

GlobalGeoTree: A Multi-Granular Vision-Language Dataset for Global Tree Species Classification

Yang Mu^{1,6}, Zhitong Xiong^{1,6}, Yi Wang¹, Muhammad Shahzad², Franz Essl³, Holger Kreft⁴, Mark van Kleunen⁵, and Xiao Xiang Zhu^{1,6}

¹Technical University of Munich, Munich, Germany

²University of Reading, Reading, UK

³University of Vienna, Vienna, Austria

⁴University of Göttingen, Göttingen, Germany

⁵University of Konstanz, Konstanz, Germany

⁶Munich Center for Machine Learning

Correspondence: Xiao Xiang Zhu (xiaoxiang.zhu@tum.de)

Abstract. Global tree species mapping using remote sensing data is vital for biodiversity monitoring, forest management, and ecological research. However, progress in this field has been constrained by the scarcity of large-scale, labeled datasets. To address this, we introduce GlobalGeoTree, a comprehensive global dataset for tree species classification. GlobalGeoTree comprises 6.3 million geolocated tree occurrences, spanning 275 families, 2,734 genera, and 21,001 species across the hierarchical taxonomic levels. Each sample is paired with Sentinel-2 image time series and 27 auxiliary environmental variables, encompassing bioclimatic, geographic, and soil data. The dataset is partitioned into *GlobalGeoTree-6M*, a large subset for model pretraining, and curated evaluation subsets, primarily *GlobalGeoTree-10kEval*, a benchmark for zero-shot and few-shot classification. To demonstrate the utility of the dataset, we introduce a baseline model, GeoTreeCLIP, which leverages paired remote sensing data and taxonomic text labels within a vision-language framework pretrained on *GlobalGeoTree-6M*. Experimental results show that GeoTreeCLIP achieves substantial improvements in zero- and few-shot classification on *GlobalGeoTree-10kEval* over existing advanced models. By making the dataset, models, and code publicly available, we aim to establish a benchmark to advance tree species classification and foster innovation in biodiversity research and ecological applications. The code is publicly available at <https://github.com/MUYang99/GlobalGeoTree>, and the GlobalGeoTree dataset is available at <https://huggingface.co/datasets/yann111/GlobalGeoTree> (Mu et al., 2025b).

1 Introduction

Forests cover approximately 31% of the global land surface (Hansen et al., 2013) and provide essential ecosystem services, including carbon sequestration (Jenkins et al., 2003), biodiversity conservation (Lindenmayer et al., 2006), and climate regulation (Bonan, 2008). Accurate and large-scale mapping of tree species plays an increasingly vital role in addressing pressing environmental challenges (Mu et al., 2025a), including effective biodiversity monitoring (Felton et al., 2020), informed forest management practices (Franklin, 2001), and comprehensive ecological research aimed at understanding the complex impacts of climate change (Hamann and Wang, 2006).




Figure 1. Overview of the GlobalGeoTree dataset, which includes 6.3 million samples spanning 21,001 tree species across 221 countries/regions. The map illustrates the geographic coverage, with color intensity representing the number of samples in each $1^\circ \times 1^\circ$ latitude/longitude grid. Each sample is paired with remote sensing data, including Sentinel-2 time series, auxiliary environmental variables, and hierarchical taxonomic labels spanning from functional type to species level.

Traditional ground-based forest monitoring methods (Wellbrock et al., 2018), while providing detailed information, are often limited in their spatial and temporal coverage, making it challenging to obtain a comprehensive understanding of global forest composition and dynamics. In contrast, remote sensing has emerged as a key technology for large-scale forest monitoring, offering non-invasive and cost-effective approaches to tree species classification (Hermosilla et al., 2022). Despite significant advancements in this field, progress has been constrained by the limited availability of comprehensive, high-quality, and accurately labeled datasets that capture the global diversity of tree species (Bountos et al., 2025). Existing datasets typically focus on specific geographic regions or limited taxonomic coverage, hampering the development of models with global applicability (Ouaknine et al., 2025).

To bridge these gaps, we present GlobalGeoTree, a large-scale dataset comprising 6.3 million remote sensing samples paired with multi-level taxonomic labels. This dataset integrates time-series satellite imagery from Sentinel-2 with 27 bioclimatic, geographic, and soil variables, offering a rich multimodal representation of tree species within their environmental contexts. The taxonomic hierarchy spans family, genus, and species levels, enabling classification across various scales of biological organization.

In addition to the dataset, we introduce GeoTreeCLIP, a vision-language model specifically designed for tree species classification. Drawing on frameworks like CLIP (Radford et al., 2021), our approach aligns satellite imagery with taxonomic

labels to learn nuanced representations. Unlike traditional classifiers treating labels as discrete categories, vision-language models can inherently process label hierarchical structure, generalizing to unseen species through representations of related genera or families (Stevens et al., 2024). This enables robust zero-shot and few-shot learning of GeoTreeCLIP, which are critical for addressing the vast scale of global biodiversity, ever-evolving species catalogs, and the practical impossibility of exhaustive data collection for all taxa.

GeoTreeCLIP leverages domain-specific pretraining on *GlobalGeoTree-6M*, main part of the dataset tailored for model pretraining, and evaluated on a specialized benchmark, *GlobalGeoTree-10kEval*, which enables a comprehensive assessment of model performance across multiple taxonomic levels. Through the open availability of GlobalGeoTree, its associated models, and evaluation protocols, we seek to establish a community-driven benchmarking standard that will accelerate the development of generalizable models for tree species mapping and deepen our understanding of global forest biodiversity.

2 Related Work

2.1 Open datasets for tree species classification

Table 1 provides an overview of notable open datasets that have contributed to tree species classification, detailing their geographic coverage, size, and taxonomic diversity. For instance, the Seu Nico Forest dataset (Gastauer et al., 2015) from Brazil provides geolocated samples for 228 species but is geographically constrained. Similarly, the Maraca Ecological Station dataset (Farias et al., 2020) includes 110 species but is also region-specific. In Europe, the EUForest dataset (Mauri, 2017) offers broader coverage with data for 242 species. On a global scale, datasets such as Tallo (Jucker et al., 2022) provide significant taxonomic diversity, covering 5,163 species across 187 families. However, these datasets lack integration with remote sensing or environmental variables, limiting their application in ecological modeling.

Table 1. Overview of publicly available datasets for tree species classification.

Dataset	Geographic scope	Size	Classes	Year
Seu Nico Forest (Gastauer et al., 2015)	Brazil	2,868	54 families; 139 genera; 228 species	2015
EUForest (Mauri, 2017)	Europe	588,983	83 genera; 242 species	2017
Maraca Eco. Sta. (Farias et al., 2020)	Brazil	680	40 families; 110 species	2020
TreeSatAI (Ahlswede, 2022)	Germany	50,381	15 genera; 20 species	2022
Tallo (Jucker et al., 2022)	Global	498,839	187 families; 1,453 genera; 5,163 species	2022
Indi. Tree Point Clouds (Weiser et al., 2022)	Germany	1,491	22 species	2022
NEON Veg. Struc. (Kampe et al., 2010)	USA	N/A	949 genera; 2,826 species	2023
PureForest (Gaydon and Roche, 2025)	France	135,569	18 species	2024
Planted (Pazos-Outón et al., 2024)	Global	2,264,747	46 genera; 40 species	2024
GlobalGeoTree	Global	6,263,345	275 families; 2,734 genera; 21,001 species	2025

Advances in high-resolution imaging and lidar technologies have enabled datasets like PureForest (Gaydon and Roche, 2025) and Individual Tree Point Clouds (Weiser et al., 2022), which utilize aerial and point cloud data for species classification. While these datasets offer detailed structural information, they remain region-specific and lack the spectral and temporal depth of satellite-based datasets. The TreeSatAI dataset (Ahlsweide, 2022) combines multi-sensor data, including aerial imagery and Sentinel-1/2, for tree species classification in Germany but covers only 20 species. Similarly, the Planted dataset (Pazos-Outón et al., 2024) focuses on only 40 planted species globally, limiting its broader applicability. Our work also complements broader biodiversity benchmarks like GeoLifeCLEF (Botella et al., 2025) and GeoPlant (Picek et al., 2024). While these focus on regional, multi-lifeform plant prediction, GlobalGeoTree is the first dataset with a global scope specifically curated for over 21,000 forest tree species.

Collectively, while valuable, these existing datasets highlight the persistent need for a benchmark that synergizes global coverage, deep taxonomic information for forest tree species, and multimodal remote sensing data, a gap GlobalGeoTree aims to fill.

2.2 Vision-language models for remote sensing applications

Vision-language models (VLMs) enable the integration of visual and textual information. Among these, Contrastive Language-Image Pretraining (CLIP) (Radford et al., 2021) has demonstrated exceptional zero-shot transfer capabilities by jointly training image and text encoders through a contrastive learning objective, aligning image-text pairs within a shared embedding space. For tree species classification, CLIP’s ability to learn from image-text pairings (such as satellite imagery and taxonomic labels) offers a path to capture complex visual and semantic relationships. Its proven zero-shot capabilities are particularly suited for addressing the challenges of identifying species within dynamic and evolving catalogs (Stevens et al., 2024). Meanwhile, its few-shot capabilities tackle the issue of limited labeled data, a common obstacle in biodiversity research, offering an advantage over traditional supervised methods.

In remote sensing, VLMs have been applied to various tasks, including image classification, retrieval, and scene understanding, with domain-specific adaptations yielding significant improvements. For example, RemoteCLIP (Liu et al., 2024), the first VLM specifically tailored for remote sensing, leverages pretraining on large-scale remote sensing imagery paired with aligned text, achieving state-of-the-art performance in zero-shot classification, linear probing, and few-shot learning. Similarly, SkyCLIP (Wang et al., 2024), SkyCLIP-50 (Wang et al., 2024) and CLIP-laion-RS (He et al., 2024) extend the capabilities of CLIP through continual pretraining on semantically diverse remote sensing image-text pairs. These models demonstrate enhanced generalization and transferability, achieving substantial gains in tasks such as zero-shot scene classification, fine-grained classification, and cross-modal retrieval compared to the original CLIP model.

These advancements underscore the importance of domain-specific pretraining in adapting VLMs for remote sensing applications. Aligning models more closely with the unique characteristics of remote sensing tasks has demonstrated significant potential to advance progress in this field.

3 The GlobalGeoTree dataset

3.1 Geolocated data collection and preprocessing

The GlobalGeoTree dataset provides unprecedented global and taxonomic coverage for tree species classification using remote sensing data, and the collection involved several key steps:

3.1.1 Tree species catalog construction

We constructed a comprehensive tree species catalog by integrating two major global repositories: TreeGOER (Kindt, 2023) and GlobalTreeSearch (Beech et al., 2017), containing 48,129 and 57,681 tree species respectively. This compilation was further enriched with multiple open-source datasets documented in Table 1. The taxonomic framework was subsequently validated and expanded using the Global Biodiversity Information Facility (GBIF) Species API (Global Biodiversity Information Facility (GBIF), 2025), ensuring nomenclatural consistency and accuracy. The resulting catalog encompasses 87,845 species, representing the global diversity of tree species.

3.1.2 Geolocation sampling

For each tree species in our catalog, we queried the GBIF Occurrence API (Chamberlain and Boettiger, 2017) to retrieve global geolocations with documented occurrences. To ensure data quality and reliability, we applied strict filtering criteria, including: (1) selecting only recent observations recorded between 2015 and 2024; (2) limiting data to human observation records; (3) excluding records with geospatial issues as flagged by GBIF (e.g., country-coordinate mismatches); (4) filtering for occurrences with a "present" status; and (5) removing duplicate entries and observations with low geographic precision. Additionally, we ensured all samples conform to open data licenses (CC0 1.0, CC-BY-4.0, etc.), maintaining the dataset's accessibility and reusability for the broader research community.

