

Article

Not peer-reviewed version

Uncertainty Quantification of GEDI Clear-Sky Terrain Height Retrievals Using a Mixture Density Network

[Jonathan Sipps](#) and [Lori A. Magruder](#) *

Posted Date: 17 October 2023

doi: 10.20944/preprints202310.1028.v1

Keywords: GEDI; laser altimetry, lidar, uncertainty quantification; mixture density network; terrain elevation



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Article

Uncertainty Quantification of GEDI Clear-Sky Terrain Height Retrievals Using a Mixture Density Network

Jonathan Sippes ^{1,2,*} and Lori A. Magruder ^{1,2}

¹ Department of Aerospace Engineering and Engineering Mechanics, Cockrell School of Engineering, University of Texas at Austin, Austin, TX 78705, USA

² Center for Space Research, University of Texas at Austin, Austin, TX 78759, USA

* Correspondence: john.sippes@utexas.edu

Abstract: Early spaceborne laser altimetry mission development starts in pre-phase A design, where diverse ideas are evaluated against mission science requirements. A key challenge is predicting realistic instrument performance through forward modeling at arbitrary spatial scale. Analytical evaluations compromise accuracy for speed, while radiative transfer modeling is not applicable at global scale due to computational expense. Instead of predicting arbitrary properties of a lidar measurement, we develop a baseline theory to predict only the distribution of uncertainty specifically for the terrain elevation retrieval based on terrain slope and fractional canopy cover features through a deep neural network gaussian mixture model, also known as a mixture density network (MDN). Training data was created from differencing geocorrected GEDI L2B elevation measurements with 32 independent reference lidar datasets in the contiguous U.S. from the National Ecological Observatory Network. We trained the MDN and selected hyperparameters based on regional distribution predictive capability. On average, the relative error of equivalent standard deviation of predicted regional distributions was 15.9%, with some anomalies in accuracy due to generalization and insufficient feature diversity and correlation. As an application, we predict the percent of elevation residuals of a GEDI-like lidar within a given mission threshold from 60°S to 78.25°N, which correlate to qualitative understanding of prediction accuracy and instrument performance.

Keywords: GEDI; laser altimetry; lidar; uncertainty quantification; mixture density network; terrain elevation

1. Introduction

In the lifecycle of space-based mission design, the starting point is pre-phase A, where a broad spectrum of ideas can be explored and evaluated. Level 1 mission requirements drive a myriad of pre-phase A mission development decisions related to orbit design and instrument selection for single or multiple satellites. All three criteria are intertwined, with competing cost and “value” with respect to mission science requirements. Predicting instrument performance with respect to realistic environmental conditions can be difficult, and typically realism is sacrificed for computational feasibility, leading to sub-optimal design. As an example, missions looking to measure global elevation need to explore all scenarios of the surface structure, composition, reflectivity, and attenuation as components that impact the ability of the sensor to adequately capture the measurement. Often, laser altimetry (lidar), an active sensing technology, is selected to observe the Earth's surface in 3-dimensions and requires comprehensive analysis of the energy link budget during pre-phase A studies. The utility of measuring height at the global scale supports the Decadal Survey (DS) observational needs across a diverse number of scientific communities [1]. Although space-based lidar is a relatively new remote sensing technology for Earth observation compared to other sensing types, it has been proven to support measured needs associated with surface topography, vegetation structure, cryospheric processes, and shallow water bathymetry, with high vertical accuracy [2]. Unlike imaging systems or synthetic aperture radar (SAR), lidar can penetrate

dense canopies, revealing details in the underlying topography and vertical structure of the vegetation.

Having knowledge of laser altimetry's capabilities from previous planetary missions using the technology, NASA launched the first Earth observing satellite laser altimeter in 2003. This mission, ICESat (Ice, Cloud, Land Elevation Satellite), was primarily focused on quantifying changes in our polar ice sheets and sea ice, as hot spots of climate impact on sea level rise and the radiative balance between ocean/land and atmosphere [3]. However, the instrument operated globally so it was able to capture measurements of vegetation, oceans, atmosphere, and mid-latitude topography. ICESat was operational until decommissioning in 2009 after serving a wide range of scientists using altimetry to study Earth systems. Following the success of ICESat, ICESat-2 – operational in 2018 – continued the focus on the cryosphere, including measuring changes in sea ice freeboard and ice sheet elevation in Greenland and Antarctica [4]. The combined data products from ICESat and ICESat-2 missions have made significant contributions to a ~20-year timeseries of cryospheric response to changes in the atmosphere and ocean conditions [4].

Using the success of ICESat lidar technology but designed for mid-latitude vegetation, NASA developed the Global Ecosystem Dynamics Investigation (GEDI). GEDI science goals are focused on measuring ecosystem structure, namely canopy heights, to derive aboveground biomass and inform global carbon stocks [5]. Through full-waveform lidar technology, GEDI provides vertical profiles of forest structure ideal for tropical and temperate regions. GEDI height distributions initialize and constrain [5,6] the predictive outputs of the Ecosystem Demography model [7], directly supporting carbon flux estimation and its role in grander climate change understanding. All three space-based lidar mission products have been fused with Landsat [8], TanDEM-X [6], and other mission products [5,9] to enhance and/or calibrate optical and radar-based measurements.

Clearly, there is great utility in maximizing the output of a novel lidar-based mission concept. Striking a balance between realism and analytic estimates of lidar performance, we introduce a statistical approach to assess whether lidar instrument performance meets mission requirements. We narrowly focus our study on quantifying lidar-measured elevation uncertainty based on terrain slope and canopy fractional cover for surface topography and vegetation applications, to provide a baseline work for future studies to adapt.

Lidar modeling techniques range from macroscopic approximations of performance as pertains to measurement objectives [10–13] to highly detailed simulations of the absorption of light through the atmosphere and its reflectance off of vegetation, terrain, or urban landscapes [14–17]. Analytic estimates of architecture performance can be derived from simplifications of the lidar equation, assuming mean surface reflectance and atmospheric transmittance [10,13]. For simple reflectance geometries, explicit formulas on the change in waveform are tractable [18]. When applied to many microfacets, aggregation of differential changes in the transmitted waveform results in realistic return waveforms [18,19]. At the other end of the spectrum, 3D radiative transfer (RT) models such as the Discrete Anisotropic Radiative Transfer (DART) model [17] account for absorption, emission, and scattering of radiation, using external or RT-derived tabular data on physical characteristics of a scene, through a combination of Monte-Carlo and ray-tracing. While comprehensive, RT models are computationally demanding and may not be applicable in preliminary mission design at the global scale.

Between these two extremes, the GEDI Simulator is a data-driven approach to simulate return waveforms, purpose-built for calibration and validation of GEDI data products [20,21]. Instead of simulating radiance interactions, the GEDI Simulator aggregates high-resolution airborne lidar survey (ALS) point clouds to create waveforms, a technique originally conceptualized by Blair and Hofton [22]. A stark advantage of the GEDI Simulator is that it does not require optical properties of terrain and vegetation structure, as is necessary with RT models. However, its simulation domain is limited to the extent of available reference lidar.

Aside from classification, machine-learning (ML) methods are commonly applied in an inverse problem setting to estimate a physical property about measurements [23], such as above ground biomass, canopy height, or canopy cover. For example, the GEDI science team developed a

convolutional neural network to predict canopy heights more accurately from GEDI waveforms [24]. Contrary to traditional lidar applications of ML, we seek to model the elevation return quality based off of physical properties of Earth's surface; therefore, our technique loosely falls under forward modeling, for uncertainty quantification of GEDI terrain elevation measurements.

We model the footprint-level elevation residual distribution (ERD) as a Gaussian mixture model (GMM), predicted through a Mixture Density Network (MDN) [25,26], which is a deep neural network (DNN) with special activation functions at the output layer that produces mixing coefficients, means, and standard deviations for each component of the GMM. With a GMM, we can fit the near Cauchy-like ERD while maintaining generality for more mountainous or vegetated regions that could have less linear elevation residual variance than flatter, barren regions. Our key innovation is to fit at footprint scale, but attempt to predict elevation residuals over entire areas, estimating a so-called "regional elevation residual distribution" or RERD. We therefore train the ERD on a traditional 80/20 split but choose hyperparameters based on RERD predictions to better isolate generality from local to global scales. In Section 2, we discuss reference lidar and GEDI measurement data used to train the MDN, thoroughly discuss the development of RERD prediction, and end with a summary over the methodology. Then in Section 3, we show model validation and accuracy from a 5-fold cross-validation, and finally conclude through a discussion in Section 4.

2. Materials and Methods

2.1. Overview

We seek to train a model based on terrain slope and canopy cover input features to predict the geocorrected elevation residual in the version 2 L2B GEDI elevation measurement at non-polar latitudes as a means to explore the performance of the GEDI instrument. Our baseline model choice is a neural network based Gaussian mixture model (GMM), known as a Mixture Density Network (MDN) [25,26]. With a properly trained model, we can predict what the elevation measurement quality of a spaceborne, GEDI-like lidar will be, given any location on land, where trained features are applicable. The predicted ERDs are fundamentally meter-scale, but such resolution is computationally prohibitive and overly precise for global, conceptual mission design. We therefore downsample predictions to km-scale, outputting the distribution of elevation residuals over large spatial regions (RERDs), given the terrain and vegetation structure of those regions. We develop training data by differencing the measured GEDI elevation against airborne lidar surveys (ALS), then correlate slope and cover features with elevation residuals to train the model. Through a statistical aggregation, we predict how GEDI elevation residuals are distributed throughout entire spatial regions of Earth's land surface. In this section, we describe the reference lidar, GEDI L2B elevations, modeling theory, and implementation details.

