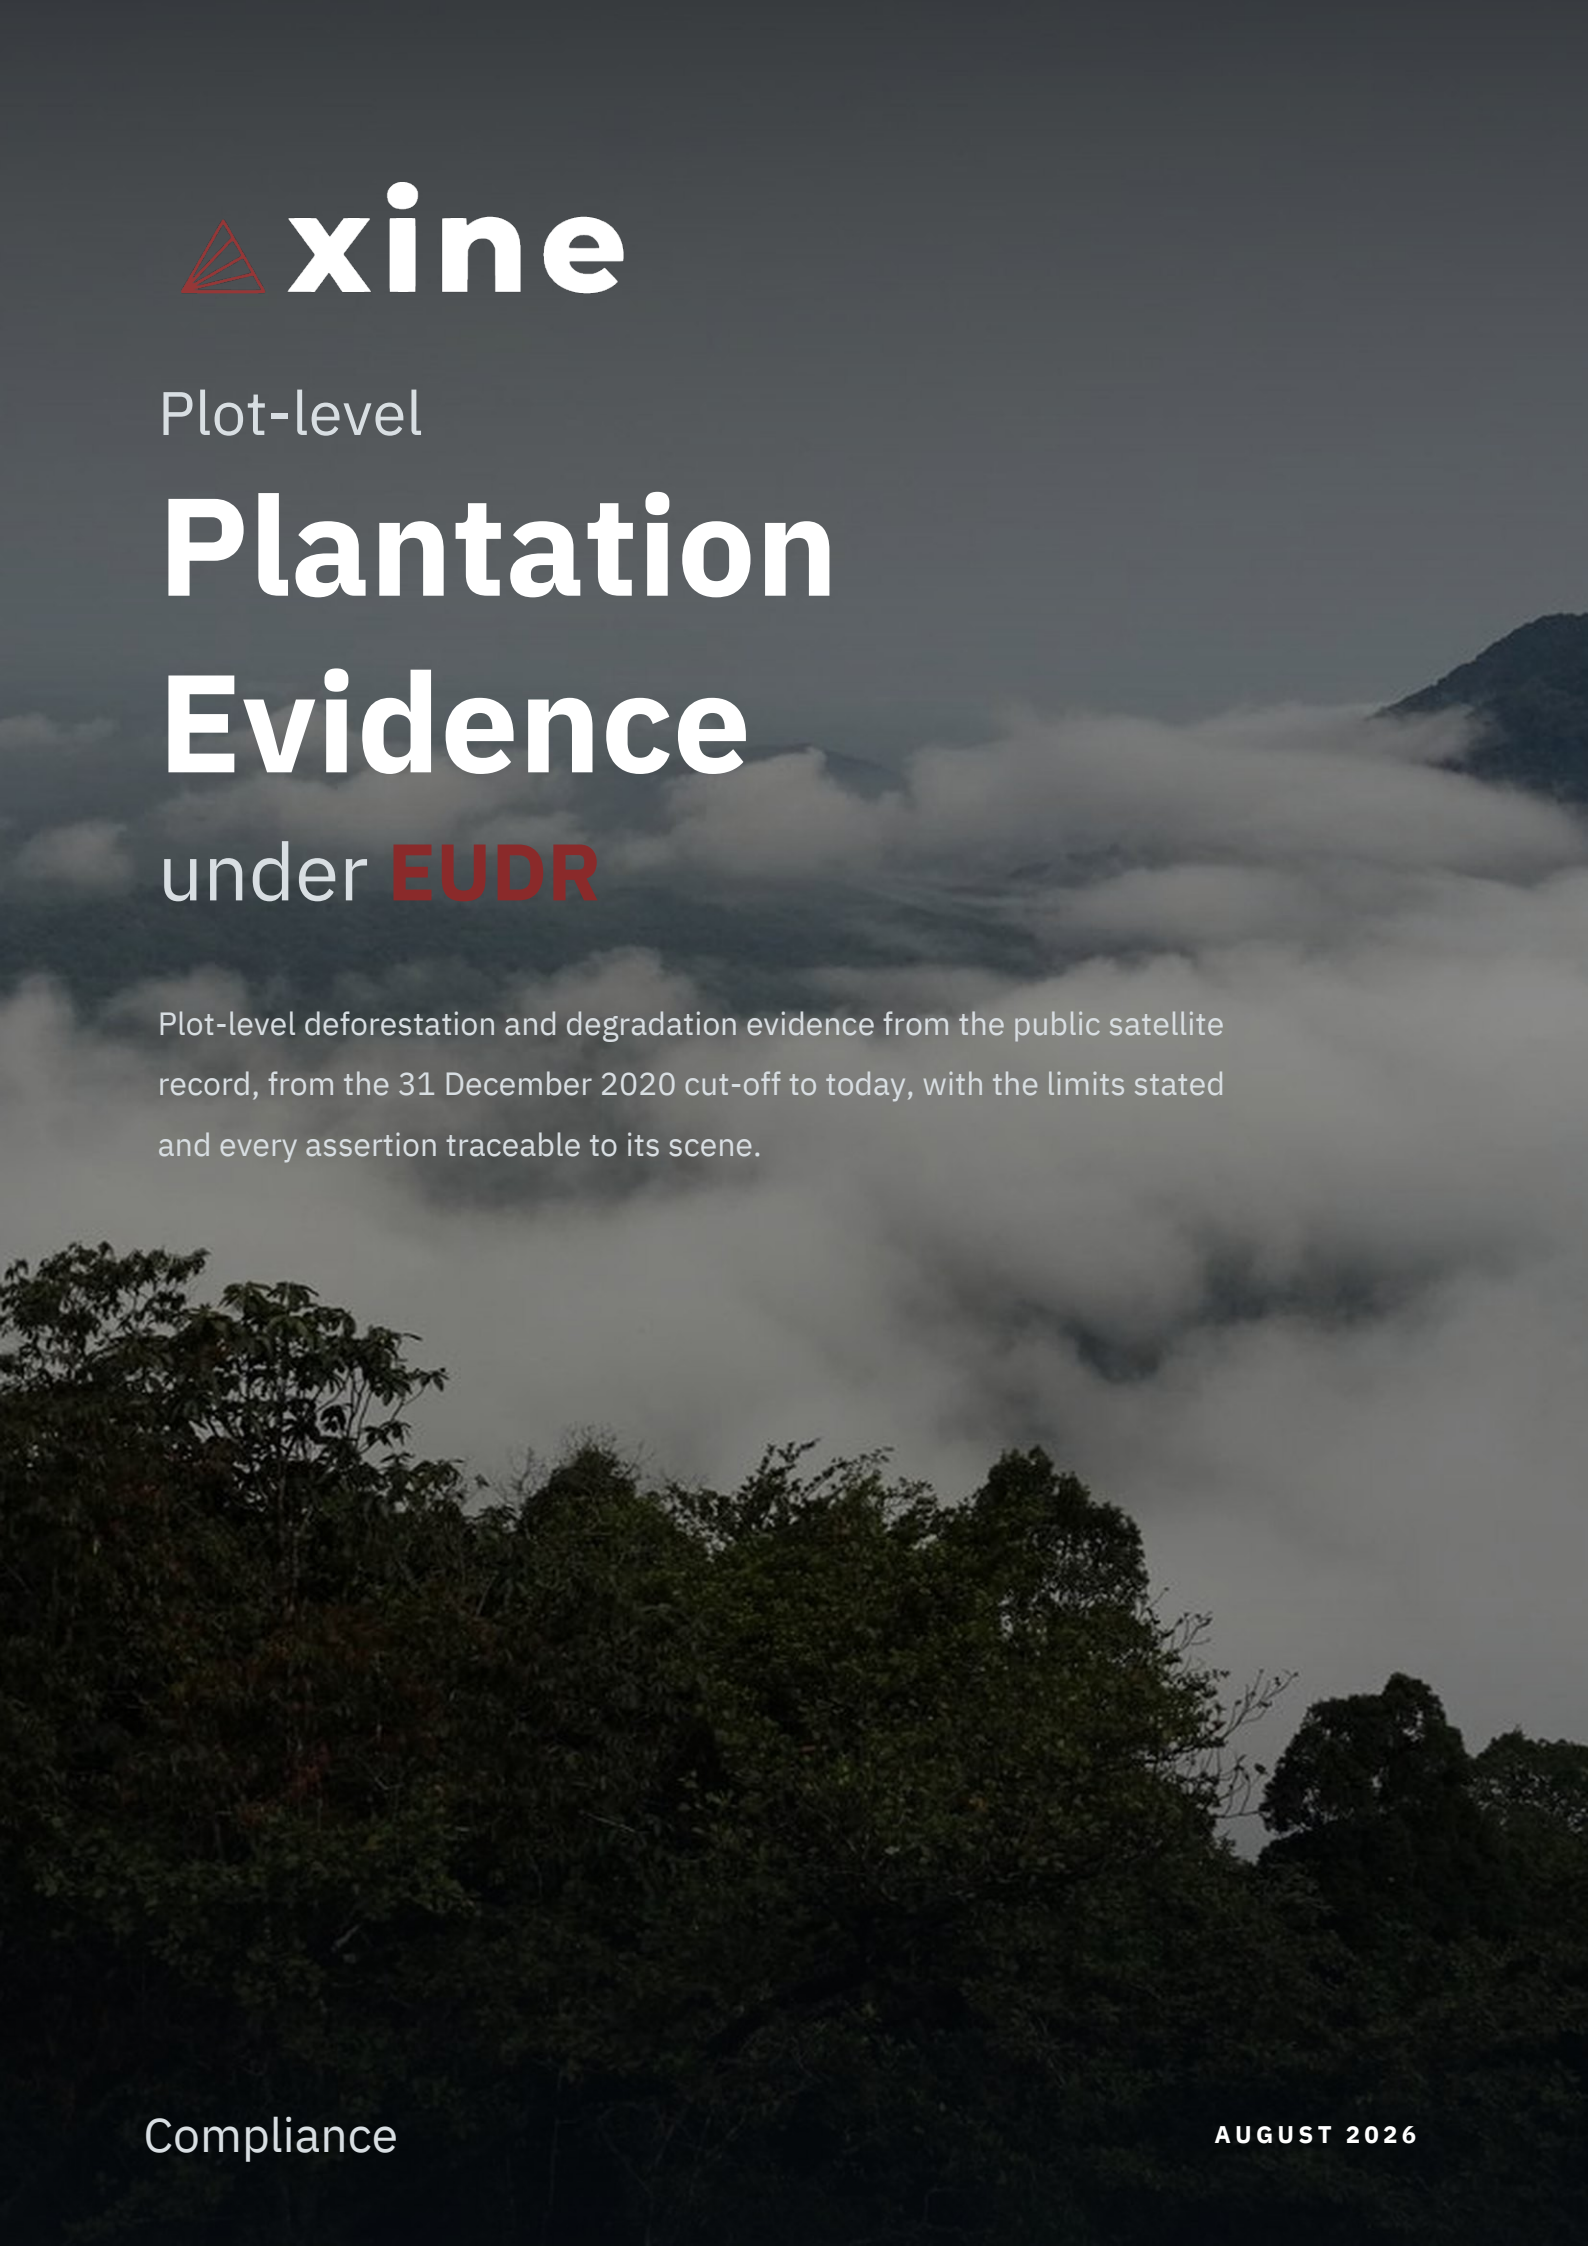

xine

Plot-level

Plantation Evidence

under EUDR

Plot-level deforestation and degradation evidence from the public satellite record, from the 31 December 2020 cut-off to today, with the limits stated and every assertion traceable to its scene.

Compliance

AUGUST 2026

Contents

Introduction

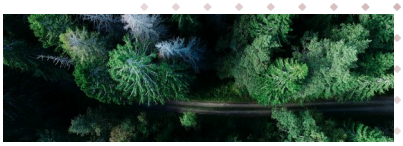

This paper sets out how satellite spectral analysis supports due diligence under the European Union Deforestation Regulation, and how Axine Laboratories applies it to plot-level verification of commodity plantations: what a plot was at the regulatory cut-off, what has changed inside it since, and what replaced the canopy where it did. It is written for compliance platforms, traders and processors filing due-diligence statements, upstream producers asked to prove their plantations, and the auditors who read the evidence.

Figure 1 — *A rubber estate on the Ivorian coast under a working dry-season sky, March 2026. Every block in it has a planting date; none of them is written on the imagery.*

Motivation: the plantation that reads as forest

Regulation (EU) 2023/1115 requires that commodities placed on the Union market be deforestation-free with reference to 31 December 2020, and that operators document how that was established for every plot of production[1]. The obligation applies to large and medium operators from 30 December 2026 and to micro and small enterprises from 30 June 2027[2]. Penalties reach four per cent of Union-wide annual turnover, with confiscation and procurement exclusion[1].

The difficulty is not the imagery; it is the reference maps. The Union's own 10 metre global forest baseline, GFC2020, classifies 31 per cent of rubber as forest, the highest overlap of any regulated commodity, against 7 per cent for oil palm and 6 per cent for cocoa; its peer review reports an 18 per cent commission error and states that co-location with mapped forest does not by itself indicate non-compliance[3]. An assessment of 21 global forest maps against the Regulation's forest definition found only two of them meeting every indicator[4]. Ground data from coffee-growing regions in Ethiopia and Tanzania puts high rates of misclassification on five global land-cover datasets, worst where the crop is shade-grown or grown in agroforestry — which is what a tree crop under canopy looks like from above[5].

A mature plantation is tree cover in every global layer, and its normal life produces exactly the signals a tree-cover-loss alert exists to catch: rubber is felled and replanted on a 25 to 35 year cycle, and it drops its leaves every dry season. Neither is deforestation under the Regulation, which excludes land predominantly under agricultural use from its definition of forest[1]. But both are flagged, and a flagged plot today is resolved by sending someone to the field with a camera.

