



Dated satellite records

Claims and Portfolios

for insurers and **lenders**


Radar flood evidence, construction progress, asset status and climate-risk tables from open satellite archives, each read dated, each method named, and the gaps in our own validation stated.



Contents

■ Introduction	01
■ Challenges in Claims and Portfolio Records	03
■ The Axine Method	05
■ The Products	06
■ Worked Examples	09
■ A construction ramp	10
■ A three-asset portfolio	11
■ A climate-risk scan	12
■ Validation and Limits	14
■ How the Record Is Used	16
■ About Axine Laboratories	17
■ References	18

Introduction



This paper describes how Axine Laboratories turns open radar and optical satellite archives into dated records for insurance claims desks and for the credit and asset teams of lenders and landholders: what moved on a parcel or an asset, by how much, and on which scenes it was read.


Before: radar reference

21 March to 19 May 2023




After: 1 to 10 days after the loss date

21 to 30 May 2023



Newly flooded

Sentinel-1 VH drop above 1.13 dB



Grey scale: VH backscatter, fixed -25 to 0 dB

■ Newly flooded area

Figure 1. Sentinel-1 VH radar over one agricultural parcel at Conselice, Emilia-Romagna, used as a test case with a loss date of 20 May 2023. Left, the radar reference from 21 March to 19 May 2023; centre, the look from 21 to 30 May 2023; right, the newly flooded area in blue (VH drop above 1.13 dB), 23.17 ha or 20.4 per cent of the parcel.

Motivation: a loss date, a polygon and a week to decide

Natural catastrophe losses reach insurers as claims and reach lenders as changes in collateral. Swiss Re Institute puts insured natural catastrophe losses in 2025 at USD 107 billion, within global economic losses of USD 220 billion of which about 49 per cent were insured^{1,2}. Secondary perils, among them wildfires, severe convective storms and floods, accounted for a record 92 per cent of insured losses². In the European Union, EIOPA and the ECB count around EUR 900 billion of direct economic losses from natural catastrophe-related extreme events between 1981 and 2023, and report that only about a quarter of the losses from extreme weather and climate-related events were insured³.


Supervisors have started to ask lenders for the same attention. The ECB guide on climate-related and environmental risks names lower collateral valuations in real estate portfolios as a result of increased flood risk as a physical-risk example, and expects institutions to monitor these risks in their portfolios⁴. In Brazil, CMN Resolution 5.267/2025 has been in force since 1 March 2026. It makes remote sensing mandatory for custeio and investimento credit operations contracted from that date with a total area above 300 hectares, applies it from before contracting until settlement, and requires auditable records of the satellite, sensor, image date, resolution, bands and method⁵.

Two readers use the same archive and ask different questions. We wrote this paper for both of them.

- + **The claims desk.** It receives a loss date, a polygon and a claimed amount, and decides within days whether to send an adjuster. The question is narrow and dated: what moved on this parcel, by how much, and when was it read?
- + **The portfolio manager.** A credit officer, a farmland fund or an analytics lead asks a standing question: of everything held, what changed since the last look, and which three of six hundred assets need a person this month?
- + **The auditor or the other side of a dispute.** Either reader may later have to show the method, the scenes and the threshold behind a number. The record is built so that this can be done from the document alone.

Six products answer these questions: flood extent, damage mapping, claims validation, asset status, construction progress and climate risk. A seventh, portfolio monitoring, collects the results of the others into one feed. The paper describes each method with its published source, then works three of them end to end on reads made on 6 October 2026, and lists what we have not validated.

Challenges in Claims and Portfolio Records



Perils arrive under cloud, losses are largely uninsured, and a number from a satellite index is only as useful as the method that can be inspected behind it.

USD 107 bn

insured natural catastrophe losses in 2025, Swiss Re Institute

92%

share of those losses from secondary perils, a record

EUR 8.5 bn

direct damage, Emilia-Romagna floods, May 2023

Perils arrive under cloud


Optical satellites need daylight and a clear sky, and floods and storms come with neither. MODIS cloud-mask data put the global cloud fraction at about 67 per cent, and about 55 per cent over land⁶. A study of European flood events found that on average 58 per cent were potentially observable by Sentinel-1 and only 28 per cent by Sentinel-2, because of cloud cover⁷. Radar passes through cloud and records at night, so a flood read has to start from radar, with optical scenes as a second line when a clear one exists.

Our Conselice test shows the effect. A radar read taken one to ten days after the loss date finds 20.4 per cent of the parcel newly flooded. The first clear optical scene came about a month later, after a persistently cloudy May, and read 0.08 per cent. We report both numbers and the reason they differ in the section on the worked example.