2.2. Reference Lidar

Airborne lidar surveys (ALS) typically produce centimeter level accurate point-clouds [27], containing high-resolution 3D structure of terrain and vegetation in a km-scale region. If temporally consistent, GEDI waveforms can be directly compared to reference lidar to validate derived datasets such as canopy height, terrain elevation, or the entire waveform [20,28]. For terrain-only applications, the GEDI waveform lowest mode elevation retrieval can be compared to high-resolution raster digital terrain models (DTMs) of ground-classified reference lidar [29] to determine the fidelity of spaceborne measurements.

In recent years, Refs. [28,30,31] have validated GEDI elevation and canopy height accuracy against the National Ecological Observatory Network (NEON) discrete return point clouds. NEON has an open-source easy-to-use interface to query reference lidar, with collection dates from 2013 to 2023, with 59 collection sites in the United States, including regions in Hawaii, Alaska, and Puerto Rico [32]. The NEON airborne observation platform deploys an Optech Gemini lidar for the discrete lidar product, maintaining measurement quality for each region of interest (ROI) to <5-15 cm horizontal accuracy (1σ) and <5-35 cm vertical accuracy (1σ) [28]. All NEON reference lidar tiles are

referenced to the WGS84 ITRF2000 ellipsoid horizontal datum, with the NAVD88 vertical datum using GEOID12A [33].

The temporal differences between measured and reference lidar can lead to additional uncertainty in comparisons [28]. GEDI launched in late 2018 on a 2-year mission, starting in April 2019 after on-orbit checkout [5]. Because a third of 2019 does not have GEDI data, and because 2020 does not have a sufficient number of NEON reference tiles, we chose the year 2021 to optimize both the amount of temporally consistent GEDI data and abundance of NEON reference data, while simplifying the comparisons to a single year of reference lidar. To compare interannual trends in ecology, NEON ALS measurements are taken during the spring and summer when vegetation is at maximum greenness [33]. In this regard, the reference lidar is limited to March - September 2021.

Out of 59 discrete return lidar sites, only 32 were selected as 3 sites (BONA, DEJU, and HEAL) were outside latitudinal bounds of GEDI data and the rest did not have available data in the year 2021. Despite those sites out of range, 32 sites offer significant diversity in biomes across the United States [28,30,31]. Figure 1 shows how NEON sites are distributed in space over terrain slope and canopy cover features.

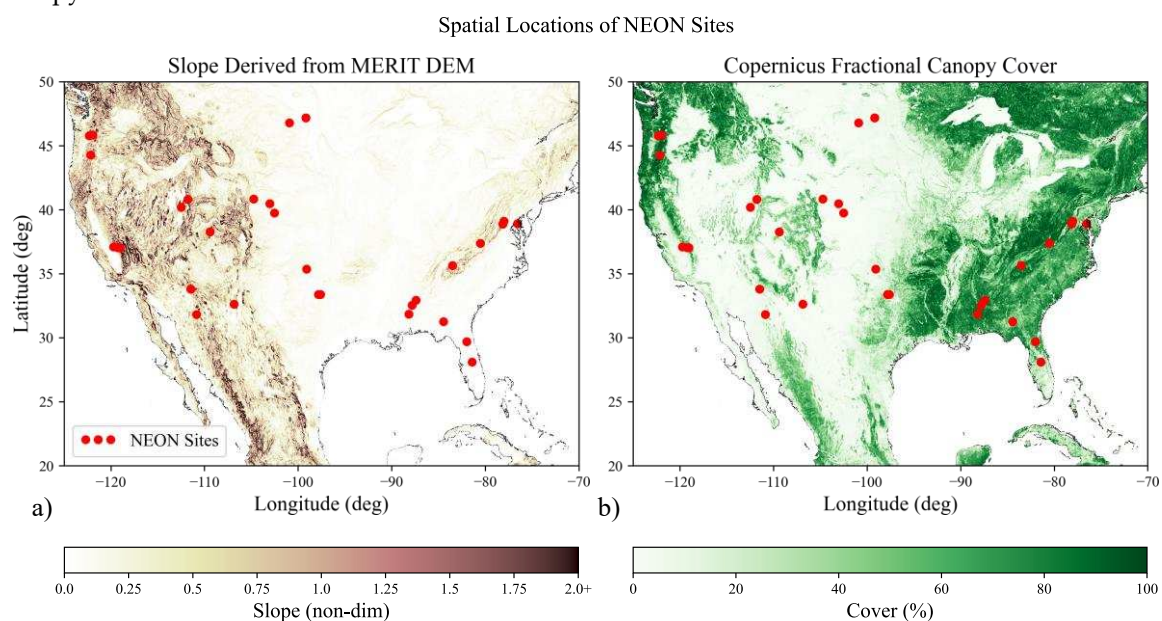


Figure 1. NEON sites with a) slope and b) cover feature diversity.

In Figure 1a, we downsampled the Multi-Error-Removed Improved-Terrain Digital Elevation Model (MERIT DEM) [34] to $1^\circ/24$, then derived demonstrative slope via a gradient, and scaled the slope to high-resolution slope obtained within each ROI. For Figure 1b, we performed a similar operation, downsampling Copernicus Global Land Cover product [35] to $1^\circ/25.2$ then averaged the value of cover for each pixel. The result is a demonstrative figure that shows where NEON ROIs are with respect to geographic features. Statistically, there are 10 ROIs with 90th percentile slope (S90) less than 0.25 and 90th percentile cover (C90) less than 50%; another 10 ROIs with S90 less than 0.25 but with C90 greater than 50%; 9 ROIs with S90 greater than 0.25 and C90 greater than 50%; and only 3 ROIs with S90 greater than 0.25 and C90 less than 50%.

2.3. GEDI L2B Data

GEDI launched in December 2018 and was positioned on the International Space Station (ISS) Japanese Experiment Module-Exposed Facility, which covers between $\pm 51.6^\circ$ latitude based on the ISS orbital inclination. The GEDI instrument provides multi-beam elevation measurements using three lasers. One “coverage” laser is split into two beams at 4.5 mJ each, which are dithered to produce four incident transects. The other two full “power” lasers are only dithered to produce an additional two transects, each at 15 mJ [28,36]. In total, there are 4 coverage and 4 power profiles (ground-tracks)

on the surface of the Earth. Full power laser energy penetrates canopy cover more easily than the lower energy density of the coverage lasers, producing more ground detections, but not necessarily better terrain elevation accuracy than coverage beams [28]. The across-track spatial coverage of the resulting 8 beam profiles is 4.2 km, with ~600 m separation between beams [5]. Each of the GEDI lasers are produced at 1064 nm wavelength and provide a 25 m diameter footprint with 60 m along-track spatial resolution [5]. The full-waveform technology of the GEDI lidar allows the system to capture a complete temporal profile of the returned laser energy along the laser line-of-sight. This capability allows for retrieval of both vegetation structure, canopy metrics, and terrain elevation in most cases.

The GEDI science team derives many datasets and four level 2 data products from the measured full-waveform signal of each footprint. Datasets range from the waveform time-series (per footprint) on L1B to more refined products such as leaf area index (LAI) or canopy coverage on L2B. This study utilizes relevant L2B GEDI granules based on coincidence to NEON ROIs through the GEDI Finder tool for version 2 data. We used the VDatum software (<https://vdatum.noaa.gov/welcome.html>) to convert longitude, latitude, and elevation to respective reference lidar projected coordinate systems.

Waveforms were filtered with the `l2b_quality_flag`, which removes returns obscured by clouds and background noise [37]. Variations in day- or nighttime collection has not been shown to have a meaningful impact on terrain elevation accuracy [28]. Likewise, both power and coverage beams provide similar terrain elevation accuracy, where the ground signal is observable through canopy cover [28].

The GEDI horizontal geolocation error at the footprint level has a standard deviation of 10.2 m [38], and must be geocorrected before comparing against reference lidar to separate instrument performance from operationally induced errors. Waveform geolocation can be corrected through two primary techniques, waveform matching [30] and terrain matching [29,30]. In waveform matching, the full waveform profile from GEDI's L1B product is compared against synthetic waveforms created from reference lidar, typically through the GEDI Simulator. The waveform geolocation is shifted horizontally until it agrees with the reference waveforms under a relevant matching criteria. We applied terrain matching instead, where a short segment (km-scale) of each beam track is shifted iteratively over a 1 m resolution reference lidar-based DTM until the mean absolute error (MAE) of elevations between the track and the DTM are minimized. We apply the same algorithm as used to geocorrect ICESat-2 tracks in Refs. [39,40] through research correspondence. In our particular terrain matching technique, tracks are shifted by up to +/-48 m in along- and cross-track directions, then compared to iteratively higher resolution DTMs made directly from the reference lidar. Resolutions iterate over 8m, 4m, 2m, then 1m, and for each resolution, we derive a correction estimate, then increase the resolution, and provide the prior correction as an initial condition for the next estimate. In this way, we robustly and efficiently geocorrect to 1 m resolution.

