2026 Volume: 02 Issue: 11 www.ijrpa.com ISSN 2456-9995 Research Article

International Journal Research Publication Analysis

Page: 01-21

CARBON INSIGHT: AN AI-BASED CARBON ESTIMATION AND VERIFICATION SYSTEM FOR BLUE CARBON ECOSYSTEMS

***Dawan N. Gowda, Gourang A. M., Hemanth R., J. Manu, Pallavi R. (Guide)**

Dept. of Computer Science and Engineering, Jyothy Institute of Technology, Bengaluru, Karnataka, India.

Article Received: 15 October 2026
Article Revised: 04 November 2026
Published on: 23 November 2026

***Corresponding Author: Dawan N. Gowda**
Dept. of Computer Science and Engineering, Jyothy Institute of Technology, Bengaluru, Karnataka, India.

DOI: https://doi-doi.org/101555/ijrpa.6019

ABSTRACT

Blue carbon projects need site-level estimates of mangrove extent and carbon stock. Field inventories are slow, and global default factors can be far from local conditions. This paper reports the implementation of Carbon Insight, a web system in which a user draws or uploads a boundary and receives a land-cover map, a carbon estimate with an uncertainty range, and a verification report for that area. A FastAPI backend builds a cloud-masked Sentinel-2 median composite in Google Earth Engine and classifies it at 10 m with a four-class U-Net trained on 688 tiles from 42 Indian regions. On five held-out regions of the Karnataka coast and Bengaluru the model reaches a mean IoU of 0.643 and a mangrove IoU of 0.587, against 0.120 for an NDVI threshold. Above-ground biomass is computed per pixel from canopy height calibrated on nearby GEDI lidar footprints; in spatial block cross-validation its median error is 20% for mangrove and 14% for tree cover, compared with 1045% and 241% for IPCC Tier 1 defaults. Every map is checked against ESA WorldCover and Global Mangrove Watch, and a Monte Carlo procedure combines map error and coefficient error into 90% intervals and a conservative creditable quantity. A restoration module ranks candidate replanting sites and projects their CO_2 removal. The outputs are screening estimates, not verified credits: at a field site from the literature the biomass carbon is still about six times the plot value.

INDEX TERMS: Blue Carbon, Mangrove Mapping, Sentinel-2, U-Net, Google Earth Engine, Gedi, Canopy Height, Uncertainty, Carbon Verification.

I. INTRODUCTION

Mangroves, tidal marshes and seagrass meadows store carbon at high density, and a large share of it sits in waterlogged soil where it can remain for centuries [1]. In the Indo-Pacific, mangrove ecosystems have been measured at roughly 1,000 Mg C per hectare when soil to depth is included [2]. These stocks are what restoration programmes and blue carbon finance try to protect, and both need numbers for a specific site, not a national average.

Two routes to such numbers are common. Field inventories with allometric equations, such as the study of Bindu *et al.* at Kunhimangalam in Kerala [3], are accurate for the plots measured but slow to repeat. IPCC Tier 1 default factors [4], [5] can be applied anywhere in seconds, yet a single biomass value per climate zone ignores the large difference between tall deltaic mangroves and the short, scrubby stands found along much of the Indian coast. Our earlier survey [6] found that most published methods sit at one of these two extremes, and that few tools report how wrong an estimate might be.

The first prototype of this project shared that weakness. It labelled every pixel with an NDVI above 0.3 as vegetation and multiplied the area by a fixed density. A binary U-Net trained later by the team scored an IoU of 0.883 on its own validation tiles, but when we tested it on regions it had not seen, its woody-cover IoU fell to 0.225. That gap shaped the present system, in which every model and coefficient is judged on data from places it was not fitted to.

The contributions of this work are:

1) a working web system that turns a drawn or uploaded polygon into a cloud-free composite, a four-class land-cover map, carbon stocks, change between years and a PDF report, with a REST interface, saved history and 64 automated tests;

2) a four-class U-Net trained on 42 Indian regions and evaluated on held-out regions and seasons against the earlier U-Net and an NDVI threshold;

3) per-pixel above-ground biomass from canopy height, re-calibrated for each analysis on nearby GEDI footprints and validated by spatial block cross-validation;

4) a verification step against two independent reference maps and a Monte Carlo uncertainty estimate that yields a conservative creditable amount; and

5) a restoration planner that ranks candidate replanting sites and projects their removals.

II. Related Work

A. Carbon estimation from field and remote-sensing data

Bindu *et al.* [3] related LISS-IV NDVI to plot-level above-ground biomass at Kunhimangalam and converted it to carbon with a root-to-shoot ratio and a carbon fraction, so the method depends on local field calibration. Dutta Roy *et al.* [7] reviewed 105 remote-sensing studies of mangrove blue carbon in the Asia-Pacific; Sentinel-2 was the most used sensor, machine-learning regressors were becoming common, and below-ground carbon remained the least studied pool. Zhang [8] combined field inventory, spaceborne lidar and machine learning to map mangrove biomass, reporting $R^2 = 0.72$ and an RMSE of 37.24 Mg/ha at continental scale. Zhang and Fan [9] used Google Earth Engine with Sentinel-2 and GEDI to classify mangrove species and estimate carbon at Dongzhaigang, China.

B. Segmentation and global reference products

U-Net [10] remains a standard architecture for pixel labelling. Maung *et al.* [11] used it with PlanetScope and Sentinel-2 imagery over Myanmar's Wunbaik mangroves and reached 94.05% overall accuracy, with results depending on the input features. Several global layers are now open at 10–30 m: ESA WorldCover [12], Global Mangrove Watch (GMW) v3 with extents from 1996 to 2020 [13], a 10 m canopy height model trained on GEDI [14], and the GEDI footprint biomass product L4A [15], [16]. Each is useful on its own; the system described here uses them together, as training labels, references and calibration data.

C. Accuracy and uncertainty

Olofsson *et al.* [17] showed that counting mapped pixels gives a biased area and recommended error-adjusted estimates with confidence intervals. The IPCC guidelines [4], [5] give low and high bounds for most default factors. We did not find a blue carbon tool in the reviewed work that carries both kinds of error through to the final carbon figure.

D. Research gap

Table I summarises the gap. Existing studies need field plots or lidar processing, cover one site, or stop at classification. Carbon Insight trades some ecological detail for coverage: it runs for any polygon in India from a browser, separates mangrove from other vegetation, and states its uncertainty and its agreement with independent maps.

