Direct mapping of plantation forest aboveground biomass change with deep learning and SAR-optical fusion

Brian Lee, Alex Rich, Nathan Thomas, Atticus Stovall, Temilola Fatoyinbo, Guillermo F Olmedo, Pablo Ramírez de Arellano, Pablo Mena-Quijada, Robert Heilmayr

Preprint status: This is a non-peer-reviewed preprint submitted to EarthArXiv. This manuscript has been submitted to *Remote Sensing of Environment* and is currently under peer review.

Author Affiliations & Contact

Brian Lee* (corresponding author)

Bren School of Environmental Science & Management, University of California, Santa Barbara, 2400 Bren Hall, Santa Barbara, CA 93106, USA | bhyleee@gmail.com

Alex Rich

Department of Computer Science, University of California, Santa Barbara, 2104 Harold Frank Hall, Santa Barbara, CA 93106, USA

Nathan Thomas

History, Geography & Social Sciences Department, Edge Hill University, St. Helens Rd., Ormskirk, Lancashire, UK

Atticus Stovall

Department of Geographical Sciences, University of Maryland, 7251 Preinkert Drive, College Park, MD 20742, USA

Temilola Fatoyinbo

MIT Media Lab, Massachusetts Institute of Technology, 75 Amherst St., Cambridge, MA 02139, USA

Guillermo F Olmedo, Pablo Ramírez de Arellano, Pablo Mena-Quijada

Investigaciones Forestales Bioforest S.A., Camino a Coronel, Km. 15, 403 0000, Concepción, Chile

Robert Heilmayr

Bren School of Environmental Science & Management, University of California, Santa Barbara, 2400 Bren Hall, Santa Barbara, CA 93106, USA

Direct mapping of plantation forest aboveground biomass change with deep learning and SAR-optical fusion

Brian Lee[a, *, 1], Alex Rich[b], Nathan Thomas[c], Atticus Stovall[d], Temilola Fatoyinbo[e], Guillermo F Olmedo[f], Pablo Ramírez de Arellano[f], Pablo Mena-Quijada[f], Robert Heilmayr[a]

[a] *Bren School of the Environmental Science & Management, University of California, Santa Barbara, 2400 Bren Hall, Santa Barbara, CA 93106, USA*

[b] *UCSB Computer Science, University of California, Santa Barbara, 2104 Harold Frank Hall, Santa Barbara, CA 93106, USA*

[c] *History, Geography & Social Sciences Department, Edge Hill University, St. Helens Rd., Ormskirk, Lancashire, UK*

[d] *Department of Geographical Sciences, University of Maryland, 7251 Preinkert Drive, College Park, MD 20742, USA*

[e] *MIT Media Lab, Massachusetts Institute of Technology, 75 Amherst St., Cambridge, MA 02139, USA*

[f] *Investigaciones Forestales Bioforest S.A., Camino a Coronel, Km. 15, 403 0000, Concepción, Chile*

Abstract

Accurately tracking changes in forest aboveground biomass (ΔAGB) is necessary for understanding global carbon dynamics. Traditional approaches estimate ΔAGB indirectly by differencing two independent biomass predictions, compounding uncertainty and reducing accuracy. Here, we develop a Mixture-of-Experts (MOE) machine learning framework that uses multi-sensor fusion of Sentinel-1 C-band SAR, ALOS PALSAR L-band SAR, and Sentinel-2 optical data to directly estimate ΔAGB and prediction uncertainty. Our training and validation data are drawn from a large, repeat forest inventory (9,087 plots resampled between 2016-2021) that characterizes a variety of management conditions causing both increases and decreases in AGB (e.g. planting, growth, pruning, thinning and harvest). Using this dataset, we trained an ensemble of three component models at 30-meter resolution: (1) a classifier that identifies change type, (2) a regression model for AGB growth, and (3) a regression model for AGB loss. This MOE approach of ΔAGB achieves high accuracies (R^2 of 0.90, RMSE of 26.89 Mg/ha, global NRMSE of 4.5%), and dramatically outperforming the indirect approach applied to both internal baselines and existing global products (reducing RMSE by at least 57%). A heteroscedastic Gaussian Negative Log-Likelihood loss function and Monte Carlo dropout provide per-pixel predictive uncertainty alongside each prediction, offering a transparent and operationally useful measure of model confidence. This study demonstrates how multi-sensor fusion,

*Corresponding author: brian.lee-4@colorado.edu

[1] Present address: *Environmental Data Science Innovation & Impact Lab, University of Colorado Boulder, SEEC Building, 4001 Discovery Drive, Boulder, CO 80309, USA*

large-scale repeat inventories, and MOE modeling can improve the accuracy and reliability of forest carbon monitoring for climate mitigation and sustainable forest management.

Keywords: Biomass estimation, direct change, machine learning, data fusion, carbon monitoring, plantation forests, uncertainty quantification

1. Introduction

Forests play an important role in regulating global carbon fluxes, with significant implications for national and international carbon budgets (Harris et al., 2021; Pugh et al., 2019). Global initiatives, such as REDD+ (Reducing Emissions from Deforestation and Forest Degradation) rely on robust Measurement, Reporting, and Verification (MRV) frameworks to help accurately monitor forests (Gibbs et al., 2007; Hawthorne et al., 2016; Herold et al., 2011). However, forests are dynamic systems subject to both natural and anthropogenic disturbances, such as timber harvesting, wildfires, and disease (Hurteau & North, 2009; McDowell et al., 2020; Pugh et al., 2019). This poses a challenge: MRV systems must track carbon in landscapes that are constantly changing, while also accounting for the uncertainty introduced during this complex monitoring process.

Recent advances in remote sensing and machine learning provide powerful tools for developing accurate forest monitoring solutions. Remote sensing allows scalable, repeatable measurements of forest structure over large areas, reducing dependence on costly field inventories (Lechner et al., 2020; Mitchell et al., 2017; Potapov et al., 2014). Fusion of data from optical and SAR sensors has significantly improved the accuracy of biomass mapping, from global forest extent products to fine-scale aboveground biomass (AGB) estimates (Duncanson et al., 2020; Joshi et al., 2016; Kellndorfer et al., 2010; Montesano et al., 2013). In parallel, a wide range of machine learning approaches have been applied to the remote sensing of AGB (Amitrano et al., 2023; Li et al., 2020; Maxwell et al., 2018; Zhang, 2010). These approaches range from pixel-wise models such as Random Forests and gradient boosting to large deep learning architectures that capture spatiotemporal patterns and landscape-level trends (L. Chen et al., 2019; Jamaluddin et al., 2024; Zeng et al., 2019). Yet despite these advances, accurately estimating changes in biomass (ΔAGB) over time remains particularly challenging, as errors from each stage of the modeling process accumulate when predictions from multiple time points are combined (McRoberts et al., 2015, 2018).

Indirect estimation–calculating ΔAGB by differencing AGB predictions across two points in time–is the dominant strategy for biomass change monitoring, underpinning widely used global products (Baccini et al., 2017; Harris et al., 2021). Yet this approach propagates errors from both predictions–compounding uncertainty from data collection, allometric modeling, and multiple model

output, which can substantially inflate the uncertainty of the differenced estimate relative to either individual prediction alone (McRoberts et al., 2015; Zolkos et al., 2013). In contrast, direct modeling of ΔAGB relies on training a single model on observed AGB differences, bypassing the need for intermediate AGB estimates and thereby reducing error propagation (McRoberts et al., 2015; Skowronski et al., 2014). However, applications of direct approaches have been limited by two barriers: 1) the scarcity of representative, repeat ground reference datasets (Bollandsås et al., 2013; Skowronski et al., 2014); and 2) the difficulty of capturing both positive growth and extreme negative biomass loss within a single regression model, given their strong imbalance and change magnitude (Kennedy et al., 2010; Næsset et al., 2013). These challenges require modeling frameworks that can explicitly separate gradual biomass accumulation from abrupt loss, while also estimating uncertainty in ΔAGB predictions rather than treating change as a simple difference between two independently modeled biomass estimates.

This study addresses key gaps in the forest biomass monitoring literature by presenting an ensemble machine learning framework for directly estimating ΔAGB across intensively managed plantation forests in central Chile. We integrate multi-sensor data resampled to 30 m resolution (Sentinel-1, Sentinel-2, ALOS PALSAR, and Copernicus DEM) with an unusually large repeat forest inventory dataset, consisting of more than 18,174 biomass observations from 9,087 re-sampled plots collected between 2016 to 2021. Our Mixture-of-Experts (MOE) modeling framework first classifies the change-type of a specific pixel, then routes it to one of two neural network-based regression models specialized to predict either AGB growth or loss, directly addressing the imbalance between gradual growth and abrupt loss. We leverage repeat observations within the structured management of industrial timber operations to assess how predictive uncertainties align with residual errors to better understand how model inaccuracies arise. Together, these advances provide a framework for monitoring biomass change in intensively managed plantation forests, with broader relevance for forest management and carbon accounting.

2. Methods

2.1. Study Area

Chile's timber industry, dominated by 3 million acres of fast-growing plantation species such as Monterey Pine (*Pinus radiata*) and *Eucalyptus* species (primarily *Eucalyptus nitens* and *Eucalyptus globulus*), provides an ideal context for developing and validating methodologies for ΔAGB estimation. The controlled management interventions in clear, spatially defined parcels include harvest, pruning, thinning, and replanting events. These can serve as potential proxies for broader forest processes such as deforestation, degradation, and regeneration. Additionally, the timber industry's systematic monitoring practices yield high-quality forest inventory datasets that, when

combined with extensive multi-sensor remote sensing data, enable the development of scalable and accurate AGB monitoring approaches (Olmedo et al., 2020; Salas et al., 2016).

Our study area is located in central Chile, extending from the capital city, Santiago, southward to the city of Puerto Montt, encompassing eight of Chile's sixteen municipal regions (UTM Zone 18S, between 5,400,000 m N and 6,200,000 m N). This area spans approximately 1,500 kilometers and covers a diverse geography, from coastal plains to the Andes mountains (*Fig. 1*). The area is characterized by steep gradients in elevation, ranging from sea level to over 3,000 meters. Climate variability is significant, with a semi-arid Mediterranean climate dominating the Northern part and a temperate rainforest climate in the Southern portion. Forest plantations and native forests, including the Valdivian temperate rainforest, are both subject to various pressures, including timber harvesting,

human development, and natural disturbances such as wildfires and drought, all of which contribute to biomass fluctuations that this study aims to track (Carrasco et al., 2022; Cordero et al., 2024).

Figure 1. *Spatial density of forest inventory plots across the central Chile study area, summarized as the number of plots per 3 km x 3 km grid cell. Darker colors indicate higher plot density. The insets show a) the location of the study area within South America, and b) the arrangement of inventory plots within a 1 km x 1 km cell.*

2.2. Reference data

Celulosa Arauco y Constitución (ARAUCO) is one of the world's largest timber producers, operating more than 1.6 million hectares of plantation forest across South America (ARAUCO, 2022). In Chile, their plantations are dominated by *Pinus radiata* and *Eucalyptus* species. As part of their routine operations, ARAUCO surveyed 745,930 forest inventory plots between the years of 2000 and 2021, generating 955,398 visits (plot-by date measurements). Each visit includes tree- and plot-level attributes such as location, plot ID, parcel ID, species, mean diameter at breast height of trees (DBH), mean height of trees, number of trees in plot, number of trees in parcel, total basal area, inventory age, and basal area. More recent inventories (2017-2021) also include dry AGB estimates from allometric equations derived from destructive sampling of 1,093 pine trees conducted within ARAUCO's plantations (Olmedo et al., 2020; Sandoval et al., 2022). We applied these same allometric equations to earlier (2016-2017) visits to calculate the plot-level AGB estimates.

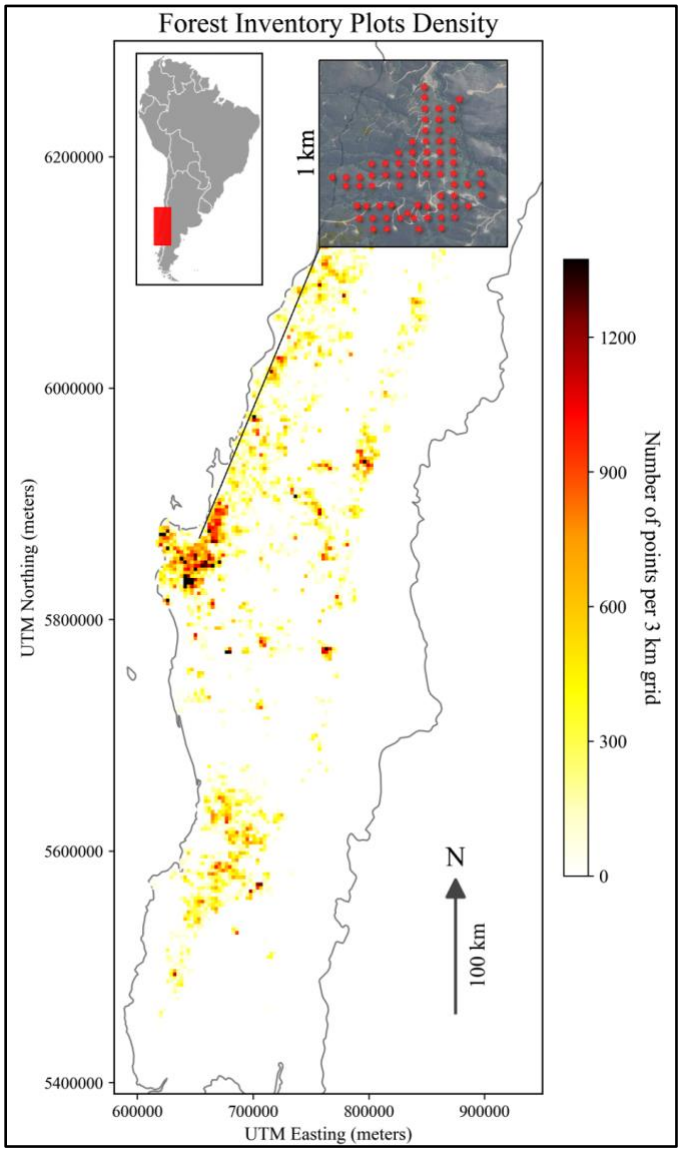

For this study, we restricted the ARAUCO dataset to 2016-2021, matching

the availability of overlapping observations from our chosen sensors (Sentinel-1, Sentinel-2, and ALOS PALSAR 2). Coordinates were reprojected to UTM (WGS 84 UTM 18S). Biomass values were standardized from kilograms of C per tree (kg C/tree) to megagrams of C per hectare (Mg C/ha) using plot-level measurements of forest density. Given our study's focus on measuring ΔAGB, we retained only plots with at least two visits. Due to slight geolocation errors inherent in field inventory collection, we defined repeat visits as those occurring within 40m of each other from different year-month combinations, consistent with the nominal plot radius used in ARAUCO's inventory protocol. For plots with more than two visits, only the first and last were retained. Each pair of visits to the same plot was used to derive a single change observation–a paired ($t_1 \rightarrow t_2$) ΔAGB estimate formatted as a 30m pixel, consistent with remote sensing predictors. For instance, if a plot was visited in 2019 (350 Mg C/ha) and again in 2021 following a clear-cut (0 Mg C/ha), the resulting change observation would receive a label of -350 Mg C/ha.

Given the asymmetry between subtle growth and abrupt losses, we classified each change observation into one of several change-type categories based on its ΔAGB values and the stand age at t_2. Large, negative ΔAGB (> 150 Mg/ha lost) were labeled "harvest", smaller negative ΔAGB was considered thinning or pruning, positive ΔAGB as new growth (0-2 years), early growth (3-7 years) and mid/late growth (> 8 years) based on the stand age of the final observation (See SI 7.1). These categories were subsequently used in the MOE ML framework, where a classifier first predicts change-type before using specialized regression models to estimate growth or loss.

Within our study time frame, 7,310 plots had at least two visits, yielding the same number of change observations. These plots are distributed across 567 unique forest parcels and a mean revisit period was 1.4 years. Due to the operational management of eucalyptus stands, repeat visits to eucalyptus plots were rarely observed in our repeat data and we therefore restricted our analysis to pine plantations. Across all change observations, the mean ΔAGB between all years is +19.06 Mg/ha, with a standard deviation of 34.48 Mg/ha and a range of -144.22 Mg/ha AGB lost and +201.52 Mg/ha AGB gained. Note that this range reflects only change observations derived from repeated visits; because ARAUCO rarely revisits plots following harvest or records AGB over bare earth, large negative and positive ΔAGB values associated with the majority of harvest and new growth events are instead represented by augmented change observations (see *Reference data augmentation* below).