The Regulation itself never names satellite imagery. Its standard is "adequately conclusive and verifiable information", and the Commission's guidance describes its own map as non-exclusive, non-mandatory and carrying no legal value, a support tool that does not dispense operators from their duty of diligence[6]. Satellite evidence is therefore not required, in practice unavoidable, and entirely unspecified: no layer, resolution or error rate is prescribed, while the burden of proof sits with the operator and competent authorities re-check plots. That gap is what plot-level spectral evidence fills.

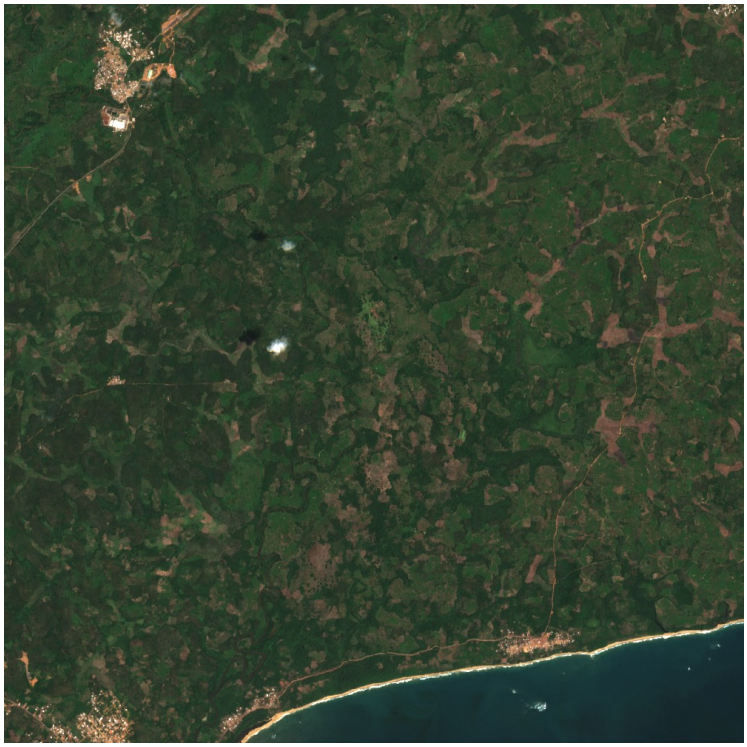

Figure 2 — *The identical frame in the window that straddles the cut-off, January 2021. Neither view says which blocks were forest and which were already plantation, and that is the compliance question.*

What is spectral change evidence?

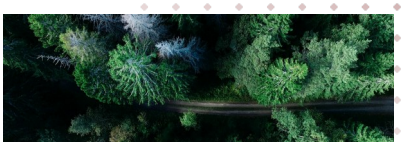

A canopy reflects sunlight in a pattern set by its pigments, its water and its structure, and the ground it stands on reflects differently again. Green leaves absorb red light and reflect strongly in the near infrared. Water in living tissue absorbs in the short-wave infrared, so an intact canopy is dark there and dry material, felled stems, slash and litter, is bright. Bare soil is bright across the visible and short-wave infrared alike. Shadow, whether from terrain, cloud or the canopy itself, darkens every band together.

The Sentinel-2 and Landsat missions have recorded these bands over every plantation on Earth every few days since 2015 and every sixteen days since the 1980s, at 10 and 30 metres, in a public archive. A plot's history at the cut-off is already in that record. The task is to read it in a way that separates what an auditor needs to know, canopy present or removed, replaced by what, from what does not matter, illumination, season and atmosphere.

Fractions, not a greenness number

A single vegetation index reports that a pixel became less green. It cannot say why, and it saturates over dense canopy so that thinning barely moves it. Spectral mixture analysis instead resolves each pixel into the fractions of it that are green vegetation, non-photosynthetic vegetation (felled stems, slash, litter), soil and shade. The Normalized Difference Fraction Index, NDFI, reads canopy condition off those fractions with shade normalised out, so a shaded canopy is still a canopy. Intact closed canopy sits near one; a stand opened by felling, with slash on the ground and soil on skid trails and landings, falls toward zero and below[7]. This is the operational degradation index of tropical forest monitoring, the basis of CLASlite and of the CODED algorithm used in national forest monitoring[8], [9].

Band group	Wavelength (nm)	What it responds to	Significance for plantation evidence
Blue, green, red	450–680	Pigment absorption; soil brightness	Bare ground, roads and landings against canopy
Near infrared	780–900	Leaf cell structure	Canopy presence and density; saturates when closed
SWIR1	1550–1700	Leaf and soil water	Canopy moisture; dry debris against living canopy
SWIR2	2100–2300	Cellulose, lignin, dry matter	Slash and felled stems; the disturbance signal a greenness index misses

Band group	Wavelength (nm)	What it responds to	Significance for plantation evidence
All bands together	450–2300	Mixture of green, dead, soil, shade	Fractions and NDFI: what replaced the canopy, with shadow removed

Time, and the two things it separates

Two comparisons answer the compliance question. The first is between a window before the cut-off and a window straddling it: a plot that was already established plantation on 31 December 2020 is not a conversion, whatever happens to it later. The second is between the cut-off and today: what has changed inside the plot since, and whether what replaced the canopy is harvest debris on a working estate or cleared ground. A third reading, the plot's seasonal signature over a run of years, is what distinguishes a deciduous crop from evergreen forest, and is the subject of ongoing calibration described in the validation section.

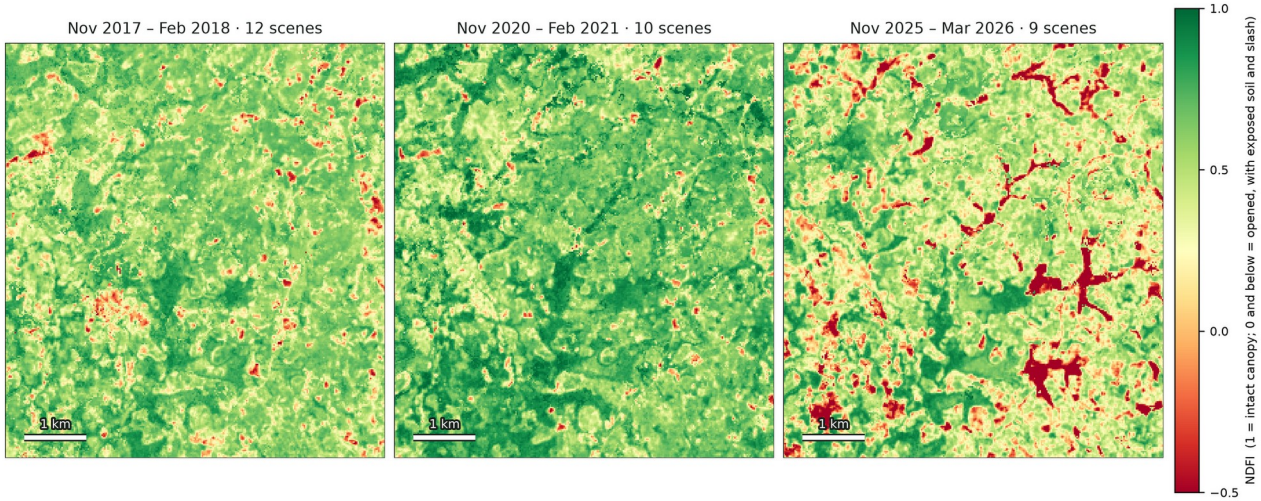

Figure 3 — *NDFI over three dry-season windows of the same ground, built from twelve, ten and nine scenes. Red is opened canopy — exposed soil and slash — and the blocks felled in the last five years read red in the current window and in neither of the earlier two.*