TABLE I: COMPARISON WITH RELATED WORK.

Work	Approach	Gap addressed here
Bindu [3]	LISS-IV NDVI, field plots	Needs plots; one site
Dutta Roy [7]	Review of 105 studies	Few working tools
Zhang [8]	Lidar and ML for AGB	Needs lidar processing
Zhang, Fan [9]	GEE, Sentinel-2, GEDI	Single site study
Maung [11]	U-Net land cover	No carbon estimate
Lang [14]	Global canopy height	Height, not carbon
Olofsson [17]	Error-adjusted area	Method, not a tool
This work	**U-Net, height AGB, checks, intervals**	**—**

III. PROBLEM STATEMENT AND OBJECTIVES

For any polygon in India, the system must estimate the mangrove and other woody cover inside it, their carbon, and the reliability of both, within seconds and without remote-sensing expertise. The objectives are:

O1: accept a polygon of up to 5,000 ha, drawn on the map or uploaded as GeoJSON or KML;

O2: build cloud-free Sentinel-2 imagery for the area automatically;

O3: separate mangrove, tree cover, other vegetation and non-vegetation at 10 m, with accuracy measured on regions unseen in training;

O4: estimate biomass, soil and total carbon and CO_2 equivalent, with every coefficient traceable to a source;

O5: verify each map against independent references and attach an uncertainty interval;

O6: support change detection between years and restoration planning; and

O7: store results, produce PDF reports and expose a documented, tested REST interface.

IV. SYSTEM ARCHITECTURE

Carbon Insight (repository name BlueCarbonAI) uses the four-layer client–server design shown in Fig. 1. The user interface is a React 19 single-page app, bundled with Vite and styled with Tailwind CSS, with home, analysis, report, history and about pages. Its react-leaflet map offers satellite, street and topographic base layers; Leaflet-Geoman handles drawing, Turf.js computes the geodesic area, and GeoJSON or KML boundaries can be uploaded. Results appear as map overlays and panels.

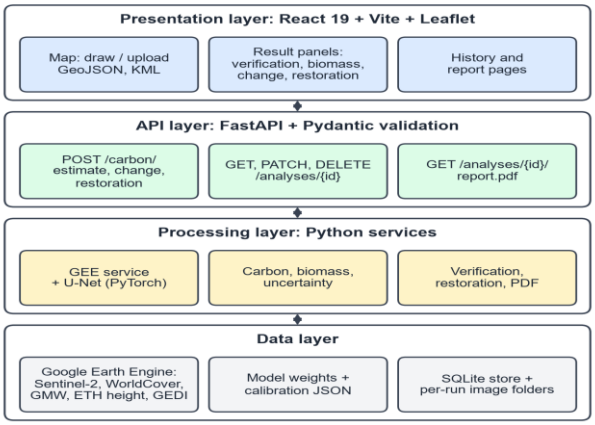

Fig. 1. Layered architecture of Carbon Insight as implemented.

The API layer is a FastAPI application. Request bodies are validated by Pydantic models, long analyses run in a worker thread so the server stays responsive, and errors are mapped to HTTP codes: 404 when no usable imagery exists, 422 for invalid input, 502 when Earth Engine fails and 500 otherwise. Processing is split into the Python services of Table II. The data layer holds the Earth Engine catalogue [18], the model weights and calibration files, a SQLite database of saved analyses, and one folder of PNG layers per run; folders not linked to a saved analysis are deleted after seven days.

Table III lists the REST interface. An estimate request carries a GeoJSON polygon, its area in hectares and the method (unet or ndvi). A change request adds two calendar years, and a restoration request adds a planting success rate (0.05–1) and a horizon of 5–30 years. Every successful result is saved automatically; a storage failure is logged but does not lose the response. Fig. 2 shows how one estimate request passes through these services.

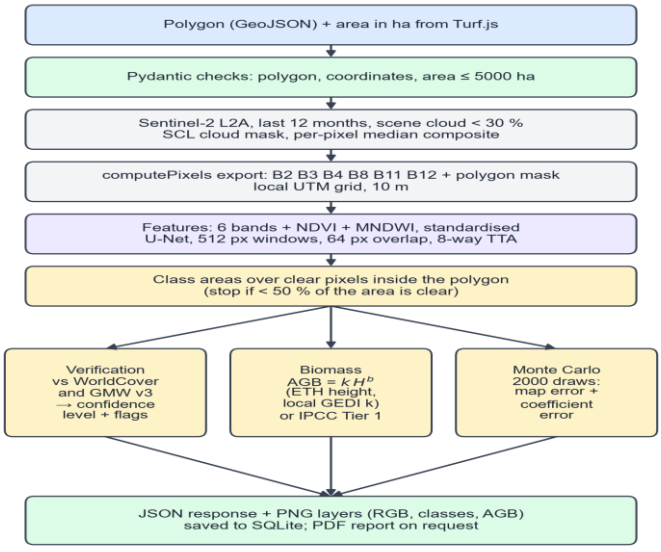

Fig. 2. Processing pipeline for one estimate request.

TABLE II: SOFTWARE COMPONENTS.

Component	Role	Technology
Map, drawing	Draw, edit, measure	react-leaflet, Geoman, Turf
Upload	Read GeoJSON and KML	Browser File API
Charts	Areas, carbon, change	Chart.js 4
API	Validation, routing, errors	FastAPI 0.115, Pydantic
Imagery	Composite and export	earthengine-api 1.7
Classifier	Four-class U-Net, TTA	PyTorch 2.13 [19]
Carbon, biomass	Stocks, height model	NumPy
Uncertainty	Monte Carlo, 2,000 draws	NumPy
Restoration	Suitability, patches	SciPy ndimage
Storage	Saved analyses, history	SQLite
Reports	A4 PDF report	ReportLab 5.1

TABLE III: REST INTERFACE OF THE BACKEND.

Endpoint	Behaviour
POST /carbon/estimate	Classes, carbon, biomass, checks, intervals
POST /carbon/change	Two-year transition matrix, carbon change
POST /carbon/restoration	Ranked sites, CO_2 projection
GET /analyses/	List saved analyses (paged)
GET, PATCH, DELETE /analyses/{id}	Reopen, rename, delete
GET /analyses/{id}/report.pdf	Server-made A4 PDF
GET /health/; /outputs/*	Status; PNG layers of runs

V. IMPLEMENTATION METHODOLOGY

A. Area input and validation

The browser sends the polygon and its Turf.js area. The backend accepts only a GeoJSON Polygon with at least three corners and valid longitude and latitude, rejects areas above 5,000 ha, and for change requests requires $2018 \le y_1 < y_2 \le$ the current year, since 2018 is the first full year of Sentinel-2 surface reflectance over India.