The age distribution of change observations reveals an imbalance in the underlying dataset: the median age at the first visit (t_1) was 9 years, while the median age at the second visit (t_2) was 11.3 years. Few plots were visited before the age of 6 (< 8%), and no plots captured newly planted stands. As a result, change observations are dominated by mid/late growth transitions (> 8 years, n = 3,975), with fewer capturing thinning/pruning (n = 1,458), or early growth (n = 1,877). This pattern is consistent with operational priorities: young stands are rarely prioritized for inventory until they reach an age

where management decisions are required, and pre-harvest stands, while inventoried, are often cleared before repeat measurements can be collected. This bias in sampling introduces systematic underrepresentation of regeneration dynamics and abrupt harvest events, the very transitions that are critical for ΔAGB monitoring and carbon accounting. Without correction, models trained on these change observations would be disproportionately tuned to detect gradual stand maturation and management interventions while performing poorly on early growth or complete removals. (See SI 7.1)

Reference data augmentation

To address the class imbalance of the reforestation and harvest change classes and to extend the temporal scope of our project, we augmented the reference dataset with synthetic change observations. For plots without a prior bare-ground observation, we introduced a synthetic t_0 state representing the pre-planting or t_1 representing post-harvest, enabling paired ($t_0 \rightarrow t_1$) change observations that would otherwise be absent from the dataset. These synthetic states enabled additional paired change observations, including bare ground to observed biomass for reforestation events. Candidate synthetic observations were manually validating using very high-resolution Google Earth Pro imagery (2013-2020), identifying growth events where bare ground preceded the inventory and harvest events where it followed (Olofsson et al., 2014; Potapov et al., 2012). Specifically, we added 717 validated bare-ground-to-inventory change observations to capture new growth, and 1,060 validated harvest change observations. Bare-earth pixels from 2016-2021 were aggregated from across 38 parcels and summarized using their mean and variance of spectral and SAR indices, creating a spectral library of bare-ground conditions. These statistics were then used to impute features across our remotely sensed predictors for earlier zero-biomass plots (2010-2015) when direct sensor observations were unavailable. This approach, inspired by both oversampling methods such as SMOTE (Chawla et al., 2002), spectral endmember modeling (Powell et al., 2007; Roberts et al., 1998), and uni-temporal change estimation (Puliti et al., 2021), allowed us to significantly increase coverage of transitions to and from zero AGB, resulting in a final reference dataset of 9,087 change observations.

2.3. Remotely sensed data

We used complementary datasets from Sentinel-1 (C-band SAR), ALOS PALSAR (L-band SAR), Sentinel-2 (optical multispectral), and Copernicus GLO-30 DEM (Table 1). Sentinel-1 backscatter provides sensitivity to upper canopy structure, ALOS PALSAR penetrates deeper into the canopy for a better understanding of the lower canopy, and Sentinel-2 supplies high-resolution spectral and vegetation index information (ESA, 2022; Li et al., 2020; Peregon & Yamagata, 2013; Puliti et al., 2021). We derived elevation, slope, and aspect from the Copernicus DEM (ESA, 2022).

All data were accessed and pre-processed in Google Earth Engine (GEE; Gorelick et al., 2017). Sentinel-1 imagery was corrected for border noise and terrain effects and mosaiced using VV, VH, and VV/VH ratio in dB (Mullissa et al., 2021). ALOS PALSAR annual mosaics (JAXA) were orthorectified and slope-corrected, with HH, HV, and their HH/HV ratio extracted in dB (Rosenqvist et al., 2007). Sentinel-2 Level-1C imagery was cloud-masked, composited, and used to generate standard vegetation indices (NDVI, EVI, MSAVI, MSI, NBR, NDMI) in addition to the original bands (Drusch et al., 2012). For both Sentinel-1 and Sentinel-2, we calculated median composites for each year between December and April, corresponding with the generally cloud-free Chilean summer. To better capture parcel-level structural homogeneity, we derived median texture filters for these bands using a 70 m radius. The Copernicus GLO-30 DEM provided terrain attributes of slope, elevation, and aspect. All composites were resampled to 30 m and reprojected to WGS84/UTM18S.

We harmonized the multi-year dataset by calibrating all bands to a 2020 reference year using manually-derived pseudo-invariant features across stable landcover types, ensuring comparability across acquisitions and years (Bao et al., 2012). Each change observation (t_2- t_1) was linked to the corresponding annual composites, and predictor variables were derived as differences between t_2 and t_1 (Δ features). For the indirect annual AGB approach, single-year composites were extracted at each plot visit rather than differenced, using the same 52 features drawn from the corresponding annual composite. In total, 52 features were used, including 32 difference bands, along with static terrain variables and base t_1 RS features, consistent with established change-detection modeling features (Potapov et al., 2012). This approach aligns RS inputs directly with inventory-based biomass changes and provides a foundation for model training and evaluation. (Details of RS data processing, calibration, bare-earth imputation, and quality control are provided in SI 7.2)

Sensor type	***Sensor***	***Channels***	***Bands***	***Resolution/ Footprint***	***Temporal resolution***	***Date availability***
Optical	*Sentinel-2*	*13*	*R, G, B, NIR,*	*10-60m*	*5 days*	*2015 - Present*
SAR	*Sentinel-1*	*2+*	*C-band*	*10m*	*12 days*	*2014 - Present*
SAR	*ALOS PALSAR 2*	*2+*	*L-band*	*25m*	*Composite*	*2015 - 2024*
DEM	*Copernicus GLO-30*	*3*	*Elevation, aspect, slope*	*30m*	*Not Applicable*	*2011 - 2015*

Table 1. *The various remotely sensed datasets used and their capabilities.*

2.4. Modeling approach

To estimate pixel-level changes in forest carbon (ΔAGB) from satellite and inventory data, we developed and compared three modeling approaches, each representing a different strategy for learning from the data (Figure 2). The simplest approach, the indirect model, follows a traditional approach used by many existing global biomass products: it predicts total biomass at each year separately using satellite observations, then subtracts the two estimates to infer change. While intuitive, this strategy has a known limitation--errors from both annual predictions combine, potentially amplifying uncertainty in the final change estimation. To address this, our baseline model takes the

difference in satellite observations between two years and predicts biomass change directly in a single step, avoiding the error compounding of the indirect approach.

Figure 2. *Schematic comparison of indirect and direct approaches to forest carbon change estimation. In the indirect approach (top), separate models are trained to predict absolute forest carbon at each time point using field inventory measurements and satellite imagery; change is then inferred by differencing the two independent estimates, which compounds uncertainty across both predictions. In the direct approach, the model instead takes the difference in field measurements and satellite observations between time points as inputs, and predicts carbon change in a single modeling framework. This approach reduces error propagation and allows the model to learn change-specific signatures, such as the contrast between a harvest and growing stand, rather than relying on absolute biomass estimates at each date.*

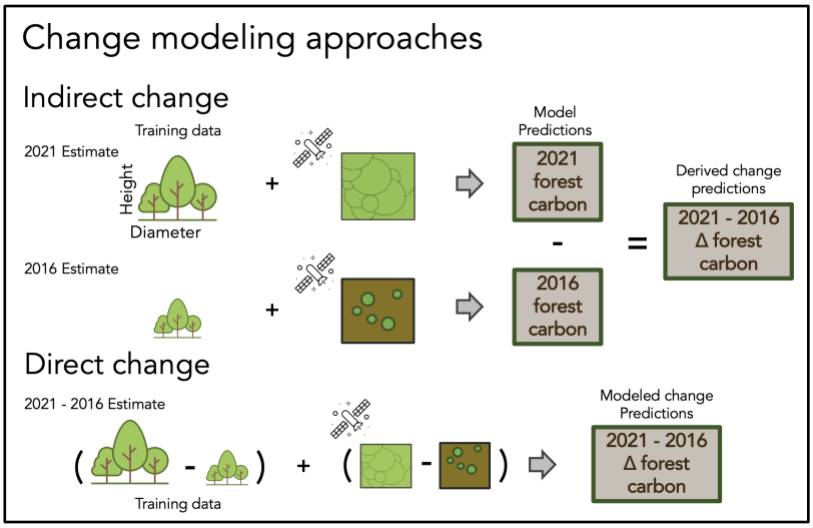

Modeling ΔAGB presents a unique challenge: gradual forest growth produces many observations of slow, positive change, while large-scale clearcuts produce sudden, dramatic losses. A single model must handle both types of change simultaneously. To address this, we implemented a Mixture-of-Experts (MOE) approach as our primary modeling approach, which relies on an initial model to route tasks to specialized "expert" sub-models. This approach is widely used in other domains to split complex prediction tasks into more manageable sub-problems (Y. Liu et al., 2025; Masoudnia & Ebrahimpour, 2014; Yuksel et al., 2012). In our specific use-case, we incorporated a change-type classifier that predicts whether a pixel is more likely to represent growth or loss, which then routes predictions to one of the two specialized sub-models: a growth regressor trained exclusively on positive ΔAGB change observations, or a loss regressor trained on negative ΔAGB change observations. This helps the MOE accurately identify trends of carbon loss (deforestation and thinning/pruning) and gradual growth.

In preliminary experiments, neural network models--specifically multilayer perceptrons (MLP)--consistently outperformed tree-based methods (Random Forest, XGBoost; see SI 7.4) at learning the complex relationships between satellite observations and biomass change (Caraballo-Vega

et al., 2023; Nitze et al., 2012). The indirect, baseline, and MOE modeling frameworks were therefore implemented as feed-forward MLPs with architectures optimized using Optuna, with the indirect and baseline models serving as benchmarks against which the MOE's accuracy and uncertainty estimates were evaluated.

In general, an MLP with L hidden layers is defined as:

$$h^{(l)} = f\left(W^{(l)}h^{(l-1)} + b^{(l)}\right), \quad l = 1, \ldots, L \tag{1}$$

$$y = W^{(\text{out})}h^{(L)} + b^{(\text{out})} \tag{2}$$

where $h^{(0)} = x$ is the standardized input feature vector, $f(\cdot)$ is a nonlinear activation (*tanh, SiLU, GELU*), and dropout/batch normalization are applied between layers. Our indirect annual AGB model used 52 input features, 6 hidden layers of 256 units, SiLU activation, and 0.057 dropout. Our baseline direct change MLP had 2 hidden layers of 256 units, tanh activation, and 0.0 dropout. Our MOE and baseline models both used 52 features, including both years' RS bands in addition to the 32 differenced bands. More complex variants were employed for the MOE classifier, growth regressor, and loss regressor (See SI 7.4).

Model calibration

Model selection and hyperparameter optimization were performed using five-fold spatial cross-validation on the training set, with the folds defined at the forest parcel level to reduce spatial autocorrelation between training and validation observations (A. Ramezan et al., 2019; Pohjankukka et al., 2017). This parcel-level grouping ensures that observations from the same stand are never split across training and validation folds, preventing information leakage from spatially correlated measurements.

Model validation

To evaluate model performance, we reserved a spatially independent held-out test set prior to any model training or hyperparameter optimization. The test area comprises a geographically contiguous region in the central portion of the study area, bounded by coordinates [-73.56°, -37.53°] and [-71.73°, -38.88°], and was selected to ensure complete spatial separation from training parcels. No test set observations were used at any stage of model development, cross-validation, or hyperparameter tuning. The test set contains 1,242 change observations drawn from 91 parcels, distributed across all five change-type classes: new growth (n=220), early growth (n=151), mid/late growth (n=531), thinning/pruning (n=232), and harvest (n=108). The full test set of 1,242 observations includes

augmented change observations with synthetic bare-earth RS features extending temporal coverage back to 2013, prior to Sentinel-1 and Sentinel-2 availability. Comparison involving the indirect MLP and external products (Planet Carbon Diligence, ESA CCI Biomass) are restricted to a subset of 784 observations corresponding to change intervals where real sensor observations exist at both t_1 and t_2 (2016-2021), as these approaches require independent per-year RS inputs rather than pre-computed differenced features. The test area spans comparable ranges of stand age, elevation, and climate conditions to the broader study area, supporting generalization of reported accuracy metrics. In addition to parcel-level accuracy metrics, we conducted parcel-level validation by aggregating predictions across plots within each of the 91 test parcels and applying paired Wilcoxon signed-rank tests to assess whether MOE predictions differed significantly from inventory reference values. This parcel-level assessment directly reflects the operational scale at which forest managers make carbon accounting decisions, and provides a statistically grounded complement to pixel-level metrics.

We evaluated all models using the Coefficient of Determination (R^2), Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Normalized RMSE (NRMSE). Each metric highlights different aspects of predictive performance within our analysis framework. R^2 measures the proportion of variance in observed ΔAGB explained by the model and reflects its ability to capture variability in biomass change. MAE reports the average magnitude of prediction errors, regardless of sign, and therefore represents the typical deviation between observed and predicted ΔAGB across both subtle growth and loss events. RMSE places greater emphasis on large deviations, making it particularly sensitive to abrupt disturbances such as harvest events, which often dominate forest ΔAGB dynamics in our study system. Because absolute RMSE is scale-dependent, we also reported normalized metrics. Relative RMSE (RRMSE), which divides RMSE by the mean of the dependent variable, is widely used in AGB estimation studies (Verrelst et al., 2013; Yue et al., 2018). However, it is not appropriate here because our ΔAGB labels include both gains (forest growth) and losses (harvest, thinning), yielding a low mean value (32.2 Mg/ha across all change observations) that artificially inflated RRMSE. Instead, we use NRMSE, which normalizes RMSE by the observed range of ΔAGB values, which in our study ranges between a minimum loss of -400 Mg/ha, and a positive gain of +200 Mg/ha. This choice avoids inflated values caused by low means and provides a balanced evaluation across disturbance classes. NRMSE values closer to zero indicate better model performance, and when multiplied by 100, can be interpreted as an error percentage relative to the observed range. By incorporating these complementary metrics, we gain a comprehensive understanding of model performance across our different modeling approaches.

2.5. Comparison to traditional ΔAGB differencing

To further evaluate model performance, we compared our MOE outputs against both internal baselines (direct and indirect MLP models; see Section 2.5) and external products, including the Planet Labs Carbon Diligence product, and the ESA CCI Biomass product. For each of these datasets, we extracted annual AGB estimates over our independent testing area in central Chile and aligned them with the same spatial and temporal reference points used in our model validation. We then compared predictions to inventory-based reference labels using common error metrics (R^2, MAE, RMSE, NRMSE) and residual-based uncertainty , though these comparisons were restricted to a smaller sample (784-805 change observations) because augmented change observations with t_0 synthetic states predating Sentinel-1 and Sentinel-2 availability (2010-2015) could not be matched to direct acquisitions from external data products. Focusing upon this subset of observations allowed for a consistent comparison of these different modeling approaches and their suitability for MRV applications in intensively managed plantation landscapes.

Planet Labs Carbon Diligence product

The Planet Labs Carbon Diligence product provides global, annual aboveground carbon estimates at 30 m resolution by fusing GEDI LiDAR Level 4A with multispectral and SAR data using a hybrid deep learning and boosted regression framework (Planet Carbon Diligence Technical Specification v1.1, 2024). We obtained AGB rasters for 2016-2021 over our test area, extracted values at reference plot locations, and computed ΔAGB for each change observation (t_2- t_1). Because the product reports aboveground carbon rather than biomass, we converted values to AGB using a standard carbon fraction (0.47) prior to comparison (Requena Suarez et al., 2019). In addition to the pixel-wise calibrated uncertainty intervals and quality flags, we estimated residual-based uncertainties derived from the differences between Planet ΔAGB and our inventory reference values for cross-model comparison.

ESA CCI Biomass product

The European Space Agency's (ESA) Climate Change Initiative (CCI) Biomass project produces global, annual AGB estimates from 2015-2022 by integrating GEDI LiDAR with a SAR and multispectral sensors approach similar to ours (Santoro et al., 2024). Predictions are provided at 100 m resolution with pixel-level standard deviations and quality flags. For comparison with our study, we resampled the 100m layers to 30 m and extracted values at reference plot locations. As with our earlier comparison with the Planet product, we calculated ΔAGB for each remeasurement interval and model performance was evaluated using standard accuracy metrics and residual-derived uncertainties relative to inventory-based reference values.

