstiel, Joao Lucas de Sousa Almeida, Pedro Henrique de Oliveira, Paolo Fraccaro, Francesc Marti Escofet, Daniela Szwarcman, Naomi Simumba, Romeo Kienzler, and Bianca Zadrozny. Terratorch: The geospatial foundation models toolkit. *arXiv preprint arXiv:2503.20563*, 2025. 2
- [28] Nora Gourmelon, Thorsten Seehaus, Matthias Braun, Andreas Maier, and Vincent Christlein. Calving fronts and where to find them: A benchmark dataset and methodology for automatic glacier calving front extraction from sar imagery. *Earth System Science Data Discussions*, 2022:1–37, 2022. 4
- [29] Nora Gourmelon, Thorsten Seehaus, Matthias Braun, Andreas Maier, and Vincent Christlein. Calving fronts and where to find them: A benchmark dataset and methodology for automatic glacier calving front extraction from sar imagery. *Earth System Science Data Discussions*, 2022:1–37, 2022. 12
- [30] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 5
- [31] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. 3
- [32] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models (2022). *arXiv preprint arXiv:2203.15556*, 2022. 8
- [33] Jeremy Irvin, Hao Sheng, Neel Ramachandran, Sonja Johnson-Yu, Sharon Zhou, Kyle Story, Rose Rustowicz, Cooper Elsworth, Kemen Austin, and Andrew Y Ng. Forestnet: Classifying drivers of deforestation in indonesia using deep learning on satellite imagery. *arXiv preprint arXiv:2011.05479*, 2020. 4
- [34] Johannes Jakubik, Felix Yang, Benedikt Blumenstiel, Erik Scheurer, Rocco Sedona, Stefano Maurogiovanni, Jente Bosmans, Nikolaos Dionelis, Valerio Marsocci, Niklas Kopp, Rahul Ramachandran, Paolo Fraccaro, Thomas Brunschweiler, Gabriele Cavallaro, Juan Bernabe-Moreno, and Nicolas Longépé. Terramind: Large-scale generative multimodality for earth observation, 2025. 7, 8
- [35] Kartik Jindgar and Grace W Lindsay. Improving satellite imagery segmentation using multiple sentinel-2 revisits. *arXiv preprint arXiv:2409.17363*, 2024. 4, 16
- [36] Hannah Kerner, Snehal Chaudhari, Aninda Ghosh, Caleb Robinson, Adeel Ahmad, Eddie Choi, Nathan Jacobs, Chris Holmes, Matthias Mohr, Rahul Dodhia, et al. Fields of the world: A machine learning benchmark dataset for global agricultural field boundary segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 28151–28159, 2025. 4
- [37] Hannah Kerner, Snehal Chaudhari, Aninda Ghosh, Caleb Robinson, Adeel Ahmad, Eddie Choi, Nathan Jacobs, Chris Holmes, Matthias Mohr, Rahul Dodhia, et al. Fields of the world: A machine learning benchmark dataset for global agricultural field boundary segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 28151–28159, 2025. 13
- [38] Alexandre Lacoste, Nils Lehmann, Pau Rodriguez, Evan Sherwin, Hannah Kerner, Björn Lütjens, Jeremy Irvin, David Dao, Hamed Alemohammad, Alexandre Drouin, et al. Geo-bench: Toward foundation models for earth monitoring. *Advances in Neural Information Processing Systems*, 36:51080–51093, 2023. 1, 4, 6, 15, 16, 17
- [39] Hao Li, Yang Zou, Ying Wang, Orchid Majumder, Yusheng Xie, R Manmatha, Ashwin Swaminathan, Zhuowen Tu, Stefano Ermon, and Stefano Soatto. On the scalability of diffusion-based text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9400–9409, 2024. 8
- [40] Xiang Li, Yong Tao, Siyuan Zhang, Siwei Liu, Zhitong Xiong, Chunbo Luo, Lu Liu, Mykola Pechenizkiy, Xiao Xiang Zhu, and Tianjin Huang. Reobench: Benchmarking robustness of earth observation foundation models. *arXiv preprint arXiv:2505.16793*, 2025. 1
- [41] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 5
- [42] Valerio Marsocci, Yuru Jia, Georges Le Bellier, David Kerekes, Liang Zeng, Sebastian Hafner, Sebastian Gerard, Eric Brune, Ritu Yadav, Ali Shibli, et al. Pangaea: A global and inclusive benchmark for geospatial foundation models. *arXiv preprint arXiv:2412.04204*, 2024. 1, 5, 8
- [43] Francesc Marti Escofet, Benedikt Blumenstiel, Linus Scheibenreif, Paolo Fraccaro, and Konrad Schindler. Fine-tune smarter, not harder: Parameter-efficient fine-tuning for geospatial foundation models. In *Joint European Conference on Machine*