Applying the same terrain matching technique as ICESat-2 to GEDI could affect geocorrection quality, as ICESat-2 is a standalone satellite with much higher platform stability than the ISS, and the ICESat-2 footprint spacing is only 0.7 m as opposed to GEDI's 60 m [4]. Schleich et al. [29] determined that, for GEDI, the along-track segment length should be lower than the structural vibration period of the ISS (1-10 sec), choosing a segment length no lower than 13 footprints and no larger than 50 footprints (3 km, or 0.43 sec of along-track motion). We initially processed 24 ROIs with no fewer than 10 footprints and <25 km segment lengths, and processed 8 ROIs with <5 km lengths, in light of the preprint by Schleich et al. [29]. In post-processing, we found a systematic bias in 5 ROIs in particularly mountainous and forested terrain; upon re-processing at <5 km, much of the systematic errors were reduced.

After estimating a horizontal geocorrection for every segment of each GEDI beam, we subtract out the vertical difference in elevation between the GEDI transect and the reference lidar, fully georectifying the GEDI beam transect with reference lidar in three dimensions. Then we can finally determine the elevation residual, which is the remaining difference between geocorrected, range-corrected, measured GEDI elevations and the true elevation, approximated by reference lidar. That

is, we separate the errors in GEDI elevation measurement from geolocation and ranging accuracy, and isolate errors only due to atmospheric effects, terrain characteristics, and vegetation structure.

Processed data is filtered by quality footprints, whether there are more than 10-13 footprints in a “track”, and where the geocorrection algorithm diverged for a given track. Through all filters, we obtained 624,795 footprints to train the model. The smallest ROI has 2,660 footprints while the largest has 40,924. Out of processed footprints, 48.2% are in day and 51.8% are at night. In addition, 56.9% are from the power beam, while 43.1% are from the coverage beam.

2.4. Model Input Features

Numerous factors affect GEDI terrain elevation accuracy. Wang et al. [31] compiled a review of 9 studies since 2020, each focusing on factors influencing both terrain retrieval and relative canopy height accuracy. According to the extensive performance survey conducted by Liu et al. [28], in order of importance, the most impactful factors for terrain elevation retrieval are slope, canopy cover, canopy height, land cover type, and beam sensitivity. This study prioritizes the development of a general methodology that can scale to as many features (or factors influencing elevation residuals) as is necessary. As such, we have chosen the two most dominant features – terrain slope and canopy cover – to investigate elevation residual prediction.

Terrain slope – hereafter “slope” – is the tangent of the angle that terrain makes as it changes in elevation over a small area, with respect to both an x- and y-direction, such as along- and cross-track, or Easting and Northing in a Universal Transverse Mercator (UTM) projection. In a mathematical sense, slope is merely the magnitude of the gradient of the terrain surface. But in application to a real DTM, the calculation requires a user-specified “small area”. In the context of GEDI, the appropriate area size is the footprint, which is modeled as a circular disk with a diameter of 25 m projected vertically onto the DTM, centered on geocorrected footprint geolocation.

Canopy fractional cover – or just “cover” – is equivalent to the percent of ground covered by a vertical projection of canopy material. The GEDI data product “cover” on L2B estimates cover, given leaves, branches, and stems [41]. As a feature, dense cover (e.g., in jungle/forest) correlates with lower terrain elevation accuracy because photons cannot pierce the canopy, resulting in a lower signal-to-noise ratio in ground detection.

While we can derive slope and cover from a reference DTM, the features are only applicable where reference data exists, preventing prediction elsewhere. Instead, we train the model on ancillary global feature datasets that are available at prediction. We derive slope through terrain elevation estimated by MERIT DEM [34]. The MERIT DEM is at 3 arcsecond resolution (~90 m at the equator), and covers latitudes from 60°S to 90°N, referenced to WGS84. The MERIT DEM is subset to the given ROI latitude and longitude bounds, then upsampled to 25 m (by interpolation) to reflect GEDI footprint size, then slope is derived by the gradient of the DEM, averaged over the GEDI footprint area. In this way, we extract a slope derived from MERIT per GEDI footprint.

The process to extract canopy cover is similar, except we need not apply a post-processed gradient operation. Cover is represented by the Copernicus Global Land Cover product [35] “Tree_CoverFraction_layer” dataset for the year 2019, at ~3.6 arcsecond resolution (1°/1008), or ~100 m at the equator. The dataset extent is from 60°S to 78.25°N, and represents forest cover annually in 2019, the latest year available. The cover dataset is subset by the given ROI, upsampled to 25 m to coincide with GEDI footprint size, then cover is averaged over GEDI footprint area. In this way, we extract Copernicus forest canopy cover per GEDI footprint.

Slope is mostly temporally invariant, but canopy cover varies significantly by season. Unfortunately, there are not many high spatial and temporal resolution, global canopy cover datasets available. “Cover” is a dataset on the GEDI L2B product that maps the exact forest cover to the GEDI footprint, which is spatially and temporally coincident. However, this is only applicable in training, and we cannot interpolate outside of GEDI footprint measurements. Were a higher resolution, more temporally consistent canopy cover dataset to become available, we could replace Copernicus forest cover in the year 2019 with a better feature dataset. The result is that the “cover” feature from

Copernicus is less indicative of cover's effect on elevation residual, as it would be with the GEDI L2B "cover" dataset.

Because the feature datasets are limited in latitudinal extent, our "global" predictive capability is limited to the overlap of MERIT and Copernicus datasets, or 60°S to 78.25°N. While GEDI footprints only spatially cover 51.6° north and south, the feature space may cover further. That is, the model trained on slope and cover within $\pm 51.6^\circ$ may be indicative of slope and cover at higher/lower latitudes, assuming that slope and cover in other areas of the world are indicative of slope and cover within the U.S.-based NEON sites considered.

2.5. Theory

For a single point on Earth, there are only two scalar features (slope and cover), and the prediction, or target variable, is the entire ERD. The ERD quantifies the error of terrain elevation measurements – or the last peak of the waveform – not any other elements in the waveform. To be clear, the ERD is the culmination of the error in many thousands (or more) of ground returns and should not be confused with the distribution-like character of the return lidar waveform. If the ERD were modeled as Gaussian, we could parameterize the predictions as two scalars: mean and standard deviation. However, we found the Gaussian distribution to be insufficient for the accuracy required. In fact, the measured error (approximately true) distributions very closely resemble a Cauchy distribution. However, assuming Cauchy also introduced errors into prediction, and we settled on a Gaussian mixture model (GMM).

Mathematically, the elevation residual is

$$\Delta z_i = z_i - g(\vec{p}_i) \quad (1)$$

where z_i is the i th measured ground elevation, g is a function representing the reference DTM surface, \vec{p}_i is the geocorrected position in reference coordinates, and Δz_i is the elevation residual. The elevation residual Δz (scalar) is modeled as a random process f , dependent on feature vector x :

$$\Delta z \sim f(x) \quad (2)$$

For our study, $x = (m, c)$ is a feature pair of slope m and canopy cover c . That is, for every slope and cover, we have a *different* distribution. If f were Gaussian, the mean and standard deviation would change with respect to each feature pair. Note that f does not depend on spatial location, only slope and cover (or features in general).

Ultimately, the use case of our model is to predict how well GEDI retrieves elevation over a *region*, not just at a given point. Let X be the set of all features in a given ROI. Then through Monte-Carlo integration, the RERD is

$$f_{ROI}(\Delta z; X) \approx \frac{1}{M} \sum_{i=1}^M f(\Delta z; x_i) \quad (3)$$

where $M = |X|$ is the number of feature points (or pairs) in the ROI (at training or prediction).

2.6. Model

To clarify model discussion, Figure 2 shows an example ROI (DCFS) including spatial location of the ROI, feature distributions for slope and canopy cover, and the corresponding measured RERD directly from comparisons against the reference lidar. In addition, Figure 2 shows Cauchy and Gaussian fits of the RERD.

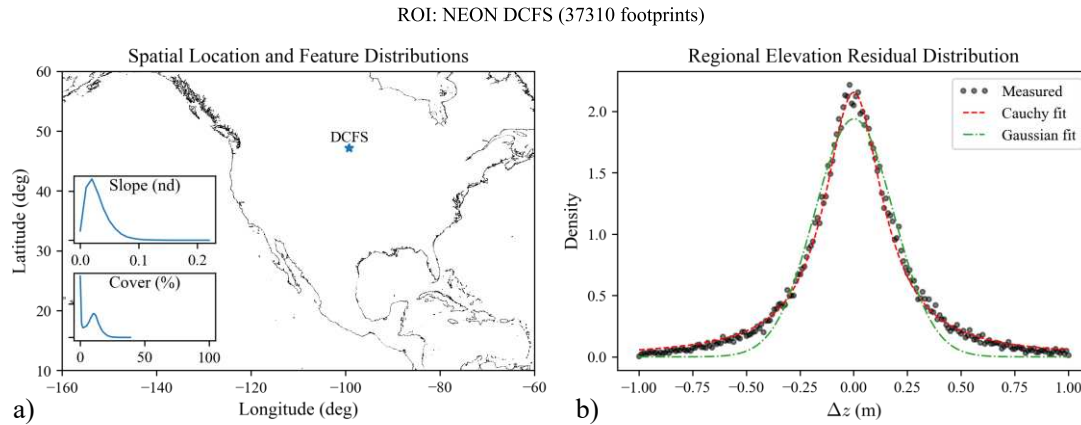


Figure 2. a) Spatial location of DCFS ROI and slope/cover distributions, and b) the resulting RERD for DCFS ROI.