B. Imagery retrieval and compositing

The Earth Engine service filters the COPERNICUS/S2_SR_HARMONIZED collection [20] to scenes that intersect the polygon, fall in the previous 12 months and have under 30% scene cloud. Pixels whose scene-classification (SCL) class is no data, saturated, cloud shadow, cloud, cirrus or snow are masked, and the per-pixel median of the remaining observations forms the composite. Bands B2, B3, B4, B8, B11 and B12 and a polygon mask are exported with computePixels on the local UTM grid at 10 m (coarser only if a side would exceed 2,500 pixels). A pixel is counted only if it lies inside the polygon and has data; if fewer than half the pixels are clear, the request stops with a message instead of reporting on a partly visible area.

C. Input features

Reflectances are scaled to the 0–1 range and two indices are appended: the normalised difference vegetation index [21] and the modified normalised difference water index [22], which helps separate inundated mangrove from dry forest:

$$\mathrm{NDVI} = (\rho_{\mathrm{NIR}} - \rho_{\mathrm{R}}) / (\rho_{\mathrm{NIR}} + \rho_{\mathrm{R}}), \quad (1)$$

$$\mathrm{MNDWI} = (\rho_{\mathrm{G}} - \rho_{\mathrm{SWIR1}}) / (\rho_{\mathrm{G}} + \rho_{\mathrm{SWIR1}}), \quad (2)$$

Where NIR, R, G and SWIR1 are bands B8, B4, B3 and B11. The eight channels are standardised with the mean and standard deviation stored in the model checkpoint, so training and inference always use the same preprocessing.

D. Segmentation Network

The classifier is a U-Net [10] with batch normalisation (Fig. 3). Three encoder stages of 32, 64 and 128 channels, each two 3×3 convolutions with batch normalisation and ReLU followed by 2×2 max pooling, lead to a 256-channel bottleneck with 10% spatial dropout. The decoder upsamples with 2×2 transposed convolutions, concatenates the matching encoder features, and a 1×1 convolution outputs four classes: non-vegetation, mangrove, tree cover and other vegetation. The network has 1,928,548 parameters. Areas larger than 512 pixels are processed in 512-pixel windows that overlap by 64 pixels, with class probabilities averaged where windows meet. By default the probabilities are also averaged over the eight flips and rotations of each window (test-time augmentation). The probability of the chosen class is kept as a per-pixel confidence.

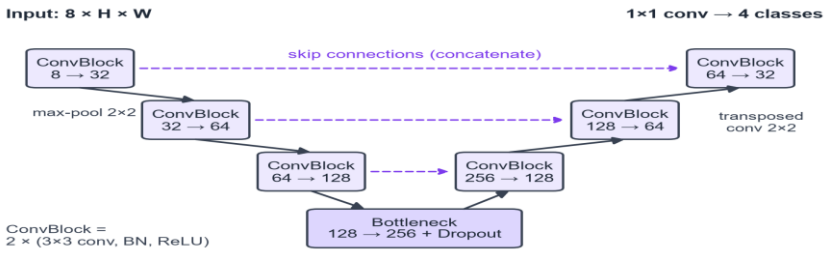

Fig. 3. Four-class U-Net used for land-cover mapping. Numbers give input and output channels of each block.

E. Training data and procedure

The dataset builder cuts 256×256-pixel tiles at 10 m from 42 regions (Table IV). For each tile it exports the clearest Sentinel-2 scene of each season (December–February, March–May,

June–September and October–November) between 2020 and 2022, keeping scenes with at most 10% cloud in the tile. Labels come from ESA WorldCover 2021 [12]: tree cover becomes *tree cover*, mangroves become *mangrove*, shrubland, grassland, cropland, herbaceous wetland and moss become *other vegetation*, and built-up, bare, snow and water become *non-vegetation*; cloud and shadow pixels are ignored in the loss. In mangrove regions 70% of tile centres are sampled on mangrove pixels so that the class is not swamped by its surroundings. Regions, not tiles, are assigned to the splits: the test set is the Karnataka coast (Kali, Aghanashini, Kundapura, Netravati estuaries) and Turahalli forest near Bengaluru, none of which contributes any training tile.

TABLE IV: TRAINING DATA AND CONFIGURATION.

Item	Value
Regions	42 (34 train, 3 validation, 5 test)
Tiles	688 (587 / 41 / 60), 256 × 256 px at 10 m
Imagery	Sentinel-2 L2A, 2020–2022, four seasons
Inputs	B2, B3, B4, B8, B11, B12, NDVI, MNDWI
Labels	ESA WorldCover 2021 v200, four classes
Augmentation	192 px crops, rotations, flips, gain, haze
Loss	Class-weighted cross-entropy + Dice
Optimiser	AdamW, lr 10^{-3}, decay 10^{-4}, one-cycle
Batch, epochs	8, 120, mixed precision
Hardware	NVIDIA GeForce RTX 3050, ≈4.3 s per epoch
Selected model	Epoch 110, validation mIoU 0.675

Training minimises the sum of a class-weighted cross-entropy and a soft Dice loss [23]:

$$L = L_{\mathrm{CE}}^{\mathrm{w}} + L_{\mathrm{Dice}}, \qquad (3)$$

$$L_{\mathrm{Dice}} = 1 - (1/K)\,\Sigma_{\mathrm{k}}\,(2I_{\mathrm{k}} + 1) / (U_{\mathrm{k}} + 1), \qquad (4)$$

Where I_{k} is the sum over pixels of the product of the predicted probability and the one-hot label for class k, U_{k} is the sum of both, and $K = 4$. The cross-entropy weights are the inverse square roots of the class frequencies, normalised to a mean of one. Optimisation uses AdamW [24] with a one-cycle schedule [25]; augmentation applies random crops, rotations, flips, brightness and per-band gain changes, and haze. The checkpoint with the best validation mIoU is kept (Fig. 4). A second training round added 14 regions of small estuarine mangroves and dry scrub, which raised test mIoU from 0.617 to 0.643.

