



Per-field record

Forty-Three Seasons

of one **field**, read from open imagery

A weekly record for US row-crop fields built from public satellite archives, with crop type, land cover, boundaries, yield and management zones scored against public references and every failed gate printed.



Contents

■ Introduction	01
■ Why This Is Hard	03
■ The Axine Method	05
■ Results and Controls	07
■ Limits and What Is Not Evaluated	16
■ How the Record Is Used	17
■ About Axine Laboratories	18
■ References	19

Introduction



This paper describes the weekly field record that Axine Laboratories builds from open satellite archives for the people who advise, insure and lend on US row crops, and reports how each product on top of it performed against public references that nobody at Axine built.




Figure 1. Season peak NDVI for the Seward pivot, one bar per year from 1984 to 2026, coloured by the sensor that carried the peak. The vertical axis starts at 0.60.

We offer one cloud-screened look at each field every week, the field's own history back to 1984 to read it against, a plain-word status, and a document per field in which every figure traces to its source scene. The floor of the service (the register, the weekly look, weather and history) has been running since 30 September 2026 and carries no accuracy claim, because it records what the imagery shows and predicts nothing.

Several products sit on top of that floor: crop type, seasonal crop classification, land cover, field boundaries, yield, and a terrain-plus-vegetation management-zone map. Each was built against a named published method and scored against a public reference. This paper gives the result of each as a numerator over a denominator, names what failed and says what changed. Among them, field boundaries clear all three of their gates on the one gated Kansas box and still ship as candidates for review. Yield clears its gates in two of eight state-crop cells. The rest ship as a record with no accuracy claim, or labelled as screening.

Motivation: advisers and lenders ask for dated records

Remote sensing has moved from a convenience to a requirement in some parts of agricultural finance. In Brazil, Resolution 5.267 of the National Monetary Council took effect on 1 March 2026. It makes remote-sensing monitoring mandatory for rural credit operations for current costs (custeio) and investment (investimento) that are contracted from that date and cover a total area above 300 hectares, runs from before contracting until settlement, and requires auditable records of the satellite, sensor, image date, resolution, bands and method¹. In the United States, the 2026 Prevented Planting Standards Handbook added a paragraph titled "TALV Remote Inspections (Pilot Process)", which lets remote-sensed evidence enter the inspection of a prevented-planting claim². The 2025 edition had no such text.

- + **A large and variable base.** NASS estimated 95.3 million acres of corn and 85.4 million acres of soybeans planted in the United States for 2026³. Disease reduced corn yield by an estimated 6.0 percent across the United States in 2024 and by 7.0 percent in 2025 (United States and Ontario together)^{4,5}, so the field-level difference between seasons is large enough to matter to a claims desk.
- + **Insurance runs on acreage.** For crop year 2025 the Risk Management Agency reports 561,264 thousand net reported acres and \$12,155,226 thousand in indemnities across all of its lines of business⁶. The acres are RMA's net reported quantity and not a count of insured farmland.
- + **Fields are small units.** The 2022 Census of Agriculture counts 1.9 million farms and ranches with an average size of 463 acres⁷. That average covers all farms, including small and livestock operations, and gives no row-crop field size.



A green maize field with a high-voltage pylon and power lines under grey-white cumulus cloud.

Why This Is Hard



A map that looks right to a reader can still lose to a baseline that guesses the largest class, and the public references that a product is scored against are themselves imperfect.

Two readers, two questions

The agronomist, whether a retail agronomy lead or an independent crop consultant, is measured on acres covered per agronomist. Their question of any new reading is narrow: if this flags a field, is it worth the drive? That is why we score every product as a numerator over a denominator against a public reference, at the size and in the places a real Corn Belt or Plains field sits, and avoid a single headline accuracy number. An adviser who passes a record to a grower also needs to know the record will hold up in front of that grower, so a feature labelled as candidates for review carries no accuracy figure on any customer surface, and a screening product says so in the catalogue and on the page.

The claims desk, whether a crop insurance claims lead or an independent adjuster, asks something more exacting, because a federal pilot sets out what a remote-sensed file has to contain. As we read paragraph 62 of the 2026 handbook, it names imagery within the insurance period, each scene's source and capture time, a map with the field boundary overlaid, and a narrative that ties the record to the final planting date². That is close to the shape of the weekly record described in the last section. Water use per look, the one product that has passed its validation bar (section on yield and zones), is the reading we would hand to such a filing. A screening-grade reading does not go in.