2.6. Uncertainty estimation

Understanding uncertainty is critical for trustworthy AGB monitoring and MRV applications. In this context, we distinguish between error – the deviation between a model prediction and a reference value – and predictive uncertainty – the model's own estimate of its expected range of deviation for a given input (Q. Chen et al., 2015; Lu et al., 2016). Unlike error, which can only be computed where reference data exist, predictive uncertainty can be generated for any pixel at inference, making it operationally useful for identifying where model outputs should be interpreted with caution. In this study, our use of a heteroscedastic GNLL loss function combined with Monte Carlo (MC) dropout produces per-pixel predictive uncertainty estimates alongside ΔAGB predictions, and we compare these against residual-based uncertainty derived from the test set.

Within the ARAUCO dataset, forest inventory plot measurements list their parcel ID, which reflects the same species planted at the same time in close proximity, allowing us to leverage within-parcel variability to empirically estimate measurement variability (Chave et al., 2014). Using a MC approach, tree diameters and heights were repeatedly sampled from normal distributions defined by observed within-parcel standard deviations (Réjou-Méchain et al., 2017). Each realization was propagated through the species-specific allometric equation, which included published model error, to generate a distribution of biomass estimates per plot (Sandoval et al., 2022). The variance of these simulated biomass estimates was then used as the propagated label uncertainty for AGB (Equation 3).

$$\sigma^2_{AGB} = \frac{1}{M-1}\sum_{m=1}^{M}(B^{(m)} - \bar{B})^2 \tag{3}$$

where M is the number of Monte Carlo realizations, $B^{(m)}$ is the biomass estimate from the m^{th} realization, and $\underline{B}$ is the mean biomass across all realizations.

All regression models, including the indirect annual network, baseline direct ΔAGB network and the MOE growth and loss regressors were trained with a heteroscedastic GNLL loss function (Equation 4), which enables the network to predict both the mean and variance of ΔAGB for each pixel (Schoups & Vrugt, 2010). This predicted variance captures input-dependent uncertainty in the model's predictions. To further account for model-based variance, we applied MC dropout at inference, running multiple stochastic forward passes and computing the variance across predictions (Z. Liu et al., 2025). Combined, these return a per-pixel measurement of total predictive uncertainty directly associated with model outputs.

$$L_{\mathrm{NLL}} = \frac{1}{2}\left[\frac{(y-\mu)^2}{\sigma^2} + \log\sigma^2\right] \tag{4}$$

For the indirect approach, we trained a single regressor to predict annual AGB estimates at the two test years (t_1, t_2), each producing a mean and an input-dependent variance under the GNLL (Equation 4). Variance was again estimated via MC dropout at inference. Assuming independence between year-specific errors, the variance of the differenced AGB was computed as the sum of the per-year variances (Equation 5), from which we derived the pixel-level standard deviation of indirect change estimates.

$$Var(\Delta AGB) = Var(AGB_{t_2}) + Var(AGB_{t_1}) - 2 \cdot \mathrm{Cov}(AGB_{t_2}, AGB_{t_1}) \tag{5}$$

To benchmark our predictive uncertainty directly associated with the different modeling approaches, we also computed a residual-based uncertainty metric, an empirical measure of total unexplained error derived from the variance of prediction residuals (Equation 6; Lu et al., 2016). For each change observation within our test area, residuals were calculated and summarized as their standard deviation. This metric provides a model-agnostic measure of realized predictive error, enabling direct comparison of uncertainty estimates across methods and datasets (Atkinson & Foody, 2002; Gneiting & Raftery, 2007).

$$\hat{\sigma}_{\mathrm{resid}} = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left[(\hat{y}_i - y_i) - \frac{1}{n}\sum_{j=1}^{n}(\hat{y}_j - y_j)\right]^2} \tag{6}$$

2.7. Estimating parcel-scale changes in biomass

To demonstrate the application of our framework at operational scales, we implemented two parcel-level biomass change analyses using ARAUCO's forest property boundaries. Model predictions of ΔAGB were first generated across the independent testing area between 2016 and 2021. The resulting raster outputs were then spatially clipped to specific ARAUCO parcel boundaries, which delineates the extent and management attributes of individual plantation blocks. For each selected parcel, ΔAGB values were then aggregated by calculating the mean pixel and total biomass change across all included pixels. This analysis allowed parcel-level summaries of forest growth and loss, allowing us to directly compare model outputs with ARAUCO's forest inventory records (Figure 3).

Restricting the analysis to parcels with available inventory information ensured that the change estimates reflected only pine plantations under active management, minimizing uncertainty introduced by non-plantation forest land cover types. To quantify parcel-level uncertainty, we applied

the law of total variance across MC aggregates. Pixel-level predictive uncertainty was propagated to the parcel scale using standard variance aggregation, and we report 95% confidence intervals derived from the total predictive uncertainty.

3. Results

3.1. Direct estimation of ΔAGB

The primary Mixture-of-Experts model accurately predicts changes in aboveground biomass (Figure 3). Across the 1,242 independent change observations reserved for testing, the MOE model reached an R^2 of 0.90, an RMSE of 26.90 Mg/ha, and a global NRMSE of 4% when bounded by a maximum positive growth of +200 Mg/ha and a minimum negative loss of -400 Mg/ha (See Section 2.4). For comparison, the baseline MLP returned similarly high accuracy metrics (R^2 of 0.87, RMSE 31.07 Mg/ha, NRMSE of 5.2%, 85% class accuracy). While both models performed similarly, the MOE provided improvements in predictive accuracy and error reduction across all metrics.

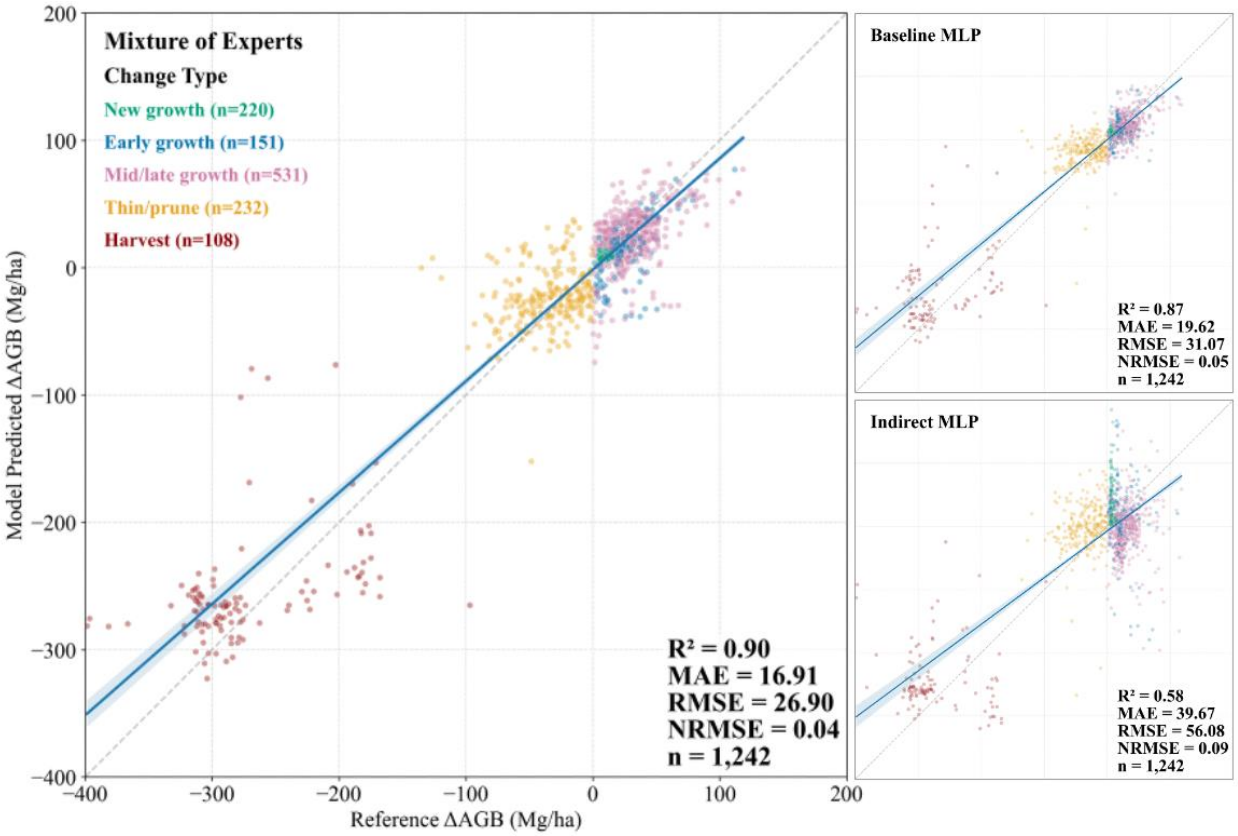

Figure 3. *Mixture-of-Experts predictions of change in ΔAGB (Mg/ha) on the common independent test set shared by the MOE, baseline MLP, and indirect MLP models (n = 1,242). Each point is colored by change type (New growth: green, Early growth: blue, Mid/late growth: purple, Thin/prune: orange;*

Harvest: red). The dashed gray line is the 1:1 reference; the solid blue line represents the fitted trend with 95% CI band. The MOE shows strong overall agreement with the reference data (R^2 = 0.90, MAE = 16.91 Mg/ha, RMSE = 26.9 Mg/ha, NRMSE = 0.04), outperforming both the baseline MLP and indirect MLP shown on the right. Harvest and new growth align closely with the 1:1 line, reflecting the added coverage from our synthetic early growth labels. Thin/prune and mid/late growth exhibit wider residual spread and underpredict positive ΔAGB. The axis limits are fixed at -400 to 200 Mg/ha.

The MOE's strong overall performance is partly attributable to the gating classifier (accuracy = 0.86, precision = 0.88, recall = 0.86, F1 = 0.87), which routes each change observation to either the growth or loss regression expert before prediction. Although the classifier operates on a binary growth/loss distinction, we evaluated its performance against the full five-class change typology to better understand where errors occur (Figure 3). Classification performance was highest for classes with distinct structural signals–all new growth and nearly all harvest change observations (n = 220, n = 108) were correctly identified, reflecting the model's ability to separate bare-earth transitions from established stands. More subtle change types caused higher confusion; thinning/pruning activities were more likely to be misclassified as early growth or mid/late growth. These errors reflect the more subtle spectral and structural difference associated with gradual growth and interventions compared to abrupt disturbances such as harvest or establishment.

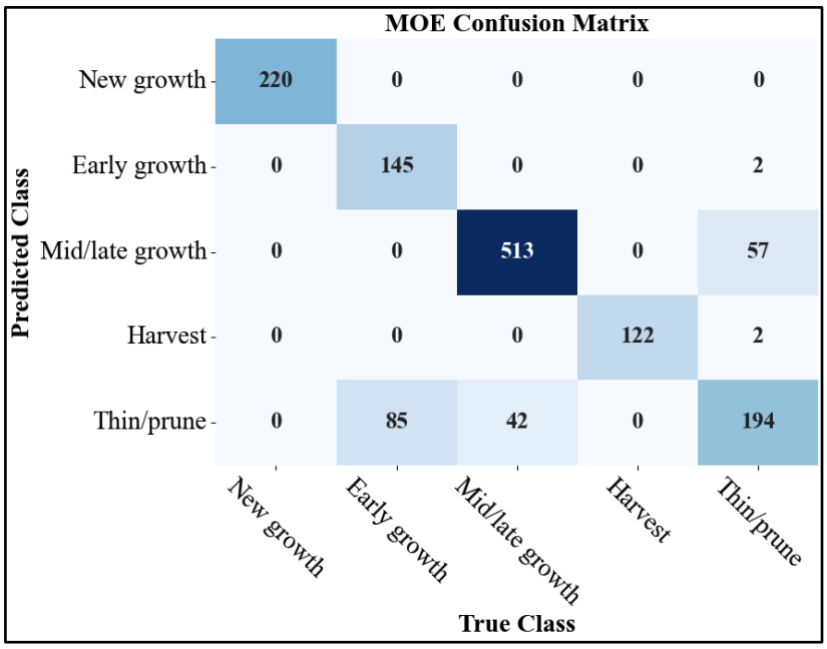

Figure 3. *Confusion matrix for the MOE classifier showing predicted versus true reference forest change types. The classifier achieved high accuracy overall (86.4%), with strong performance for abrupt transitions such as new growth and harvest. Misclassifications were more frequent for subtle transitions, including confusion between early growth, mid/late growth, and thinning/pruning.*

Per-class prediction accuracy of Mixture-of-Experts Model						
	MAE (Mg/ha)	*RMSE (Mg/ha)*	*Bias (Mg/ha)*	*Residual SD (Mg/ha)*	*Relative MAE*	*N*
Overall	*16.91*	*26.90*	*0.15*	*26.91*	*0.36*	*1242*
New growth	*2.0*	*2.90*	*17.34*	*2.63*	*0.28*	*220*
Early growth	*14.12*	*21.20*	*-42.31*	*18.28*	*0.55*	*151*
Mid/late growth	*15.31*	*20.84*	*-21.01*	*19.95*	*0.53*	*531*
Harvest	*42.93*	*56.60*	*6.22*	*54.26*	*0.16*	*108*
Thin/prune	*24.43*	*32.97*	*34.64*	*30.65*	*0.69*	*232*

Table 2. *Per-class prediction accuracy of the Mixture-of-Experts model across the 1,242 independent test change observations. Metrics include MAE and RMSE (Mg/ha), bias, residual standard deviation (SD, Mg/ha), and relative MAE (unitless; MAE normalized by the mean absolute reference ΔAGB within each class). Relative MAE enables comparison of error magnitude across change types that differ substantially in characteristic scale.*

Evaluation metrics across modeling approaches (constant observations across data products)							
	R^2	*MAE (Mg/ha)*	*RMSE (Mg/ha)*	*NRMSE (%) [-400, 200]*	*Residual mean (Mg/ha)*	*Residual std*	*N*
Mixture-of-Experts	*0.90*	*19.13*	*30.46*	*5%*	*-1.02*	*27.44*	*784*
Baseline MLP	*0.89*	*22.45*	*34.93*	*6%*	*-2.75*	*30.90*	*784*
Indirect MLP	*0.67*	*44.07*	*60.20*	*10.7%*	*-1.53*	*64.53*	*784*
Planet CD	*0.34*	*51.80*	*86.51*	*14.2%*	*-4.67*	*85.32*	*784*
ESA CCI	*0.14*	*74.67*	*98.79*	*16.4%*	*27.23*	*95.02*	*784*

Table 3. *Comparison of various ΔAGB datasets across 784 independent test change observations, including the primary MOE model, baseline change MLP, indirect differencing approach, and two widely used external datasets: Planet Labs Carbon Diligence and ESA CCI. Accuracy metrics for ΔAGB include R2, RMSE (Mg/ha), MAE (Mg/ha), NRMSE (%), and residual-based uncertainty. Sample size for this comparison is restricted to a smaller temporal subset due to the data augmentation predating Sentinel-1, -2 availability.*

3.2. Comparison with indirect AGB

The primary MOE model consistently outperformed the indirect approaches, including the indirect MLP trained on the reference labels as well as the Planet and ESA CCI derived carbon products. Due to differences in data availability, these sample sizes are smaller than those used for the direct MOE and baseline MLP comparisons (n = 1,242). The indirect differencing approach achieved an R^2 = 0.58, an RMSE of 56.07 Mg/ha, and NRMSE of 9.34%, (n = 1,242). The external datasets performed less favorably. The Planet Labs dataset systematically underestimated both annual estimates of AGB and ΔAGB, R^2 = 0.33, an RMSE of 85.40 Mg/ha, and NRMSE of 14.2%, (n = 805). Similarly, the ESA CCI Biomass product showed the weakest alignment with the reference data, with R^2 = 0.13, RMSE = 98.79 Mg/ha, and an NRMSE of 16.5% (n = 784). These results suggest that, despite their global coverage, the existing AGB approaches fail to capture fine-scale ΔAGB dynamics required for management or monitoring applications in this specific setting.

