International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 6, Nov-Dec-2026, ISSN (Online): 2395-566X

Carbon Credit Marketplace with Multi-Modal Verification for Smallholder Agroforestry - Satellite NDVI, Ground-Photo AI and Soil Carbon

Rudramurthy V C[1], Chiranthan S V[2], Dheemanth J Nirvani[2], Samarth Kumble[2],
[1]Dept. of ISE, BMS College of Engineering, Bengaluru, Karnataka, India
[2]Dept. of CSBS, BMS College of Engineering, Bengaluru, Karnataka, India

Abstract — Smallholder agroforestry plots rarely market, because the measurement, reporting and verification (MRV) step that backs each credit still rests on field visits that are too expensive to run on a one- or two-acre holding, in many respects; indeed, the platform is a FastAPI backend, a React web app and an Expo mobile app, with 139 automated tests covering it, to a significant degree. A rule-based fraud engine, role-separated farmer, auditor and buyer portals. The check is fusing satellite vegetation data (NDVI computed from real Sentinel-2 imagery inside the boundary the farmer draws), a ground-photo classifier (MobileNetV3-Small, 98.4% accuracy on 1,420 held-out real photos) and a soil organic carbon (SOC) reading into one weighted score, 0.40 NDVI + 0.35 photo + 0.25 SOC. A score of 75 or above is permitted, 55 to 75 is routed to a human auditor; indeed, we built a working prototype of a carbon-credit marketplace in which a plot can mint credits only after it clears an automated, explainable three-part check. And anything under 55 is rejected; if any single piece of indication is missing the plot stays pending, and no value is ever invented, and an Ethereum-style ledger (Ganache) are logging each credit from issuance, through the transfer to retirement. In an end-to-end run on live data, a real areca plot in Shivamogga, Karnataka created an NDVI of 0.794 from a Sentinel-2 scene dated 27 September 2026, was authorized at 84.9/100, minted a 20 tCO2 credit; and thus, as, well, that credit was then bought and retired.

Keywords— carbon credits, agroforestry, measurement reporting and verification, Sentinel-2, NDVI, MobileNetV3, soil organic carbon, fraud detection, blockchain.

I. INTRODUCTION

A carbon credit is one tonne of carbon dioxide equal (tCO_2) pulled out of, or prevented from entering, the atmosphere; these credits are bought by firms to cover emissions they can't yet cut. On one or two acres that cost is shared across the very few credits and is rarely worth paying. Each credit has to sit behind measurement, reporting and verification (MRV), so on top of that, site visits by accredited verifiers are meant by standard MRV. Trees grown on farmland alongside crops (agroforestry, e.g., areca, coconut and teak) lock up carbon year after year, so smallholders could in tenet earn money from credits. What's more, in the real world, small farms almost never qualify.

1. Problem Statement

A cheap way is being needed by smallholder plots to show that the trees they're listing actually exist, are healthy, and stand on the land they're saying they do, without someone having to visit every plot, on the whole. The proof has to be hard to fake (re-used photos, inflated counts, wrong coordinates), in real-world scenarios. As, well, each decision has to be something an auditor and a buyer can, to some extent, follow step by step.

2. Contributions

The advancements of this work are:

- A three-modality verification engine that's fusing satellite NDVI, ground-photo category and SOC into a single explainable score, with fixed rules for approval, review and rejection, in actual application, and what's more, one rule it never bends - it doesn't replace a default for missing proof, broadly speaking.
- A satellite pipeline that computes NDVI from real Sentinel-2 L2A pixels falling inside the polygon the farmer draws, with cloud masking, retries and recorded provenance.
- A multi-photo indication model that's holding up against copy and unrelated images, paired with a rule-based fraud and risk engine, in real-world scenarios.
- A full-stack marketplace with separate farmer, auditor and buyer portals, whole-lot purchase and retirement of credits, and a smart-contract ledger that's keeping a tamper-evident history, in many respects.
- An end-to-end check on live data, covering a photo classifier that's reaching 98.4% accuracy on a held-out real-photo test set, to a significant degree.

