farmers for forests

AI Based Automated Calculation of Carbon Stock in Agroforestry Plantations

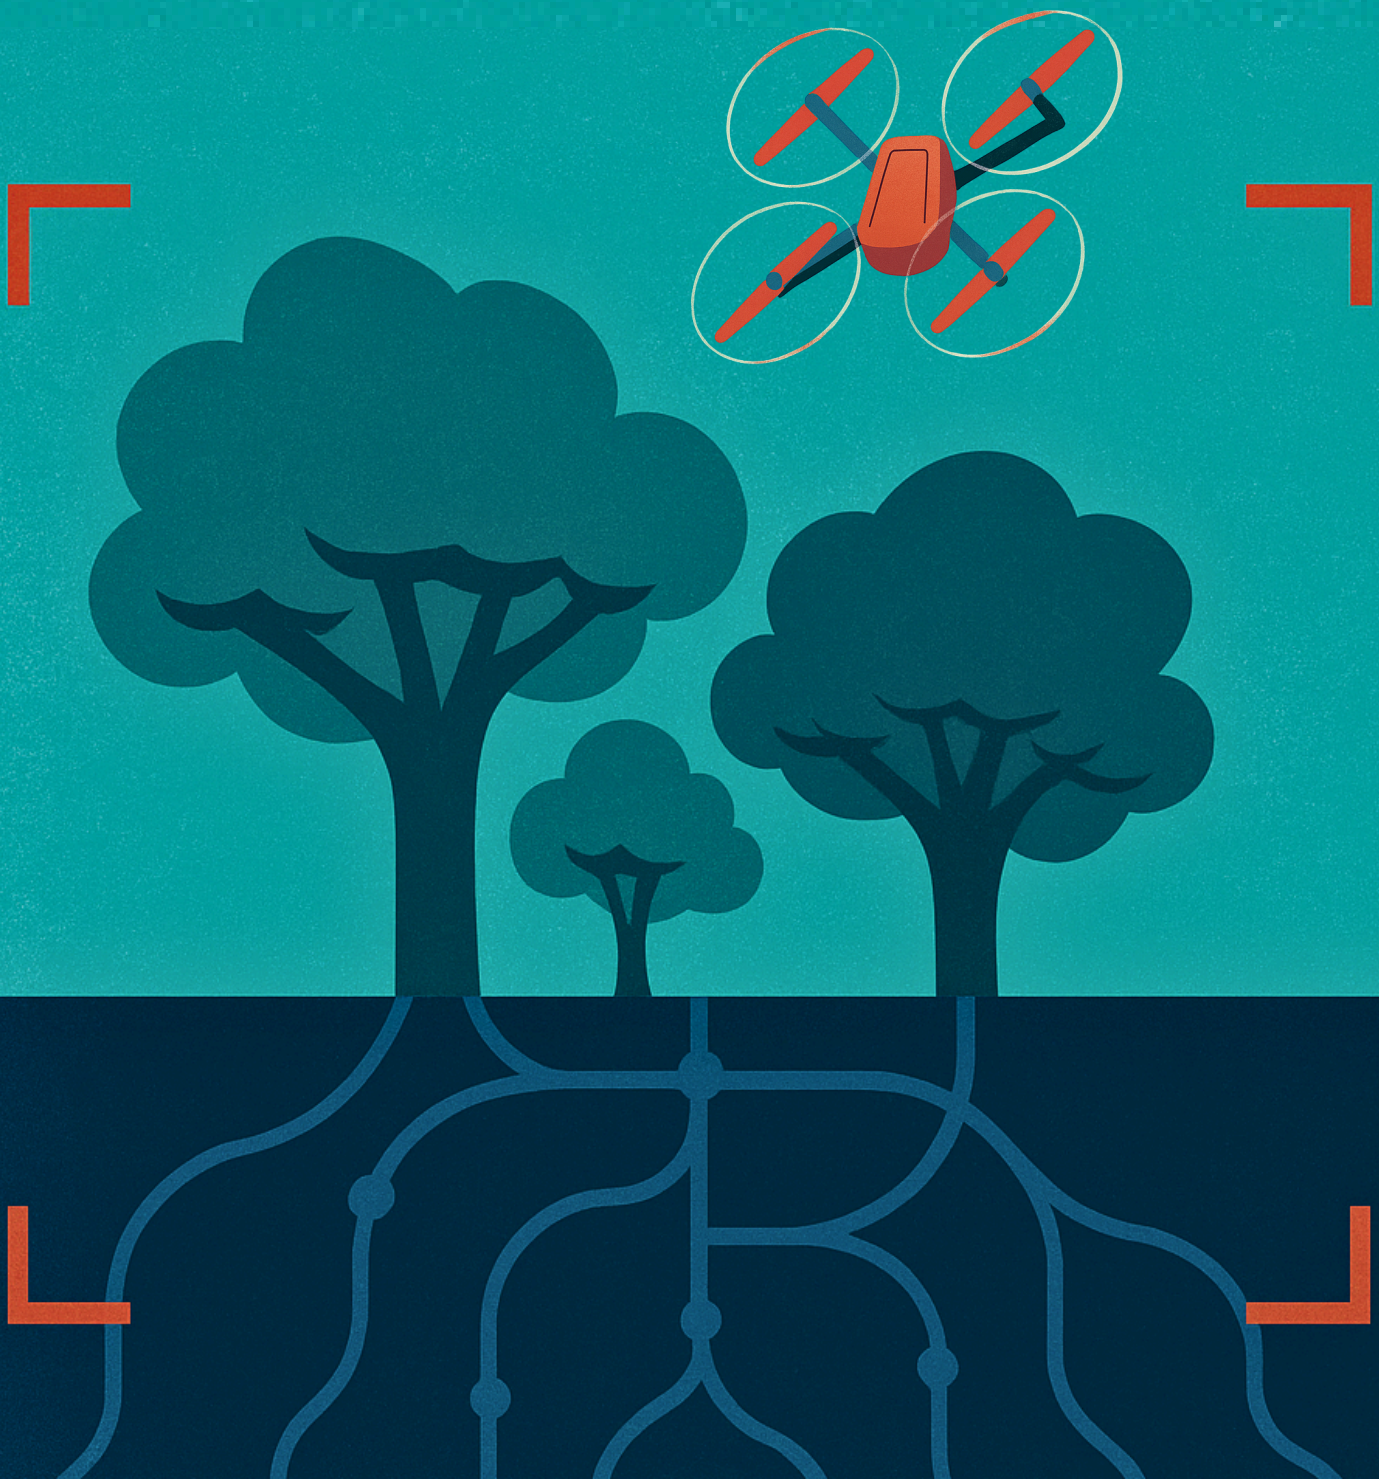

Authored By

Dr. Abhishek Gupta
Ramanujan Faculty Fellow at Indian Institute of Technology (IIT) Goa

Makarand Datar
Vice President, Wildlife & Tech at Farmers for Forests

Pravin Mulay
Vice President, Research & Tech at Farmers for Forests

August 2025
Research Paper Version 1.0

Acknowledgements

A huge shoutout to Jagdish Bade and Amol Phade, Farmers for Forests' (F4F) amazing drone pilot team for braving the sun, wind, and the occasional curious cow to capture our drone data.

To Manish Patil, Abhijeet Pawar, Vijendra Singh, Thomas Mampilly, Nikhil Raj, and Mohd. Shamoon - the impeccable F4F tech team - for quietly pulling the strings behind the scenes and making sure the research paper gets completed.

To Aditya Avinashe, whose strategic leadership in streamlining drone operations and on-ground coordination made this work possible.

To Krutika Ravishankar and Arti Dhar for stress-testing how the model would actually be used on land, and connecting it back to the bigger picture of carbon finance and impact.

To our philanthropic funders - thank you for seeing beyond just "a cool AI project" and believing in our bigger vision: building open tools that don't just work for us, but democratize how land, trees, and carbon are measured everywhere.

And finally, to the F4F field team - thank you for collecting all the data that informed our model. You collected valuable information, sometimes through mud and scorching heat - but almost always with humor and a smile.

Table of Content

1. Introduction

The urgency of climate change underscores the pressing need for carbon offset initiatives alongside efforts to reduce greenhouse gas emissions. Interventions in agriculture and forestry are effective strategies for achieving carbon offsets. These interventions sequester carbon while enhancing climate resilience for farmers and local communities. The voluntary carbon market provides a promising source of financing for such interventions. With the increasing demand for carbon offsets, carbon credit projects are being implemented on a larger scale. Accurate estimation of emission reductions without the usual pitfalls of overestimation, double counting, and lack of reproducible documentation/evidence is crucial for the success of these projects. Emission reduction calculations and methodologies vary depending on the type of intervention. This initiative focuses on tree-based carbon sequestration and emission reduction and lays out a tech stack that can provide estimates for carbon sequestration with uncertainty quantification along with geo-located photographic documentation.

2. Carbon Credit and Quantification of Carbon Sequestrations

Carbon credits emerged from debt-for-nature swaps in the 1980s, eventually evolving into a market-based tool for reducing greenhouse gas emissions. The Kyoto Protocol required developed nations to cap emissions and allowed them to trade Assigned Amount Units (AAUs). This led to the creation of a global carbon market where businesses and countries could purchase credits to offset emissions by investing in renewable energy, forest conservation, and reforestation projects. One carbon credit represents the reduction or removal of one metric ton of carbon dioxide or its equivalent, with reductions facilitated through mechanisms such as Joint Implementation (JI), the Clean Development Mechanism (CDM), and International Emissions Trading (IET).

Agricultural and forestry projects play a vital role in carbon sequestration by capturing and storing atmospheric CO_2 in trees and soil. Measuring carbon sequestration in these projects involves various methodologies including biomass sampling and soil carbon analysis. Assessing changes in tree biomass is essential for quantifying sequestration potential. Tree biomass is estimated using physical parameters such as diameter, height, and wood density. Numerous allometric equations have been developed to estimate tree biomass based on these parameters. These equations are based on either destructive sampling (tree harvesting) or non-invasive techniques such as LiDAR.

2.1. Need for Digital Measurement, Reporting and Verification (MRV)

As demand for carbon offsets grows, carbon credit projects are being implemented at larger scales. This has led to practical challenges in the monitoring, validation and verification of the projects. Having multiple plantation locations, multiple stakeholders along with multiple potential buyers, it is becoming increasingly tough to accurately measure, report and verify the claimed carbon sequestration credits. In most of the carbon projects the data for monitoring is collected by a manual plot survey. So, for large projects it is difficult to collect all of the plot level as well as tree level data.. Some small percentages of the representative samples are selected and actual data for the number of trees and tree parameters is collected. The number of trees and total carbon stock in a project is typically estimated by extrapolating data collected from sample plots. However, this method presents two major challenges: the difficulty of on-ground data collection and the lack of verifiable ground truth, which undermines trust and transparency in carbon markets.

This creates a pressing need for tools that can establish credible ground truth, enable remote monitoring of the plantation plots, track growth over time, and support more accurate carbon stock estimation.

Recent advancements in remote sensing and machine learning offer promising solutions for reliable and scalable Monitoring, Reporting, and Verification (MRV) of carbon sequestration.

Digital MRV, in particular, shows strong potential to become a cornerstone of voluntary carbon markets - especially for nature-based solutions (NBS) such as reforestation, afforestation, and sustainable land management. Digital MRV systems leverage advanced technologies like remote sensing, drone imaging, satellite imagery, and AI-powered analytics to provide precise, near-real-time insights on carbon sequestration and emissions reductions.

By improving transparency and accountability, digital MRV systems help ensure that carbon credits reflect genuine and verifiable climate benefits. In addition, they reduce the costs and operational complexity of data collection and reporting, allowing e NBS projects to scale more efficiently and participate more widely in the voluntary carbon markets.

2.2. Options available for Digital MRV

Accurate carbon sequestration estimation by trees is dependent on the tree biomass. Tree biomass estimation requires data on key tree dimensions like diameter at breast height (DBH), tree crown size and tree height along with some knowledge about non-dimensional properties like the wood density. The different data sources for obtaining tree dimensions include:

1. Optical infrared satellite
2. Satellite-based LiDAR
3. Satellite-based photogrammetry
4. UAV-based LiDAR
5. Terrestrial LiDAR
6. UAV-based photogrammetry

Satellite based data is suitable for contiguous patches of forest and large trees but trees of small and medium size don't show up effectively even in the high resolution satellite data. Thus in the early years of plantation when the trees are small, their monitoring for documentation and survival rates is not possible using the satellite data. Further, even when the trees are large enough to be eligible for a carbon credit assessment they are still too small for the satellite data to accurately extract their dimensions. Additionally, satellite data is available in large land swathes (tiles). With scattered plantations on small land holdings (1-2 acres) a very small percentage of the data in the large tile is of interest and this makes it cost prohibitive even when the trees are large enough to be seen in the satellite imagery. Owing to these deficiencies, for NBS solutions like small holder agroforestry projects, drone (UAV) based data collection is an effective alternative.

