

On the Foundations of Earth and Climate Foundation Models

Xiao Xiang Zhu^{1,2*}, Zhitong Xiong^{1†}, Yi Wang^{1†}, Adam J. Stewart^{1†},
Konrad Heidler^{1†}, Yuanyuan Wang¹, Zhenghang Yuan¹, Thomas Dujardin¹,
Qingsong Xu¹, Yilei Shi³

¹Chair of Data Science in Earth Observation, Technical University of Munich, Arcisstraße 21, 80333 Munich, Germany.

²Munich Center for Machine Learning, 80333 Munich, Germany.

³School of Engineering and Design, Technical University of Munich, Arcisstraße 21, 80333 Munich, Germany.

*Corresponding author(s). E-mail(s): xiaoxiang.zhu@tum.de;

†These authors contributed equally to this work.

Abstract

Foundation models have enormous potential in advancing Earth and climate sciences, however, current approaches may not be optimal as they focus on a few basic features of a desirable Earth and climate foundation model. Crafting the ideal Earth foundation model, we define eleven features which would allow such a foundation model to be beneficial for any geoscientific downstream application in an environmental- and human-centric manner. We further shed light on the way forward to achieve the ideal model and to evaluate Earth foundation models. What comes after foundation models? Energy efficient adaptation, adversarial defenses, and interpretability are among the emerging directions.

As deep learning continues to mature, a general trend towards larger datasets (with millions to billions of data points¹) and larger model architectures (with billions to trillions of parameters²) has emerged. In the past decade in particular, we have witnessed a paradigm shift from single-purpose models to general-purpose models, and from supervised pre-training to self-supervised pre-training. This has been led by the rise in popularity of models like CLIP³, LLaMA⁴, GPT⁵, and SAM⁶ spawning billion dollar companies and revolutionizing many aspects of daily life.

Bommasani et al.⁷ proposed the term “foundation model” (FM) to encompass “any model that is trained on broad data (generally using self-supervision at scale) that can be adapted (e.g., fine-tuned) to a wide range of downstream tasks” in order to study this phenomenon. The majority of FMs like CLIP and GPT focus on the image and text domains. In this work, we specifically focus on “data” and “downstream tasks” relating to the Earth and its climate system, as shown in Fig. 1.

We choose to limit the scope of our work to the Earth’s surface and atmosphere for three reasons. First, the Earth’s surface and troposphere are our home, and include the majority of processes that directly impact and are impacted by human activity. Second, the “big data” revolution required to pre-train FMs can currently only be provided at scale using remote sensing data, which is only available at the surface. Third, the timescale at which the deep ocean and mantle operate is too slow to study in combination with a human timescale, and can be considered to be more or less static.

The Potential of Earth Foundation Models

Earth and climate science has inevitably entered the era of big data. This data, largely unlabeled and vast in volume, can offer novel insights into our planet’s systems. Earth observation (EO) data-informed Earth and Climate modeling offers perspectives towards higher resolution and less biased future projections.

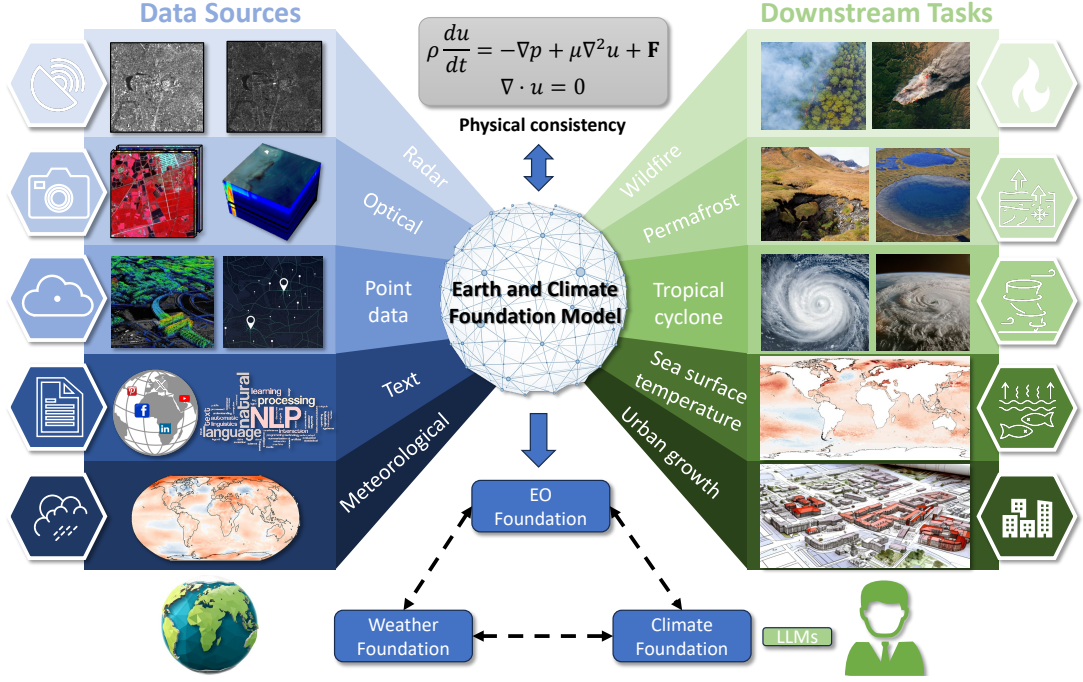


Fig. 1: The scheme of an Earth and climate FM. It should be trained on common data modalities, including imagery (radar, optical), non-Euclidean data (point clouds, text), and meteorological data. It should also provide consistency with physical laws. The FM is task-agnostic, exemplified by five possible downstream tasks. Because of the large difference of the characteristics among Earth observation, weather, and climate data, the resulting FM may consist of a pool of expert models, where feedback loops exist among each of them. For attribution of figure elements, please see Supplementary Information.

However, the sheer volume and redundancy of this information calls for more sophisticated analytical frameworks⁸. FMs are proving to be a glimmer of hope in this context. This section delves into the critical reasons underpinning the need for FMs in the Earth and climate sciences, illustrating their potential to revolutionize our understanding of the Earth system towards its sustainable management.

- **Unlocking the value of big EO and climate data:** The primary allure of FMs lies in their ability to harness the vast, often unlabeled, reservoirs of EO and climate data by means of learning general representations from them⁷. Traditional models falter under the weight of this information overload, struggling to process and extract meaningful insights without extensive, manually labeled datasets. FMs, on the other hand, can learn task-agnostic general feature representations from massive amounts of unlabelled data. This capability not only improves our understanding of Earth's systems, but also accelerates the pace of discovery from the learned general features.
- **Enhancing label efficiency/boosting performance in downstream tasks:** Downstream EO tasks often suffer from a lack of labeled data; a bottleneck that can significantly impact the training and performance of models. FMs address this limitation by being inherently more efficient. They are trained through self-supervised learning on large, unlabeled datasets, which can then be fine-tuned with smaller, task-specific datasets, reducing the reliance on extensive labeled data⁹. This efficiency not only accelerates the applications of artificial intelligence (AI) for EO, but also improves their accuracy, robustness and reliability.
- **Reducing carbon footprints:** The environmental impact of training AI models from scratch for every new use case cannot be overstated¹⁰. Each training cycle consumes considerable computational resources, contributing to the carbon footprint of AI research and applications. FMs offer a sustainable alternative by allowing for the reuse and adaptation of pre-trained models across multiple use cases. This approach significantly reduces the computational demand and, consequently, the environmental impact of AI model training in the Earth and climate sciences.
- **Bridging EO and climate science:** The synergy between EO and climate science is pivotal for the development of informed climate models, in particular, given the severe spatial and time-span mismatch between high resolution observations and climate modeling and projections. Earth and climate FMs facilitate this integration by providing a framework through which EO data can directly inform and enhance climate models. This is essential for improving the resolution and

prediction accuracy of climate models, and reducing their biases and uncertainty, and thus leads to more effective and sustainable climate change mitigation and adaptation strategies.

- **Improving Earth system modeling:** The study of climate change trends and the prediction of future weather and climate scenarios have traditionally relied on general circulation models and numerical weather prediction models. However, the complexity and scale of Earth system modeling (ESM) pose challenges to computational capacity and efficiency with these conventional methods. FMs offer scale invariance and efficient feature representations, presenting an opportunity for a unified approach to ESM¹¹. Furthermore, ESM’s complexity and uncertainties require an interdisciplinary methodology to advance scientific knowledge discovery. Generative FMs show promise in analyzing diverse ESM data modalities, including images and sequences, to facilitate tasks such as physical law discovery^{2,12}.


To this end, advances in Earth and climate FMs would lead to a paradigm shift in how we approach the vast and complex observational and simulated data characterizing our planet. By unlocking the value of big EO and climate data, enhancing label efficiency, reducing carbon footprints, and bridging the gap between EO and climate science, FMs pave the way for a more sustainable, efficient, and comprehensive understanding of the Earth system. Their potential to revolutionize Earth and climate sciences is immense, in particular when it comes to inform critical decisions in the face of climate change, e.g. by improving impact modeling and ESM.

Big (Earth) Data


The development of Earth and climate FMs relies on vast and varied data for self-supervised pre-training. Below, we delve into the most important data sources for both training and evaluation.

- **EO data archive:** To date, there are more than 1000 active remote sensing satellites in space²⁰. Various EO archives provide a huge volume of data sources for building EO FMs. As Fig. 2a shows, the data volume of EO missions whose data is open and free drastically increased over the past decade, surging to nearly 100 PB today, not including commercial data with comparable volume. These archives provide a continuous, high-quality, and open stream of satellite imagery, capturing a wide range of Earth’s features at various spatial, spectral, and temporal resolutions. For example, the Sentinel series²¹ includes multiple satellites with different sensors designed to monitor the land, ocean, and atmosphere, and the Landsat series²² commands attention for its extensive temporal reach, providing an invaluable long-term perspective on Earth’s surface changes.
- **Weather patterns and climate models:** In meteorology, various weather and atmospheric science agencies, including ECMWF and NOAA, provide accurate and comprehensive datasets for understanding the Earth’s climate system. Reanalysis products such as ERA5²³, JRA-55²⁴, NCEP²⁵, and HRRR²⁶ are already widely used for FM pre-training. High-resolution, historical, and near-real-time data on atmospheric conditions make it possible to build large-scale weather nowcasting and medium-range forecasting systems. Furthermore, climate models like CMIP6²⁷ and MERRA-2²⁸ offer simulations of past, present, and future climate conditions built on strong physical understandings of the climate system that are valuable to climate FMs. Fig. 2a highlights a number of these datasets.
- **Curated ML-ready datasets:** The rise of machine learning (ML) in the past decade has driven the EO community to curate and publish over 500 ML-ready datasets¹⁹. Integrated with human annotation, these supervised datasets cover a wide range of applications, such as land-use land-cover mapping²⁹, change detection³⁰, object detection³¹, and disaster monitoring³², as illustrated in Fig. 2b. The diverse array of applications along with semantic labels can steer the FMs toward nuanced understanding and broad applicability. Several recent works on FMs have also released a set of curated pre-training datasets without labels^{33–36}. Compared to the raw EO archive, these datasets are compact yet comprehensive, crafted with strategic sampling strategies that enhance the efficiency of model training. Nevertheless, in contrast to EO data archive, existing labelled datasets in EO amounts to less than 0.1 PB to date (see the right y -axis of Fig. 2a), i.e., only about 0.1% of the archive of open EO data, highlighting the huge discrepancy between them.

Apart from the above, other sources such as natural language texts, cadastral records, and historical in-situ measurements are also valuable for the training of Earth and climate FMs.



(a) Historical, current, and estimated future volume of open EO and climate data archives in petabytes^{13–19}. Labeled dataset volumes are shown by a dashed line and use a different y -axis scale. Note that the volume of labeled data is less than 0.1% of the total volume of unlabeled data.



(b) Evolution of labeled EO datasets over the last 16 years¹⁹. Bubble size and y -axis height indicate the volume of labeled data in each dataset. Bubble color denotes the type of dataset.