- Learning and Knowledge Discovery in Databases*, pages 516–532. Springer, 2025. 5
- [44] Andrea Nascetti, Ritu Yadav, Kirill Brodt, Qixun Qu, Hongwei Fan, Yuri Shendryk, Isha Shah, and Christine Chung. Biomasters: A benchmark dataset for forest biomass estimation using multi-modal satellite time-series. *Advances in Neural Information Processing Systems*, 36:20409–20420, 2023. 4
- [45] Andrea Nascetti, Ritu Yadav, Kirill Brodt, Qixun Qu, Hongwei Fan, Yuri Shendryk, Isha Shah, and Christine Chung. Biomasters: A benchmark dataset for forest biomass estimation using multi-modal satellite time-series. *Advances in Neural Information Processing Systems*, 36:20409–20420, 2023. 12
- [46] Esther Rolf, Jonathan Proctor, Tamma Carleton, Ian Bolliger, Vaishaal Shankar, Miyabi Ishihara, Benjamin Recht, and Solomon Hsiang. A generalizable and accessible approach to machine learning with global satellite imagery. *Nature communications*, 12(1):4392, 2021. 9, 14
- [47] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. 5
- [48] Vivien Sainte Fare Garnot, Loic Landrieu, and Nesrine Chehata. Multi-modal temporal attention models for crop mapping from satellite time series. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2022. 4
- [49] Michael Schmitt, Seyed Ali Ahmadi, Yonghao Xu, Gülşen Taşkin, Ujjwal Verma, Francescopaolo Sica, and Ronny Hänsch. There are no data like more data: Datasets for deep learning in earth observation. *IEEE Geoscience and Remote Sensing Magazine*, 11(3):63–97, 2023. 2
- [50] Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025. 6, 7, 8
- [51] Gencer Sumbul, Marcela Charfuelan, Begüm Demir, and Volker Markl. Bigearthnet: A large-scale benchmark archive for remote sensing image understanding. In *IGARSS 2019-2019 IEEE international geoscience and remote sensing symposium*, pages 5901–5904. IEEE, 2019. 12
- [52] Daniela Szwarcman, Sujit Roy, Paolo Fraccaro, Þorsteinn Elfi Gíslason, Benedikt Blumenstiel, Rinki Ghosal, Pedro Henrique de Oliveira, Joao Lucas de Sousa Almeida, Rocco Sedona, Yanghui Kang, et al. Prithvi-eo-2.0: A versatile multi-temporal foundation model for earth observation applications. *arXiv preprint arXiv:2412.02732*, 2024. 4, 7, 16
- [53] Aysim Toker, Lukas Kondmann, Mark Weber, Marvin Eisenberger, Andrés Camero, Jingliang Hu, Ariadna Pregel Hoderlein, Çağlar Şenaras, Timothy Davis, Daniel Cremers, et al. Dynamicearthnet: Daily multi-spectral satellite dataset for semantic change segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21158–21167, 2022. 4
- [54] Aysim Toker, Lukas Kondmann, Mark Weber, Marvin Eisenberger, Andrés Camero, Jingliang Hu, Ariadna Pregel Hoderlein, Çağlar Şenaras, Timothy Davis, Daniel Cremers, et al. Dynamicearthnet: Daily multi-spectral satellite dataset for semantic change segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21158–21167, 2022. 12
- [55] Adam Van Etten and Daniel Hogan. The spacenet multi-temporal urban development challenge. In *NeurIPS 2020 Competition and Demonstration Track*, pages 216–232. PMLR, 2021. 14
- [56] Adam Van Etten, Dave Lindenbaum, and Todd M Bacastow. Spacenet: A remote sensing dataset and challenge series. *arXiv preprint arXiv:1807.01232*, 2018. 4
- [57] Adam Van Etten, Daniel Hogan, Jesus Martinez Manso, Jacob Shermeyer, Nicholas Weir, and Ryan Lewis. The multi-temporal urban development spacenet dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6398–6407, 2021. 14
- [58] Yi Wang, Conrad M Albrecht, Nassim Ait Ali Braham, Chenying Liu, Zhitong Xiong, and Xiao Xiang Zhu. Decoupling common and unique representations for multimodal self-supervised learning, 2024. 7
- [59] Yi Wang, Zhitong Xiong, Chenying Liu, Adam J Stewart, Thomas Dujardin, Nikolaos Ioannis Bountos, Angelos Zavras, Franziska Gerken, Ioannis Papoutsis, Laura Leal-Taixé, and Xiao Xiang Zhu. Towards a unified copernicus foundation model for earth vision. *arXiv preprint arXiv:2503.11849*, 2025. 1
- [60] Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Siddhartha Naidu, et al. Livebench: A challenging, contamination-free llm benchmark. *arXiv preprint arXiv:2406.19314*, 4, 2024. 8
- [61] Ross Wightman. Pytorch image models. <https://github.com/huggingface/pytorch-image-models>, 2019. 7
- [62] Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, et al. The fair guiding principles for scientific data management and stewardship. *Scientific data*, 3(1):1–9, 2016. 3
- [63] Zhitong Xiong, Yi Wang, Fahong Zhang, Adam J. Stewart, Joëlle Hanna, Damian Borth, Ioannis Papoutsis, Bertrand Le Saux, Gustau Camps-Valls, and Xiao Xiang Zhu. Neural plasticity-inspired multimodal foundation model for earth observation, 2024. 7
- [64] Zongzhe Xu, Ritvik Gupta, Wenduo Cheng, Alexander Shen, Junhong Shen, Ameet Talwalkar, and Mikhail Khodak. Specialized foundation models struggle to beat supervised baselines. In *The Thirteenth International Conference on Learning Representations*, 2025. 8
- [65] Xiao Xiang Zhu, Jingliang Hu, Chunping Qiu, Yilei Shi, Jian Kang, Lichao Mou, Hossein Bagheri, Matthias Haberle, Yuansheng Hua, Rong Huang, et al. So2sat lc42: A benchmark data set for the classification of global local climate zones [software and data sets]. *IEEE Geoscience and Remote Sensing Magazine*, 8(3):76–89, 2020. 4
- [66] Xiao Xiang Zhu, Zhitong Xiong, Yi Wang, Adam J Stewart, Konrad Heidler, Yuanyuan Wang, Zhenghang Yuan, Thomas Dujardin, Qingsong Xu, and Yilei Shi. On the foundations of earth and climate foundation models. *arXiv preprint arXiv:2405.04285*, 2024. 8