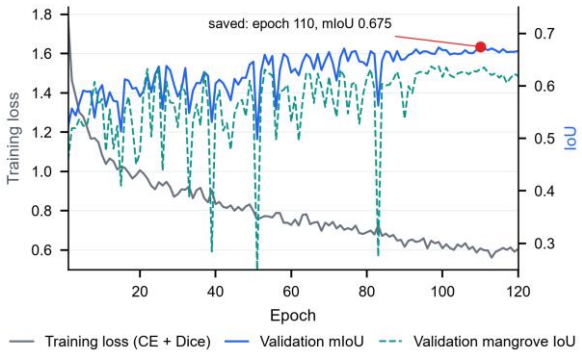

Fig. 4. Training loss and validation IoU over 120 epochs (second round). The validation set has only 41 tiles from three regions, which explains the noise.

F. Class areas and carbon stock

Let N be the number of counted pixels and n_k the number assigned to class k. The area of class k is its share of the polygon area A:

$$A_k = A \cdot n_k / N. \qquad (5)$$

Biomass carbon C_B and soil carbon C_S are summed over the woody classes, and CO_2 equivalent uses the molar mass ratio 44/12:

$$C_B = \Sigma_k A_k \cdot AGB_k \cdot (1 + R_k) \cdot CF_k, \qquad (6)$$

$$CO_2e = (44/12) \cdot (C_B + A_{man} \cdot SOC), \qquad (7)$$

Where AGB is above-ground biomass in t dry matter per hectare, R the root-to-shoot ratio, CF the carbon fraction and SOC the mangrove soil carbon of the top metre. Only mangrove and tree cover carry carbon; other vegetation is small or short-lived and is not counted. Table V gives the IPCC values used for R, CF and SOC and, when canopy height is unavailable, for AGB. Each is stored as a low, best and high value, and the response reports all three scenarios. The *potential credits* shown in the interface are biomass carbon expressed as t CO_2e.

TABLE V: IPCC TIER 1 COEFFICIENTS. (LOW / BEST / HIGH)

Class	AGB (t/ha)	*R*	CF	SOC (t/ha)
Mangrove	92/192/204	0.29/0.49/0.51	0.429/0.451/0.471	286/386/471
Tree cover	130/180/280	0.28/0.24/0.37	0.44/0.47/0.49	not counted

Mangrove: [5], Ch. 4, Tables 4.2, 4.3, 4.5 and 4.11 (low = tropical dry, best = tropical wet, high = upper 95% bound). Tree cover: [4], Vol. 4, Tables 4.3, 4.4 and 4.7, Asia continental forests.

G. Biomass from canopy height

Tier 1 assigns 192 t/ha to every tropical wet mangrove, while the median GEDI L4A footprint over Indian mangroves is 13.8 Mg/ha. The system therefore replaces the default AGB with a per-pixel estimate from the ETH 10 m canopy height map H [14]:

$$\mathrm{AGB}(x) = k \cdot H(x)^{b}, \;\; k = \Sigma_j \,\mathrm{AGB}_j \,/\, \Sigma_j \, H_j^{b}. \quad (8)$$

The exponent b was fitted once per class as a log–log regression with Duan's smearing correction [26] on quality-filtered GEDI L4A footprints [16] (quality flag 1, not degraded, beam sensitivity above 0.95), with height averaged over each 25 m footprint: 37,680 mangrove footprints from 27 regions gave b = 1.211, and 168,758 tree-cover footprints from 42 regions gave b = 1.087. For each analysis the scale k is re-computed from up to 5,000 footprints j of the same class within 25 km of the polygon, which absorbs regional differences in wood density and stand structure. With fewer than 30 footprints the India-wide fit is used. The uncertainty of AGB combines a bootstrap of the local footprints with a log-normal model error whose spread was measured in validation (Section VII-C).

H. Verification against reference maps

Each class map is compared, over the counted pixels, with WorldCover 2021 [12] and the GMW v3 2020 mangrove extent [13]. The system reports the mangrove area according to all three maps, the IoU [27] of the model's mangrove with each reference and of the two references with each other, and the IoU of woody cover with WorldCover. A confidence score combines agreement G, the clear-pixel share V and the mean class probability $\bar{P}$:

$$S = 0.50\, G + 0.25\, V + 0.25\, \bar{P}. \quad (9)$$

When either reference shows at least 1 ha of mangrove, G is the mean of the two mangrove IoUs averaged with the woody IoU; otherwise G is the woody IoU. The level is High for $S \geq 0.70$, Medium for $S \geq 0.50$ and Low below that. Flags warn about cloud gaps, low model confidence, small areas, mangrove mapped where neither reference has any, and large over- or under-estimates. Because WorldCover also supplied the training labels, agreement with it is not fully independent; GMW is.

I. Uncertainty propagation

Pixel counts are converted to bias-corrected areas with the stratified estimator of Olofsson *et al.* [17]:

$$\hat{A}_j = \Sigma_i A_i \cdot \hat{p}_{ij}, \ \hat{p}_{ij} = n_{ij} / n_{i\bullet}, \quad (10)$$

Where n_{ij} counts pixels mapped as class i and labelled j in the 60 held-out test tiles. In each of 2,000 Monte Carlo draws the tiles are resampled with replacement before $\hat{p}$ is computed, so the spread reflects how error rates vary from place to place. In the same draw, AGB comes from the biomass model and R, CF and SOC from triangular distributions over the IPCC low, best and high values. The response gives the median and the 5th–95th percentile of every area and carbon quantity. The *conservative creditable amount* is the 5th percentile of biomass CO_2e.

J. Change detection

For two years y_1 and y_2 the system builds a calendar-year composite of each on the same grid, classifies both, keeps pixels that are clear in both, and forms a 4×4 transition matrix in hectares. Gains and losses per class are the off-diagonal column and row sums. Carbon change is the difference of the two stock estimates for each scenario, and a change map colours mangrove and tree gain and loss.

TABLE VI: RESTORATION SUITABILITY CRITERIA.

Criterion	Weight	Data	Score
Former mangrove	0.35	GMW v3, 1996	1 if mangrove in 1996
Tidal flooding	0.25	JRC occurrence	1 at 5–40%, 0 at 75%, 0.25 if never
Low elevation	0.20	GLO-30 DEM	1 at ≤ 3 m, 0 at 10 m
Near mangrove	0.20	Model + GMW 2020	exp(−d / 500 m)

Weights are expert judgement and are kept in one place in the code so they can be revised.