Fig. 2: Volume of EO and climate data archives and curated benchmark datasets.

Crafting the Ideal Earth and Climate Foundation Model

An ideal Earth and climate FM (see Fig. 1) should take domain-specific characteristics in Earth and climate sciences into account, and thus **must** feature:

1. **Geolocation embedding:** EO and climate data are linked to a geolocation. Embedding geolocation into the FM, either in a hard-coded way by importing the latitude and longitude as inputs, or in a smoother manner by contrasting features according to their geographical distances, can improve feature representation or even facilitate the study of localized events, such as heatwaves or hurricanes. Researchers can potentially pinpoint the exact locations of environmental phenomena, enabling precise analysis of regional climate patterns and their global implications.
2. **Balanced geographical representations:** To date, ML models for Earth and climate sciences have been trained with a bias towards data-rich regions. To take urban applications as an example, it is a common practice to consider open source OpenStreetMap³⁷ data for training models, leading to a severe bias in favor of Europe and North America. A balanced geographical coverage implies

a careful selection of the data used for training the FMs, to ensure that the FMs accurately reflect the diversity of Earth’s surfaces, climates, and ecosystems.

3. **Scale awareness:** The spatial resolution of EO and climate data varies from sub-meter for high-resolution aerial images to hundreds of kilometers for global climate models. The ideal FM should have scale awareness of different input data and learn spatial features referring to a certain characteristic scale. Practically, this is a mandatory feature for the model to deal with scale variances of the input data. More importantly, the FM should be capable of examining large-scale patterns, such as ocean circulations, while also delving into finer-scale processes, like land cover changes.
4. **Wavelength embedding:** Every individual satellite has a unique specification in terms of wavelengths and number of bands. The ideal FM must incorporate data across different wavelengths, including visible light, infrared, ultraviolet, and microwave, and possess different numbers of channels, ranging from one (e.g., panchromatic) to several hundred (e.g., hyperspectral). Fine-tuning pre-trained models should be possible with input data of any wavelength and number of bands, including completely new and unseen sensors.
5. **The time variable:** EO and climate time-series data offer valuable new capabilities in monitoring the changing planet and making predictions with unprecedented spatial and temporal resolution. Therefore, the ideal FM must be able to learn sequential relations of data within its spatial context. This feature is vital, e.g., for distinguishing between natural variability and anthropogenic impact on the Earth surface, detecting anomalies, and identifying emerging patterns.
6. **Multisensory:** Multi-source and multisensory EO and climate data offer complementary information in order to serve various tasks. The FM must be capable of learning general feature representations across all data sources while retaining the unique characteristics of each source. This dual capability ensures that the model can draw upon a broad range of data and explore their joint and source-specific features to build a comprehensive picture of the Earth and its climate system. Meanwhile, preserving source-specific features allows for, e.g., the accurate attribution of climate phenomena to their respective causes.
7. **Task-agnostic:** In Earth and climate sciences, AI models have been trained to serve various tasks, such as image classification, semantic segmentation, detection, regression, and forecasting. As a FM, it must be task-agnostic yet beneficial for a wide spectrum of downstream use cases. Potential benefits are, for example, enhanced label efficiency, boosted performance, improved generalizability and transferability, and better robustness against label noise.
8. **Carbon minimized:** Given the urgency of addressing climate change, it is indispensable that the development and operation of FMs minimize their carbon footprint. This can be achieved through unified models for more general purposes, efficient computing practices, the use of renewable energy sources, and the optimization of model architectures. By prioritizing sustainability in its design, the FM aligns with the broader goals of reducing greenhouse gas emissions and promoting environmental stewardship.

In addition to the above features that are considered must-haves, we also consider a number of features that are **highly desirable** for the ideal Earth and climate FM.


Uncertainty quantification: Quantifying uncertainty and detecting out-of-distribution instances is a crucial ability for FMs, especially when adapted to extreme Earth events. In this situation, data scarcity and unexpected situations can significantly impact the performance and robustness of downstream tasks. Considering this, enabling FMs with uncertainty quantification and out-of-distribution detection ability can make them more reliable in critical applications, such as disaster response.

Physical consistency: Purely data-driven FMs can make predictions based on non-physical relationships or spurious correlations¹¹. Incorporating principles of physical consistency, such as conservation, symmetry, and causality, improves the pre-training and fine-tuning of FMs, enhancing their transferability to new domains and increasing their transparency. Physically consistent models, alternatively termed knowledge-guided or physics-aware ML^{38,39}, represent a promising research direction to address physical consistency of Earth and climate FMs.

AI assistants: Developing ML systems based on the Earth and climate FMs to understand and interpret complex EO data in a human-like manner would provide researchers and decision-makers with intuitive insights and recommendations tailored to specific scenarios. This approach can further enrich common user interactions through interfaces with the incorporation of natural language, making the model more user-friendly and facilitating the communication of EO and climate science to a wider audience.

1. Geolocation embedding
2. Balanced geographical representations
3. Scale awareness
4. Wavelength embedding
5. The time variable
6. Multisensory
7. Task-agnostic
8. Carbon minimized

9. Uncertainty quantification
10. Physical consistency
11. AI assistants



Recent Advances and Gaps




Fig. 4: Representative EO and climate FMs. Representative models are chosen based on both popularity and novelty. Note that no single model excels at all evaluation criteria, as most models focus on only one or two of all “must-have” features.

Pioneering studies has been conducted on the development of Earth and climate FMs. However, compared to the ideal FM we describe above, existing models are still in their infancy, focusing on two or three of the above “must-have” features but rarely approaching coverage of all eight or addressing any “highly desirable” features. In this section, we discuss the state of the art and current limitations with a focus on EO and climate FMs, which are typically developed separately and with different downstream

applications in mind. A general view of several representative EO and climate models is shown in Fig. 4a, with the spatial scale potentially linking the two fields in monitoring the same planet. As it stands, there exists a clear gap in the spatial resolution at which EO and climate models operate. In order to create a unified EO and climate model, not only does the spatial connection need to be consolidated, but also other aspects such as the time dimension need to be considered.

Earth observation foundation models

EO has a relatively rich record of FM development, owing in part to the many similarities between remotely-sensed and natural images and the abundance of unlabeled data (Fig. 2a). Existing architectures, self-supervised techniques, and data augmentation methods can be conveniently adopted, leading to the widespread availability of EO FMs⁴⁰. Similar to computer vision models, EO models demonstrate *task-agnostic* representation capabilities when trained on million image-scale datasets^{36,41}. Nevertheless, the unique characteristics of EO data and applications call for domain-specific considerations⁴², leading to many exciting recent advances.

Many researchers start with careful curation of pre-training datasets in order to achieve *geographical balance*. SeCo³³ and SSL4EO^{34,35} rely on sampling from regions with high spatial variability, while others like SatlasPretrain³⁶ focus on scale, with imagery covering a total area of 21 million km², roughly 10% of the total land surface. Other researchers focus on developing novel self-supervised techniques designed to incorporate *geolocation embedding*, including Tile2Vec⁴³ via contrastive learning and GASSL⁴⁴ via an explicit pretext task. Scale-MAE⁴⁵ uses the ground sampling distance as input to the ViT positional encoder, allowing for explicit *scale awareness*. A number of recent FMs, including SatMAE⁴⁶, Presto⁴⁷, Prithvi⁴⁸, SatlasPretrain³⁶, and SkySense⁴¹, also incorporate *time*.

Support for *multisensor* inputs has also required a number of architectural and optimization modifications^{41,49–52}, owing to differences between the number of spectral bands and wavelengths captured by each sensor. Among them, DOFA⁵² currently supports the widest range of sensor modalities, including SAR, RGB, MSI, and HSI. It incorporates *wavelength embedding* using a wavelength-based dynamic weight generation strategy, allowing it to adapt to unseen spectral signals. A number of recent papers explore integration between vision and language models in EO^{53–56}, paving the way for *AI assistants*.

Weather and climate foundation models

Compared to EO, the use of ML for weather forecasting and climate modeling is relatively recent. One of the greatest successes of climate FMs has been *minimizing the carbon footprint* of weather forecasting. As an example, FourCastNet⁵⁷ is capable of generating a week-long global forecast in less than two seconds on a single GPU, requiring 12,000 times less energy than ECMWF’s IFS⁵⁸. Faster prediction times also enable larger ensembles, bolstering capabilities for *uncertainty quantification*.

The majority of modern weather and climate FMs are trained on *multisensor* reanalysis products like ERA5 at a global scale, and thus innately offer *geolocation embedding* and *balanced geographical representations*. Although ERA5 limits many models to a spatial resolution of 0.25°, a notable exception is the recent FengWu-GHR⁵⁹ model, which brings the spatial resolution down to 0.09° (about 10km) using priors from pre-trained high resolution models.

Weather and climate forecasting also happen at a number of *temporal* scales, including nowcasting (0–2 hours), medium-range forecasting (2 hours–2 weeks), subseasonal (2 weeks–2 months) to seasonal (2 months–2 years) forecasting, and true “climate” modeling (decadal scale). The majority of weather FMs fall under “medium-range forecasting”, and predict a number of *task-agnostic* atmospheric and surface phenomena every 6 hours. The first weather FM to outperform traditional numerical models, Pangu-Weather⁶⁰, combines multi-resolution models to allow for predictions every 1 hour. Subseasonal models like DLWP⁶¹ and true “climate” models like ACE⁶² are less common, and even fewer models are able to handle multiple temporal scales like ClimaX⁶³.

In terms of *physical consistency*, NeuralGCM represents a unique and promising modeling direction, with a fully differentiable general circulation model encompassing many physics-based atmospheric dynamics. The rapid progress of large language models has also impacted climate models, making them more user-friendly and widely applicable through natural language integration. For instance, ClimateBert⁶⁴ facilitates climate text analysis, and ClimateGPT⁶⁵ serves as an *AI assistant* for climate change and sustainability topics.

Gaps towards the ideal model

Despite significant advances towards many of the “must-have” and “highly desirable” features of an ideal FM, large gaps still remain. Works like SeCo and SSL4EO that focus on pre-training dataset curation tend to oversample from densely-populated regions like cities and undersample from underpopulated regions like rainforests, polar regions, and oceans, each critical for ecology, climatology, and oceanography. Furthermore, pre-training datasets tend to focus on medium- to high-resolution RGB and MSI, with few large-scale datasets for SAR, HSI, or Lidar data or for very high- or low- resolution imagery⁴⁰. Multisensor models like DOFA are rare, and most models need to be trained from scratch on every new imaging platform encountered, making it difficult to take advantage of the thousands of satellites in orbit. Time-series EO modeling is also in its infancy, with few models able to handle variable length image sequences with irregular spacing. Integrating uncertainty quantification and physical consistency in EO remains underexplored.

Weather and climate FMs have come a long way in a short period of time, yet still show similar gaps in many aspects. Although inference is several orders of magnitude faster than traditional numerical models, training can be a bottleneck for many researchers. For example, NeuralGCM required 10 days to train on 128 TPU v5e, equivalent to roughly 130K hours on a single NVIDIA A100 GPU. Most FMs are currently limited by the 0.25° resolution of ERA5, and struggle to approach higher resolutions. Medium-range forecasting has seen great progress, while nowcasting, subseasonal to seasonal (S2S) forecasting, and climate projection remain underexplored. Weather FMs suffers from overly smooth forecasts and bias drift as lead time increases⁶⁶, making long-range forecasting an active challenge.

As seen in Fig. 4a, there remains a large gap in spatial resolution between the EO and climate domains, despite obvious connections between surface processes and weather. Table 1 illustrates current progress towards the ideal FM with several prominent examples.