© 2026 IJSRET

3. Scope

This is a prototype, not an accredited carbon registry; it targets small agroforestry plots in India (the examples come from Dakshina Kannada and Shivamogga, Karnataka), broadly speaking. What's more, the blockchain is a local development chain; in turn, carbon quantities rest on a straightforward per-tree estimate, payments are currently copied, to a significant degree.

Background

Carbon Credits and MRV

Believable carbon credits hinge on three features: the carbon is real, extra, and not double-counted [1]; forestry and land-use carbon accounting guidance is set out in [2], in real-world scenarios; MRV is what guarantees the first. What the work here automates is the verification step with remote sensing and machine learning, leaves a human auditor in the loop for uncertain cases, and employs a ledger to head off double counting, on the whole.

NDVI and Sentinel-2

The Normalised Difference Vegetation Index [3] reads vegetation density off two spectral bands:

$$NDVI = (NIR - Red) / (NIR + Red) \quad (1)$$

Healthy leaves bounce back a lot of near-infrared (NIR) light and soak up red light, so a packed canopy sits around 0.7 to 0.9 on the NDVI scale. Already, bare soil and buildings fall under 0.3. Sentinel-2 is a pair of ESA satellites that revisit each spot roughly every five days at 10 m resolution [4]. Its Level-2A product reflects for band B04 (red), band B08 (NIR) and a scene classification layer (SCL) that tags each pixel as vegetation, cloud, shadow, water and so on; these files are pulled from Microsoft Planetary Computer, a free public STAC catalogue [5].

Ground Evidence and Soil Carbon

Satellite pixels are 10 m extensive and can't tell an areca plantation apart from any other green cover, in real-world scenarios. A ground photo is showing what's actually increasing on the plot. A compact convolutional network, MobileNetV3-Small [6], is sorting each photo into the plantation, non-plantation or unclear, in real-world scenarios. Soil organic carbon (SOC, given as a percentage of soil mass) monitors stored carbon and soil health; characteristic Indian farm soils sit between roughly 0.5% and 2.5%. The SOC value is being supplied from a soil test by the farmer, along with sampling depth and soil type.

Blockchain Registries

A blockchain is an append-only ledger whose past can't be changed without leaving a trace [7], in real-world scenarios; registries are using it to record who owns. What's more, who has retired each credit, so the same credit can't be sold twice, in real-world scenarios; and thus, in this effort a Solidity smart contract is rolled out on Ganache, a local Ethereum development chain.

II. SYSTEM ARCHITECTURE

Fig. 1 lays out the architecture. Every business rule (verification, fraud checks, credit issuance, marketplace eligibility) lives in the backend, so the web and mobile clients have no way around them, in real-world scenarios.

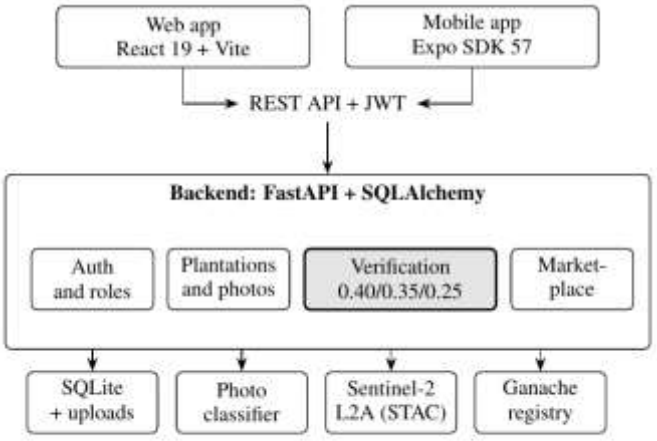


Figure 1. System architecture. The verification engine is the only path from a registered plot to an issued credit.