3.1.3 Forest layer filtering

To ensure that each geolocated observation corresponds to a valid tree, we performed an additional forest cover verification step. We utilized the EC JRC Global Map of Forest Cover 2020 (Version 2) (Bourgoin et al., 2025), which has a 10m spatial resolution. Each geolocated point from GBIF was cross-referenced with this map, and only samples located within forest areas were retained. This filtering not only served as an essential quality control measure to enhance data reliability but also defined 2020 as the target year for our study. This allowed us to subsequently acquire the Sentinel-2 time series for 2020, ensuring a direct temporal correspondence between the verified ground observations and the satellite imagery. While we utilized the binary JRC mask for this study, future operational pipelines could integrate multi-class land cover products such as GLC_FCS10 (Zhang et al., 2025) to constrain candidate species based on mapped forest types.

3.2 Dataset Composition and Structure

The filtering process results in the final GlobalGeoTree dataset, which comprises 6,263,345 high-quality samples distributed across 221 countries and regions. Approximately 70% of the records are sourced from "Research Grade" iNaturalist observations, which require identification consensus from at least two independent contributors. Other samples are mainly from authoritative datasets available through GBIF, including National Forest Inventory (NFI) data from countries such as Sweden, France, and Colombia. The full list of source datasets and their provenance is traceable via the unique GBIF Derived Dataset DOI of GlobalGeoTree: <https://doi.org/10.15468/dd.9qxqyy>. A core feature of this dataset is its multi-granular taxonomic hierarchy, which is essential for developing models that can generalize across different levels of biological classification. This hierarchical structure is visualized in Fig. 2, illustrating the relationships branching from four broad functional types at the center (level 0), through 275 families (level 1) and 2,734 genera (level 2), out to the 21,001 individual species at the periphery (level 3).




Figure 2. The taxonomic hierarchy of the GlobalGeoTree dataset. The visualization shows the nested relationships, branching from the four functional types (e.g., *Deciduous Broadleaf*) at the center, through families (e.g., *Fabaceae*), genera, and out to the 21,001 species at the outermost ring. This multi-level structure is a core feature that enables multi-granular classification tasks.

Each of the 6.3 million samples is structured as a single record that includes this full taxonomic lineage, its geolocation, and source metadata. This entire collection is made available as a downloadable CSV file (`GlobalGeoTree.csv`) for full transparency and accessibility. Table 2 presents several example records to illustrate the data structure. A comprehensive statistical breakdown and accessibility details of the dataset is provided in Appendix A.

Table 2. Example records from the `GlobalGeoTree.csv` file, demonstrating the structure of each geolocated data point.

country_code	level0	level1_family	level2_genus	level3_species	source	longitude	latitude
ES	Deciduous Broadleaf	Fagaceae	<i>Quercus</i>	<i>Quercus rotundifolia</i>	iNaturalist Res.-grade Obs.	-3.74024	40.49546
US	Evergreen Needleleaf	Cupressaceae	<i>Thuja</i>	<i>Thuja plicata</i>	iNaturalist Res.-grade Obs.	-122.168	48.33392
IE	Deciduous Broadleaf	Fagaceae	<i>Castanea</i>	<i>Castanea sativa</i>	iNaturalist Res.-grade Obs.	-9.524	52.05556
PT	Deciduous Broadleaf	Fagaceae	<i>Quercus</i>	<i>Quercus rotundifolia</i>	Forestry Inventory 2015	-7.85145	39.42701
AU	Deciduous Broadleaf	Fabaceae	<i>Acacia</i>	<i>Acacia platycarpa</i>	ALA species sightings	132.9895	-14.818

3.3 Paired remote sensing data

Each sample in the GlobalGeoTree dataset is paired with a rich set of multimodal data to enable robust modeling. Table 3 provides a comprehensive overview of these features, including paired Earth Observation (EO) data and auxiliary environmental variables derived from remote sensing sources.

3.3.1 EO data

The Sentinel-2 data for each sample consists of a time series of 12 monthly median composites from January to December 2020. The full-year temporal coverage enables models to capture the distinct seasonal patterns of different tree species, such as leaf-out, full canopy, and senescence (Mu et al., 2025a). These phenological features are particularly discriminative for deciduous species.

For each month, all L2A Sentinel-2 images with less than 30% cloud cover were collected, and the median composites were generated from these images. This process ensures a robust, largely cloud-free representation for each month, effectively mitigating issues caused by transient atmospheric conditions in single-date imagery.

Each composite includes a 5×5 pixel patch centered on the geolocation of the tree species. This patch size was chosen to account for typical crown sizes (10–30 m) (Jucker et al., 2022) and to align with other public datasets, such as PureForest (Gaydon and Roche, 2025) and TreeSatAI (Ahlsweide, 2022). The chosen 5×5 patch size represents a trade-off: it is sufficiently large enough to capture the crown of a mature tree and its immediate context, while remaining compact enough to minimize noise from surrounding unrelated species.

3.3.2 Auxiliary data

The auxiliary data enriches the EO data for each sample by providing additional contextual environmental information. Geographic variables, such as elevation, slope, and aspect, are derived from the USGS SRTM dataset (Jarvis et al., 2008), while

Table 3. Overview of features in each sample in the GlobalGeoTree dataset.

Feature Name	Type	Description
Remote Sensing Data		
Sentinel-2 Time Series	float	12 monthly median composites; Includes RGB, NIR, Vegetation Red Edge, and SWIR bands; dimensions: (12, 10, 5, 5).
Geographic Variables	float	Latitude and longitude, as well as elevation, slope, and aspect derived from USGS (SRTM) (30m resolution).
Soil Variables	float	3 Volumetric Water Content data at 33kPa across depths: 0-5 cm, 5-15 cm, 15-30 cm (250m resolution).
Bioclim Variables	float	19 climatic variables from WorldClim (1km resolution).
Text Labels		
level0	string	Functional type of the species (e.g., Evergreen Broadleaf).
level1_family	string	Taxonomic family of the species (e.g., Myrtaceae).
level2_genus	string	Taxonomic genus of the species (e.g., <i>Eucalyptus</i>).
level3_species	string	Scientific name of the species (e.g., <i>Eucalyptus kybeanensis</i>).
Meta Data		
location	string	Geographic location of the sample (e.g., Australia).
country_code	string	ISO country code of the sample location.
source	string	Source of the sample record (e.g., iNaturalist Research-grade Observations).
species_key	float	Unique identifier for the species in the GBIF database.
record_year	int	The year when the record was collected.

soil data, including volumetric water content at 33 kPa across three depths (0–5 cm, 5–15 cm, 15–30 cm), are obtained from SoilGrid (Poggio et al., 2021). Additionally, 19 bioclimatic variables are sourced from WorldClim (Fick and Hijmans, 2017). Due to the relatively coarse spatial resolution of these datasets (ranging from 30m to 1km), only the values corresponding to the exact coordinates of each occurrence are extracted to ensure precision and relevance (Gillespie et al., 2024). A detailed description of all 27 auxiliary variables and their sources is provided in Appendix B.

3.4 Dataset partitioning

For effective model development and evaluation, the GlobalGeoTree dataset was partitioned into *GlobalGeoTree-6M* and curated evaluation subsets, primarily *GlobalGeoTree-10kEval*.

GlobalGeoTree-6M comprises the vast majority of the samples and is specifically designed for model pretraining. This large size allows models to learn robust and generalizable representations of tree species and their associated environmental contexts.




Figure 3. Species in GlobalGeoTree are categorized into Frequent, Common and Rare groups based on the number of samples per species.

GlobalGeoTree-10kEval is a carefully curated dataset designed to benchmark model performance across taxonomic levels and species frequency categories. To address the characteristic long-tail distribution (detailed in Appendix A3), we categorized species into three groups based on sample frequency: Frequent (> 1500 samples), Common (100–1500 samples), and Rare (< 100 samples), as shown in Fig. 3.




Figure 4. Geographic distribution of *GlobalGeoTree-10kEval*. This benchmark includes species selected from Frequent, Common, and Rare categories, as described in the text.

The primary *GlobalGeoTree-10kEval* dataset includes 30 species from each of these three categories, resulting in a total of 90 species. The sample proportions within this evaluation set are 12% for Rare species, 33% for Common species, and 55% for Frequent species, culminating in around 10,000 samples. Fig. 4 shows the geographical distribution of *GlobalGeoTree-10kEval*, which spans diverse regions across the globe. This global distribution ensures that the dataset captures a wide range of ecological and environmental contexts, making it representative of real-world scenarios. By focusing on a diverse set of species with varying levels of representation, *GlobalGeoTree-10kEval* serves as a robust evaluation benchmark for assessing the ability of models to tackle challenges posed by the long-tail distribution of species and shifts in geographical domains.

To further evaluate model robustness and scalability across broader taxonomic scopes, we constructed two additional evaluation subsets: *GlobalGeoTree-10kEval-300* and *GlobalGeoTree-10kEval-900*, containing 100 and 300 species per category, respectively. Crucially, all samples within these evaluation sets, were subjected to the rigorous filtering criteria detailed in Sect. 3.1.2 and are strictly excluded from the *GlobalGeoTree-6M* pretraining set to ensure fair evaluation. Details of all evaluation subsets (Appendix C) and the corresponding evaluation results (Appendix D) are also provided. Given the complexity of global tree species classification, our primary analysis focuses on the 90-species *GlobalGeoTree-10kEval*, which serves as a practical starting point for systematic benchmarking.

4 Data Quality and Validation

To ensure the reliability of GlobalGeoTree, we implemented a multi-tiered validation framework combining fine-scale visual inspection, global-scale cross-verification with independent land cover products, and a public platform for community auditing.

4.1 Visual Inspection and Community Validation

We first quantified land-cover accuracy through a visual validation study. A random subset of 300 locations was sampled from *GlobalGeoTree-10kEval*. For each sample, we retrieved Very High Resolution (VHR) satellite imagery to verify whether the geolocated point fell within forest or tree cover. The inspection confirmed that 98.3% of the samples were correctly located in forested areas, while a small minority were ambiguous (e.g., forest edges), validating the efficacy of our automated filtering pipeline.

To further enhance transparency and enable large-scale qualitative validation by the research community, we developed the GlobalGeoTree Explorer, a web-based application hosted on Google Earth Engine (Fig. 5). This tool allows users to visualize the spatial distribution of all 6.3 million samples, overlay them on high-resolution satellite basemaps, and inspect detailed taxonomic attributes for any individual point. The application is publicly accessible at <https://ee-yangm.projects.earthengine.app/view/globalgeotree-explorer>.

4.2 Cross-Validation with Global Land Cover Products

To validate taxonomic consistency at the functional type level, we performed a cross-comparison against the Copernicus Global Land Cover Layers (CGLS-LC100, Collection 3) (Buchhorn et al., 2020). Since no global dataset currently provides




Figure 5. The GlobalGeoTree Explorer App. The interface allows users to visualize the global distribution of tree samples (bottom map) and inspect individual points against high-resolution satellite imagery (top panels a-f). Detailed attributes for selected samples are displayed in the sidebar, facilitating transparent community validation.

high-resolution distribution maps at the species or genus level, the CGLS-LC100 product represents the most detailed global baseline available, offering discrete land cover layers at 100m resolution that distinguish between broadleaf and needleleaf forest types.