Zhou et al. (2022) used UAV-LiDAR for estimating the 3D green volume and above-ground biomass of urban forest trees, achieving the highest accuracy with the voxel coupling convex hull by slices algorithm with percentage root mean square error (RMSE) on biomass of 15.12%. Jayathunga, Owari, & Tsyuki (2018) used fixed-wing UAV photogrammetry with LiDAR derived Digital Terrain Model (DTM) to estimate tree volume and biomass, finding percentage RMSE values of 17.4% for photogrammetry and 16.7% for LiDAR. Mayamanikandan, Reddy, & Jha (2019)used terrestrial LiDAR data for volume estimation of teak trees in Central Indian Forests, achieving high accuracy (volume R^2 = 0.958) with an average volume bias of 5.13% compared to manual measurements.

Literature indicates (Mayamanikandan, Reddy, & Jha, 2019) that terrestrial LiDAR data has a lowest error (5.13%) in wood volume and biomass measurement but is costly and

impractical for large-scale implementation. Aerial LiDAR and photogrammetry show comparable accuracy, with photogrammetry being more cost-effective and easier to implement. Therefore, in our digital-MRV approach, aerial photogrammetry is the method of choice for remote monitoring, tracking, and estimation of tree carbon sequestration.

2.3. Estimation of Tree Carbon Sequestration Using Aerial Photogrammetry

According to the Verra module VMD0001 (Verified Carbon Standard, 2022), tree carbon sequestration is a function of tree biomass. Carbon sequestration can be calculated from tree biomass, carbon fraction, and the molecular weight ratio of CO2 to carbon. The equation is as follows:

*Carbon sequestration (CO2) = Tree Biomass * Carbon Fraction * (44/12)*.

Tree biomass is calculated using allometric and volumetric equations from forestry literature.
Brahma et al. (2021) reviewed forest biomass estimation equations in India, evaluating species-specific and multi-species equations for trees, bamboos, palms, and bananas. They identified 85 species-specific and six multi-species equations, with half based on a power-law function using diameter at breast height (DBH). (Sileshi, 2014) found that tree height (H) and tree wood density (ρ) are essential co-variables with DBH for accurate biomass estimation. Therefore, measuring tree diameter at breast height (DBH), tree height, and wood density is critical for accurate tree biomass estimation. Here, we are not mentioning which type of equation we want to proceed with.

In photogrammetry technology, aerial data collected from the drone can be converted into the form of the orthorectified image named as orthomosaic. When flying over an area where the data needs to be collected, the drone captures a mosaic of individual images that are then stitched together to form the orthomosaic. Additionally, using SfM (structure from motion) and MVS (multi-view stereo) techniques, a point cloud can be generated from the image mosaic that can then be interpolated and projected to generate a Digital Elevation Model (DEM) corresponding to the orthomosaic. While an orthomosaic is a three channel RGB (Red-Green-Blue) image, the DEM is a single layer data that stores elevation values per pixel. Both the orthomosaics and the DEM are stored as geo-referenced GeoTiff files. This geo-referencing allows for extraction of geo-location data when the GeoTiffs are further processed. To make use of the data (orthomosaic and the DEM) generated from drone based photogrammetry, it is essential to extract key tree dimensions that can be then fed

into a suitable allometric equation for biomass calculation. . The first step in this process is the detection of individual trees in the orthomosaic image. For every detected tree, its crown size can be calculated by knowing the number of image pixels it occupies in the NS and EW direction along with the ground sampling distance (GSD). GSD is the ground distance each image pixel represents. The region of detection for every tree in the orthomosaic can be mapped to a corresponding location in the DEM to extract the height of each tree. A regression model trained on field measurements can then be used to predict the DBH for every tree as an output with tree height and its crown size as inputs. Once the DBH, crown size and tree height are known, an allometric equation of choice (either one that uses only DBH or one that uses both DBH and tree height) can be used to estimate the biomass and carbon sequestered in every tree. Summing up the data over all trees generates a plot level biomass estimate. With the geo-referenced GeoTiff, the geolocations of every tree in the plot can also be reported.

3. Tree Crown Detection

In literature, there are two major approaches for tree detection:

1. Rule-based detection: Tree segmentation, template matching
2. Detection using deep learning methods

There is plenty of literature available on both the approaches. Ke & Quackenbush (2011) evaluate three tree crown detection and delineation algorithms—watershed segmentation, region growing, and valley-following—using high-resolution aerial and satellite imagery. The region growing method showed the highest accuracy for softwood stands, with a 70% accuracy rate and a 15% root mean square error for crown diameter estimation. Zheng et al. (2020)developed a Multi-level Attention Domain Adaptation Network (MADAN) for oil palm tree detection, achieving an 84.81% F1-score. Yao et al. (2021)developed four deep convolutional neural networks for tree detection, with the Encoder-Decoder Network achieving an average accuracy greater than 91.58%.

Neupane, Horanont, & Nguyen (2019)used deep learning for banana plant detection and counting, achieving up to 96% accuracy. Skole et al. (2025)developed a deep learning model for monitoring trees outside forests in South Asia using high-resolution satellite images. Weinstein et al. (2020, 2020)introduced a cost-effective machine vision system for tree identification and mapping using UAV-captured RGB images and convolutional neural networks (CNN).

Literature indicates Khan & Gupta (2018) that deep learning models are preferred for tree crown detection due to their superior accuracy and ability to handle complex datasets. Unlike rule-based models that rely on predefined rules and manual feature extraction, deep learning models can automatically learn and extract relevant features from the raw data, making them more adaptable to diverse and challenging environments. Additionally, deep learning models are highly scalable, allowing them to effectively manage large datasets and provide more reliable detection results.

Hence, we prefer an object detection deep learning approach for tree detection.

4. Tree Detection Model Training

In an automated flow of biomass sequestration calculation on agroforestry plots, the first step is to gather the drone orthomosaic of the plot (RGB data) along with a digital elevation model or DEM. Mavic M series of DJI made drones is used for this purpose. This drone uses image overlap to generate elevation maps using structure from motion technique. The ground sampling distance of the orthomosaic and of the generated DEM is between 2-3 cm when the flight height is set to approximately ABCD meters.. DEM is a single channel data with one elevation value per pixel of the DEM. A sample orthomosaic along with corresponding DEM is shown in Figure 1.

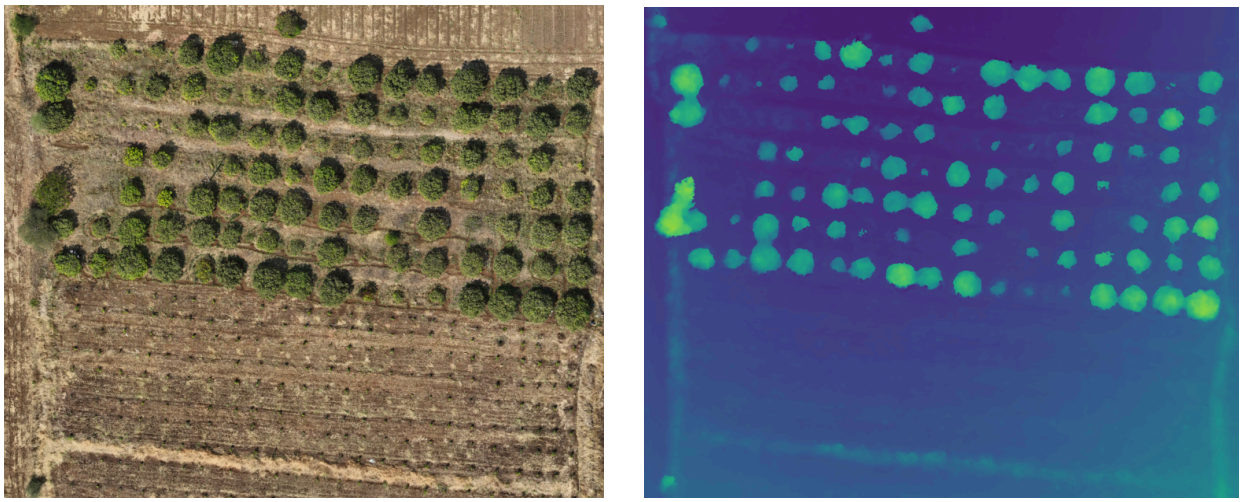

Figure 1: Sample drone orthomosaic and DEM data for an agroforestry plot

Although in the initial years of plantation of an agroforest, the trees are not large enough for carbon credit issuance, the detection of very small sized saplings is still important for various reasons.

1. It allows for documentation of plantations which leads to robustness of MRV procedures required in the carbon credit market
2. Via detection of tree saplings, survival rates of samplings can be determined and documented leading to replantation decisions and farmer feedback
3. Detection of trees on georeferenced tif images allows for extraction of GPS locations of trees; also important for MRV documentation

For these reasons, an automated, machine learning based, tree detection system needs to be established for detecting both small and large trees starting with a drone orthomosaic as the input.

4.1. Data Annotation

There are different types of annotations used in the field of computer vision: bounding boxes, Polygonal Segmentation, Semantic Segmentation, 3D cuboids, Key-Point, Landmark, Lines and Splines. Bounding boxes are easy to annotate, widely used, and most suitable for our purpose and so they are a natural choice for annotating trees in our images.

For a computationally optimal model training which is both memory efficient and requires a practically small annotation cost, training the model on fully annotated orthomosaics is not an ideal solution. The orthomastics are large in size (of the order of 10^6 to 10^8 pixels with sizes ranging from to 1 GB) and hence makes them a poor choice to be used in ML model training. Further, each orthomosaic can contain anywhere from 300 - 800 objects (trees) that would need to be fully annotated before training. Doing this much annotation on many orthomosaics is a very costly proposition in terms of the time required. Ideally, annotating only a small number of trees from an orthomosaic should suffice since each tree in the orthomosaic and the background (ground) is mostly similar and this repeated information in an orthomosaic means that a smaller subset of information should be able to represent almost all of the information contained in the full orthomosaic. For small trees, even with a high resolution of the drone orthomosaic, the image covers a fairly large swathe of the ground (100-200 meters in one direction) and at this scale the small size of saplings make them difficult to spot and annotate.

To overcome these challenges, the training-test-validation sets are prepared by cropping small tiles from the entire image (500 px by 500 px to 800 px by 800 px) which capture the

crux of how the trees/saplings look on that particular land, the soil type, the intercropping pattern and other geophysical parameters. This approach drastically reduces the burden of annotation (by having to annotate only a small subset of trees from a plot) and provides high visibility of samplings on a small tile leading to accurate complete annotations without missing sapling objects. Crops are very small in size and can easily be pushed on to a GPU and hence allows parallel processing and significant computational gains during model training as well as during model inference.

Some examples of annotation on sub-tiles of full orthomosaics are shown in Table 1.