Large losses, small insured shares

The Emilia-Romagna floods of May 2023 are a European example of the gap. The Commission's proposal for the EU Solidarity Fund put total direct damage at EUR 8.5 billion⁸. An early estimate from the EU Space Programme, published on 22 May 2023, had been EUR 5 billion, with 14 deaths and more than 30,000 residents displaced⁹. Guy Carpenter projected industry property losses of EUR 500 million to EUR 1 billion, with motor adding 10 to 15 per cent¹⁰. The EEA reports that most of the countries it covers had over 50 per cent of their weather and climate-related losses uninsured, and in many cases over 90 per cent¹¹.

Copernicus Emergency Management Service activated EMSR659 for the event on 3 May 2023, with a second activation, EMSR664, on 16 May¹². The first activation text names the breaking of the Sillaro River embankment near Massa Lombarda and Conselice. These products map a flood for the emergency response. An insurer or lender needs the answer at the level of one parcel or one asset, dated to a loss.



A tidal river with moored boats, a road bridge and terraced houses, seen from directly above. No flooding is visible in this image.

A number nobody can check


Satellite indices already price some payouts. On 25 September 2026 the presidents of four French livestock unions (FNB, FNPL, FNO and FNEC) wrote to the agriculture minister and the insurers about the Airbus grassland production index, asking for more visibility on loss rates. One sentence of the letter reads: "Chaque éleveur français doit pouvoir connaître précisément l'estimation de sa perte telle qu'elle est calculée par votre indice satellitaire."^{13,14} Regional unions in the Massif central asked for the loss rates to be published commune by commune¹³. At the Space show on 16 September 2026 the minister had called it excessive to say the index does not work, and said he wanted to make it more reliable¹⁵.

We take two points from this dispute. A payout that rests on a dated satellite record draws questions about the method once the amounts matter, and a desk that publishes the method, the scenes and the threshold with every number can answer them. The sections that follow describe what we publish.

Open archives and their schedule

Sentinel-1 now flies as Sentinel-1C and 1D. The Copernicus Data Space notice of 24 June 2026 records that the nominal six-day repeat was restored, and Sentinel-1A ended operations at the end of June 2026¹⁶. A single Sentinel-1 satellite repeats every 12 days and a pair every 6¹⁷. The Conselice reference stack in this paper, from 2023, shows the 12-day pattern: ten acquisitions between 23 March and 11 May 2023 in five pairs of adjacent days. Sentinel-2 offers a revisit of about five days with two satellites and is free and open, as are the radar data¹⁸.

The Axine Method



Each product reads a published method, calls it directly on the parcel or asset, and writes the measurement together with its own denominator and the scenes it came from.





Figure 2. From open archives to a dated record that a claims desk or portfolio manager reads and acts on.

Data and sources

Every read in this paper comes from Copernicus or another open archive, fetched on 6 October 2026 through the public STAC catalogue of Microsoft Planetary Computer or through a named public service. No commercial or restricted-licence imagery appears in any figure. Nothing in the paper is simulated. Each figure and table value was produced that day by calling the production code for the product itself against the live archive, and none is a replay of an earlier run.

Source	What it gives	Used by
Sentinel-1 RTC (ESA, via Planetary Computer)	VV and VH backscatter, amplitude only; no interferometric coherence on this stack	Flood extent, claims validation, asset status, construction progress, damage mapping
Sentinel-2 L2A (ESA, via Planetary Computer)	Optical reflectance and the indices NDBI, NDVI, MNDWI and NBR	All six products
JRC Global Surface Water ¹⁹	Permanent-water seasonality, used to exclude lakes and rivers from a flood mask	Flood extent, claims validation
Copernicus DEM GLO-30	Elevation and slope, for the terrain flood-vulnerability class	Flood extent
WRI Aqueduct Floods ²⁰	A five-class basin-level flood-hazard score, served through Resource Watch and CARTO	Climate risk
JRC GWIS burnt-area archive (MODIS and VIIRS) ²¹	Matched historical fire records, back to at least 2012	Climate risk
ERA5 and ERA5-Land (via Open-Meteo)	Daily rain and temperature, for SPI and KBDI	Climate risk

The Products



Flood extent

We follow the UN-SPIDER recommended practice for radar flood mapping²², generalised in the way Twele, Cao, Plank and Martinis describe²³. The read takes the drop in VH backscatter between a pre-event Sentinel-1 reference composite and the nearest post-event look. It thresholds the drop with Otsu's method²⁴ on the histogram of the area of interest, and floors the threshold at the UN-SPIDER fixed cut, a ratio of 1.25 in linear power, which is 0.97 dB. No area can therefore pick a threshold more permissive than the published practice. Permanent water is excluded with JRC Global Surface Water, and steep ground with a Copernicus DEM slope mask. Flood vulnerability is reported as a simplified height-above-nearest-drainage class^{25,26}, a terrain exposure class that carries no return period.