## 6. Appendix

### 6.1. Datasets

The following sections introduce the different datasets as well as the respective preprocessing schemes chosen to generate the benchmark version of the dataset.

#### 6.1.1. BigEarthNet V2

BigEarthNet V2 [15] is an updated version of the popular BigEarthNet [51] multi-label classification dataset with Sentinel 1 and 2 imagery over Europe. The classification labels are derived from the CORINE Land Cover (CLC) labels. In BigEarthNet V2, the Sentinel 2 images were updated through a newer atmospheric correction scheme, the labels were revised with the most recent CLC labels, and a new geographic train, validation, and test split was introduced that reduces their spatial correlation.

#### 6.1.2. BioMassters

The BioMassters dataset [45] consists of multi-temporal Sentinel 1 and 2 imagery over Finland and has pixel-wise above-ground biomass (AGB) annotations that were derived from airborne LiDAR measurements. The dataset was released with a designated train, validation, and test set but without geospatial information.

#### 6.1.3. CaFFe

The Calving Fronts and Where to Find them (CAFFe) dataset [29] includes SAR imagery of glaciers from Antarctica, Greenland, and Alaska, with zone segmentation labels of rock, glacier, and ocean or ice melange. In the designated train, validation, and test set, the test set locations are disjoint from the other sets. The dataset was released as single-channel PNG




Figure 5. Training Set Examples for BigEarthNet V2 dataset.