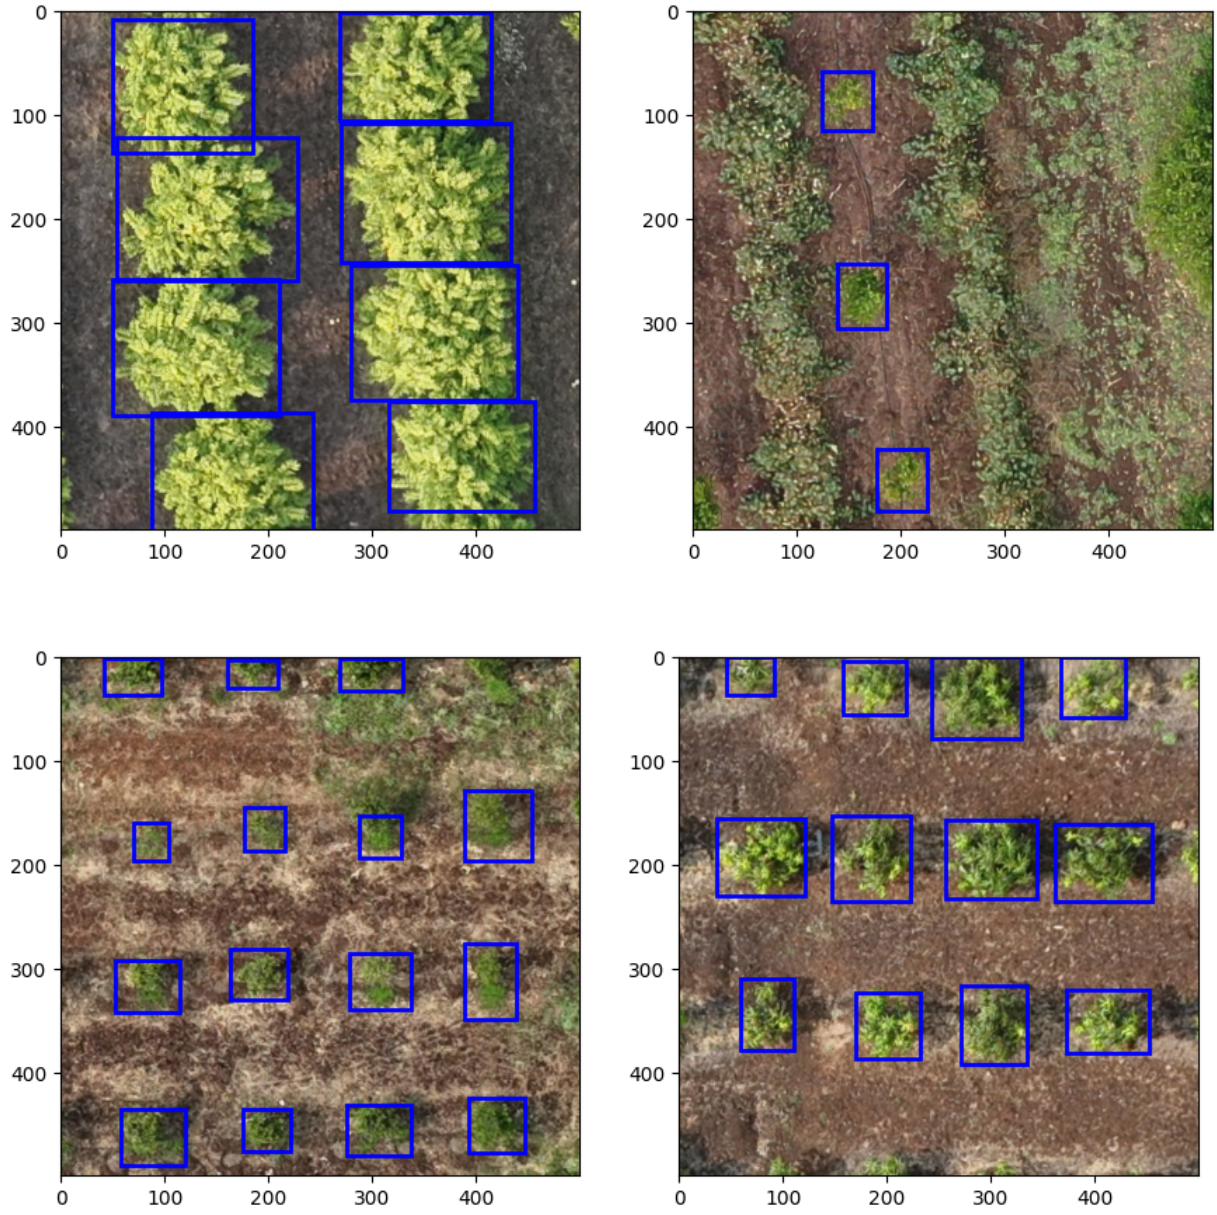

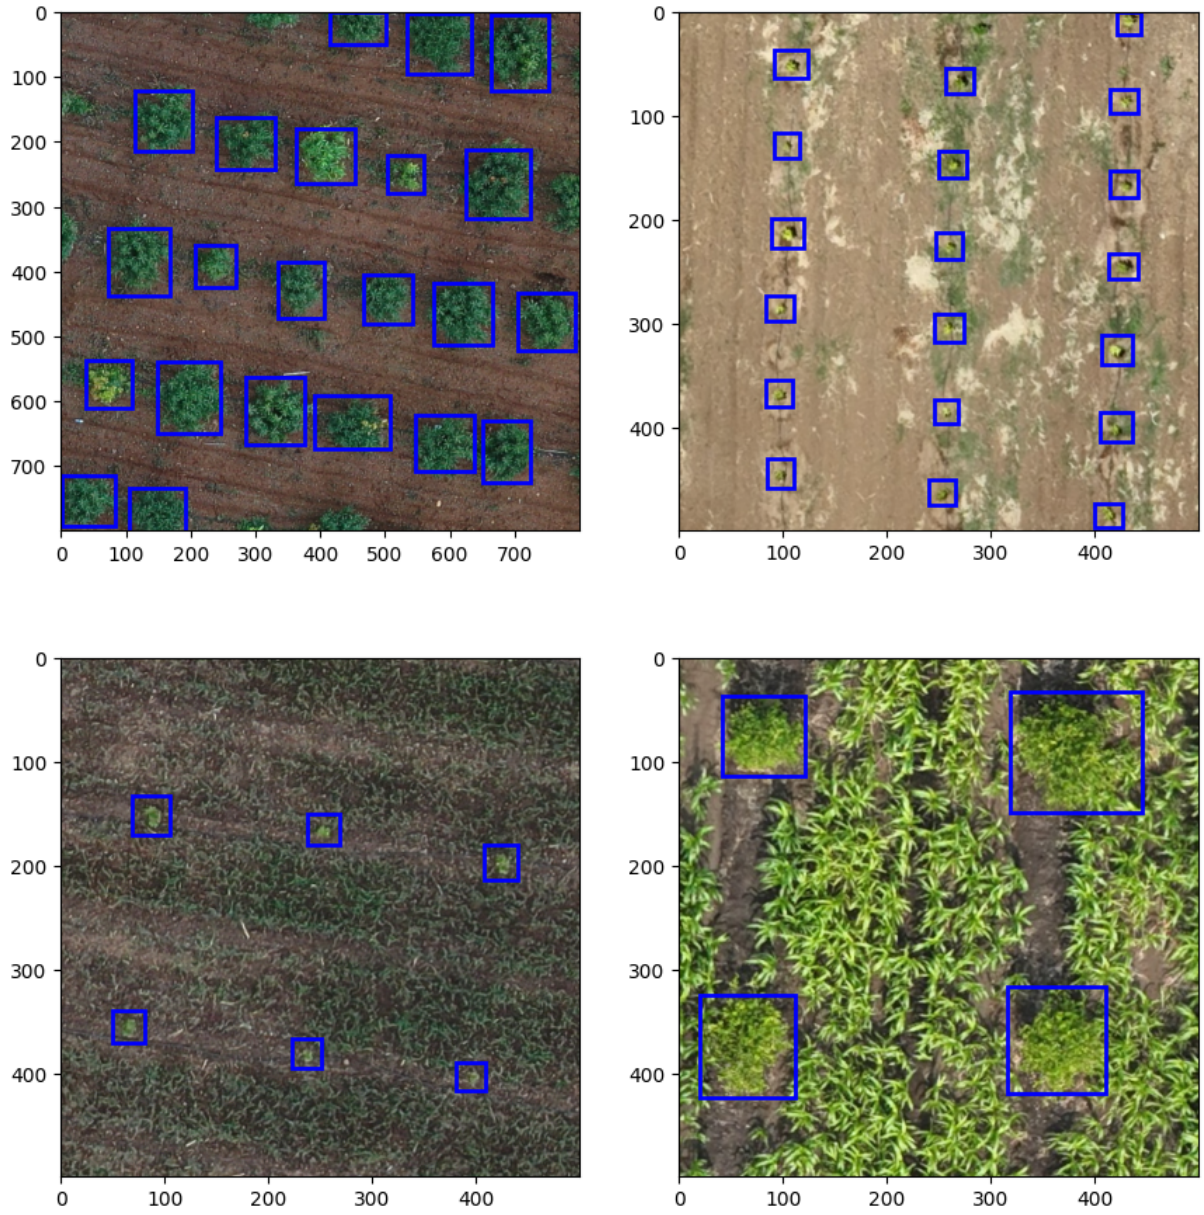

Table 1: Samples of image tiles with annotations

Out of all the annotated crops, 80 percent are used for the training and 10 percent each are used in the validation and test set. Train, validation and test sets are mutually exclusive with approximately 1830 images used in the train set and 230 images each in the validation and the test set.

4.2. Training process

For the tree detection model training, we used Detectron-2 as a base platform. Detectron2 is a PyTorch-based deep learning framework developed by Facebook AI Research (FAIR) for object detection and segmentation tasks. Various object detection and segmentation models provided by Detectron have been pre-trained on real life scenes and are capable of detecting standard real world objects off-the-shelf. However, the standard model is unable to detect tree crowns in the drone imagery as it has not previously been trained on such objects. As a result, after fetching one of the standard models from the model library, it is further trained using the annotated train set.

We use faster_rcnn_R_50_FPN_3x from the standard Detectron-2 library. This is a Faster R-CNN object detection model with a ResNet-50 backbone and Feature Pyramid Network (FPN), trained on the COCO dataset using a 3x training schedule (36 epochs).

Annotations are in coco-json format and a single annotation file containing annotations for all the images is provided to the faster_rcnn_R_50_FPN_3x model for further training. We use Google's collaborative platform for model training.

4.3. Model Selection

During the training process, versions of the model are saved after every few epochs and training loss and validation losses are recorded. As seen in Figure 2,the training loss continues to decrease throughout the period of model training while the loss function on the validation set initially falls and then starts to go up indicating that the model is starting to overfit to the training dataset. A model corresponding to the number of epochs for the lowest point on the validation curve is selected and with this model, the test is evaluated to understand the model performance.

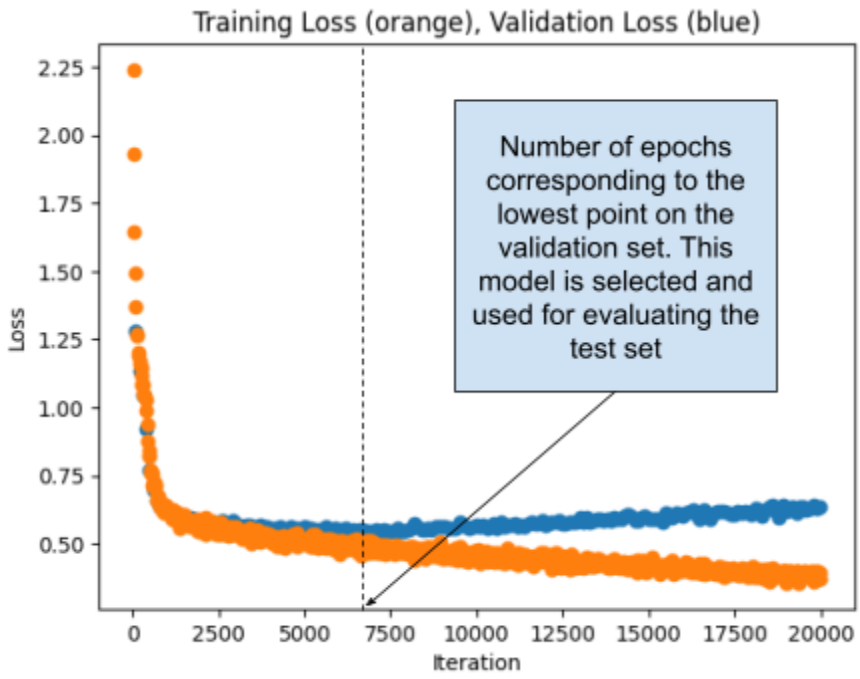

Figure 2: Model selection based on train-Validation loss

4.4. Model Evaluation

The model is evaluated on a set of images previously unseen by it both in the training and validation set. The model has a precision of 0.89 and a recall of 0.89. Some sample model predictions and annotations are shown in Table 2.

