

AGBD: A Global-scale Biomass Dataset

GhJulia Sialelli¹, Torben Peters¹, Jan D. Wegner², Konrad Schindler¹

¹ Photogrammetry and Remote Sensing, ETH Zurich - (gsialelli, tpeters, schindler)@ethz.ch

² EcoVision Lab, DM3L, University of Zurich - jandirk.wegner@uzh.ch

Keywords: Remote Sensing, Biomass Estimation, Machine Learning, Dataset.

Abstract

Accurate estimates of Above Ground Biomass (AGB) are essential in addressing two of humanity's biggest challenges: climate change and biodiversity loss. Existing datasets for AGB estimation from satellite imagery are limited. Either they focus on specific, local regions at high resolution, or they offer global coverage at low resolution. There is a need for a machine learning-ready, globally representative, high-resolution benchmark dataset. Our findings indicate significant variability in biomass estimates across different vegetation types, emphasizing the necessity for a dataset that accurately captures global diversity. To address these gaps, we introduce a comprehensive new dataset that is globally distributed, covers a range of vegetation types, and spans several years. This dataset combines AGB reference data from the GEDI mission with data from Sentinel-2 and PALSAR-2 imagery. Additionally, it includes pre-processed high-level features such as a dense canopy height map, an elevation map, and a land-cover classification map. We also produce a dense, high-resolution (10 m) map of AGB predictions for the entire area covered by the dataset. Rigorously tested, our dataset is accompanied by several benchmark models and is publicly available. It can be easily accessed using a single line of code, offering a solid basis for efforts towards global AGB estimation.

1. Introduction

Measuring Above Ground Biomass (AGB) is crucial to combat biodiversity loss and address the ongoing climate crisis. Accurate AGB estimates enable the quantification of carbon stocks, which play a pivotal role in carbon offsetting schemes. Furthermore, AGB is correlated with various biodiversity metrics, providing valuable insights into the structure of biodiversity hotspots (Vance-Chalcraft et al., 2010, Sonkoly et al., 2019, Li et al., 2018). Traditional methods for measuring AGB are based on field work, either via destructive sampling, where sample trees are cut and weighed; or with non-destructive techniques that use tree dimensions (like tree height and trunk diameter at breast height) in tailored regression equations. Increasingly, LiDAR campaigns are replacing manual measurements with Remote Sensing (RS), becoming the preferred method. However, despite their precision, the high costs and time requirements limit their global scalability. As a result, previous work has mainly focused on localized AGB estimation. Launched in 2019, the NASA GEDI mission uses an on-orbit laser altimeter aboard the International Space Station to scan the globe between 51.6° north and south. Calibrated on coincident airborne LiDAR data and ground plot field inventories, the mission provides sparse AGB estimates across its observation domain. These estimates represent the largest existing reference AGB dataset, which can be combined with Machine Learning (ML) methods and RS data to advance AGB estimation across the globe. We provide a ML-ready dataset consisting of coincident GEDI AGB estimates and various RS data products. Although freely accessible, obtaining remotely sensed data globally is tedious, and the sheer scale of it is prohibitive for ML research. The corresponding global dataset would require approximately 70TB of storage, ≈ 60 times more than the complete ImageNet. For AGB estimation that data will need extensive preprocessing, and it would be too large to repeatedly train and test within reasonable time during method development. To reduce the scale, but to the greatest possible extent maintain the representativeness of the dataset, we derive a subset of regions whose vegetation distribution emulates the

global one. The complex interactions between AGB and various vegetation types (Chen et al., 2023) motivate this choice.

The dataset serves multiple purposes. (i) **Providing a globally representative high-res testbed for AGB estimation:** We show that all existing datasets are either too localized to generalize around the globe, or have only low resolution. Our collection covers all biomes and is globally representative at high resolution, enabling the training of AGB estimation models with improved performance and better generalization abilities. (ii) **Improving Regional Performance:** Prior research has shown that combining GEDI data with local reference data yields better results than using only the local data alone (Bullock et al., 2023). Researchers can leverage our comprehensive dataset for initial training and then fine-tune their models with local data specific to their regions, enhancing the accuracy and performance of their analyses.

To highlight the accessibility of our dataset, we offer a fully pre-processed version that is compatible with all major ML frameworks, such as TensorFlow and PyTorch. It is hosted on HuggingFace (Lhoest et al., 2021) and can be downloaded and used with just the following lines of code:

```
#!pip install datasets
from datasets import load_dataset
dataset = load_dataset("prs-eth/AGBD", streaming=True)["train"] # or test, val
```

Our contributions are:

- We offer a machine learning-ready, easily accessible, globally representative dataset comprising coincident AGB estimates together with various remote sensing data products.
- We conduct a comprehensive analysis of our dataset and apply several standard models to it to validate its accuracy and reliability and to serve as baselines.
- We generate and provide a dense, high-resolution (10 m) map of AGB predictions across the dataset's coverage area with the best-performing baseline.

- We release all benchmark models and pretrained weights.

All code is hosted on Github, moreover the original raw data, the dense predictions, and the model weights are hosted on the ETH Research Collection.

2. Background

In recent years, remote sensing (RS) and machine learning (ML) have advanced various mapping tasks, ranging from human population (Metzger et al., 2023) to canopy height (Lang et al., 2023). Biomass estimation also has seen significant research, with (Zhang et al., 2019, Araza et al., 2022) providing a review of existing AGB datasets and maps, supplemented by Table 1. Currently, global-scale AGB datasets suffer from low spatial resolution (at best 100 meters), while high-resolution datasets (10-30 meters) are limited to restricted geographic areas. This dichotomy forces researchers to strike a compromise between generalization across space and spatial resolution. Notably, (Araza et al., 2022) highlight the challenge of achieving consistent accuracy across continental and global-scale AGB maps, due to the absence of a comprehensive global reference dataset. Furthermore, (Leonhardt et al., 2022) underscore the necessity for a dataset featuring globally distributed biomass reference records to improve regional studies. Similarly, (Nascetti et al., 2024) advocate for the release of more openly accessible, deep learning-ready data. (Tolan et al., 2024) highlight another aspect of this issue when they apply a correction based on GEDI to their recent 1-meter resolution canopy height maps for São Paulo (Brazil) and California (USA), aiming to compensate for the limited geographic diversity of their reference data. Taken together, these observations point at a critical gap in the literature: there is a need for a robust, globally diverse dataset to support high-resolution AGB mapping.