The RERD, or the ERD aggregated over the ROI, is built by a normalized histogram of all elevation residuals in the ROI, without respect to features. In Figure 2a, one can see “Slope” and “Cover” feature distributions. The key task is to correlate the distribution of slope and cover with the RERD, shown in Figure 2b, and ultimately predict a new RERD given features in a new ROI.

A GMM can theoretically fit any distribution (given enough components), and it served as our baseline model, the MDN [25]. Expressing f in terms of a GMM,

$$f(\Delta z; x) \approx \sum_{j=1}^n \alpha_j N(\Delta z; \mu_j, \sigma_j^2) \quad (4)$$

or Δz is distributed over f , dependent on a feature pair x , which is approximately equal to the weighted sum of n Gaussian components; α_j , μ_j , and σ_j are respectively the coefficient, the mean, and the standard deviation for the j th component. Explicitly, α_j , μ_j , and σ_j are functions dependent on x .

The MDN is a DNN with an output, mixture layer, that combines hidden layer outputs into coefficients, means, and standard deviations for the GMM. The MDN applied to our problem is shown in Figure 3.

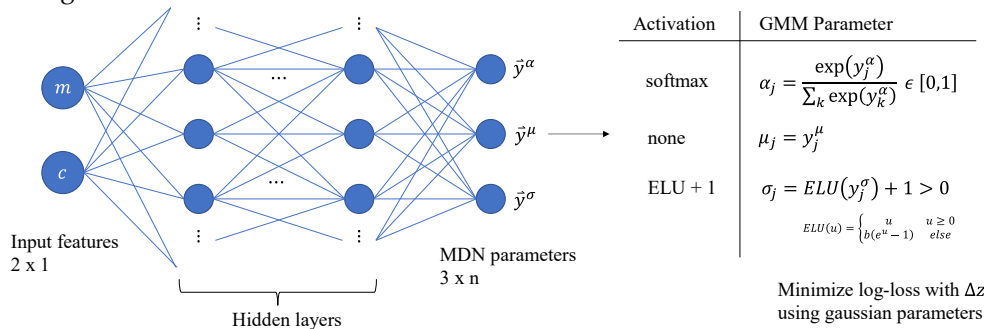


Figure 3. The MDN takes slope (m) and cover (c) as inputs, and outputs MDN parameters, which are transformed by relevant activation functions to produce GMM parameters. Note there is an entire distribution prediction per pair of (m, c) .

The difference between the MDN and a traditional DNN is the addition of the output layer, and the corresponding selection of activation functions, shown in the table in Figure 3. The softmax activation is applied to α_j as the coefficients must sum to 1. The means have no constraints, and therefore no special activation function. The standard deviations must be positive. The activation was originally the exponential function [25], but was later adapted to the exponential linear unit (ELU) + 1 with unit slope ($b = 1$ in Figure 3) [26], which exponentially approaches zero for DNN output dummy variable $u < 0$, and remains linear when $u \geq 0$. In this way, σ predictions cannot diverge

as easily, better conditioning the model. The loss function driving the MDN is the maximum likelihood of the negative log-loss of the GMM [25]:

$$L(\vec{\theta}) = -\frac{1}{M_T} \sum_{i=1}^{M_T} \log \left(\sum_{j=1}^n \alpha_{ij} N(\Delta z_i; \mu_{ij}, \sigma_{ij}^2) \right) \quad (5)$$

where $\vec{\theta}$ is a vector of $3n$ GMM parameters, including all α , μ , and σ over all components, computed over all training features M_T . Note that realizations Δz_i (training data) are input instead of the elevation residual axis, Δz .

After all components in $\vec{\theta}$ are estimated (the model is trained), we have the capability to predict at footprint-scale through Equation (4), or equivalently, through f . However, the RERD is of more importance to our study. Combining Equation (4) with Equation (3), the aggregate prediction is

$$f_{ROI}(\Delta z; X_P) \approx \frac{1}{M_P} \sum_{i=1}^{M_P} \sum_{j=1}^n \alpha_{ij} N(\Delta z; \mu_{ij}, \sigma_{ij}^2) \quad (6)$$

or Δz is distributed over the RERD prediction f_{ROI} , given all prediction features in the ROI X_P ($M_P = |X_P|$), which is approximately equal to the aggregate of Gaussian mixtures at *every* feature point in the ROI.

The MDN can predict an entire distribution over an arbitrary region, but from a mission planning perspective, the distribution is not as important as whether elevation residuals fall within a confidence interval threshold. For example, mission planners might specify that “90% of elevation residuals shall be within +/-2m.” Then we can specify the elevation threshold Δz_0 (e.g., 2m), and the percentage of elevation residuals within a predicted RERD that meet mission requirements is

$$P_{\Delta z} = \Pr(|\Delta z| < \Delta z_0 | X_P) = \int_{-\Delta z_0}^{\Delta z_0} f_{ROI}(\Delta z; X_P) \Delta z \quad (7)$$

or the probability that elevation residuals Δz are within the mission threshold Δz_0 , based on the new prediction ROI features X_P , is equal to the integral over the predicted RERD in the new ROI. In short, we quantify this confidence interval as $P_{\Delta z}$, which is directly informative in terms of mission requirements.

2.7. Training and Prediction

Training directly on quality footprints from each ROI can introduce bias in the model if the ROIs are not equally weighted. Figure 4 shows how the number of measurements is significantly different from one ROI to another. To balance the occurrences of the elevation residual per ROI, we set a maximum number of training points to the total number of footprints available (624,795 samples), then uniformly sample each ROI dataset equally. In this way, ROIs are represented equally in training.

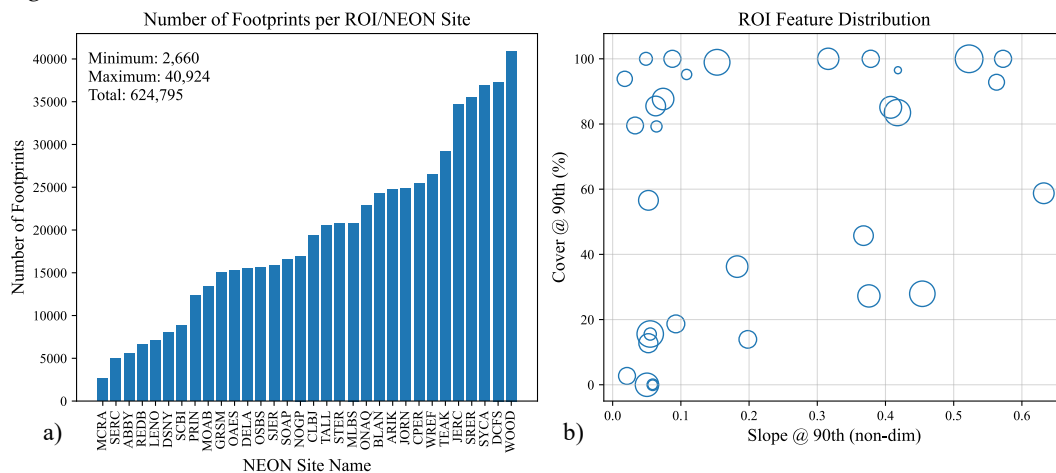


Figure 4. a) Number of footprints per ROI and b) 90th percentile feature distribution by NEON site, where in b), larger circles denote more data.

About half (19/32) of the measured RERDs have a mode of nearly zero (<1 cm), but some ROIs with large slopes and dense canopy cover reached modes of up to 7 cm. If the geocorrection pipeline were perfect, we should see modes of zero. We attribute the error to the culmination of segment length, geocorrection resolution of 1 m, potential misclassifications in reference ground signal at the base of trees, and temporal differences between NEON and the GEDI footprints. For a worst-case ROI (GRSM), we reduced the maximum segment length from 25 km to 5 km, and the corresponding mode offset changed from 21 cm to 7 cm, supporting the hypothesis that mode offsets are due to processing errors, not true GEDI elevation residual. Therefore, before training on any elevation residuals, we fit a Cauchy distribution to the observed RERD, and subtract the mode of the peak from all elevation residuals in the ROI, thereby removing processing errors from measured elevation residuals.

The model is trained by splitting the elevation residuals and features into training and testing sets, with 80% and 20% of the data, respectively. To determine the hyperparameters to train on, we performed a k-fold cross-validation. Because we are ultimately predicting RERDs, we split total data sets by ROI during cross-validation, not by footprint. For example, with a 5-fold cross-validation, the first fold may contain 25 ROIs of data and another 7 ROIs are set aside for validation. We then train on 80% of the 25 ROIs, predict on 20% of the same 25 ROIs to fit the model, then compare new predictions against the 7 holdout ROIs for validation. In this way, we determine how the model generalizes to data it has not seen. Because we have a sparse set of ROIs (only 32), a uniformly random split across all the ROIs could produce models that are not trained on significant portions of the feature space; therefore, we performed a k-fold cross-validation with random selection over subsets of the feature space. Table 1 lists the distribution of ROIs per feature subset.

Table 1. Cross-validation subset distribution.

Quadrant	Slope (non-dim)	Cover (%)	Number of ROIs
1	>0.25	>50	9
2	<0.25	>50	10
3	<0.25	<50	10
4	>0.25	<50	3

For each k-fold, we randomly select 2 ROIs from quadrants 1-3 and randomly select 1 ROI from quadrant 4 as test ROIs (7 total), then train on the rest.