Healthy leaves reflect near-infrared (NIR) light strongly and absorb red light, so dense canopy scores about 0.7–0.9 while bare soil and buildings score below 0.3. Sentinel-2 is a pair of European Space Agency satellites imaging each location about every five days at 10 m resolution [4]. Its Level-2A product provides surface reflectance for band B04 (red), band B08 (NIR) and a

Figure 1. System architecture. The verification engine is the only path from a registered plot to an issued credit.

1. User Roles

Four functions are supported. Auditors are working through the verification queue, inspect every piece of proof. Admins can do all of that, plus create auditor accounts and issue credits themselves, so buyers browse listed credits, buy whole lots, retire them and track their portfolio. Farmers register plots with

© 2026 IJSRET

a drawn boundary, upload up to 10 photos and soil data, run verification, and issue credits from authorized plots. Either approve, hold in review, or reject with notes. The auditor portal shows verification work only; buyer accounts, buys and sales figures stay hidden. What's more, any such request is being rejected by the backend coming from an auditor.

2. End-to-End Flow

(1) The farmer draws the plot boundary on a satellite map; the backend stores the polygon and works out its subfield, in real-world scenarios. (5) An Authorized plot mints one credit lot, which is recorded on the blockchain when the node is up, to a significant degree. (2) The farmer uploads ground photos and a SOC value, and (4) Plots in REVIEW go to an auditor. (6) A buyer is buying the credit (ownership moves on chain) and may, in practice, retire it later. (3) The backend fetches a cloud-free Sentinel-2 scene, computes NDVI over the polygon, classifies every photo, scores SOC, runs fraud checks, and returns a score with a decision, to a significant degree.

3. Technology Stack

The backend runs on Python, FastAPI, Pydantic v2 and SQLAlchemy 2 over SQLite, with JWT bearer tokens and bcrypt hashing, on the whole; and thus, testing is employing pytest, oxlint, the TypeScript compiler and Playwright. The web client is React 19, Vite 8, Tailwind CSS 3 and Leaflet; the mobile client is Expo SDK 57 and React Native 0.86 with TypeScript, in many respects; satellite access goes through the Planetary Computer STAC API, an application programming interface, with rasterio/GDAL windowed reads. The photo model is built in PyTorch and torchvision, with imagehash for perceptual hashing; so, the ledger is a Solidity 0.8.20 contract on Ganache 7.9, reached through web3.py.

III. VERIFICATION METHODOLOGY

A plot is scored only when all three kinds of indication are provide and measurable, in real-world scenarios; and thus, the finding is then decided by the score. What's more, that outcome can (in practical terms) be pushed by a handful of safety rules only stricter, never looser, in actual application.

1. Required Evidence

Verification is needing (i) a plot boundary (a drawn polygon, or a centre point and field), (ii) at least one valid ground photo, (iii) a SOC value within 0 to 10%, and (iv) an NDVI value, either computed from Sentinel-2 or reported by the farmer with source and date, broadly speaking. No cloud-free scene, or the photo model won't run), the verification is PENDING, all scores stay empty and the reason is logged, if any item is missing or can't be measured (say, to a significant degree; nothing is filled in with a default.

2. Composite Score

A score on a 0 to 100 scale is being yielded by each mode, and the overall score is

$$S = 0.40 \cdot S_{NDVI} + 0.35 \cdot S_{photo} + 0.25 \cdot S_{SOC} \qquad (2)$$

Of course, supplied, Sphoto is probably the lower median of the unique photos' scores, when several photos are, broadly speaking, and Table I maps raw proof onto modality scores.