FMs	1. Geolocation Embedding	2. Balanced Distribution	3. Scale Awareness	4. Wavelength Embedding	5. The Time variable	6. Multisensory	7. Task-agnostic	8. Carbon-minimized	9. Uncertainty	10. Physical Consistency	11. Language Alignment
SatMAE ⁴⁶	✗	✗	✗	✓	✓	✗	✓	✗	✗	✗	✗
Scale-MAE ⁴⁵	✗	✗	✓	✗	✗	✗	✓	✗	✗	✗	✗
SatlasPretrain ³⁶	✗	✗	✗	✗	✓	✓	✓	✗	✗	✗	✗
SSL4EO-L ³⁵	✗	✓	✗	✗	✗	✓	✓	✗	✗	✗	✗
SkySense ⁴¹	✓	✗	✗	✗	✓	✓	✓	✗	✗	✗	✗
DOFA ⁵²	✗	✗	✗	✓	✗	✓	✓	✓	✗	✗	✗
FourCastNet ⁵⁷	✓	✓	✗	-	✓	✓	✓	✓	✓	✗	✗
Pangu-Weather ⁶⁷	✓	✓	✗	-	✓	✓	✓	✓	✓	✗	✗
GraphCast ⁶⁸	✓	✓	✗	-	✓	✓	✓	✓	✗	✗	✗
ClimaX ⁶³	✓	✓	✗	-	✓	✓	✓	✗	✗	✗	✗
NeuralGCM ⁶⁹	✓	✓	✓	-	✓	✓	✓	✗	✗	✓	✗
FengWu-GHR ⁵⁹	✓	✓	✓	-	✓	✓	✓	✗	✗	✗	✗

Table 1: Representative EO and climate FMs and which of the 11 features of an ideal FM have been addressed in their model design or training procedure. Wavelength embedding is ignored for climate FMs due to the lack of a wavelength component in atmospheric measurement data. Note that many climate FMs are carbon-minimized at inference time (compared to numerical weather prediction models), but still require substantial computing resources during training.

The Way Forward

To achieve the ideal Earth and climate FMs, two factors should be carefully considered: 1) design a comprehensive data construction process to ensure the high-quality of training data; 2) develop model architectures with key features tailored to the specific challenges of these domains. Fig. 5 presents a general workflow of different specific modules, which we will dig into in the following.

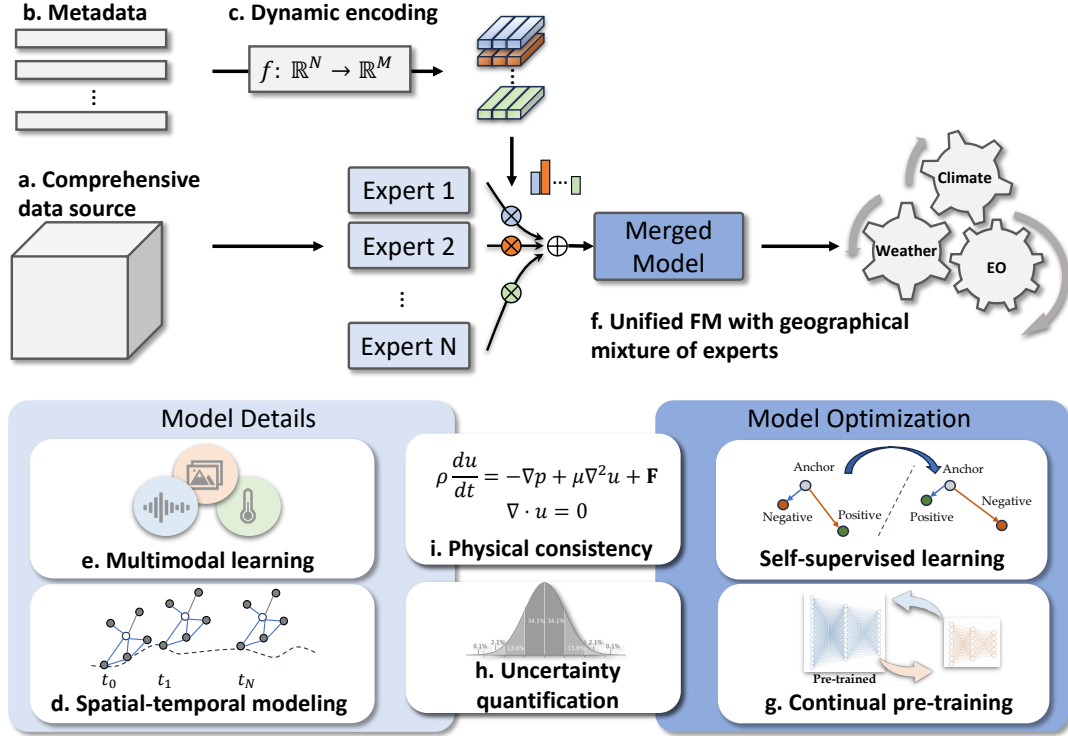


Fig. 5: The way forward towards the ideal Earth and climate FM, from data curation to model design and training. We illustrate the design principles to achieve a unified FM, including dynamic encoding, spatial-temporal modeling, geographical mixture of experts, task-agnostic multimodal learning, carbon minimization, physical consistency, uncertainty quantification, and language alignment.

Data curation

- a. **Comprehensive data sources:** Diversity is the foundation of an effective large-scale pretraining dataset. Satellite imagery, ground-based sensors, weather stations, and climate models are all important sources of data for model training. The diverse data sources ensure a holistic representation of the Earth’s system, and capture phenomena across scales and conditions. To ensure geographical diversity, it is vital to include data across evenly distributed geographical regions to avoid data bias. To improve modality diversity, different data modalities including language data are crucial to capturing complex data relationships and dynamic processes.
- b. **Metadata integration and standardization:** Metadata is the crucial bridge linking diverse observation data sources, playing a key role in organizing extensive data observed from Earth’s surface and atmosphere. Metadata like the data source, collection time, modality, resolution, cloud coverage, spectral bands, and preprocessing details are critical for enabling FMs with features like geolocation embedding, scale awareness, and wavelength embedding. As for weather and climate phenomena, it is important to employ standardized ontologies for each data point to maintain consistency across datasets.

Model design and training

- c. **Dynamic encoder with conditional computation:** Dynamic encoder that employs conditional computation techniques allows the model to efficiently manage a broad spectrum of inputs flexibly.

Among the promising techniques to process varying input sizes are: an adaptive input projection layer designed to dynamically adjust according to the quantity of input features⁵²; Transformers with recurrent mechanism to process sequences of variable lengths; an input chunking strategy, complemented by mechanisms that preserves context across segments. These sophisticated design strategies bolster the model’s capability to analyze complex, multi-scale data and support features like geolocation embedding, wavelength embedding, and scale awareness.

- d. **Spatio-temporal analysis across multiple scales:** EO data is captured at distinct spatial and temporal resolutions. To effectively process this complex data, it is essential to design effective spatiotemporal attention mechanisms to simultaneously capture the temporal and spatial dependencies. It facilitates an understanding of how conditions evolve and vary across different regions, thereby enhancing the model’s performance in understanding Earth’s dynamic systems.
- e. **Multi-modal learning framework:** An ideal FM, pre-trained with data from multiple sources, must skillfully exploit the synergy among diverse data modalities, weaving together complementary insights and common knowledge⁴⁹. To accomplish this, methods that concentrate on two essential aspects of multi-modal learning need to be designed^{50,70}. One is the meta-modality representation learning, which extracts and synthesizes information across various modalities, creating a meta-representation that captures overarching patterns. The other is modality-specific representation learning, which utilizes modality-specific encoders to learn rich, modality-specific features.
- f. **Unified FM with geographical mixture of experts:** Designing an effective Geographical Mixture of Experts (GeoMoE) strategy is pivotal to harnessing the collective strengths of FMs in Earth and climate fields. GeoMoE integrates specialized models across various domains, enabling the framework to precisely allocate the right expertise. GeoMoE can not only enhance performance by adeptly adapting to diverse situations but also reduce computational costs by judiciously deploying resources. The adaptable nature of GeoMoE allows for the seamless integration of emerging scientific discoveries, ensuring the FMs remain at the forefront of technological and scientific progress.
- g. **Continual pre-training for carbon-minimized FMs:** Incorporating continual learning mechanisms⁷¹, regularization techniques, and transfer learning strategies is critical to creating a carbon-minimized and adaptable Earth and climate FM. These strategies allow the model to persistently update its knowledge, enhancing performance and generalization in response to new and evolving scenarios. Techniques like elastic weight consolidation⁷² to prevent catastrophic forgetting, and knowledge distillation for efficient model scaling are pivotal.
- h. **Uncertainty quantification:** In EO and climate modeling, where handling petabyte-scale datasets is a trend, efficient uncertainty quantification is critical. Sparse Gaussian processes⁷³ offer a scalable Bayesian approach, using inducing points to manage the computational load. Ensemble methods utilize multiple models to capture a broad spectrum of outcomes, enhancing performance and providing robust uncertainty measures. Quantile regression directly computes uncertainty intervals, offering insights into prediction reliability efficiently. These candidate techniques have the potential to offer the FMs with uncertainty quantification ability required in Earth and climate sciences.
- i. **Physical consistency:** Physical consistency can be achieved through the integration of physics into datasets, architectures, or loss functions, thereby augmenting the FM’s generalization and interpretability. For instance, physical knowledge is usually represented by differential equations with analytical or numerical solutions. Physics-informed neural networks^{74,75} incorporate these equations into the neural network’s loss function, enabling unsupervised learning for understanding physical processes. Furthermore, employing a hybrid learning framework that integrates physical models and FMs within a differentiable neural network can result in the establishment of a structured geodynamics FM.
- j. **Alignment with large language models:** Merging EO, weather, and climate FMs with language processing capabilities is vital for enhancing their analytical depth and facilitating user interaction⁷⁶. By designing effective projection modules, the Earth and climate FMs can learn to align with language capabilities to gain open-vocabulary learning ability⁷⁷, and enrich the understanding of an extensive range of terms and concepts.

Evaluating Earth Foundation Models

Task-specific models are usually evaluated by a specific test dataset. Summarizing model performance through numerical metrics such as accuracy or F1-score allows for easy comparison of such models. FMs, however, are supposed to be beneficial for a wide range of tasks, posing challenges to the comparability of FMs. In natural language processing, FMs are evaluated through so-called *benchmarks*, which specify a

collection of evaluation tasks. The tasks within a benchmark are selected to evaluate different aspects of downstream applications⁷⁸. In computer vision, benchmarks are mostly assembled ad-hoc from existing evaluation datasets⁷⁹.

Useful Earth FMs must learn robust representations for diverse EO, weather, and climate data. Evaluating these models requires standardized benchmarking suites similar to the ones proposed in natural language processing to assess model performance across various tasks and datasets. This section discusses key requirements for such benchmarks and examines existing benchmarks, highlighting their strengths and limitations. Ultimately, we identify gaps where future benchmarks can offer even more comprehensive evaluations of Earth FMs.

Requirements for Earth benchmarks

- **Data Diversity:** Benchmarks for Earth FMs need to evaluate the model on diverse data. For EO, various data modalities need to be covered, such as multispectral, SAR, and hyperspectral data. For weather and climate data, various variables should be included, like atmospheric pressure, temperature, and precipitation. Another aspect of data diversity is spatio-temporal coverage of the evaluation tasks, which should represent most parts of the globe and range from near real-time applications to long time-series including historical data. For climate data, even future projections should be included.
- **Representative Tasks:** Furthermore, the tasks in the benchmark need to be representative of downstream tasks that the models could be applied for. For EO, this includes tasks like scene classification, object detection, segmentation, change-detection, time-series analysis, and regression. For the weather and climate domains, possible tasks are forecasting, extreme event detection, downscaling, or emulation of physical climate models and corresponding climate scenario analysis.
- **Interaction between components:** Ideally, a benchmark for Earth FMs should not only evaluate the model on its individual aspects of EO, weather, and climate, but also evaluate their integration. For this, tailored tasks shall be defined that require the interaction between components. Possible examples for such tasks are the downscaling weather data using satellite imagery, combining weather and satellite data to detect extreme events like droughts, or improving long-term climate projections with knowledge about short-term weather patterns.