We aligned the `level0` labels of GlobalGeoTree with CGLS classes by aggregating them into two primary leaf types: Needleleaf and Broadleaf. To ensure global representativeness and mitigate regional data density biases, we employed a stratified sampling strategy, selecting 10,000 samples per continent.

Figure 6 presents the results of this cross-validation. The confusion matrix (left) reveals an overall agreement of 80.38% at the leaf level. We observed high consistency for Broadleaf samples (81%), while some Needleleaf samples (22%) showed discrepancies with the coarser 100m product, likely due to mixed pixels in transition zones. The regional breakdown (right) demonstrates exceptionally high agreement in South America (99.3%) and Africa (98.7%), where broadleaf forests dominate. Lower agreement rates in North America and Europe are expected given the higher prevalence of mixed forests and the resolution gap between our point-based data and the 100m validation raster.




Figure 6. Cross-validation results against the CGLS-LC100 product. **Left:** Normalized confusion matrix at the Leaf Level (Needleleaf vs. Broadleaf), showing an overall agreement of 80.38%. **Right:** Agreement rates broken down by continent, highlighting strong consistency in tropical regions.

4.3 Taxonomic Reliability

Beyond spatial validation, taxonomic reliability is ensured through our strict data provenance pipeline. We exclusively utilize observations which require community consensus for species identification, and 100% of these samples have passed the GBIF taxonomic backbone check (GBIF Secretariat, 2023).

5 Benchmarks

5.1 GeoTreeCLIP model

To establish a strong baseline on the GlobalGeoTree dataset, we developed GeoTreeCLIP, a vision-language model (VLM) specifically tailored for tree species classification. Our choice of the VLM paradigm was motivated by the inherent challenges of this task: preliminary experiments showed that traditional supervised models struggle with the vast, long-tailed, and hierarchical label space of GlobalGeoTree, especially at the fine-grained genus and species levels (see Sect. 6.1).

In contrast, a VLM framework treats taxonomic labels not as discrete indices, but as structured text. This allows the model to leverage the rich semantic relationships between species, genera, and families. By jointly training on image-text pairs, GeoTreeCLIP learns powerful, transferable representations that align visual features with their corresponding taxonomic context. This approach is particularly effective for zero-shot and few-shot scenarios, which are critical for addressing the immense scale of global biodiversity and the practical impossibility of collecting exhaustive data for all taxa (Stevens et al., 2024).

As shown in Fig. 7, the GeoTreeCLIP model architecture consists of the following components:




Figure 7. Architecture of the GeoTreeCLIP baseline model. It processes Sentinel-2 time series and auxiliary data through visual encoder and MLP, and hierarchical taxonomic labels through a text encoder. The resulting multimodal visual and text features are then aligned using a contrastive loss.

- **Visual Encoder:** A ViT-B/16 backbone (Dosovitskiy et al., 2020) adapted to process the multi-spectral, time-series Sentinel-2 data. To handle the 10-channel, 5x5 pixel input patches, we modify the initial patch embedding layer of the ViT. Each 5x5 patch is first upsampled to 16x16 pixels using bicubic interpolation. The modified embedding layer then processes these 10 channels and projects the entire 16x16 patch into a single visual token. This approach allows the model to capture the holistic spectral information of the patch at each timestep. These 12 monthly tokens are then fed into a temporal attention module (Vaswani et al., 2017) to learn and fuse key phenological patterns across the year. This mechanism allows the model to dynamically weigh the importance of different months, implicitly learning which parts of the phenological cycle are most informative for identifying a particular species.
- **Auxiliary Feature Integration:** A multi-layer perceptron (MLP) (Rosenblatt, 1958; Gillespie et al., 2024) designed to process bioclimatic, soil, and geographic data. The MLP consists of several feature-specific linear layers (hidden dimension: 256) followed by LayerNorm and ReLU activation. The encoded features are then fused and projected into a 768-dimensional embedding, which is added to the visual token embedding from the Visual Encoder before the final contrastive projection. This allows the model to integrate environmental context with visual information. We adopt a dual-branch architecture comprising a Visual Encoder and a separate MLP for auxiliary features, rather than stacking all inputs into a single data cube. This design addresses the fundamental heterogeneity of the data modalities. First, Sentinel-2 data consists of a dynamic 12-month time series capturing phenology, whereas auxiliary variables such as bioclimatic and topographic metrics are static. Integrating these static variables into the visual branch would require replicating them across the temporal dimension, introducing significant data redundancy. Second, the spatial resolution of auxiliary

variables is typically much coarser ($\sim 1\text{km}$ or 250m) than the 10m Sentinel-2 imagery. Treating them as image bands would necessitate upsampling, creating spatially invariant feature maps that are computationally inefficient for the Vision Transformer. By processing these modalities separately, the MLP branch efficiently encodes environmental context as a conditional prior, which is subsequently fused with spatiotemporal visual features in the shared embedding space.

- **Text Encoder:** A 77-token causal autoregressive transformer (Vaswani et al., 2017; Radford et al., 2019) that processes the hierarchical taxonomic labels. For each sample, the text input is structured to convey the full taxonomic lineage, for example, "Evergreen Broadleaf, Family Myrtaceae, Genus Eucalyptus, Species Eucalyptus kybeanensis". By encoding this structured text, the model learns embeddings that capture the relationships between different taxonomic levels. This allows GeoTreeCLIP to understand, for instance, that two different species within the same genus (*Eucalyptus*) are semantically closer than species from different families. This learned semantic structure is key to its strong generalization capabilities, especially for rare or unseen species.

5.1.1 Pretraining details

Both the Visual Encoder and Text Encoder are initialized using OpenAI’s publicly available CLIP checkpoint (Radford et al., 2021) and further pre-trained on the *GlobalGeoTree-6M* using Distributed Data Parallel (DDP) across 5 NVIDIA 3090 (24GB) GPUs. We employed a batch size of 384 per GPU, with gradient accumulation over 2 steps, resulting in an effective batch size of 768 per GPU (3840 globally). The AdamW optimizer (Loshchilov and Hutter, 2017) was used with a base learning rate of 1×10^{-5} for the visual and auxiliary encoders, and a reduced learning rate of 1×10^{-6} for the pretrained text encoder. A weight decay of 1×10^{-4} was applied. A linear learning rate warmup was implemented for the first 5 epochs. Following the warmup, a Cosine Annealing with Warm Restarts (Loshchilov and Hutter, 2022) scheduler was used, with $T_0 = 10$ epochs and $T_{mult} = 2$, and a minimum learning rate of 1×10^{-7} . Gradients were clipped to a maximum L2 norm of 1.0. The loss function was the standard CLIP contrastive loss (Oord et al., 2018). Mixed-precision training (Micikevicius et al., 2017) was enabled. A full 25-epoch training run required approximately 2 days, with peak GPU memory consumption observed at roughly 14 GB per GPU.

5.2 Experimental setup

We evaluated GeoTreeCLIP against two advanced pretrained vision-language models: CLIP (Radford et al., 2021) and RemoteCLIP (Liu et al., 2024), a specialized VLM for remote sensing applications. To ensure a fair comparison with models not explicitly designed for time-series data, we employed an ensemble-like approach for CLIP and RemoteCLIP. Specifically, features or probabilities were computed independently for each of the 12 monthly images, and the final results were obtained by averaging these values. Additionally, the results of other baseline models, including SkyCLIP-50 (Wang et al., 2024) and CLIP-laion-RS (He et al., 2024), are provided in Appendix E.

All models were evaluated on the *GlobalGeoTree-10kEval* benchmark using zero-shot and few-shot learning settings. Performance was measured using top-1 and top-5 prediction accuracy, with separate evaluations for each taxonomic level

(family, genus, and species). To ensure robustness, we repeated each experiment 5 times using different random seeds and reported the mean accuracy and variance.

5.2.1 Zero-shot evaluation

For zero-shot evaluation, we assess each model’s ability to classify samples from the evaluation sets (*GlobalGeoTree-10kEval*, *-10kEval-300*, and *-10kEval-900*) without any fine-tuning. Our protocol is framed as an "in-domain" zero-shot classification task, following the setup in recent remote sensing benchmarks (Wang et al., 2024). This approach tests a model’s transfer capability to new, unseen samples of species, rather than to entirely unseen categories. Although the species categories in the evaluation sets are present in the training data, this setup strikes a balance between preserving the model’s ability to classify all tree species after pre-training on *GlobalGeoTree-6M* and ensuring a valid zero-shot scenario.

5.2.2 Few-shot evaluation

For few-shot evaluation, we explored scenarios such as one-shot learning, where the model is provided with only one labeled example per species. To implement this, we adopted a fine-tuning-based approach (Parnami and Lee, 2022) using the pre-trained model. Specifically, we randomly sampled k labeled examples per class (e.g., $k = 1$, $k = 3$) to form the support set and fine-tuned the visual encoder of the pre-trained model on this set. During fine-tuning, most of the visual encoder’s parameters were frozen, with only the last four transformer layers and the classification-related parameters remaining trainable. The text encoder was entirely frozen, leveraging the pre-trained textual embeddings for class labels. The fine-tuning was conducted for 10 epochs to balance adaptation and prevent overfitting due to the small support set. Afterward, the model was evaluated on the query set, which consisted of the remaining examples in the dataset.

For each query image, predictions were made by computing similarity scores between its visual embedding (extracted by the fine-tuned visual encoder) and the textual embeddings of the class labels. The class with the highest similarity score was assigned as the predicted label. Classification accuracy on the query set was then used to evaluate the model’s performance. This fine-tuning approach enables the model to adapt to the few-shot setting while retaining the benefits of the pre-trained representations.

This evaluation framework highlights the model’s ability to generalize effectively from limited labeled data, which is a crucial capability for real-world applications in biodiversity monitoring where obtaining large amounts of labeled data for every species is often infeasible.

5.3 Experimental results

5.3.1 Zero-shot evaluation

The results of zero-shot evaluation are presented in Table 4, clearly demonstrating the substantial improvements achieved by GeoTreeCLIP across all taxonomic levels. At the family level, GeoTreeCLIP achieves a top-1 accuracy of 20.99% and a top-5 accuracy of 56.88%. The performance gap is even more pronounced at the genus level. GeoTreeCLIP achieves a top-1

Table 4. Zero-shot evaluation on *GlobalGeoTree-10kEval*. Results are presented as mean accuracy (%) $\pm$ standard deviation (%) over 5 runs.

Taxon.	CLIP		RemoteCLIP		GeoTreeCLIP	
	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
Family	10.80 $\pm$ 0.03	25.32 $\pm$ 0.05	1.11 $\pm$ 0.01	10.55 $\pm$ 0.04	20.99 $\pm$ 0.28	56.88 $\pm$ 0.42
Genus	1.09 $\pm$ 0.01	9.34 $\pm$ 0.01	1.11 $\pm$ 0.01	6.25 $\pm$ 0.02	18.39 $\pm$ 0.26	50.98 $\pm$ 0.41
Species	1.09 $\pm$ 0.01	7.02 $\pm$ 0.02	1.11 $\pm$ 0.01	6.25 $\pm$ 0.02	16.71 $\pm$ 0.25	47.52 $\pm$ 0.37

accuracy of 18.39% and a top-5 accuracy of 50.98%, outperforming CLIP (1.09% top-1, 9.34% top-5) and RemoteCLIP (1.11% top-1, 6.25% top-5) by a large margin. At the most challenging species level, GeoTreeCLIP still shows significant superiority, achieving a top-1 accuracy of 16.71% and a top-5 accuracy of 47.52%.