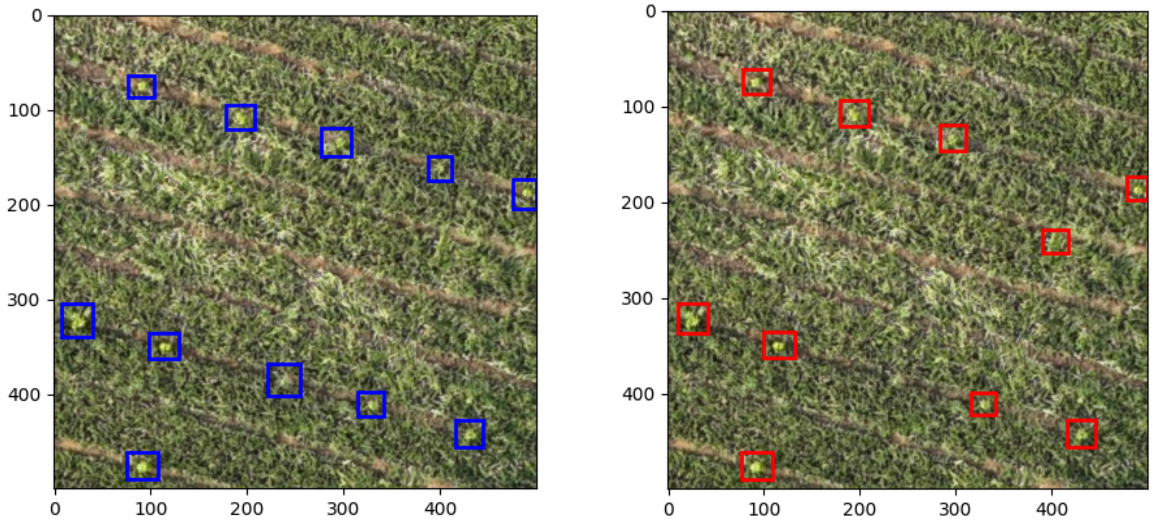

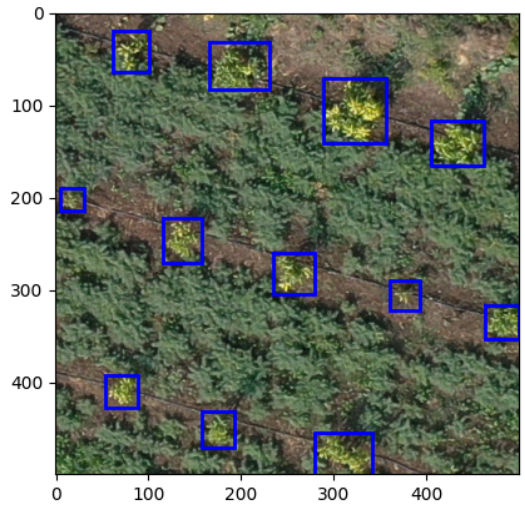
0
100
200
300
400
0
100
200
300
400

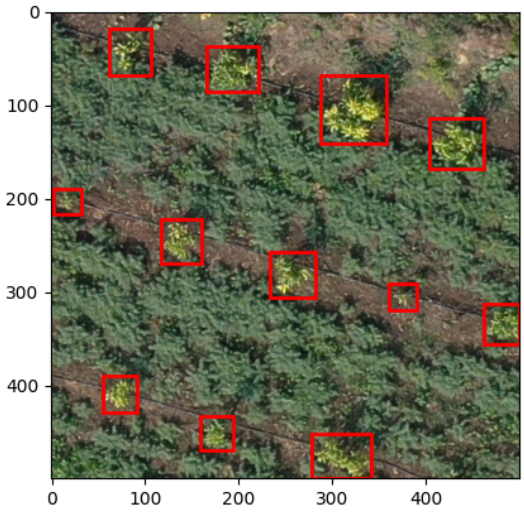
0
100
200
300
400
0
100
200
300
400

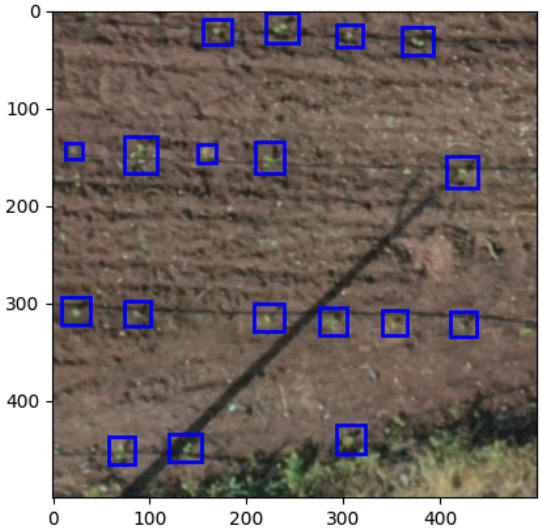
0
100
200
300
400
0
100
200
300
400

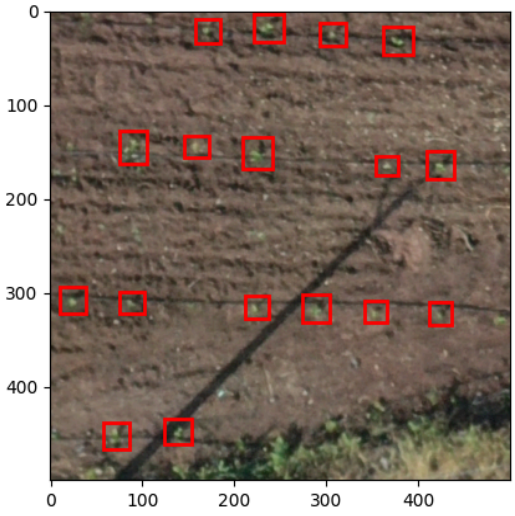
0
100
200
300
400
0
100
200
300
400

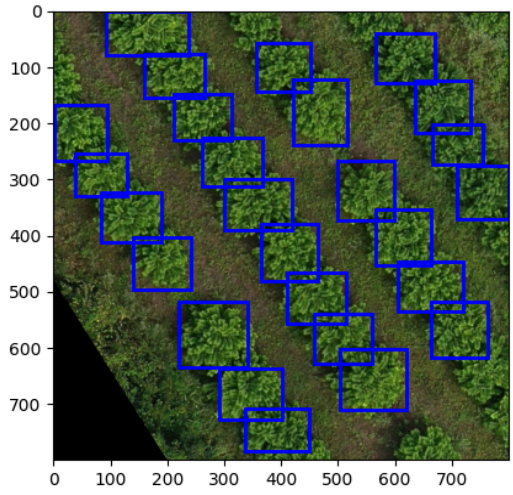
0
100
200
300
400
500
600
700
0
100
200
300
400
500
600
700

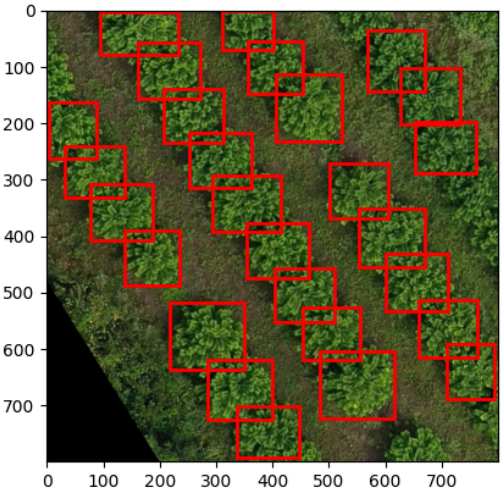
0
100
200
300
400
500
600
700
0
100
200
300
400
500
600
700

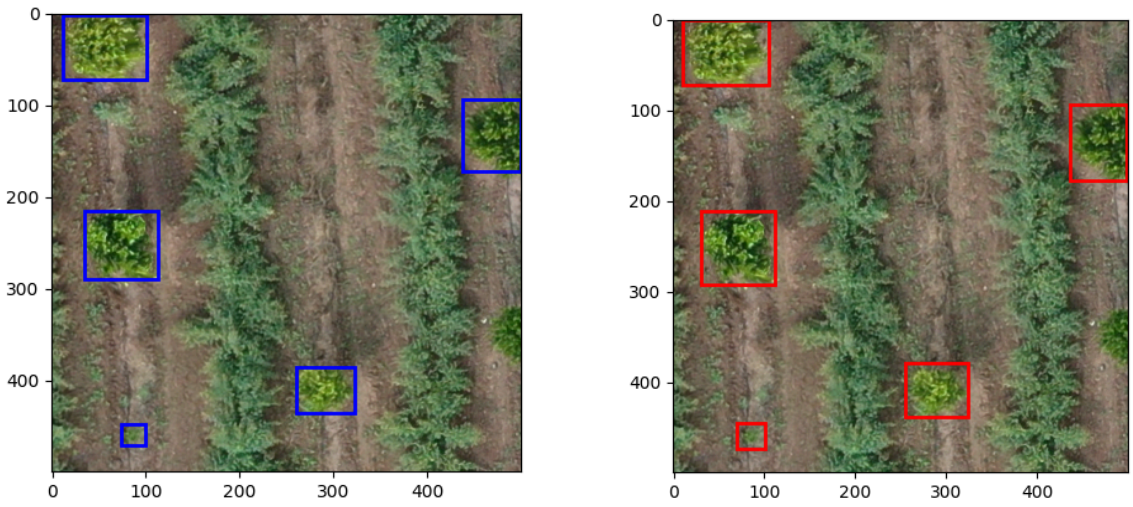

Table 2: Sample annotated and predicted images (Left images with blue boxes are annotated images and images on right with red boxes are predicted images)

5. Estimation of the tree trunk diameter (DBH) with uncertainty quantification: Gaussian Process Regression

A typical empirical method for estimating biomass involves predicting the AGB (above ground biomass) of the tree using its DBH (diameter at breast height) or DBH in combination with the tree height. In this work, we will consider the biomass estimation using only tree DBH. In order to perform biomass estimation on an agroforestry plot, an orthomosaic of the plantation is acquired using drone imagery. Along with this RGB image of the plot, an elevation model or Digital Elevation Map (DEM) is also obtained. This DEM gives the height of each pixel in the image. Once a bounding box containing each tree is obtained using a tree detection algorithm, the crown size (referred to henceforth as TC: Tree Crown) of the tree can be easily calculated based on the size of the bounding box in pixels and the ground sampling distance (GSD). Mapping this bounding box on to DEM, tree height (TH) can also be extracted. However, the DBH of the tree can not be directly obtained from the drone imagery as this feature is hidden underneath the tree canopy and can not be directly observed from an overhead drone image.

This necessitates development of a regression model that can predict the DBH of the tree as a function of known features: Tree Crown and Tree Height. Along with this, due to this additional uncertainty introduced by the regression model (direct measurement of DBH vs

predicting it using TC and TH), some methodology needs to robustly capture the uncertainty in this DBH prediction and propagate it at the biomass level in order to allow uncertainty quantification about the biomass estimates at the plot level.

Here, instead of fixed estimate regression models like various types of linear, polynomial regressions and neural networks, a Gaussian Process Regression is a natural model of choice as it is able to not only provide inference about the expected value of a query point but also the expected uncertainty in the inference.

5.1. Data collection for DBH model training

On field measurements were conducted for various agroforestry species that are involved in the operations. For each measured datapoint, TC, TH and DBH value was obtained. A total of 3500 such datapoints are collected.

5.2. Assumptions

In this work, both the estimation of biomass from DBH using the allometric equations as well as fitting of gaussian process regressor to predict DBH from TH and TC are considered species agnostic. This simplification, though has its limitations, allows for a robust practical implementation and automated cascading of data pipeline from drone imagery to biomass estimation without needing to identify the tree species at the time of tree detection. The uncertainties stemming from this assumption are embedded in and quantified by the gaussian process regressor as presented in the future sections.

5.3. Model Training

Based on the collected data with three variables recorded for each tree: Tree Crown, Tree Height and DBH, the natural choice for gaussian process model training is to learn DBH as a function of Tree Crown and Tree Height. However, in this work, it is more convenient to instead learn the natural logarithm of DBH as a function of TC and TH. There are two main reasons behind this. The AGB is defined as:

$$AGB = exp(-1.996 + 2.32 * ln(DBH))$$

If DBH is learnt via a gaussian process, at inference time, on a query datapoint the value of DBH sampled from the learnt mean and covariance matrix can become negative. Although

this happens less often, it can occur in certain cases and can as such be expected to occur when sampling from a Gaussian distribution with a positive mean value that is closer to zero and a higher standard deviation. As DBH can not be a negative value, it is instead more convenient and robust to learn the term $log_e(DBH)$ which can freely take negative values when sampled from a gaussian distribution.

Another reason for choosing log DBH as the target variable to be learnt via gaussian process is the tendency of DBH to have a skewed distribution as observed by Bliss & Reinker (1964)and Lima et al. (2015). Authors describe the DBH to have a lognormal distribution. A target variable with skewed distribution can potentially pose trouble for the gaussian process regression. Hence converting the lognormally distributed DBH to normally distributed log(DBH) makes the problem better conditioned for the gaussian process regressor.

The DBH data gathered for this work shows a similar trend as to one quoted in the papers above. Log(DBH) is normally distributed and DBH shows a skewed lognormal distribution. The left side plot on Figure 3 shows a highly skewed distribution of DBH data while the right hand side plot shows a more normal distribution of the transformed quantity log(DBH).

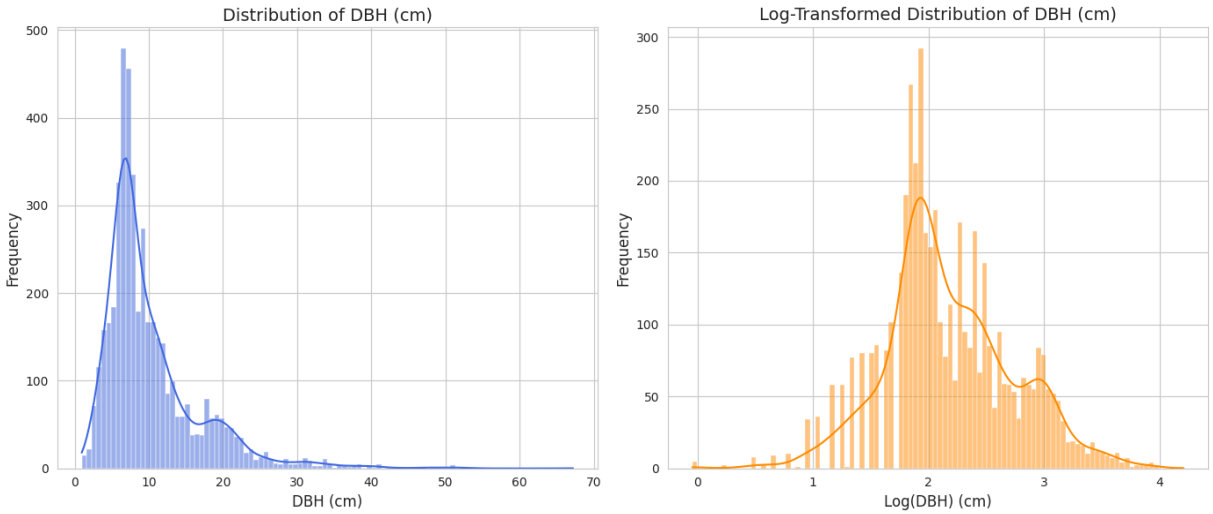

Figure 3: Distribution of DBH and log(DBH)

The nature of the DBH data in relation to the available features to define it (TC and TH) suggests that there would be a fair amount of aleatoric uncertainty that would need to be captured by the gaussian regressor. In contrast to the epistemic uncertainty that can potentially be reduced by gathering additional data points where there exist gaps in the current training data, the aleatoric uncertainty can not be reduced by fetching more

training data while remaining within the current feature space involving only TC and TH. The aleatoric nature of uncertainty can also be seen in the vertical stacks of datapoints when the DBH is visualized against each feature as shown in Figure 4.