Existing benchmark datasets

For evaluating EO models, SustainBench⁸⁰ focuses on the United Nations’ sustainable development goals, encompassing 15 tasks. The tasks are representative, covering classification, segmentation, regression and time-series analysis. It evaluates on a number of optical sensors like Landsat, Sentinel-2 and MODIS. However, SAR data is not considered in this benchmark. Finally, owing to the design of the dataset as a sustainability benchmark, the considered tasks are limited to agricultural and humanitarian applications. More recently, GEO-Bench⁸¹ is a benchmarking suite designed specifically for evaluating EO FMs. It includes diverse Earth-relevant tasks and has standardized evaluation protocols. By limiting evaluation to classification and segmentation tasks, however, it omits important EO tasks, such as regression or time-series analysis. MDAS⁸³ and PhilEOBench⁸⁴ are suites that trade larger datasets for a smaller number of tasks. Finally, Prithvi⁴⁸ and SkySense⁴¹ have proposed new benchmarks alongside their FMs. Table 2a summarizes these benchmarks by the evaluated modalities, tasks, resolutions and timespan.

For weather data, the WeatherBench⁸⁵ and WeatherBench 2⁸⁶ collections provide a comprehensive benchmark for medium-range weather forecasts on a global scale. These datasets include both historical records and model forecasts of essential meteorological variables, allowing for standardized evaluation of weather FMs. An increasing number of weather FMs are directly evaluated on ERA5^{67,91,92}. ExtremeWeather⁸⁷ specifically evaluates weather models on extreme weather events. Nowcasting is typically done only for precipitation⁹³, for which highly resolved datasets such as MRMS⁹⁰ exist. As a first multi-task benchmark, ClimateLearn⁸⁸ evaluates not only weather forecasting, but also extreme weather prediction and downscaling.

ClimateBench⁸⁹ collects physical simulations under different forcing scenarios to aid the development and evaluation of data-driven climate emulation models. Similar to the trend of evaluating weather models directly on ERA5, some climate models also directly evaluated against CMIP6 simulation outputs²⁷. Weather and climate benchmarks are summarized by evaluation tasks, resolution and timespan in Table 2b.

Name	# Tasks	RGB	Multispectral	Hyperspectral	SAR	Time Series	Classification	Regression	Segmentation	Object Det.	Change Det.	Resolution	Timespan
SustainBench ⁸⁰	15	✓	✓	✗	✗	✓	✓	✓	✓	✗	✗	0.6 m–30 m	1996–2019
GEO-Bench ⁸¹	12	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	0.1 m–15 m	2001–2021
FoMo-Bench ⁸²	16	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	0.01 m–60 m	2011–2023
MDAS ⁸³	3	✓	✓	✓	✓	✗	✗	✗	✓	✗	✗	0.3 m–30 m	2018
PhilEO Bench ⁸⁴	3	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	10 m	N/A
SkySense ⁴¹	17	✓	✓	✗	✓	✓	✓	✗	✓	✓	✓	0.05 m–30 m	2002–2019
Prithvi ⁴⁸	4	✗	✓	✗	✓	✓	✗	✗	✓	✗	✗	30 m	2017–2022

(a) Benchmarks for EO FMs

Name	Nowcasting	Medium-Range	S2S Forecasting	Climate Projection	Extreme Events	Downscaling	Resolution (Lat)	Resolution (Lon)	Resolution (Time)	Timespan
ERA5 ^{† 23}	✓	✓	✓	✓	✗	✗	0.25°	0.25°	1 h	1959–2024
WeatherBench ⁸⁵	✗	✓	✗	✗	✗	✗	5.63°	5.63°	6 h	1979–2018
WeatherBench 2 ⁸⁶	✗	✓	✗	✗	✗	✗	1.5°	1.5°	6 h	1979–2023
ExtremeWeather ⁸⁷	✗	✗	✗	✗	✓	✗	0.23°	0.31°	6 h	1979–2005
ClimateLearn ⁸⁸	✗	✗	✗	✗	✓	✓	5.63°	5.63°	6 h	1979–2018
CMIP6 ^{† 27}	✗	✓	✓	✓	✗	✗	1.25°	2.5°	24 h	1850–2100
ClimateBench ⁸⁹	✗	✗	✗	✓	✗	✗	1.89°	2.5°	1 yr	1850–2100
MRMS ^{† 90}	✓	✓	✓	✗	✗	✓	0.01°	0.01°	2 min	2014–2024

[†]Raw dataset without associated evaluation protocol

(b) Benchmarks for weather and climate FMs

Table 2: Overview of existing benchmark datasets for EO FMs (a) and weather/climate FMs (b)

Gaps in existing benchmarks

Currently, there is **no consensus on how to adapt a pre-trained FM for a downstream task**. Some works fine-tune the entire model, updating all parameters. Other works only train a single final network layer in a process called linear probing. Finally, parameter-efficient fine-tuning methods are another option. Benchmark results will only be comparable when models have been adapted to the downstream tasks using the same fine-tuning method. In particular, benchmarks are missing for few- and zero-shot settings.

However, **certain regions are not represented** in existing EO benchmarks. This limits the potential for cross-disciplinary research and holistic understanding of global climate systems. Including benchmarks for the missing regions such as the oceans⁹⁴ or polar areas⁹⁵ in evaluations would facilitate better modeling of phenomena like sea-level rise or ocean circulation patterns.

Further, evaluation datasets are often filtered for specific acquisition conditions like minimal cloud cover, which is **not representative of real-world data**. This can introduce artificial biases in the evaluation scores, making the models appear more capable than they truly are when applied to the diverse and often imperfect data encountered in actual use. Similarly, weather and climate benchmarks are often based on reanalysis products and might benefit from evaluations on observational data as well.

For evaluating Earth FMs, the main gap in existing benchmarks is **the lack of evaluation tasks that integrate EO data with weather and climate data**. This limits our ability to assess how well these models can capture the complex interplay between land, ocean, atmosphere, and human activity.

The rapid development of Earth FMs currently faces a challenge with **the lack of standardized evaluation benchmarks**. With each model employing its own unique set of tasks, objective comparisons

and assessments are difficult. A convergence towards standardized models is needed to ensure fair and transparent evaluations, and allow researchers to determine the strengths and weaknesses of each model. This, in turn, would greatly aid collaboration and accelerate the advancement of FMs for the Earth sciences.

Perspectives

What comes after the ideal Earth FM? Identifying the most relevant research questions is crucial for maximizing the utility of the model in advancing our understanding of Earth and climate science. We give seven recommendations, as follows:

- **Energy-efficient Adaptation:** Upon acquiring FMs, enhancing their efficiency for downstream tasks is vital for broad applications. Future research should explore strategies to boost computational, parameter, and memory efficiencies without compromising predictive performance. Possible research directions include quantization to reduce parameter precision, pruning to eliminate redundant parameters for faster inference, and knowledge distillation to transfer insights from larger to smaller models efficiently. These efforts aim not only to broaden the models’ applicability but also to advance toward more sustainable and energy-efficient ML development.
- **Foundation model enabled sciences:** Once an ideal Earth and climate FM becomes available, the most prominent research question is EO-informed weather and climate forecasting. As the FM is jointly trained from EO, weather and climate data, it can be explored to address the scale mismatch between observational data and the physical model of the Earth system and its climate. Exploiting feedback at higher spatial resolution can enable new exciting forecast and management applications. One additional example is to detect and predict extremes. As the FM is learned from massive data with “normal” patterns, it can flag deviations with much high accuracy, helping detect “extremes”, such as wildfires, floods, or other environmental disturbances.
- **Machine unlearning:** An ideal EO FM utilizes multimodal data, such as high-resolution and multispectral images, which can violate privacy by exposing sensitive information about individuals and infrastructures. Even basic optical imagery can pose a threat to privacy. Addressing this, the concept of “machine unlearning”, coined in 2015⁹⁶, aims to enable these models to effectively “forget” specific subsets of a dataset while minimizing negative effects on downstream tasks performances. This is crucial for compliance with regulations, as models’ generalization abilities partly rely on memorization⁹⁷, posing a trade-off between data protection and model performance.
- **Continual learning:** Efficiently updating FMs is also important. This includes integrating data from new sensors or modalities, upgrading outdated old data with better-quality new data, and ensuring up-to-date knowledge or constraints of the model with the evolving world. A natural technique is continual learning, which aims to allow models to learn new information without forgetting the past. However, compared to continual learning on task-specific applications, continual learning with generic FMs is more challenging due to model complexity, data size, computational cost, and task-agnostic objectives. To solve these, it’s crucial to develop new model architectures, training objectives, and more sophisticated optimization mechanisms.
- **Adversarial defenses:** Furthermore, the continuous learning of data integration process should be complemented with incorporating adversarial ML techniques to mitigate risks on FMs. This approach trains models to resist adversarial data, which are deliberately manipulated inputs designed to mislead the FM into making incorrect predictions⁹⁸. For instance, adversarial inputs might include subtly modified images, despite being barely noticeable, which can lead to inaccuracies in predictions and affect performance. By integrating adversarial ML, FMs become more robust and better equipped to adapt to new patterns when deployed in real-world scenarios.
- **Interpretability:** FMs with billions of parameters trained on massive unlabeled data present particular challenges in explaining their outputs, compared to smaller, specialized ML systems. The nature of FMs, i.e., serving as the backbone for a wide range of downstream applications, makes research on their interpretability of crucial importance. This can be achieved to a certain extent, e.g. by examining token prioritization, directly adjusting model weights, interpretable distillation, and large language model-based explanation techniques. Nevertheless, these tools may introduce biases due to inappropriate model setup, non-physical reasoning, and the absence of counterfactual reasoning capabilities and shall therefore be utilized with caution.
- **Cross-disciplinary inspiration:** Concepts from image FMs could inspire future improvements for Earth FMs, such as new architectures based on state-space models⁹⁹ or new ways of training FMs, like label-free scene understanding¹⁰⁰. Following similar requirements for multimodal FMs, medical

AI is currently moving towards combining highly heterogeneous sets of modalities within FMs¹⁰¹. For text FMs, retrieval augmented generation¹⁰² is currently opening up many new applications. Future Earth FMs could similarly retrieve data that is geospatially or temporally close to a given query to improve their responses.

Declarations

Acknowledgements. This work is jointly supported by the European Commission through the project “ThinkingEarth—Copernicus Foundation Models for a Thinking Earth” under the Horizon 2020 Research and Innovation program (Grant Agreement No. 101130544), by the German Federal Ministry for the Environment, Nature Conservation, Nuclear Safety and Consumer Protection (BMUV) based on a resolution of the German Bundestag (grant number: 67KI32002B; Acronym: *EKAPEx*), by the German Federal Ministry of Education and Research (BMBF) in the framework of the international future AI lab “AI4EO – Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond” (grant number: 01DD20001), by German Federal Ministry for Economic Affairs and Climate Action in the framework of the “national center of excellence ML4Earth” (grant number: 50EE2201C), by the Excellence Strategy of the Federal Government and the Länder through the TUM Innovation Network EarthCare and by Munich Center for Machine Learning.

Author contributions. X.Z. conceived the work and created an initial outline. All authors wrote and edited the manuscript.

Competing interests. The authors declare no competing interests.

Data availability. Data used to create figures in this paper are available at <https://github.com/zhu-xlab/EarthFoundationModels>.

Additional information. Supplementary information is available for this paper.

Correspondence and requests for materials should be addressed to X.Z.

Reprints and permissions information is available at www.nature.com/reprints.

Supplementary Information

Acronyms

The following standard acronyms are used in this paper, in order of appearance:

FM	Foundation Model
EO	Earth observation
AI	Artificial Intelligence
ML	Machine Learning
ESM	Earth System Model
PB	Petabytes
ECMWF	European Centre for Medium-Range Weather Forecasts
NOAA	National Oceanic and Atmospheric Administration
SAR	Synthetic-Aperature Radar
RGB	Red-Green-Blue
MSI	Multispectral Imagery
HSI	Hyperspectral Imagery
GPU	Graphics Processing Unit
TPU	Tensor Processing Unit
GeoMoE	Geographic Mixture of Experts
S2S	Subseasonal to Seasonal

The volume of big Earth data

To plot Fig. 2a, the accumulated volume of big Earth data, we search for available statistics about popular EO and climate archives, and fit an exponential curve for missing years and future trends. Main references include:

- **Sentinel:** Copernicus Sentinel Data Access Annual Reports¹³.
- **Landsat:** The 50-year Landsat collection 2 archive¹⁴, Soille *et al.*¹⁵.
- **MODIS:** Soille *et al.*¹⁵.