In their comprehensive review of deep learning for remote sensing of forest inventories, (Hamedianfar et al., 2022) identify two critical avenues for progress. The first is the integration of data from multiple scales and sources, which has been shown to enhance predictive accuracy. For instance, incorporating diverse data modalities, such as DEM-derived Canopy Height Models (CHM) (Puliti et al., 2020) and Synthetic Aperture Radar (SAR) data (Liu et al., 2023b), not only alleviates issues of data saturation but also improves the predictive accuracy for high biomass values (Zhang et al., 2019). The second lever is the recognition and differentiation of various tree species, which can significantly refine the granularity and precision of forest inventory analyses. These strategies underscore the potential of deep learning to revolutionize forest monitoring and management through richer data integration and species-specific insights.

Our dataset fills in these gaps: it has a nominal resolution of 10 m, is representative of the global land cover and vegetation distribution, and assembles a range of data sources whose synergies have — to the best of our knowledge — not yet been systematically investigated for the task at hand.

3. Dataset

In this section, we discuss the rationale for creating the dataset, detail its components, and describe the process of its generation. For detailed information on data formats and downloading instructions, please refer to the supplementary material and the to-be-released project website.

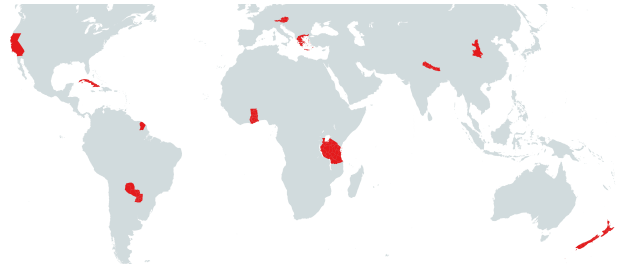


Figure 1. The regions of interest: California (USA), Cuba, Austria, Greece, Nepal, Shaanxi (China), French Guiana, Paraguay, Ghana, Tanzania, New Zealand.

3.1 Vegetation types analysis

Land Cover Classification Map (LC). To correctly represent the world's vegetation types, we turn to the Copernicus Global Land Service Dynamic Land Cover map (Buchhorn et al., 2020), which delivers a global map of the land cover at 100 m spatial resolution. It consists of a discrete classification system with 21 classes (following UN-FAO's Land Cover Classification System) that differentiates between open (O) and closed (C) forest types, with the following sub-types: evergreen needle-leaved (ENL), deciduous needle-leaved (DNL), evergreen broad-leaved (EBL), deciduous broad-leaved (DBL), mixed type (MT), and unknown type (UT). Additional land cover classes include shrubland (Shrubs), herbaceous vegetation (HV), herbaceous wetland (HW), moss and lichen (Moss), bare or sparse vegetation (Bare), cropland (Crops), built-up areas (ignored), snow and ice (ignored), and permanent water bodies (ignored). We refer the reader to (Buchhorn et al., 2020) for detailed descriptions of each class. We base our analysis on the 2019 iteration of the product (latest available). Note that while the CGLS-LC100 map is global, we only consider vegetation types that are found within the GEDI coverage, as only there we have reference AGB data. As a consequence, vegetation types found only outside of the GEDI coverage are not accounted for (the grey biomes in Figure 2). These are: open and closed DNL only found in Siberia; and moss and lichen, mostly found in Northern Canada, Siberia and Greenland.




Figure 2. Per-biome distribution of GEDI AGBD values and ESA CCI residuals. Gray biomes are outside of GEDI coverage.

A globally representative vegetation distribution. The ESA CCI map is the most recent and adequate global AGB map available. It is provided by the European Space Agency (ESA), through its Climate Change Initiative (CCI) Biomass project. Global maps of AGB are provided for five epochs (2010, 2017, 2018, 2019 and 2020) at 100 m resolution. We observe that this map performs variably across different vegetation types, as shown in Figure 2. We utilise CCI AGB estimates and GEDI footprints from the year 2019 to analyze residuals across various vegetation types within the GEDI coverage. Our analysis revealed significant variances, with underestimations in moss and

	Method(s)	Geographical Extent	Resolution	Limitation	Reference data	Input data
(Nascetti et al., 2024)	U-TAE, Swin U-Net TTransformer, U-Net++	Finland	10m	Localized	Airborne LiDAR and allometry	S1, S2
(Liu et al., 2023b)	RF, SVM, CNN	Southwest China	10m	Localized	Destructive AGB measurements	S1, S2
(Puliti et al., 2020)	RF	Northern Norway	10m	Localized	Field measurements and allometry	ArcticDEM, S2
(Liu et al., 2023a)	UNet	Europe	30m	Localized	Field measurements, airborne LiDAR, allometry	PlanetScope, LiDAR
(Leonhardt et al., 2022)	CGAN	Northwestern USA	30m	Localized	US Carbon Monitoring System AGB map	Landsat-8, ALOS-2 PALSAR-2
(Xu et al., 2017)	Maximum Entropy estimator	Democratic Republic of Congo	100m	Localized, low-resolution	Airborne LiDAR and allometry	Landsat-8, ALOS-2 PALSAR-2, SRTM
(Pascarella et al., 2023)	UNet	Vietnam, Myanmar	100m	Localized, low-resolution	ESA CCI Map	S2

Table 1. Recent works and limitations.




Figure 3. Land cover distribution across the world (left) and across our subset (right).

lichen and overestimations in closed evergreen broadleaf forests, compelling us to adopt a design that accounts for differences between vegetation types. In the absence of a truly global dataset and facing the logistical barrier of large-scale remote sensing data downloads, we argue for a reasonably-sized dataset that mirrors the global vegetation distribution. To that end we select the following regions: California (USA), Cuba, Austria, Greece, Nepal, Shaanxi (China), French Guiana, Paraguay, Ghana, Tanzania and New Zealand. They are strategically chosen for their geographic and ecological diversity, as illustrated in Figures 1 and 3. This selection was based on a careful process. We chose moderate-sized countries or large administrative regions to avoid areas that were too large or too small. We considered the number of available GEDI footprints, as regions further from the equator have higher revisit times, affecting data density. To ensure diverse ecological representation, we used the Shannon entropy and Simpson's diversity index, based on land cover pixel values, to measure the biome diversity. Finally, regions were selected from various continents, with a balance above, below, and on the equator. Both the vegetation types and the ESA CCI residuals within our subset are in fairly good agreement with the global ones (Figure 2).