Figure 6. Training Set Examples for Biomasssters dataset.

files, but was reprocessed as TIF files with merged location information from a metadata table.

#### 6.1.4. Cloudsen12

Introduced by [8], the Cloudsen12+ dataset forms a comprehensive global collection of Sentinel-1 and Sentinel-2 imagery for cloud and shadow detection. Samples were assigned to a train, validation, or test set based on a spatially stratified block split strategy. For our benchmark, we only select the processed samples with 512x512 pixels and the "high-quality" label tag.

#### 6.1.5. DynamicEarthNet

The DynamicEarthNet dataset [54] is a global multi-temporal dataset of high-resolution Planet Lab imagery and Sentinel-2




Figure 7. Training Set Examples for CaFFe dataset.




Figure 8. Training Set Examples for CloudSen12 dataset.

imagery with seven land use and land cover (LULC) classes. Based on 75 different areas of interest, across six continents, the dataset provides daily Planet imagery across these locations from 2018 to 2020. The samples released within the publicly available subset were processed in the following way. First, the 1024x1024 tiles were split into four 512x512 patches. Subsequently, an eight-by-eight grid binning strategy was used to assign samples to a train, validation or test set. The resulting splits neither overlap in space nor time.




Figure 9. Training Set Examples for DynamicEarthNet dataset.

### 6.1.6. EverWatch

The EverWatch [24] is a bird detection dataset with high-resolution imagery over the Everglades region. The aerial drone imagery was manually annotated with seven different bird species labels. The dataset comes with a designated train, validation, and test set and only partially complete geospatial information. The variable-sized tiles were resized to patches of size 512x512.




Figure 10. Training Set Examples for EverWatch dataset.

### 6.1.7. Fields of The World

The Fields of the World (FTW) dataset [37] is a global dataset of Sentinel-2 imagery for agricultural field boundary delineation. Each sample consists of two temporal views from different time stamps and a field mask. The dataset was filtered by country to only include data consistent with the open CC-BY-SA license. For this subset, the original train, validation, and test split was used.




Figure 11. Training Set Examples for FTW dataset.






Figure 12. Training Set Examples for FLAIR2 dataset.

### 6.1.8. FLAIR2

The French Land cover from Aerospace ImageRy (FLAIR) Version 2 dataset [21] is an updated version of FLAIR [20] that now also contains Sentinel-2 time-series data next to the high-resolution aerial imagery with thirteen semantic land cover classes across France. The new version also includes an updated test split from distinct spatial domains across the country.

### 6.1.9. KuroSiwo

The KuroSiwo dataset [12] is a global dataset of SAR imagery from a broad range of global flooding events with annotations for permanent water and flood water. We use the designated train, validation, and test set.

### 6.1.10. PASTIS-R

The Panoptic Agricultural Satellite Time Series (PASTIS) dataset [25] is a multi-temporal dataset with Sentinel-1 SAR






Figure 13. Training Set Examples for KuroSiwo dataset.



(a) Training Set Examples for PASTIS Semantic Segmentation dataset.



(b) Training Set Examples for PASTIS Instance Segmentation dataset.

Figure 14. Training Set Examples for PASTIS datasets.

and Sentinel-2 imagery with crop type annotations across France. We utilize the PASTIS-R version, which is a superset of the original PASTIS dataset and includes the SAR imagery. The dataset comes with a designated train, validation, and test split.

### 6.1.11. SpaceNet2

The second edition of the SpaceNet dataset series [55] consists of high-resolution aerial imagery across four different cities (Las Vegas, Paris, Shanghai, Khartoum) with binary building footprint annotations. As a competition dataset, the publicly available data that includes annotations is only designated for training. To have separate training and evaluation sets, we use the following procedure: For each city, we use a checkerboard style split, inspired by MOSAICS [46], that overlays a grid structure and all samples within a grid are assigned to a train, validation and test set such that their percentage of the total samples is roughly 70-20-10 respectively.