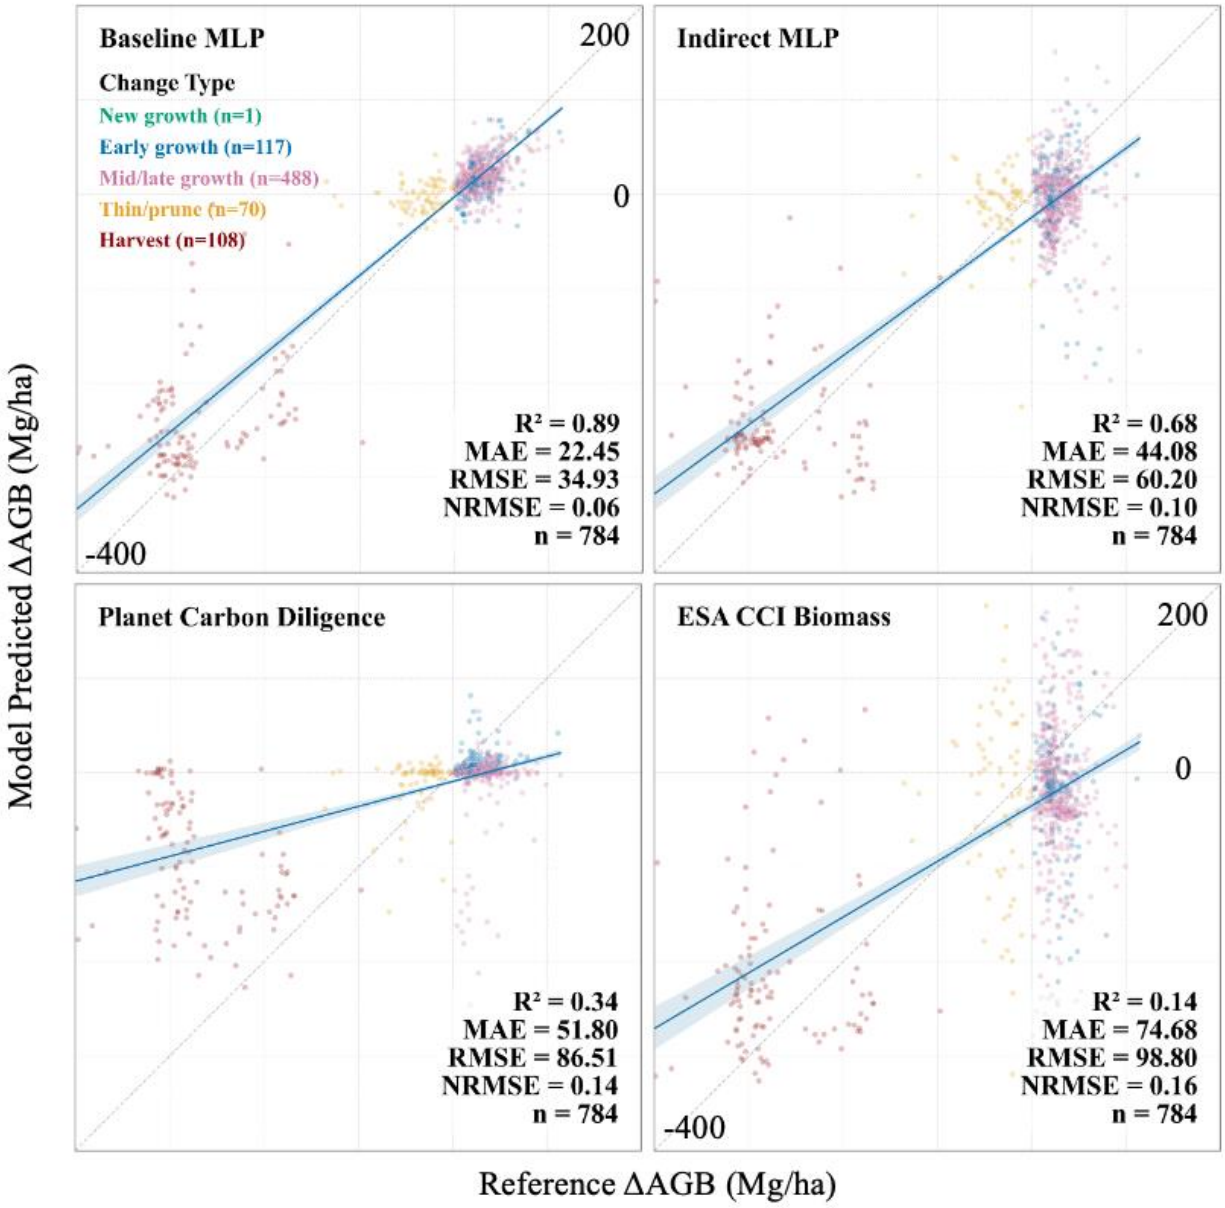

Figure 4. *Model-predicted vs reference change in ΔAGB (Mg/ha) across model outputs, including our Baseline direct MLP, Indirect MLP, Planet Carbon Diligence data product, and ESA CCI Biomass data product. Points are colored by change type; the gray dashed line marks the 1:1 reference and the blue line with shading is the least-squares fit with 95% CI. All panels share the same axes and number of observations (-400 to 200 Mg/ha; n = 784). Inset boxes report R^2, MAE, RMSE, and NRMSE. Note: due to data availability and our data augmentation approach, the number of observations is significantly reduced across all models compared to the MOE and baseline modeling approach.*

3.3. Uncertainty analysis

The direct MOE produced per-pixel predictive uncertainty estimates derived from the GNLL. A formal reliability analysis comparing empirical coverage against nominal confidence levels revealed systematic overconfidence in the MOE: at a nominal 90% confidence interval, empirical coverage was 50.2% overall, indicating that predictive intervals are roughly half as wide as the true error distribution requires. Calibration performance was strongly class-dependent in the MOE (SI Figure 8). New growth transitions were nearly perfectly calibrated, achieving 92.7% empirical coverage at a nominal 90% interval–consistent with the strong and unambiguous bare-earth-to-canopy signal that defines this class. In contrast, mid/late growth (40.5%), harvest (39.3%), and thinning/pruning (33.2%) showed systematic overconfidence. Critically, the rank ordering of uncertainty across classes was preserved: the model assigned lowest predictive SD to new growth (2.9 Mg/ha) and highest to harvest (18.4 Mg/ha), matching the residual SD pattern (2.2 Mg/ha and 48.5 Mg/ha respectively; SI Table 4). This consistent ordering indicates that although absolute calibration is imperfect, the model's uncertainty signal captures meaningful variation in prediction confidence across forest change types.

3.4. Parcel-level change

Aggregating pixel-level predictions to the parcel scale demonstrates the operational utility of the MOE framework for forest management applications. Across the 91 parcels in the independent test area, MOE predictions showed strong agreement with inventory reference values at the parcel scale. A paired Wilcoxon signed-rank test comparing parcel-level predictions to inventory measurements found that 67% of parcels (61/91) showed no statistically significant difference ($p > 0.05$) between MOE predictions and parcel-level reference values. Mean bias across all parcels was 0.12 Mg/ha, with nearly equal numbers of parcels showing positive (n = 16) and negative (n = 14) significant bias, indicating no systematic directional error. The following examples illustrate model behavior across two contrasting management scenarios–harvest and regrowth–that represent the primary change types relevant to plantation carbon monitoring.

Within Parcel 20081, ARAUCO collected data from 18 inventory plots in March, 2017 before harvesting 28.08 ha of Monterey Pine in 2018, with the aggregated inventory plots indicating a total loss of -8,404.2 Mg (95% CI [-9,429.3, -7379.0] Mg; mean = -299.3 Mg/ha; Figure 5). The MOE raster estimated a total harvest of -6,909.4 Mg AGB (mean -246.1 Mg/ha), which falls just outside the inventory 95% CI and is a total difference of -1,494.7 Mg (17.8% of the inventory total). Within the parcel, predictive MC uncertainty (t = 30 MC passes) averaged 16.7 Mg/ha across parcel pixels; when scaled to the total parcel area, this corresponds to σ = 111 Mg (95% CI $\pm$400 Mg around the MOE total). In comparison, the empirical residual standard deviation within the parcel was 60.1 Mg/ha (SE = 14.2 Mg/ha). These results suggest that while the model captures the direction of biomass loss

during harvest, it systematically underestimates the magnitude, and the MC uncertainty substantially underestimates the true predictive error at the parcel level.

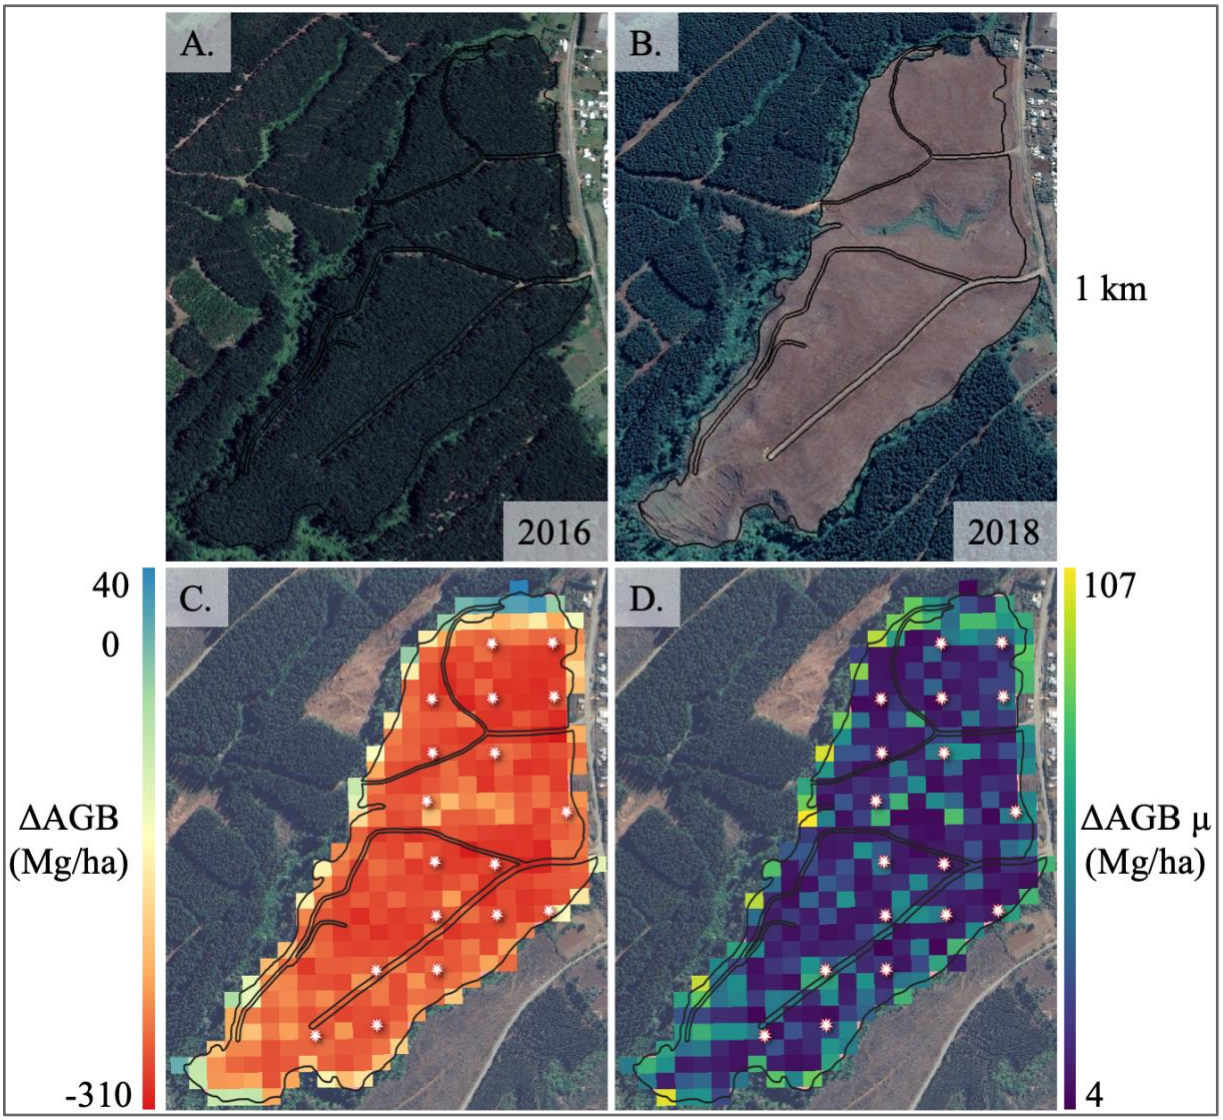

Figure 5. *AGB loss between 2016 and 2018 within Parcel 20081 inside the independent testing area. Panels show (A) very high-resolution Google Earth imagery from 2016, (B) Google Earth imagery from 2018, (C) predicted ΔAGB from the MOE model, and (D) associated uncertainty. The parcel, delineated by ARAUCO property boundaries, experienced a harvest event in 2018, resulting in a total inventory-estimated AGB loss of -8,404.2 Mg AGB, and a model-predicted loss of -6,909.4 Mg AGB. Higher uncertainty is associated with parcel boundaries. The white stars indicate repeat visit inventory plots from 2016 and 2018.*

Parcel 31539 was planted in the year 2013 and inventoried in June, 2018 and again in September, 2019. Between these years, the model captured strong biomass accumulation associated with mid-age stand growth (Figure 6). At the 52 revisited inventory plots across the 116.4 ha parcel, there was an average observed gain of 37.1 Mg/ha and a parcel-level accumulation of 4,322.0 Mg (95% CI [4,058.0, 4,586.0] Mg). The MOE output predicted a mean gain of 33.0 Mg/ha, equivalent to a total of 3,839.3 Mg, which underestimates the inventory total by 482.7 (11.2%). Note that the composite image corresponding to the 2019 inventory date was assembled from imagery acquired in December, 2019 through March, 2020, rather than earlier in 2019, as harvesting activity within the parcel occurred prior to this window. Predictive MC uncertainty (t = 30) averaged 12.4 Mg/ha across parcel pixels; when propagated to the full area, this corresponds to ~199.4 Mg of uncertainty (95% CI, 391 Mg around the MOE total). In comparison, the empirical residual standard deviation with the parcel was 7.3 Mg/ha (SE = 1.0 Mg/ha), yielding a parcel-level residual uncertainty of 135 Mg (95% CI, 264 Mg). These results demonstrate the operational utility of the MOE framework for regrowth monitoring, capturing the direction and general magnitude of biomass accumulation with a model underestimation of 11.2%.

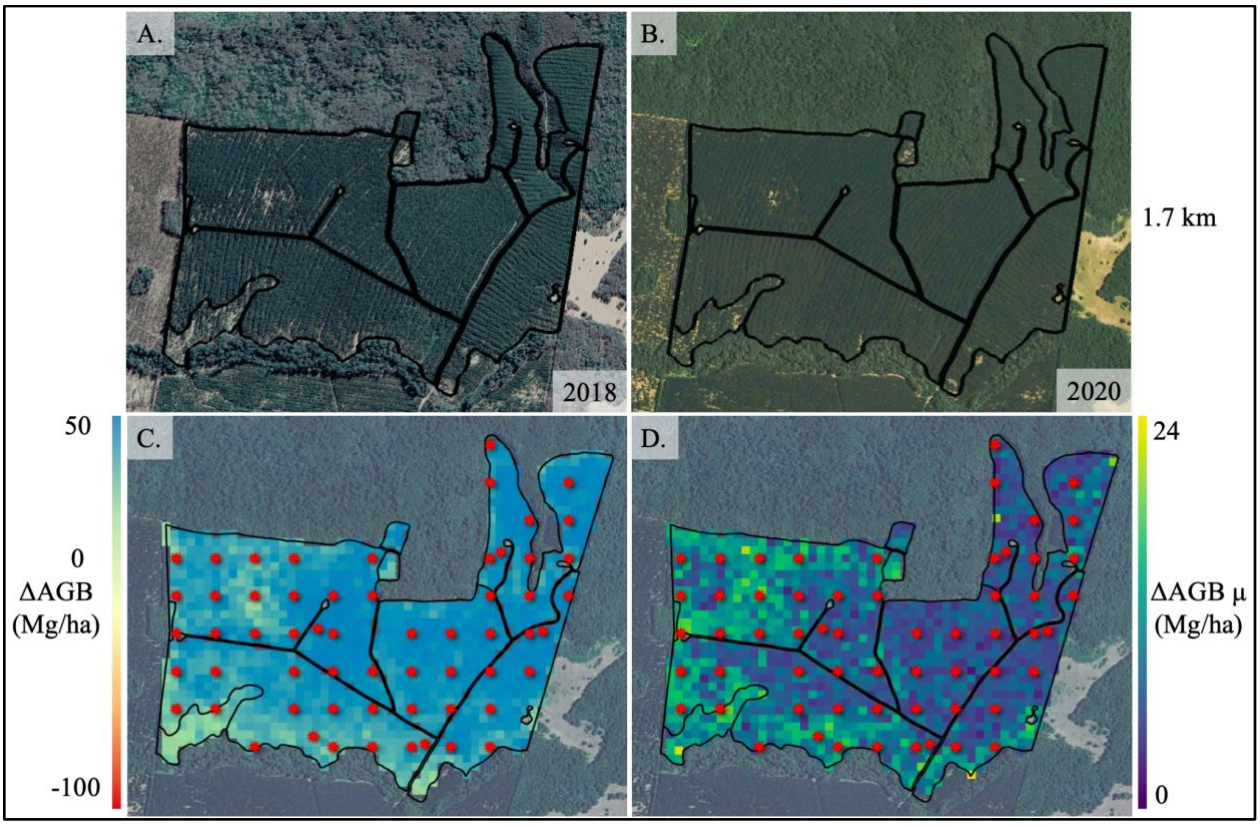

Figure 6. *AGB gain within Parcel 31539 between 2018 and 2020. The 52 repeat inventory plots (red stars) are highlighted within the parcel boundaries. Panels show (A) very high-resolution Google Earth imagery from 2018, (B) Google Earth imagery from 2020 (to capture inventory points from September*

2019), (C) predicted ΔAGB from the MOE model, and (D) associated uncertainty. In the ΔAGB raster, red represents biomass loss and blue represents biomass gain, with a color scale ranging from -100 to +50 Mg/ha.