We compare predictions to the measured RERD to tune the model, using the Jensen-Shannon Divergence (JSD) [42], which can be defined by the average of the Kullback-Liebler (KL) divergence D_{KL} between two distributions P and Q :

$$JSD(P||Q) = \frac{1}{2}D_{KL}(P||R) + \frac{1}{2}D_{KL}(Q||R) \quad (8a)$$

$$R = \frac{1}{2}(P + Q) \quad (8b)$$

The combination gives a “measure of the total divergence to the average distribution” [43], where, in our case, P is the measured RERD and Q is the predicted RERD.

We set the number of hidden layers to 2 with 100 neurons each. We used the python package TensorFlow to implement a DNN and modified it to fit the MDN structure, using the Adam optimizer. Then we applied a grid-search based 5-fold cross-validation to vary the number of Gaussian components over [3,5,7,10], batch size over [64,128,256], and learning rate over [0.1, 0.01, 0.001, 0.0001].

During training and cross-validation, we randomly sampled each ROI to preserve feature diversity while maintaining a sufficiently high number of training and prediction samples. We predict through Equations (6) and (7). However, on a global scale, prediction becomes computationally intractable. Relative to feature space, it is not vital to aggregate predictions for every single feature pair in a new ROI. Instead, to predict at scale, we downsample the feature space in a

new ROI, predicting on weighted feature bins instead of individual feature points, approximating the RERD prediction (Appendix A).

2.8. Method Summary

Summarizing training, we

- Choose and download 32 NEON reference lidar sites;
- Download GEDI L2B version 2 granules over those sites and extract quality footprints;
- Geocorrect GEDI L2B footprints with DTMs built from ground classified NEON reference lidar;
- Difference geocorrected elevations with DTMs to obtain elevation residuals;
- Extract slope derived from MERIT DEM and canopy cover from Copernicus over each NEON site;
- Remove mode/processing bias in RERDs from measured elevation residuals;
- And finally train the MDN, correlating slope and cover to elevation residual.

Summarizing the regional prediction, we

- Choose a user-defined ROI somewhere on Earth within 60°S to 78.25°N;
- Extract slope derived from MERIT DEM and canopy cover from Copernicus over the new ROI;
- Downsample the feature space into weighted feature bins for faster prediction;
- Predict on weighted feature bins, obtaining approximate GMM parameters;
- And finally aggregate all predicted GMMs over the new ROI.

3. Results

3.1. Model Validation and Accuracy

To train the model, we minimize the negative log-loss of the GMM. To optimally choose hyperparameters in cross-validation, we minimize the JSD between predicted and measured RERDs. The JSD was minimized to 0.0098 for 7 mixture components, with a batch size of 64 and a learning rate of 0.001. Divergence is a convenient measure of how different a predicted distribution is from a validation (or measured) distribution. However, divergence is not easily interpretable in terms of model accuracy. The model assumes a non-Gaussian distribution, but we can present accuracy in terms of an “equivalent standard deviation” σ_{EQ} or the standard deviation of a Gaussian fit to the measured or predicted distributions. In this regard, we model with no Gaussian assumption, but present cross-validation prediction accuracy in a way that is more easily understood. Note that our prediction variable is not a kind of mean, it is a type of variance (through the predicted distribution), because we are quantifying uncertainty in the GEDI elevation retrieval. It would not make sense to present results on an “equivalent mean” because the predicted mean is nearly zero, always, because we estimate only the error in georectified elevation measurement. In Table 2, we present cross-validation prediction accuracy over all folds.

Table 2. Cross-validation error statistics.

Row	Error Type	Using all feature points	Using feature bins
A	Average error in σ_{EQ}	-0.06 m	-0.04 m
B	Standard deviation of σ_{EQ}	0.18 m	0.19 m
C	Average of relative error of σ_{EQ}	15.9%	26.7%

In Table 2 row A, we show that the mean of the difference in σ_{EQ} of our predictions vs. the measured data is -0.06 m, with no feature bin downsampling. That is, our σ_{EQ} prediction is consistent with measured RERDs. In Table 2 row B, we evaluate the standard deviation of our predicted σ_{EQ} , which is 0.18 m. That is, we may, on average predict σ_{EQ} with a 6 cm bias, but the true error may vary between +/-18 cm at 1σ . And finally in Table 2 row C, we show the relative error of our prediction for σ_{EQ} , which is 15.9%, or our prediction of σ_{EQ} is off, on average, by 15.9%. In the last column, we instead predict by downsampling the dataset in feature space (Appendix A). Rows

A and B remain consistent, but row C is increased because there are several predicted distributions that have small variance (in flat, semi-barren areas), and with a lower prediction accuracy due feature binning, their relative error is high. Because Table 2 results are based on cross-validation, the values are representative of generalization error.

Alternatively, we look at the in-sample relative error of σ_{EQ} by ROI. That is, we train on all data, and test on all data to conduct an intercomparison of each ROI's relative error, shown in Figure 5.

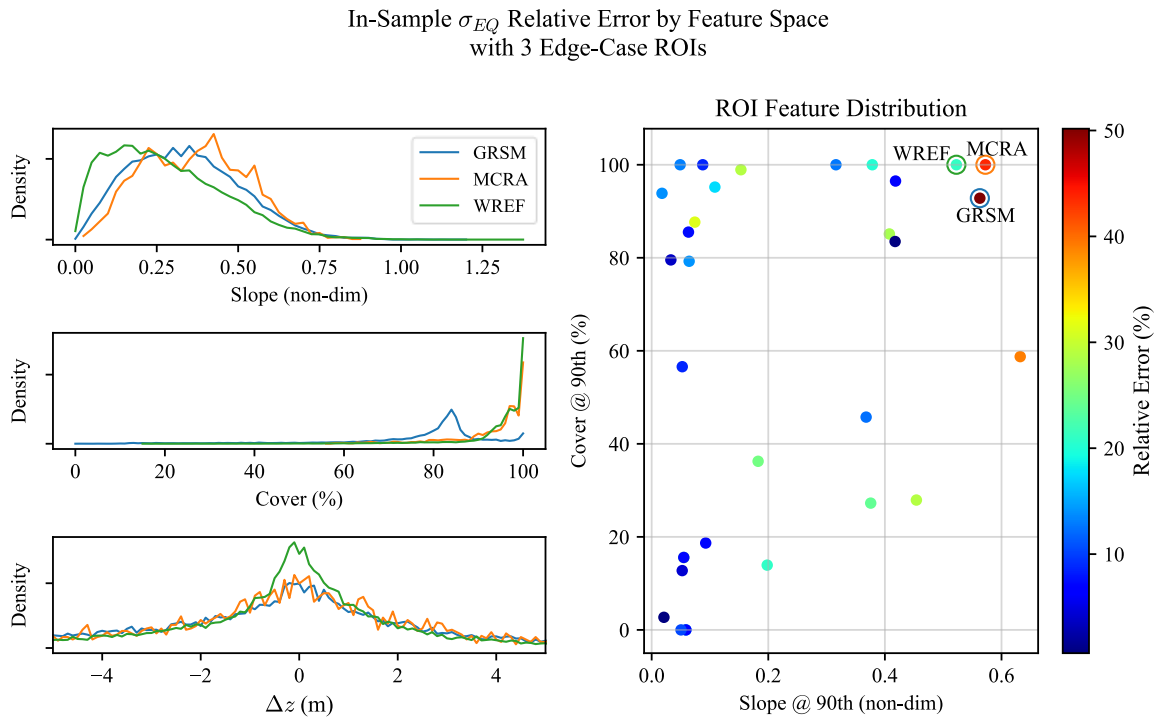


Figure 5. Relative error distribution over feature space.

Figure 5 shows the relative error change across the feature space for each ROI. In addition, we have highlighted 3 different ROIs at high slope and high canopy cover: WREF, MCRA, and GRSM, with slope, cover, and RERDs on the left side of Figure 5. The relative error of GRSM is 50.2%, but WREF and MCRA are 21.1% and 43.7% respectively. The RERD of GRSM nearly matches that of MCRA, but the slope and canopy cover of MCRA is higher. According to slope and cover, the model predicts GRSM as a sharper distribution than is measured, more similar to WREF. Likely, slope and canopy cover are not enough to quantify the variation in GRSM vs. MCRA. Similarly, 9 other ROIs are predicted with relative error in excess of 25% throughout the feature space. The outlier error suggests that the model is not overfitting the training data and that the error is likely a result of predicting from only two features and generalizing the model during cross-validation.

In light of the errors, we exchanged the cover feature for the GEDI L2B cover dataset to predict in-sample errors, using the same cross-validated model hyperparameters, as a quantitative investigation. The ROI GRSM relative error reduced from 50.2% to 30.3%. WREF increased by 7.5% to 28.7% relative error, and MCRA reduced by 30.1% to 13.6% relative error. The average relative error remained similar to Table 2 row C, at 16.3%, implying that outlier ROIs are predicted better given higher quality cover. While 72% (23/32) of ROIs are within 25% relative error given cover derived from Copernicus, this comparison against the GEDI cover dataset indicates that the Copernicus cover dataset – and the choice of only two features – is insufficient to reliably predict mountainous, forested regions at meter-scale elevation error.

3.2. Application: Global Prediction

The presented methodology offers a technique to predict the error in GEDI-like elevation measurement from footprint to global scale, given clear-sky conditions. As an example application, we predict $P_{\Delta z}$ at 0.25x0.25 deg. resolution globally (60°S to 78.25°N). We first predict $P_{\Delta z}$ at 10x10 km² in the sinusoidal equal area projection, which produces unbiased statistics regardless of latitude [44], then resample outputs to 0.25x0.25 deg. resolution on a latitude-longitude geographic projection with the WGS84 ellipsoid. In addition, because MERIT DEM includes inland lakes and rivers, we masked $P_{\Delta z}$ outputs by resampled Ocean and Permanent Water Bodies classifications at 0.25x0.25 deg. from the Copernicus Moderate Dynamic Land Cover 100m dataset (version 3) [45].