### 6.1.12. SpaceNet7

The seventh edition of the SpaceNet dataset series [57] is a multi-temporal building footprint dataset of 101 different locations across the globe. For each location, there is a sequence




Figure 15. Training Set Examples for SpaceNet2 dataset.

of Planet Lab Dove RGB imagery with fine-grained building annotations. While the intended purpose was to track building changes over time, we cast it as a static building segmentation dataset. The original 1024x1024 tiles were separated into 4 512x512 patches, and samples were assigned to a train, validation, or test split such that their locations are disjoint and their percentage of total samples is roughly 70-20-10 respectively.

#### 6.1.13. TreeSatAI

The TreeSatAI dataset [4] originally included single-timestamp aerial, Sentinel-1, and Sentinel-2 imagery with corresponding multi-label tree species classes. [7] extended the dataset to include a time series for the Sentinel modalities.




Figure 16. Training Set Examples for Spacenet7 dataset.




Figure 17. Training Set Examples for TreeSatAI dataset.

The samples again were assigned to a train, validation, or test set with a checkerboard style split strategy, where a 10x10 grid was overlaid on the dataset area of northern Germany, and samples within each block belong to one split such that there is roughly a 70-20-10 distribution across splits.

#### 6.1.14. NZCattle

The NZCattle is a dataset included in GEO-Bench [38]. It includes high-resolution aerial RGB images to detect cattle in New Zealand. Originally, the dataset was presented as a segmentation task, which here we repurpose it as an object detection task. Original splits were retained.



Figure 18. Training Set Examples for NZCattle dataset.

#### 6.1.15. BurnScars

This dataset contains Harmonized Landsat and Sentinel-2 imagery of burn scars and the associated masks for the years 2018-2021 over the contiguous United States [52]. There are 804 512x512 images. For the benchmark we took the original dataset and just repurposed it in TACO format with no modifications.




Figure 19. Training Set Examples for BurnScars dataset.

#### 6.1.16. Substation

This dataset was curated by TransitionZero and sourced from publicly available data repositories, including OpenStreetMap



Figure 20. Training Set Examples for Substation dataset.

and Copernicus Sentinel data [35]. The dataset consists of Sentinel-2 228x228 pixels images from 27k+ global locations; the original task was to semantically segment power substations. For the benchmark, we converted the original labels into instance segmentation labels. Furthermore, as most locations had 4-5 images taken at different timesteps (i.e., revisits), but beyond clouds occlusions, there is, in principle, no need for multiple timestamps to complete the task, we created a cloud-free composite for each location. Data was sub-sampled for each split to reduce computational requirements.

#### 6.1.17. So2sat

So2Sat is a dataset designed for multimodal classification. The aim is to classify local Climate Zones globally. We downloaded the modified version of this dataset from GEO-Bench [38] and repurposed it in TACO format with no modifications.

#### 6.1.18. Forestnet

Forestnet is a curated dataset of Landsat 8 satellite images of known forest loss events paired with driver annotations from






Figure 21. Training Set Examples for So2Sat dataset.




Figure 22. Training Set Examples for Forestnet dataset.

expert interpreters in South Asia. We downloaded the modified version of this dataset from GEO-Bench [38] and repurposed it in TACO format with no modifications.

#### 6.1.19. Dataset Geographical Distribution by continent

Figure 23 shows the geographical distribution of dataset samples across continents.

#### 6.1.20. Dataset Statistics

Z-score normalization statistics were computed modality and band-wise for each dataset on the training set split and subsequently used to normalize all inputs. For the BioMasters




Figure 23. Distribution of samples weighted by source dataset size across continents. Each dataset’s continent distribution is normalizing by the dataset size before aggregating across all dataset.

dataset that contains continuous pixel-wise regression labels, z-score normalization was also used for the labels.

## 6.2. Detailed experimental setup

To ensure reliable results, the evaluation protocol employs a combination of hyperparameter optimization (HPO) followed by repeated seed experiments. For segmentation, regression, and classification tasks, we conduct 16 trials during the HPO phase. Due to higher computational requirements, detection experiments use 10 trials. After the HPO process for a given model and dataset combination, the set of parameters yielding the best validation set metric is selected and subsequently used for repeated experiments with five different random seeds. The final performance is then calculated on the test for each seed. Both HPO and repeated experiments are configured through the TerraTorch Iterate library, which automates the benchmarking workflow. HPO employs Optuna [6] with Bayesian optimization and a Tree Structured Parzen Estimator (TPE) for efficient sampling. The general optimization search space was fixed across all models and datasets:

- Learning Rate: A real value ranging from  $1e-6$  to  $1e-3$ .
- Batch Size: An integer value selected from the list:  $\{8, 16, 32\}$ .