Both readers share one rule, which we follow throughout: a claim is a figure printed in a dated science note with its denominator and quoted with its scope.

Public references are imperfect

Every reference available for free at field scale is either a classifier or an old snapshot.

- + **The Cropland Data Layer is a classifier.** NASS publishes it annually from 2008 for the conterminous United States as public domain data, and it moved from 30 m to 10 m resolution beginning in 2024⁸. Its reported overall crop accuracy for 2025 is 75.4 percent, against 81 to 82 percent for 2021 to 2023⁸. We read 2021 labels in the seasonal tests below, so they predate that change.
- + **WorldCereal learned from the same labels.** The ESA WorldCereal 10 m 2021 product reports user's and producer's accuracies of 88.5 percent and 92.1 percent for temporary-crop extent, and its own validation paper records that CDL points were among its training data in the United States⁹. Agreement between the two is therefore agreement between related classifiers.

- + **WorldCover sets a ceiling.** The ESA WorldCover 2021 v200 map has a global overall accuracy of 76.7 ± 0.5 percent¹⁰. A model trained on its labels cannot show more skill than the labels carry.
- + **No current field-boundary layer is public.** No current USDA Common Land Unit layer is public (Section 1619 of the 2008 Farm Bill). The only free CLU-derived reference we found for either demonstration county is Kansas Irrigated Cropland c.2007, which covers irrigated fields only and is a dated snapshot¹¹.
- + **Boundary models carry licence limits.** Fields of The World, the benchmark behind the boundary model we use, spans 24 countries on four continents and has no US data. Its labelled data keep their source licences, and some checkpoints are restricted to non-commercial use¹².



An aerial view of green pasture fields bounded by dark hedgerows, with one wind turbine and a small herd of cattle in the lower field.

Why a headline number misleads

Our first seasonal crop classifier run scored worse than guessing the largest class. Overall accuracy flatters a model on unbalanced ground, so we print each result beside its majority-class baseline, count a class recall only where at least 100 held-out pixels stand behind it, and publish a failed gate as failed.

The Axine Method



One cloud-screened look per field per week, read against that field's own history, with every number drawn from open archives and tied to the scene that produced it.

Data and sources

Sentinel-2 L2A (ESA, 10 m, 2016 onward) carries the weekly look. Landsat 4 through 9 Collection 2 Level 2 (USGS, 30 m, 1984 onward) extends a field's history back before Sentinel-2 existed; Landsat 5 returned data from its launch on 1 March 1984 until 5 June 2013¹³. Weather comes from the ERA5-backed historical API of Open-Meteo, with its forecast API covering the most recent days (CC BY 4.0). None of the data is commercial imagery.

Layer	Source	Terms and coverage
Weekly look	Sentinel-2 L2A, 10 m	ESA, 2016 onward
History	Landsat 4 to 9 Collection 2 Level 2, 30 m	USGS, 1984 onward
Crop type	WorldCereal 10 m 2021 v100; USDA Cropland Data Layer	CC BY 4.0, 2021 only; public domain, annual, 2008 onward
Land cover	ESA WorldCover 10 m 2021 v200	CC BY 4.0
Field boundaries	Fields of The World v3.1, CC-BY-4.0 checkpoint	Model weights
Yield and terrain	USDA NASS Quick Stats county yields; Copernicus DEM GLO-30	Public




Figure 2. From open archives to a record for one field.

The 43-season record

Figure 1 plots the full record of season peak NDVI for the Seward pivot, a 47.43 ha centre-pivot field that the Cropland Data Layer for 2021 labels 98 percent corn. There is one bar for each calendar year from 1984 to 2026, coloured by the sensor that carried that year's peak. The season of 2026 is the highest on record, at a peak NDVI of 0.952 on 1 August (Sentinel-2), ahead of 2018 (0.937) and 2012

(0.935, Landsat 7). The weakest seasons are 1993 (0.683), 2015 (0.689) and 1992 (0.728). Across the record the peak was read once from Landsat 4, 19 times from Landsat 5, 10 from Landsat 7, 2 from Landsat 8 and 11 from Sentinel-2.