3.2 Modalities

In the following we describe the data sources, how they were processed, and how we assembled them into the final dataset. An overview is provided in Figure 4. A more detailed explanation of each individual data field can be found in the Appendix (Table 3 and Section 8.1).

GEDI, an altimeter mission (2018-2023) on the International Space Station, produced high-resolution laser observations of the Earth's 3D structure. Each observation consists of a 25m diameter footprint, with an associated AGB estimate. Due to its sampling pattern, GEDI has a (mostly) global yet sparse coverage: individual footprints from 51.6°N to 51.6°S. We downloaded all GEDI L4A (Dubayah et al., 2022) (v2.1) footprints for the regions of interest from the start of the mission until Dec. 31st 2020. We only kept biomass values in the range [0, 500] Mg/ha, as suggested by previous literature (Carreiras et al., 2017); we only retain the "power" beams, as previous work (Lanfranchi et al., 2022, Lahssini et al., 2022) has found that they provide more reliable estimates than the "coverage" beams; and as suggested by the GEDI team, we discard footprints where $sensitivity < 0.95$, $14_quality_flag \neq 1$, and $selected_algorithm = 10$. Details on the process by which GEDI estimates the AGB can be found in the Appendix, in Section 8.6.

Sentinel-2 (S2), a European multi-spectral imaging mission, delivers high-resolution imagery with a revisit frequency of 5 days at the equator. We downloaded selected Level-1C products with less than 20% cloud coverage from October 2018 to December 2020 from PEPS, and processed them to Level-2A using Sen2Cor (Main-Knorn et al., 2017) (v2.11). The pre-processing of the products involves the conversion of raw digital numbers (DN) to surface reflectance values via $SR = \frac{DN + BOA_offset}{10,000}$, with BOA_offset equal to 0 prior to January 2022, and -1000

after that. We use all 12 available spectral bands as input features, upsampling the 20m and 60m ones to 10m resolution. We use the (upsampled) SCL band for patch selection purposes.

PALSAR-2 (P2), on board the ALOS-2 satellite, is a Synthetic Aperture Radar (SAR) that can capture images at night and through clouds, making it a 24-hour, all-weather technology. The 25 m PALSAR-2 yearly mosaic (Shimada et al., 2014) is a seamless global image in WGS84 (EPSG:4326) projection, created by merging strips of SAR imagery for a year, featuring both HH and HV polarization backscatter. It uses a different gridding system from ESA, so for each year we download the necessary ALOS tiles, then mosaic, crop and reproject them to align with Sentinel-2 tiles. The provided digital numbers (DN) are linear backscatter amplitudes. Pre-processing involves converting DN values to γ_0 values in decibels via $\gamma^0 = 10 \cdot \log_{10}(DN^2) - 83$.

Digital Elevation Model (DEM). The ALOS Global Digital Surface Model (AW3D30) (Tadono et al., 2014) offers 30-meter resolution elevation data. As for P2 mosaics, we downloaded, mosaicked, and reprojected the DEM to align with Sentinel-2 tiles.

Land Cover (LC). We downloaded the Copernicus Global Land Service Dynamic Land Cover map (Section 3.1) at 100-meter resolution, upsampled, cropped and reprojected it to match the Sentinel-2 tiles. The map includes two layers: the discrete classification, and the probability of the discrete classification (which acts as a quality indicator).

Canopy Height (CH). (Lang et al., 2023) provide source code for canopy height mapping from Sentinel-2 imagery, using probabilistic deep learning. The model estimates CH for each input Sentinel-2 L2A image and merges redundant predictions from different dates to create a yearly CH map for each Sentinel-2 tile, along with a map of the associated standard deviations (STD). We follow this procedure to obtain yearly CH maps for 2019 and 2020 for all Sentinel-2 tiles within our regions of interest.

Coordinates. We additionally provide encoded latitude and longitude information to the model.

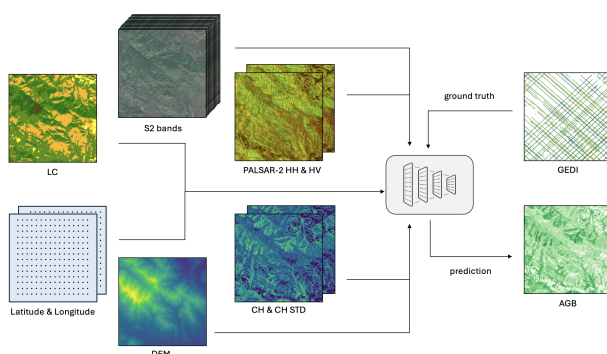


Figure 4. Visualization of all data sources along with our training pipeline. Note the sparse GEDI labels, which are upgraded to dense AGB maps by the prediction.

3.3 Dataset Assembly and Train/Test Split

To create an analysis-ready dataset, we crop all data into *patches* as follows. First, all data sources are up-sampled to a common resolution of 10 m, using nearest neighbor for categorical

variables (the S2 SCL band, and the LC probability band) and bi-linear interpolation for continuous variables (all others). We then iterate over each GEDI footprint, and fetch the corresponding data: for S2, the closest product within the six months preceding the footprint acquisition; for the yearly available data products (CH and P2), the data from the year of the footprint acquisition; and the product itself for non time-dependent data products (DEM and LC). Once all data sources are temporally aligned, we crop all raster data to 25×25 pixel squares centered on the GEDI footprint, and store them in a HDF5 file.