Challenges in commodity compliance

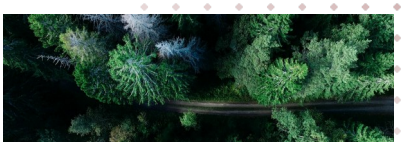

Operators, their upstream suppliers and the platforms that serve them are being asked to geolocate every plot, prove its history, file it, and defend it, at a scale that field verification cannot reach. The constraints below are the ones plot-level spectral evidence is positioned to relieve.

A cut block with the stems still stacked on it. On the ground the difference between a harvest and a clearance is obvious; from orbit it is the fractions — slash and litter against bare soil — that tell the two apart.

+ **Reference layers that flag plantations.** Global forest maps classify mature tree crops as forest and tree-cover-loss alerts fire on replanting and leaf-fall. The cost lands on the producer, who must prove a plantation is a plantation.
+ **Ground verification at scale.** A field visit with a photograph resolves one plot at the cost of a day. A supplier with thousands of smallholder plots cannot be verified that way, and a competent authority re-checking a sample will not accept that it was.
+ **Establishment date.** The question is not whether there is tree cover now but whether there was forest on 31 December 2020. That is a question about the past, answerable only from a record that already existed then.
+ **Small plots.** Smallholder plots under a hectare are a handful of pixels at 10 metres. Under the Regulation, plots over 4 hectares need polygons and smaller plots a point[1]; the evidence has to say when a plot is too small to assess rather than assess it anyway.
+ **Cloud.** The producing regions are the cloudiest on Earth. A method that summarises whatever it can see will summarise cloud; a compliant one reports the dates it could not measure.
+ **Mixed plots.** Agroforestry and crop mosaics are neither plantation nor forest at pixel scale, and a binary verdict on them is wrong in one direction or the other.

+ **Evidence that survives an audit.** A verdict without its scenes, dates, threshold and area measured is an opinion. Due-diligence statements carry a five-year record obligation[1]; the evidence behind them has to be reproducible for as long.
+ **Volume.** A statement may carry up to 10,000 plots as GeoJSON in WGS-84[1]. Verification has to be a pipeline, not a project.

Figure 4 — *Plot sizes drawn to scale on Sentinel-2 at its native 10 metres. The Regulation's 0.5 hectare forest threshold is about fifty pixels; a 0.05 hectare smallholder plot is five. The first can be assessed, the second is reported as below the mapping unit or escalated to high-resolution imagery.*

The common thread is that each of these is a question the imagery can answer only if the method is prepared to say "not evaluated". A product that always returns a verdict is, in this setting, a liability generator.

Spectral evidence in action

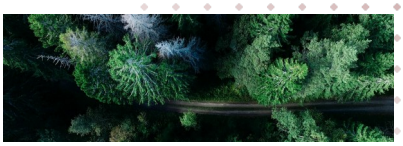

Across a plot's life under the Regulation, spectral evidence supports three capabilities.

Cloud sitting in a forested valley. Optical spectroscopy reads the illuminated surface, so ground under cloud is returned as not measured rather than inferred from the dates either side of it.

+ **Baseline.** The plot's state in a window straddling the cut-off, compared with a window some years before it. A stable plantation at both dates rules out conversion after the cut-off regardless of what a global layer says about the present.
+ **Change.** What happened inside the plot between the cut-off and today, as an area against the pair's own noise floor, with the patches listed and the fractions saying what replaced the canopy.
+ **Watch.** Continuous monitoring of the plot on every new clear acquisition, each reading compared with the same season a year earlier so a deciduous crop's cycle is not an alert and a year-on-year loss is, with every check on the record including the ones that found no usable scene.

What a plot gets back

Output	Contents	Format
Verdict	Change, no change, or not evaluated with the reason written	Portal, API, PDF
Change map	Decrease and increase patches with area and mean change; the two dates as layers	GeoJSON, GeoTIFF, tiles
Figures	Area measured, area beyond threshold, area in patches, noise floor, threshold and its source, scene offset	CSV, PDF

Output	Contents	Format
Provenance	Source scenes and acquisition dates or composite window and scene count, sensor, processing version	All exports
Fractions	Green vegetation, dead plant material, soil and shade on each date	Layers and GeoTIFF
Watch record	Every check: acquisition, cloud, reading, reference, status, and the reason for any refusal	Portal, API

Public forest layers versus plot-level evidence

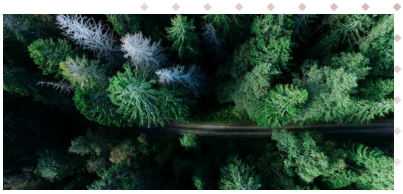

Global forest layers are the right tool for the question they were built for, national and global extent and loss over time, and the wrong tool for the question the Regulation asks of a plot. The difference is one of purpose, not of quality.

Property	Public forest layers and alerts	Plot-level spectral evidence
Unit of answer	A pixel class, or a loss alert, on a global grid	A verdict on the plot polygon, with figures for that polygon
State at the cut-off	Forest / non-forest as of the map year	The plot's own pre-cut-off and cut-off windows compared
What replaced the canopy	Not resolved	Fractions of dead plant material, soil and regrowth
Replanting and leaf-fall	Flagged as loss	Read against the plot's baseline and the same season a year earlier
Cloud and gaps	Interpolated or silent	Reported as not evaluated, with the fraction seen
Provenance	Map version	Scenes, dates, threshold, area measured, per verdict
Resolution	10–30 m, fixed	10–30 m public archive; high-resolution tasking for escalated plots
Legal standing	Non-mandatory, no legal value	Evidence the operator documents under Article 10

The two are complementary in practice: the public layer says which plots will be questioned; the plot-level evidence answers the question.