K. Restoration planner

The planner looks for land where mangrove could be replanted. Pixels that are already mangrove (in the model, GMW 2020 or WorldCover), tree cover, built-up, permanently flooded (water in at least 90% of observations) or higher than 10 m are excluded. The rest receive a suitability score from the four criteria of Table VI:

$$s = 0.35\,F + 0.25\,T + 0.20\,E + 0.20\,e^{-d/500}, \quad (11)$$

Where F marks GMW mangrove in 1996, following the restoration-potential idea of Worthington and Spalding [28], T scores tidal flooding from JRC surface-water occurrence [29], E scores elevation from the Copernicus GLO-30 DEM [30], and d is the distance in metres to living mangrove. Pixels with $s \geq 0.45$ are grouped into 8-connected patches, patches under 0.2 ha are dropped, and the rest are ranked by score times area; mean scores of 0.60 or more are labelled High. Each patch lists its current cover, why it scored well, and cautions to check in the field.

For one replanted hectare, above-ground biomass grows towards the mature biomass M of the surrounding mangrove, taken from local GEDI footprints:

$$\mathrm{AGB}(t) = M\,(1 - \mathrm{e}^{-Gt/M}), \quad (12)$$

$$\mathrm{CO_2e}(t) = \eta\,(44/12)\,[\mathrm{AGB}(t)(1+R)\,\mathrm{CF} + r_s t], \quad (13)$$

Where G is the initial growth rate (3.3, 9.9 or 18.1 t/ha/yr; [5], Table 4.4), r_s = 1.62 t C/ha/yr is soil accumulation (Table 4.12, 95% CI 1.3–2.0) and η is the planting success rate set by the user (default 0.5). Monte Carlo draws over M, G, R, CF and r_s give the median and 90% interval for each year. These are gross removals: carbon already on site and natural regrowth without planting are not subtracted.

L. Persistence and reporting

Each result is stored in a SQLite table with its polygon, a short summary and the full JSON response, and the History page lists, reopens, renames and deletes analyses. A ReportLab PDF report gives the imagery details, class areas, carbon scenarios, biomass, verification, uncertainty, coefficient sources and map layers.

VI. Application Workflow

Fig. 2 traced one request through the backend. From the user's side, the workflow is to open the Analysis page, draw a polygon or upload a boundary, and press *Analyze Selected Area*. During development an analysis typically took 5–20 s, mostly waiting for Earth Engine. The result shows class areas, carbon for the three scenarios, a confidence badge with flags, the biomass comparison with Tier 1 and the uncertainty intervals, while land-cover and biomass layers appear on the map. The user can then compare two years, run the restoration planner, or download the PDF report; every run appears on the History page.

VII. Testing, Results and Observations

A. Automated tests

The backend has 64 pytest cases (52 test functions, some parametrised) in nine modules: request validation (5), API status codes, saving and PDF reports (9), the analysis service including cloud rejection, change detection and fallbacks (10), carbon arithmetic (6), biomass calibration (4), uncertainty (3), verification flags and levels (6), restoration (7) and the predictor (2). They use synthetic band arrays and a stubbed Earth Engine service, so they run offline; the full suite passed in 7.3 s on the development machine. A smoke script also runs the model and carbon calculation on a bundled Sentinel-2 tile without network access.

B. Land-cover accuracy on held-out regions

Table VII compares three classifiers on the 60 test tiles, which cover all four seasons. Woody cover means mangrove plus tree cover, the only classes that carry carbon. The four-class U-Net reaches an mIoU of 0.643 and a pixel accuracy of 81.4%. For mangrove, its IoU of 0.587 is five times that of the NDVI threshold, which finds 94% of mangrove pixels but labels so much other vegetation that its precision is only 0.121. The original binary U-Net shows the opposite failure: it is precise (0.953) but finds only 22.7% of woody pixels in regions it was not trained on, although it scored 0.883 on tiles from its own region. Mangrove IoU per test estuary ranges from 0.551 to 0.626, except at Netravati (0.278), where mangroves occur as narrow fringes only a few pixels wide.

TABLE VII: Segmentation Results on Held-Out Test Regions.

Metric	U-Net (ours)	Old U-Net	NDVI > 0.3
Mangrove IoU	**0.587**	n/a	0.120
Mangrove precision/recall	0.789/0.696	n/a	0.121/0.944
Woody IoU	**0.721**	0.225	0.616
Woody precision/recall	0.811/0.866	0.953/0.227	0.630/0.966
Woody IoU, by season	0.701–0.756	0.111–0.353	0.563–0.691
IoU: non-veg./tree/other	0.809/0.689/0.487	n/a	n/a
Four-class mIoU	0.643	n/a	n/a
Pixel accuracy	81.4%	n/a	n/a

Reference: ESA WorldCover 2021 labels on 60 tiles from 5 test regions. n/a: the method has no such class.

C. Biomass validation

TABLE VIII: BIOMASS VALIDATION BY SPATIAL BLOCK CROSS-VALIDATION.

Method	Mangr. error %	Mangr. RMSE	Tree error %	Tree RMSE
IPCC Tier 1	1045	165.8	241	122.5
India-wide height	61.6	26.4	53.0	38.1
Local GEDI mean	**17.6**	16.8	16.7	28.4
Local GEDI + height	20.0	**15.3**	**14.5**	**21.8**

Error is the median absolute percentage error per 50 km block; RMSE in Mg/ha. The reference is GEDI L4A, itself a model product, not field plots.

The biomass methods were compared by spatial block cross-validation [31]: India was divided into 50 km blocks, and for each block the model was calibrated without that block's footprints and then compared with the block's mean GEDI AGB. Table VIII and Fig. 5 show the results for 186 mangrove and 745 tree-cover blocks. Tier 1 overstates biomass by a median of 1045% for mangrove and 241% for tree cover. Local calibration removes most of that bias. For mangrove, the plain mean of nearby footprints has a slightly lower median error (17.6%) than the height model (20.0%), but the height model has the lower RMSE and the higher correlation (r = 0.91 against 0.89), and it gives a per-pixel map. For tree cover the height model is best on all three measures (r = 0.92).