Table 1. Modality Scoring Bands

Modality	Score formation (0–100)
NDVI (0.40)	Mean NDVI of cloud-free pixels: ≥ 0.70 → 80–100; 0.50–0.70 → 65–80; 0.35–0.50 → 45–65; < 0.35 → 10–45
Photo (0.35)	Plantation → 70–98; non-plantation → 10–40; unclear → 15–35 (scaled by classifier confidence)
SOC (0.25)	≥ 2.2% → 85–100; 1.5–2.2% → 70–85; 0.8–1.5% → 52–70; < 0.8% → 15–52; adjusted by up to 10% for soil type

3. Decision Thresholds

The engine decision is

$$D = \text{APPROVED, if } S \geq 75;\ \text{REVIEW, if } 55 \leq S < 75;\ \text{REJECTED, if } S < 55 \qquad (3)$$

4. Safety Rules

An otherwise Authorized finding is normally capped at REVIEW. As, well, sent to an auditor whenever any of the following holds: (i) the risk score is typically 61 or above (Section 5.4); (ii) the NDVI was reported by the farmer rather than computed from satellite pixels; (iii) one photo is mainly classified plantation, in real-world scenarios; at the same time, another is confidently non-plantation; or (iv) the median photo, or most photos, come back with classifier confidence below 60%, in real-world scenarios.

5. Auditor Decisions and Evidence Snapshot

An auditor can maybe approve, hold in review, or reject the latest verification on a plot. Overriding the engine's decision writtenly notes, which are usually stored with the verification and in the audit log. What's more, approval is generally refused unless all three modality scores are often introduce. Once a credit has been issued from a plot, its indication and decision are, in effect, locked. Every verification also is keeping a snapshot of exactly what it employed: the boundary, tree count,

© 2026 IJSRET

soil values, the SHA-256 of each photo, every photo's classifier result, the satellite scene ID, acquisition date, cloud cover, clear-pixel share, NDVI minimum/mean/maximum, and the reasons behind the decision. A later reviewer can therefore still see why a plot was permitted, even if the farmer changes the indication afterwards.

IV. MODULE IMPLEMENTATION

1. Satellite NDVI Pipeline

NDVI is, in practice, computed from real Sentinel-2 pixels inside the plot in six stages.

- Scene search. A STAC query finds Sentinel-2 L2A scenes that intersect the polygon, acquired in the last 120 days with scene cloud cover below 20%. The least cloudy scenes are tried first (up to three).
- Access token. A single SAS token for the collection is fetched, cached until five minutes before expiry, and appended to each band URL.
- Windowed read. Only the window of each band covering the plot is downloaded (B04 and B08 at 10 m; SCL at 20 m, resampled to 10 m).
- Masking. Pixels outside the polygon, and pixels the SCL flags as no-data, saturated, cloud shadow, medium or high cloud, cirrus or snow, are thrown out. At least 50% of the plot has to be clear, or the scene is rejected.
- NDVI. Digital numbers are converted to reflectance (including the -1000 offset in use since processing baseline 04.00), NDVI is computed per pixel, and the mean, minimum, maximum and the share of pixels with NDVI at or above 0.40 are reported.
- Provenance. The result records the source (SENTINEL2_COMPUTED), scene ID, acquisition date, cloud cover, pixel counts and a catalogue link.

Timeouts and HTTP 429/5xx responses are typically retried up to three times with backoff, an expired token is getting one refresh, and every failure message names both the stage and the deviation. NDVI is perhaps unavailable and the plot stays PENDING, if no usable scene turns up. A farmer can (to some degree) instead enter a reported NDVI with its source and date; it gets scored but is always capped at REVIEW, to a significant degree.

2. Multi-Photo Evidence

Each photo has to decode as JPEG, PNG, WEBP or TIFF (the extension isn't trusted), be at most 15 MB, and have at least 64 px on its shorter side, broadly speaking; indeed, exact re-uploads are rejected, and up to 10 photos are normally kept per plot. The classifier is running on every valid photo; indeed, as, well, its class and confidence are stored per photo. The logic is deliberately conservative - one good photo can't lift a few weak or unrelated ones, and scores are often never summed, broadly speaking; the plot's photo score is the lower median of the unique photos' scores, and near-duplicate photos (perceptual-hash [9] distance of 6 or less) count once.