The Axine analytical method

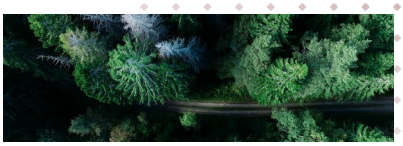

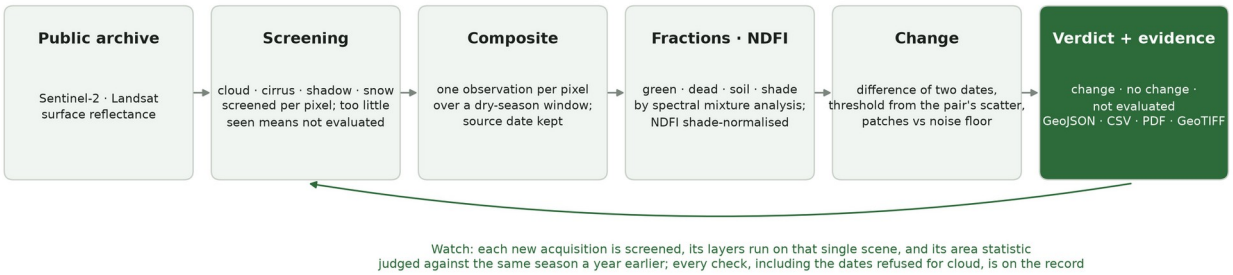

Figure 5 — *From the public archive to a verdict a client can forward. The watch is the loop back to screening: each new acquisition is screened, the same layers run on that single scene, and its area statistic is judged against the same season a year earlier. Every check goes on the record, including the dates refused for cloud.*

Reflectance and screening

Imagery is read from the public Sentinel-2 and Landsat archives as surface reflectance. Cloud, cirrus, cloud shadow and snow are screened per pixel from the mission's own quality layers, and a date on which too little of the plot is clear is reported as not evaluated with the fraction seen, rather than summarised from the part that was.

Cloud-free composites

Each date in a comparison is a composite over a dry-season window. For every pixel one real observation from one acquisition is chosen, so the spectrum is one that was actually recorded and the source date of every pixel is kept as a layer. Windows are matched season for season, so that a comparison is between two dry seasons and not between leaf-on and leaf-off.

Fractions and NDFI

Each pixel is resolved into green vegetation, non-photosynthetic vegetation, soil and shade by spectral mixture analysis against published generic endmembers, and NDFI is read off the fractions with shade normalised out[7]. The fractions are published as layers in their own right. A pixel the cloud screen missed is caught by its own fraction and excluded rather than read as bright soil.

Change against a noise floor

The difference between two dates is thresholded from the pair's own scatter rather than a fixed cut, so a noisier pair sets itself a higher bar. A uniform scene-wide shift between the dates is measured and reported, and can be removed. Pixels beyond the threshold are grouped, and only coherent patches with a substantial mean change are reported, which removes speckle. Beside the reported area sits the area that noise alone would produce at that threshold, so a change figure is never quoted without its floor.

The watch

On each new acquisition that clears the cloud limit, the same layers are computed for the plot and the reading is compared with the same season a year earlier. A reading with no such reference becomes next year's reference and claims nothing. Every check, including the dates on which no usable scene existed, is on the record with its reason.

The evidence pack

Each verdict ships as GeoJSON with the patches and the direction of subtraction on the collection, CSV with every figure and the threshold's provenance, a PDF report, and the difference raster as GeoTIFF. Composite deliverables carry their window and scene count; single-scene analyses carry the acquisition. Monitor events are hash-sealed with the rule they fired under, so an alert can be audited against the rule that produced it after the rule has changed.

Refusal

A plot the method cannot assess, for cloud, season, size or a missing layer, is returned as not evaluated with the reason. It is never returned as clear. This is the single design rule the rest of the product is built around.

Parameters, endmembers, thresholds, window lengths and the full processing chain are documented in the method statement, provided under non-disclosure agreement.

Worked example: a rubber estate in Côte d'Ivoire

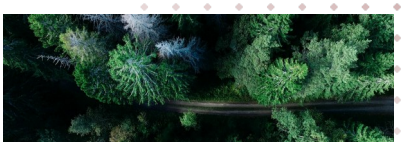

A 40.5 km² area of the industrial rubber estate landscape north-east of Grand-Béréby, Côte d'Ivoire, the world's third-largest natural-rubber producer since 2023, ahead of Vietnam[13]. Three dry-season composites were built from the Sentinel-2 archive and two comparisons run on the NDFI layer, with the fraction layers carried alongside. The estate is a working plantation on a rotational felling cycle, which makes it the right test: a method that calls its replanting deforestation fails the case the market actually has.

Window	Scenes	Mean NDFI over the area	Role
Nov 2017 – Feb 2018	12	0.54	Pre-cut-off state
Nov 2020 – Feb 2021	10	0.63	Cut-off baseline
Nov 2025 – Mar 2026	9	0.38	Current

Stable at the cut-off

2.49 km² **change in patches, Nov 2017 to Nov 2020**	**2.02 km²** **noise floor at this pair's threshold**

In the three years before the cut-off the reported change barely exceeds what noise alone would produce. The mean NDFI over the area rises from 0.54 to 0.63 across that pair and the verdict does not rest on it: a uniform shift between two dates is measured as a scene offset and removed, and stability is read from the area of thresholded change patches against the pair's own noise floor, not from the mean. The area was established plantation on 31 December 2020. Nothing that follows can be conversion of forest after the cut-off.

Replanting since

4.92 km² **change in patches, Nov 2020 to Mar 2026; 3.25 km² with the scene offset removed**	**3.29 km²** **of the area changed in dead plant material**	**0.86 km²** **changed in soil**

Roughly a tenth of the estate has been opened in five years, and the fractions say what replaced the canopy: felled stems and slash on the ground far more than bare soil. That is harvest debris, the signature of rotational felling, not of clearing for roads or settlement. On a plantation with a 25 to 35 year cycle, a tenth in five years is the cycle itself. Under the Regulation that is agricultural use, not

deforestation, and the evidence for the reading is the plot's own record on the page rather than an assertion.

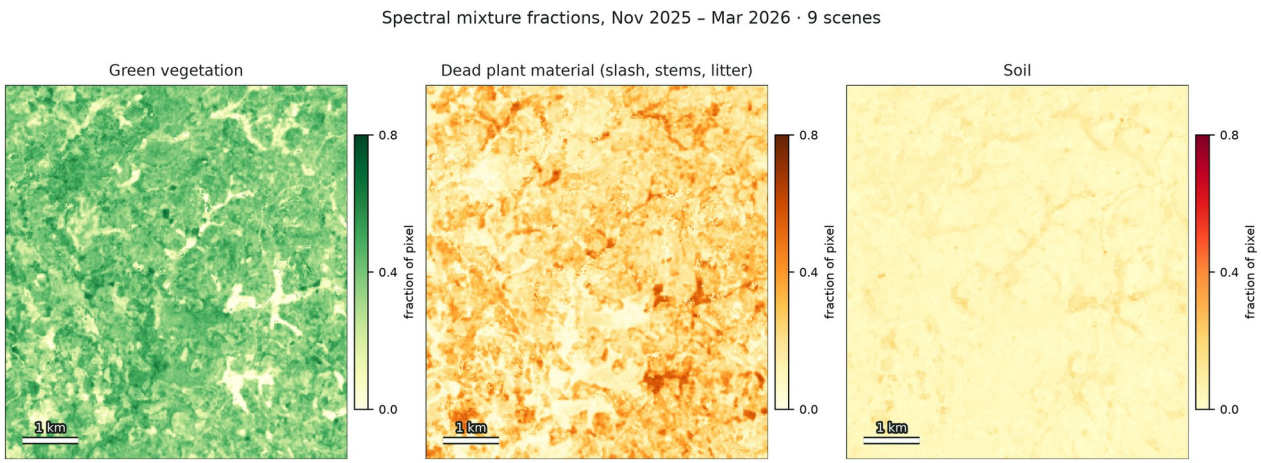

Figure 6 — *The current window resolved into fractions. Where the canopy has gone it is dead plant material that has replaced it — felled stems, slash and litter — while the soil fraction barely moves across the whole frame. That is harvest debris rather than cleared ground.*