The experimental results reveal two key patterns. First, accuracy consistently declines as the taxonomic level becomes finer, reflecting the growing challenge of distinguishing closely related classes. This trend is observed across all models but is particularly pronounced for CLIP and RemoteCLIP, which perform poorly at the genus and species levels. In contrast, GeoTreeCLIP demonstrates stronger performance at these fine-grained levels, likely due to its ability to learn and leverage the hierarchical relationships in taxonomic labels, as supported by feature embedding visualizations (see Sect. 6.2 and Figure 8). Second, the significant performance gap between GeoTreeCLIP and the baseline models underscores synergistic effect of domain-specific pretraining and its tailored architecture. Unlike general-purpose models such as CLIP and RemoteCLIP, our approach effectively leverages spatiotemporal and multispectral information to enhance classification capabilities. Additional zero-shot benchmark results, including evaluations of SkyCLIP-50 (Wang et al., 2024) and CLIP-laion-RS (He et al., 2024), are provided in Appendix E.

5.3.2 Performance Analysis by Species Rarity

To assess model robustness against the long-tail distribution inherent in global biodiversity data, we analyzed zero-shot performance at the species level across three frequency groups: Rare, Common, and Frequent. Table 5 details the mean accuracy and standard deviation over 5 runs on the *GlobalGeoTree-10kEval* benchmark.

Table 5. Zero-shot species-level performance breakdown by rarity group on *GlobalGeoTree-10kEval*. Results are mean accuracy (%) $\pm$ standard deviation (%) over 5 runs.

Rarity Group	CLIP		RemoteCLIP		GeoTreeCLIP	
	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
Rare	0.00 $\pm$ 0.00	3.61 $\pm$ 0.05	0.00 $\pm$ 0.00	2.43 $\pm$ 0.13	15.25 $\pm$ 0.04	46.65 $\pm$ 0.47
Common	0.64 $\pm$ 0.04	4.52 $\pm$ 0.03	0.00 $\pm$ 0.00	3.45 $\pm$ 0.00	15.93 $\pm$ 0.08	47.51 $\pm$ 0.08
Frequent	3.00 $\pm$ 0.03	9.27 $\pm$ 0.02	3.35 $\pm$ 0.01	12.90 $\pm$ 0.04	17.52 $\pm$ 0.36	51.88 $\pm$ 0.12

The experimental results indicate substantial differences in model generalization across frequency groups. Both baseline models show significant performance degradation on data-scarce classes. Specifically, CLIP yields 0.00% top-1 accuracy for the Rare group, while RemoteCLIP records 0.00% top-1 accuracy across both Rare and Common categories. Conversely, GeoTreeCLIP exhibits consistent performance stability across the frequency spectrum. The model achieves a top-1 accuracy of 15.25% on Rare species, showing a relatively small decrease compared to the 17.52% accuracy observed for Frequent species. This limited performance disparity suggests that GeoTreeCLIP facilitates knowledge transfer through hierarchical taxonomic relationships, enabling the identification of rare species by leveraging features learned from more frequent, related taxa.

5.3.3 Performance Analysis by Geographic Region

To evaluate model generalizability across different geographical regions, we disaggregated the zero-shot results on *GlobalGeoTree-10kEval* by continent. Table 6 presents the species-level accuracy for each region.

Table 6. Zero-shot species-level performance by continent on *GlobalGeoTree-10kEval*. Results are mean accuracy (%) $\pm$ standard deviation (%) over 5 runs. (n indicates the sample count per region in the evaluation set).

Region	CLIP		RemoteCLIP		GeoTreeCLIP	
	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
Africa ($n = 558$)	17.20 ± 0.05	19.12 ± 0.05	0.00 ± 0.00	0.00 ± 0.00	17.31 ± 0.21	56.63 ± 1.51
Asia ($n = 1120$)	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	4.91 ± 0.09	8.31 ± 0.19	39.61 ± 1.60
Europe ($n = 1648$)	0.00 ± 0.00	19.08 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	19.96 ± 0.19	49.91 ± 0.24
North America ($n = 4442$)	0.00 ± 0.00	1.15 ± 0.02	0.00 ± 0.00	6.81 ± 0.01	18.98 ± 0.33	52.79 ± 0.45
Oceania ($n = 812$)	0.92 ± 0.06	4.83 ± 0.07	13.25 ± 0.00	24.64 ± 0.18	7.64 ± 0.11	24.83 ± 0.02
South America ($n = 1150$)	0.00 ± 0.00	14.76 ± 0.01	0.00 ± 0.00	4.60 ± 0.25	16.83 ± 0.75	41.89 ± 1.19

GeoTreeCLIP demonstrates relatively consistent zero-shot capabilities across diverse geographic regions, with top-1 accuracies ranging from 7.64% to 19.96%. The highest performance is observed in Europe and North America, which aligns with the higher density of training samples available for these regions in the GlobalGeoTree dataset. Conversely, performance is lower in Asia and Oceania, likely reflecting the higher species diversity and relative scarcity of labeled data in these areas.

Notably, baseline models exhibit extreme regional bias. CLIP, for example, achieves high accuracy in Africa (17.20%) but drops to 0.00% in most other continents. This suggests that without domain-specific pretraining, models may rely on spurious correlations with broad landscape features (e.g., savannas) rather than learning discriminative species-level traits. GeoTreeCLIP’s ability to maintain performance across all continents highlights its generalization capabilities.

Table 7. Few-shot evaluation on *GlobalGeoTree-10kEval*. Results are presented as mean accuracy (%) $\pm$ standard deviation (%) over 5 runs.

Taxon.	CLIP		RemoteCLIP		GeoTreeCLIP	
	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
<i>One-Shot Evaluation</i>						
Family	2.95 $\pm$ 0.01	15.06 $\pm$ 0.03	11.25 $\pm$ 0.01	23.85 $\pm$ 0.03	29.37 $\pm$ 0.07	69.38 $\pm$ 0.34
Genus	2.43 $\pm$ 0.01	8.14 $\pm$ 0.02	2.31 $\pm$ 0.01	7.68 $\pm$ 0.02	27.70 $\pm$ 0.11	64.40 $\pm$ 0.29
Species	1.67 $\pm$ 0.01	6.59 $\pm$ 0.03	1.94 $\pm$ 0.00	6.59 $\pm$ 0.02	25.80 $\pm$ 0.15	62.43 $\pm$ 0.25
<i>Three-Shot Evaluation</i>						
Family	6.19 $\pm$ 0.01	23.50 $\pm$ 0.03	4.44 $\pm$ 0.02	17.02 $\pm$ 0.04	37.77 $\pm$ 0.23	75.49 $\pm$ 0.25
Genus	4.04 $\pm$ 0.01	12.76 $\pm$ 0.03	2.77 $\pm$ 0.03	11.70 $\pm$ 0.02	36.19 $\pm$ 0.22	72.50 $\pm$ 0.23
Species	3.41 $\pm$ 0.01	11.46 $\pm$ 0.03	1.88 $\pm$ 0.03	9.40 $\pm$ 0.06	33.67 $\pm$ 0.24	71.53 $\pm$ 0.23

5.3.4 Few-shot evaluation

The results in Table 7 demonstrate that providing even a small amount of labeled data for fine-tuning generally improves performance compared to the zero-shot setting across all models and taxonomic levels. GeoTreeCLIP consistently achieves the highest accuracies in both one-shot and three-shot scenarios. For instance, in the one-shot setting at the species level, GeoTreeCLIP reaches a top-1 accuracy of 25.80%, substantially outperforming CLIP (1.67%) and RemoteCLIP (1.94%). This advantage becomes even more pronounced with three shots, where GeoTreeCLIP’s species-level top-1 accuracy increases to 33.67%, while CLIP and RemoteCLIP show more modest gains to 3.41% and 1.88%, respectively.

Our few-shot experiments reveal distinct patterns across models when increasing from one to three shots. GeoTreeCLIP demonstrates substantial improvements across all taxonomic levels (species top-1 accuracy rising from 25.80% to 33.67%), while CLIP shows consistent but smaller gains. RemoteCLIP exhibits mixed results, including a slight decrease in family-level accuracy, suggesting difficulties effectively utilizing additional examples. Both baseline models demonstrate limited capacity to leverage few-shot supervision compared to GeoTreeCLIP, with only marginal improvements over their zero-shot performances (Table 4), particularly at finer taxonomic levels. This indicates that general pretraining approaches may not align sufficiently with the specific challenges of fine-grained tree species classification from remote sensing data, even when provided with in-domain examples.

The poor performance of CLIP and RemoteCLIP can likely be attributed to their design, which is optimized for RGB three-channel data and lacks the capability to process time-series information. Additionally, these models struggle with the small-patch classification tasks required for tree species identification. These limitations further emphasize the importance of introducing this benchmark for the global tree species classification task. More benchmark results on larger evaluation subsets can be found in Table D1 and Table D2 in Appendix D.

6 In-depth Analysis

6.1 Comparison with Supervised Learning Paradigm

To further contextualize the performance of our contrastive learning-based GeoTreeCLIP model, we conducted an additional experiment using a traditional supervised learning paradigm. This allows for a more direct comparison of learning objectives (contrastive vs. supervised) while keeping core architectural components and training settings as consistent as possible.

6.1.1 Supervised Model Architecture and Training

We designed a supervised model, termed SupervisedGeoTree, which retains the visual processing pathway of GeoTreeCLIP, including the `VisualEncoder` for Sentinel-2 time series and the `AuxiliaryEncoder` for environmental variables. The features from these two encoders are projected, normalized, and then fused via concatenation followed by a fusion layer, similar to the visual feature preparation in GeoTreeCLIP.

However, unlike GeoTreeCLIP, the SupervisedGeoTree model does not include a text encoder or employ a contrastive loss. Instead, the fused visual-auxiliary features are fed into four independent classification heads (fully connected layers), each dedicated to predicting labels for one of the hierarchical taxonomic levels: functional type (`level0`), family (`level1_family`), genus (`level2_genus`), and species (`level3_species`). The number of output neurons for each head corresponds to the number of unique classes at that respective taxonomic level in the GlobalGeoTree (4 for `level0`, 275 for family, 2,734 for genus, and 21,001 for species).

The model was trained on the *GlobalGeoTree-6M* dataset. The loss function employed was a sum of standard Cross-Entropy losses, calculated independently for each of the four taxonomic levels. The contributions of each level’s loss to the total loss were equally weighted. All other training hyperparameters, including the learning rate (1×10^{-5}), optimizer (AdamW), weight decay, number of epochs (25), warmup strategy (5 epochs), learning rate scheduler (Cosine Annealing with Warm Restarts), and batch size, were kept identical to those used for pretraining GeoTreeCLIP to ensure a fair comparison of the learning paradigms.

6.1.2 Comparison of the Supervised Model

After training on the *GlobalGeoTree-6M* dataset, the SupervisedGeoTree model was evaluated on the *GlobalGeoTree-10kEval* subset. Since this model is trained with fixed classification heads for the classes seen during training, its ability to perform "zero-shot" classification in the same sense as a CLIP-style model (i.e., classifying entirely new, unseen samples provided at test time) is inherently limited. However, for this comparison, we evaluate its performance on the classes within *GlobalGeoTree-10kEval* that were also part of the *GlobalGeoTree-6M* training vocabulary for each respective taxonomic head. If a class in *GlobalGeoTree-10kEval* was not in the training vocabulary for a specific head, it cannot be correctly predicted by that head.