3. Soil Organic Carbon Scoring

The SOC value is mainly mapped onto a 0 to 100 score in four bands (Table I) and adjusted by up to 10% resting on the soil type's carbon retention (e.g., clay or laterite versus sand), in real-world scenarios; values outside 0 to 10% are generally rejected as invalid.

4. Fraud and Risk Engine

The risk engine adds penalty points for warning signs (Table II) and grades the total: below 31 is LOW, 31 to 60 is MEDIUM, to a significant degree. As, well, 61 or above is HIGH, which caps the decision at REVIEW.

Table 2. Fraud and Risk Penalties

Check	Points
Photo reused from another plot (perceptual hash)	+35
Photo classified non-plantation	+40
Photo classified unclear	+25
Plantation photo with confidence under 60%	+15
Dense canopy (NDVI ≥ 0.70) but photo score under 40	+30
Dense foliage in photo but satellite NDVI under 0.35	+25
Tree density above 2,500 per hectare	+25
Tree density below 20 per hectare	+15
SOC above 3.5% on sandy soil, or above 5.5% on any soil	+20
Coordinates overlap an already registered plot	+30

5. Carbon Estimation and Credit Issuance

A transparent per-tree estimate is being employed by the credit quantity,

$$Q = N \times 0.05 \times f_s \times f_p \quad [tCO_2e/year] \quad (4)$$

where N is the tree count, fs a species factor (e.g., teak 1.15, bamboo 1.30, mango 0.95) and fp a practice factor (e.g., 1.12 for biochar and composting), broadly speaking; in fact, a lot is currently minted only from the plot's latest APPROVED, fully

© 2026 IJSRET

scored verification; a zero-carbon result mints nothing, and each plot issues at most one lot, broadly speaking. Credits are, by design, priced at INR 1,500 per tCO2e.

6. Marketplace and Blockchain Ledger

A credit is listed only if it's available, its plot is checked; in fact, as, well, it comes from the plot's latest authorized, fully scored verification with recorded NDVI provenance, to a significant degree. Amounts are stored in kilograms of CO2e so that fractions stay exact; and thus, the CarbonCreditRegistry contract (Solidity 0.8.20) entries three events - issue, transfer, and retire. No real money moves. Buyers can retire a credit, which is ending its life cycle; buys are whole-lot only and run as an atomic database update, so two buyers can't grab the same credit (tested with five happening together buyers), broadly speaking. The same rule is re-checked at buy. Transaction hashes are shown only when a transaction was actually mined; if the chain is offline the credit is marked NOT_RECORDED, and what's more, unsold credits can (in actual operation) later be registered by a sync script. In this prototype, user wallets are custodial.

7. Security and Data Integrity

Parts. Every router checks the JWT function server-side; the frontend's role routing only is hiding views, in real-world scenarios. Farmers are seeing only their own plots, and buyers only their own buys. Auditor isolation. Transaction and holdings endpoints return 403 for reviewers and admins, and the auditor audit-log feed leaves out buyer accounts and purchase events; automated tests enforce this, broadly speaking. Private photos. Photos are stored under the server-generated names and served, through the short-lived HMAC-SHA256 signed links, only to users permitted to see that plot. No invented data. NDVI is recorded as UNAVAILABLE with the reason; farmer-reported NDVI is labelled REPORTED and can never generate an automatic approval; photo classes are shown as model predictions, and EXIF time and GPS as unverified, when imagery can't be fetched. Tamper evidence. Evidence snapshots, once-only minting per verification, a complete audit log and ledger records are making every state change traceable.