The sign, the threshold direction, the water mask and the terrain logic are tested on constructed arrays with a known right answer, and 59 of 59 tests pass. No benchmark at area-of-interest scale against a named Copernicus EMS activation or the Global Flood Database has been run. We probed both portals directly. The EMS rapid-mapping pages are rendered in the browser and we found no unauthenticated bulk vector endpoint in the time available. The per-event rasters of the Global Flood Database are an Earth Engine asset, which our production stack does not use. We state this gap again in the section on limits.

Claims validation

Four published methods cover four loss types, each called directly on the claim polygon as a single stand or parcel.

- + **Storm.** A windthrow method for forest stands^{27,28}. The same rule was checked against Copernicus EMS activation EMSN231 for forest stands.
- + **Fire.** A burn-ratio rule with refitted severity breaks, which reached 0.9412 fire-weighted binary agreement against MTBS, as the primary read²⁹. A second line keeps the older median-dNBR read against the USGS FIREMON severity bins from the same composite.
- + **Flood.** The flood extent rule above as the primary read, run on the claim's own grid over a 60-day pre-event radar reference and a look 1 to 10 days after the loss date. MNDWI³⁰ stays as a second line and no longer makes the call.
- + **Other.** Before and after medians of NDVI, MNDWI and NBR, with no severity class, because no literature bin applies to a loss type that nobody named.

The flood primary read is reported as hectares flooded over the parcel's own total hectares: 23.17 of about 113.6 ha at Conselice. No accuracy figure is printed anywhere in the report. Hectares, percent of parcel and decibels are measurements with a stated method. The fire and flood primary reads inherit whatever their sibling products have measured, and neither sibling has an independent control of its own at the scale of a claim. We say so on the page because a stale statement that a method is validated does more harm than none.

Asset status

The read uses the change in NDBI³¹ and NDVI between a before window and an after window, with Sentinel-1 VV and VH as a corroborator that decides nothing. The change class compares the read with the asset's own trailing-year history. An asset reads **above** when this read's changed-pixel fraction exceeds the larger of its own highest quiet-year fraction and a fixed floor, and **within** when it does not. The same rule carries our Copernicus EMS validation for forest stands and is applied here to buildings and sites by analogy. We have not built a separate control for it.

The window, quiet-range-shift and report-contract rules are tested on constructed cases, and 15 of 15 pass. After the live run we repeated the read three times on the same three-asset portfolio. Each repeat gave the same result with no duplicate asset records, and the history of each asset carried forward from one run to the next.

Construction progress

We form one calendar-month Sentinel-2 composite for each month since the site's start date and compute NDBI³¹ against the published built and bare line, an NDBI of 0.0. The line is not calibrated locally against a ground survey, and the report says so. When the customer uploads a planned footprint, the same month's built mask is read a second time inside the footprint's own pixels. Sentinel-1 VV and VH are read once for the first month and once for the latest, and printed as a corroborating series.

Built percent is the share of the site's valid 10 m Sentinel-2 pixels whose NDBI crossed the published line in that month. It is a direct measurement against a published threshold and makes no accuracy claim. The 17 self-tests on synthetic month windows, masks and percentages pass, and the live run read 17 of 18 calendar months with a clear Sentinel-2 look.

Climate risk

One job returns seven hazards, each a published, named index. SPI-12^{32,33} is computed on ERA5 and ERA5-Land rainfall. SPI needs a long record: WMO guidance allows as little as 20 years of data and recommends at least 30 where possible. KBDI³⁴ covers fire weather, and ETCCDI-style day counts³⁵ cover heat and frost. Vegetation trend comes from our annual peak-NDVI reader over Landsat and

Sentinel-2, capped at 8 years to bound compute. Yield volatility comes from USDA NASS county yield where a key is configured, and from the interannual variation of the vegetation trend otherwise. Two lookups read public services: WRI Aqueduct Floods and the JRC GWIS burnt-area archive.

Each hazard's arithmetic is checked against its own defining identity: SPI of noise around a constant mean returns a mean of 0 and a standard deviation of 1, KBDI rises under heat with no rain and falls after a soaking rain, and a portfolio feature with no geometry is skipped without a crash. All 15 of 15 checks pass. No accuracy figure is printed for this product. The values are published indices and two public services read through, and no detector has been scored against a reference map.