4. Discussion

4.1. The value of directly modeling changes in AGB

Our results show that direct estimation of ΔAGB with a Mixture-of-Experts (MOE) framework substantially improves predictive performance over traditional indirect difference approaches. The MOE approach achieved an R^2 = 0.90 and RMSE = 26.89 Mg/ha, outperforming a traditional indirect model (R^2 = 0.58, RMSE = 56.07 Mg/ha). This result is consistent with the broader distinction in the literature between direct estimation of biomass change and indirect estimation obtained by differencing two biomass predictions. McRoberts et al. (2015) note that direct estimation is often considered preferable because it requires only a single set of prediction errors to be accommodated, although they emphasize that the relative precision of direct and indirect approaches is context dependent. In our study, the strong performance of a baseline direct MLP (R^2 = 0.87, RMSE = 31.06 Mg/ha) further suggests that much of the improvement arises from treating ΔAGB as the response variable itself rather than inferring change from two separate AGB predictions.

Although much of the improvement appears to stem from modeling ΔAGB directly, the additional gain achieved by the MOE suggests that model structure also matters. Prior work has emphasized that biomass change is difficult to model because responses may be both positive and negative, and that separating loss and gain processes may be beneficial (Bollandsås et al., 2018; Næsset et al., 2013). Our classifier-gated MOE is consistent with this logic, as it allows different experts to specialize across distinct change processes. In doing so, the model is better able to represent the contrasting signatures of harvest, regrowth, and intermediate change than a single regression model.

While this study demonstrates direct ΔAGB estimation in managed plantation forests, the approach has broader relevance for natural forest systems where national forest inventory (NFI) data are available at scale. To our knowledge, applications of direct change estimation in natural forests remain largely unexplored–Puliti et al. (2021) is a notable exception, using Norwegian NFI data, though limited to optical datasets. Several countries maintain dense NFI programs with repeat measurements that could support this modeling approach at national scales (Lister et al., 2020; McRoberts et al., 2010). Extending this framework to such systems, potentially through transfer learning from data-rich plantation environments, represents a promising direction for broadening the geographic and ecological scope of direct ΔAGB estimation.

4.2. Multi-sensor data fusion for direct change detection

Previous projects have directly estimated ΔAGB using a single-sensor approach, primarily using optical datasets such as Landsat and Sentinel-2 (Arévalo et al., 2023; Healey et al., 2006; Puliti et al., 2021). While effective, these approaches are constrained by cloud cover and limited structural sensitivity. Our results show that combining multiple sensors, particularly SAR and optical, enhances predictive power (SI, Figure 7). Fusing Sentinel-1, Sentinel-2, and ALOS PALSAR provided the most accurate results across all change types, supporting evidence that multi-sensor fusion can capture a more comprehensive picture of forest dynamics (Duncanson et al., 2020; Joshi et al., 2016; Zhang, 2010).

Harmonizing data temporally proved to be challenging, particularly when aligning temporal windows with continuous forestry operations. Forest inventories in Chilean plantations are collected year-round, whereas various sensors have their own acquisition cycles, and cloud cover further constrains optical image availability. To minimize this misalignment, we restricted our reference labels and predictor variables to a narrow December-March window, which corresponds to the period of least cloud cover and the fewest inventory measurements. We also implemented rules for pairing remote sensing data with inventory years, sampling imagery from before the first inventory and after the second inventory to better reflect the timing of observed changes. Despite this, we encountered instances where sensor timing was imperfect: for example, Parcels 30454 and 30151 were harvested some time during the Chilean summer months of 2019, but ALOS PALSAR acquisitions showed standing forest while Sentinel-2 data indicated cleared areas. These mismatches introduce noise to the training labels and likely contribute to residual uncertainty in our predictions, highlighting the need for better temporal harmonization frameworks in future work.

4.3. Local calibration and direct estimation outperform existing AGB products

Our comparisons with external AGB products highlight two distinct limitations of existing datasets for biomass change monitoring. The first is methodological: these products generally do not estimate ΔAGB directly, but instead require change to be inferred by differencing separate biomass maps using the indirect approach. The second is empirical: they are not always well calibrated to local forest conditions, species, management practices, and region-specific relationships between remote sensing signals and biomass change.

These two issues are highlighted when comparing the three model settings: external products that are neither locally calibrated nor designed for direct estimation of change, our locally calibrated indirect model, and our locally calibrated direct MOE. Within the independent test set, the locally calibrated indirect model (RMSE = 64.51 Mg/ha, NRMSE = 11%) outperformed Planet's Forest Carbon Diligence estimates of change (RMSE = 84.4 Mg/ha, NRMSE = 14.2%) and ESA's CCI

Biomass product (RMSE = 98.8 Mg/ha, NRMSE = 16.4%). The contrast between the external products and our indirect model indicates that local calibration alone provides a substantial improvement. Even without directly modeling change, the locally trained indirect model was considerably better aligned with the independent reference data than either Planet's Forest Carbon Diligence estimates of change, or ESA's CCI Biomass product, even though the three share similar modeling approaches and remote sensing datasets. This result underscores the importance of region-specific training data in plantation forests, where species composition, stand structure, harvest timing, and local management dynamics strongly influence how biomass change is captured in optical and SAR observations.

The comparison between the locally calibrated indirect model and the direct MOE indicates that direct estimation provides an additional improvement once local calibration is already in place. Together, these comparisons suggest that the underperformance of global and commercial products in this setting reflects both limited local calibration and reliance on indirect differencing. The scatterplot comparisons (Figure 3) reinforce this interpretation: Planet estimates tended to compress the magnitude of change toward zero across change types, whereas ESA exhibited high variance and weak alignment with observed harvest and growth dynamics. By contrast, both locally trained models were better matched to the reference data, with the MOE showing the tightest fit and lowest spread.

4.4. Predictive uncertainty and the challenge of model calibration

Our MOE approach systematically underestimates predictive uncertainty relative to observed residuals–a pattern consistent with known limitations of heteroscedastic GNLL training and MC dropout in settings where label noise is not fully captured during training (Seitzer et al., 2022). In our case, this likely reflects temporal misalignment between inventory measurement dates and annual remote sensing composites (in the case of ALOS PALSAR), which potentially introduces residual error that the model cannot anticipate at inference. At a nominal 90% CI, the MOE achieves only 50.2% empirical coverage overall, indicating substantial overconfidence in absolute terms. Despite this, the MOE's uncertainty estimates retain practical value in two ways. First, calibration is class-structured in a physically interpretable way: new growth transitions, defined by a strong bare-earth-to-canopy signal, are nearly perfectly calibrated (92.7% empirical coverage at 90% nominal). Thinning/pruning and mid/late growth classes, where spectral signatures are subtle and temporal misalignment is more likely, show the greatest overconfidence. This pattern is consistent across both predictive and residual uncertainty (SI Figure 8), confirming that the model's uncertainty signal reflects real difference in prediction difficulty rather than arbitrary variation. Second, at the parcel scale, parcel-level aggregation of pixel predictions demonstrated practical utility across 91 independent test parcels, with 67% showing no statistically significant difference between MOE predictions and inventory reference

values (Wilcoxon signed-rank test, $p > 0.05$) and a mean bias of 0.12 Mg/ha across all parcels. The two case study parcels (Section 3.4) illustrate how model behavior manifests in specific operational contexts--a harvest event and a mid-age regrowth stand--but are not intended to be representative of uncertainty behavior across all change types or parcel configurations.

These results suggest that the current uncertainty estimates are suitable for relative confidence flagging–identifying where predictions should be interpreted with greater caution–and for directing future sampling campaigns toward high-uncertainty areas. Furthermore, the high degree of certainty associated with new growth suggests that the current framework may already be suitable for formal probabilistic reporting in reforestation monitoring contexts, where bare-earth-to-canopy transitions produce strong signals. However, broader deployment for Tier 2 or Tier 3 MRV reporting requires improved absolute calibration. Post-hoc calibration via temperature scaling (Guo et al., 2017; Kuleshov et al., 2018) represents the most direct path forward, and should be explored in future research.

4.5. Implications for forest management and carbon monitoring

Direct AGB estimation with explicit uncertainty provides operational value for forest managers and policy relevance for MRV systems. At the parcel scale, the MOE captured the direction and approximate magnitude of real-world management events. In Parcel 20081, the MOE correctly identified substantial biomass loss but underestimated the inventory total by 17.8% (inventory = -8,404.2 Mg; MOE = -6,909.4 Mg), while in Parcel 31539 it accurately classified regrowth but underestimated total accumulation by ~11.2%. Parcel-level aggregation can reduce pixel-level noise, though uncertainty remains high in regrowth contexts where mid-rotation stand dynamics are underrepresented in training data. These examples demonstrate how model outputs can support management decisions, including quantifying harvest removals, validating compliance with harvest reporting requirements and supporting Tier 2 and 3 greenhouse gas inventory frameworks, and assessing regrowth trajectories.

For MRV and REDD+ frameworks, direct estimation with per-pixel predictive uncertainty has the potential to provide transparent benchmarks for carbon accounting. Spatial patterns in predictive uncertainty can guide future sampling campaigns and improve monitoring system design by identifying where accurate data needs are highest. This framework performs best for abrupt, structurally unambiguous transitions–particularly harvest and new growth–where SAR-optical fusion provides clear and consistent signals. Harvest and deforestation events are directly relevant to REDD+ MRV frameworks, where accurate quantification of emissions from forest loss is required. New growth transitions, showing near-perfect uncertainty calibration, are well-suited for reforestation verification under voluntary carbon markets, where crediting of carbon removals requires verifiable

evidence of biomass accumulation. Conversely, performance is more limited for gradual or spectrally subtle changes such as thinning and pruning, where temporal misalignment or remote sensing composites potentially introduce label noise that can reduce accuracy and calibration quality. The open-source framework we developed is also scalable: transfer learning approaches allow models trained on large, high-quality datasets in data-rich areas to be fine-tuned with limited local data in regions where the same reference data are not available (Ma et al., 2024). In the context of AGB estimation, this capability is particularly useful, as many areas lack extensive forest inventory data (Amitrano et al., 2023; Astola et al., 2021; L. Chen et al., 2023).

A major strength of this study lies in the use of large, repeat forest inventory plots, which provide a robust foundation for monitoring AGB and ΔAGB. These data capture forest structure and dynamics across multiple stages of growth, disturbance, and recovery, offering valuable ground-truth data that is difficult to obtain from remote sensing alone and rarely available at scale. Such information is essential to track regeneration rates, evaluate conservation and restoration efforts, and monitor progress toward global carbon sequestration goals (Chazdon & Guariguata, 2016). Furthermore, our approach might be useful for identifying and tracking subtle forms of degradation (Bullock et al., 2020). This framework's reliance on freely available satellite data and an openly transferable architecture makes it adaptable for monitoring in other managed forest systems, and future work can explore the effectiveness in natural forests.

4.6. Limitations and future improvements

Despite these advances, several limitations remain. The ARAUCO inventory data, while extensive, varies in focus and collection methods across years and regions. Because the dataset was primarily collected to support timber production, data collection focused on specific forest stand ages, leading to an overrepresentation of older early-growth and mid-growth forests and an underrepresentation of very early growth stages and mature forests. This imbalance introduces inherent biases that, while we tried to address through data augmentation, may limit the model's ability to generalize across all forest stand ages.

Incorporating more sophisticated models that can account for both the temporal and spatial context may enhance the ability to capture temporal dependencies in biomass dynamics, particularly for long-term monitoring (Crisóstomo de Castro Filho et al., 2020; Gao, 2019). Advanced modeling approaches are notoriously data-hungry, and the size of this reference dataset might be suitable for an approach that leverages both the spatial and temporal contexts, such an LSTM or Transformer. Additionally, integration with broader-scale climate and land-use data could provide insights into the general trends of biomass accumulation and forest health that might be impacted by conditions such as megadrought and extreme precipitation events (González et al., 2018), as well as the interactions

between plantation forests and natural forests within the broader landscape (Alvarez-Garreton et al., 2019).

While our framework provides per-pixel predictive uncertainty through GNLL and MC dropout, a formal reliability analysis confirmed systematic overconfidence across most change classes (Section 4.4). Addressing this through post-hoc calibration techniques such as temperature scaling is a priority for future work, particularly before deploying uncertain estimates for formal Tier 2 or Tier 3 MRV reporting. More broadly, uncertainty decomposition, class-structured calibration, and transfer learning across forest types warrant future exploration.

5. Conclusion

By integrating a Mixture-of-Experts architecture with multisensor fusion and an unusually large repeat forest inventory dataset, we demonstrate that direct estimation of ΔAGB can achieve substantially higher accuracy than traditional indirect differencing approaches – reducing RMSE by 57% relative to a comparable indirect model and outperforming widely used global products. ARAUCO's extensive plantation inventory addresses the critical scarcity of repeat ground reference data, while the classifier-gated MOE framework addresses the challenge of capturing both gradual growth and abrupt loss within a single framework. At both the pixel- and parcel-scale, the model accurately reproduced real-world management events including harvest and regrowth, with per-pixel predictive uncertainty estimates that are class-structured in a physically interpretable way that establishes a foundation for future calibration work.

These results highlight the value of region-specific repeat inventory data and temporally aligned multi-sensor fusion for forest carbon monitoring – a combination that global products currently cannot replicate at fine spatial scales. The open-source, transferable architecture suggests potential for broader deployment through transfer learning in data-limited regions. Future work should explore application in natural forest systems, integration of spatiotemporal modeling approaches, and improved uncertainty calibration to further close the gap between predictive and realized error. Together, these advances offer a scalable and transparent pathway toward the high-accuracy ΔAGB monitoring required for climate mitigation, sustainable forest management, and robust MRV frameworks.

Acknowledgements
BL was supported by NASA MUREP Fellowship (80NSSC20K1475) and generous support from the Schmidt Family Foundation. RH and BL acknowledge funding from the Arnhold UC Santa Barbara-Conservation International Climate Solutions Collaborative. AS was supported by the NASA Terrestrial Ecology Program under A.60 GEDI ST (80NSSC24K0603 P00003) and by the Early Career Investigators Program (80NSSC26K0143). TF was supported by A.60 GEDI ST (80NSSC24K0603 P00003). Planet Carbon Diligence data were provided by Planet Labs PBC; we thank David Marvin for facilitating access.

Author Contribution
B.L. and R.H. conceived the study with input from A.R., N.T., A.S., and T.F.. B.L., A.R., and R.H. led the investigation and methodology development, with additional methodological contributions from N.T. and A.S.. B.L. performed the formal analysis, software development, visualization, and validation, with software contributions from R.H.. G.F.O., P.R.A., P.M.Q., and B.L. contributed to data curation. B.L. and R.H. managed project administration and resources. R.H. acquired funding. B.L., A.R., and R.H. wrote the original draft. All authors contributed to review and editing.