V. GROUND-PHOTO CLASSIFIER

1. Dataset

The classifier (version 2.0.0) is trained on real photographs from the public Intel Image Classification dataset [8], pinned to a fixed version for reproducibility, broadly speaking; and thus, the split is 6,127 training, 683 validation and 1,420 held-out test photos, the last of which is never employed in training, on the whole. The plantation class is built from "forest" (tree canopy) photos; non_plantation from "buildings", "street", "sea" and "glacier" photos; and unclear_evidence from real photos with added blur, darkness, overexposure or occlusion, in actual application.

2. Training Setup

The model is usually MobileNetV3-Small [6] with three output classes, taking 224x224 RGB input with ImageNet normalisation, and trains for 6 epochs with batch size 32 employing AdamW (learning rate 3e-4, weight decay 1e-4) on a cosine schedule, with label smoothing of 0.05 and random-resized-crop, horizontal-flip and colour-jitter augmentation, on the whole. On a laptop CPU, the principal processor, training takes about 12 minutes.

3. Results

Per-class findings and the confusion matrix are in Tables III and IV; so, the first version, trained only on computer-generated images, scored 49.6% accuracy, on the whole; precision is 98.38% on the held-out test set. As, well, macro F1 is 0.984; the best validation accuracy was 99.85%.

Table 3. Per-Class Test Results (1,420 Photos)

Class	Prec.	Recall	F1	Photos
plantation	0.981	0.992	0.986	474
non_plantation	0.977	0.981	0.979	472
unclear_evidence	0.994	0.979	0.986	474

Table 4. Confusion Matrix (Rows: True, Columns: Predicted)

True \ Pred.	non_plant	plant.	unclear
non_plantation	463	7	2
plantation	3	470	1
unclear_evidence	8	2	464

4. Limitations

The dataset holds no areca, coconut or agroforestry photos, so the model has learned "plantation" as tree canopy in broad. As well, its exactness on real plantation field photos hasn't been measured. Finally-to-end run it classified a real coconut plantation photo as plantation with 61.3% confidence, just above the 60% low-confidence cutoff. A set of labelled photos from real farms is the most valuable next step.

VI. TESTING AND RESULTS

1. Automated Tests

The backend pytest suite is running against a short-term SQLite database; and thus, as, well, 2 skipped tests (the two skipped

© 2026 IJSRET

tests are needing live network access to Planetary Computer), in actual application. The latest full run yielded 139 passed. The suite is covering the verification engine (weights, thresholds, PENDING on missing proof, and every REVIEW cap), satellite access (token caching and refresh, retries with backoff, filters, UNAVAILABLE provenance), multi-photo handling (limits, validation, signed links, duplicate detection, lower-median aggregation), auditor privacy, marketplace operations including once-only minting, and role checks on every protected route, in actual application.

2. Live Sentinel-2 Validation

A script is running each stage (token, STAC search, band reads, cloud masking, NDVI) separately against the real Planetary Computer service and reports where it crashes, in actual application; indeed, findings on real plots are in Table V. The scene ID returned for a plot was verified against the catalogue to validate that the image and date were genuine, to a significant degree.

Table 5. Live NDVI Validation on Real Plots

Plot	NDVI	Note
Kumble plantation	0.794	Dense canopy, consistent with photos
thanush farms	0.785	Dense canopy
Demo point plot	0.343	Sparse vegetation

3. End-to-End Demonstration

The full flow was run locally with real accounts; in fact, the auditor portal was supported to show only the overview and verification queue, with no buyer information, broadly speaking. For the Kumble plantation, the farmer registered the plot with photos, live NDVI was 0.794 and the composite score was 84.9; the auditor permitted it, a 20 tCO2e credit was minted, bought by a buyer and retired, to a significant degree; indeed, for thanush farms, NDVI was 0.785 and the score 81.8, but the risk engine flagged duplicate photos; these were counted once and the case was surfaced to the auditor.