No year from 1984 on lacks a clear look (at least 70 percent of the field cloud-free) over any of the three Kansas and Nebraska demonstration fields behind this record.

The Landsat 7 scan-line lesson

Building the record exposed a defect in our own peak rule. We first counted a look toward the season peak when it saw at least 30 percent of the field, the gate the weekly register uses. On that gate Landsat 7, whose Scan Line Corrector failed in 2003 and left diagonal data gaps on every scene from 2004 on, carried 11 to 13 bars per field. Comparing 1999 to 2002 (corrector working) with 2004 to 2011 (corrector failed) showed Landsat 7 reading as much as 0.22 NDVI higher than Landsat 5 in the failed years on one field, Pivot east in Gray County, Kansas.

The stripes were already declared nodata and dropped out correctly, so the gate was the cause. A Landsat 7 look that saw only part of a non-uniform field could see its greener part: on 26 April 2006 a look covering 51 percent of Pivot east read NDVI 0.80, while whole-field looks on either side of it read 0.25. A peak takes the maximum over the year, so that one partial look won the year whenever it existed. Landsat 7 after 2003 produces far more partial looks than any other sensor, so the effect showed up there first.

The shipped rule is a 70 percent gate. It brought the Landsat 7 minus Landsat 5 difference down to between -0.03 and +0.05 across three fields, corrector working or not. A gate tighter than 70 percent costs Landsat 7 too many of its remaining looks, and its peak then starts reading low from undersampling. We set the value on three fields and twelve years; a wider campaign may move it.

A boundary effect we left in place

We left a second effect uncorrected. Sentinel-2 reads 0.03 to 0.05 lower than Landsat on two demonstration fields whose square boundary is drawn around a circular pivot. The corners are not cropped ground, and Sentinel-2's 10 m grid samples them at their own value while Landsat's 30 m grid mixes corner and circle. A third field, drawn to its true rectangular edge, shows none (mean -0.002). Nothing is transformed between sources on account of it, and every bar in the register and in Figure 1 names its own sensor, so a reader can see the seam.

The prior-season baseline

A field's weekly status (steady, moved, no clear look) is judged against the same ISO week of up to five earlier seasons when the crop type matches, and against the field's own trailing weeks when it does not. The dryland strip's wheat harvest still reads "moved" against its own trailing weeks, because in nine of its eleven Sentinel-2 years its annual peak falls in autumn (a different crop, sorghum or triticale in the Cropland Data Layer), and only two years, 2021 and 2026, show the spring wheat pattern a same-week comparison expects. Fixing this needs a crop type read for each season.

Results and Controls



Each product is described by the published method followed, the public reference used, the result as a numerator over a denominator, and what failed and what changed.

Product	Public reference	Status as published
Crop type	WorldCereal 2021, Cropland Data Layer	Agreement between two classifiers, no accuracy claim of our own
Seasonal crop classification	Cropland Data Layer, 5 km blocks	Agreement on held-out ground, 6 classes with 100+ pixels
Land cover	ESA WorldCover 2021	Agreement on held-out ground, 3 km blocks
Field boundaries	Kansas Irrigated Cropland c.2007	Candidates for review, no accuracy figure
Yield	NASS county yields	2 of 8 state-crop cells pass; corn passes nowhere
Management zones	Terrain and vegetation indices	No accuracy figure
Water use per look	22 independent sites	All seven pre-registered gates passed

Crop type: WorldCereal and the Cropland Data Layer

We read ESA WorldCereal's published global product (10 m, 2021, v100) for each field by agro-ecological zone tile, and the Cropland Data Layer for each field through the CropScape service (annual, 2008 onward). Neither is fitted by us. Both are classifiers and neither is a field visit, and the CDL points that WorldCereal trained on are the overlap noted earlier, so a comparison between them is agreement and cannot be an independent accuracy figure.

Four of four demonstration fields agree between WorldCereal 2021 and the Cropland Data Layer 2021. That agreement surfaced two findings about our own register. Two fields registered as corn sit on grassland or pasture in both maps in every year checked (CDL 2021 to 2025, WorldCereal 2021). A third field's rectangle mixes sorghum, grassland, alfalfa and triticale and is not the wheat field it was registered as. With WorldCereal published for 2021 only, no field had two earlier seasons of a matching crop type, so the baseline fix above could not yet be exercised. We shipped the Cropland Data Layer, annual and public domain, as a second named US source to close that gap.