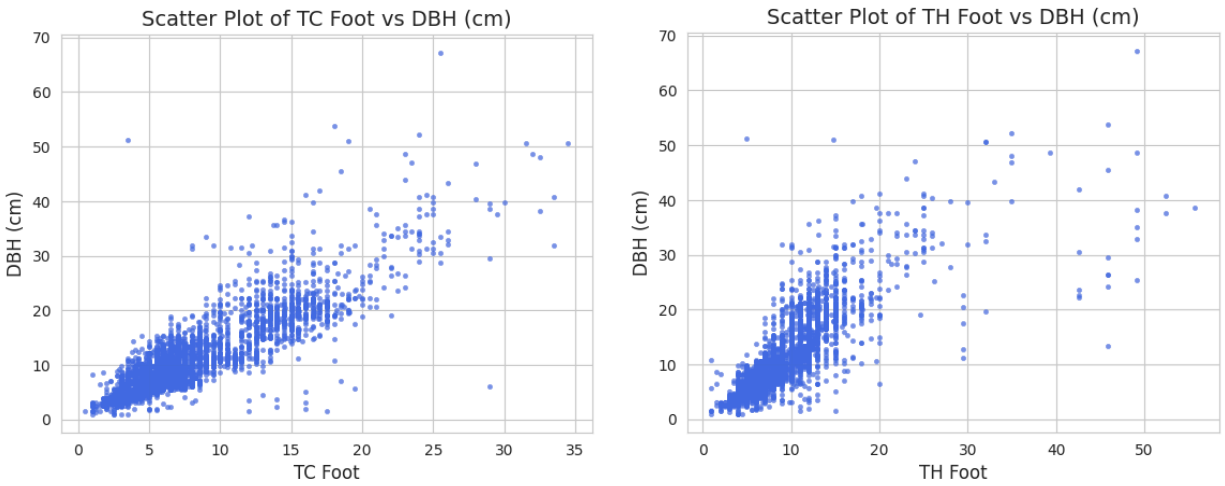

Figure 4: DBH plotted vs TH and TC

Datapoints stacking vertically indicate the presence of several distinct observations for the same input feature values. This is further confirmed when looking at the total feature space including both TC and TH. For example, there are 110 data points where TC = 5.5 feet and TH = 6 feet while the DBH_cm value varies for each datapoint. 10 datapoints out of these 110 are shown in Table for clarification. Many such instances exist in the dataset where for a given combination of input features (TC, TH) multiple distinct outputs exist (DBH_cm).

Number	TC_foot	TH_foot	DBH_cm
1	5.5	6	6.366198
2	5.5	6	9.55414
3	5.5	6	8.917197
4	5.5	6	10.191083
5	5.5	6	7.324841
6	5.5	6	7.006369
7	5.5	6	7.324841
8	5.5	6	9.55414
9	5.5	6	6.687898
10	5.5	6	6.369427

Table 3: Aleatoric uncertainty in DBH data

This observation further confirms the presence of aleatoric uncertainty in the data within the available feature space that can not be removed or reduced by acquisition of additional training data. In other words, the data indicates features beyond TC and TH (but unavailable and or unknown) that influence the DBH value of trees and need to be dealt with through the uncertainty quantification mechanism offered by the gaussian processes.

Entire data is divided into three disjoint sets namely: Train, Validation and Test sets. After randomly shuffling the datapoints, 70% of the gathered data is picked for training and 15% of the data goes into the Validation and Test sets respectively. An Radial Basis Function (RBF) kernel of the form shown below is used

$$K(x_i, x_j) = theta * exp(-||x_i - x_j||^2/L)$$

Where, theta represents a scalar multiplier and L represents the length scale of the function. Both these are hyperparameters and are learnt by the gaussian process regressor during run time. Epistemic uncertainty at unknown query points previously unseen during the training process can be captured by a variance matrix of the trained gaussian process that depends only on the covariances between the train-train, test-train and test-test sets via the kernel function. However, in order to capture the aleatoric uncertainty, an additional variance term $\sigma^2_{\text{aleatoric}}$ is introduced in the variance matrix as described by Eyke Hüllermeier and Willem Waegeman (Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods). This $\sigma^2_{\text{aleatoric}}$ is an additional hyperparameter that is specified during the initialization of the gaussian process regressor and learnt directly from the Y_train under the assumption of a homoscedastic gaussian process.

5.4. Model selection using validation set

From the family of trained models, the best model is selected by looking at the the performance of each model on the validation set on two factors. Mean Absolute Percentage Error and slope of coverage. Coverage of a gaussian process regression is plotted as described below.

For every trained model, inference is made on all points in the validation set (mu and standard deviation) using the trained gaussian process regressor. Then using the standard deviation for every prediction, it is checked whether the true value lies within the confidence interval given by the predicted standard deviation applied around the predicted

mean or expected value. The percentage of total points lying within the given interval is noted and this process is repeated for a range of confidence intervals. For example, with a well calibrated gaussian process regressor, it is expected that 60 percent of the true values from validation set lie within the 60% confidence interval around expected mean prediction of each data point. When plotted over a range of confidence intervals, the coverage plot is expected to have a slope of close to 1 if the model is well calibrated. An under confident model has a slope less than one and an overconfident model has a slope greater than one.

The best model among the trained models, shows a slope of 1.08 and a mean absolute percentage error of 8.59. This suggests a well trained and well calibrated model that can now be used for inference on the test set. The coverage plot for the model is shown in Figure 5.

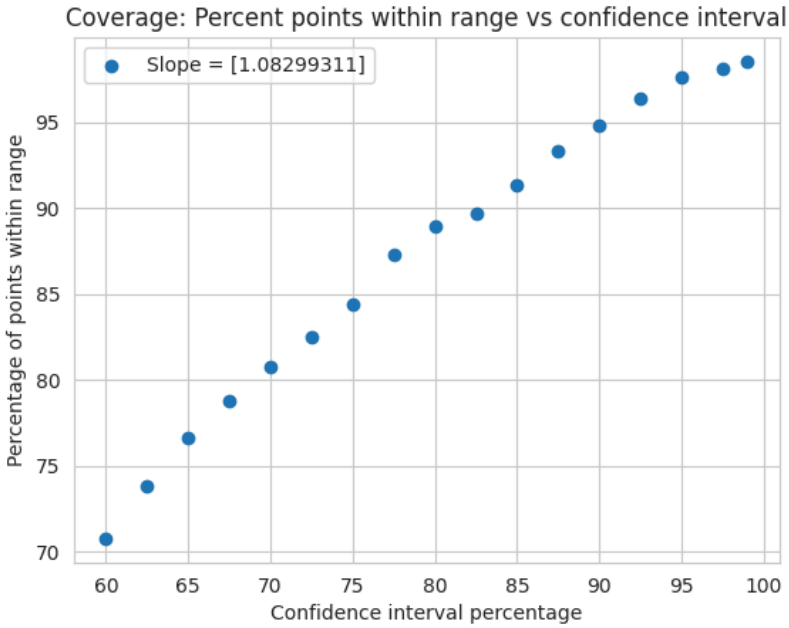

Figure 5: Coverage plot for a well calibrated Gaussian Regressor

5.5. Inference on test set

After fitting, the mean and variance from the gaussian process regressor take the following form.

$$\mu = K(x_{\text{test}}, X_{\text{train}}) \left(K(X_{\text{train}}, X_{\text{train}}) + \sigma^2_{\text{aleatoric}} I_{\text{ntrain}} \right)^{-1} y_{\text{train}}$$

$$\sigma^2 = K(x_{\text{test}}, x_{\text{test}}) + \sigma^2_{\text{aleatoric}} I_{\text{ntest}} - K(x_{\text{test}}, X_{\text{train}}) \left(K(X_{\text{train}}, X_{\text{train}}) + \sigma^2_{\text{aleatoric}} I_{\text{ntrain}} \right)^{-1} K(X_{\text{train}}, x_{\text{test}})$$

x_{test} is the test set with size n_{test} by $n_{features}$. n_{test} is the number of datapoints in the test set and we have $n_{features}$ = 2 since we have 2 features (TH and TC). X_{train} is of size n_{train} by $n_{features}$. y_{train} is of size n_{train} by 1. $\sigma_{aleatoric}$ is a scalar. $K(x_{test}, X_{train})$ is a matrix of dimensions n_{test} by n_train. $K(X_{train}, X_{train})$ is a matrix of size n_{train} by n_{train}. I_{ntrain} is an identity matrix of size n_{train}. μ is a vector of size n_{test} by 1.

$K(x_{test}, x_{test})$ is a matrix of dimension n_{test} by n_{test}. I_{ntest} is an identity matrix of size n_{test} by n_{test}. $K(x_{test}, X_{train})$ is a matrix of size n_{test} by n_{train}. $K(X_{train}, x_{test})$ is a matrix of size n_{train} by n_{test}. σ^2 is a covariance matrix of size n_{test} by n_{test}.

The trained Gaussian Process regressor is evaluated for its performance using the test set that is mutually exclusive from both the train and validation set. The mean absolute error on the test set is 6.16. The plot of true log(DBH) and the predicted log(DBH) by the model is shown in Figure 6.

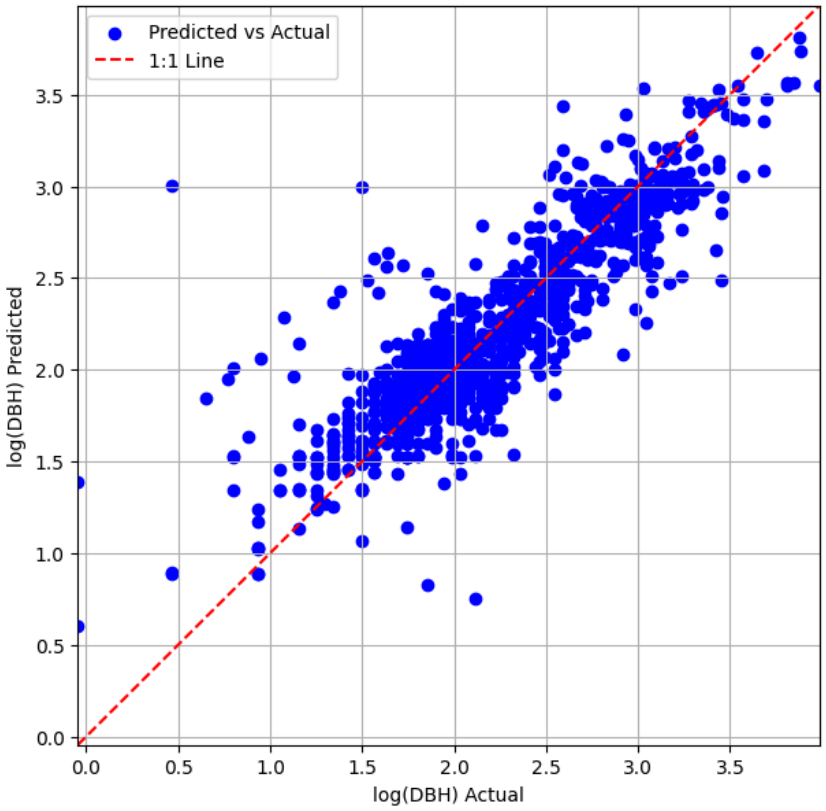

Figure 6: DBH model evaluation on the test set

For comparison, 10 sample data points for true log(DBH) and expected log(DBH) value as predicted by the gaussian processor regressor are shown in Table 4.

No.	True log(DBH)	Predicted log(DBH)	Percentage error
1	1.15836229	1.70089886	-46.83651865
2	2.43929614	2.33557051	4.252277052
3	2.2897644	2.45832627	-7.36153772
4	1.15836229	1.34381548	-16.00994711
5	2.3215131	2.08671689	10.11393001
6	2.66243969	2.34831182	11.79849711
7	2.63996683	2.5016821	5.238123768
8	2.91622021	2.96576166	-1.69882404
9	2.18798171	2.1020281	3.928442802
10	2.35228476	2.26882753	3.547922064

Table 4: True vs predicted log(DBH)

6. Tree crown Estimation and validation

As discussed above, the DBH is a function of tree crown size and tree height. For training the DBH model, the data gathered for tree crown size is the average of tree crown width in NS and EW direction. The same needs to be extracted from the object detection model for the detected trees. This is applicable only for the larger trees (3 years and older) as they are annotated by drawing boxes tangential to their crowns while smaller trees are annotated with boxes that are slightly larger due to their very small size and the detection purpose is only for survival and documentation in the case of small trees.