Damage mapping

The product gives one adjuster-shaped read of a single area and a single event of any type. Sentinel-2 and Sentinel-1 are compared before and after. When the stated event is a flood or a fire, the read reuses the flood extent or burnt-area rule, and for a storm it uses a fall in NBR as a windthrow proxy²⁸. Built-up damage candidates (NDBI brightness and radar amplitude, limited to the WorldCover built class) and contamination candidates (a turbidity index after Lacaux et al.³⁶ and an oil-spill index) are always attempted, whatever the event type. Each is reported as a candidate and never as a finding, because a repaved roof or ordinary turbid sediment can trigger either screen with nothing abnormal underneath.


All 21 self-tests pass, including four end-to-end runs (flood, fire, storm, other) on a constructed grid with a known damaged subset. The built-up and contamination screens have no area-scale benchmark of their own and inherit none from a sibling product.

Portfolio monitoring

This product adds no new detector. It wraps the four watched products (forest change, burnt area, flood extent and asset status) and reads what each has already found for an asset's own footprint, so no product is re-run at every check. An asset reads **above** when at least one watched product's own figures cross its change rule.

The aggregation logic passes 10 of 10 tests. A live run against a scratch database with demonstration assets and real completed jobs gave, on its second run, 5 assets, 4 changed, 0 with no change and 1 not evaluable. Every watched product tripped **above** on its own asset. The fifth asset had no coverage and read "no baseline", where "no change" would have been wrong. The run needed no code change on either attempt.

Worked Examples



A parcel at Conselice in the Emilia-Romagna floods of May 2023, read with the claims validation flood rule on 6 October 2026.

A flood claim at Conselice

The parcel is an agricultural plot at Conselice, the town named in the activation text for EMSR659¹². We use it as a test case with a loss date of 20 May 2023. Figure 1 comes from calling the claims validation flood reading directly on this parcel and loss date, which is the same function the earlier live run used, and reading Sentinel-1 afresh from Planetary Computer. The parcel measures about 113.6 ha.


The reference composite holds 10 Sentinel-1 acquisitions from 23 March to 11 May 2023. The post-event look holds 2 acquisitions, on 22 and 23 May 2023, which fall 1 to 10 days after the loss date. Otsu's method picked a threshold of 1.13 dB from the histogram of this parcel's own VH drop. That sits above the 0.97 dB floor, so the contrast of the parcel itself decides the cut and the floor does not.

23.17 of about 113.6 ha, or 20.4 per cent, read as newly flooded.

The first clear Sentinel-2 look after the loss date came about a month later. A persistent cloud cover held over the region through most of May, which is the time a flood happens. That scene read only **0.08 per cent newly flooded** by MNDWI, because the crop canopy had closed over the water by the time a clear optical scene existed. Both numbers are genuine reads of the same ground. They differ because the optical method is structurally late for a flood.

Copernicus EMS activated EMSR659 for this event, and the radar reading is consistent with the EMS public account of the area as flooded. We did not retrieve a parcel-level grade from the EMSR659 vector product in this session, because the EMS portal offers no unauthenticated bulk vector download that we could find in the time available. The 20.4 per cent and 0.08 per cent figures were read directly from Sentinel-1 and Sentinel-2 and do not depend on EMS in either direction. Neither number has been checked against a ground survey of the flooded area.

A construction ramp



Monthly built percent at a factory construction site near Grünheide, Brandenburg, with a drawn sub-area standing in for a planned footprint.

We read the trailing 18 calendar months from February 2021 to July 2022 for a site with a start date of 1 June 2020. The read calls the construction progress month windows, built mask and percent functions directly. A sub-area of the site, drawn by us, is read a second time as a stand-in planned footprint. It is not an official site plan.





Figure 3. Monthly built percent at the Grünheide site, February 2021 to June 2022 (17 of 18 months read), for the whole site and for a sub-area we drew as a stand-in planned footprint, with true-colour Sentinel-2 chips for February 2021 and June 2022. July 2022 was not read and is not plotted.

17 of 18 months had a clear Sentinel-2 look. The July 2022 slice was only one day wide because it is the last slice of the trailing window. The series moves with cloud and with the number of looks in each composite, from 1 to 8. Whole-site built percent started at 39.6 per cent in February 2021 and ran between 42.3 and 50.4 per cent from March to July 2021. It fell to 21.3 per cent in November 2021 on one look, 27.9 per cent in December 2021 on one look and **19.8 per cent in January 2022** on four. It recovered through the spring, with 47.4 per cent in March 2022 on eight looks, and reached **44.2 per cent in June 2022**.