What may be said about crop type is WorldCereal's own published figures (user's and producer's accuracy of 88.5 and 92.1 percent for temporary-crop extent, against about 50,000 expert-interpreted sample sites⁹), attributed to WorldCereal and with their caveat about spatial bias in the

validation sample. The 4-of-4 check is a sanity check and carries no accuracy claim. We give no crop type for a season that neither source has published.

Seasonal crop classification

We followed Kussul, Lavreniuk, Skakun and Shelestov, who showed that a multi-temporal feature stack over a growing season separates crop types that a single date cannot, because crops that look alike on one date differ in when they green up and senesce¹⁴. We built a seasonal composite of up to six Sentinel-2 looks, NDVI maximum, minimum and amplitude as phenology features, optional Sentinel-1 radar, and a classifier trained fresh for each job on the area's own public labels (the Cropland Data Layer for a US area, WorldCereal for a European one). The reference is the Cropland Data Layer with an exact-year match, so 2021 imagery is scored against the 2021 layer, with 5 km spatial blocks held out so that no training block touches a test block.

What failed and what changed

The first complete run scored worse than guessing the largest class: balanced accuracy 0.1269 against a majority baseline of 0.3642, with corn and soybean recall both around 17 percent. The cause was a label grid built south-up by its own documented convention, while every feature grid we read is north-up. Every Cropland Data Layer label was mirrored top to bottom against the Sentinel-2 pixels it trained on. We confirmed it by printing the north-south term of each grid's transform on the same area (-0.00043038 for the feature grid and +0.00043038 for the label grid, opposite signs, as the fault required), fixed it by flipping the label grid after every read, and pinned the fix with a regression test.




Figure 3. Seasonal crop classification in Story, Boone and Hamilton counties, Iowa, 2021. Left: balanced accuracy before and after the label-grid fix, with the majority-class baseline marked on each. Right: recall by class after the fix, with classes under 100 held-out pixels shown as grey diamonds.

After the fix we scored 65,814 held-out cells (season 1 May to 30 September 2021). Corn recall is $13,728/16,853 = 0.8146$ and soybean recall $13,906/16,870 = 0.8243$, both well above the 100-pixel floor below which a recall is not a measurement. Non-crop and other land reaches 0.8930 (n 28,966),

alfalfa 0.4409 (n 2,098), oats 0.0778 (n 527) and other hay at 0.0432 (n 486). Two classes, other crop (n 9) and winter wheat (n 5), have too few held-out pixels for their recall to mean anything, and both read 0. We exclude them from the headline figure, so balanced accuracy is reported two ways: 0.3867 over all 8 classes, and 0.5156 over the 6 classes with at least 100 held-out pixels. The majority baselines are 0.4401 and 0.4403, so the 6-class figure carries a skill of +0.0753, and we regard that as the honest one.

Our earlier single-date crop-type card, trained once on a far larger pixel count on the same area and imagery, measures a producer's accuracy of 0.96 for both corn and soybeans on its own spatial split. The seasonal classifier's 0.81 and 0.82 sits lower but in the same range. Sentinel-1 radar and the WorldCereal label path were built and were not exercised in this run, so both remain open measurements.

Land cover

The method is a seasonal Sentinel-2 composite classified against public reference data, with area reported in hectares and percent and a two-period change path. We train for each job against the area's own ESA WorldCover read and do not republish WorldCover's fixed 2021 map, which is what makes a date other than 2021 classifiable at all. The reference is ESA WorldCover 10 m v200 (2021, CC BY 4.0), whose global overall accuracy of 76.7 percent is a ceiling on what a model trained on its labels can show¹⁰.

What failed and what changed

The land-cover reader had the same grid-orientation fault as the crop classifier and received the same fix. After it, on the Sacramento Valley, California (season 1 July to 31 October 2021, 3 km spatial blocks, 44,569 held-out cells), the two weakest classes were the two rarest. Bare or sparse ground reached a recall of 0.1915 (n 632) and tree 0.6251 (n 4,372), against grass at 0.8709 (n 18,906) and built-up at 0.8799 (n 9,432). That is the signature of an unweighted fit being outvoted by whichever class covers the most ground.