An Adam optimizer is used with a fixed weight decay of 0.01. Training incorporates early stopping with a patience of 10 epochs. For object detection, the batch size is fixed at 8, and HPO is conducted solely on the learning rate.

## 6.3. Additional Results

### 6.3.1. Main Raw Results by Dataset

The raw results obtained using these base experimental settings for each of the 19 datasets are presented in Figure 24. Beyond the main experimental settings, the succeeding sections give results from the various ablations conducted to understand the effect of changing settings on benchmarking.





Figure 24. Main results: Raw performance by dataset.

### 6.3.2. Frozen Encoder Ablation

To assess the impact of freezing the encoder during training, we performed an ablation where the encoder remains fixed while all other settings match the main experiments. This ablation was applied to all pixel-wise and classification datasets using a subset of models. Figure 25 shows a side-by-side comparison of full fine-tuning versus frozen encoder aggregated performance for each model, while Figure 26 presents the raw results per dataset. A systematic drop in performance is observed when the encoder is frozen, suggesting that all models benefit from some fine-tuning. Importantly, a change in model ranking is also observed. Figure 4 shows the aggregate change in model rank due to this ablation along with rank changes due to other ablations.

It should be noted that the drop in performance is not the same across all datasets. To investigate the effect size by dataset, a paired t-test is performed and the Cohen’s D effect size is computed for datasets that have a statistically significant change. Figure 27

shows the Cohen’s D effect size for datasets with a statistically significant drop in test metric due to frozen backbones (ag-




Figure 25. Aggregated Frozen Vs Unfrozen Ablation: freezing the encoder results in a systematic drop in performance.



Figure 26. Frozen Vs Unfrozen Encoder Ablation: Raw results per dataset.




Figure 27. Cohen's D effect size for Unfrozen/Frozen Ablation: Per dataset reduction in performance due to freezing encoder

gregated across all models). Datasets such as treesat and benv2 have the biggest drop in performance with encoder freezing.

### 6.3.3. Linear Decoder Ablation

To evaluate the impact of using a linear decoder for pixel-wise tasks, we ran an ablation replacing the UNet decoder with a single linear layer, keeping all other settings identical to the main experiments. This was applied to all pixel-wise datasets

(segmentation and regression) using a subset of models.

Figure 28 compares aggregated performance of UNet versus Linear decoder for each model, while Figure 29 shows the raw results per dataset. Again, a linear decoder shows lower performance and substantially affects scoring. Figure 30 shows variation in Cohen's D effect size per dataset for datasets with statistically significant reduction. In particular,




Figure 28. Aggregated UNet Vs Linear Decoder Ablation: linear decoder lowers performance for all models.




Figure 29. UNet Vs Linear Decoder Ablation: Raw results per dataset.




Figure 30. Cohen's D effect size for UNet/Linear Decoder Ablation: Per dataset reduction in performance due to linear decoder.

fotw and spacenet7 datasets have the largest decrease. Figure 4 demonstrates that while there is an aggregate change in model rank due to this ablation, the difference in rank is smaller than changes observed due to other ablations.

### 6.3.4. Multi-Spectral Vs RGB Data Ablation

To assess the role of multi-spectral data, we ran an ablation using only GeoFM encoders on datasets with additional bands beyond RGB. Unlike the main experiments, which used all available bands, this ablation restricted inputs to RGB only, while keeping all other settings unchanged.

Figure 32 shows the raw per-dataset results, while 31 shows aggregated performance. A general drop in performance is ob-

served in both figures when multi-spectral bands are excluded, thus highlighting the importance of multi-spectral data. Similar to the other ablations, the Cohen's D effect size is computed for datasets that have a statistically significant change according to a paired t-test. Datasets with a statistically significant drop in performance due to excluding multi-spectral bands (aggregated across all models) are highlighted in Figure 33. The burn scar dataset suffers the biggest reduction by far, possibly due to some burn scars not being detectable with simple RGB data. This emphasizes the importance of multi-spectral data for some down-




Figure 31. Aggregated Multi-Spectral Vs RGB Data Ablation: excluding multi-spectral data tends to reduce performance.