Notably, the *GlobalGeoTree-6M* dataset was designed to retain nearly all tree species categories, as our goal was to pretrain a model capable of classifying the full spectrum of tree species. Consequently, the zero-shot evaluation here can be seen as

measuring the model’s “zero-shot” transfer capability on unseen datasets, aligning with the concept of "in-domain" zero-shot classification defined in Wang et al. (2024).

Table 8. Evaluation on *GlobalGeoTree-10kEval*. The table compares our vision-language model, GeoTreeCLIP, against the CLIP and a traditional supervised baseline (SupervisedGeoTree). Results are mean accuracy (%) $\pm$ standard deviation (%) over 5 runs.

Taxon.	CLIP		SupervisedGeoTree		GeoTreeCLIP	
	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
Family	10.80 $\pm$ 0.03	25.32 $\pm$ 0.05	9.55 $\pm$ 0.10	27.41 $\pm$ 0.07	20.99 $\pm$ 0.28	56.88 $\pm$ 0.42
Genus	1.09 $\pm$ 0.01	9.34 $\pm$ 0.01	1.19 $\pm$ 0.09	8.18 $\pm$ 0.14	18.39 $\pm$ 0.26	50.98 $\pm$ 0.41
Species	1.09 $\pm$ 0.01	7.02 $\pm$ 0.02	0.00 $\pm$ 0.00	0.28 $\pm$ 0.02	16.71 $\pm$ 0.25	47.52 $\pm$ 0.37

The results in Table 8 indicate that a consistent trend across all models is the decline in classification accuracy as the taxonomic level becomes finer (from family to species). This reflects the inherent increase in difficulty when distinguishing between more closely related taxa. However, the extent of this performance degradation varies significantly between models.

The SupervisedGeoTree model, which employs traditional supervised classification heads for each taxonomic level, achieves reasonable accuracy at the family level (9.55% top-1). However, its performance drops sharply for genus (1.19% top-1) and becomes negligible at the species level (0.00% top-1). This drastic decline underscores the challenge of fine-grained classification when relying solely on visual and auxiliary features without leveraging the semantic relationships embedded in textual taxonomic labels, and the inherent limitation of generalizing to a large number of specific classes in a purely supervised manner.

In contrast, the proposed GeoTreeCLIP model demonstrates substantial improvements over all other baselines across every taxonomic level. At the family level, GeoTreeCLIP achieves a top-1 accuracy of 20.99%, more than doubling the performance of CLIP and significantly outperforming SupervisedGeoTree. This advantage is even more pronounced at the finer-grained levels: GeoTreeCLIP obtains 18.39% top-1 accuracy for genus and 16.71% for species identification. These results strongly indicate the power of contrastive vision-language learning to align visual features with rich, hierarchical taxonomic text labels. This capability allows GeoTreeCLIP to learn more nuanced and generalizable representations, leading to its superior performance in this challenging zero-shot evaluation. The significant gap, especially compared to SupervisedGeoTree at the species level, highlights the efficacy of the contrastive learning paradigm for handling large, structured label spaces.

6.2 Qualitative Analysis: t-SNE Feature Embeddings Visualization

To qualitatively assess and compare the learned feature representations from different modeling paradigms, we performed t-SNE visualizations (Van der Maaten and Hinton, 2008). We used zero-shot image embeddings (or the final fused visual-auxiliary features for SupervisedGeoTree before the classification heads) extracted from a subset of the *GlobalGeoTree-10kEval-300* dataset. This analysis includes our proposed GeoTreeCLIP, the CLIP pretrained by OpenAI, and the SupervisedGeoTree model. Given the large number of classes and the hierarchical nature of the labels, we adopted a selective visualization strategy: first examining embeddings at the family level for five randomly selected families (Anacardiaceae, Berberidaceae, Cactaceae,




Figure 8. t-SNE visualization of feature embeddings from *GlobalGeoTree-10kEval-300* at different taxonomic levels, comparing GeoTreeCLIP (top row), CLIP (middle row), and SupervisedGeoTree (bottom row). Columns from left to right represent visualizations at the Family level (selected: Anacardiaceae, Berberidaceae, Cactaceae, Hernandiaceae, Rosaceae), Genus level (selected within Rosaceae: *Prunus*, *Pyrus*, *Rhodotypos*, *Rubus*, *Spiraea*), and Species level (selected within *Prunus*: *Prunus avium*, *Prunus caroliniana*, *Prunus ilicifolia*, *Prunus laurocerasus*, *Prunus obtusata*).

Hernandiaceae, and Rosaceae); then focusing on five genera within the Rosaceae family (*Prunus*, *Pyrus*, *Rhodotypos*, *Rubus*, and *Spiraea*); and finally, visualizing five species within the *Prunus* genus (*Prunus avium*, *Prunus caroliniana*, *Prunus ilicifolia*, *Prunus laurocerasus*, and *Prunus obtusata*).

The comparative t-SNE visualizations are presented in Figure 8. Across all three taxonomic levels (Family, Genus, and Species, shown as columns), GeoTreeCLIP (top row) consistently demonstrates the most effective separation and formation

of distinct clusters. At the family level (left column), GeoTreeCLIP clearly distinguishes between the selected families. The CLIP (middle row) exhibits considerable overlap, particularly for the diverse Rosaceae family. SupervisedGeoTree (bottom row) shows some separation but less defined clusters compared to GeoTreeCLIP, with Rosaceae still forming a very broad distribution.

This pattern of superior clustering by GeoTreeCLIP continues at the genus level within Rosaceae (middle column). GeoTreeCLIP forms relatively distinct groups for genera like *Prunus*, *Pyrus*, and *Rhodotypos*. Both CLIP and SupervisedGeoTree struggle more, with CLIP showing a highly condensed and overlapping structure, while SupervisedGeoTree offers some separation but with less clarity than GeoTreeCLIP. The most striking difference is observed at the species level within the *Prunus* genus (right column). GeoTreeCLIP achieves remarkable separation, forming visually distinct clusters for each of the five *Prunus* species. In contrast, both CLIP and SupervisedGeoTree largely fail to differentiate these closely related species, with their embeddings heavily intermingled.

These visualizations qualitatively affirm that GeoTreeCLIP, through its contrastive vision-language learning approach tailored with hierarchical taxonomic information and domain-specific data, learns more semantically meaningful and discriminative representations across all taxonomic ranks. This aligns with its superior quantitative performance in classification tasks compared to both general-domain VLMs and a traditional supervised approach.

6.3 Ablation Study on Input Modalities

To quantify the contribution of different input modalities, we conducted an extensive ablation study. We trained several variants of GeoTreeCLIP (for 10 epochs on *GlobalGeoTree-6M*) and evaluated their zero-shot performance on the *GlobalGeoTree-10kEval* benchmark. The results are summarized in Table 9.

Table 9. Ablation study on input modalities. Performance is reported as mean Top-1 accuracy (%) over 5 runs on *GlobalGeoTree-10kEval*.

Model Variant	Family Acc.	Genus Acc.	Species Acc.
GeoTreeCLIP (Full Input)	18.66%	16.30%	13.95%
- Image Only (10-band S2)	15.83%	14.51%	12.62%
- Auxiliary Only (27 vars)	13.66%	10.50%	9.52%
- RGB + Geolocation Only	7.92%	6.04%	5.50%
- RGB Only	6.38%	5.08%	4.43%

The results lead to several key insights. First, the full input consistently outperforms all ablated versions, highlighting the synergistic effect of integrating spectral-temporal imagery with environmental context. Second, comparing the "Image Only (10-band)" and "RGB Only" variants reveals the critical importance of non-visible bands (e.g., NIR, SWIR) for vegetation analysis, as they provide a substantial performance boost (e.g., 15.83% vs. 6.38% at the family level). Third, adding geolocation to RGB imagery ("RGB + Geolocation") improves performance over "RGB Only", confirming that spatial priors help capture

species' geographical ranges. The reason "Image Only (10-band)" outperforms "RGB + Geolocation" is that the rich information from the 7 additional spectral bands is more discriminative than the signal provided by geolocation alone.

7 Ethics, Limitations and Impact

7.1 Ethics

Our data collection process was designed to adhere to ethical standards for data reuse. All geolocated occurrence records sourced from GBIF (Lane and Edwards, 2007) were filtered to include only those with permissive licenses (e.g., CC0, CC-BY) explicitly chosen by the original data contributors. We also respected the privacy safeguards of the source platforms; for instance, the coordinates of sensitive or threatened species are often automatically obscured by platforms like iNaturalist, and our dataset preserves these privacy-preserving modifications.

The Sentinel-2 L2A data (Spoto et al., 2012) were accessed via Google Earth Engine. The auxiliary bioclimatic, geographic, and soil data were sourced from WorldClim (Fick and Hijmans, 2017), USGS SRTM (Jarvis et al., 2008), and SoilGrids (Poggio et al., 2021), respectively, all of which are publicly available.

7.2 Limitations

While GlobalGeoTree represents one of the largest and most comprehensive datasets of its kind, users should be aware of its inherent limitations. As with large part of datasets derived from citizen-science observations, it contains geographic and taxonomic biases. Data coverage is higher in regions with active observer communities (e.g., North America, Europe) and is skewed towards more common or easily identifiable species, resulting in a long-tail distribution. Furthermore, since we assimilated available official inventory data (e.g., NFIs) into the training set to maximize coverage, independent external validation remains challenging, and shared observation biases may influence evaluation metrics. The ambiguous boundary between trees and shrubs in botanical classification further complicates the dataset, as some samples may represent shrubs rather than trees. Tree species taxonomy is also subject to frequent revisions driven by new genetic evidence, which may misalign dataset labels with updated classifications over time. Although CLIP-based models can identify unseen species to some extent, such taxonomic shifts may still affect model interpretability and evaluation consistency. Additionally, the dataset relies on Sentinel-2 data from 2020, restricting its ability to capture long-term vegetation dynamics or recent disturbances. Future versions incorporating multi-year observations could better account for phenological changes and climate-driven impacts.

7.3 Potential impact and Operational Deployment

GlobalGeoTree holds significant potential for advancing forest monitoring, biodiversity conservation, and climate change mitigation. Beyond benchmarking, the dataset and GeoTreeCLIP model support practical operational workflows. For large-scale mapping, users can employ a sliding-window inference approach on Sentinel-2 tiles to generate pixel-wise species maps. To handle prediction uncertainty in diverse ecosystems, we recommend a hierarchical inference strategy, where the model defaults

to genus or family-level classifications if species-level confidence falls below a threshold. Furthermore, to address regional domain shifts, the pretrained model serves as a robust foundation that can be efficiently fine-tuned with small sets of local field data. This "local adaptation" workflow allows practitioners to leverage global knowledge while tailoring predictions to specific forest management needs.

8 Code and Data Availability

The GlobalGeoTree dataset is openly available under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. The complete dataset is archived on Hugging Face at <https://huggingface.co/datasets/yann111/GlobalGeoTree> under DOI <https://doi.org/10.15468/dd.9qxqyy> (Mu et al., 2025b). The pretraining and evaluation datasets are provided in WebDataset format (Aizman et al., 2019). This format enables efficient online data streaming to train models without requiring full dataset downloads, facilitating large-scale machine learning workflows. It also integrates seamlessly with popular deep learning frameworks, improving accessibility and usability for researchers.

All code is available in our github repository <https://github.com/MUYang99/GlobalGeoTree> under the Apache License 2.0, which provides comprehensive tools for using these resources. We are dedicated to maintaining and enhancing the dataset, addressing issues, and incorporating updates like new data sources or taxonomic revisions in future versions. The code and Huggingface repository will serve as the primary channels for updates and community feedback.