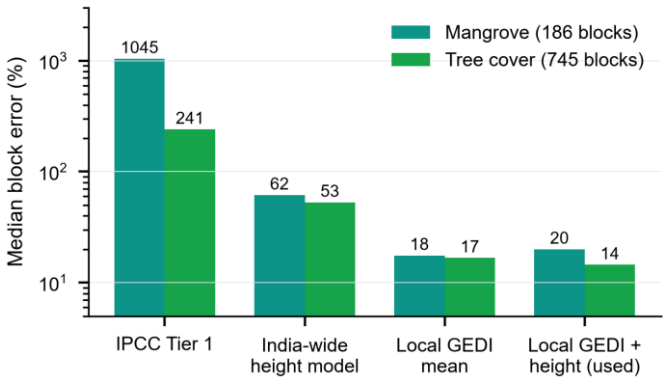

Fig. 5. Median block error of four biomass methods (log scale).

D. Case studies

Table IX summarises full analyses of three mangrove sites, all run with the final model on imagery from November 2025 to November 2026. The Kali estuary near Karwar lies in a test region, Kunhimangalam in a validation region and Bhitarkanika in a training region, so only Kali is a fully independent case. Fig. 6 shows the Kali composite, land cover, biomass map and a 2018–2026 change map.

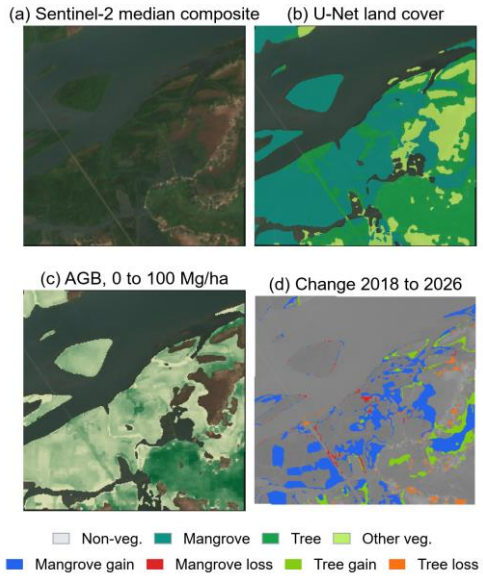

Fig. 6. Kali estuary, Karwar (685 ha): (a) median of 42 Sentinel-2 scenes; (b) U-Net classes; (c) per-pixel AGB from canopy height; (d) land-cover change between the 2018 and 2026 composites.

TABLE IX: RESULTS AT THREE MANGROVE SITES.

Quantity	Kali	Bhitarkan.	Kunhiman.
Area (ha); scenes	685; 42	400; 87	200; 34
Mangrove: model (ha)	211.7	296.2	27.9
Mangrove: WorldCover (ha)	241.0	253.2	36.4
Mangrove: GMW 2020 (ha)	174.7	341.1	38.3
Mangrove IoU: model, GMW	0.555	0.771	0.381
Mangrove IoU: WC, GMW	0.547	0.662	0.347
Woody IoU: model, WC	0.852	0.944	0.891
Confidence (score)	High (0.82)	High (0.88)	High (0.77)
Mangrove AGB (Mg/ha)	22.2	31.2	20.4
GEDI footprints, mangrove	46	2,988	282
Total carbon, best (t C)	89,742	122,598	14,369
Monte Carlo total C (t C)	75,645	98,309	13,466
90% interval ($\times 10^3$)	63.4–88.2	82.8–114.9	11.0–16.3
Biomass CO_2e, height (t)	29,422	30,279	13,162
Biomass CO_2e, Tier 1 (t)	154,311	163,799	53,913
Conservative credit (t)	17,301	14,764	6,967

Kali estuary: test region; Kunhimangalam: validation region; Bhitarkanika: training region. WC: WorldCover. Credits are conservative biomass CO_2e (5th percentile).

Three observations stand out. First, at Kali the model agrees with the independent GMW map (IoU 0.555) as well as WorldCover, its own label source, does (0.547); the references disagree with each other as much as the model disagrees with them. Second, the height-based

biomass CO_2e is 4.1–5.4 times lower than Tier 1 because the mapped mangroves average 20–31 Mg/ha of AGB rather than 192, so mangrove soil carbon, which still uses the Tier 1 range, makes up 91% of the best-case total at Kali. Third, the bias correction lowers the Kali mangrove area from 211.7 ha to a median of 178.3 ha (90% interval 164.2–191.9 ha), because on the test tiles the model tends to extend mangrove into neighbouring classes. The conservative credit is therefore 59% of the best-case biomass CO_2e. At an inland control site (Turahalli, 11.8 ha) the model and both references mapped no mangrove, and the system reported the estimate as terrestrial tree cover only.

E. Comparison with field measurements

Bindu *et al.* [3] measured 12.67 t of biomass carbon on 4.86 ha of plots at Kunhimangalam, about 2.6 t C/ha. We drew a 4.86 ha square at the study site; the model mapped 3.79 ha of it as mangrove with High confidence (mangrove IoU 0.93 with WorldCover and 0.78 with GMW). Its mangrove biomass carbon density is 15.3 t C/ha, against 129.0 t C/ha with Tier 1: about six and fifty times the field value, respectively. Local calibration thus removes most, but not all, of the overestimate. Likely causes are that the square is not the plot boundary, that canopy height is from 2020, and that GEDI L4A has no mangrove-specific allometry, whereas mangrove wood density and form differ from those of terrestrial trees [32].

F. Change detection

For the Kali estuary the 2018 composite (17 scenes) and the 2026 composite (34 scenes) give a net mangrove increase from 145.4 to 219.7 ha, with 77.6 ha gained and 3.3 ha lost, while tree cover changed little (−4.5 ha net; Fig. 6(d)). We have no reference change data for this site, and the 2018 composite rests on half as many scenes, so the figure is indicative only.

G. Restoration planning

Table X lists the planner's output for four sites with a 50% planting success rate and a 20-year horizon. Only Thane Creek contains much former mangrove: 19.0 of its 52.3 suitable hectares were mangrove in 1996, and its two highest-ranked patches are ponds within today's mangrove (Fig. 7a). Because mature local mangroves hold only 17.6 Mg/ha of AGB there (879 footprints), biomass reaches 90% of that level within about 4 years (90% interval 2.6–7.3), and after 20 years 74% of the 81.0 t CO_2e removed per hectare is in the soil (Fig. 7b). The narrow intervals reflect coefficient uncertainty only; they do not include the uncertainty of the suitability map itself.