The drawn sub-area reads higher throughout: 49.5 per cent in February 2021, 28.3 per cent in January 2022 and 48.0 per cent in June 2022. The sub-area excludes the unbuilt perimeter and the access roads that the whole-site number still carries. A lender or insurer who follows a site month by month needs to know how many looks sit behind each point. We hold no site programme or survey for this location, so the series is a measurement against the published NDBI line.

A three-asset portfolio



Three assets read against their own history with the asset status product.

We read a warehouse near Bologna, a solar farm near Montalto di Castro and a vineyard near Montalcino. These are real coordinates, the same three assets as in our database run of 5 October 2026. The table is an independent fourth read made by calling the asset status function for each asset directly over a 30-day before and after window that ended in the last days of September 2026. No database is involved.

Asset	Coordinates	Status	Event fraction	Sentinel-1 corroborator	Last look
Warehouse near Bologna	44.5115 N, 11.3370 E	Within	0.0176	-0.00 dB	29 September 2026
Solar farm near Montalto di Castro	42.3350 N, 11.8160 E	Within	0.0000	1.56 dB	26 September 2026
Vineyard near Montalcino	43.0530 N, 11.4935 E	Within	0.0030	-0.20 dB	26 September 2026

All three assets read within their own quiet range. The event fraction is the share of an asset's valid Sentinel-2 pixels that moved in NDBI and NDVI between the before and after windows. The values match, to four decimal places, the 0.0176, 0.0 and 0.003 recorded for the same three assets in the earlier run on a different day. That is what a stable, repeatable read against an asset's own history should look like.

The Sentinel-1 corroborator read -0.00, 1.56 and -0.20 dB on the three assets and decided nothing on any of them, by design. One earlier rendering of this read met a transient 403 response from the archive on one quiet-year radar acquisition for the warehouse. The primary read completed anyway, and this rerun met no error. A degraded corroborator that does not stop the read is the intended behaviour: an absent guard weakens the mask and does not stop the product.

A climate-risk scan



The two networked lookups of the climate risk product, read live at six points.

Both service lookups were corrected against live responses on 5 October 2026. The flood table name we first used did not exist on the CARTO service, and the burnt-area host we first used had been decommissioned. The live replacement answers only through a pixel query on the layer, and its response is plain text, which we parse as such. We then exercised the two lookups against the six points used in that correction round: three flood-hazard points and three fire-history points.

Flood point	Coordinates	Basin	Hazard class	Label returned by the service
Rhine floodplain, near Cologne	50.933 N, 6.980 E	Rhine	Low	Low (0 to 1 in 1,000)
Po valley, near Piacenza	45.050 N, 9.700 E	Po	Low	Low (0 to 1 in 1,000)
Sacramento valley	38.580 N, 121.490 W	California	Moderate	Low - Medium (1 in 1,000 to 2 in 1,000)

The flood classes are identical to those of the 5 October verification, read fresh. We print the label exactly as the service returns it.

The burnt-area lookup is a point query, and which of several overlapping fire records a given call returns is not perfectly stable between calls. The table gives the largest record among the matches in each point's surrounding area, next to the named record matched on 5 October.


Fire point	Matched events	Coordinates	Largest record, 6 October 2026	Named record, 5 October 2026
Evros, Dadia forest, Greece	9 in the last 5 years	41.10 N, 26.30 E	242 ha, 9 March 2025	Fire 12288764, 82,210 ha, 19 August to 2 September 2023
Tenerife	24	28.23 N, 16.60 W	16,337 ha, 16 to 23 August 2023	The same record, 16,337 ha
Fox Lake area, Alberta	4	58.00 N, 114.50 W	408 ha, 30 May 2023	Fire 12107418, 10,547 ha, 6 to 18 July 2023

The two named 2023 fires at Evros and Fox Lake were matched directly on 5 October and checked against the public record for those events, and both remain true. The point query of 6 October did not return those two records among its matches at the same two locations, and it did return the Tenerife record. This behaviour comes from a single-pixel query against a dense archive of overlapping polygons. The earlier and later figures are two independently real readings of the same points on different days. A reader should neither average them nor try to reconcile them.



An aircraft releasing a white spray over forested hills in haze.

Validation and Limits



Axine publishes a gate that has not been run as not run. This section lists what each product has been checked against and what remains unevaluated.