In Figure 6, we show $P_{\Delta z}$ from 60°S to 78.25°N for $\Delta z_0 = 1$ m and $\Delta z_0 = 3$ m. For $\Delta z_0 = 1$ m, the lower 10th percentile of $P_{\Delta z}$ is 60.7%, the median is 88.5%, and the 90th percentile is 96.6%. For $\Delta z_0 = 3$ m, the lower 10th percentile of $P_{\Delta z}$ is 83.4%, the median is 96.4%, and the 90th percentile is 99.3%. In Figure 6a, $\Delta z_0 = 1$ m, and there are mountainous areas such as the Rocky Mountains, Himalayas, and the Andes Mountains, where the percent of elevation residuals within 1m is ~50%. Further, in areas of dense vegetation, $P_{\Delta z}$ is also quite low, such as throughout the equatorial latitudes in South America, Africa, and Indonesia. In Figure 6b, $\Delta z_0 = 3$ m, so there is more room for elevation error, and $P_{\Delta z}$ is in general much higher than in Figure 6a. That is, it is much more likely that elevation residuals are within 3m than 1m. And in flatter, barren regions, $P_{\Delta z}$ is high for both cases. These trends follow qualitative understanding that, in regions with high slopes and dense cover, geocorrected elevation retrieval has more error.

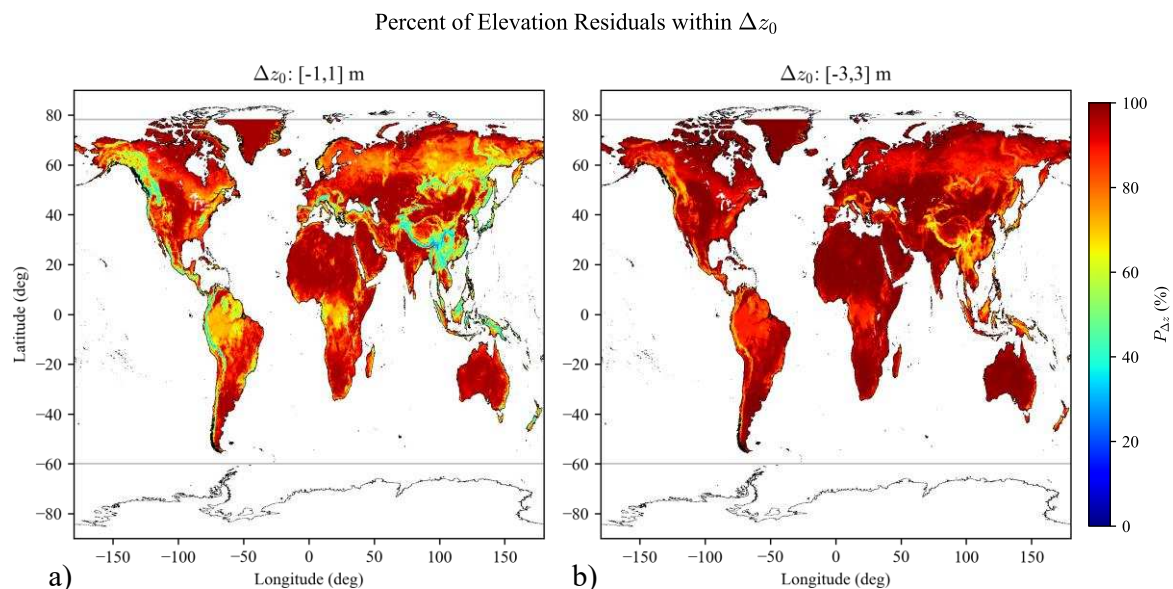


Figure 6. Global prediction with a) $\Delta z_0 = 1$ m and b) $\Delta z_0 = 3$ m.

4. Discussion and Conclusions

This study has shown how to map slope and canopy cover to clear-sky measured elevation error for a spaceborne, 1064 nm, 25 m footprint full-waveform lidar. Since we have removed platform-specific and processing errors such as georectification, we have fit the terrain retrieval error only considering the waveform itself. While the immediate applications of the trained MDN are limited, the predicted distributions could be coupled together with simulated errors of spaceborne platform positional accuracy and geolocation accuracy. In addition, optical depth could be queried from external cloud predictions to remove simulated footprints due to obscuration. In essence, we have modeled a challenging piece of the entire terrain extraction pipeline, and by convolving predictive outputs with other, realistic error sources, mission planners can generate a data-informed uncertainty model.

We have trained the MDN at the footprint scale and developed a statistical technique through feature binning to predict at global scale while remaining computationally feasible. Through cross-

validation results, predicting from only slope and canopy cover features has 15.9% mean relative error of σ_{EQ} . From in-sample analyses, the MDN predicts for flat, barren areas with less than 25% relative error of σ_{EQ} . For more mountainous and forested regions, the relative error can reach 50.2%. The error is understandable considering that only average forest canopy cover is used from Copernicus 2019. Furthermore, GEDI footprints are not filtered by leaf on/leaf off conditions and footprints are queried from 2019 through 2021. Canopy cover could be enhanced by including Moderate Resolution Imaging Spectroradiometer (MODIS) and Landsat tree cover datasets [46,47], which are on the GEDI L2B data product per footprint (for training) and available at global scale via external datasets for prediction. Potentially, low spatial resolution canopy cover with higher temporal resolution might sufficiently supplement a high spatial resolution canopy cover dataset, effectively breaking temporal and spatial cover into separate features. MERIT DEM derived slope is likely sufficient and is clearly impactful to the global prediction, however it could be supported by other external terrain datasets, such as the Forest and Buildings Removed Digital Elevation Model (FABDEM) [48]. Besides terrain slope and canopy cover, there are other factors that influence GEDI elevation measurement fidelity, such as canopy height, land cover type, solar illumination, and beam sensitivity [28], and the MDN predictions would likely fare better including these features.

While the MDN is an effective choice to model a GMM, it is somewhat heavy for the task of predicting a univariate, Cauchy-like distribution. Experimenting with Gaussian and Cauchy distribution models, we found them in much more error than the GMM. There may be an in-between model based on regularized maximum likelihood to predict a Student's T-distribution, that would suffice over the MDN. With a simpler model, predictions could be faster, and the model could be less prone to overfitting. Or, instead of training at footprint scale and aggregating distributions to predict regionally, one can directly train σ_{EQ} regionally through linear regression. The drawback is that predictions may be fixed to only the spatial scale trained on, whereas the intent of the MDN strategy is applicability to a wide response domain, such as predictions of elevation residual quality from footprints to continents.

Author Contributions: Conceptualization, LAM.; methodology, JS; software, JS; validation, JS; formal analysis, JS; investigation, JS; resources, JS and LAM; data curation, JS; writing—original draft preparation, JS; writing—review and editing, LAM; visualization, JS; supervision, LAM; project administration, LAM; funding acquisition, LAM. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by a NASA Space Technology Graduate Research Opportunity (NSTGRO), grant number 80NSSC22K1213.

Data Availability Statement: The GEDI L2B data used in this study are available at https://lpdaac.usgs.gov/products/gedi02_bv002/. MERIT DEM is available at https://hydro.iis.u-tokyo.ac.jp/~yamadai/MERIT_DEM/. The Copernicus Global Land Cover product is available at <https://zenodo.org/record/3939050>.

Acknowledgments: The authors would like to thank Dr. Tan Bui for insights on neural network modeling direction. Additionally, we would like to acknowledge the support of the University of Texas at Austin leadership from the Aerospace Engineering and Engineering Mechanics Department at the Cockrell School of Engineering and the research staff at the Center for Space Research.

Conflicts of Interest: The authors declare no conflicts or competing interests associated with this manuscript.

Appendix A

For MERIT DEM 90m resolution and Copernicus land cover of 100m resolution, there are about 15 billion feature points to predict globally, no matter the resolution over which the prediction takes place. For every feature point, we predict $3n$ parameters, or 21 parameters for $n = 7$ Gaussian components. Predicting every feature globally, we would generate ~315 billion parameters, which is computationally infeasible, and far too detailed for global scale.

To create a heatmap that is usable in a pre-phase A mission development concept, we desire resolutions on the order of $10 \times 10 \text{ km}^2$. Let a "grid cell" be a $10 \times 10 \text{ km}^2$ areal region on Earth's land surface. Figure A1 shows the feature space of an arbitrary grid cell. In this ROI, we see every feature

point $x \in X$, where X is the set of all feature points in the ROI. Most slopes are below 1.5 ($\sim 56^\circ$) and canopy cover extends from 0% to $\sim 100\%$. The number of points is M .