The splitting into training, validation and test sets is not done at the level of GEDI footprints, but at the level of the Sentinel-2 tiles, meaning that all patches within the same $100 \times 100 \text{ km}^2$ tile will belong to the same set. To align with our geographical distribution and vegetation types, we split each country/region into training, validation and test portions and aim to match the distribution of vegetation types within each portion to that of the whole region. To maintain our geographical distribution and vegetation types distribution, we perform a train/validation/test split for each region. We enforce that the vegetation types distribution within each region's split should match the whole region's distribution. This is done by randomly exploring the combinatorial assignment of tiles to the three splits until a sufficiently close match has been found. In each region 65% of the tiles are used for training, 15% for validation, and the remaining 20% for testing. The detailed assignment per tile can be found on our project website.

4. Experiments

4.1 Models

Early works on biomass estimation employed linear regression models, but the relationship between satellite observations and AGB is more complex and can be better modeled with high-capacity ML methods (Li et al., 2020, Liu et al., 2023b). We provide a selection of pertinent ML models as benchmarks for the task.

GBDT. We start with a `lightgbm` implementation of Gradient Boosted Decision Trees (GBDT). This model does not leverage the spatial context information of the patches but performs regression based only on the central pixel. The model was trained by minimising the RMSE loss, with early stopping enabled. Model hyper-parameters can be found in the Appendix.

For the following deep learning methods, we always use a patch-wise training procedure: each patch (of size 25×25 or 15×15) has *one* ground-truth pixel, its center. The model emits a prediction for each pixel in the patch, but only the central pixel prediction contributes to the loss. The optimisation of the network parameters is done in the standard way with mini-batch stochastic gradient descent, using the ADAM (Kingma and Ba, 2014) variant and batch size 256. Each model is trained five times with different random seeds to account for stochastic variability, and we report the mean RMSE and standard deviation over the five runs. All models run on consumer hardware, specifically we used NVIDIA GeForce RTX 2080 Ti GPUs available through our institution's high-performance cluster. The training pipeline is illustrated in Figure 4.

FCN (0.5M parameters). This model is a straightforward Fully Convolutional Neural Network (FCN) consisting of convolutional layers with batch normalization and ReLU activations.

UNet (10M parameters). A standard UNet implementation with an initial double convolution block, followed by symmetric

downsampling and upsampling layers with skip connections to preserve spatial detail, and a final output convolution to produce a single-channel output. The number of down- and up-sampling layers depends on the patch size, with larger patches going through more layers.

Lang et al. (13M parameters). The architecture developed by (Lang et al., 2023) for CH estimation, but trained (from scratch) for AGB estimation, with standard ℓ_2 -loss (rather than the original log-likelihood loss). It consists of a series of residual blocks with separable convolutions *without* any downsampling. We use 728 filters per block (instead of 256).

4.2 Results

All quantitative results are reported in Table 2. The lowest error is achieved by the Lang et al. architecture, trained on all considered input channels but CH. The best models achieve RMSE values slightly below 60 Mg/ha, consistent with previous literature. We elaborate on various aspects in the following paragraphs.

Role of input features. In order to study the impact of each data source, we conduct a series of experiments. We vary the input features available to each model: all features, only the canopy height (**CH**), only the 10 m RGB (B02, B03, B04) and NIR (B08) S2 bands (**RGBN**), all S2 bands (**S2 (all)**), all S2 bands plus the LC map (**LC**), all S2 bands plus the DEM (**DEM**), all S2 bands and the two ALOS-2 PALSAR-2 bands (**P2**), and all features except for the CH. Per-pixel latitude and longitude are always provided as inputs. In each setting, the 5 models with the highest validation scores are applied to the held-out test set to obtain the mean RMSE values and associated standard deviations reported in Table 2. The best model is the one which has access to additional input features (i.e., on top of S2); we attribute this to the comprehensive, complementary evidence that the model can rely on (namely optical, SAR, elevation, and vegetation).

Patch size. We consider patches of size consistent with the literature (Lang et al., 2023, Tolan et al., 2024). Namely, to study how varying the patch-size affects model performance, we train and evaluate with two different patch sizes: 15×15 and 25×25 pixels (see Table 2). Our experiments showed minimal performance differences, favoring the 25×25 pixels patch size. Given the limited spatial context of most existing biomass retrieval methods, it is interesting that increasing the context window yields better results. At present it is unclear whether this is due to real context information or simply a data augmentation effect.

Binned residuals analysis. Expanding on the previous analysis, we also analyse the residual distribution across various AGB intervals, see Figure 5. Besides our baseline models described in Section 4.1 we also include the ESA CCI AGB map, upsampled to 10 m resolution, in this analysis. That comparison is directly relevant, since ESA CCI also uses GEDI to calibrate their model, according to their Algorithm Theoretical Basis Document (Santoro and Cartus, 2023, p.23). We point out that the distribution of AGB values is skewed: most (50%) of all GEDI labels fall into the first bin ($[0, 50]$).

Per biome residuals analysis. To emphasize how much AGB estimates depend on the biome at hand, we plot the residuals across the biomes, in Figure 6.




Figure 5. Binned test residuals for the best-performing model of each architecture, and for the ESA CCI predictions.




Figure 6. Per-biome test residuals for the best-performing model of each architecture, and for the ESA CCI predictions.

Saturation analysis. A known issue for AGB estimation is *saturation*, which happens when beyond a certain AGB threshold, the RS information no longer reflects the change in AGB, making it challenging to estimate high values. To investigate whether including additional data sources mitigates with this issue, we include the binned residuals of the Lang et al. models trained on the various feature combinations. The results are shown in Figure 7, and indicate that having access to the additional features makes the AGB estimates better particularly in the higher bins.




Figure 7. Binned mean test RMSE (\downarrow) and associated standard deviation per model, with various inputs.

Qualitative results inspection. For a visual impression, Figure 8 shows the predictions of the best-performing model versus the CCI map for a selected Sentinel-2 tile (30NXM, located in Ghana). As the GEDI L4A reference footprints are too small and sparse to visualize, we include the GEDI L4B 1km product (Dubayah et al., 2023), which provides statistically inferred estimates of mean AGB from the footprints located within the borders of each 1 km cell. This figure illustrates how the high spatial resolution of our estimates is able to preserve spatial details that are not captured in the CCI product, meaning that they would not be restored by conventional upscaling of existing maps.

5. Discussion

Results. Our various analyses indicate that including input features beyond Sentinel-2 is beneficial for the task of AGB estimation, and that more complex models tend to better capture the complex relationship between input features and AGB.