A second fix on 6 October 2026 added NDVI maximum and minimum, a brightness feature and the Bare Soil Index, and weighted every fit by the inverse frequency of each pixel's class in the training set. Bare or sparse recall more than tripled, from $632 \times 0.1915 = 121$ to $632 \times 0.6661 = 421$ correct of 632 (0.19 to 0.67), and tree recall rose from 0.63 to 0.76 (n 4,372). Balanced accuracy rose from 0.5844 to 0.7060 against a 0.4242 majority baseline, a skill of +0.1602 rising to +0.2818. Overall accuracy fell from 0.8141 to 0.7990, a small and expected cost: some grass, built-up and cropland pixels near a class boundary are now called something else because the minority classes no longer concede the decision boundary by default. The new 0.7060 sits close to the 0.7768 balanced accuracy of the existing single-date card, which we name as the ceiling that a four-look classifier trained for each job is not expected to beat.




Figure 4. Land-cover recall by class in the Sacramento Valley before the second fix (unweighted) and after it. The majority-class baseline is 0.424, and balanced accuracy rose from 0.584 to 0.706.

Field boundaries

Boundary candidates clear all three gates on the one gated Kansas box after round two, and they still ship with a human accept step and no accuracy figure.

Method and reference

We use Fields of The World v3.1, the "PRUE" U-Net with an EfficientNet-B7 encoder, in its CC-BY-4.0 checkpoint trained only on CC-BY and CC0 data¹². Delineate Anything (AGPL-3.0 weights and runtime)¹⁵ and the benchmark's own "Standard" checkpoints (non-commercial training data) were ruled out on licence and never imported into the served pipeline. The model runs as ONNX on CPU, and a 12 by 12 km read cost 73 to 81 seconds in production at the time of round one.

The only gated reference is Kansas Irrigated Cropland c.2007, covering irrigated fields only¹¹. For Nebraska, Iowa and Illinois no CLU-derived layer exists, and USDA NASS Crop Sequence Boundaries, itself a classifier output, is read as agreement with no gate.

Round one, 30 September 2026

The gated box is Gray County, Kansas, season 2018, with 96 reference fields and 160 candidates offered. Object recall at an intersection-over-union (IoU) above 0.5 was $36/96 = 0.375$, which passed its gate of 0.35 or more. Conditional precision, the share of candidates that overlap a reference field by 10 percent or more that are the right field, was $36/65 = 0.554$, which failed its gate of 0.58 or more. Median per-object IoU was 0.369, which failed its gate of above 0.50. On the 20 to 60 ha fields that dominate this window, 34 of 71 were recalled and 33 of 36 of those were the right field (conditional precision 0.92). Fields under 20 ha were almost never found (2 of 20) and fields over 60 ha fared worst of all (0 of 5 recalled, conditional precision 1 of 13). The verdict was candidates for review only, with no accuracy figure on any customer surface.

A lever that failed, 5 October 2026

We hypothesised that the thin-boundaries option of the benchmark, which folds the boundary class back into the field mask and cuts only a one-pixel Zhang-Suen skeleton¹⁶, would help the over-60 ha bucket. We tested it on synthetic shapes, then re-scored it on the real checkpoint and the real Gray County box. It made both failing gates worse: conditional precision fell from 0.554 to 0.500 and median IoU from 0.369 to 0.331, and the over-60 ha bucket did not move at all. A direct check found dense, correctly placed boundary pixels already separating the large reference fields, so the lever had no gap to recover, and its tendency to bridge thin boundary bands at a field's short ends and at junctions cost slightly more than it gained. It ships off by default, and we kept the implementation and its self-tests.

Round two, 6 October 2026

Four further levers were pre-registered and scored on the real checkpoint across three boxes: K (Gray County, Kansas, the only gated box), I (Benton County, Iowa, agreement with Crop Sequence Boundaries only) and L (Piatt County, Illinois, agreement only). Figure 5 plots recall, conditional precision and median IoU for every lever on box K, with conditional precision by field-size bucket for three of them.