Product	Self-tests on constructed cases	Independent control at area scale
Flood extent	59 of 59	None run against an EMS activation or the Global Flood Database
Claims validation	Flood and fire reads inherit their sibling products	None of its own for flood or fire; the storm rule was checked against EMSN231 for forest stands
Asset status	15 of 15	Carried by analogy from the forest-stand control, with none of its own
Construction progress	17 of 17	None; the built and bare line is a published convention
Climate risk	15 of 15	None; no loss-ratio or claims calibration has been run
Damage mapping	21 of 21	None for the built-up and contamination screens
Portfolio monitoring	10 of 10	A live run on five demonstration assets

- + **No area-scale flood benchmark.** Flood extent and the flood primary read of claims validation have not been scored for intersection over union, omission or commission against a named Copernicus EMS activation or the Global Flood Database. Neither portal offered an unauthenticated bulk vector download in the time available. No intersection-over-union number appears in this paper, because we would be printing a figure with no denominator behind it.
- + **The Conselice divergence is a measurement.** The 20.4 per cent from radar and 0.08 per cent from optical are two live reads. Neither has been checked against an independent ground survey of the flooded area, and neither is an accuracy figure for its method.
- + **Own-history change rule by analogy.** The rule used by asset status and claims validation carries the Copernicus EMS validation of our forest-stand product. That control was built and run over forest stands, and not over buildings, solar farms or vineyards. The within reads in the portfolio table are real, repeatable measurements against each asset's own history. They are not an accuracy claim against an independent reference.
- + **Burnt-area lookup stability.** The single-pixel query against overlapping polygons does not return the same record on every call, as the fire table shows at Evros and Fox Lake.

- + **Built and bare line not calibrated locally.** The NDBI line of 0.0 follows Zha, Gao and Ni³¹ and has not been calibrated against the roofing, soil or climate of the Grünheide site. The 44.2 per cent figure is a direct measurement against that line.
- + **The planned footprint is a drawn stand-in.** The sub-area at Grünheide is our own illustrative boundary. It shows how a percent measured against a planned footprint reads higher than the whole-site percent once a shell exists and an unbuilt perimeter still lowers the wider number.
- + **No accuracy figure for climate risk.** SPI, KBDI and the ETCCDI-style counts are published indices checked here against their own mathematical identities. They have not been scored against a loss-ratio or claims record, and that calibration campaign has not been run.
- + **Damage mapping screens are unvalidated proxies.** They are reported as candidates only.
- + **No in-situ measurement.** Every comparison in this paper is between two satellite reads, or between a satellite read and the same asset's own history. None is made against a field survey.

A fifth figure, flood validation intersection over union for each event against a named EMS activation or the Global Flood Database, is missing from this paper. It has not been measured. We will add it when a confirmed vector download path exists for one of the two references.

How the Record Is Used



Every product writes the same document: dated sections, a table with its denominator, and a list of the scenes behind the numbers.

Each report holds a method section that names the published source, and a results section whose table carries the measurement and its own denominator, so a bare percentage never appears. Where the method produced a chip, the file carries the date it was taken and the caption states it. A provenance list names every archive and dataset read, and a scene list gives the acquisition dates and granule identifiers. A vocabulary check runs on every report before it is rendered. The vocabulary is fixed across all products. The report never says confirmed or denied, and it says what moved, by how much, on which scenes, read on which day. A claims desk reads the record and decides. The record does not decide for it.

Reader	Question	What the record gives
Claims desk	Send an adjuster to this parcel?	Hectares and percent of parcel with the loss date, the radar reference and look windows, the threshold and the scene list
Portfolio manager	What changed since the last look?	An above or within class for each asset against its own history, with event fraction and last look
Lender under a monitoring rule	Is there an auditable file for each loan?	Satellite, sensor, scene dates, bands and method for every read, kept with the asset
Auditor or dispute	How was this number produced?	The published method, the threshold and the scenes, with the failures of the method stated in the same document

For the Brazilian rule the scene list answers directly the record items the resolution names: the satellite and sensor, the image date, the bands and the method⁵. A lender in a market with a looser rule can still keep the same file for each asset, and a supervisor that asks about physical risk in a portfolio⁴ receives the same dated answer.

We have noted that Annex III of the EU AI Act lists, as high-risk, AI systems intended for risk assessment and pricing in relation to natural persons in the case of life and health insurance³⁷. Flood or crop claims handling is not named there. Each customer makes its own legal assessment of its use.

A radar read can say that an area was flooded and cannot say who is liable. An asset that reads within may still have a loss that a satellite cannot see, such as damage inside a building. Each report states these limits.

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. 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. A plot or field the method cannot assess is reported as not evaluated, because a silent clear is the most expensive output a record can give.