Figure 7 — *NDFI change between the cut-off window and the current one, drawn over the current composite. 4.92 km² of the 40.5 km² area changed against a noise floor of 2.02 km² taken from the pair's own scatter: 101 decrease patches and 9 increase. Each patch carries its area and mean change, and the collection carries both windows, the direction of subtraction and where the threshold came from.*

The watch, and what a deciduous crop teaches it

A Sentinel-2 watch on the same area reads NDFI on every clear acquisition. Its first full season of checks:

Acquisition	Cloud	Reading	Reference	Status
22 Nov, 17 Dec, 22 Dec 2024; 5 Feb 2025	17–33%	—	—	Not measured: 21–57% of the area seen
24 Mar 2025	37%	0.371	—	Reference for next year
17 Nov 2025	31%	—	—	Not measured: 54% seen
17 Dec 2025	25%	0.694	—	No same-season reference
20 Feb 2026	30%	0.491	0.371	Clear
12 Mar 2026	1%	0.320	0.371	Flagged: 0.05 below last March

Read consecutively, the readings fall from 0.69 in December to 0.32 in March, and a rule that compared each check with the previous one flagged February and March as change. Nothing had been felled: rubber drops its leaves every dry season, the deciduous signature that distinguishes it from evergreen forest[10]. Read against the same season a year earlier, February is clear and March is flagged by the amount one year of rotational felling would remove, with the cloudier reference named in the event's caveat. The wet-season gap is on the record as five dates the watch could not measure, not as silence.

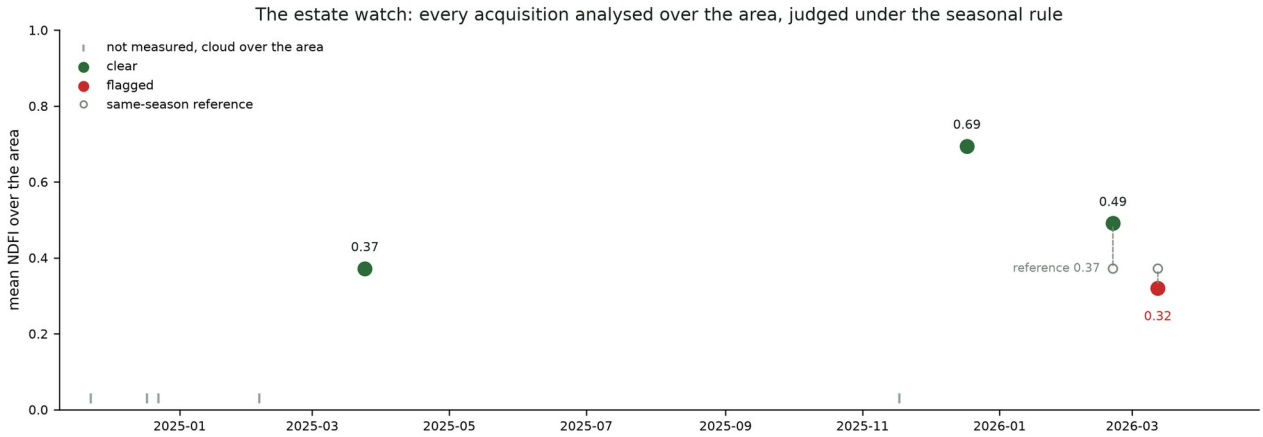

Figure 8 — *The first season of the watch. Each acquisition over the area is judged against the same season a year earlier rather than against the check before it: tick marks are dates cloud made unmeasurable, green points cleared, and the red point is the one that fell below its year-earlier reference. Read consecutively the readings drop from 0.69 to 0.32, and a previous-check rule would have called two of them change.*

Validation and confidence

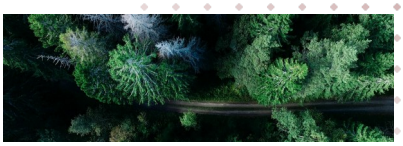

Axine treats validation as a product feature, and in compliance it is the product. An error rate an operator can put behind a due-diligence statement has to have been measured against an independent reference, under a protocol fixed before the result was seen, and published whether or not it passed.

Protocol

+ **Pre-registration.** The test sample, the reference and the pass criteria are committed to the repository under a version hash before any result is examined.
+ **Frozen parameters.** Nothing is tuned after the reference is seen; a change to the method needs a fresh sample.
+ **Independent references.** A reference produced by another classifier over the same imagery measures agreement, not accuracy, and is labelled as such.
+ **Sealed evidence.** Monitor events are hashed with the rule they fired under; every published comparison is reproducible against the exact inputs that produced it.
+ **Failures published.** The record keeps the runs that failed their gates, with the diagnosed cause.

Posture on every layer

Every layer renders with an explicit posture. Deterministic band arithmetic quotes no accuracy figure because none is claimed for it. Fitted or modelled layers are screening until a pre-registered campaign against independent ground truth says otherwise, and the fraction and NDFI layers ship as screening today: the method is published and reproduced from a public implementation[7], [9], and has not yet been scored by Axine against field data on rubber.

The precedent

On the lane where this discipline is longest established, burn extent and severity from Landsat were scored against the United States federal reference across 30 fires and 2.98 million pixels, with the sample committed before scoring. Extent agreement reached 0.94 and passed its gate; the severity class reached a chance-corrected agreement of 0.42 against a 0.45 gate, failed, and is published as failed with its cause. The same reporting standard governs every lane, including this one.

What a trial produces

Axine does not hold a rubber-specific error rate and will not quote one. A trial on a client's plots, with the planting years they hold, produces it under the protocol above: detection rate and false-positive rate against the plots' known history, refusal rate, and the reason for each refusal, plot by plot. The

same plots calibrate the seasonal-signature model that separates a deciduous crop from evergreen forest, which is the reason the trial is offered at no cost: the data it yields is the input that model cannot be built without.