9 Conclusion and future work

In this paper, we introduced GlobalGeoTree, a large-scale, globally comprehensive dataset and benchmark for tree species classification. The dataset includes over 6 million geolocated tree occurrences spanning 21,001 species, paired with Sentinel-2 time series data and a rich set of auxiliary environmental variables. We also proposed GeoTreeCLIP, a baseline vision-language model specifically designed for this task, leveraging domain-specific pretraining on *GlobalGeoTree-6M*. Experimental results demonstrate that GeoTreeCLIP significantly outperforms existing advanced models in classification accuracy across all taxonomic levels, highlighting both the effectiveness of our approach and the importance of introducing this benchmark for global tree species classification.

Future work could explore several promising directions. Expanding the GlobalGeoTree with more recent data, additional satellite sensors (e.g., SAR data for structural information), and a broader range of auxiliary variables could enhance its utility. Specifically, incorporating global structural datasets, such as GEDI LiDAR metrics or Meta's 1-m Canopy Height Map, could provide critical vertical information to distinguish between ecologically similar tree and shrub species, a limitation of current spectral-only approaches. Investigating alternative vision-language model architectures, pretraining strategies, and methods for addressing the long-tail distribution could further improve classification accuracy, especially at the species level. Developing techniques for uncertainty estimation and improving model explainability are also critical areas for future work. Moreover,

applying GlobalGeoTree and GeoTreeCLIP to real-world applications in biodiversity monitoring, conservation, and forest management could provide practical support and holds great potential.

Appendix A: GlobalGeoTree Dataset Statistics

The GlobalGeoTree dataset is a large-scale, multimodal resource for tree species classification. This section provides detailed statistics complementing the overview in the main paper.

A1 Data Records and Taxonomic Catalog

To provide full transparency and facilitate detailed exploration of the dataset, we have made two key files publicly available in our Hugging Face repository (<https://huggingface.co/datasets/yann111/GlobalGeoTree/tree/main/files>).

- **Complete Taxonomic Catalog (`Tree_species_catalog.csv`):** This file contains the complete list of all 275 families, 2,734 genera, and 21,001 species included in GlobalGeoTree, along with their corresponding unique species keys used in the GBIF database. This catalog serves as a comprehensive reference for the taxonomic scope of our dataset.
- **Full Georeferenced Records (`GlobalGeoTree.csv`):** This file contains all 6.3 million geolocated records. Each row includes the coordinates, hierarchical taxonomic labels (from functional type to species), the original data source (e.g., iNaturalist Research-grade Observations), and other metadata as detailed in Table 3. This allows for complete traceability and enables users to perform custom geographical or taxonomic subsetting.

A2 Basic Statistics

The dataset encompasses a comprehensive collection of geolocated tree occurrences:

- **Total Samples:** 6,263,345
- **Countries/Regions¹ Covered:** 221
- **Taxonomic Coverage:** Families: 275, Genera: 2,734, Species: 21,001

A3 Long-tail Distribution Analysis

The dataset exhibits a characteristic long-tail distribution across taxonomic levels as shown in Fig. A1. This highlights the challenge of classifying both common and rare taxa:

- **Family Level:** The top 20% of families (55 families) account for 91.01% of all samples. Conversely, 24 families (8.73% of total families) have fewer than 10 samples each.
- **Genus Level:** The top 20% of genera (546 genera) cover 96.65% of the samples. A significant portion, 760 genera (27.80% of total genera), have fewer than 10 samples.

¹The term "Countries/Regions" is used to align with global data standards (e.g., from GBIF) and accurately represents both sovereign states and other geographic entities such as overseas territories or dependencies, ensuring comprehensive coverage.

- **Species Level:** The distribution is most skewed at the species level, where the top 20% of species (4,200 species) comprise 97.21% of the samples. A majority of species, 11,611 species (55.29% of total species), have fewer than 10 samples.

This long-tail distribution underscores the importance of evaluation strategies, like those employed for *GlobalGeoTree-10kEval*, that explicitly consider species rarity.




Figure A1. Long-tail distribution across taxonomic levels in GlobalGeoTree.

A4 Detailed Categorical Statistics

Below are statistics for key categorical attributes within the GlobalGeoTree dataset, illustrating geographical and taxonomic diversity and distribution.

A4.1 Location (location)

- Number of unique countries/regions: 221
- Top 5 most frequent locations:
 1. United States of America: 1,932,465
 2. Australia: 506,179
 3. Canada: 429,266
 4. Colombia: 330,896
 5. Russian Federation: 209,019

A4.2 Functional Type (level10)

- Number of unique functional types: 4
- Distribution of functional types:
 - Deciduous Broadleaf: 3,582,456
 - Evergreen Broadleaf: 2,208,578
 - Evergreen Needleleaf: 447,568
 - Deciduous Needleleaf: 24,743

A4.3 Taxonomic Family (level1_family)

- Number of unique families: 275
- Top 5 most frequent families:
 1. Ericaceae: 423,365
 2. Fabaceae: 384,038
 3. Fagaceae: 362,317
 4. Rosaceae: 355,950
 5. Pinaceae: 320,415

A4.4 Taxonomic Genus (level2_genus)

- Number of unique genera: 2,734
- Top 5 most frequent genera:
 1. Cornus: 288,678

2. Quercus: 210,104
3. Pinus: 168,917
4. Vaccinium: 158,362
5. Prunus: 125,604

A4.5 Taxonomic Species (level3_species)

- Number of unique species: 21,001
- Top 5 most frequent species:
 1. Cornus acuminata: 180,120
 2. Securidaca volubilis: 103,441
 3. Cupania sylvatica: 99,797
 4. Bourreria cumanensis: 96,304
 5. Fagus sylvatica: 75,503

A5 Representative Species Profile

Table A1 provides a structured summary of representative species from the GlobalGeoTree dataset, highlighting the diversity across different functional types and geographic regions.

Table A1. Structured summary of representative tree species in GlobalGeoTree, spanning major functional types and biogeographical realms.

Functional Type	Family	Genus	Species	Dominant Region	Ecological Context
Evergreen Needleleaf	Pinaceae	Picea	Picea abies	Europe	Dominant boreal forest conifer
Evergreen Needleleaf	Pinaceae	Pseudotsuga	Pseudotsuga menziesii	North America	Key timber species of Pacific Northwest
Deciduous Broadleaf	Fagaceae	Quercus	Quercus robur	Europe	Keystone species of temperate woodlands
Deciduous Broadleaf	Sapindaceae	Acer	Acer saccharum	North America	Major component of northern hardwood forests
Evergreen Broadleaf	Myrtaceae	Eucalyptus	Eucalyptus globulus	Oceania	Fast-growing, widely distributed in Australia
Evergreen Broadleaf	Urticaceae	Cecropia	Cecropia peltata	South America	Pioneer species in neotropical rainforests
Deciduous Needleleaf	Pinaceae	Larix	Larix decidua	Europe	High-altitude conifer of the Alps

Appendix B: Auxiliary Environmental Variables

The 27 auxiliary variables used in GlobalGeoTree were sourced from three publicly available datasets, accessible via Google Earth Engine (GEE):

- **WorldClim V1 (Bioclimatic Variables):** GEE Dataset ID: WORLDCLIM/V1_BIO
- **SoilGrids 250m v2.0 (Soil Variables):** GEE Dataset ID: ISRIC/soilgrids/v02
- **USGS SRTM GL1 v003 (Geographic Variables):** GEE Dataset ID: USGS/SRTMGL1_003

Table B1 provides a detailed description of a subset of these variables.

Table B1: Description of the 19 bioclimatic variables from WorldClim (BIO1-BIO19) used in the GlobalGeoTree dataset.

Variable	Units	Scale Factor	Description
bio01	°C	0.1	Annual Mean Temperature
bio02	°C	0.1	Mean Diurnal Range (Mean of monthly (max temp - min temp))
bio03	%	1	Isothermality ($BIO2/BIO7 \times 100$)
bio04	°C	0.01	Temperature Seasonality (Standard Deviation $\times 100$)
bio05	°C	0.1	Max Temperature of Warmest Month
bio06	°C	0.1	Min Temperature of Coldest Month
bio07	°C	0.1	Temperature Annual Range (BIO5 - BIO6)
bio08	°C	0.1	Mean Temperature of Wettest Quarter
bio09	°C	0.1	Mean Temperature of Driest Quarter
bio10	°C	0.1	Mean Temperature of Warmest Quarter
bio11	°C	0.1	Mean Temperature of Coldest Quarter
bio12	mm	1	Annual Precipitation
bio13	mm	1	Precipitation of Wettest Month
bio14	mm	1	Precipitation of Driest Month
bio15	Coeff. of Variation	1	Precipitation Seasonality
bio16	mm	1	Precipitation of Wettest Quarter
bio17	mm	1	Precipitation of Driest Quarter
bio18	mm	1	Precipitation of Warmest Quarter
bio19	mm	1	Precipitation of Coldest Quarter

Appendix C: Details of Evaluation Subsets

C1 Overview and Construction

To enable robust benchmarking across various taxonomic diversity scales and species rarity, we constructed three evaluation subsets: *GlobalGeoTree-10kEval*, *GlobalGeoTree-10kEval-300*, and *GlobalGeoTree-10kEval-900*. These subsets were created by first categorizing all species in the GlobalGeoTree dataset into Rare, Common, and Frequent groups based on available sample counts (Sect. 3.4), then randomly selecting 30, 100, and 300 species per category, respectively. The primary *GlobalGeoTree-10kEval* benchmark (90 species) is featured in the main paper, while the larger subsets enable assessment of model scalability and performance on increasingly complex tasks. Detailed overviews of each subset’s composition are provided in Table C1, Table C2, and Table C3. The geographical distribution of the two additional evaluation sets is shown in Figure C1a and Figure C1b.

C2 *GlobalGeoTree-10kEval-300* and *GlobalGeoTree-10kEval-900*

Table C1. Overview of the *GlobalGeoTree-10kEval* evaluation subset. This subset comprises 30 species selected from each of the three rarity categories (Rare, Common, Frequent), totaling 90 unique species and 9,930 geolocated samples.

Category	Example Species	Num of Samples
Rare	<i>Acacia platycarpa</i>	40
	<i>Adenanthos cuneatus</i>	40
	<i>Adenocarpus decorticans</i>	40
	...	...
Rare (Total)	30 species selected	1,200
Common	<i>Abies religiosa</i>	110
	<i>Aloe marlothii</i>	110
	<i>Alternanthera sessilis</i>	110
	...	...
Common (Total)	30 species selected	3,300
Frequent	<i>Acer glabrum</i>	181
	<i>Arctostaphylos glandulosa</i>	181
	<i>Ardisia paniculata</i>	181
	...	...
Frequent (Total)	30 species selected	5,430
Total	90 species total	9,930

Table C2. Overview of the *GlobalGeoTree-10kEval-300* evaluation subset. This subset comprises 100 species selected from each of the three rarity categories (Rare, Common, Frequent), totaling 300 unique species and 10,000 geolocated samples.