Data availability
Forest inventory data used in this study were collected by ARAUCO and are proprietary; requests for access should be directed to ARAUCO. All remote sensing inputs are publicly available through the Copernicus programme (Sentinel-1, Sentinel-2), the USGS Earth Explorer (Landsat), and freely accessible through Google Earth Engine, which the code provides. Pre-trained model weights, model architecture, and model outputs are available from the corresponding author upon reasonable request.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Claude Code (Anthropic) in order to assist with code refactoring and software development tasks. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

6. References

A. Ramezan, C., A. Warner, T., & E. Maxwell, A. (2019). Evaluation of Sampling and Cross-Validation Tuning Strategies for Regional-Scale Machine Learning Classification. *Remote Sensing*, *11*(2), Article 2. https://doi.org/10.3390/rs11020185

Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). *Optuna: A Next-generation Hyperparameter Optimization Framework* (arXiv:1907.10902). arXiv. https://doi.org/10.48550/arXiv.1907.10902

Alvarez-Garreton, C., Lara, A., Boisier, J. P., & Galleguillos, M. (2019). The Impacts of Native Forests and Forest Plantations on Water Supply in Chile. *Forests*, *10*(6), Article 6. https://doi.org/10.3390/f10060473

Amitrano, D., Giacco, G., Marrone, S., Pascarella, A. E., Rigiroli, M., & Sansone, C. (2023). Forest Aboveground Biomass Estimation Using Machine Learning Ensembles: Active Learning Strategies for Model Transfer and Field Sampling Reduction. *Remote Sensing*, *15*(21), Article 21. https://doi.org/10.3390/rs15215138

ARAUCO. (2022). *2022 ARAUCO Integrated Report*. https://na.arauco.com/en/resources/download/2022_ARAUCO_Sustainability_Report

Arévalo, P., Baccini, A., Woodcock, C. E., Olofsson, P., & Walker, W. S. (2023). Continuous mapping of aboveground biomass using Landsat time series. *Remote Sensing of Environment*, *288*, 113483. https://doi.org/10.1016/j.rse.2023.113483

Astola, H., Seitsonen, L., Halme, E., Molinier, M., & Lönnqvist, A. (2021). Deep Neural Networks

with Transfer Learning for Forest Variable Estimation Using Sentinel-2 Imagery in Boreal Forest. *Remote Sensing*, *13*(12), Article 12. https://doi.org/10.3390/rs13122392

Atkinson, P. M., & Foody, G. M. (2002). Uncertainty in Remote Sensing and GIS: Fundamentals. In *Uncertainty in Remote Sensing and GIS* (pp. 1–18). John Wiley & Sons, Ltd. https://doi.org/10.1002/0470035269.ch1

Baccini, A., Walker, W., Carvalho, L., Farina, M., Sulla-Menashe, D., & Houghton, R. A. (2017). Tropical forests are a net carbon source based on aboveground measurements of gain and loss. *Science*, *358*(6360), 230–234. https://doi.org/10.1126/science.aam5962

Bao, N., Lechner, A. M., Fletcher, A., Mulligan, D., Mellor, A., & Bai, Z. (2012). Comparison of relative radiometric normalization methods using pseudo-invariant features for change detection studies in rural and urban landscapes. *Journal of Applied Remote Sensing*, *6*(1), 063578. https://doi.org/10.1117/1.JRS.6.063578

Bollandsås, O. M., Ene, L. T., Gobakken, T., & Næsset, E. (2018). Estimation of biomass change in montane forests in Norway along a 1200 km latitudinal gradient using airborne laser scanning: A comparison of direct and indirect prediction of change under a model-based inferential approach. *Scandinavian Journal of Forest Research*, *33*(2), 155–165. https://doi.org/10.1080/02827581.2017.1338354

Bollandsås, O. M., Gregoire, T. G., Næsset, E., & Øyen, B.-H. (2013). Detection of biomass change in a Norwegian mountain forest area using small footprint airborne laser scanner data. *Statistical Methods & Applications*, *22*(1), 113–129. https://doi.org/10.1007/s10260-012-0220-5

Bullock, E. L., Woodcock, C. E., & Olofsson, P. (2020). Monitoring tropical forest degradation using spectral unmixing and Landsat time series analysis. *Remote Sensing of Environment, Time Series Analysis with High Spatial Resolution Imagery*, *238*, 110968. https://doi.org/10.1016/j.rse.2018.11.011

Caraballo-Vega, J. A., Carroll, M. L., Neigh, C. S. R., Wooten, M., Lee, B., Weis, A., Aronne, M., Alemu, W. G., & Williams, Z. (2023). Optimizing WorldView-2, -3 cloud masking using machine learning approaches. *Remote Sensing of Environment*, *284*, 113332. https://doi.org/10.1016/j.rse.2022.113332

Carrasco, G., Almeida, A. C., Falvey, M., Olmedo, G. F., Taylor, P., Santibañez, F., & Coops, N. C. (2022). Effects of climate change on forest plantation productivity in Chile. *Global Change Biology*, *28*(24), 7391–7409. https://doi.org/10.1111/gcb.16418

Chave, J., Réjou-Méchain, M., Búrquez, A., Chidumayo, E., Colgan, M. S., Delitti, W. B. C., Duque, A., Eid, T., Fearnside, P. M., Goodman, R. C., Henry, M., Martínez-Yrízar, A., Mugasha, W. A., Muller-Landau, H. C., Mencuccini, M., Nelson, B. W., Ngomanda, A., Nogueira, E. M., Ortiz-Malavassi, E., … Vieilledent, G. (2014). Improved allometric models to estimate the aboveground biomass of tropical trees. *Global Change Biology*, *20*(10), 3177–3190. https://doi.org/10.1111/gcb.12629

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, *16*, 321–357. https://doi.org/10.1613/jair.953

Chazdon, R. L., & Guariguata, M. R. (2016). Natural regeneration as a tool for large-scale forest restoration in the tropics: Prospects and challenges. *Biotropica*, *48*(6), 716–730. https://doi.org/10.1111/btp.12381

Chen, L., Lin, H., Long, J., Liu, Z., Yang, P., & Zhang, T. (2023). Evaluating the Transferability of Spectral Variables and Prediction Models for Mapping Forest Aboveground Biomass Using Transfer Learning Methods. *Remote Sensing*, *15*(22), Article 22. https://doi.org/10.3390/rs15225358

Chen, L., Wang, Y., Ren, C., Zhang, B., & Wang, Z. (2019). Assessment of multi-wavelength SAR and multispectral instrument data for forest aboveground biomass mapping using random forest kriging. *Forest Ecology and Management*, *447*, 12–25. https://doi.org/10.1016/j.foreco.2019.05.057

Chen, Q., Vaglio Laurin, G., & Valentini, R. (2015). Uncertainty of remotely sensed aboveground biomass over an African tropical forest: Propagating errors from trees to plots to pixels. *Remote Sensing of Environment*, *160*, 134–143. https://doi.org/10.1016/j.rse.2015.01.009

Cordero, R. R., Feron, S., Damiani, A., Carrasco, J., Karas, C., Wang, C., Kraamwinkel, C. T., & Beaulieu, A. (2024). Extreme fire weather in Chile driven by climate change and El Niño–Southern Oscillation (ENSO). *Scientific Reports*, *14*(1), Article 1. https://doi.org/10.1038/s41598-024-52481-x

Crisóstomo de Castro Filho, H., Abílio de Carvalho Júnior, O., Ferreira de Carvalho, O. L., Pozzobon de Bem, P., dos Santos de Moura, R., Olino de Albuquerque, A., Rosa Silva, C., Guimarães

Ferreira, P. H., Fontes Guimarães, R., & Trancoso Gomes, R. A. (2020). Rice Crop Detection Using LSTM, Bi-LSTM, and Machine Learning Models from Sentinel-1 Time Series. *Remote Sensing*, *12*(16), Article 16. https://doi.org/10.3390/rs12162655

Drusch, M., Del Bello, U., Carlier, S., Colin, O., Fernandez, V., Gascon, F., Hoersch, B., Isola, C., Laberinti, P., Martimort, P., Meygret, A., Spoto, F., Sy, O., Marchese, F., & Bargellini, P. (2012). Sentinel-2: ESA's Optical High-Resolution Mission for GMES Operational Services. *Remote Sensing of Environment, The Sentinel Missions - New Opportunities for Science*, *120*, 25–36. https://doi.org/10.1016/j.rse.2011.11.026

Duncanson, L., Neuenschwander, A., Hancock, S., Thomas, N., Fatoyinbo, T., Simard, M., Silva, C. A., Armston, J., Luthcke, S. B., Hofton, M., Kellner, J. R., & Dubayah, R. (2020). Biomass estimation from simulated GEDI, ICESat-2 and NISAR across environmental gradients in Sonoma County, California. *Remote Sensing of Environment*, *242*, 111779. https://doi.org/10.1016/j.rse.2020.111779

ESA. (2022). *Copernicus DEM Product Handbook*. https://spacedata.copernicus.eu/documents/20123/122407/GEO1988-CopernicusDEM-SPE-002_ProductHandbook_I5.0+%281%29.pdf/706ee17d-2cce-f1fa-a73e-1686d28f09dd?t=1679657087883

Gao, S. (2019). Deforestation Prediction Using Time Series and LSTM. *2019 International Conference on Information Technology and Computer Application (ITCA)*, 95–99. https://doi.org/10.1109/ITCA49981.2019.00029

Gibbs, H. K., Brown, S., Niles, J. O., & Foley, J. A. (2007). Monitoring and estimating tropical forest carbon stocks: Making REDD a reality. *Environmental Research Letters*, *2*(4), 045023. https://doi.org/10.1088/1748-9326/2/4/045023

Gneiting, T., & Raftery, A. E. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, *102*(477), 359–378. https://doi.org/10.1198/016214506000001437

González, M. E., Gómez-González, S., Lara, A., Garreaud, R., & Díaz-Hormazábal, I. (2018). The 2010–2015 Megadrought and its influence on the fire regime in central and south-central Chile. *Ecosphere*, *9*(8), e02300. https://doi.org/10.1002/ecs2.2300

Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., & Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment, Big Remotely Sensed Data: Tools, Applications and Experiences*, *202*, 18–27. https://doi.org/10.1016/j.rse.2017.06.031

Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On Calibration of Modern Neural Networks. *Proceedings of the 34th International Conference on Machine Learning*, 1321–1330. https://proceedings.mlr.press/v70/guo17a.html

Harris, N. L., Gibbs, D. A., Baccini, A., Birdsey, R. A., de Bruin, S., Farina, M., Fatoyinbo, L., Hansen, M. C., Herold, M., Houghton, R. A., Potapov, P. V., Suarez, D. R., Roman-Cuesta, R. M., Saatchi, S. S., Slay, C. M., Turubanova, S. A., & Tyukavina, A. (2021). Global maps of twenty-first century forest carbon fluxes. *Nature Climate Change*, *11*(3), Article 3.

https://doi.org/10.1038/s41558-020-00976-6

Hawthorne, S., Boissière, M., Felker, M. E., & Atmadja, S. (2016). Assessing the Claims of Participatory Measurement, Reporting and Verification (PMRV) in Achieving REDD+ Outcomes: A Systematic Review. *PLOS ONE*, *11*(11), e0157826. https://doi.org/10.1371/journal.pone.0157826

Healey, S. P., Yang, Z., Cohen, W. B., & Pierce, D. J. (2006). Application of two regression-based methods to estimate the effects of partial harvest on forest structure using Landsat data. *Remote Sensing of Environment*, *101*(1), 115–126. https://doi.org/10.1016/j.rse.2005.12.006

Herold, M., Román-Cuesta, R. M., Mollicone, D., Hirata, Y., Van Laake, P., Asner, G. P., Souza, C., Skutsch, M., Avitabile, V., & MacDicken, K. (2011). Options for monitoring and estimating historical carbon emissions from forest degradation in the context of REDD+. *Carbon Balance and Management*, *6*(1), 13. https://doi.org/10.1186/1750-0680-6-13

Hurteau, M., & North, M. (2009). Fuel treatment effects on tree-based forest carbon storage and emissions under modeled wildfire scenarios. *Frontiers in Ecology and the Environment*, *7*(8), 409–414. https://doi.org/10.1890/080049

Jamaluddin, I., Chen, Y.-N., & Fan, K.-C. (2024). Spatial–Spectral–Temporal Deep Regression Model With Convolutional Long Short-Term Memory and Transformer for the Large-Area Mapping of Mangrove Canopy Height by Using Sentinel-1 and Sentinel-2 Data. *IEEE Transactions on Geoscience and Remote Sensing*, *62*, 1–17. https://doi.org/10.1109/TGRS.2024.3362788

Joshi, N., Baumann, M., Ehammer, A., Fensholt, R., Grogan, K., Hostert, P., Jepsen, M. R., Kuemmerle, T., Meyfroidt, P., Mitchard, E. T. A., Reiche, J., Ryan, C. M., & Waske, B. (2016). A Review of the Application of Optical and Radar Remote Sensing Data Fusion to Land Use Mapping and Monitoring. *Remote Sensing*, *8*(1), Article 1. https://doi.org/10.3390/rs8010070

Kellndorfer, J. M., Walker, W. S., LaPoint, E., Kirsch, K., Bishop, J., & Fiske, G. (2010). Statistical fusion of lidar, InSAR, and optical remote sensing data for forest stand height characterization: A regional-scale method based on LVIS, SRTM, Landsat ETM+, and ancillary data sets. *Journal of Geophysical Research: Biogeosciences*, *115*(G2). https://doi.org/10.1029/2009JG000997

Kennedy, R. E., Yang, Z., & Cohen, W. B. (2010). Detecting trends in forest disturbance and recovery using yearly Landsat time series: 1. LandTrendr — Temporal segmentation algorithms. *Remote Sensing of Environment*, *114*(12), 2897–2910. https://doi.org/10.1016/j.rse.2010.07.008

Kuleshov, V., Fenner, N., & Ermon, S. (2018). *Accurate Uncertainties for Deep Learning Using Calibrated Regression* (arXiv:1807.00263). arXiv. https://doi.org/10.48550/arXiv.1807.00263

Lechner, A. M., Foody, G. M., & Boyd, D. S. (2020). Applications in Remote Sensing to Forest Ecology and Management. *One Earth*, *2*(5), 405–412. https://doi.org/10.1016/j.oneear.2020.05.001

Li, Y., Li, M., Li, C., & Liu, Z. (2020). Forest aboveground biomass estimation using Landsat 8 and Sentinel-1A data with machine learning algorithms. *Scientific Reports*, *10*(1), Article 1. https://doi.org/10.1038/s41598-020-67024-3

Lister, A. J., Andersen, H., Frescino, T., Gatziolis, D., Healey, S., Heath, L. S., Liknes, G. C., McRoberts, R., Moisen, G. G., Nelson, M., Riemann, R., Schleeweis, K., Schroeder, T. A., Westfall, J., & Wilson, B. T. (2020). Use of Remote Sensing Data to Improve the Efficiency of National Forest Inventories: A Case Study from the United States National Forest Inventory. *Forests*, *11*(12), 1364. https://doi.org/10.3390/f11121364

Liu, Y., Wu, C., Guan, M., & Wang, J. (2025). MixtureRS: A Mixture of Expert Network Based Remote Sensing Land Classification. *Remote Sensing*, *17*(14), 2494. https://doi.org/10.3390/rs17142494

Liu, Z., Ma, Q., Zhang, T., Zhao, S., Gao, X., Sun, T., & Dai, Y. (2025). Quantitative modeling and uncertainty estimation for small-sample LIBS using Gaussian negative log-likelihood and monte carlo dropout methods. *Optics & Laser Technology*, *181*, 111720. https://doi.org/10.1016/j.optlastec.2024.111720

Lu, D., Chen, Q., Wang, G., Liu, L., Li, G., & Moran, E. (2016). A survey of remote sensing-based aboveground biomass estimation methods in forest ecosystems. *International Journal of Digital Earth*, *9*(1), 63–105. https://doi.org/10.1080/17538947.2014.990526

Ma, Y., Chen, S., Ermon, S., & Lobell, D. B. (2024). Transfer learning in environmental remote sensing. *Remote Sensing of Environment*, *301*, 113924.

https://doi.org/10.1016/j.rse.2023.113924

Masoudnia, S., & Ebrahimpour, R. (2014). Mixture of experts: A literature survey. *Artificial Intelligence Review*, *42*(2), 275–293. https://doi.org/10.1007/s10462-012-9338-y