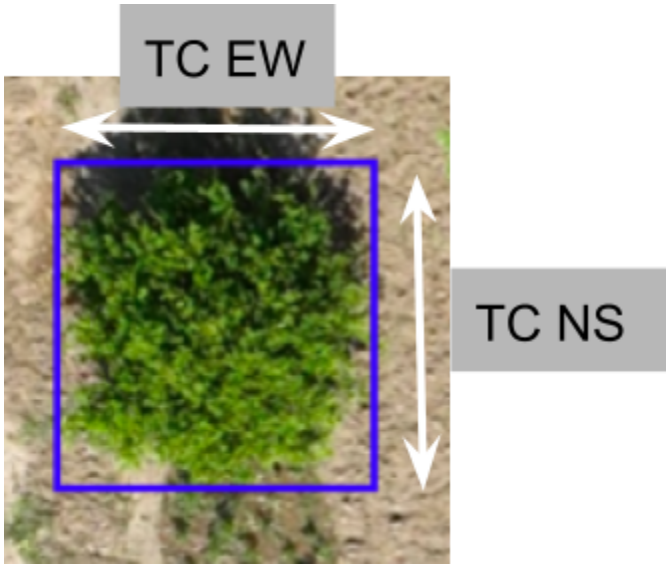

For each detected tree, the crown size in EW and NS direction is calculated as the length of the bounding box in pixels in that direction multiplied by the ground sampling distance. Ground sampling distance is the amount of length on the ground covered by each pixel of the image. Once TC EW and TC NS are calculated, they are averaged to get the size of the tree crown. This data extraction process is validated using ground data for 10 trees and is presented in Table 5.

No.	Crown size in foot from image analysis	Crown size in foot from field measurement	% error in crown size
1	10.93513229	11.7	6.537330892
2	13.35768467	12.85	-3.950853434
3	18.20278943	16	-13.76743396
4	15.44377144	15.15	-1.939085428
5	16.48681483	15	-9.912098858
6	19.01030689	17.15	-10.84727052
7	18.64019472	17.5	-6.515398422
8	15.17459895	14.5	-4.652406575
9	14.29978837	13.7	-4.378017306
10	12.75204657	11.55	-10.40732963

Table 5: Tree crown size measured on ground vs derived from drone image

7. Tree Height Estimation and validation

Aerial photogrammetry gives the output as the RGB image and Digital Elevation Model (DEM). DEM is a single channel data with a single number representing height of every pixel. Visualization of a DEM for a plot with trees is shown in Figure 1 above.

Height of each predicted tree can be calculated from the digital elevation model. First, the bounding box of a tree detected in the orthomosaic image is mapped to a corresponding location in the DEM file. From here, the highest elevation in that region (corresponding to the top of the tree crown) along with the lowest elevation in the immediate neighborhood of the tree can be used to calculate the tree height. Table 6 shows comparison of the actual measured tree height by field survey and height deduced using drone based approach.

No.	Tree height in foot from image analysis	Tree height in foot from field measurement	% error in tree height
1	17.88962915	16	-11.81018216
2	19.59312682	17	-15.25368717
3	16.59483475	15.8	-5.030599659
4	11.50116244	12.5	7.990700488
5	13.23009143	13.8	4.129772259
6	14.67206707	13.9	-5.55443932

Table 6: Tree height measured on ground vs derived from drone image

8. Estimation of Carbon Sequestration: Methodology

In the forestry literature, the equations for the tree biomass estimation are developed based on the actual measurement of the tree biomass and other input parameters(DBH, Height) and find the relation between them. We have used such equations for the estimations of the biomass. For the Carbon Credit Projects, Verra has listed some of the region-specific equations. We have used the equation from Verra Literature.

$AGB = \exp\{-1.996 + 2.32 * \ln(DBH)\}$
AGB - Above-ground biomass in Kg
DBH - Diameter at breast height in cm

BGB - Below-ground biomass in Kg

Carbon Stock in a tree
= 50% of the total biomass of tree
= 0.5 * (AGB + BGB)
= 0.5 * (AGB + Root shoot Ratio*ABG)

With Root Shoot Ratio assumed at 0.27
Carbon Stock in a tree
= 0.5 * 1.27 * AGB

Considering molecular weight ratio for Carbon,
Carbon Sequestration by tree = 44/12 * 0.5 * 1.27 * AGB

Hence, all the calculations can be cascaded from knowing the AGB and this can be calculated by applying the above mentioned formula once we have an estimate for the DBH of every tree.

9. Data Collection

9.1. Drone Data collection

- Orthomosaic and DEM have a ground sampling distance of approximately 2.5 cm and this is maintained across the datasets.
- Orthomosaic and DEM are cropped to the exact plot boundary for automatic cascading of the entire pipeline. This prevents trees outside of the plot boundaries being included in the carbon calculation.
- Drone flying speed is set to around 10 m/s and this gives a good balance of time needed to gather the data and data quality.

9.2. Data Variations

There are three main categories of trees when it comes to tree crown detection:

1. Fully Grown trees
2. Medium-sized trees of height 10-20 feet.
3. Small Trees of height less than 10 feet.

Since the model needs to work well on all these scenarios, care is taken to include imagery for trees from all these classes. Other types of variations expected in the model as shown in Figure 7.

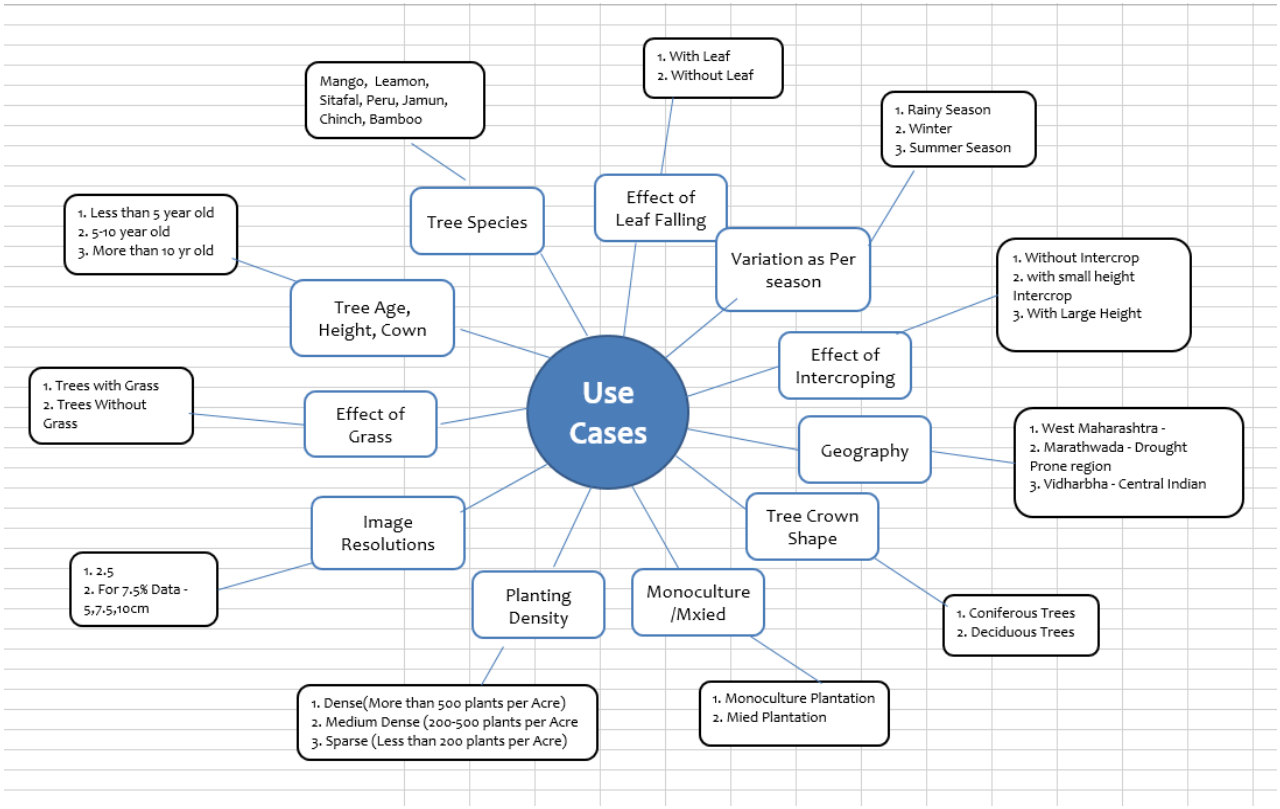

Figure 7: Different data variations in the trees

Some of these variations are much more significant than others. So, the data is collected for the variations mentioned below and used for the training of the model:

- Tree Species
- Tree Age, Height, and Crown
- Planting Density
- Monoculture/Mixed
- Variations as per Season

Other variations will be included in the future versions of the model for more coverage in terms of geography and on ground differences in how plantations/trees look.

9.3. Geographies where data is collected

1. Existing Plantation of Farmers in our Implementation area - Shrigonda, Parner, Karjat, Jamkhed block:

- Plantation of age 1-3yr - Farmers who are part of the carbon credit project.
- Plantation of age: 3-10yr - Other Farmers in F4F's implementation area (most of this will be monoculture)

2. F4F Plantation of 2022-23 - mixed plantation of 1 yr old.

10. Results

Here, we present three case studies by implementing our drone based biomass estimation technique on land parcels where agro-forestry has been implemented and compare the results with biomass estimation done using the manual approach.

With drone approach, we follow the following steps:

1. Gathering drone data: Orthomosaic and Digital Elevation Model
2. Tree detection
3. Tree detection cleanup
4. Extracting tree heights and crown sizes
5. Estimating DBH using Gaussian processes
6. Estimating the biomass distribution mean and standard deviation using Gaussing processes

For manual measurement, we selected a few plots and conducted manual measurements with following steps:

1. Measurement of tree height, tree crown, DBH on a subset of trees in the plot
2. Calculating the biomass using DBH or DBH and tree height for all the measured trees
3. Summing up the biomass of measured trees to get total measured biomass
4. Extrapolating this to the entire plot based on the total number of trees in the plot and number of trees measured

These two approaches can then be compared to understand how the drone based biomass detection approach performs.

Case Study 1


0
1000
2000
3000
4000
5000
6000
0
500
1000
1500
2000
2500
3000
3500

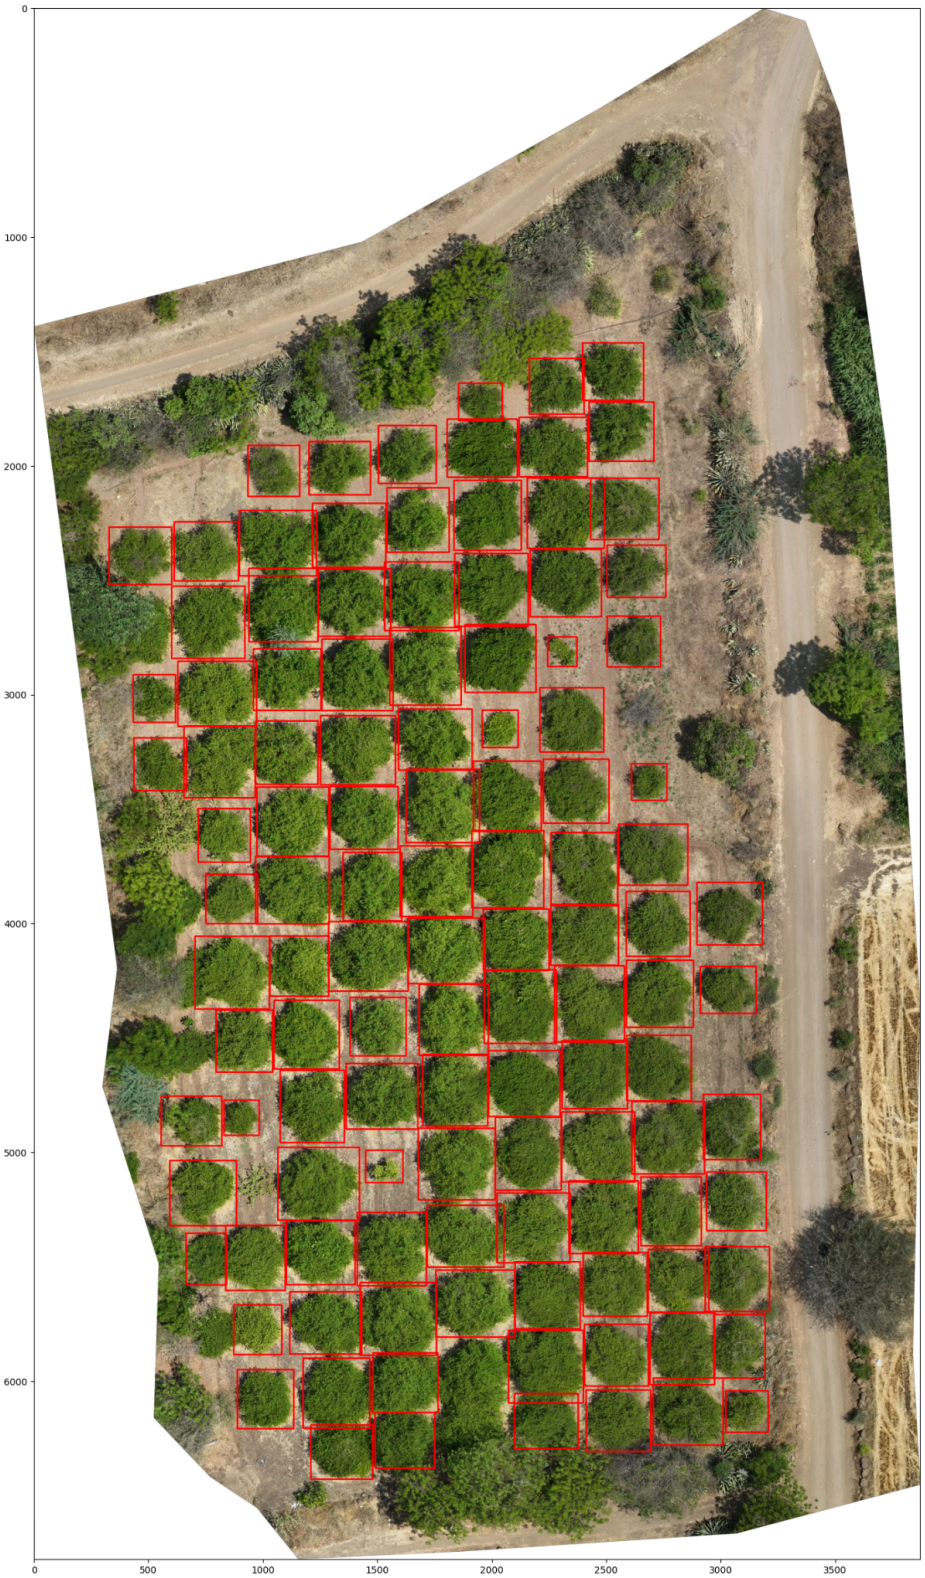
0
1000
2000
3000
4000
5000
6000
0
500
1000
1500
2000
2500
3000
3500