Category	Example Species	Num of Samples
Rare	<i>Abutilon wrightii</i>	12
	<i>Acacia georgensis</i>	12
	<i>Acacia loroloba</i>	12
	...	...
Rare (Total)	<i>100 species selected</i>	1,200
Common	<i>Acacia confusa</i>	33
	<i>Acacia mucronata</i>	33
	<i>Achyranthes spec</i>	33
	...	...
Common (Total)	<i>100 species selected</i>	3,300
Frequent	<i>Acacia dealbata</i>	55
	<i>Acacia decurrens</i>	55
	<i>Acacia polyphylla</i>	55
	...	...
Frequent (Total)	<i>100 species selected</i>	5,500
Total	<i>300 species total</i>	10,000




Figure C1. Geographic distributions of *GlobalGeoTree-10kEval-300* and *GlobalGeoTree-10kEval-900*.

Table C3. Overview of the *GlobalGeoTree-10kEval-900* evaluation subset. This subset comprises 300 species selected from each of the three rarity categories (Rare, Common, Frequent), totaling 900 unique species and 10,200 geolocated samples.

Category	Example Species	Num of Samples
Rare	<i>Abies hickelii</i>	4
	<i>Abutilon auritum</i>	4
	<i>Acacia adunca</i>	4
	...	...
Rare (Total)	<i>300 species selected</i>	1,200
Common	<i>Abies fraseri</i>	11
	<i>Acacia echinula</i>	11
	<i>Acacia falcata</i>	11
	...	...
Common (Total)	<i>300 species selected</i>	3,300
Frequent	<i>Abies amabilis</i>	19
	<i>Abies balsamea</i>	19
	<i>Acacia dealbata</i>	19
	...	...
Frequent (Total)	<i>300 species selected</i>	5,700
Total	<i>900 species total</i>	10,200

Appendix D: Model Performance on *GlobalGeoTree-10kEval-300* and *GlobalGeoTree-10kEval-900*

The evaluation results on the larger *GlobalGeoTree-10kEval-300* (Table D1) and *GlobalGeoTree-10kEval-900* (Table D2) subsets align with trends from the primary benchmark *GlobalGeoTree-10kEval* (Table 4 and 7). GeoTreeCLIP consistently outperforms CLIP and RemoteCLIP across all settings (zero-shot, one-shot, three-shot) and taxonomic levels (Family, Genus, Species). Despite lower absolute accuracies on these challenging subsets, especially *GlobalGeoTree-10kEval-900*, GeoTreeCLIP maintains a significant performance edge, highlighting the advantages of its domain-specific pretraining and tailored architecture for multimodal tree species classification.

Table D1. Zero-shot and Few-shot evaluation on *GlobalGeoTree-10kEval-300*. Results are presented as mean accuracy (%) $\pm$ standard deviation (%) over 5 runs.

Taxon.	CLIP		RemoteCLIP		GeoTreeCLIP	
	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
<i>Zero-Shot Evaluation</i>						
Family	7.07 $\pm$ 0.02	20.14 $\pm$ 0.03	1.35 $\pm$ 0.01	7.28 $\pm$ 0.08	12.55 $\pm$ 0.20	40.25 $\pm$ 0.30
Genus	2.34 $\pm$ 0.01	5.59 $\pm$ 0.03	0.57 $\pm$ 0.01	2.10 $\pm$ 0.05	9.26 $\pm$ 0.24	28.34 $\pm$ 0.25
Species	0.46 $\pm$ 0.00	2.06 $\pm$ 0.01	0.56 $\pm$ 0.01	1.70 $\pm$ 0.02	7.87 $\pm$ 0.20	25.20 $\pm$ 0.21
<i>One-Shot Evaluation</i>						
Family	1.55 $\pm$ 0.00	6.80 $\pm$ 0.02	2.43 $\pm$ 0.01	10.30 $\pm$ 0.09	18.58 $\pm$ 0.21	50.41 $\pm$ 0.21
Genus	0.87 $\pm$ 0.01	3.59 $\pm$ 0.01	0.66 $\pm$ 0.01	3.91 $\pm$ 0.07	14.57 $\pm$ 0.26	41.92 $\pm$ 0.27
Species	0.63 $\pm$ 0.00	2.35 $\pm$ 0.01	0.61 $\pm$ 0.01	2.75 $\pm$ 0.03	13.31 $\pm$ 0.22	38.39 $\pm$ 0.26
<i>Three-Shot Evaluation</i>						
Family	5.28 $\pm$ 0.01	16.12 $\pm$ 0.04	3.79 $\pm$ 0.01	11.49 $\pm$ 0.05	23.91 $\pm$ 0.26	57.97 $\pm$ 0.29
Genus	1.79 $\pm$ 0.01	6.02 $\pm$ 0.02	1.11 $\pm$ 0.02	3.60 $\pm$ 0.04	19.22 $\pm$ 0.28	50.35 $\pm$ 0.27
Species	1.33 $\pm$ 0.00	4.17 $\pm$ 0.01	0.74 $\pm$ 0.02	2.64 $\pm$ 0.02	17.90 $\pm$ 0.23	47.54 $\pm$ 0.29

Table D2. Zero-shot and Few-shot evaluation on *GlobalGeoTree-10kEval-900*. Results are presented as mean accuracy (%) $\pm$ standard deviation (%) over 5 runs.

Taxon.	CLIP		RemoteCLIP		GeoTreeCLIP	
	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
<i>Zero-Shot Evaluation</i>						
Family	4.77 $\pm$ 0.04	13.17 $\pm$ 0.05	1.10 $\pm$ 0.00	10.08 $\pm$ 0.02	7.62 $\pm$ 0.13	27.25 $\pm$ 0.05
Genus	0.69 $\pm$ 0.01	2.58 $\pm$ 0.02	0.04 $\pm$ 0.00	0.55 $\pm$ 0.02	4.36 $\pm$ 0.10	16.32 $\pm$ 0.11
Species	0.12 $\pm$ 0.00	0.59 $\pm$ 0.00	0.04 $\pm$ 0.00	0.51 $\pm$ 0.02	3.40 $\pm$ 0.09	11.83 $\pm$ 0.05
<i>One-Shot Evaluation</i>						
Family	3.44 $\pm$ 0.02	9.76 $\pm$ 0.04	1.89 $\pm$ 0.04	4.69 $\pm$ 0.06	14.16 $\pm$ 0.26	41.81 $\pm$ 0.31
Genus	1.53 $\pm$ 0.01	2.82 $\pm$ 0.02	0.25 $\pm$ 0.01	1.21 $\pm$ 0.02	10.24 $\pm$ 0.21	31.45 $\pm$ 0.32
Species	0.45 $\pm$ 0.01	1.27 $\pm$ 0.01	0.22 $\pm$ 0.01	0.83 $\pm$ 0.01	8.18 $\pm$ 0.15	25.17 $\pm$ 0.28
<i>Three-Shot Evaluation</i>						
Family	4.39 $\pm$ 0.02	12.93 $\pm$ 0.04	5.43 $\pm$ 0.09	13.29 $\pm$ 0.02	16.00 $\pm$ 0.14	46.07 $\pm$ 0.34
Genus	2.08 $\pm$ 0.01	5.25 $\pm$ 0.02	0.77 $\pm$ 0.02	4.22 $\pm$ 0.04	12.50 $\pm$ 0.21	37.42 $\pm$ 0.20
Species	1.37 $\pm$ 0.01	3.97 $\pm$ 0.01	0.34 $\pm$ 0.02	1.84 $\pm$ 0.06	10.23 $\pm$ 0.25	31.79 $\pm$ 0.17

Appendix E: Additional Baselines Evaluated on *GlobalGeoTree-10kEval*

To further contextualize GeoTreeCLIP’s performance, we extended our zero-shot evaluation on the *GlobalGeoTree-10kEval* subset to include two additional publicly available vision-language models: SkyCLIP-50 (Wang et al., 2024) and CLIP-laion-RS, a CLIP model pretrained on the remote sensing subset of LAION-2B (Schuhmann et al., 2022). These models were evaluated on the *GlobalGeoTree-10kEval* subset under the same zero-shot protocol used for CLIP and RemoteCLIP (features extracted from individual monthly images, probabilities averaged). For ease of comparison, their performance alongside our GeoTreeCLIP is presented in Table E1.

Table E1. Zero-shot evaluation of SkyCLIP-50, CLIP-laion-RS, and GeoTreeCLIP on *GlobalGeoTree-10kEval*. Results are presented as mean accuracy (%) $\pm$ standard deviation (%) over 5 runs.

Taxon.	SkyCLIP-50		CLIP-laion-RS		GeoTreeCLIP	
	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
Family	2.33 $\pm$ 0.01	18.40 $\pm$ 0.08	1.15 $\pm$ 0.01	17.83 $\pm$ 0.16	20.99 $\pm$ 0.28	56.88 $\pm$ 0.42
Genus	1.10 $\pm$ 0.03	6.48 $\pm$ 0.04	1.12 $\pm$ 0.01	7.33 $\pm$ 0.03	18.39 $\pm$ 0.26	50.98 $\pm$ 0.41
Species	1.10 $\pm$ 0.03	6.36 $\pm$ 0.04	1.12 $\pm$ 0.01	7.27 $\pm$ 0.02	16.71 $\pm$ 0.25	47.52 $\pm$ 0.37

The results in Table E1 show that both SkyCLIP-50 and CLIP-laion-RS, despite their pretraining on remote sensing imagery, achieve zero-shot accuracies on *GlobalGeoTree-10kEval* that are substantially lower than our GeoTreeCLIP. For instance, at the species level, SkyCLIP-50 obtains a top-1 accuracy of 1.10% and CLIP-LAION-RS achieves 1.12%, in contrast to GeoTreeCLIP’s 16.71%. As indicated in the main text for similar baseline models (CLIP, RemoteCLIP), such performance can partly be attributed to their design, which is often optimized for RGB data and lacks effective handling of time-series information or small-patch classification crucial for tree species identification. These limitations further emphasize the value of the GlobalGeoTree benchmark and the effectiveness of our tailored GeoTreeCLIP approach in advancing global tree species classification research.

The comparison highlights that general remote sensing pretraining alone is insufficient for the nuanced task of global fine-grained tree species identification. The domain-specific dataset characteristics, multimodal input integration (including time-series and auxiliary data), and the tailored contrastive learning approach of GeoTreeCLIP appear critical for achieving strong performance on this challenging benchmark.

Author contributions. Conceptualization: Y.M., Z.X., S.M., X.Z.; methodology: Y.M., Z.X., Y.W., S.M., X.Z.; software: Y.M., Z.X., Y.W., X.Z.; results validation: Y.M., Z.X., Y.W., S.M., F.E., H.K., M.K., X.Z.; analysis: Y.M., Z.X., X.Z.; data collection: Y.M., Y.W.; writing, and original draft preparation: Y.M.; paper revision: Y.M., Z.X., Y.W., S.M., F.E., H.K., M.K., X.Z.; visualization: Y.M., Z.X.; supervision: S.M., X.Z.; funding acquisition: X.Z.; project administration: X.Z.; resources: X.Z.

Competing interests. none

Disclaimer. none

Acknowledgements. The work is jointly supported by the German Federal Ministry for Economic Affairs and Climate Action in the framework of the "national center of excellence ML4Earth" (grant number: 50EE2201C), by the European Union's Horizon Europe research and innovations actions programme in the framework of the "Multi-source and Multi-scale Earth observation and Novel Machine Learning Methods for Mineral Exploration and Mine Site Monitoring (MultiMiner)" (grant number: 101091374) and "ThinkingEarth—Copernicus Foundation Models for a Thinking Earth" (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 European Research Council (ERC) (grant agreement No. [ERC-2022-PoC-101101093], Acronym: *EO4FoodSecurity*) by the Excellence Strategy of the Federal Government and the Länder through the TUM Innovation Network EarthCare and by Munich Center for Machine Learning.