Maxwell, A. E., Warner, T. A., & Fang, F. (2018). Implementation of machine-learning classification in remote sensing: An applied review. *International Journal of Remote Sensing*, *39*(9), 2784–2817. https://doi.org/10.1080/01431161.2018.1433343

McDowell, N. G., Allen, C. D., Anderson-Teixeira, K., Aukema, B. H., Bond-Lamberty, B., Chini, L., Clark, J. S., Dietze, M., Grossiord, C., Hanbury-Brown, A., Hurtt, G. C., Jackson, R. B., Johnson, D. J., Kueppers, L., Lichstein, J. W., Ogle, K., Poulter, B., Pugh, T. A. M., Seidl, R., … Xu, C. (2020). Pervasive shifts in forest dynamics in a changing world. *Science*, *368*(6494), eaaz9463. https://doi.org/10.1126/science.aaz9463

McRoberts, R. E., Næsset, E., Gobakken, T., & Bollandsås, O. M. (2015). Indirect and direct estimation of forest biomass change using forest inventory and airborne laser scanning data. *Remote Sensing of Environment*, *164*, 36–42. https://doi.org/10.1016/j.rse.2015.02.018

McRoberts, R. E., Stehman, S. V., Liknes, G. C., Næsset, E., Sannier, C., & Walters, B. F. (2018). The effects of imperfect reference data on remote sensing-assisted estimators of land cover class proportions. *ISPRS Journal of Photogrammetry and Remote Sensing*, *142*, 292–300. https://doi.org/10.1016/j.isprsjprs.2018.06.002

McRoberts, R. E., Tomppo, E. O., & Næsset, E. (2010). Advances and emerging issues in national forest inventories. *Scandinavian Journal of Forest Research*, *25*(4), 368–381.

https://doi.org/10.1080/02827581.2010.496739

Mitchell, A. L., Rosenqvist, A., & Mora, B. (2017). Current remote sensing approaches to monitoring forest degradation in support of countries measurement, reporting and verification (MRV) systems for REDD+. *Carbon Balance and Management*, *12*(1), 9. https://doi.org/10.1186/s13021-017-0078-9

Montesano, P. M., Cook, B. D., Sun, G., Simard, M., Nelson, R. F., Ranson, K. J., Zhang, Z., & Luthcke, S. (2013). Achieving accuracy requirements for forest biomass mapping: A spaceborne data fusion method for estimating forest biomass and LiDAR sampling error. *Remote Sensing of Environment*, *130*, 153–170. https://doi.org/10.1016/j.rse.2012.11.016

Mullissa, A., Vollrath, A., Odongo-Braun, C., Slagter, B., Balling, J., Gou, Y., Gorelick, N., & Reiche, J. (2021). Sentinel-1 SAR Backscatter Analysis Ready Data Preparation in Google Earth Engine. *Remote Sensing*, *13*(10), Article 10. https://doi.org/10.3390/rs13101954

Næsset, E., Bollandsås, O. M., Gobakken, T., Gregoire, T. G., & Ståhl, G. (2013). Model-assisted estimation of change in forest biomass over an 11 year period in a sample survey supported by airborne LiDAR: A case study with post-stratification to provide "activity data." *Remote Sensing of Environment*, *128*, 299–314. https://doi.org/10.1016/j.rse.2012.10.008

Nitze, I., Schulthess, U., & Asche, H. (2012). *COMPARISON OF MACHINE LEARNING ALGORITHMS RANDOM FOREST, ARTIFICIAL NEURAL NETWORK AND SUPPORT VECTOR MACHINE TO MAXIMUM LIKELIHOOD FOR SUPERVISED CROP TYPE CLASSIFICATION*. 7.

Olmedo, G. F., Guevara, M., Gilabert, H., Montes, C. R., Arellano, E. C., Barría-Knopf, B., Gárate, F., Mena-Quijada, P., Acuña, E., Bown, H. E., & Ryan, M. G. (2020). Baseline of Carbon Stocks in Pinus radiata and Eucalyptus spp. Plantations of Chile. *Forests*, *11*(10), Article 10. https://doi.org/10.3390/f11101063

Olofsson, P., Foody, G. M., Herold, M., Stehman, S. V., Woodcock, C. E., & Wulder, M. A. (2014). Good practices for estimating area and assessing accuracy of land change. *Remote Sensing of Environment*, *148*, 42–57. https://doi.org/10.1016/j.rse.2014.02.015

Peregon, A., & Yamagata, Y. (2013). The use of ALOS/PALSAR backscatter to estimate above-ground forest biomass: A case study in Western Siberia. *Remote Sensing of Environment*, *137*, 139–146. https://doi.org/10.1016/j.rse.2013.06.012

Pohjankukka, J., Pahikkala, T., Nevalainen, P., & Heikkonen, J. (2017). Estimating the prediction performance of spatial models via spatial k-fold cross validation. *International Journal of Geographical Information Science*, *31*(10), 2001–2019. https://doi.org/10.1080/13658816.2017.1346255

Potapov, P. V., Dempewolf, J., Talero, Y., Hansen, M. C., Stehman, S. V., Vargas, C., Rojas, E. J., Castillo, D., Mendoza, E., Calderón, A., Giudice, R., Malaga, N., & Zutta, B. R. (2014). National satellite-based humid tropical forest change assessment in Peru in support of REDD+ implementation. *Environmental Research Letters*, *9*(12), 124012. https://doi.org/10.1088/1748-9326/9/12/124012

Potapov, P. V., Turubanova, S. A., Hansen, M. C., Adusei, B., Broich, M., Altstatt, A., Mane, L., &

Justice, C. O. (2012). Quantifying forest cover loss in Democratic Republic of the Congo, 2000–2010, with Landsat ETM + data. *Remote Sensing of Environment, Landsat Legacy Special Issue*, *122*, 106–116. https://doi.org/10.1016/j.rse.2011.08.027

Powell, R. L., Roberts, D. A., Dennison, P. E., & Hess, L. L. (2007). Sub-pixel mapping of urban land cover using multiple endmember spectral mixture analysis: Manaus, Brazil. *Remote Sensing of Environment*, *106*(2), 253–267. https://doi.org/10.1016/j.rse.2006.09.005

Pugh, T. A. M., Lindeskog, M., Smith, B., Poulter, B., Arneth, A., Haverd, V., & Calle, L. (2019). Role of forest regrowth in global carbon sink dynamics. *Proceedings of the National Academy of Sciences*, *116*(10), 4382–4387. https://doi.org/10.1073/pnas.1810512116

Puliti, S., Breidenbach, J., Schumacher, J., Hauglin, M., Klingenberg, T. F., & Astrup, R. (2021). Above-ground biomass change estimation using national forest inventory data with Sentinel-2 and Landsat. *Remote Sensing of Environment*, *265*, 112644. https://doi.org/10.1016/j.rse.2021.112644

Réjou-Méchain, M., Tanguy, A., Piponiot, C., Chave, J., & Hérault, B. (2017). biomass: An r package for estimating above-ground biomass and its uncertainty in tropical forests. *Methods in Ecology and Evolution*, *8*(9), 1163–1167. https://doi.org/10.1111/2041-210X.12753

Requena Suarez, D., Rozendaal, D. M. A., De Sy, V., Phillips, O. L., Alvarez-Dávila, E., Anderson-Teixeira, K., Araujo-Murakami, A., Arroyo, L., Baker, T. R., Bongers, F., Brienen, R. J. W., Carter, S., Cook-Patton, S. C., Feldpausch, T. R., Griscom, B. W., Harris, N., Hérault, B., Honorio Coronado, E. N., Leavitt, S. M., … Herold, M. (2019). Estimating aboveground net

biomass change for tropical and subtropical forests: Refinement of IPCC default rates using forest plot data. *Global Change Biology*, *25*(11), 3609–3624. https://doi.org/10.1111/gcb.14767

Roberts, D. A., Gardner, M., Church, R., Ustin, S., Scheer, G., & Green, R. O. (1998). Mapping Chaparral in the Santa Monica Mountains Using Multiple Endmember Spectral Mixture Models. *Remote Sensing of Environment*, *65*(3), 267–279. https://doi.org/10.1016/S0034-4257(98)00037-6

Rosenqvist, A., Shimada, M., Ito, N., & Watanabe, M. (2007). ALOS PALSAR: A Pathfinder Mission for Global-Scale Monitoring of the Environment. *IEEE Transactions on Geoscience and Remote Sensing*, *45*(11), 3307–3316. IEEE Transactions on Geoscience and Remote Sensing. https://doi.org/10.1109/TGRS.2007.901027

Salas, C., Donoso, P. J., Vargas, R., Arriagada, C. A., Pedraza, R., & Soto, D. P. (2016). The Forest Sector in Chile: An Overview and Current Challenges. *Journal of Forestry*, *114*(5), 562–571. https://doi.org/10.5849/jof.14-062

Sandoval, S., Montes, C. R., Olmedo, G. F., Acuña, E., & Mena-Quijada, P. (2022). Modelling above-ground biomass of Pinus radiata trees with explicit multivariate uncertainty. *Forestry: An International Journal of Forest Research*, *95*(3), 380–390. https://doi.org/10.1093/forestry/cpab048

Santoro, M., Cartus, O., Quegan, S., Kay, H., Lucas, R. M., Araza, A., Herold, M., Labrière, N., Chave, J., Rosenqvist, Å., Tadono, T., Kobayashi, K., Kellndorfer, J., Avitabile, V., Brown,

H., Carreiras, J., Campbell, M. J., Cavlovic, J., Bispo, P. da C., ... Seifert, F. M. (2024). Design and performance of the Climate Change Initiative Biomass global retrieval algorithm. *Science of Remote Sensing*, *10*, 100169. https://doi.org/10.1016/j.srs.2024.100169

Schoups, G., & Vrugt, J. A. (2010). A formal likelihood function for parameter and predictive inference of hydrologic models with correlated, heteroscedastic, and non-Gaussian errors. *Water Resources Research*, *46*(10). https://doi.org/10.1029/2009WR008933

Seitzer, M., Tavakoli, A., Antic, D., & Martius, G. (2022). *On the Pitfalls of Heteroscedastic Uncertainty Estimation with Probabilistic Neural Networks* (arXiv:2203.09168). arXiv. https://doi.org/10.48550/arXiv.2203.09168

Skowronski, N. S., Clark, K. L., Gallagher, M., Birdsey, R. A., & Hom, J. L. (2014). Airborne laser scanner-assisted estimation of aboveground biomass change in a temperate oak–pine forest. *Remote Sensing of Environment, Special Issue on 2012 ForestSAT*, *151*, 166–174. https://doi.org/10.1016/j.rse.2013.12.015

Verrelst, J., Rivera, J. P., Moreno, J., & Camps-Valls, G. (2013). Gaussian processes uncertainty estimates in experimental Sentinel-2 LAI and leaf chlorophyll content retrieval. *ISPRS Journal of Photogrammetry and Remote Sensing*, *86*, 157–167. https://doi.org/10.1016/j.isprsjprs.2013.09.012

Yue, J., Feng, H., Yang, G., & Li, Z. (2018). A Comparison of Regression Techniques for Estimation of Above-Ground Winter Wheat Biomass Using Near-Surface Spectroscopy. *Remote Sensing*, *10*(1), 66. https://doi.org/10.3390/rs10010066

Yuksel, S. E., Wilson, J. N., & Gader, P. D. (2012). Twenty Years of Mixture of Experts. *IEEE Transactions on Neural Networks and Learning Systems*, *23*(8), 1177–1193. https://doi.org/10.1109/TNNLS.2012.2200299

Zeng, N., Ren, X., He, H., Zhang, L., Zhao, D., Ge, R., Li, P., & Niu, Z. (2019). Estimating grassland aboveground biomass on the Tibetan Plateau using a random forest algorithm. *Ecological Indicators*, *102*, 479–487. https://doi.org/10.1016/j.ecolind.2019.02.023

Zhang, J. (2010). Multi-source remote sensing data fusion: Status and trends. *International Journal of Image and Data Fusion*, *1*(1), 5–24. https://doi.org/10.1080/19479830903561035

Zolkos, S. G., Goetz, S. J., & Dubayah, R. (2013). A meta-analysis of terrestrial aboveground biomass estimation using lidar remote sensing. *Remote Sensing of Environment*, *128*, 289–298. https://doi.org/10.1016/j.rse.2012.10.017

7. Supplementary Information

7.1 Reference data processing & label creation

Reference data processing proceeded in two stages: first, harmonization and quality control of the raw ARAUCO inventory data; second, spatial-temporal pairing of repeat observations to generate ΔAGB labels. The following sections document the key decisions at each stage, including outlier control thresholds, temporal alignment logic, and change-type assignment rules. These steps were designed to balance data completeness against label quality in an operational inventory dataset not originally collected for change detection purposes.

A. Harmonization, unit handling, and QA

- Projection and geometry: All plot coordinates are transformed to UTM 18S (EPSG: 32718). Integer meter x/y are retained for reproducibility; derived plot buffers use a radius of 80 m for spatial tests and pairing.
- Allometrics: *Pinus radiata* AGB is computed per tree via the Sandoval component equations; per-tree means and SDs are summed in quadrature and scaled to Mg/ha using plot density. The script stores both analytic plot-level SD and later MC SD.
- Robust outlier control (within parcel x year x month): Groups with < 5 plots are removed. Groups > 5 use a MAD filter ($z<2.5$) on AGB to drop spurious records while avoiding cross-month mixing.
- De-duplication priority: When exact duplicates exist for parcel x year x month x x x y, recent inventories are kept over historic.

B. Bare-earth and “no-change” controls

- Mapped bare-earth joins: Plot buffers intersect a mapped bare-earth polygon layer; matched records are cloned as bare points (AGB = 0) with month fixed to June for that year.
- Synthetic bare at t1 for very young stands: Plots with age < 5 (2017-2021) are duplicated to create a bare-at-t1 record by shifting inv_year = inv_year - age and zeroing tree/stand metrics. This increases representation of establishment transitions.

C. Parcel-year uncertainty via Monte Carlo

- Aggregation unit: Uncertainty is summarized at (parcel, year, month, age).
- Simulation: For each group, the script performs N=1000 draws. Each draw samples a plot in the group, then draws a tree-level biomass from N(mu_tree, sd_tree) and scales by density; kg/ha are converted to Mg/ha.
- Outputs: For every group: agb_ha_mean, agb_ha_std_mc, and 5th/95th percentiles (lower/upper). These MC quantities are later propagated into pairing/label QA and serve as propagated label uncertainty.

D. Time alignment and “RS year” logic (DJF anchor)

- Anchor: A remote-sensing year (RS year) is defined as Dec-Mar (DJF) around January 31 of year Y.
- Before/after rules (per inventory date):
    - Before image for a plot at (inv_year, inv_month): rs_year_t1 = inv_year - 1 if month < Jan, else inv_year.
    - After image: rs_year_t2 = inv_year if month < Mar, else inv_year + 1
- Change-type adjustment: For growth classes (new/early/mid-late), rs_year_t2 is shifted -1 (to use the rio DJF composite); for loss classes it is not.
- Ordering guarantees: Enforce rs_year_t2 > rs_year_t1; if equal after adjustments, t2 = t1 + 1.

E. Spatial–temporal pairing and coherence checks

- Nearest with topology: Within each parcel, observations are paired forward in time using nearest-neighbor joins subject to centroid distance ≤ 60 m and buffer IoU ≥ 0.05.
- Minimum temporal gap: 30 days between paired observations (when timestamps exist).
- Age-calendar coherence: For pairs with known ages, require non-decreasing age and $|(age_2-age_1) - (year_2-year_1)|$ within a tolerance (base ±2 yr, scaled by calendar gap, capped at ±3 yr). A relaxed rule applies when t_1 and t_2 are in the same RS year but separated by ≥ 30 days.

- "Large signal" overrides: Keep pairs that clearly indicate large change even if coherence is marginal:
  (i) deforest-like: $\Delta AGB \leq -150$ Mg ha^{-1} or $t_2 \approx$ bare;
  (ii) new growth from bare with $\Delta AGB \geq 5$ Mg ha^{-1};
  (iii) $|\Delta AGB| \geq 120$ Mg ha^{-1}.
- Growth-rate realism: Keep all losses; cap positive rates at ≤ 50 Mg ha^{-1} yr^{-1}.
- Deforest spatial sanity: Remove any deforest labels whose t_2 buffered point does not intersect mapped bare-earth polygons (prevents commission).