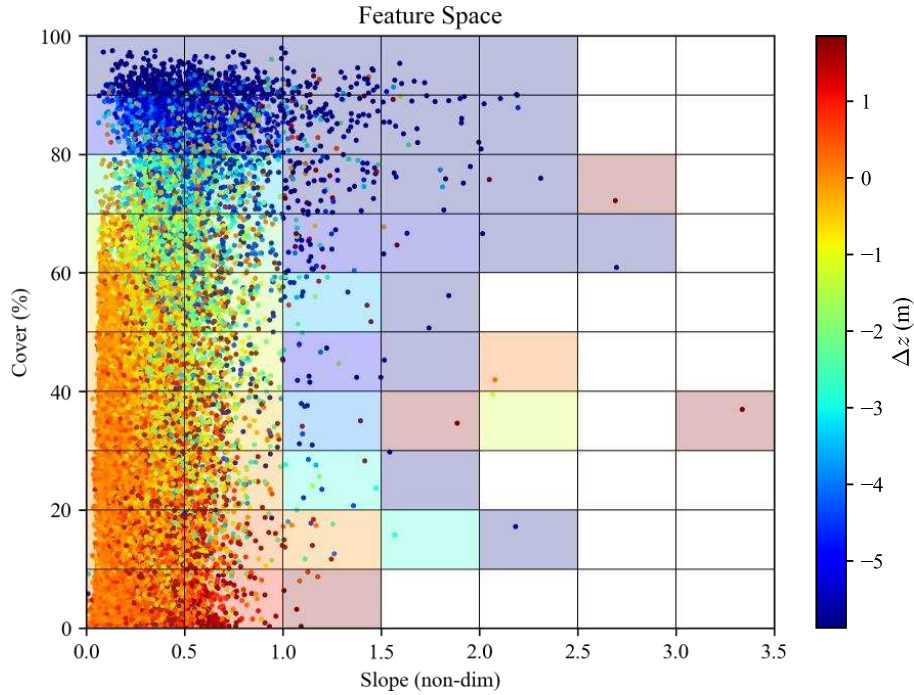


Figure A1. Pictorial representation of pointwise feature space for an arbitrary grid cell.

To downsample f_{ROI} , we raster the feature space with bins (colored in Figure A1), and weigh predictions by the number of points in each bin. Expanding Equation (3),

$$f_{ROI}(\Delta Z; X) \approx \frac{1}{M} \left[\sum_{i \in I_1} f(\Delta Z; x_i) + \sum_{i \in I_2} f(\Delta Z; x_i) + \dots + \sum_{i \in I_L} f(\Delta Z; x_i) \right] \quad (A1a)$$

where

$$I_k = \{i: x_i \in X_k\} \quad (A1b)$$

$$M_k = |I_k| \quad (A1c)$$

$X_k \subset X$ is a subset of all feature points, such as the points within colored bins in Figure A1, and M_k is the number of feature points in bin k . We assume that, within each bin, the distribution f does not change much. Therefore,

$$\sum_{i \in I_k} f(\Delta Z; x_i) \approx M_k f(\Delta Z; \bar{x}_k) \quad (A2)$$

where \bar{x}_k is the geometric center of a feature bin k (e.g., at $m = 0.75$, $c = 5$). f_{ROI} is approximated as

$$f_{ROI}(\Delta Z; X) \approx \sum_{k=1}^L \frac{M_k}{M} f(\Delta Z; \bar{x}_k) \quad (A3)$$

where $L \ll M$. Given thousands of grid cells and millions of feature points, f_{ROI} defined via L approximately constant distributions is more scalable to calculate. In our global prediction, we chose a bin size of 0.1×10.0 . Based on the GMM, f_{ROI} is approximated as

$$f_{ROI}(\Delta Z; X_P) \approx \sum_{k=1}^{L_P} \frac{M_{P,k}}{M_P} \sum_{j=1}^n \bar{\alpha}_{kj} N(\Delta Z; \bar{\mu}_{kj}, \bar{\sigma}_{kj}^2) \quad (A4)$$

where $L_p \ll M_p$ is the number of feature bins in prediction, and overbars indicate predictions made from the centers of feature bins, not feature points.

References

1. Committee on the Decadal Survey for Earth Science and Applications from Space; Space Studies Board; Division on Engineering and Physical Sciences; National Academies of Sciences, Engineering, and Medicine *Thriving on Our Changing Planet: A Decadal Strategy for Earth Observation from Space*; National Academies Press: Washington, D.C., 2018; p. 24938; ISBN 978-0-309-46757-5.
2. Donnellan, A.; Harding, D.; Lundgren, P.; Wessels, K.; Simard, M.; Parrish, C.; Jones, C.; Lou, Y.; Stoker, J.; Ranson, K.J.; et al. *Observing Earth's Changing Surface Topography & Vegetation Structure: A Framework for the Decade*; NASA's Surface Topography and Vegetation Incubation Study Team Report; Jet Propulsion Laboratory, California Institute of Technology, NASA Goddard Space Flight Center, George Mason University, Oregon State University, United States Geological Survey, 2021;
3. Webb, C.E.; Jay, Z.H.; Abdalati, W. *The Ice, Cloud, and Land Elevation Satellite (ICESat) Summary Mission Timeline and Performance Relative to Pre-Launch Mission Success Criteria*; NASA Goddard Space Flight Center, 2012;
4. Markus, T.; Neumann, T.; Martino, A.; Abdalati, W.; Brunt, K.; Csatho, B.; Farrell, S.; Fricker, H.; Gardner, A.; Harding, D.; et al. The Ice, Cloud, and Land Elevation Satellite-2 (ICESat-2): Science Requirements, Concept, and Implementation. *Remote Sensing of Environment* **2017**, *190*, 260–273, doi:10.1016/j.rse.2016.12.029.
5. Dubayah, R.; Blair, J.B.; Goetz, S.; Fatoyinbo, L.; Hansen, M.; Healey, S.; Hofton, M.; Hurtt, G.; Kellner, J.; Luthcke, S.; et al. The Global Ecosystem Dynamics Investigation: High-Resolution Laser Ranging of the Earth's Forests and Topography. *Science of Remote Sensing* **2020**, *1*, 100002, doi:10.1016/j.srs.2020.100002.
6. Choi, C.; Cazcarra-Bes, V.; Guliaev, R.; Pardini, M.; Papathanassiou, K.P.; Qi, W.; Armston, J.; Dubayah, R.O. Large-Scale Forest Height Mapping by Combining TanDEM-X and GEDI Data. *IEEE J. Sel. Top. Appl. Earth Observations Remote Sensing* **2023**, *16*, 2374–2385, doi:10.1109/JSTARS.2023.3244866.
7. Ma, L.; Hurtt, G.; Ott, L.; Sahajpal, R.; Fisk, J.; Lamb, R.; Tang, H.; Flanagan, S.; Chini, L.; Chatterjee, A.; et al. Global Evaluation of the Ecosystem Demography Model (ED v3.0). *Geosci. Model Dev.* **2022**, *15*, 1971–1994, doi:10.5194/gmd-15-1971-2022.
8. Liang, M.; Duncanson, L.; Silva, J.A.; Sedano, F. Quantifying Aboveground Biomass Dynamics from Charcoal Degradation in Mozambique Using GEDI Lidar and Landsat. *Remote Sensing of Environment* **2023**, *284*, 113367, doi:10.1016/j.rse.2022.113367.
9. Ren, C.; Jiang, H.; Xi, Y.; Liu, P.; Li, H. Quantifying Temperate Forest Diversity by Integrating GEDI LiDAR and Multi-Temporal Sentinel-2 Imagery. *Remote Sensing* **2023**, *15*, 375, doi:10.3390/rs15020375.
10. Hancock, S.; McGrath, C.; Lowe, C.; Davenport, I.; Woodhouse, I. Requirements for a Global Lidar System: Spaceborne Lidar with Wall-to-Wall Coverage. *R. Soc. open sci.* **2021**, *8*, 211166, doi:10.1098/rsos.211166.
11. Hansen, J.N.; Hancock, S.; Prade, L.; Bonner, G.M.; Chen, H.; Davenport, I.; Jones, B.E.; Purslow, M. Assessing Novel Lidar Modalities for Maximizing Coverage of a Spaceborne System through the Use of Diode Lasers. *Remote Sensing* **2022**, *14*, 2426, doi:10.3390/rs14102426.
12. Crisp, N.H.; McGrath, C.N.; Roberts, P.C.E.; Edmondson, S.; Haigh, S.J.; Holmes, B.E.A.; Rojas, A.M.; Oiko, V.T.A.; Sinpetru, L.A.; Smith, K.L.; et al. Very Low Earth Orbit Constellations for Earth Observation. In *Proceedings of the 73rd International Astronautical Congress*; September 15 2022.
13. McGrath, C.; Lowe, C.; Macdonald, M.; Hancock, S. Investigation of Very Low Earth Orbits (VLEOs) for Global Spaceborne Lidar. *CEAS Space J* **2022**, *14*, 625–636, doi:10.1007/s12567-022-00427-2.
14. Alvarado, M.J.; Payne, V.H.; Mlawer, E.J.; Uymin, G.; Shephard, M.W.; Cady-Pereira, K.E.; Delamere, J.S.; Moncet, J.-L. Performance of the Line-By-Line Radiative Transfer Model (LBLRTM) for Temperature, Water Vapor, and Trace Gas Retrievals: Recent Updates Evaluated with IASI Case Studies. *Atmos. Chem. Phys.* **2013**, *13*, 6687–6711, doi:10.5194/acp-13-6687-2013.
15. Gastellu-Etchegorry, J.-P.; Yin, T.; Lauret, N.; Cajgfinger, T.; Gregoire, T.; Grau, E.; Feret, J.-B.; Lopes, M.; Guilleux, J.; Dedieu, G.; et al. Discrete Anisotropic Radiative Transfer (DART 5) for Modeling Airborne and Satellite Spectroradiometer and LIDAR Acquisitions of Natural and Urban Landscapes. *Remote Sensing* **2015**, *7*, 1667–1701, doi:10.3390/rs70201667.
16. Wang, Y.; Kallel, A.; Yang, X.; Regaieg, O.; Lauret, N.; Guilleux, J.; Chavanon, E.; Gastellu-Etchegorry, J.-P. DART-Lux: An Unbiased and Rapid Monte Carlo Radiative Transfer Method for Simulating Remote Sensing Images. *Remote Sensing of Environment* **2022**, *274*, 112973, doi:10.1016/j.rse.2022.112973.
17. Yang, X.; Wang, Y.; Yin, T.; Wang, C.; Lauret, N.; Regaieg, O.; Xi, X.; Gastellu-Etchegorry, J.P. Comprehensive LiDAR Simulation with Efficient Physically-Based DART-Lux Model (I): Theory, Novelty, and Consistency Validation. *Remote Sensing of Environment* **2022**, *272*, 112952, doi:10.1016/j.rse.2022.112952.