4. Discussion

Autonomous signals are agreeing on real plantations. On the thin demo plot NDVI dropped to 0.34, pulling the score down as intended; combining three sources means one weak or manipulated input can't on its own push a plot, through to approval, in real-world scenarios, and what's more, the composite scores (84.9 and 81.8) cleared the approval threshold. On both thick plots, satellite NDVI (around 0.79) and the photo classifier pointed the same way, in many respects. A likely number is beaten by honest failure. What early tests showed was satellite NDVI as unavailable on some machines, caused by network-level token, timeout and temporary read deviations rather than the plots themselves. Or better yet than fall back to an estimate, the framework now retries, records why a fetch failed. As, well, keeps the case out of auto-approval, so every permitted credit can be traced back to a real Sentinel-2 scene.

The lower-median rule, the single counting of near-duplicates, and the REVIEW cap on disagreeing photos together mean a farmer can't lift a score by uploading many copies of one good picture, as the thanush farms are running showed in practice, in real-world scenarios. Robustness to photo manipulation.

Role separation. Keeping buyer data out of the auditor portal leaves the verifier autonomous of who buys the credits and at what price, which mirrors the separation of verification from trading in real registries, in actual application.

Conservative credits. The 0.05 t per tree per year base rate is a transparent rule of thumb, not a biomass inventory, so it's fitting a pilot rather than a certified offset market, to a significant degree.

Limitations and Future Work

The central shortcomings and the enhancements we plan are being listed by Table VI, on the whole.

Table 6. Limitations and Possible Improvements

Limitation	Improvement
SOC is a farmer-supplied soil-test value	Use SoilGrids or signed lab reports
Fixed per-tree carbon factors	Species allometric equations or canopy-height data (e.g., GEDI)
NDVI saturates in dense canopy and varies with season	Add EVI/NDRE, use time series and a baseline year for additionality
Photo model trained on general scenes	Fine-tune on auditor-labelled Indian field photos
EXIF GPS and time can be edited	Device-signed location and time in the mobile app
SQLite, single server	PostgreSQL/PostGIS and object storage
Local Ganache chain	Public testnet or permissioned ledger; third-party signatures
Simulated payments	Payment gateway and escrow
One auditor per decision	Two-auditor sign-off and periodic re-verification

© 2026 IJSRET

VII. CONCLUSION

We workingly marketed in which carbon credits are minted only after a plantation clears three autonomous checks: live Sentinel-2 NDVI, classified ground photos and soil organic carbon. With autonomous soil data and better biomass estimates, the same architecture could carry a real pilot run with farmer groups. Each decision is explainable, with the score breakdown, satellite scene, photo results and risk flags stored alongside it. Farmers, buyers and reviewers each are getting a portal restricted to what their function needs, and issued credits can be bought, retired and, when a local chain is running, recorded on a smart contract. Automated tests, a live satellite check and an end-to-end run on real plots are showing the pipeline working as designed. And when a source is unavailable, the framework won't invent data.

REFERENCES

1. Verra, Verified Carbon Standard (VCS) Program Guide.
2. IPCC, 2006 IPCC Guidelines for National Greenhouse Gas Inventories, Vol. 4: Agriculture, Forestry and Other Land Use (2019 Refinement).
3. J. W. Rouse, R. H. Haas, J. A. Schell and D. W. Deering, “Monitoring vegetation systems in the Great Plains with ERTS,” NASA Special Publication 351, 1974.
4. European Space Agency, Sentinel-2 User Handbook, ESA Standard Document, 2015.
5. Microsoft, “Planetary Computer: Sentinel-2 Level-2A collection and STAC API documentation.” [Online]. Available: planetarycomputer.microsoft.com
6. A. Howard et al., “Searching for MobileNetV3,” in Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV), 2019.
7. G. Wood, “Ethereum: A secure decentralised generalised transaction ledger,” Yellow Paper.
8. P. Bansal, “Intel Image Classification dataset,” Kaggle, 2018.
9. C. Zauner, “Implementation and benchmarking of perceptual image hash functions,” 2010.

© 2026 IJSRET