- **ERA:** The ERA5 global reanalysis^{16,17,23}.
- **CMIP:** Coupled Model Intercomparison Project Phase 6^{18,103}.
- **Labeled datasets:** EarthNets¹⁹.

Additional details of state-of-the-art foundation models

Earth observation foundation models

Extended Data Table 1 lists a set of representative EO FMs, from which we selected 6 to include in the figures. Data size takes the number of pixels. To create Fig. 4b, *modality diversity* gets the number of modalities; *resolution diversity* gets the number of stripes among {very high, high, medium, low}; *time series* gets the capacity level of processing time sequences among {unable, fixed, flexible}; *spatial coverage* gets the coverage level of pretraining data among {local, regional, global}; model size gets the log scaled number of parameters. These numbers are then scaled to 0-1.0 by dividing by the maximum for each category.

Weather and climate foundation models

Extended Data Table 2 lists the architecture, spatiotemporal resolution and range, model size, and computational requirements of several popular weather and climate FMs. Time span is the maximum attempted lead time used in each paper, and does not necessarily suggest the range over which forecasts can be made with any accuracy guarantee. Cluster size and training time are for pre-training, not inference. SEEDS is not an autoregressive model, and therefore has no temporal resolution. NeuralGCM is a family of models evaluated at different spatiotemporal resolutions. In order to create Fig. 4c, resolution is inverted, the log of all numbers is taken, and all numbers are scaled to the 0.2–1.0 range.

Vision language foundation models for Earth observation

Vision language model architectures

When the FM’s training dataset contains text, or semantic information that can be turned into text, existing vision-language models can easily be fine-tuned using this data, and a frozen instance of the FM. This allows for human-model interaction by the means of a large language model, and can even improve performances on downstream tasks, as empirical evidence suggests¹⁰⁴. Vision-language models such as, for instance, CLIP³, BLIP-2¹⁰⁵, and CogVLM¹⁰⁶ integrate visual and textual data to enhance their understanding and the human interaction capabilities they offer. CLIP processes image-text pairs by separately embedding them using modality-specific encoders and projecting them into a shared embedding space. CLIP performs "contrastive learning": when the image and the text are indeed paired (positive pair), their embeddings are brought closer through the maximization of their cosine similarity. When it is not the case (negative pair), the cosine similarity is instead minimized. CLIP offers basic human-model interaction, such as cross-modal image retrieval, and is commonly used as a text-enriched vision encoder. BLIP-2 improves over this by using an alignment module that generates embeddings of such quality that they can be directly projected into the hidden space of a large language model (LLM) along with a text prompt. This capability effectively facilitates the 'interrogation' of images, allowing for more sophisticated querying and interaction. CogVLM takes it further with a more detailed design, allowing for even better embeddings, for deeper and more meaningful human-model interactions.

In EO, RemoteCLIP⁵³ and SkyScriptCLIP⁵⁴ showcase their ability to handle EO data without overfitting, and require far fewer parameters than off-the-shelf CLIP models pretrained on large amounts of image-text data. Similarly, RS5M¹⁰⁷ pushes this further by displaying high performances on simple downstream tasks (but not on complex image-text tasks yet), such as zero-shot land cover classification, thus successfully training the current largest CLIP model in EO, three times larger than RemoteCLIP and SkyScriptCLIP. However, challenges remain in training models of such size to perform well on advanced visual reasoning tasks within remote sensing, highlighting an open problem related to the scarcity of textual data and its mixed quality in this domain. More recent initiatives, relying on much larger models, such as EarthGPT⁵⁶ developed for multisensor RS data comprehension, exhibits high performance across different tasks compared to some specialist models.

Curated vision-language datasets

Advances in vision-language models have sparked interest in the remote sensing community, owing to their promising results in visual reasoning. As argued above, there is a need for datasets that combine visual data with semantically rich text annotations. Extended Data Table 3 shows vision-language datasets for EO. The most explored image-text task in EO is image captioning^{108–110}, requiring the creation of sentences that describe the corresponding image. More complex tasks include visual question answering^{111–113} and visual grounding^{114,115}, which involve answering questions based on images and identifying specific objects within those images, respectively.

Complex, exhaustive text data is crucial for training vision-language models, highlighting the importance of rich, detailed image-text combinations to improve their visual reasoning skills. Despite the potential benefits, the availability of such datasets is limited, primarily due to the need for manual or costly Multimodal Large Language Model annotating, the exclusive ownership of some data by specific entities, the demonstrated impossibility to collect remote sensing image-text pairs at scale using web-crawled large-scale datasets¹¹⁶, and model filtering⁵⁴. To remedy to these issues, efforts have been made to generate remote sensing image-text pairs through indirect methods, such as leveraging semantic segmentation and object detection datasets to derive semantic information in the form of text^{53,111,113,115,115}, cross-referencing OpenStreetMap geotags with Google Earth Engine images to craft comprehensive captions or question/answer pairs^{54,112,114}. However, these methods yield a very limited fraction of usable data compared to the vast collections available for natural images, such as the extensive LAION-5B dataset¹¹⁶, on top of often producing non-natural descriptions. Recent initiatives like ChatEarthNet¹¹⁷ or RS5M¹⁰⁷ demonstrate the ongoing efforts to construct EO image-text datasets at scale by leveraging the capabilities of advanced vision-language models, such as GPT4-V¹¹⁸ that performs well in image captioning, but that is costly.

Distribution perspectives of foundation models

The majority of users who could potentially benefit from a FM are domain experts (as opposed to ML researchers) and non-technical policy experts with limited ML experience. It is paramount that a unified Earth and climate FM is as easy to use as ChatGPT and SAM, and does not require a Ph.D. in computer science to use. The base model architecture and weights can be distributed through existing popular geospatial ML libraries like Raster Vision¹¹⁹ or TorchGeo¹²⁰, with an additional graphical user interface for non-technical users. These libraries are already well tested across a variety of Python versions and platforms and have large dedicated userbases. An interactive interface that allows users to create labels for new benchmark datasets similar to SAM could greatly speed up research on a number of important topics.

Another concern is the license under which the model weights are released. Many potential EO data sources like Planet and Maxar data may prevent the public release of pretrained model weights and cannot be used to train the model. Ideally, the weights should be released under a permissive open source license like MIT, Apache, or BSD so as to benefit the greatest number of people. However, such a FM can also be used for harm, including applications like illegal surveillance, military aggression, and resource exploitation¹²¹. Given these concerns, it may make sense to consider ethical software licenses like the Hippocratic License that prohibit such activities, although this can be challenging if not impossible to legally enforce.

Attributions

Fig. 1: Background modified from template by PresentationGO (www.presentationgo.com). Icons from Adobe Stocks (stock.adobe.com): 126670692, by KundraK; 764196510, by AbtoCreative; 509096452, by Barudak Lier; 444274525, by vectorsanta; 734797956, SkyLine.

Fig. 3: Icons from Adobe Stocks: 398214492, by Vectorina; 589636564, by Gopal; 489357227, by SilenceVideo; 622171957, by lovemask; 734797956, by SkyLine; 265484642, by alekseyvanin; 583742888, by barks; 534436638, AmethystStudio; 593497565, LeonART. Sentinel-1 image from ESA (©ESA). Sentinel-2, -5p image from Wikipedia.

Extended Data

Extended Data Table 1: State-of-the-art EO FMs. The slash sign / indicates separate models for each; the range sign - indicates a rough range as the exact numbers are not available; the comma sign , indicates a combined coverage.


Model	Pre-training Dataset	Data Size (# pixels)	Modality	Spatial Resolution	Spatial Coverage (km ²)	Model Size (# params)	Temporal Coverage	Sequence Length
SeCo ³³	SeCo	70B	MS	10m	Global (1.4M)	23M	N/A	1
SatMAE ⁴⁶	fMoW	50B	MS / RGB	0.5m / 10m	Global (1.7M)	305M	5 years	3
RVSA ¹²²	MillionAID	4B	RGB	0.5m-150m	Global (18K)	89M	N/A	1
Scale-MAE ⁴⁵	fMoW-RGB	80B	RGB	0.3-10m	Global (1.7M)	305M	N/A	1
CROMA ⁴⁹	SSL4EO-S12	140B	MS, SAR	10m	Global (1.8M)	305M	N/A	1
DeCUR ⁵⁰	SSL4EO-S12 / GeoNRW	140B	MS, SAR / RGB, DEM	10m / 1m	Global (1.8M)	23M	N/A	1
Presto ⁴⁷	Presto	4B	MS, SAR, ERA5, DEM, etc.	10m	Global (2K)	400K	2 years	24
Prithvi ⁴⁸	HLS	158B	MS	10m, 30m	US	100M	1+ years	3
SpectralGPT ¹²³	fMoW+BigEarthNet	12B	MS	10m	Global (1.7M)	86M	N/A	1
SkySense ⁴¹	SkySense	97T	MS, RGB, SAR	0.3m, 10m	Global (8.8M)	2B	?	20
OFA-Net ⁵¹	Satlas. + 5B pix. + HyspecNet.	13B	MS, RGB, SAR, HS	1m, 4m, 10m, 30m	Global (260K)	86M	N/A	1
DOFA ⁵²	Satlas. + 5B pix. + HyspecNet.	105B	MS, RGB, SAR, HS	1m, 4m, 10m, 30m	Global (2.6M)	305M	N/A	1
SSL4EO-S12 ³⁴	SSL4EO-S12	140B	MS, SAR	10m	Global (1.8M)	23M	N/A	1
SSL4EO-L ³⁵	SSL4EO-L	348B	MS	30m	Global (15.6M)	23M	N/A	1
SatlasPretrain ³⁶	SatlasPretrain	17T	MS, RGB, SAR	0.5-2m, 10m	Global (21M)	88M	11 years	10
SkyScriptCLIP ⁵⁴	SkyScript	130B	MS, RGB, text	0.5m, 1-2m, 10-30m	Global (5M)	305M	N/A	1

Extended Data Table 2: State-of-the-art weather and climate FMs.