F. Change-type assignment and consensus

- Label rules (first-match wins):
  new_growth (t_1 bare → t_2 non-bare, $\Delta AGB > 0$),
  early_growth ($\Delta AGB > 0$, t_1 age 1–7),
  mid_late_growth ($\Delta AGB > 0$, t_1 age > 7),
  deforest ($\Delta AGB \ll 0$ or t_2 bare),
  thin_prune ($\Delta AGB < 0$ otherwise), else no_change (rare after filters).
- Parcel-year consensus filter: Within each *(parcel, y_1m_1, y_2m_2)* group, compute class votes and retain groups with majority_share ≥ 0.40; within those, keep only plots matching the majority class (drops minority labels while preserving spatial density).

G. Final identifiers and housekeeping

- Canonical unique_id:
  unique_id = parcel_norm + "_" + y1m1 + "_" + y2m2 + "_" + x_t2_5 + "_" + y_t2_5,
  where x_t2_5/y_t2_5 are 5-m snapped coordinates; y1m1/y2m2 are YYYYMM tags.
- Duplicate control: Exact duplicates by unique_id are dropped at the end of processing.
- Sanity constraints: Keep rows with inv_year ≥ 2016 (bare rows exempt), and winsorized outlier AGB to (-600, 400) Mg ha^{-1} during early QA to avoid pathological inputs.

7.2 Remote sensing data processing

All remote sensing inputs were processed in Google Earth Engine (GEE) using a standardized pipeline to generate annual composites at 30 m resolution in UTM 18S (EPSG:32718). This pipeline integrates

Sentinel-1 (C-band SAR), Sentinel-2 (optical), ALOS PALSAR (L-band SAR), and Copernicus GLO-30 DEM. Below we summarize the key pre-processing and harmonization steps.

Sentinel-1 (SAR)

- Source: COPERNICUS/S1_GRD.
- Pre-processing: border noise removal, terrain flattening (SRTM DEM), conversion to linear power.
- Bands: VV, VH, VV/VH ratio (linear + dB).
- Annual composites: median of acquisitions during December–March (DJFM) window.
- Texture features: 70 m neighborhood median filters.

Sentinel-2 (optical)

- Source: COPERNICUS/S2_HARMONIZED Level-1C.
- Cloud mask: QA60 bitmask (cirrus + opaque).
- Bands: B2–B8A, B11, B12.
  Indices: NDVI, NDMI, NBR, MSI, EVI, SAVI, MSAVI.
- Annual composites: median DJFM composites; fallback all-masked if no valid acquisitions.
- Texture features: 70 m neighborhood medians for key SWIR/NIR bands.
- Level-1C was chosen over Level-2A for temporal coverage (2015 onward); atmospheric correction was omitted given limited marginal gains in ML settings (Rumora et al., 2020).

ALOS PALSAR (SAR)

- Source: JAXA/ALOS/PALSAR/YEARLY/SAR_EPOCH.
- Pre-processing: orthorectified, slope-corrected annual mosaics.
- Conversion: DN^2 scaled to linear power; HH, HV, HH/HV ratio exported in linear and dB.
- Temporal alignment: calendar year matching RS "t1/t2" rules, with previous-year fallback if no scene available.

Copernicus DEM

- Source: COPERNICUS/DEM/GLO30.
- Products: elevation, slope, aspect, all reprojected to 30 m.

7.3 Sensor selection and model testing

To assess the impact of the different sensors on biomass change estimation, we systematically tested combinations of the three sensors using a baseline MLP. The performance evaluation of

different sensor combinations highlighted that including multiple sensors improved biomass change estimation accuracy. Sentinel-2 alone performed well with a RMSE of 109.48 Mg/ha, while Sentinel-1 and ALOS PALSAR showed slightly higher error with RMSE values of 117.18 Mg/ha and 150.13 Mg/ha, respectively. Combining Sentinel-2 with Sentinel-1 slightly improved accuracy, reducing the RMSE to 109.09 Mg/ha. A combination of Sentinel-2 and ALOS PALSAR further improved results, with a RMSE of 108.13 Mg/ha. However, the best performance came from using all three sensors—Sentinel-1, Sentinel-2, and ALOS PALSAR—yielding the lowest RMSE of 106.90 Mg/ha.

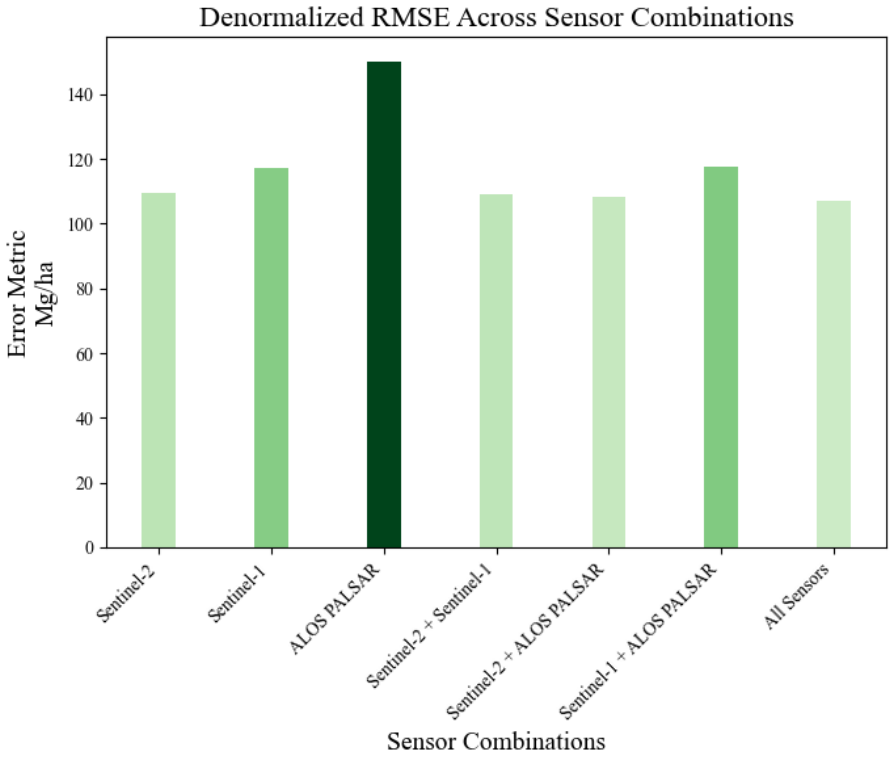

Figure 7. *Accuracy measured in RMSE (de-normalized from the original scaled output of the MLP) across various individual remote sensors and sensor combinations. Results show that a combination of all three sensors most accurately predicts changes in aboveground biomass.*

7.4 Machine learning model descriptions

For the Random Forest model, we utilized the RandomForestRegressor from Scikit-Learn. We performed hyperparameter tuning through grid search, where we explored different combinations of the number of estimators (n_estimators), maximum tree depth (max_depth), and the minimum number of samples required to split an internal node (min_samples_split and min_samples_leaf). The best-performing model was found with the following hyperparameters: 2,000 trees (n_estimators), a maximum depth of 20 (max_depth), a minimum leaf size of 2 (min_samples_leaf), and a minimum node split of 2 (min_samples_split). The goal was to find the best set of hyperparameters that

minimized the mean squared error (MSE). After training, the best performing model was evaluated on the test set using R-squared and RMSE.

We used the XGBoost python package to train and tune the XGBoost model. This involved using the xbg.train function to train the model with hyperparameters including max_depth, min_child_weight, and subsample. The best-performing model had a learning rate of 0.2, a maximum tree depth of 10, a minimum child weight of 10, 1,500 estimators, and a subsample ratio of 0.7. The XGBoost algorithm is optimized for both speed and accuracy, making it ideal for our dataset. After finding the best hyperparameters using the minimized MSE, we evaluated the final model using R-squared and RMSE.

For the MLP model architecture, we designed feed-forward neural networks using PyTorch, with tunable hyperparameters to optimize performance for predicting aboveground biomass changes. The hyperparameters, including the number of hidden layers, hidden units, dropout rate, learning rate, weight decay, and output layer bias, were tuned using the Optuna framework. Optuna systematically searched the hyperparameter space to minimize MSE loss during training, resulting in different optimal architectures for each of the four models created. These variations allowed the models to adapt to the specific characteristics of the data and task at hand.

The final architectures for the models varied based on the hyperparameter tuning results. For instance, one model consisted of three fully connected hidden layers with 96 neurons each, while others had different numbers of layers and neurons tailored to their respective tasks. All models employed the hyperbolic tangent (tanh) activation function to accommodate both positive and negative biomass change values. Dropout regularization was applied with rates ranging from 0.2 to 0.5, depending on the model, to prevent overfitting. The output layer in each model used a linear transformation to predict a single biomass change value per pixel, with the bias initialized to values such as -0.0386, as determined by the tuning process. Optimization was performed using the Adam optimizer, with learning rates ranging from 0.0017 to 0.01, and weight decay rates varying between 1×10^{-6} and 1×10^{-3} to enhance generalization. These distinct architectures highlight the flexibility of the hyperparameter tuning process in adapting to varying forest dynamics and remote sensing inputs, ensuring robust performance across diverse scenarios.

$$y = W_{out} \cdot \text{dropout}(\tanh(W_3 \cdot \text{dropout}(\tanh(W_2 \cdot \text{dropout}(\tanh(W_1 x + b_1), 0.2) + b_2), 0.2) + b_3), 0.2) + b_{out}$$

With the optimal parameters found:

- 3 hidden layers, each with 96 neurons.

- $W_1, W_2, W_3 \in \mathbb{R}^{96\times 23}, b_1, b_2, b_3 \in \mathbb{R}^{96}$
- $W_{out} \in \mathbb{R}^{1\times 96}, b_{out} \in \mathbb{R}$ with initial bias $b_{out} = -0.0386$

We first compared the performance of short-term ΔAGB models, where we calculated the biomass change between 1-2 years in the reference dataset. The RF model demonstrated reliable performance, with an RMSE of 127.53 Mg/ha and NRMSE of 0.196 or 19.6%. The XGBoost model underperformed the RF in reducing the overall RMSE to 135.63 Mg/Ha and NRMSE of 20.8%. Lastly, the MLP architecture further improved predictions by achieving a test RMSE of 86.11 Mg/ha and a NRMSE of 14%.

7.5 Uncertainty analysis and calibration

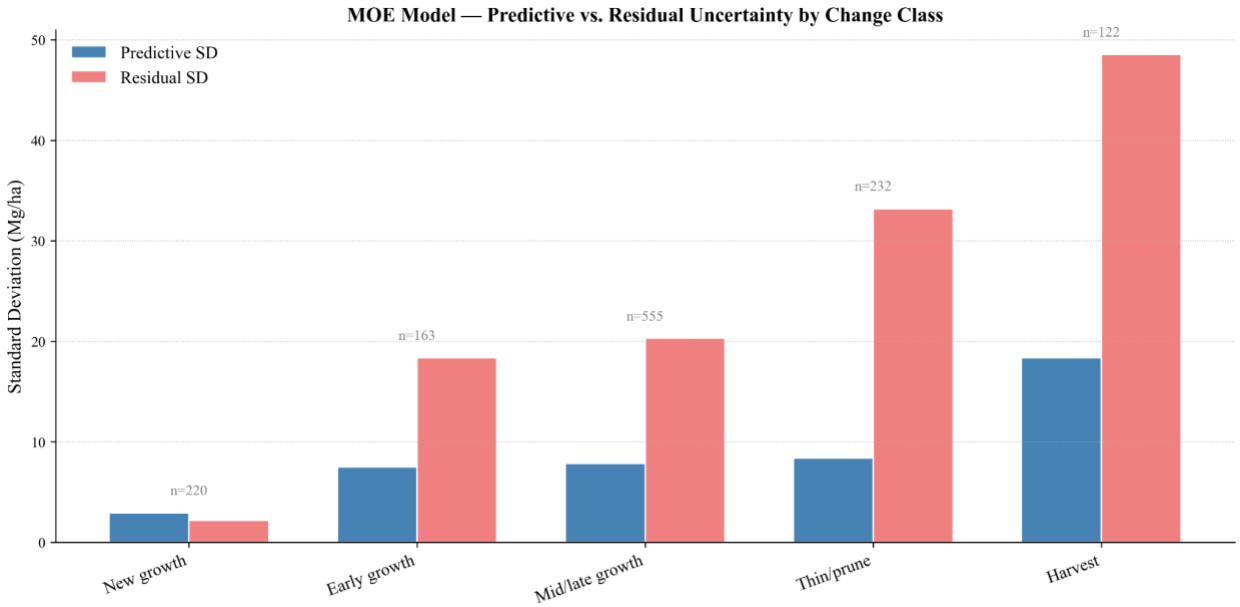

Figure 8. *Predictive uncertainty (mean predictive SD from GNLL and MC dropout) versus residual uncertainty (empirical residual SD from the test set) by change class for the MOE model. For each class, the gap between bars reflects the degree of overconfidence. In the New growth change class, predictive and residual SD are in close agreement, while all other classes show systematic overconfidence. Sample sizes (n) are shown above each class pair.*

Coverage by change class at nominal 90% confidence interval						
	Empirical %	*Gap %*	*Predicted SD (Mg/ha)*	*Residual SD (Mg/ha)*	*Relative MAE*	*N*
New growth	*92.7*	*2.7*	*2.29*	*2.20*	*0.24*	*220*
Early growth	*57.7*	*-32.3*	*7.48*	*18.36*	*0.47*	*163*
Mid/late growth	*40.5*	*-49.5*	*7.85*	*20.29*	*0.48*	*555*
Harvest	*39.3*	*-50.7*	*18.36*	*48.54*	*0.15*	*122*
Thin/prune	*33.2*	*-56.8*	*8.37*	*33.17*	*0.75*	*232*

Table 4. *Uncertainty calibration and prediction error by change class for the MOE model. Empirical coverage is reported at a nominal 90% confidence interval; Gap (%) is the difference between empirical and nominal coverage, where negative values indicate overconfidence. Predicted SD is the mean per-pixel predictive SD from the GNLL and MC dropout inference (t = 30). Residual SD is the empirical SD of prediction errors on the independent test set. N reflects the number of independent test set observations per class after spatial holdout.*

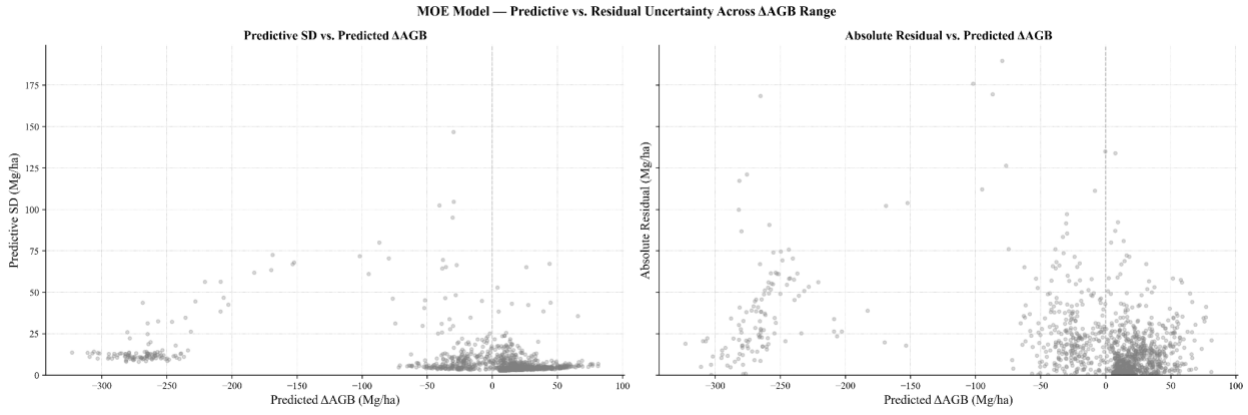

Figure 9. *Predictive SD (left) and absolute residual (right) versus predicted ΔAGB for the MOE model across the independent test set, plotted on a shared y-axis to facilitate direct comparison. Both panels show the same qualitative asymmetry: uncertainty and residuals increase with loss magnitude on the negative side while remaining low and flat across all growth predictions. The primary difference appears in the moderate loss zone (-50 to 0 Mg/ha), where absolute residuals are elevated but predictive SD remains low.*

The compressed scale of the left panel relative to the right reflects systematic underestimation of predictive uncertainty across all change types.