Every record ships with the scenes, dates and parameters behind it, so a client can check any figure against the public archive it came from.

Axine Laboratories

axinelabs.com | contact@axinelabs.com

References



1. Swiss Re Institute (2026). *sigma* 1/2026: Natural catastrophes in 2025.
<https://www.swissre.com/institute/research/sigma-research/sigma-2026-01-natcat-2025-wildfire-storm-risk.html>
2. Swiss Re Institute (2026). Press release: wildfires, storms and floods contribute to a record 92% of global insured losses in 2025. <https://www.swissre.com/press-release/Wildfires-storms-floods-contribute-to-record-92-of-global-insured-losses-in-2025-says-Swiss-Re-Institute/7b39b1a5-b878-4a55-a5ff-bf5aa561a675>
3. EIOPA and European Central Bank (2024). *Towards a European system for natural catastrophe risk management*. Joint paper, December 2024.
https://ecb.europa.eu/pub/pdf/other/ecb.climateinsuranceprotectiongap_EIOPA202412~6403e0de2b.sv.pdf
4. ECB Banking Supervision (2020). *Guide on climate-related and environmental risks*, November 2020.
<https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.202011finalguideonclimate-relatedandenvironmentalrisks~58213f6564.en.pdf>
5. Conselho Monetario Nacional (2025). Resolucao CMN n. 5.267 of 28 November 2025, amending Section 7 (Monitoramento e Fiscalizacao), Chapter 2, of the Manual de Credito Rural.
https://www.normasbrasil.com.br/norma/resolucao-5267-2025_487044.html
6. King, M.D. et al. (2013). Spatial and temporal distribution of clouds observed by MODIS onboard the Terra and Aqua satellites. *IEEE Transactions on Geoscience and Remote Sensing* 51(7). https://atmosphere-imager.gsfc.nasa.gov/sites/default/files/ModAtmo/King_et_al.2013.pdf
7. Tarpanelli, A., Mondini, A.C. and Camici, S. (2022). Effectiveness of Sentinel-1 and Sentinel-2 for flood detection assessment in Europe. *Natural Hazards and Earth System Sciences* 22, 2473.
<https://nhess.copernicus.org/articles/22/2473/2022/nhess-22-2473-2022.html>
8. European Parliament Legislative Observatory (2023). Summary of the Commission proposal for an EU Solidarity Fund contribution, Italy (Emilia-Romagna floods): total direct damage EUR 8.5 billion.
<https://oeil.europarl.europa.eu/oeil/en/document-summary?id=1787774>
9. EU Space Programme (22 May 2023). *Historic floods hit Emilia-Romagna, Italy*. Image of the day. <https://eu-space.europa.eu/components/earth-observation-copernicus/image-of-day/historic-floods-hit-emilia-romagna-italy>
10. Guy Carpenter (June 2023). *2023 Mid-May Emilia-Romagna Flood*.
https://www.guycarp.com/insights/2023/06/Italy_Emil-Romagna_Flood_2023-05.html
11. European Environment Agency. *Economic losses and fatalities from weather- and climate-related extremes*, 1980 to 2023. <https://www.eea.europa.eu/en/analysis/publications/economic-losses-from-climate-extremes>
12. Joint Research Centre. Copernicus Emergency Management Service activation EMSR659, Flood in Emilia-Romagna, Italy (catalogue record, with the second activation EMSR664). DOI 10.2905/JRC.MDHXV50.
<https://data.jrc.ec.europa.eu/dataset/63248754-b1a1-4e52-b10e-112fb86dea18>
13. Réussir (28 September 2026). *Assurance prairies: appel à la mobilisation en Auvergne-Rhône-Alpes et dans tout le Massif central*. <https://www.reussir.fr/articles/assurance-prairies-appel-mobilisation-en-auvergne-rhone-alpes-et-tout-le-massif-central>
14. web-agri.fr (25 September 2026). *Assurance prairie: les éleveurs de ruminants exigent une totale transparence*.
<https://www.web-agri.fr/fnpl/t260>