Figure 32. Multi-Spectral Vs RGB Data Ablation: Raw results per dataset.




Figure 33. Cohen's D effect size for Multi-Spectral/RGB Ablation: Per dataset reduction in performance due to excluding multi-spectral data.

stream tasks in EO. The Multi-Spectral-dependent capability is comprised of datasets with a statistically significant change.

### 6.3.5. Ground Sampling Distance (GSD) Ablation

This ablation applies a 14x14 kernel filter to increase the Ground Sampling Distance (GSD) of validation and test images, while train images are kept at the original resolution. The best HPO settings are used for training with 5 repeated seeds for each model on each dataset.

All other settings used in the main experiments were maintained. Datasets with a resolution under 10m are used with a subset of models. Ground sampling distance is an important factor in satellite data, determining the level of detail that can be observed in a satellite image. A lower GSD (high resolution) implies that each pixel covers a smaller area on the ground and finer details are clearer when looking at the overall image. Figure 34 shows raw results per dataset from this ablation, while 35 shows aggregated performance. Both figures highlight that increasing GSD tends to lower test metrics across all models.






Figure 34. Increasing Ground Sampling Distance Ablation: Raw results on 5 datasets.




Figure 35. Aggregated Original GSD Vs Increased GSD: lower GSD can improve performance in some cases.




Figure 36. Cohen's D effect size for Increasing Ground Sampling Distance Ablation: Per dataset reduction in performance due to increased GSD.

This is expected as the images with lower (original) GSD have more details. However, it is notable that some datasets exhibit a smaller effect size as demonstrated by the results Cohen's D effect sizes of Figure 36.

### 6.3.6. Multi-Modality Ablation

To investigate the effect of multi-modal inputs, this ablation applies both Optical and SAR data simultaneously to each model versus only using optical data (e.g. Sentinel 2 or Planet data). All other settings used in the main experiments were maintained. The raw results of Figure 37 show that this ablation is particularly dependent on the dataset and model.

For the pastis dataset, multi-modal inputs reduce test metrics regardless of the model, suggesting that the combination of the




Figure 37. Multi-modal Vs Optical only: change in performance with multi-modal input vs Optical input only. Ablation on all multi-modal datasets and models only

multi-modal inputs may not be beneficial for this dataset. Conversely, the biomassers dataset consistently shows an improvement across models. Benv2 dataset has more nuanced results: while TerraMind-V1-Large increases test metrics with multi-modal inputs, the remaining models have either similar or lower performance. Overall, the effect of multi-modal inputs is inconclusive as performance appears to vary by model and dataset.

### 6.3.7. Multi-Temporal Vs Mono-Temporal Ablation

To evaluate the effect of using a single timestamp, we ran an ablation on multi-temporal datasets where inputs were restricted to one timestamp instead of multiple (up to seven in the main experiments), keeping all other settings unchanged.

Applied to a subset of models, this ablation revealed substantial performance drops across some datasets such as pastis, thus confirming the importance of temporal information for these particular datasets in the multi-temporal capability. Raw per-dataset results are shown in Figure 38. Additionally, Figure 4 also shows that changing from multi-temporal to mono-temporal inputs causes the biggest change in rank, second only to the effect of freezing the encoder.

### 6.3.8. Cosine Annealing Ablation

In this ablation, we applied a cosine annealing schedule to adjust the learning rate during training. Figure 39 shows the raw per-dataset results. Outcomes varied across datasets and




Figure 38. Multi-Temporal Vs Mono-Temporal Ablation: raw results on 4 datasets.




Figure 39. Cosine annealing Learning rate Vs Fixed Learning rate: raw results on 3 datasets are inconclusive.




Figure 40. Linear Warm Up on Learning rate Vs Fixed Learning rate: raw results on 3 datasets are inconclusive.

models, providing no definitive conclusion.

### 6.3.9. Linear Warm Up Ablation

In this ablation, we applied linear warm-up to adjust the learning rate during the first five training epochs. Figure 40 shows the raw per-dataset results. Outcomes varied across datasets and models not showing any clear trend.