Features						Methods						
CH	RGBN	S2 (all)	LC	DEM	P2	GBDT (1×1)	FCN (15×15) (25×25)	UNet (15×15) (25×25)	Lang et al. (15×15) (25×25)			
×						70.73 ±0.18	69.29 ±0.16	69.33 ±0.07	69.08 ±0.08	68.54 ±0.12	68.54 ±0.27	68.11 ±0.24
						72.92 ±0.18	64.97 ±0.25	64.75 ±0.11	63.14 ±0.29	61.97 ±0.09	61.60 ±0.13	60.39 ±0.19
×						69.41 ±0.11	63.52 ±0.22	63.48 ±0.08	62.00 ±0.09	60.69 ±0.05	60.36 ±0.13	59.15 ±0.22
						66.48 ±0.05	62.40 ±0.19	62.47 ±0.12	61.09 ±0.18	59.99 ±0.12	59.65 ±0.10	58.82 ±0.14
×						67.57 ±0.06	62.43 ±0.24	61.78 ±0.69	60.47 ±0.67	59.33 ±0.40	58.55 ±0.19	57.90 ±0.21
						67.27 ±0.08	61.95 ±0.19	61.81 ±0.18	60.59 ±0.14	59.27 ±0.11	59.12 ±0.11	58.22 ±0.23
×						65.16 ±0.02	60.97 ±0.05	60.97 ±0.19	59.84 ±0.10	58.50 ±0.16	58.16 ±0.15	57.14 ±0.29
						63.76 ±0.09	60.98 ±0.23	60.92 ±0.23	60.18 ±0.36	58.33 ±0.47	58.04 ±0.31	57.52 ±0.41

Table 2. Mean test RMSE (\downarrow) and associated standard deviation per model, with various inputs. Crosses denote the *presence* of a feature, values in brackets denote patch size. Values are colored from dark red (higher) to dark blue (lower).




Figure 8. Sentinel-2 tile 30NXM (Ghana), our best model's prediction, the corresponding ESA CCI map, and GEDI L4B product. Global view (top) and zoomed in views (bottom).

Stand-alone task. As AGB is highly dependent on CH (via allometric equations), one could wonder whether CH information is enough to estimate AGB. We conclude that it is not the case: our results show that learning AGB "from scratch" (using S2, P2, LC and DEM) yields better results than learning from CH directly. While this applies to the specific Lang et al. CH product at hand, the most recent ones published (Pauls et al., 2024, Tolan et al., 2024) also rely on GEDI CH estimates, so we expect similar results. We thus argue that AGB estimation should be tackled as its own task.

Limitations. Despite the nominal 10 m resolution of our estimates, the effective spatial resolution at which biomass is identified is lower. Additionally, the uncertainty in the footprint-level GEDI AGBD product used as reference data influences the accuracy and uncertainties of the model predictions. We elaborate on these points in the Appendix, Section 8.7.

6. Conclusion

We provide a ML-ready, unrestricted, easily accessible dataset for high-resolution biomass estimation from remote sensing data. This dataset consists of $\approx 16 \cdot 10^6$ patches, covering a diverse and representative selection of geographic locations with a total of $\approx 500 \cdot 10^6$ hectares and two different years (2019 and 2020). Furthermore, we complement the dataset with a suite of baseline architectures and trained weights for others to build upon. We hope that our dataset may help to reduce the geographical and ecosystem-related biases of biomass models, enhance the accuracy of future biomass maps, and perhaps contribute towards a regular, operational high-resolution monitoring system.

7. Acknowledgments

The data products used for this paper have been collected from various sources:

- the Global PALSAR-2/PALSAR/JERS-1 Mosaics and Forest/Non-Forest Maps of the Japan Aerospace Exploration Agency;
- the ALOS Global Digital Surface Model "ALOS World 3D - 30m (AW3D30)" of the Japan Aerospace Exploration Agency;
- the Land Cover 2019 (v3) of the European Union's Copernicus Land Monitoring Service;
- the Sentinel-2 L1C products, courtesy of Copernicus, downloaded via the French PEPS platform;
- the GEDI L4A product, courtesy of the NASA GEDI Team, downloaded from Oak Ridge National Laboratory DAAC.

We thank them for making this data openly available for academic purposes.