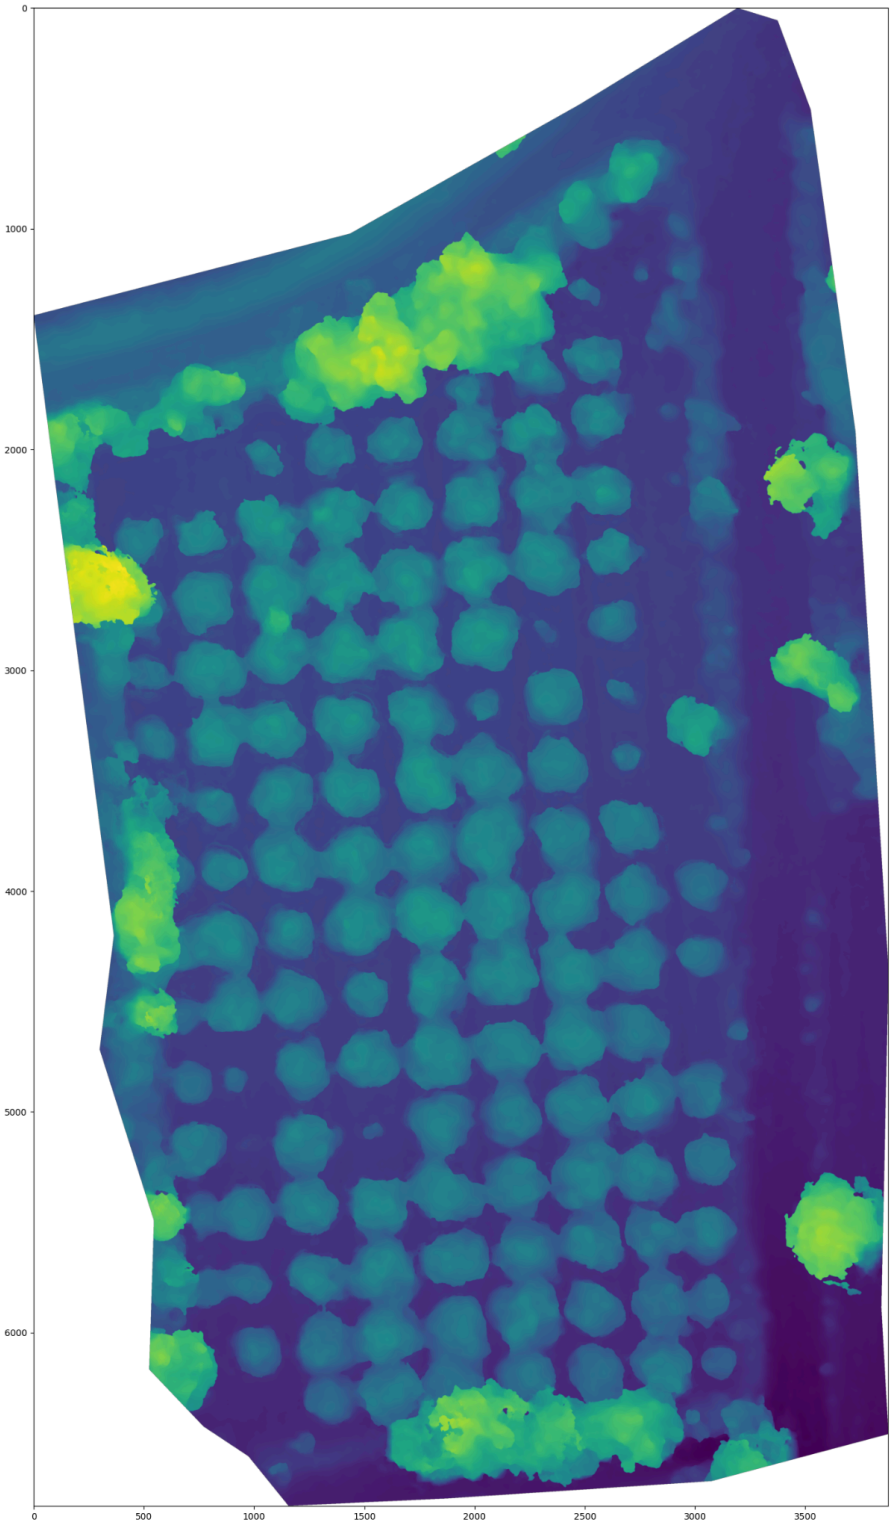

Number of trees measured on the ground = 30
Number of trees on the plot = 115

Total amount of AGB calculated and summed up for the 30 manually measured trees = 4732
Total amount of AGB extrapolated to the plot from the measurement of 30 trees = 18142.

Distribution of total AGB from Gaussian processes:

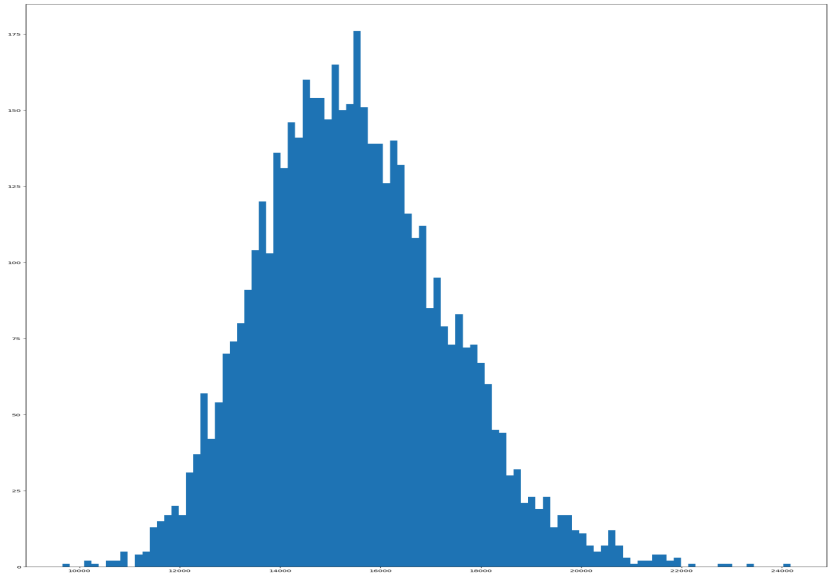

Mean AGB from the gaussian process distribution: 15480
Standard deviation from the gaussian process distribution: 1856

The on ground extrapolated AGB is contained within the +- 20% window of the mean expected biomass predicted by the gaussian process.

Case Study 2

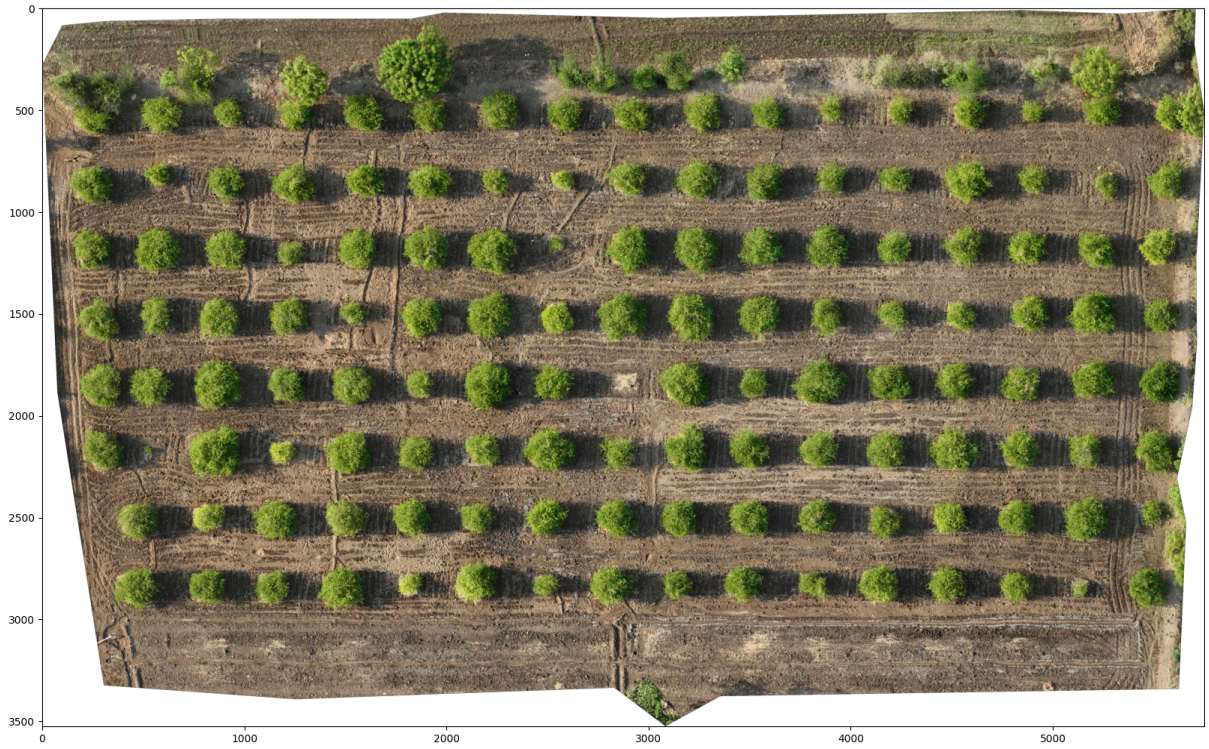
0
500
1000
1500
2000
2500
3000
3500
0
1000
2000
3000
4000
5000