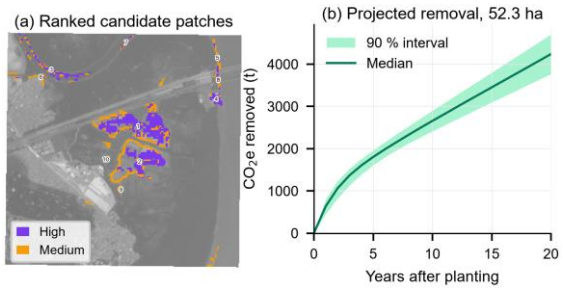

Fig. 7. Restoration planner at Thane Creek, Mumbai: (a) candidate patches ranked by score and area; (b) projected CO_2e removal of the 52.3 suitable hectares.

TABLE X: RESTORATION PLANNER RESULTS. (20 YEARS, 50% SUCCESS)

Site (ha)	Suitable (High), ha	Former mangr., ha	Patches	CO_2e removed, t (90% CI)
Thane (1,000)	52.3 (31.3)	19.0	23	4,238 (3,768–4,703)
Kali (685)	44.3 (17.3)	0.1	37	4,503 (4,018–4,997)
Kunhiman. (200)	17.6 (6.7)	0.0	13	1,687 (1,525–1,847)
Bhitarkan. (400)	9.0 (2.3)	0.3	12	823 (747–907)

VIII. DISCUSSION

Evaluating on unseen regions. The original U-Net looked excellent on its own validation tiles and failed elsewhere. Splitting by region, not by tile, was the single most useful decision in this project, because neighbouring tiles share scenes, season and landscape and make a model look better than it is. The same reasoning led to block cross-validation for biomass. Against NDVI, which only measures greenness, the network gains from the shortwave-infrared bands, MNDWI and spatial context such as nearby water and creeks.

Default factors and Indian mangroves. The median GEDI mangrove AGB in India is 13.8 Mg/ha, and near 2 Mg/ha in arid Gujarat creeks, against a Tier 1 value of 192, so Tier 1 overstates such stands by an order of magnitude. Local calibration closes most of that gap; the remaining factor of about six at Kunhimangalam points to the lack of a mangrove model in GEDI L4A [16].

What verification can and cannot say. At our sites the two reference maps agree with each other only at IoU 0.35–0.66. A High confidence level therefore means that the model is consistent with the best available maps and with clear imagery, not that it is correct. The weights in (9) are design choices and the score is not a calibrated probability.

Uncertainty and credits. The intervals omit the spatial variation of soil carbon, the age gap between the 2020–2021 reference layers and current imagery, and GEDI's own error, so they are likely too narrow. The conservative amount is also a stock, whereas credits are issued for

additional, permanent removals against a baseline, net of leakage. None of these is modelled, and the interface states that its figures are not verified credits.

IX. Applications and Real-World Relevance

The system is suited to screening rather than certification. Coastal and forest departments can compare sites and track change in mangrove cover, and project developers can check a candidate area and decide where field surveys or replanting should start.

X. Limitations and Future Work

Training and test labels come from WorldCover, so the reported IoU measures agreement with that product, not with ground truth, and inherits its errors. The test set is 60 tiles from five regions in one state, and only one field site was available. Soil carbon uses the Tier 1 range for every mangrove pixel. Seagrass and salt marsh are not mapped, because optical imagery sees little of submerged vegetation and no suitable labels were available. The restoration weights have not been tested against restoration outcomes. Storage is a single-user SQLite file without authentication, and the Docker set-up has not yet been tested.

Future work will (1) collect field plots at several Karnataka estuaries to validate both area and biomass; (2) fit a mangrove-specific height–biomass relation; (3) add Sentinel-1 radar for cloudy monsoon months; (4) map soil carbon spatially instead of using one range; (5) validate change detection and the restoration ranking against GMW change data and restored sites; and (6) add user accounts and a server database for multi-user deployment. In short, every module in Fig. 1 is implemented; area, biomass and segmentation are validated as reported here, while change detection, the restoration ranking and the Docker deployment are not yet validated.

XI. Conclusion

Carbon Insight turns a polygon on a map into a mangrove map, a carbon estimate with an uncertainty range and a verification report from open data. Testing on regions the model never saw showed that a learned four-class model separates mangrove far better than NDVI (IoU 0.587 against 0.120), and that local GEDI calibration of canopy height cuts biomass error from over 1000% to about 20% relative to GEDI. The comparison with field plots shows that the estimates are still several times too high for scrub mangroves, so the system is best used for screening, monitoring change and choosing where to measure, with field plots remaining the reference.

ACKNOWLEDGMENT

The authors thank their guide, Mrs. Pallavi R, the project coordinator, Dr. Abhishek Kulkarni, and the Head of the Department, Dr. Prabhanjan S, of Jyothy Institute of Technology for their guidance. We thank ESA and Copernicus, NASA, ETH Zurich, the Global Mangrove Watch consortium, the JRC and Google Earth Engine for open data and computing. An AI assistant (Claude, Anthropic) was used to help write and refactor parts of the software, to run and summarise experiments, and to draft and edit the text of this paper. The authors checked every number against the project's code and output files and take full responsibility for the content.

REFERENCES

1. E. Mcleod *et al.*, "A blueprint for blue carbon: Toward an improved understanding of the role of vegetated coastal habitats in sequestering CO_2," *Front. Ecol. Environ.*, vol. 9, no. 10, pp. 552–560, 2011, doi: 10.1890/110004.
2. D. C. Donato, J. B. Kauffman, D. Murdiyarso, S. Kurnianto, M. Stidham, and M. Kanninen, "Mangroves among the most carbon-rich forests in the tropics," *Nature Geosci.*, vol. 4, no. 5, pp. 293–297, 2011, doi: 10.1038/ngeo1123.
3. G. Bindu, P. Rajan, E. S. Jishnu, and K. A. Joseph, "Carbon stock assessment of mangroves using remote sensing and geographic information system," *Egypt. J. Remote Sens. Space Sci.*, vol. 23, no. 1, pp. 1–9, 2020, doi: 10.1016/j.ejrs.2018.04.006.
4. IPCC, *2006 IPCC Guidelines for National Greenhouse Gas Inventories*, vol. 4, *Agriculture, Forestry and Other Land Use*, H. S. Eggleston, L. Buendia, K. Miwa, T. Ngara, and K. Tanabe, Eds. Hayama, Japan: IGES, 2006.
5. IPCC, *2013 Supplement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories: Wetlands*, T. Hiraishi *et al.*, Eds. Geneva, Switzerland: IPCC, 2014.
6. D. N. Gowda, A. M. Gourang, R. Hemanth, J. Manu, and R. Pallavi, "Carbon Insight: A comprehensive survey of AI-based carbon estimation and verification systems for blue carbon ecosystems," *Int. J. Res. Publ. Anal.*, vol. 2, no. 5, pp. 1–12, 2026.
7. A. Dutta Roy *et al.*, "Remote sensing-based mangrove blue carbon assessment in the Asia-Pacific: A systematic review," *Sci. Total Environ.*, vol. 938, Art. no. 173270, 2024, doi: 10.1016/j.scitotenv.2024.173270.
8. K. Zhang, "Mangrove biomass estimation through remote sensing and machine learning based approaches," Ph.D. dissertation, School Geogr. Earth Sci., Univ. Glasgow, Glasgow, U.K., 2025.