References

- Araza, A., de Bruin, S., Herold, M., Quegan, S., Labriere, N., Rodriguez-Veiga, P., Avitabile, V., Santoro, M., Mitchard, E. T., Ryan, C. M., Phillips, O. L., Willcock, S., Verbeeck, H., Carreiras, J., Hein, L., Schelhaas, M.-J., Pacheco-Pascagaza, A. M., da Conceição Bispo, P., Laurin, G. V., Vieilledent, G., Slik, F., Wijaya, A., Lewis, S. L., Morel, A., Liang, J., Sukhdeo, H., Schepaschenko, D., Cavlovic, J., Gilani, H., Lucas, R., 2022. A comprehensive framework for assessing the accuracy and uncertainty of global above-ground biomass maps. *Remote Sensing of Environment*, 272, 112917. <http://dx.doi.org/10.1016/j.rse.2022.112917>.
- Buchhorn, M., Smets, B., Bertels, L., Lesiv, M., Tsendbazar, N.-E., Masiliunas, D., Linlin, L., Herold, M., Fritz, S., 2020. Copernicus global land service: Land cover 100m: Collection 3: epoch 2019: Globe (version v3.0.1) [data set].
- Bullock, E. L., Healey, S. P., Yang, Z., Acosta, R., Villalba, H., Insfrán, K. P., Melo, J. B., Wilson, S., Duncanson, L., Næsset, E., Armston, J., Saarela, S., Ståhl, G., Patterson, P. L., Dubayah, R., 2023. Estimating aboveground biomass density using hybrid statistical inference with GEDI lidar data and Paraguay's national forest inventory. *Environmental Research Letters*, 18(8), 085001. <http://dx.doi.org/10.1088/1748-9326/acdf03>.
- Carreiras, J., Quegan, S., Le Toan, T., Ho, D., Tong Minh, D. H., Saatchi, S., Carvalhais, N., Reichstein, M., Scipal, K., 2017. Coverage of high biomass forests by the ESA BIOMASS mission under defense restrictions. *Remote Sensing of Environment*, 196, 154–162.
- Chen, X., Luo, M., Larjavaara, M., 2023. Effects of climate and plant functional types on forest above-ground biomass accumulation. *Carbon Balance and Management*, 18(1). <http://dx.doi.org/10.1186/s13021-023-00225-1>.
- Dubayah, R., Armston, J., Healey, S., Yang, Z., Patterson, P., Saarela, S., Stahl, G., Duncanson, L., Kellner, J., Bruening, J., Pascual, A., 2023. Gedi 14b gridded aboveground biomass density, version 2.1.
- Dubayah, R., Armston, J., Kellner, J., Duncanson, L., Healey, S., Patterson, P., Hancock, S., Tang, H., Bruening, J., Hofton, M., Blair, J., Luthcke, S., 2022. Gedi 14a footprint level aboveground biomass density, version 2.1.
- Hamedianfar, A., Mohamedou, C., Kangas, A., Vauhkonen, J., 2022. Deep learning for forest inventory and planning: a critical review on the remote sensing approaches so far and prospects for further applications. *Forestry: An International Journal of Forest Research*, 95(4), 451–465. <http://dx.doi.org/10.1093/forestry/cpac002>.
- Hancock, S., Armston, J., Hofton, M., Sun, X., Tang, H., Duncanson, L. I., Kellner, J. R., Dubayah, R., 2019. The GEDI Simulator: A Large-Footprint Waveform Lidar Simulator for Calibration and Validation of Spaceborne Missions. *Earth and Space Science*, 6(2), 294–310. <http://dx.doi.org/10.1029/2018EA000506>.
- Hunka, N., Santoro, M., Armston, J., Dubayah, R., McRoberts, R. E., Næsset, E., Quegan, S., Urbazaev, M., Pascual, A., May, P. B., Minor, D., Leitold, V., Basak, P., Liang, M., Melo, J., Herold, M., Málaga, N., Wilson, S., Durán Montesinos, P., Arana, A., Ernesto De La Cruz Paiva, R., Ferrand, J., Keoka, S., Guerra-Hernández, J., Duncanson, L., 2023. On the NASA GEDI and ESA CCI biomass maps: aligning for uptake in the UNFCCC global stocktake. *Environmental Research Letters*, 18(12), 124042. <http://dx.doi.org/10.1088/1748-9326/ad0b60>.
- Kellner, J. R., Armston, J., Duncanson, L., 2023. Algorithm Theoretical Basis Document for GEDI Footprint Above-ground Biomass Density. *Earth and Space Science*, 10(4). <http://dx.doi.org/10.1029/2022EA002516>.
- Kingma, D. P., Ba, J., 2014. Adam: A method for stochastic optimization.
- Lahssini, K., Baghdadi, N., le Maire, G., Fayad, I., 2022. Influence of GEDI acquisition and processing parameters on canopy height estimates over tropical forests. *Remote Sens. (Basel)*, 14(24), 6264.
- Lanfranchi, C., Kalischek, N., Lang, N., Schindler, K., Wegner, J. D., 2022. Global biomass mapping and uncertainty estimation from gedi lidar data using bayesian deep learning.
- Lang, N., Jetz, W., Schindler, K., Wegner, J. D., 2023. A high-resolution canopy height model of the Earth. *Nature Ecology & Evolution*, 7(11), 1778–1789. <http://dx.doi.org/10.1038/s41559-023-02206-6>.
- Leonhardt, J., Drees, L., Jung, P., Roscher, R., 2022. *Probabilistic Biomass Estimation with Conditional Generative Adversarial Networks*. Springer International Publishing, 479–494.
- Lhoest, Q., Villanova del Moral, A., Jernite, Y., Thakur, A., von Platen, P., Patil, S., Chaumond, J., Drame, M., Plu, J., Tunstall, L., Davison, J., Šaško, M., Chhablani, G., Malik, B., Brandeis, S., Le Scao, T., Sanh, V., Xu, C., Patry, N., McMillan-Major, A., Schmid, P., Gugger, S., Delangue, C., Matussière, T., Debut, L., Bekman, S., Cistac, P., Goehring, T., Mustar, V., Lagunas, F., Rush, A., Wolf, T., 2021. Datasets: A community library for natural language processing. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 175–184.
- Li, S., Su, J., Lang, X., Liu, W., Ou, G., 2018. Positive relationship between species richness and aboveground biomass across forest strata in a primary Pinus Kesiya forest. *Scientific Reports*, 8(1). <http://dx.doi.org/10.1038/s41598-018-20165-y>.
- Li, Y., Li, M., Li, C., Liu, Z., 2020. Forest aboveground biomass estimation using Landsat 8 and Sentinel-1A data with machine learning algorithms. *Scientific Reports*, 10(1). <http://dx.doi.org/10.1038/s41598-020-67024-3>.