Model	Affiliation	Model Architecture	Temporal Resolution	Time Span	Spatial Resolution	Spatial Coverage	Model Size (# params)	Cluster Size	Training Time
KeislerNet ¹²⁴	Descartes	GNN	6 hrs	6 days	1°	Global	6.7M	1 A00	5.5 days
FourCastNet ¹²⁵	NVIDIA	Transformer/FNO	6 hrs	200 hrs	0.25°	Global	433M–2.19B	64 A100	16 hrs
Pangu-Weather ⁶⁷	Huawei	Transformer	1 hr	7 days	0.25°	Global	256M	192 V100	15 days
GraphCast ⁶⁸	Google	GNN	6 hrs	10 days	0.25°	Global	36.7M	32 TPU v4	4 wks
ClimaX ⁶³	UCLA	Transformer	6 hrs	1 month	0.25°	Global	108–111M	80 V100	3 days
FengWu ¹²⁶	Shanghai	Transformer	6 hrs	14 days	0.25°	Global	427M	32 A100	17 days
W-MAE ¹²⁷	Chengdu	Transformer	6 hrs	9 days	0.25°	Global	96.5M	8 A800	223.4 hrs
FuXi ⁹¹	Fudan	Transformer	6 hrs	15 days	0.25°	Global	1.56B	8 A100	30 hrs
SEEDS ¹²⁸	Google	Diffusion	N/A	16 days	2°	Global	114M	16 TPU v4	18 hrs
NeuralGCM ⁶⁹	Google	Differential GCM	3.75–12 mins	15 days	0.7–2.8°	Global	115–311M	128 TPU v5e	10 days
Stormer ¹²⁹	UCLA	Transformer	6 hrs	14 days	1.4°	Global	100–600M	128 A100	24 hrs
GenCast ⁹²	Google	Diffusion	12 hrs	15 days	1°	Global	60M	32 TPU v4	5 days
FengWu-GHR ⁵⁹	Shanghai	Transformer	6 hrs	10 days	0.09°	Global	4B	16 A100	3 days

Extended Data Table 3: Vision-language datasets in EO

Dataset	Year	# Images	# Sentences	Annotation	Image Size	Resolution	Source
UCM-Captions ¹⁰⁹	2016	2100	10500	Manual	256×256	0.3 m	USGS
Sydney-Captions ¹⁰⁹	2016	613	3065	Manual	500×500	0.6 m	Google Earth
RSICD ¹⁰⁸	2017	10921	54605	Manual	224×224	0.15 m – 15 m	Google Earth, Baidu
NWPU-Captions ¹¹⁰	2022	31500	157500	Manual	500×500	0.2 m – 0.3 m	Google Earth
RSITMD ¹³⁰	2022	4743	23715	Manual	224×224	various	Google Earth, RSICD
RSICap ¹³¹	2023	2585	?	Manual	512×512	various	GF-2, JL-1
RS5M ¹⁰⁷	2023	5M	25M	Manual/Auto	512×512	various	Google Earth, GF-2
SkyScript ⁵⁴	2023	2.6M	2.6M	OSM (auto)	various	0.1 m – 40 m	Google Earth Engine
ChatEarthNet ¹¹⁷	2024	163,488	163,484	ChatGPT (auto)	256×256	10 m	Google Earth Engine
LHRS-Align ¹³²	2024	1.15M	1.15M	Existing + GE (manual/auto)	768×768	1 m	Sentinel-2
LAION-EO ¹³³	2023	24,933	24,933	CLIP (auto)	Avg. 633×843	various	various
RSVQA-LR ¹¹²	2020	772	77,232	OSM (manual/auto)	256×256	10 m	Internet
RSVQA-HR ¹¹²	2020	10,659	1.07M	OSM (manual/auto)	512×512	0.15 m	Sentinel-2
RSVQA x BEN ¹¹³	2021	590,326	14.76M	Land cover	cf. BEN ²⁹	cf. BEN	USGS HRO
RSIVQA ¹¹¹	2022	37,264	111,134	Manual/Auto	various	0.1 m – 1 m	cf. BEN
FloodNet VQA ¹³⁴	2020	3,200	11,000	Manual (UAV)	224×224 / 3000×4000	UAV	various
CD-VQA ¹³⁵	2022	2968 pairs	122000 Q&A	Automatic	512×512	0.5 m – 3 m	Aerial platforms and sensors
RSIEval ¹³¹	2023	100	936	Manual	512×512	HR	Google Earth, GF-2, JL-1
LEVIR-CC ¹³⁶	2023	10077 pairs	50385	Manual	256×256	0.5 m	Google Earth
DIOR-RSVG ¹¹⁵	2023	17402	38320	Automatic	800×800	0.5 m – 30 m	Varied
VGRSI ¹¹⁴	2022	4239	7933	Automatic	1024×1024	Various	Google Earth, various sensors
MMRS-1M ⁵⁶	2024	1005842	1005842	Automatic	Various	Various	Varied, includes SAR/infrared



Extended Data Fig. 1: Visual reasoning datasets from the fields of Visual Question Answering, Remote Sensing Image Captioning, Change Detection Visual Question Answering, and Image-Text Cross-Modal Retrieval in remote sensing, with data point sizes proportional to citation counts. The curves represent CLIP models’ learning and overfitting ”frontiers”: CLIP-ViT-L-14 shows a transition from overfitting to learning for all sorts of image-text downstream tasks, as shown by SkyScript and RemoteCLIP^{53,54} and illustrated by the green curve. On the far right, CLIP-ViT-H-14 has the potential for learning adequate multimodal representations across simple downstream tasks without overfitting, but is still prone to poorer performances than smaller models as the downstream tasks get more complex, as demonstrated empirically by RS5M¹⁰⁷. This is indicative of overfitting. The y -axis represents the extent to which the dataset creation relied on automatic methods. Manual annotations are given a note of 1, machine-assisted human annotations and fully automatized annotations are respectively given a 2 and a 3. Note that the datasets from the figure might overlap.

References

- [1] Dehghani, M., Djolonga, J., Mustafa, B., Padlewski, P., Heek, J., Gilmer, J., Steiner, A.P., Caron, M., Geirhos, R., Alabdulmohsin, I., *et al.*: Scaling vision transformers to 22 billion parameters. In: International Conference on Machine Learning, pp. 7480–7512 (2023)
- [2] Hudson, N.C., Pauloski, J.G., Baughman, M., Kamatar, A., Sakarvadia, M., Ward, L., Chard, R., Bauer, A., Levental, M., Wang, W., *et al.*: Trillion parameter AI serving infrastructure for scientific discovery: A survey and vision. In: Proceedings of the IEEE/ACM International Conference on Big Data Computing, Applications and Technologies, pp. 1–10 (2023)
- [3] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., *et al.*: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning, pp. 8748–8763 (2021)
- [4] Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., *et al.*: LLaMA: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971 (2023)
- [5] Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., *et al.*: GPT-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
- [6] Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.-Y., *et al.*: Segment anything. In: Proceedings of the IEEE/CVF International

- Conference on Computer Vision, pp. 4015–4026 (2023)
- [7] Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., et al.: On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258 (2021)
 - [8] Zhu, X.X., Tuia, D., Mou, L., Xia, G.-S., Zhang, L., Xu, F., Fraundorfer, F.: Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine* **5**(4), 8–36 (2017)
 - [9] Wang, Y., Albrecht, C.M., Zhu, X.X.: Self-supervised vision transformers for joint SAR-optical representation learning. In: *IEEE International Geoscience and Remote Sensing Symposium*, pp. 139–142 (2022)
 - [10] Strubell, E., Ganesh, A., McCallum, A.: Energy and policy considerations for deep learning in NLP. In: *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pp. 3645–3650 (2019)
 - [11] Chen, S., Long, G., Jiang, J., Liu, D., Zhang, C.: Foundation models for weather and climate data understanding: A comprehensive survey. arXiv preprint arXiv:2312.03014 (2023)
 - [12] Wang, H., Fu, T., Du, Y., Gao, W., Huang, K., Liu, Z., Chandak, P., Liu, S., Van Katwyk, P., Deac, A., et al.: Scientific discovery in the age of artificial intelligence. *Nature* **620**(7972), 47–60 (2023)
 - [13] Castriotta, A.G., Volpi, F., Knowelden, R.: Copernicus sentinel data access annual report. Technical report, Serco Gael consortium for ESA
 - [14] Crawford, C.J., Roy, D.P., Arab, S., Barnes, C., Vermote, E., Hulley, G., Gerace, A., Choate, M., Engebretson, C., Micijevic, E., et al.: The 50-year Landsat collection 2 archive. *Science of Remote Sensing* **8**, 100103 (2023)
 - [15] Soille, P., Burger, A., De Marchi, D., Kempeneers, P., Rodriguez, D., Syrris, V., Vasilev, V.: A versatile data-intensive computing platform for information retrieval from big geospatial data. *Future Generation Computer Systems* **81**, 30–40 (2018)
 - [16] Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., et al.: ERA5 hourly data on single levels from 1940 to present. Copernicus Climate Change Service Climate Data Store (2018)
 - [17] Bell, B., Hersbach, H., Simmons, A., Berrisford, P., Dahlgren, P., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Radu, R., Schepers, D., et al.: The ERA5 global reanalysis: Preliminary extension to 1950. *Quarterly Journal of the Royal Meteorological Society* **147**(741), 4186–4227 (2021)
 - [18] Petrie, R., Denvil, S., Ames, S., Levavasseur, G., Fiore, S., Allen, C., Antonio, F., Berger, K., Bretonnière, P.-A., Cinquini, L., et al.: Coordinating an operational data distribution network for CMIP6 data. *Geoscientific Model Development Discussions* **2020**, 1–22 (2020)
 - [19] Xiong, Z., Zhang, F., Wang, Y., Shi, Y., Zhu, X.X.: EarthNets: Empowering AI in Earth observation. arXiv preprint arXiv:2210.04936 (2022)
 - [20] sp@ceboy, et al.: Orbiting Now: Active satellite orbit data. <https://orbit.ing-now.com/>
 - [21] Malenovsky, Z., Rott, H., Cihlar, J., Schaepman, M.E., García-Santos, G., Fernandes, R., Berger, M.: Sentinels for science: Potential of Sentinel-1,-2, and-3 missions for scientific observations of ocean, cryosphere, and land. *Remote Sensing of environment* **120**, 91–101 (2012)
 - [22] Wulder, M.A., Roy, D.P., Radeloff, V.C., Loveland, T.R., Anderson, M.C., Johnson, D.M., Healey, S., Zhu, Z., Scambos, T.A., Pahlevan, N., et al.: Fifty years of Landsat science and impacts. *Remote Sensing of Environment* **280**, 113195 (2022)

- [23] Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., *et al.*: The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* **146**(730), 1999–2049 (2020)
- [24] Kobayashi, S., Ota, Y., Harada, Y., Ebata, A., Moriya, M., Onoda, H., Onogi, K., Kamahori, H., Kobayashi, C., Endo, H., *et al.*: The JRA-55 reanalysis: General specifications and basic characteristics. *Journal of the Meteorological Society of Japan. Ser. II* **93**(1), 5–48 (2015)
- [25] Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaven, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., *et al.*: The NCEP/NCAR 40-year reanalysis project. *Renewable Energy*, 146–194 (2018)
- [26] Dowell, D.C., Alexander, C.R., James, E.P., Weygandt, S.S., Benjamin, S.G., Manikin, G.S., Blake, B.T., Brown, J.M., Olson, J.B., Hu, M., *et al.*: The high-resolution rapid refresh (HRRR): An hourly updating convection-allowing forecast model. Part I: Motivation and system description. *Weather and Forecasting* **37**(8), 1371–1395 (2022)
- [27] O’Neill, B.C., Tebaldi, C., Van Vuuren, D.P., Eyring, V., Friedlingstein, P., Hurtt, G., Knutti, R., Kriegler, E., Lamarque, J.-F., Lowe, J., *et al.*: The scenario model intercomparison project (ScenarioMIP) for CMIP6. *Geoscientific Model Development* **9**(9), 3461–3482 (2016)
- [28] Gelaro, R., McCarty, W., Suárez, M.J., Todling, R., Molod, A., Takacs, L., Randles, C.A., Darmenov, A., Bosilovich, M.G., Reichle, R., *et al.*: The modern-era retrospective analysis for research and applications, version 2 (MERRA-2). *Journal of Climate* **30**(14), 5419–5454 (2017)
- [29] Sumbul, G., Charfuelan, M., Demir, B., Markl, V.: BigEarthNet: A large-scale benchmark archive for remote sensing image understanding. In: *IEEE International Geoscience and Remote Sensing Symposium*, pp. 5901–5904 (2019)
- [30] Toker, A., Kondmann, L., Weber, M., Eisenberger, M., Camero, A., Hu, J., Hoderlein, A.P., Şenaras, Ç., Davis, T., Cremers, D., *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*, pp. 21158–21167 (2022)
- [31] Xia, G.-S., Bai, X., Ding, J., Zhu, Z., Belongie, S., Luo, J., Datcu, M., Pelillo, M., Zhang, L.: DOTA: A large-scale dataset for object detection in aerial images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3974–3983 (2018)
- [32] Xu, Q., Shi, Y., Zhu, X.X.: DisasterNets: Embedding machine learning in disaster mapping. In: *IEEE International Geoscience and Remote Sensing Symposium*, pp. 2536–2539 (2023)
- [33] Manas, O., Lacoste, A., Giró-i-Nieto, X., Vazquez, D., Rodriguez, P.: Seasonal Contrast: Unsupervised pre-training from uncured remote sensing data. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9414–9423 (2021)
- [34] Wang, Y., Braham, N.A.A., Xiong, Z., Liu, C., Albrecht, C.M., Zhu, X.X.: SSL4EO-S12: A large-scale multimodal, multitemporal dataset for self-supervised learning in Earth observation. *IEEE Geoscience and Remote Sensing Magazine* **11**(3), 98–106 (2023)
- [35] Stewart, A., Lehmann, N., Corley, I., Wang, Y., Chang, Y.-C., Ait Ali Braham, N.A., Sehgal, S., Robinson, C., Banerjee, A.: SSL4EO-L: Datasets and foundation models for Landsat imagery. In: *Advances in Neural Information Processing Systems*, vol. 36 (2024)
- [36] Bastani, F., Wolters, P., Gupta, R., Ferdinando, J., Kembhavi, A.: SatlasPretrain: A large-scale dataset for remote sensing image understanding. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16772–16782 (2023)
- [37] OpenStreetMap contributors: Planet dump retrieved from <https://planet.osm.org> . <https://www.openstreetmap.org> (2017)