9. R. Zhang and J. Fan, “Classification and carbon-stock estimation of mangroves in Dongzhaigang based on multi-source remote sensing data using Google Earth Engine,” *Remote Sens.*, vol. 17, no. 6, Art. no. 964, 2025, doi: 10.3390/rs17060964.
10. O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *Proc. MICCAI*, LNCS vol. 9351, 2015, pp. 234–241, doi: 10.1007/978-3-319-24574-4_28.
11. W. S. Maung, S. Tsuyuki, and Z. Guo, “Improving land use and land cover information of Wunbaik mangrove area in Myanmar using U-Net model with multisource remote sensing datasets,” *Remote Sens.*, vol. 16, no. 1, Art. no. 76, 2024, doi: 10.3390/rs16010076.
12. D. Zanaga *et al.*, “ESA WorldCover 10 m 2021 v200,” Zenodo, 2022, doi: 10.5281/zenodo.7254221.
13. P. Bunting *et al.*, “Global mangrove extent change 1996–2020: Global Mangrove Watch version 3.0,” *Remote Sens.*, vol. 14, no. 15, Art. no. 3657, 2022, doi: 10.3390/rs14153657.
14. N. Lang, W. Jetz, K. Schindler, and J. D. Wegner, “A high-resolution canopy height model of the Earth,” *Nature Ecol. Evol.*, vol. 7, no. 11, pp. 1778–1789, 2023, doi: 10.1038/s41559-023-02206-6.
15. R. Dubayah *et al.*, “The Global Ecosystem Dynamics Investigation: High-resolution laser ranging of the Earth’s forests and topography,” *Sci. Remote Sens.*, vol. 1, Art. no. 100002, 2020, doi: 10.1016/j.srs.2020.100002.
16. L. Duncanson *et al.*, “Aboveground biomass density models for NASA’s Global Ecosystem Dynamics Investigation (GEDI) lidar mission,” *Remote Sens. Environ.*, vol. 270, Art. no. 112845, 2022, doi: 10.1016/j.rse.2021.112845.
17. P. Olofsson, G. M. Foody, M. Herold, S. V. Stehman, C. E. Woodcock, and M. A. Wulder, “Good practices for estimating area and assessing accuracy of land change,” *Remote Sens. Environ.*, vol. 148, pp. 42–57, 2014, doi: 10.1016/j.rse.2014.02.015.
18. N. Gorelick, M. Hancher, M. Dixon, S. Ilyushchenko, D. Thau, and R. Moore, “Google Earth Engine: Planetary-scale geospatial analysis for everyone,” *Remote Sens. Environ.*, vol. 202, pp. 18–27, 2017, doi: 10.1016/j.rse.2017.06.031.
19. A. Paszke *et al.*, “PyTorch: An imperative style, high-performance deep learning library,” in *Proc. NeurIPS*, vol. 32, 2019, pp. 8024–8035.
20. M. Drusch *et al.*, “Sentinel-2: ESA’s optical high-resolution mission for GMES operational services,” *Remote Sens. Environ.*, vol. 120, pp. 25–36, 2012, doi: 10.1016/j.rse.2011.11.026.
21. C. J. Tucker, “Red and photographic infrared linear combinations for monitoring vegetation,” *Remote Sens. Environ.*, vol. 8, no. 2, pp. 127–150, 1979, doi: 10.1016/0034-4257(79)90013-0.

22. H. Xu, “Modification of normalised difference water index (NDWI) to enhance open water features in remotely sensed imagery,” *Int. J. Remote Sens.*, vol. 27, no. 14, pp. 3025–3033, 2006, doi: 10.1080/01431160600589179.
23. F. Milletari, N. Navab, and S.-A. Ahmadi, “V-Net: Fully convolutional neural networks for volumetric medical image segmentation,” in *Proc. 4th Int. Conf. 3D Vision (3DV)*, 2016, pp. 565–571, doi: 10.1109/3DV.2016.79.
24. I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2019.
25. L. N. Smith and N. Topin, “Super-convergence: Very fast training of neural networks using large learning rates,” in *Proc. SPIE 11006, Artif. Intell. Mach. Learn. Multi-Domain Oper. Appl.*, 2019, Art. no. 1100612.
26. N. Duan, “Smearing estimate: A nonparametric retransformation method,” *J. Amer. Statist. Assoc.*, vol. 78, no. 383, pp. 605–610, 1983.
27. M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, “The PASCAL visual object classes (VOC) challenge,” *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, 2010, doi: 10.1007/s11263-009-0275-4.
28. T. Worthington and M. Spalding, “Mangrove restoration potential: A global map highlighting a critical opportunity,” Univ. Cambridge and The Nature Conservancy, Cambridge, U.K., Tech. Rep., 2018.
29. J.-F. Pekel, A. Cottam, N. Gorelick, and A. S. Belward, “High-resolution mapping of global surface water and its long-term changes,” *Nature*, vol. 540, no. 7633, pp. 418–422, 2016, doi: 10.1038/nature20584.
30. European Space Agency, “Copernicus Global Digital Elevation Model (GLO-30),” 2021. [Online]. Available: https://dataspace.copernicus.eu
31. D. R. Roberts *et al.*, “Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure,” *Ecography*, vol. 40, no. 8, pp. 913–929, 2017, doi: 10.1111/ecog.02881.
32. A. Komiyama, S. Poungparn, and S. Kato, “Common allometric equations for estimating the tree weight of mangroves,” *J. Trop. Ecol.*, vol. 21, no. 4, pp. 471–477, 2005, doi: 10.1017/S0266467405002476.