15. Terre-net (16 September 2026). *Au Space, l'assurance des agriculteurs en question*. <https://www.terre-net.fr/assurances/article/903512/au-space-l-assurance-des-agriculteurs-en-question>
16. Copernicus Data Space Ecosystem (24 June 2026). Notice: Sentinel-1C and 1D final orbital configuration achieved. <https://dataspace.copernicus.eu/node/4010>
17. SentiWiki. Sentinel-1 mission. <https://sentiwiki.copernicus.eu/web/s1-mission>
18. Microsoft Planetary Computer. STAC collections *Sentinel 1 Radiometrically Terrain Corrected (RTC)* and *Sentinel-2 Level-2A*. <https://planetarycomputer.microsoft.com/api/stac/v1/collections/sentinel-2-l2a>
19. Pekel, J.-F., Cottam, A., Gorelick, N. and Belward, A.S. (2016). High-resolution mapping of global surface water and its long-term changes. *Nature* 540, 418-422. <https://www.nature.com/articles/nature20584>
20. World Resources Institute. *Aqueduct Floods Hazard Maps*. <https://www.wri.org/data/aqueduct-floods-hazard-maps>
21. Joint Research Centre. Global Wildfire Information System, data and services (burnt-area layers). <https://gwis.jrc.ec.europa.eu/applications/data-and-services>
22. UN-SPIDER. Recommended practice: flood mapping and damage assessment using Sentinel-1 SAR data in Google Earth Engine, step-by-step. <https://un-spider.org/advisory-support/recommended-practices/recommended-practice-google-earth-engine-flood-mapping/step-by-step>
23. Tuele, A., Cao, W., Plank, S. and Martinis, S. (2016). Sentinel-1-based flood mapping: a fully automated processing chain. *International Journal of Remote Sensing* 37(13), 2990-3004. DOI 10.1080/01431161.2016.1192304
24. Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics* 9(1), 62-66. DOI 10.1109/TSMC.1979.4310076
25. Nobre, A.D. et al. (2011). Height Above the Nearest Drainage, a hydrologically relevant new terrain model. *Journal of Hydrology* 404(1-2), 13-29. DOI 10.1016/j.jhydrol.2011.03.051
26. Rennó, C.D. et al. (2008). HAND, a new terrain descriptor using SRTM-DEM: mapping terra-firme rainforest environments in Amazonia. *Remote Sensing of Environment* 112(9), 3469-3481.
27. Dalponte, M., Marzini, S., Solano-Correa, Y.T., Tonon, G., Vescovo, L. and Gianelle, D. (2020). Mapping forest windthrows using high spatial resolution multispectral satellite images. *International Journal of Applied Earth Observation and Geoinformation* 93, 102206. DOI 10.1016/j.jag.2020.102206
28. Ruetschi, M., Small, D. and Waser, L.T. (2019). Rapid detection of windthrow events using Sentinel-1 C-band SAR data. *Remote Sensing* 11(2), 115.
29. Key, C.H. and Benson, N.C. (2006). Landscape assessment: ground measure of severity, the Composite Burn Index, and remote sensing of severity, the Normalized Burn Ratio. In Lutes, D.C. et al., *FIREMON: Fire Effects Monitoring and Inventory System*, USDA Forest Service RMRS-GTR-164.
30. Xu, H. (2006). Modification of normalised difference water index (NDWI) to enhance open water features in remotely sensed imagery. *International Journal of Remote Sensing* 27(14), 3025-3033. DOI 10.1080/01431160600589179
31. Zha, Y., Gao, J. and Ni, S. (2003). Use of normalized difference built-up index in automatically mapping urban areas from TM imagery. *International Journal of Remote Sensing* 24(3), 583-594. DOI 10.1080/01431160304987
32. McKee, T.B., Doesken, N.J. and Kleist, J. (1993). The relationship of drought frequency and duration to time scales. Eighth Conference on Applied Climatology, Anaheim, 17-22 January 1993. <https://climate.colostate.edu/pdfs/relationshipofdroughtfrequency.pdf>

- 33.** Svoboda, M., Hayes, M. and Wood, D. (2012). *Standardized Precipitation Index User Guide*. World Meteorological Organization, WMO-No. 1090.
- 34.** Keetch, J.J. and Byram, G.M. (1968). A drought index for forest fire control. USDA Forest Service Research Paper SE-38. <https://research.fs.usda.gov/treesearch/40>
- 35.** Expert Team on Climate Change Detection and Indices (ETCCDI). List of the 27 core indices. https://etccdi.pacificclimate.org/list_27_indices.shtml
- 36.** Lacaux, J.P., Turre, Y.M., Vignolles, C., Ndione, J.A. and Lafaye, M. (2007). Classification of ponds from high-spatial resolution remote sensing: application to Rift Valley Fever epidemics in Senegal. *Remote Sensing of Environment* 106(1), 66-74.
- 37.** Regulation (EU) 2024/1689 (Artificial Intelligence Act), Annex III, point 5. <https://artificialintelligenceact.eu/annex/3/>



CLAIMS AND PORTFOLIO RECORDS