0
500
1000
1500
2000
2500
3000
3500
0
1000
2000
3000
4000
5000

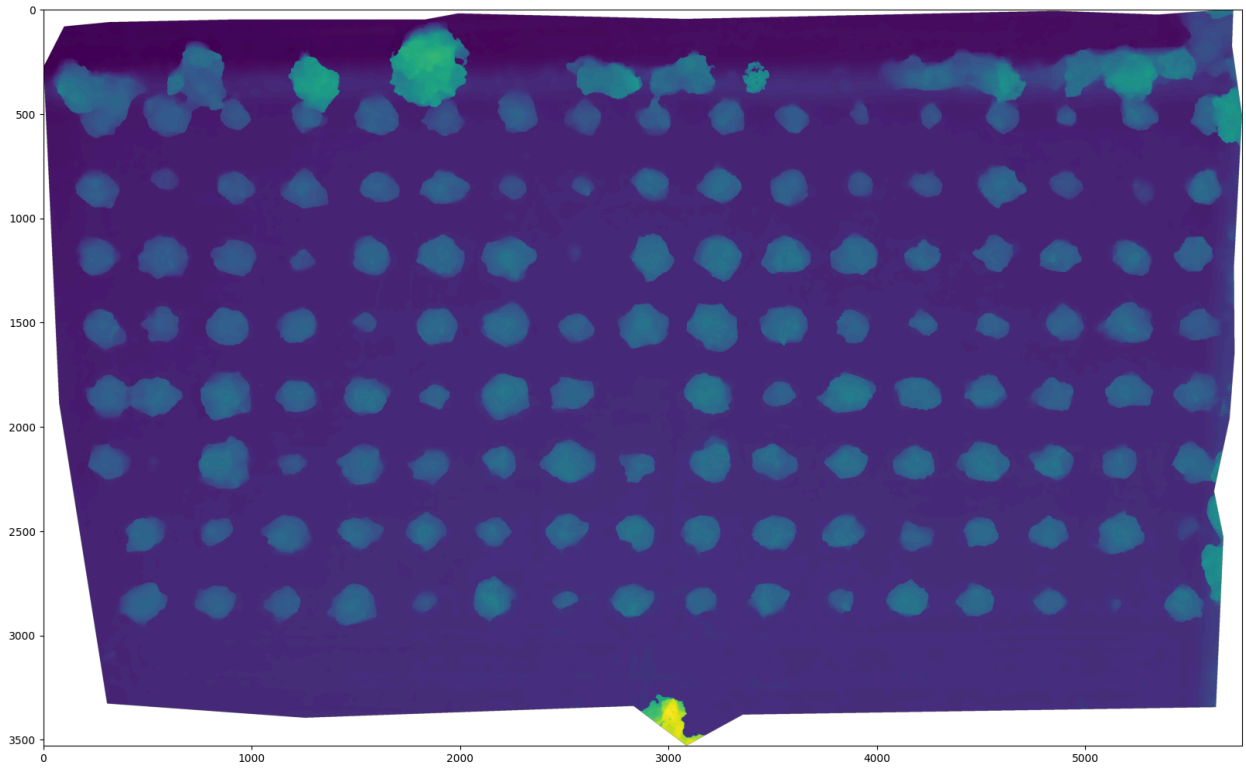

Number of trees measured on the ground = 30
Number of trees on the plot = 130

Total amount of AGB calculated and summed up for the 30 manually measured trees = 1313
Total amount of AGB extrapolated to the plot from the measurement of 30 trees = 5693.

Distribution of total AGB from Gaussian processes:

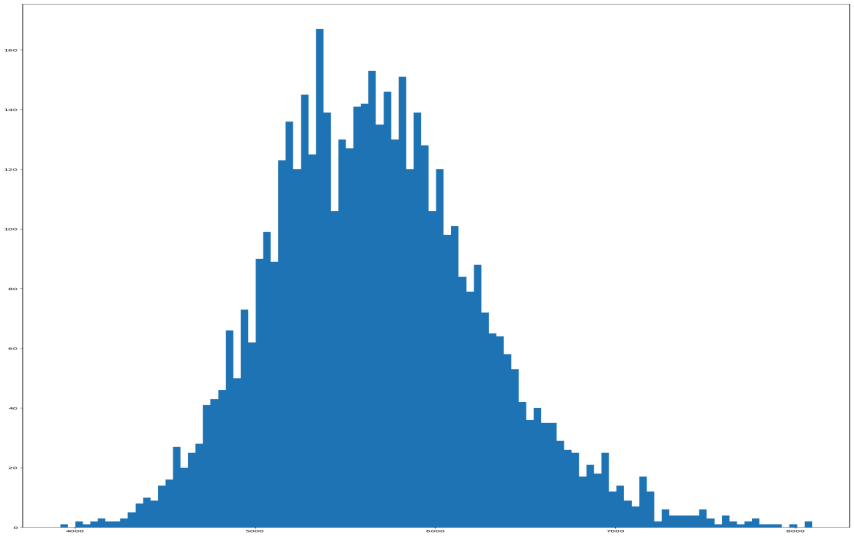

Mean AGB from the gaussian process distribution: 5680
Standard deviation from the gaussian process distribution: 590

The on ground extrapolated AGB is contained within the +- 20% window of the mean expected biomass predicted by the gaussian process.

11. Limitations

While our Detectron2-based tree detection model achieves 89% precision and recall, manual verification remains necessary before biomass calculations. This limitation is inherent to current computer vision systems where 100% accuracy is unattainable, particularly in complex agroforestry environments with varying tree sizes, overlapping canopies, and diverse species compositions.

The Gaussian Process Regression model for DBH prediction, despite achieving a mean absolute error of 6.16 on log(DBH), can, in some scenarios, exhibit significant uncertainty bounds that must be properly communicated to end users. The model's performance varies considerably across different tree sizes and species, with larger uncertainties observed for types and sizes of trees that are either unseen or seen infrequently by the training dataset. With continued addition of diverse data and regular model retraining, this model will continue to improve.

Our approach makes several simplifying assumptions that limit its generalizability:

Species-Agnostic Modeling: Both the allometric equations and Gaussian Process models treat all tree species uniformly. While this enables automated processing without species identification, it introduces systematic errors where species-specific growth patterns deviate significantly from the generalized model. The uncertainty quantification partially captures this limitation, but explicit species modeling would improve accuracy. However, this is not an inherent limitation of the overall technical stack. With improvement in species identification, species specific allometry can be easily adopted. Work is currently underway along this line.

Allometric Equation Dependency: The biomass estimation relies on pre-existing allometric equations from Verra's VMD0001 methodology. These equations were developed for specific regional contexts and may not accurately represent the diverse agroforestry systems we encounter.

Geographic Scope: Training data was collected exclusively from Maharashtra's semi-arid agroforestry systems. Model performance in different climatic zones, soil types, or management practices remains unvalidated. The Gaussian Process uncertainty estimates may not adequately capture epistemic uncertainty when applied to significantly different geographic contexts. A solution to this is to add training images from varying geographical regions. Eventually the goal of the model is to make it applicable across India/globe.

Seasonal and Temporal Bias: Data collection occurred during specific seasons, potentially missing important phenological variations that could affect crown size measurements and tree health assessments. Deciduous species during leaf-off periods, or drought-stressed trees, may not be accurately represented in our training data. Similar to geographical scope, the model's temporal coverage can be improved by adding images from various seasons.

Scale Limitations: While our approach works well for plot-level analysis (1-5 hectares), computational and logistical constraints limit its application to larger landscapes without significant infrastructure investment. Processing very large orthomosaics (>10GB) requires substantial cloud computing resources.

Processing Pipeline Dependencies: The current pipeline requires significant manual oversight at multiple stages - from flight planning to data preprocessing to model output validation. While automation reduces labor compared to ground-based methods, the system is not yet robust enough for fully autonomous operation.

Conservative Bias Implementation: Our recommendation to claim biomass credits for 80% of estimated values, while ensuring conservative estimates, may systematically undervalue projects. This trade-off between accuracy and conservatism requires careful consideration in carbon market applications.

Temporal Validation Gap: Case studies demonstrate reasonable agreement between drone-based and manual measurements at single time points, but longitudinal validation over multiple years remains limited. Growth rate predictions and multi-temporal change detection accuracy require extended validation periods.

Several technical limitations point toward necessary future developments, which we will continue to work on and share in subsequent editions of our white paper, including:

- 3D Modeling Integration: Incorporating LiDAR data for volumetric biomass estimation rather than relying solely on allometric equations
- Multi-spectral Analysis: Expanding beyond RGB imagery to include near-infrared and other spectral bands for improved tree health and species discrimination

- Automated Quality Control: Developing robust automated validation systems to reduce manual oversight requirements
- Edge Case Handling: Improving model performance for challenging scenarios such as very dense plantations, mixed-age stands, and stressed vegetation

These limitations do not invalidate the methodology, but highlight areas where users should exercise caution and where future research should focus to improve system robustness and accuracy.

12. Conclusion

The automated calculation of biomass using drone imagery is able to estimate the total biomass on an agroforestry plantation by extracting the relevant data about tree dimensions from drone Orthomosaics and the DEM. The gaussian process is able to capture the uncertainties in DBH predictions and cascade them down to the biomass estimation level. The case studies show that on field measurements based biomass calculations are captured within 20% boundaries of the mean expected biomass as per the drone image processing via gaussian processes. Claiming the biomass credits for 80% of the biomass estimated by the gaussian process is assured to not overestimate the biomass numbers and helps in establishing a conservative yet reasonably accurate estimates for the biomass. Although there are limitations to be overcome, key pieces of the tech stack are in place that can be independently improved to make the overall tech stack more robust.

References

1. Zhou, L., Li, X., Zhang, B., Xuan, J., Gong, Y., Tan, C., Huang, H., & Du, H. (2022). Estimating 3D Green Volume and Aboveground Biomass of Urban Forest Trees by UAV-Lidar. Remote Sensing, 14(20), 5211. https://doi.org/10.3390/rs14205211
2. Jayathunga, S., Owari, T., & Tsuyuki, S. (2018). The use of fixed–wing UAV photogrammetry with LiDAR DTM to estimate merchantable volume and carbon stock in living biomass over a mixed conifer–broadleaf forest. International Journal of Applied Earth Observation and Geoinformation, 73, 767–777. https://doi.org/10.1016/j.jag.2018.08.017
3. T. Mayamanikandan, R. Suraj Reddy and C. S. Jha, "Non-Destructive Tree Volume Estimation using Terrestrial Lidar Data in Teak Dominated Central Indian Forests," 2019 IEEE Recent Advances in Geoscience and Remote Sensing : Technologies, Standards and Applications (TENGARSS), Kochi, India, 2019, pp. 100-103, doi: 10.1109/TENGARSS48957.2019.8976068.
4. Verified Carbon Standard. (2022). VMD0001 Estimation of carbon stocks in the above- and below ground biomass in live tree and non-tree pools (CP-AB), v1.2. Verra. https://verra.org
5. Brahma, B., Nath, A. J., Deb, C., Sileshi, G. W., Sahoo, U. K., & Das, A. K. (2021). A critical review of forest biomass estimation equations in India. Trees, Forests and People, 5, 100098. https://doi.org/10.1016/j.tfp.2021.100098
6. Sileshi, G. W. (2014). A critical review of forest biomass estimation models, common mistakes and corrective measures. Forest Ecology and Management, 329, 237–254. https://doi.org/10.1016/j.foreco.2014.06.026
7. Ke, Y., & Quackenbush, L. J. (2011). A comparison of three methods for automatic tree crown detection and delineation from high spatial resolution imagery. International Journal of Remote Sensing, 32(13), 3625–3647. https://doi.org/10.1080/01431161003762355
8. Zheng, J., Fu, H., Li, W., Wu, W., Zhao, Y., Dong, R., & Yu, L. (2020). Cross-regional oil palm tree counting and detection via a multi-level attention domain adaptation network. ISPRS Journal of Photogrammetry and Remote Sensing, 167, 154–177. https://doi.org/10.1016/j.isprsjprs.2020.07.002
9. Yao, L., Liu, T., Qin, J., Lu, N., & Zhou, C. (2021). Tree counting with high spatial-resolution satellite imagery based on deep neural networks. Ecological Indicators, 125, 107591. https://doi.org/10.1016/j.ecolind.2021.107591
10. Neupane B, Horanont T, Hung ND (2019) Deep learning based banana plant detection and counting using high-resolution red-green-blue (RGB) images collected from unmanned aerial vehicle (UAV). PLoS ONE 14(10): e0223906. https://doi.org/10.1371/journal.pone.0223906

11. Skole, D. L., Samek, J., Mehra, S., Bajaj, R., & Tamay, T. (2025). Monitoring the extent of trees outside of forests in South Asia: Nature-based solutions for climate change mitigation. In Remote sensing of land cover and land use changes in South and Southeast Asia (Vol. 1, 1st ed., p. 28). CRC Press. https://doi.org/10.1201/9781003396253
12. Weinstein BG, Marconi S, Aubry-Kientz M, Vincent G, Senyondo H, White EP. DeepForest: A Python package for RGB deep learning tree crown delineation. Methods Ecol Evol. 2020; 11: 1743–1751. https://doi.org/10.1111/2041-210X.13472
13. Weinstein, B. G., Marconi, S., Bohlman, S. A., Zare, A., & White, E. P. (2020). Cross-site learning in deep learning RGB tree crown detection. Ecological Informatics, 56, 101061. https://doi.org/10.1016/j.ecoinf.2020.101061
14. Khan, S. and Gupta, P. K.: Comparative Study of Tree Counting Algorithms in Dense and Sparse Vegetative Regions, Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci., XLII-5, 801–808, https://doi.org/10.5194/isprs-archives-XLII-5-801-2018, 2018.
15. C. I. Bliss, K. A. Reinker, A Lognormal Approach to Diameter Distributions in Even-Aged Stands, Forest Science, Volume 10, Issue 3, September 1964, Pages 350–360, https://doi.org/10.1093/forestscience/10.3.350
16. Renato Augusto Ferreira de Lima, João Luís Ferreira Batista, Paulo Inácio Prado, Modeling Tree Diameter Distributions in Natural Forests: An Evaluation of 10 Statistical Models, Forest Science, Volume 61, Issue 2, April 2015, Pages 320–327, https://doi.org/10.5849/forsci.14-070