Operational benefits

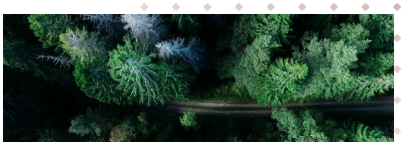

- **The plot's own history, not a global map's opinion.** Baseline state at the cut-off from the record that existed at the cut-off, for every plot, at once.
- **Replanting told from clearance.** Fractions say what replaced the canopy; the baseline and the seasonal reference say whether it is the plantation's cycle or a loss.
- **No field visit for the plots that need none.** Ground verification is concentrated on the plots the evidence cannot settle, and the evidence says which those are.
- **Evidence-grade output.** Plot-level, time-stamped, reproducible, with scenes, dates, threshold and area measured, in the formats a due-diligence statement and its auditor accept.
- **Refusal as a feature.** Cloud, season, size and mixed cover are reported, not smoothed over. A statement built on this evidence carries no silent clears.
- **Continuous rather than periodic.** The watch reads every new clear acquisition and compares like season with like, so the record between filings exists and is honest about its gaps.
- **Pipeline economics.** Polygons in through upload or API, verdicts and files out, from the public archive; high-resolution tasking only where a plot is escalated and only then priced.

Stacked stems at a landing. Felling and replanting is the normal life of a plantation on a 25 to 35 year rotation, and a tree-cover-loss alert built for natural forest fires on it every time round.

Inherent product properties and limits

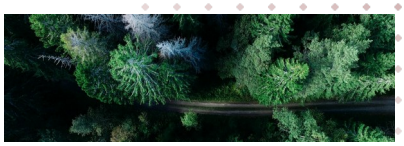

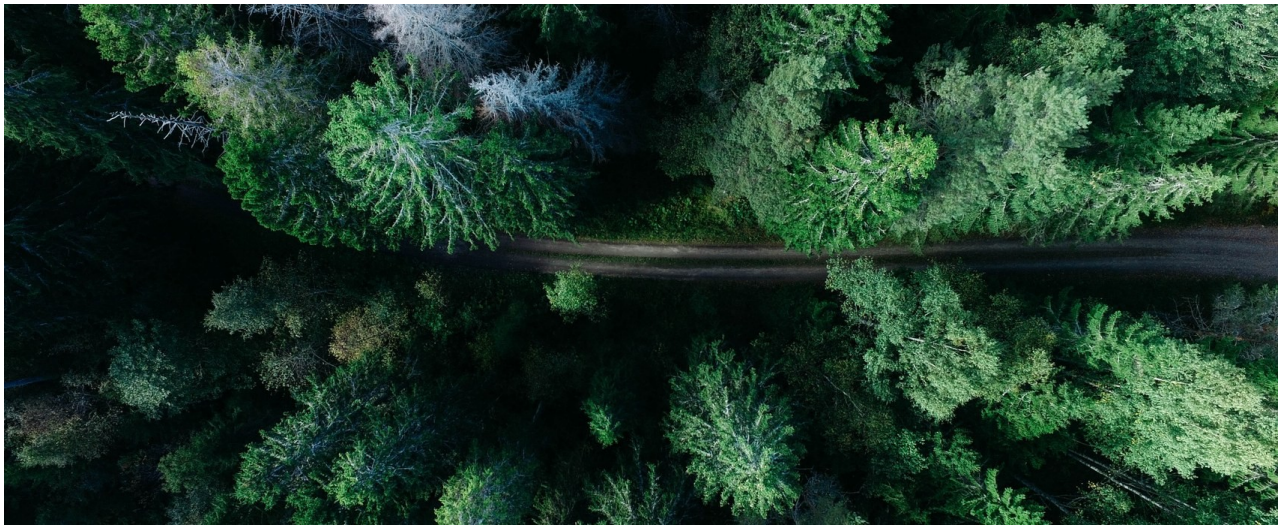

Closed canopy from directly above. At 10 metres a single pixel spans several crowns, which is where the mapping unit comes from and why a plot boundary has to be respected rather than approximated by its bounding box.

Properties

+ **Archive depth.** Sentinel-2 from 2015 and Landsat from the 1980s, so the cut-off and years before it are in the record for every plot.
+ **Physics-based and reproducible.** Published fractions and index, fixed parameters, no trained classifier in the verdict path; the same inputs give the same numbers.
+ **Season-matched.** Comparisons are dry season against dry season, and the watch compares against the same season a year earlier.
+ **Explicit data integrity.** Screened ground is not evaluated; every change figure carries its noise floor; every deliverable carries its provenance.
+ **Analysis-ready delivery.** GeoJSON, CSV, PDF and cloud-optimised GeoTIFF per verdict, portal and API access.

Limits

+ **Plantation versus natural forest from one image.** Not claimed. It comes from the plot's pre-cut-off record and from the crop's seasonal signature over a time series, which is under calibration and needs plots with known planting years.
+ **Mixed plots.** Agroforestry and mosaics are reported as mixed at 10 metres, not resolved into crop and forest.
+ **Minimum mapping unit.** About 0.1 hectare at 10 metres. Smaller plots are reported as below the unit or escalated to high-resolution imagery, quoted separately.

+ **Cloud.** Optical imagery needs a clear view. The tropics' wet season is reported as dates not measured; radar for the cloudy months is on the roadmap and not in the product today.
+ **Licensing.** Verdicts and derived layers from commercial imagery are delivered; the commercial pixels themselves are not, and provider attribution is carried on every deliverable built from them.

Optimal operating conditions

Optical spectroscopy reads the illuminated surface, so the method performs best over dry-season acquisitions with the plot largely clear of cloud, which in the producing regions means a window rather than a date. Plantations with a marked dry-season signature are the strongest case for the seasonal reference; evergreen tree crops without one are handled by the baseline and change comparisons alone. Plot polygons in WGS-84, as the Regulation requires for filing, are the right input; a point for a small plot is buffered to the mapping unit and reported accordingly.

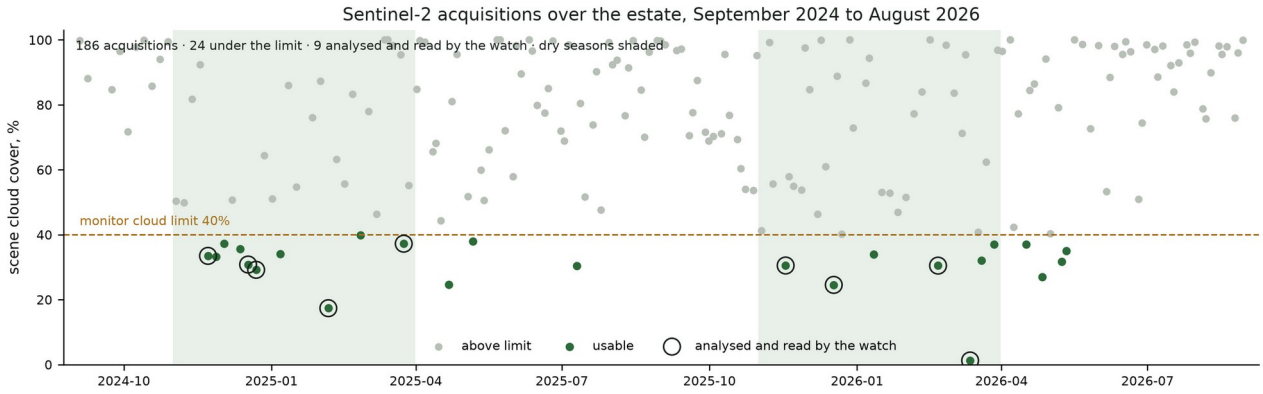

Figure 9 — *Two years of Sentinel-2 over the estate against the monitor's 40 per cent cloud limit, dry seasons shaded. Of 186 acquisitions 24 clear the limit and 9 were read by the watch. The usable dates cluster in the dry season and everything else is reported as not measured rather than summarised.*