- Liu, S., Brandt, M., Nord-Larsen, T., Chave, J., Reiner, F., Lang, N., Tong, X., Ciais, P., Igel, C., Pascual, A., Guerra-Hernandez, J., Li, S., Mugabowindekwe, M., Saatchi, S., Yue, Y., Chen, Z., Fensholt, R., 2023a. The overlooked contribution of trees outside forests to tree cover and woody biomass across Europe. *Science Advances*, 9(37). <http://dx.doi.org/10.1126/sciadv.adh4097>.
- Liu, W., Xu, C., Zhang, Z., De Boeck, H., Wang, Y., Zhang, L., Xu, X., Zhang, C., Chen, G., Xu, C., 2023b. Machine learning-based grassland aboveground biomass estimation and its response to climate variation in Southwest China. *Frontiers in Ecology and Evolution*, 11. <http://dx.doi.org/10.3389/fevo.2023.1146850>.
- Main-Knorn, M., Pflug, B., Louis, J., Debaecker, V., Müller-Wilm, U., Gascon, F., 2017. Sen2Cor for sentinel-2. L. Bruzzone, F. Bovolo, J. A. Benediktsson (eds), *Image and Signal Processing for Remote Sensing XXIII*, SPIE.
- Metzger, N., Daudt, R. C., Tuia, D., Schindler, K., 2023. High-resolution population maps derived from sentinel-1 and sentinel-2.
- Nascetti, A., Yadav, R., Brodt, K., Qu, Q., Fan, H., Shendryk, Y., Shah, I., Chung, C., 2024. BioMassters: A Benchmark Dataset for Forest Biomass Estimation using Multi-modal Satellite Time-series. *Advances in Neural Information Processing Systems*, 36. https://papers.nips.cc/paper_files/paper/2023/file/40daf2a00278c4bea1b26cd4c8a654f8-Paper-Datasets_and_Benchmarks.pdf.
- Pascarella, A. E., Giacco, G., Rigioli, M., Marrone, S., Sansone, C., 2023. ReUse: REgressive Unet for Carbon Storage and Above-Ground Biomass Estimation. *Journal of Imaging*, 9(3), 61. <http://dx.doi.org/10.3390/jimaging9030061>.
- Pauls, J., Zimmer, M., Kelly, U. M., Schwartz, M., Saatchi, S., Ciais, P., Pokutta, S., Brandt, M., Gieseke, F., 2024. Estimating canopy height at scale.
- Puliti, S., Hauglin, M., Breidenbach, J., Montesano, P., Neigh, C., Rahlf, J., Solberg, S., Klingenberg, T., Astrup, R., 2020. Modelling above-ground biomass stock over Norway using national forest inventory data with ArcticDEM and Sentinel-2 data. *Remote Sensing of Environment*, 236, 111501. <http://dx.doi.org/10.1016/j.rse.2019.111501>.
- Santoro, M., Cartus, O., 2023. Esa biomass climate change initiative (biomass_cci): Global datasets of forest above-ground biomass for the years 2010, 2017, 2018, 2019 and 2020, v4.
- Shimada, M., Itoh, T., Motooka, T., Watanabe, M., Shiraishi, T., Thapa, R., Lucas, R., 2014. New global forest/non-forest maps from ALOS PALSAR data (2007–2010). *Remote Sensing of Environment*, 155, 13–31. <http://dx.doi.org/10.1016/j.rse.2014.04.014>.
- Sonkoly, J., Kelemen, A., Valkó, O., Deák, B., Kiss, R., Tóth, K., Miglécz, T., Tóthmérész, B., Török, P., 2019. Both mass ratio effects and community diversity drive biomass production in a grassland experiment. *Scientific Reports*, 9(1). <http://dx.doi.org/10.1038/s41598-018-37190-6>.
- Tadono, T., Ishida, H., Oda, F., Naito, S., Minakawa, K., Iwamoto, H., 2014. Precise Global DEM Generation by ALOS PRISM. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, II–4, 71–76. <http://dx.doi.org/10.5194/isprsannals-II-4-71-2014>.
- Tang, H., Stoker, J., Luthcke, S., Armston, J., Lee, K., Blair, B., Hofton, M., 2023. Evaluating and mitigating the impact of systematic geolocation error on canopy height measurement performance of GEDI. *Remote Sensing of Environment*, 291, 113571. <http://dx.doi.org/10.1016/j.rse.2023.113571>.
- Tolan, J., Yang, H.-I., Nosarzewski, B., Couairon, G., Vo, H. V., Brandt, J., Spore, J., Majumdar, S., Haziza, D., Vamaraju, J., Moutakanni, T., Bojanowski, P., Johns, T., White, B., Tiecke, T., Couprie, C., 2024. Very high resolution canopy height maps from RGB imagery using self-supervised vision transformer and convolutional decoder trained on aerial lidar. *Remote Sensing of Environment*, 300, 113888. <http://dx.doi.org/10.1016/j.rse.2023.113888>.
- Vance-Chalcraft, H. D., Willig, M. R., Cox, S. B., Lugo, A. E., Scatena, F. N., 2010. Relationship Between Above-ground Biomass and Multiple Measures of Biodiversity in Sub-tropical Forest of Puerto Rico. *Biotropica*, 42(3), 290–299. <http://dx.doi.org/10.1111/j.1744-7429.2009.00600.x>.
- Xu, L., Saatchi, S. S., Shapiro, A., Meyer, V., Ferraz, A., Yang, Y., Bastin, J.-F., Banks, N., Boeckx, P., Verbeeck, H., Lewis, S. L., Muanza, E. T., Bongwele, E., Kayembe, F., Mbenza, D., Kalau, L., Mukendi, F., Ilunga, F., Ebuta, D., 2017. Spatial Distribution of Carbon Stored in Forests of the Democratic Republic of Congo. *Scientific Reports*, 7(1). <http://dx.doi.org/10.1038/s41598-017-15050-z>.
- Zhang, Y., Liang, S., Yang, L., 2019. A Review of Regional and Global Gridded Forest Biomass Datasets. *Remote Sensing*, 11(23), 2744. <http://dx.doi.org/10.3390/rs11232744>.

8. Appendix

8.1 Dataset description

Each sample in the dataset contains a pair of pre-cropped images (*patches*) and their corresponding AGB labels. You can find an up-to-date description of the dataset on HuggingFace.

We hereby provide more information on the encoding of certain features:

- Geographical Coordinates — we perform sine/cosine encoding, as follows:
 $\text{lat_cos} = (\cos(\pi * \text{lat}/90) + 1)/2$
 $\text{lat_sin} = (\sin(\pi * \text{lat}/90) + 1)/2$
 $\text{lon_cos} = (\cos(\pi * \text{lon}/180) + 1)/2$
 $\text{lon_sin} = (\sin(\pi * \text{lon}/180) + 1)/2$
- Land Cover Information — we perform sine/cosine encoding, as follows:
 $\text{lc_cos} = (\cos(2\pi * \text{lc_class}/100) + 1)/2$
 $\text{lc_sin} = (\sin(2\pi * \text{lc_class}/100) + 1)/2$
 where lc_class is the discrete classification of the land cover.