- [38] Xu, Q., Shi, Y., Bamber, J., Tuo, Y., Ludwig, R., Zhu, X.X.: Physics-aware machine learning revolutionizes scientific paradigm for machine learning and process-based hydrology. *arXiv preprint arXiv:2310.05227* (2023)
- [39] Kashinath, K., Mustafa, M., Albert, A., Wu, J., Jiang, C., Esmacilzadeh, S., Azizzadenesheli, K., Wang, R., Chattopadhyay, A., Singh, A., et al.: Physics-informed machine learning: Case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A* **379**(2194) (2021)
- [40] Wang, Y., Albrecht, C.M., Braham, N.A.A., Mou, L., Zhu, X.X.: Self-supervised learning in remote sensing: A review. *IEEE Geoscience and Remote Sensing Magazine* **10**(4), 213–247 (2022)
- [41] Guo, X., Lao, J., Dang, B., Zhang, Y., Yu, L., Ru, L., Zhong, L., Huang, Z., Wu, K., Hu, D., et al.: SkySense: A multi-modal remote sensing foundation model towards universal interpretation for Earth observation imagery. *arXiv preprint arXiv:2312.10115* (2023)
- [42] Rolf, E., Klemmer, K., Robinson, C., Kerner, H.: Mission critical–satellite data is a distinct modality in machine learning. *arXiv preprint arXiv:2402.01444* (2024)
- [43] Jean, N., Wang, S., Samar, A., Azzari, G., Lobell, D., Ermon, S.: Tile2Vec: Unsupervised representation learning for spatially distributed data. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 3967–3974 (2019)
- [44] Ayush, K., Uz Kent, B., Meng, C., Tanmay, K., Burke, M., Lobell, D., Ermon, S.: Geography-aware self-supervised learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10181–10190 (2021)
- [45] Reed, C.J., Gupta, R., Li, S., Brockman, S., Funk, C., Clipp, B., Keutzer, K., Candido, S., Uytendaele, M., Darrell, T.: Scale-mae: A scale-aware masked autoencoder for multiscale geospatial representation learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4088–4099 (2023)
- [46] Cong, Y., Khanna, S., Meng, C., Liu, P., Rozi, E., He, Y., Burke, M., Lobell, D., Ermon, S.: SatMAE: Pre-training transformers for temporal and multi-spectral satellite imagery. In: *Advances in Neural Information Processing Systems*, vol. 35, pp. 197–211 (2022)
- [47] Tseng, G., Zvonkov, I., Purohit, M., Rolnick, D., Kerner, H.: Lightweight, pre-trained transformers for remote sensing timeseries. *arXiv preprint arXiv:2304.14065* (2023)
- [48] Jakubik, J., Roy, S., Phillips, C., Fraccaro, P., Godwin, D., Zadrozny, B., Szwarcman, D., Gomes, C., Nyirjesy, G., Edwards, B., et al.: Foundation models for generalist geospatial artificial intelligence. *arXiv preprint arXiv:2310.18660* (2023)
- [49] Fuller, A., Millard, K., Green, J.: CROMA: Remote sensing representations with contrastive radar-optical masked autoencoders. In: *Advances in Neural Information Processing Systems*, vol. 36 (2024)
- [50] Wang, Y., Albrecht, C.M., Braham, N.A.A., Liu, C., Xiong, Z., Zhu, X.X.: DeCUR: Decoupling common & unique representations for multimodal self-supervision. *arXiv preprint arXiv:2309.05300* (2023)
- [51] Xiong, Z., Wang, Y., Zhang, F., Zhu, X.X.: One for all: Toward unified foundation models for Earth vision. *arXiv preprint arXiv:2401.07527* (2024)
- [52] Xiong, Z., Wang, Y., Zhang, F., Stewart, A.J., Hanna, J., Borth, D., Papoutsis, I., Saux, B.L., Camps-Valls, G., Zhu, X.X.: Neural plasticity-inspired foundation model for observing the Earth crossing modalities. *arXiv preprint arXiv:2403.15356* (2024)
- [53] Liu, F., Chen, D., Guan, Z., Zhou, X., Zhu, J., Zhou, J.: RemoteClip: A vision language foundation model for remote sensing. *arXiv preprint arXiv:2306.11029* (2023)

- [54] Wang, Z., Prabha, R., Huang, T., Wu, J., Rajagopal, R.: SkyScript: A large and semantically diverse vision-language dataset for remote sensing. *arXiv preprint arXiv:2312.12856* (2023)
- [55] Zeng, C., Bertsimas, D.: Global flood prediction: a multimodal machine learning approach. In: *ICLR Workshop on Tackling Climate Change with Machine Learning* (2023)
- [56] Zhang, W., Cai, M., Zhang, T., Zhuang, Y., Mao, X.: EarthGPT: A universal multi-modal large language model for multi-sensor image comprehension in remote sensing domain. *arXiv preprint arXiv:2401.16822* (2024)
- [57] Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., et al.: FourCastNet: A global data-driven high-resolution weather model using adaptive Fourier neural operators. *arXiv preprint arXiv:2202.11214* (2022)
- [58] Wedi, N., Bauer, P., Denoninck, W., Diamantakis, M., Hamrud, M., Kuhnlein, C., Malardel, S., Mogensen, K., Mozdzyński, G., Smolarkiewicz, P.: The modelling infrastructure of the Integrated Forecasting System: Recent advances and future challenges (2015)
- [59] Han, T., Guo, S., Ling, F., Chen, K., Gong, J., Luo, J., Gu, J., Dai, K., Ouyang, W., Bai, L.: FengWu-GHR: Learning the kilometer-scale medium-range global weather forecasting. *arXiv preprint arXiv:2402.00059* (2024)
- [60] Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., Tian, Q.: Pangu-Weather: A 3D high-resolution model for fast and accurate global weather forecast. *arXiv preprint arXiv:2211.02556* (2022)
- [61] Weyn, J.A., Durran, D.R., Caruana, R., Cresswell-Clay, N.: Sub-seasonal forecasting with a large ensemble of deep-learning weather prediction models. *Journal of Advances in Modeling Earth Systems* **13**(7), 2021–002502 (2021)
- [62] Watt-Meyer, O., Dresdner, G., McGibbon, J., Clark, S.K., Henn, B., Duncan, J., Brenowitz, N.D., Kashinath, K., Pritchard, M.S., Bonev, B., et al.: ACE: A fast, skillful learned global atmospheric model for climate prediction. *arXiv preprint arXiv:2310.02074* (2023)
- [63] Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J.K., Grover, A.: ClimaX: A foundation model for weather and climate. *arXiv preprint arXiv:2301.10343* (2023)
- [64] Webersinke, N., Kraus, M., Bingler, J.A., Leippold, M.: ClimateBert: A pretrained language model for climate-related text. *arXiv preprint arXiv:2110.12010* (2021)
- [65] Thulke, D., Gao, Y., Pelser, P., Brune, R., Jalota, R., Fok, F., Ramos, M., Wyk, I., Nasir, A., Goldstein, H., et al.: ClimateGPT: Towards AI synthesizing interdisciplinary research on climate change. *arXiv preprint arXiv:2401.09646* (2024)
- [66] Ben-Bouallegue, Z., Clare, M.C., Magnusson, L., Gascon, E., Maier-Gerber, M., Janousek, M., Rodwell, M., Pinault, F., Dramsch, J.S., Lang, S.T., et al.: The rise of data-driven weather forecasting. *arXiv preprint arXiv:2307.10128* (2023)
- [67] Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., Tian, Q.: Accurate medium-range global weather forecasting with 3D neural networks. *Nature* **619**(7970), 533–538 (2023)
- [68] Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., et al.: Learning skillful medium-range global weather forecasting. *Science* **382**(6677), 1416–1421 (2023)
- [69] Kochkov, D., Yuval, J., Langmore, I., Norgaard, P., Smith, J., Mooers, G., Lottes, J., Rasp, S., Düben, P., Klöwer, M., et al.: Neural general circulation models. *arXiv preprint arXiv:2311.07222* (2023)
- [70] Xiong, Z., Yuan, Y., Wang, Q.: MSN: Modality separation networks for RGB-D scene recognition. *Neurocomputing* **373**, 81–89 (2020)

- [71] Mendieta, M., Han, B., Shi, X., Zhu, Y., Chen, C.: Towards geospatial foundation models via continual pretraining. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 16806–16816 (2023)
- [72] Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., *et al.*: Overcoming catastrophic forgetting in neural networks. Proceedings of the National Academy of Sciences **114**(13), 3521–3526 (2017)
- [73] Snelson, E., Ghahramani, Z.: Sparse Gaussian processes using pseudo-inputs. In: Advances in Neural Information Processing Systems, vol. 18 (2005)
- [74] Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S., Yang, L.: Physics-informed machine learning. Nature Reviews Physics **3**(6), 422–440 (2021)
- [75] Xu, Q., Shi, Y., Bamber, J., Ouyang, C., Zhu, X.X.: Large-scale flood modeling and forecasting with FloodCast. arXiv preprint arXiv:2403.12226 (2024)
- [76] Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. In: Advances in Neural Information Processing Systems, vol. 36 (2024)
- [77] Kuo, W., Cui, Y., Gu, X., Piergiovanni, A., Angelova, A.: F-VLM: Open-vocabulary object detection upon frozen vision and language models. arXiv preprint arXiv:2209.15639 (2022)
- [78] Paaß, G., Giesselbach, S.: Knowledge acquired by foundation models. In: Foundation Models for Natural Language Processing: Pre-trained Language Models Integrating Media, pp. 161–185 (2023)
- [79] Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., *et al.*: DINOv2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)
- [80] Yeh, C., Meng, C., Wang, S., Driscoll, A., Rozi, E., Liu, P., Lee, J.J., Burke, M., Lobell, D.B., Ermon, S.: SustainBench: Benchmarks for monitoring the sustainable development goals with machine learning. In: Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (2021)
- [81] Lacoste, A., Lehmann, N., Rodriguez, P., Sherwin, E., Kerner, H., Lütjens, B., Irvin, J., Dao, D., Alemohammad, H., Drouin, A., *et al.*: Geo-Bench: Toward foundation models for Earth monitoring. In: Advances in Neural Information Processing Systems, vol. 36 (2024)
- [82] Bountos, N.I., Ouaknine, A., Rolnick, D.: FoMo-Bench: a multi-modal, multi-scale and multi-task forest monitoring benchmark for remote sensing foundation models. arXiv preprint arXiv:2312.10114 (2023)
- [83] Hu, J., Liu, R., Hong, D., Camero, A., Yao, J., Schneider, M., Kurz, F., Segl, K., Zhu, X.X.: MDAS: A new multimodal benchmark dataset for remote sensing. Earth System Science Data **15**(1), 113–131 (2023)
- [84] Fibaek, C., Camilleri, L., Luyts, A., Dionelis, N., Saux, B.L.: PhilEO Bench: Evaluating geo-spatial foundation models. arXiv preprint arXiv:2401.04464 (2024)
- [85] Rasp, S., Dueben, P.D., Scher, S., Weyn, J.A., Mouatadid, S., Thuerey, N.: WeatherBench: A benchmark data set for data-driven weather forecasting. Journal of Advances in Modeling Earth Systems **12**(11), 2020–002203 (2020)
- [86] Rasp, S., Hoyer, S., Meroze, A., Langmore, I., Battaglia, P., Russel, T., Sanchez-Gonzalez, A., Yang, V., Carver, R., Agrawal, S., *et al.*: WeatherBench 2: A benchmark for the next generation of data-driven global weather models. arXiv preprint arXiv:2308.15560 (2023)
- [87] Racah, E., Beckham, C., Maharaj, T., Ebrahimi Kahou, S., Prabhat, M., Pal, C.: ExtremeWeather:

- A large-scale climate dataset for semi-supervised detection, localization, and understanding of extreme weather events. In: *Advances in Neural Information Processing Systems*, vol. 30 (2017)
- [88] Nguyen, T., Jewik, J., Bansal, H., Sharma, P., Grover, A.: ClimateLearn: Benchmarking machine learning for weather and climate modeling. In: *Advances in Neural Information Processing Systems*, vol. 36 (2024)
 - [89] Watson-Parris, D., Rao, Y., Olivié, D., Seland, Ø., Nowack, P., Camps-Valls, G., Stier, P., Bouabid, S., Dewey, M., Fons, E., Gonzalez, J., Harder, P., Jeggle, K., Lenhardt, J., Manshausen, P., Novitasari, M., Ricard, L., Roesch, C.: ClimateBench v1.0: A benchmark for data-driven climate projections. *Journal of Advances in Modeling Earth Systems* **14**(10) (2022)
 - [90] Zhang, J., Howard, K., Langston, C., Kaney, B., Qi, Y., Tang, L., Grams, H., Wang, Y., Cocks, S., Martinaitis, S., Arthur, A., Cooper, K., Brogden, J., Kitzmiller, D.: Multi-radar multi-sensor (MRMS) quantitative precipitation estimation: Initial operating capabilities. *Bulletin of the American Meteorological Society* **97**(4), 621–638 (2016)
 - [91] Chen, L., Zhong, X., Zhang, F., Cheng, Y., Xu, Y., Qi, Y., Li, H.: FuXi: A cascade machine learning forecasting system for 15-day global weather forecast. *arXiv preprint arXiv:2306.12873* (2023)
 - [92] Price, I., Sanchez-Gonzalez, A., Alet, F., Ewalds, T., El-Kadi, A., Stott, J., Mohamed, S., Battaglia, P., Lam, R., Willson, M.: GenCast: Diffusion-based ensemble forecasting for medium-range weather. *arXiv preprint arXiv:2312.15796* (2023)
 - [93] Sønderby, C.K., Espeholt, L., Heek, J., Dehghani, M., Oliver, A., Salimans, T., Agrawal, S., Hickey, J., Kalchbrenner, N.: MetNet: A neural weather model for precipitation forecasting. *arXiv preprint arXiv:2003.12140* (2020)
 - [94] Kikaki, K., Kakogeorgiou, I., Mikeli, P., Raitos, D.E., Karantzaos, K.: MARIDA: A benchmark for marine debris detection from Sentinel-2 remote sensing data. *PLOS ONE* **17**(1), 1–20 (2022)
 - [95] Gourmelon, N., Seehaus, T., Braun, M., Maier, A., Christlein, V.: Calving fronts and where to find them: a benchmark dataset and methodology for automatic glacier calving front extraction from synthetic aperture radar imagery. *Earth System Science Data* **14**(9), 4287–4313 (2022)
 - [96] Cao, Y., Yang, J.: Towards making systems forget with machine unlearning. In: *IEEE Symposium on Security and Privacy*, pp. 463–480 (2015)
 - [97] Feldman, V.: Does learning require memorization? A short tale about a long tail. In: *Proceedings of the Annual ACM SIGACT Symposium on Theory of Computing*, pp. 954–959 (2020)
 - [98] Xu, Y., Ghamisi, P.: Universal adversarial examples in remote sensing: Methodology and benchmark. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–15 (2022)
 - [99] Zhu, L., Liao, B., Zhang, Q., Wang, X., Liu, W., Wang, X.: Vision Mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417* (2024)
 - [100] Chen, R., Liu, Y., Kong, L., Chen, N., Zhu, X., Ma, Y., Liu, T., Wang, W.: Towards label-free scene understanding by vision foundation models. In: *Advances in Neural Information Processing Systems*, vol. 36, pp. 75896–75910 (2023)
 - [101] Truhn, D., Eckardt, J.-N., Ferber, D., Kather, J.N.: Large language models and multimodal foundation models for precision oncology. *NPJ Precision Oncology* **8**(1), 1–4 (2024)
 - [102] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., *et al.*: Retrieval-augmented generation for knowledge-intensive NLP tasks. In: *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474 (2020)
 - [103] PCMDI: CMIP6 - Coupled Model Intercomparison Project Phase 6. <https://pcmdi.llnl.gov/>

- [104] Srivastava, S., Sharma, G.: OmniVec: Learning robust representations with cross modal sharing. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 1236–1248 (2024)
- [105] Li, J., Li, D., Savarese, S., Hoi, S.: BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In: International Conference on Machine Learning, pp. 19730–19742 (2023)
- [106] Wang, W., Lv, Q., Yu, W., Hong, W., Qi, J., Wang, Y., Ji, J., Yang, Z., Zhao, L., Song, X., et al.: CogVLM: Visual expert for pretrained language models. arXiv preprint arXiv:2311.03079 (2023)
- [107] Zhang, Z., Zhao, T., Guo, Y., Yin, J.: RS5M: A large scale vision-language dataset for remote sensing vision-language foundation model. arXiv preprint arXiv:2306.11300 (2023)
- [108] Lu, X., Wang, B., Zheng, X., Li, X.: Exploring models and data for remote sensing image caption generation. IEEE Transactions on Geoscience and Remote Sensing **56**(4), 2183–2195 (2018)
- [109] Qu, B., Li, X., Tao, D., Lu, X.: Deep semantic understanding of high resolution remote sensing image. In: International Conference on Computer, Information and Telecommunication Systems, pp. 1–5 (2016)
- [110] Cheng, Q., Huang, H., Xu, Y., Zhou, Y., Li, H., Wang, Z.: NWPU-Captions dataset and MLCA-Net for remote sensing image captioning. IEEE Transactions on Geoscience and Remote Sensing **60**, 1–19 (2022)
- [111] Zheng, X., Wang, B., Du, X., Lu, X.: Mutual attention inception network for remote sensing visual question answering. IEEE Transactions on Geoscience and Remote Sensing **60**, 1–14 (2022)
- [112] Lobry, S., Marcos, D., Murray, J., Tuia, D.: RSVQA: Visual question answering for remote sensing data. IEEE Transactions on Geoscience and Remote Sensing **58**(12), 8555–8566 (2020)
- [113] Lobry, S., Demir, B., Tuia, D.: RSVQA Meets BigEarthNet: A new, large-scale, visual question answering dataset for remote sensing. In: 2021 IEEE International Geoscience and Remote Sensing Symposium, pp. 1218–1221 (2021)
- [114] Sun, Y., Feng, S., Li, X., Ye, Y., Kang, J., Huang, X.: Visual grounding in remote sensing images. In: Proceedings of the 30th ACM International Conference on Multimedia, pp. 404–412 (2022)
- [115] Zhan, Y., Xiong, Z., Yuan, Y.: RSVG: Exploring data and models for visual grounding on remote sensing data. IEEE Transactions on Geoscience and Remote Sensing **61**, 1–13 (2023)
- [116] Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: LAION-5B: An open large-scale dataset for training next generation image-text models. In: Advances in Neural Information Processing Systems, vol. 35, pp. 25278–25294 (2022)
- [117] Yuan, Z., Xiong, Z., Mou, L., Zhu, X.X.: ChatEarthNet: A global-scale image-text dataset empowering vision-language geo-foundation models. arXiv preprint arXiv:2402.11325 (2024)
- [118] OpenAI: GPT-4V(ision) system card. <https://openai.com/research/gpt-4v-system-card>
- [119] Azavea/Element 84, Robert Cheetham: Raster Vision: An open source library and framework for deep learning on satellite and aerial imagery (2017-2023). <https://doi.org/10.5281/zenodo.8018177>
- [120] Stewart, A.J., Robinson, C., Corley, I.A., Ortiz, A., Lavista Ferres, J.M., Banerjee, A.: TorchGeo: Deep learning with geospatial data. In: Proceedings of the International Conference on Advances in Geographic Information Systems, pp. 1–12 (2022)

- [121] Slonecker, E.T., Shaw, D.M., Lillesand, T.M.: Emerging legal and ethical issues in advanced remote sensing technology. *Photogrammetric Engineering and Remote Sensing* **64**(6), 589–595 (1998)
- [122] Wang, D., Zhang, Q., Xu, Y., Zhang, J., Du, B., Tao, D., Zhang, L.: Advancing plain vision transformer toward remote sensing foundation model. *IEEE Transactions on Geoscience and Remote Sensing* **61**, 1–15 (2022)
- [123] Hong, D., Zhang, B., Li, X., Li, Y., Li, C., Yao, J., Yokoya, N., Li, H., Jia, X., Plaza, A., et al.: SpectralGPT: Spectral foundation model. *arXiv preprint arXiv:2311.07113* (2023)
- [124] Keisler, R.: Forecasting global weather with graph neural networks. *arXiv preprint arXiv:2202.07575* (2022)
- [125] Bonev, B., Kurth, T., Hundt, C., Pathak, J., Baust, M., Kashinath, K., Anandkumar, A.: Spherical Fourier neural operators: Learning stable dynamics on the sphere. *arXiv preprint arXiv:2306.03838* (2023)
- [126] Chen, K., Han, T., Gong, J., Bai, L., Ling, F., Luo, J.-J., Chen, X., Ma, L., Zhang, T., Su, R., et al.: FengWu: Pushing the skillful global medium-range weather forecast beyond 10 days lead. *arXiv preprint arXiv:2304.02948* (2023)
- [127] Man, X., Zhang, C., Li, C., Shao, J.: W-MAE: Pre-trained weather model with masked autoencoder for multi-variable weather forecasting. *arXiv preprint arXiv:2304.08754* (2023)
- [128] Li, L., Carver, R., Lopez-Gomez, I., Sha, F., Anderson, J.: SEEDS: Emulation of weather forecast ensembles with diffusion models. *arXiv preprint arXiv:2306.14066* (2023)
- [129] Nguyen, T., Shah, R., Bansal, H., Arcomano, T., Madireddy, S., Maulik, R., Kotamarthi, V., Foster, I., Grover, A.: Scaling transformer neural networks for skillful and reliable medium-range weather forecasting. *arXiv preprint arXiv:2312.03876* (2023)
- [130] Yuan, Z., Zhang, W., Fu, K., Li, X., Deng, C., Wang, H., Sun, X.: Exploring a fine-grained multiscale method for cross-modal remote sensing image retrieval. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–19 (2021)
- [131] Hu, Y., Yuan, J., Wen, C., Lu, X., Li, X.: RSGPT: A remote sensing vision language model and benchmark. *arXiv preprint arXiv:2307.15266* (2023)
- [132] Muhtar, D., Li, Z., Gu, F., Zhang, X., Xiao, P.: LHRS-Bot: Empowering remote sensing with vgi-enhanced large multimodal language model. *arXiv preprint arXiv:2402.02544* (2024)
- [133] Czerkawski, M., Francis, A.: From LAION-5B to LAION-EO: Filtering billions of images using anchor datasets for satellite image extraction. *arXiv preprint arXiv:2309.15535* (2023)
- [134] Rahnemoonfar, M., Chowdhury, T., Sarkar, A., Varshney, D., Yari, M., Murphy, R.R.: FloodNet: A high resolution aerial imagery dataset for post flood scene understanding. *IEEE Access* **9**, 89644–89654 (2021)
- [135] Yuan, Z., Mou, L., Xiong, Z., Zhu, X.X.: Change detection meets visual question answering. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–13 (2022)
- [136] Liu, C., Zhao, R., Chen, H., Zou, Z., Shi, Z.: Remote sensing image change captioning with dual-branch transformers: A new method and a large scale dataset. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–20 (2022)