References

- Ahlswede, S., S. C. G.-C.-H. P. B. B. F. M. A. F. H. J. D. B. K. B.: TreeSatAI Benchmark Archive: A multi-sensor, multi-label dataset for tree species classification in remote sensing, *Earth Syst. Sci. Data Discuss.*, 2022, 1–22, 2022.
- Aizman, A., Maltby, G., and Breuel, T.: High performance I/O for large scale deep learning, in: 2019 IEEE International Conference on Big Data (Big Data), pp. 5965–5967, IEEE, 2019.
- Beech, E., Rivers, M., Oldfield, S., and Smith, P.: GlobalTreeSearch: The first complete global database of tree species and country distributions, *Journal of sustainable forestry*, 36, 454–489, 2017.
- Bonan, G. B.: Forests and climate change: forcings, feedbacks, and the climate benefits of forests, *science*, 320, 1444–1449, 2008.
- Botella, C., Deneu, B., Marcos, D., Servajean, M., Larcher, T., Leblanc, C., Estopinan, J., Bonnet, P., and Joly, A.: Overview of GeoLifeCLEF 2023: Species composition prediction with high spatial resolution at continental scale using remote sensing, *arXiv preprint arXiv:2509.25816*, 2025.
- Bountos, N. I., Ouaknine, A., Papoutsis, I., and Rolnick, D.: FoMo: Multi-Modal, Multi-Scale and Multi-Task Remote Sensing Foundation Models for Forest Monitoring, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, pp. 27 858–27 868, 2025.
- Bourgoin, C., Verhegghen, A., Carboni, S., Degrève, L., AMEZTOY, A. I., Ceccherini, G., Colditz, R., Achard, F., et al.: Global forest maps for the year 2020 to support the EU regulation on deforestation-free supply chains, 2025.
- Buchhorn, M., Lesiv, M., Tsendbazar, N.-E., Herold, M., Bertels, L., and Smets, B.: Copernicus global land cover layers—collection 2, *Remote Sensing*, 12, 1044, 2020.
- Chamberlain, S. A. and Boettiger, C.: R Python, and Ruby clients for GBIF species occurrence data, *Tech. rep.*, *PeerJ Preprints*, 2017.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale, *arXiv preprint arXiv:2010.11929*, 2020.
- Farias, H. L. S., Silva, W. R., de Oliveira Perdiz, R., Citó, A. C., da Silva Carvalho, L. C., and Barbosa, R. I.: Dataset on wood density of trees in ecotone forests in Northern Brazilian Amazonia, *Data in brief*, 30, 105 378, 2020.
- Felton, A., Petersson, L., Nilsson, O., Witzell, J., Cleary, M., Felton, A. M., Björkman, C., Sang, Å. O., Jonsell, M., Holmström, E., et al.: The tree species matters: Biodiversity and ecosystem service implications of replacing Scots pine production stands with Norway spruce, *Ambio*, 49, 1035–1049, 2020.
- Fick, S. E. and Hijmans, R. J.: WorldClim 2: new 1-km spatial resolution climate surfaces for global land areas, *International journal of climatology*, 37, 4302–4315, 2017.
- Franklin, S. E.: *Remote sensing for sustainable forest management*, CRC Press, 2001.
- Gastauer, M., Leyh, W., and Meira-Neto, J. A.: Tree diversity and dynamics of the forest of Seu Nico, Viçosa, Minas Gerais, Brazil, *Biodiversity Data Journal*, p. e5425, 2015.
- Gaydon, C. and Roche, F.: Pureforest: A large-scale aerial lidar and aerial imagery dataset for tree species classification in monospecific forests, in: 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp. 5895–5904, IEEE, 2025.
- GBIF Secretariat: GBIF Backbone Taxonomy, <https://doi.org/10.15468/39omei>, 2023.
- Gillespie, L. E., Ruffley, M., and Exposito-Alonso, M.: Deep learning models map rapid plant species changes from citizen science and remote sensing data, *Proceedings of the National Academy of Sciences*, 121, e2318296 121, 2024.
- Global Biodiversity Information Facility (GBIF): GBIF Species API Documentation, <https://techdocs.gbif.org/en/openapi/v1/species>, accessed: 2025-05-05, 2025.

- Hamann, A. and Wang, T.: Potential effects of climate change on ecosystem and tree species distribution in British Columbia, *Ecology*, 87, 2773–2786, 2006.
- Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., et al.: High-resolution global maps of 21st-century forest cover change, *science*, 342, 850–853, 2013.
- He, Y., Xu, X., Chen, H., Li, J., and Pu, F.: Visual global-salient guided network for remote sensing image-text retrieval, *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- Hermosilla, T., Bastyr, A., Coops, N. C., White, J. C., and Wulder, M. A.: Mapping the presence and distribution of tree species in Canada’s forested ecosystems, *Remote Sens. Environ.*, 282, 113 276, 2022.
- Jarvis, A., Reuter, H. I., Nelson, A., Guevara, E., et al.: Hole-filled SRTM for the globe Version 4, available from the CGIAR-CSI SRTM 90m Database, 2008.
- Jenkins, J. C., Chojnacky, D. C., Heath, L. S., and Birdsey, R. A.: National-scale biomass estimators for United States tree species, *For. Sci.*, 49, 12–35, 2003.
- Jucker, T., Fischer, F. J., Chave, J., Coomes, D. A., Caspersen, J., Ali, A., Loubota Panzou, G. J., Feldpausch, T. R., Falster, D., Usoltsev, V. A., et al.: Tallo: A global tree allometry and crown architecture database, *Global change biology*, 28, 5254–5268, 2022.
- Kampe, T. U., Johnson, B. R., Kuester, M. A., and Keller, M.: NEON: the first continental-scale ecological observatory with airborne remote sensing of vegetation canopy biochemistry and structure, *Journal of Applied Remote Sensing*, 4, 043 510, 2010.
- Kindt, R.: TreeGOER: A database with globally observed environmental ranges for 48,129 tree species, *Global Change Biology*, 29, 6303–6318, 2023.
- Lane, M. A. and Edwards, J. L.: The global biodiversity information facility (GBIF), *Systematics Association special volume*, 73, 1, 2007.
- Lindenmayer, D., Franklin, J., and Fischer, J.: General management principles and a checklist of strategies to guide forest biodiversity conservation, *Biological conservation*, 131, 433–445, 2006.
- Liu, F., Chen, D., Guan, Z., Zhou, X., Zhu, J., Ye, Q., Fu, L., and Zhou, J.: Remoteclip: A vision language foundation model for remote sensing, *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- Loshchilov, I. and Hutter, F.: Decoupled weight decay regularization, *arXiv preprint arXiv:1711.05101*, 2017.
- Loshchilov, I. and Hutter, F.: SGDR: Stochastic Gradient Descent with Warm Restarts, in: *International Conference on Learning Representations*, 2022.
- Mauri, A., S. G. S.-M.-A. J.: EU-Forest, a high-resolution tree occurrence dataset for Europe, *Sci. Data*, 4, 1–8, 2017.
- Micikevicius, P., Narang, S., Alben, J., Diamos, G., Elsen, E., Garcia, D., Ginsburg, B., Houston, M., Kuchaiev, O., Venkatesh, G., et al.: Mixed precision training, *arXiv preprint arXiv:1710.03740*, 2017.
- Mu, Y., Guo, J., Shahzad, M., and Zhu, X. X.: National-scale tree species mapping with deep learning reveals forest management insights in Germany, *International Journal of Applied Earth Observation and Geoinformation*, 139, 104 522, 2025a.
- Mu, Y., Xiong, Z., Wang, Y., Shahzad, M., Essl, F., Kreft, H., Kleunen, M. v., and Zhu, X. X.: GlobalGeoTree: A Multi-Granular Vision-Language Dataset for Global Tree Species Classification, <https://doi.org/10.15468/dd.9qxqyy>, 2025b.
- Oord, A. v. d., Li, Y., and Vinyals, O.: Representation learning with contrastive predictive coding, *arXiv preprint arXiv:1807.03748*, 2018.
- Ouaknine, A., Kattenborn, T., Laliberté, E., and Rolnick, D.: OpenForest: a data catalog for machine learning in forest monitoring, *Environmental Data Science*, 4, e15, 2025.
- Parnami, A. and Lee, M.: Learning from few examples: A summary of approaches to few-shot learning, *arXiv preprint arXiv:2203.04291*, 2022.

- Pazos-Outón, L. M., Vasconcelos, C. N., Raichuk, A., Arnab, A., Morris, D., and Neumann, M.: Planted: a dataset for planted forest identification from multi-satellite time series, in: IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium, pp. 7066–7070, IEEE, 2024.
- Picek, L., Botella, C., Servajean, M., Leblanc, C., Palard, R., Larcher, T., Deneu, B., Marcos, D., Bonnet, P., and Joly, A.: Geoplant: Spatial plant species prediction dataset, *Advances in Neural Information Processing Systems*, 37, 126 653–126 676, 2024.
- Poggio, L., De Sousa, L. M., Batjes, N. H., Heuvelink, G. B., Kempen, B., Ribeiro, E., and Rossiter, D.: SoilGrids 2.0: producing soil information for the globe with quantified spatial uncertainty, *Soil*, 7, 217–240, 2021.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners, *OpenAI blog*, 1, 9, 2019.
- 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, PmLR, 2021.
- Rosenblatt, F.: The perceptron: a probabilistic model for information storage and organization in the brain., *Psychological review*, 65, 386, 1958.
- Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models, *Advances in neural information processing systems*, 35, 25 278–25 294, 2022.
- Spoto, F., Sy, O., Laberinti, P., Martimort, P., Fernandez, V., Colin, O., Hoersch, B., and Meygret, A.: Overview of sentinel-2, in: *2012 IEEE international geoscience and remote sensing symposium*, pp. 1707–1710, IEEE, 2012.
- Stevens, S., Wu, J., Thompson, M. J., Campolongo, E. G., Song, C. H., Carlyn, D. E., Dong, L., Dahdul, W. M., Stewart, C., Berger-Wolf, T., et al.: Bioclip: A vision foundation model for the tree of life, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 19 412–19 424, 2024.
- Van der Maaten, L. and Hinton, G.: Visualizing data using t-SNE., *Journal of machine learning research*, 9, 2008.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I.: Attention is all you need, *Advances in neural information processing systems*, 30, 2017.
- Wang, Z., Prabha, R., Huang, T., Wu, J., and Rajagopal, R.: Skyscript: A large and semantically diverse vision-language dataset for remote sensing, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, pp. 5805–5813, 2024.
- Weiser, H., Schäfer, J., Winiwarter, L., Krašovec, N., Fassnacht, F. E., and Höfle, B.: Individual tree point clouds and tree measurements from multi-platform laser scanning in German forests, *Earth System Science Data*, 14, 2989–3012, 2022.
- Wellbrock, N., Eickenscheidt, N., Hilbrig, L., Dühnelt, P.-E., Holzhausen, M., Bauer, A., Dammann, I., Strich, S., Engels, F., and Wauer, A.: Leitfaden und Dokumentation zur Waldzustandserhebung in Deutschland, *Tech. rep.*, Thünen Working Paper, 2018.
- Zhang, X., Liu, L., Zhao, T., Zhang, W., Guan, L., Bai, M., and Chen, X.: GLC_FCS10: a global 10 m land-cover dataset with a fine classification system from Sentinel-1 and Sentinel-2 time-series data in Google Earth Engine, *Earth System Science Data*, 17, 4039–4062, 2025.