Additionally from the features described in this manuscript, users of the dataset have access to the following features (mainly GEDI L4A metadata):

- AGBD Standard Error: The uncertainty estimate associated with the aboveground biomass density prediction for each GEDI footprint.
- Elevation: The height above sea level at the location of the GEDI footprint.
- Leaf-Off Flag: Indicates whether the measurement was taken during the leaf-off season, which can impact canopy structure data.
- Plant Functional Type (PFT) Class: Categorization of the vegetation type by GEDI (e.g., deciduous broadleaf, evergreen needleleaf).
- Region Class: The geographical area where the footprint is located (e.g., North America, South Asia).
- RH98 (Relative Height at 98%): The height at which 98% of the returned laser energy is reflected, a key measure of canopy height.
- Sensitivity: The proportion of laser pulse energy reflected back to the sensor, providing insight into vegetation density and structure.
- Solar Elevation: The angle of the sun above the horizon at the time of measurement, which can affect data quality.
- Urban Proportion: The percentage of the footprint area that is urbanized, helping to filter or adjust biomass estimates in mixed landscapes.
- Date of GEDI Footprints: The specific date on which each GEDI footprint was captured, adding temporal context to the measurements.
- Date of Sentinel-2 Image: The specific date on which each Sentinel-2 image was captured, ensuring temporal alignment with GEDI data.

8.2 Hosting, licensing, and maintenance plan

The processed, ready-for-use data is hosted on HuggingFace, where it can be easily accessed:

```
#!pip install datasets
from datasets import load_dataset
dataset = load_dataset("prs-eth/AGBD", streaming=True)["train"] # or test, val
```

The raw data, dense predictions, and model weights are hosted on the ETH Research Collection, guaranteeing long-term preservation.

The data is available under a CC BY-NC 4.0 DEED license. You are free to copy and redistribute the material in any medium or format; remix, transform, and build upon the material. You must give appropriate credit, provide a link to the license, and indicate if changes were made. Importantly, you may not use the material for commercial purposes.

The dataset will be maintained by the authors to the best of their abilities and availability.

8.3 Confusion density plots

We provide the confusion density plots for the best performing model of each architecture, in Figure 9. It shows good agreement between the predictions and GEDI reference, and provides insights on the distribution of the biomass values.




Figure 9. Density performance plots for the best model of each architecture.

8.4 GBDT model initialization parameters

We hereby specify the initialization parameters for the GBDT models:

```
num_leaves=165, max_depth=-1, learning_rate=0.1,
num_iterations = 1000, min_split_gain = 0.,
min_child_weight = 1e-3, min_child_samples=20,
reg_alpha= 0, reg_lambda = 5,
bagging_fraction = 0.5, bagging_freq = 1.
```

	GEDI	Sentinel-2 (S2)	ALOS-2 PALSAR-2 (P2)	ALOS Digital Elevation Model (DEM)	Land Classification (LC)	CH
Source	NASA	Copernicus	JAXA	JAXA	Copernicus	Lang et al.
Resolution (m)	25	[10,20,60]	25	30	100	10
# of channels	-	12	2	1	2	2
Temporality	-	≈ every 5 days	yearly	once	yearly	yearly

Table 3. Data sources

8.5 Dense predictions

We generate the dense 10m resolution AGB estimates for all Sentinel-2 tiles in the dataset coverage. For each tile, we pick the Sentinel-2 product that has the least cloud coverage over all products available, and perform inference on that product. As such, our dense predictions span multiple years. On the ETH Research Collection, the file `Inference > mapping.pkl` indicates which specific product was used for inference for each tile. Note that this was a design choice, and one can generate dense predictions for specific years.

8.6 GEDI AGB Estimates

The GEDI instrument uses a LiDAR laser to repeatedly scan a 25m surface footprint, where each measurement consists of a waveform from the laser's returned signal. From the cumulative return energy, the GEDI team identifies the ground return, which is used as the reference for the relative height (RH) metrics. The RH[X] metrics correspond to the relative height at which [X] percent of the total accumulated energy is returned. From the RH metrics, the GEDI team estimates biomass. To do so, they developed their own allometric equations, based on the Plant Functional Type (PFT), the region of the world, and the RH metrics. In order to calibrate them, they leveraged world-wide crowd-sourced data. Coincident airborne LiDAR and ground plot field inventory were put together. The airborne LiDAR data was used to simulate GEDI-like waveforms and derived RH metrics. These simulated RH metrics were then used to calibrate the biomass equations with the field biomass values as reference data. Further information can be found in the Algorithm Theoretical Basis Document (Kellner et al., 2023).

8.7 GEDI Limitations

Resolution. The GEDI mission operates with a 25-meter diameter footprint on the ground. The laser pulse from GEDI's LiDAR system follows a Gaussian intensity profile, with the highest precision at the center of the footprint, gradually decreasing toward the edges (Hancock et al., 2019). Within this 25-meter footprint, GEDI measures the highest reflecting surface, which is then used to derive Relative Height metrics, from which they estimate the AGB. We then chose to rasterize these footprints at a 10-meter resolution, which is a well-established method as supported by the literature (Pauls et al., 2024, Lang et al., 2023), and aligns with the highest resolution available among the sensors we considered, such as the 10-meter bands from Sentinel-2.

Uncertainty. The uncertainty in the footprint-level GEDI AGBD product used as reference data influences the accuracy and uncertainties of the model predictions. Future work could incorporate field measurements or airborne LiDAR data. However, given the sparsity of GEDI measurements, co-registering

these with in-situ sparse measurements is challenging. Notably, in order to assess the quality of GEDI AGB estimates, the authors of (Hunka et al., 2023) had to rely on an aggregated 1km x 1km grid GEDI product (where the samples present within the borders of each 1 km cell are used to statistically infer mean AGBD). Our denser and higher resolution estimates provide a crucial step forward, as well-calibrated estimates can be used as a proxy to assess the uncertainties in GEDI data by comparing them in-situ measurements. We believe that this will ultimately need a long-term community initiative across ecology, forestry, remote sensing and ML.

Geolocation error. While the geolocation error was an issue with the first version of the data released by GEDI, they released an updated version of their dataset (that we are using) with corrected geolocation errors. Notably, the authors of (Tang et al., 2023) investigated the geolocation error of the updated version, and "did not find any significant relationship between geolocation error and performance". They find that 80.8% of the footprints have a geolocation error strictly smaller than 10m.

Sampling pattern Another point to keep in mind is that the spatial distribution of GEDI footprints is not uniform: the sampling gets progressively denser as one approaches towards the borders of the observed region at $\pm 51.6^\circ$.