Platform, access and data handling

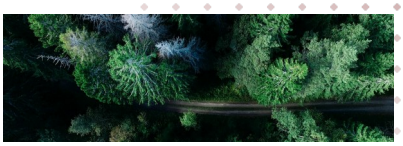

Forest and open water in one frame. Water, cloud, cirrus and their shadows are screened per pixel from the mission's own quality layers, and a date that leaves too little of a plot seen is not composited at all.

The platform

Work is ordered and delivered through a web application. Users define an area or upload plot boundaries as GeoJSON or zipped shapefile, search the archive or build a composite over a window, run the layers, compare two dates, and export. Results are served as cloud-optimised GeoTIFFs with map tiling, and any pixel can be opened to inspect the layers computed for it. Monitors run the same layers on each new acquisition and raise alerts against a seasonal reference. Reports are generated as PDF, and workspaces support teams, roles and notifications.

The client portal

A managed engagement is delivered through a client portal that shows exactly what has been published to the client and nothing else: deliverables with their maps, figures and downloads, the watch on the client's areas with every check on the record, and work still in progress. The curation boundary is enforced at the API, not the interface.

API

Plot books are pushed as polygons and returned as verdicts and files, so nothing changes in a platform's own interface. A compliance platform integrating verdicts into its plot scoring reaches the same evidence pack the portal serves.

Data handling

Plot boundaries and planting records are commercially sensitive and remain within Axine infrastructure. The client-facing assistant, Lisaris, runs entirely on Axine infrastructure rather than on a third-party model service, and any request requiring judgement is escalated to a human analyst.

Access

Offer	Coverage
Plot check	Baseline at the cut-off, change since, verdict and evidence pack per plot, from the public archive. Priced per plot, volume-banded.
Watch	Continuous monitoring of a plot book with the seasonal rule, alerts with the reading attached, every check on the record. Priced per plot per year.
Verification	High-resolution commercial imagery tasked over a plot the archive cannot settle, with the same analysis. Priced per plot on request.
Trial	Four weeks at no cost on plots from two of a client's customers, with planting years where held; per-plot outputs, a results memo covering every plot including refusals, and a method-and-limits statement under NDA.

About Axine Laboratories

Axine Laboratories builds satellite spectral analytics for forestry, agriculture, water monitoring, energy and mining. The company is registered in the United States and run by a technical team of geologists and engineers.

Our approach rests on three commitments. Detection is physics-based and reproducible, derived from measured spectra and published formulations rather than inferred by an untraceable model. Every published figure is backed by a pre-registered test against an independent reference, with the validation regime stated alongside it, and a result that fails its gate is published as failed. And a plot the method cannot assess is reported as not evaluated, never as clear, because in compliance a silent clear is the most expensive output there is.

Revenue comes from validated verdicts on the plots a client is answerable for, from continuous monitoring of them, and from high-resolution acquisition where the public archive cannot settle a plot, which aligns our interest with delivering evidence a client can defend rather than with metering access to imagery that is already public.

Axine Laboratories

axinelabs.com | contact@axinelabs.com

References

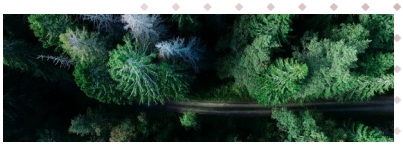

1. Regulation (EU) 2023/1115 of the European Parliament and of the Council on the making available on the Union market and the export from the Union of certain commodities and products associated with deforestation and forest degradation. Article 2 (definitions), Article 9 (geolocation: six decimal places; polygons for plots over four hectares; GeoJSON, WGS-84), Article 10 (risk assessment), Article 12 (five-year records), penalties.

2. Regulation (EU) 2024/3234 and Regulation (EU) 2025/2650 amending Regulation (EU) 2023/1115 as regards the date of application; Commission Implementing Regulation (EU) 2025/1093 on country benchmarking.

3. Bourgoin, C. et al. (2026). Mapping global forest cover of the year 2020 to support the EU regulation on deforestation-free supply chains. Earth System Science Data 18, 1331. Overall accuracy 91.5%; commission error 18%; 31% of rubber, 7% of oil palm and 6% of cocoa classified as forest.

4. Freitas Beyer, J., Köthke, M. & Lippe, M. (2025). Assessing the suitability of available global forest maps as reference tools for EUDR-compliant deforestation monitoring. Remote Sensing 17(17), 3012. Twenty-one maps assessed against the Regulation's forest-definition indicators; two met all of them.

5. Gallemore, C. et al. (2025). Avoiding access inequity due to classification errors in zero-deforestation value chains: coffee and the European Union Deforestation Regulation. Land Use Policy 157, 107609. Five global land-cover datasets tested against ground data from Ethiopia and Tanzania.

6. European Commission (2024). Guidance document for Regulation (EU) 2023/1115, C/2024/6789, section 11.5: the EU Observatory's forest cover map is non-exclusive, non-mandatory and carries no legal value.

7. Souza Jr., C. M., Roberts, D. A. & Cochrane, M. A. (2005). Combining spectral and spatial information to map canopy damage from selective logging and forest fires. Remote Sensing of Environment 98, 329–343.

8. Asner, G. P. et al. (2009). Automated mapping of tropical deforestation and forest degradation: CLASlite. Journal of Applied Remote Sensing 3, 033543.

9. 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 238, 110968.

10. Wang, Y. et al. (2023). High-resolution maps show that rubber causes substantial deforestation in Southeast Asia. Nature 623, 340–346. Sentinel-2 at 10 m; 95% ± 2% overall accuracy; 4.1 million hectares of forest cleared for rubber 1993–2016.

11. Hansen, M. C. et al. (2013). High-resolution global maps of 21st-century forest cover change. Science 342, 850–853.

12. Reiche, J. et al. (2021). Forest disturbance alerts for the Congo Basin using Sentinel-1. Environmental Research Letters 16, 024005.

13. Association of Natural Rubber Producing Countries, production data for 2023, reported January 2024. Côte d'Ivoire third worldwide at just under 1.5 million tonnes, ahead of Vietnam at 1.39 million and behind Thailand and Indonesia.

xine

PLANTATION EVIDENCE UNDER EUDR

axinelabs.com | contact@axinelabs.com