18. Hu, Y.; Hou, A.; Ma, Q.; Zhao, N.; Xu, S.; Fang, J. Analytical Formula to Investigate the Modulation of Sloped Targets Using LiDAR Waveform. *IEEE Trans. Geosci. Remote Sensing* **2022**, *60*, 1–12, doi:10.1109/TGRS.2021.3099478.
19. Hao, Q.; Cheng, Y.; Cao, J.; Zhang, F.; Zhang, X.; Yu, H. Analytical and Numerical Approaches to Study Echo Laser Pulse Profile Affected by Target and Atmospheric Turbulence. *Opt. Express* **2016**, *24*, 25026, doi:10.1364/OE.24.025026.
20. Hancock, S.; Armston, J.; Hofton, M.; Sun, X.; Tang, H.; Duncanson, L.I.; Kellner, J.R.; Dubayah, R. The GEDI Simulator: A Large-Footprint Waveform Lidar Simulator for Calibration and Validation of Spaceborne Missions. *Earth and Space Science* **2019**, *6*, 294–310, doi:10.1029/2018EA000506.
21. Huettermann, S.; Jones, S.; Soto-Berelov, M.; Hislop, S. Intercomparison of Real and Simulated GEDI Observations across Sclerophyll Forests. *Remote Sensing* **2022**, *14*, 2096, doi:10.3390/rs14092096.
22. Blair, J.B.; Hofton, M.A. Modeling Laser Altimeter Return Waveforms over Complex Vegetation Using High-Resolution Elevation Data. *Geophys. Res. Lett.* **1999**, *26*, 2509–2512, doi:10.1029/1999GL010484.
23. Coops, N.C.; Tompalski, P.; Goodbody, T.R.H.; Queinnec, M.; Luther, J.E.; Bolton, D.K.; White, J.C.; Wulder, M.A.; Van Lier, O.R.; Hermosilla, T. Modelling Lidar-Derived Estimates of Forest Attributes over Space and Time: A Review of Approaches and Future Trends. *Remote Sensing of Environment* **2021**, *260*, 112477, doi:10.1016/j.rse.2021.112477.
24. Lang, N.; Kalischek, N.; Armston, J.; Schindler, K.; Dubayah, R.; Wegner, J.D. Global Canopy Height Regression and Uncertainty Estimation from GEDI LIDAR Waveforms with Deep Ensembles. *Remote Sensing of Environment* **2022**, *268*, 112760, doi:10.1016/j.rse.2021.112760.
25. Bishop, C.M. *Mixture Density Networks*; Neural Computing Research Group; Department of Computer Science and Applied Mathematics: Birmingham, B4 7ET UK, 1994; pp. 1–25.
26. Brando Guillaumes, A. *Mixture Density Networks for Distribution and Uncertainty Estimation*. Master thesis, Universitat Politècnica de Catalunya, 2017.
27. Li, X.; Liu, C.; Wang, Z.; Xie, X.; Li, D.; Xu, L. Airborne LiDAR: State-of-the-Art of System Design, Technology and Application. *Meas. Sci. Technol.* **2021**, *32*, 032002, doi:10.1088/1361-6501/abc867.
28. Liu, A.; Cheng, X.; Chen, Z. Performance Evaluation of GEDI and ICESat-2 Laser Altimeter Data for Terrain and Canopy Height Retrievals. *Remote Sensing of Environment* **2021**, *264*, 112571, doi:10.1016/j.rse.2021.112571.
29. Schleich, A.; Durrieu, S.; Soma, M.; Vega, C. *Improving GEDI Footprint Geolocation Using a High Resolution Digital Terrain Model*; 2023;
30. Xu, Y.; Ding, S.; Chen, P.; Tang, H.; Ren, H.; Huang, H. Horizontal Geolocation Error Evaluation and Correction on Full-Waveform LiDAR Footprints via Waveform Matching. *Remote Sensing* **2023**, *15*, 776, doi:10.3390/rs15030776.
31. Wang, C.; Elmore, A.J.; Numata, I.; Cochrane, M.A.; Shaogang, L.; Huang, J.; Zhao, Y.; Li, Y. Factors Affecting Relative Height and Ground Elevation Estimations of GEDI among Forest Types across the Conterminous USA. *GIScience & Remote Sensing* **2022**, *59*, 975–999, doi:10.1080/15481603.2022.2085354.
32. NEON (National Ecological Observatory Network) Discrete Return LiDAR Point Cloud (DP1.30003.001), RELEASE-2023 2015.
33. Krause, K.; Goulden, T. *NEON L0-to-L1 Discrete Return LiDAR Algorithm Theoretical Basis Document (ATBD)*; National Ecological Observatory Network, 2015;
34. Yamazaki, D.; Ikeshima, D.; Tawatari, R.; Yamaguchi, T.; O'Loughlin, F.; Neal, J.C.; Sampson, C.C.; Kanae, S.; Bates, P.D. A High-accuracy Map of Global Terrain Elevations. *Geophysical Research Letters* **2017**, *44*, 5844–5853, doi:10.1002/2017GL072874.
35. Buchhorn, M.; Smets, B.; Bertels, L.; Roo, B.D.; Lesiv, M.; Tsendbazar, N.-E.; Li, L.; Tarko, A. *Copernicus Global Land Service: Land Cover 100m: Version 3 Globe 2015-2019: Product User Manual*; Zenodo, 2020;
36. Luthcke, S.B.; Rebold, T.; Thomas, T.; Pennington, T. Algorithm Theoretical Basis Document (ATBD) for GEDI Waveform Geolocation for L1 and L2 Products Version 1.0 2019.
37. Hofton, M.; Blair, J.B. Algorithm Theoretical Basis Document (ATBD) for GEDI Transmit and Receive Waveform Processing for L1 and L2 Products Version 1.0 2019.
38. Beck, J.; Writ, B.; Luthcke, S.B.; Hofton, M.; Armston, J. Global Ecosystem Dynamics Investigation (GEDI) Level 1B User Guide Version 2.0 2021.
39. Neuenschwander, A.L.; Magruder, L.A. Canopy and Terrain Height Retrievals with ICESat-2: A First Look. *Remote Sensing* **2019**, *11*, 1721, doi:10.3390/rs11141721.
40. Neuenschwander, A.; Guenther, E.; White, J.C.; Duncanson, L.; Montesano, P. Validation of ICESat-2 Terrain and Canopy Heights in Boreal Forests. *Remote Sensing of Environment* **2020**, *251*, 112110, doi:10.1016/j.rse.2020.112110.

41. Tang, H.; Armston, J. Algorithm Theoretical Basis Document (ATBD) for GEDI L2B Footprint Canopy Cover and Vertical Profile Metrics Version 1.0 2019.
42. Lin, J. Divergence Measures Based on the Shannon Entropy. *IEEE Trans. Inform. Theory* **1991**, *37*, 145–151, doi:10.1109/18.61115.
43. Nielsen, F. On the Jensen–Shannon Symmetrization of Distances Relying on Abstract Means. *Entropy* **2019**, *21*, 485, doi:10.3390/e21050485.
44. Snyder, J.P. *Map Projections: A Working Manual*; U.S. Government Printing Office, 1987;
45. Buchhorn, M.; Smets, B.; Bertels, L.; Roo, B.D.; Lesiv, M.; Tsendbazar, N.-E.; Li, L.; Tarko, A. *Copernicus Global Land Service: Land Cover 100m: Version 3 Globe 2015-2019: Product User Manual*; Zenodo, 2020;
46. Sexton, J.O.; Song, X.-P.; Feng, M.; Noojipady, P.; Anand, A.; Huang, C.; Kim, D.-H.; Collins, K.M.; Channan, S.; DiMiceli, C.; et al. Global, 30-m Resolution Continuous Fields of Tree Cover: Landsat-Based Rescaling of MODIS Vegetation Continuous Fields with Lidar-Based Estimates of Error. *International Journal of Digital Earth* **2013**, *6*, 427–448, doi:10.1080/17538947.2013.786146.
47. DiMiceli, C.; Carroll, M.; Sohlberg, R.; Kim, D.-H.; Kelly, M.; Townshend, J. MOD44B MODIS/Terra Vegetation Continuous Fields Yearly L3 Global 250m SIN Grid V006 2015.
48. Hawker, L.; Uhe, P.; Paulo, L.; Sosa, J.; Savage, J.; Sampson, C.; Neal, J. A 30 m Global Map of Elevation with Forests and Buildings Removed. *Environ. Res. Lett.* **2022**, *17*, 024016, doi:10.1088/1748-9326/ac4d4f.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.