Overlapping tiles with test-time augmentation (half-stride tiling, with the softmax averaged over the 8 D4 views) improved recall, conditional precision and median IoU on every box measured, and it was the only lever to do so cleanly. On K it reached a recall of 0.510 (49/96), a conditional precision of 0.754 (49/65, clearing the 0.58 gate) and a median IoU of 0.605 (clearing the 0.50 gate), at a real cost of roughly 8 CPU-minutes for each 10 by 10 km box in place of roughly 2. It is now the default.

A single end-of-season date fed twice, in place of the benchmark's own two-date stack, scored higher still on K alone (conditional precision 0.808, median IoU 0.825, recall 0.615). We did not promote it. It improved K and Iowa's agreement precision (0.244 to 0.257), but it worsened Illinois precision (0.356 to 0.248) and Iowa's median IoU slightly.




Figure 5. Boundary levers scored on the gated Kansas box. Left: recall, conditional precision and median IoU for each lever, with the 0.58 precision gate and 0.50 IoU gate marked and the licence-barred comparison model hatched. Right: conditional precision by field-size bucket for the baseline, the single end-of-season date and the adopted lever.

We state the decision rule precisely, because the rule as written ("the same direction holding on at least one of I or L") is technically satisfied by Iowa alone. The code's own documentation records the stricter bar applied in practice, which is that the direction had to hold on the other boxes measured, in the plural, and not on one. We report the gap between the rule's wording and its application because this paper's standard asks us to.

Delineate Anything v1¹⁵, scored for comparison only and unusable in the served pipeline whatever its score, reached the best median IoU of any lever on K (0.889) and narrowly missed the conditional-precision gate (0.537 against 0.58).

The size buckets behind the right-hand panel rest on small denominators. Over 60 ha, the single end-of-season date scores 0 of 3 and the adopted lever 3 of 10. Fields under 20 ha and over 60 ha remain the weakest size classes under every lever we tried.

Verdict

Boundaries ship as candidates for review with a human accept step. No accuracy figure appears on any customer-facing page.

Yield

Yield passes its pre-registered gates in two of eight state-crop cells.

We combined a trend-plus-departure design, a statistical approach to yield models of the kind discussed by Lobell and Burke¹⁷, with gradient boosting on phenology and weather features. Our first model was a peak-NDVI county regression of the kind NASS's own operational work has examined^{18, 19}. It appears below as the old line.

The reference is USDA NASS Quick Stats county yields, scored leave-one-year-out over 2018 to 2025, with gates fixed before any value was read: RMSE no higher than 20.0 bu/ac for corn and 8.0 for soybeans, and R^2 of at least 0.50. The panel has 12 counties in each of Iowa, Illinois, Nebraska and Kansas, with two 5 km boxes each.

The old model, 30 September 2026

The model was a straight line of yield on peak NDVI, one line per state and crop, with no county-specific term. It passed 0 of 8 state-crop cells. The county's own mean yield over its other years already carries everything that a county-blind model cannot see (soil, drainage, management). A model with no county information was being compared with a baseline that is almost entirely county information, and it lost in all but the weakest soybean cells.

The pooled model, 5 October 2026

The pooled model keeps the county control as a feature, defined identically and with no leakage of the scored year. It adds growing-season and mid-season growing-degree-day and precipitation anomalies for each county, keeps peak NDVI and its anomaly, and pools across all four states for each crop, so that each fold trains on roughly 600 to 670 rows in place of roughly 80.




Figure 6. Leave-one-year-out RMSE of the pooled model and of the county-mean control for each state and crop, with the gate marked (20 bu/ac for corn, 8 for soybeans). Passing cells are green, failing cells hatched red, and R^2 is printed for each cell.

Two of eight cells now pass every gate (Figure 6). Illinois soybeans reach an RMSE of 4.92 against a control of 5.04 and an R^2 of 0.56. Kansas soybeans reach 7.34 against a control of 10.29 and an R^2 of 0.53. These are the first passes this product line has produced.

Corn improves substantially in R^2 in three of four states (Illinois from -0.07 to 0.42, Kansas from 0.14 to 0.27, Nebraska from 0.25 to 0.32) but clears no state's R^2 gate of 0.50, and no corn RMSE comfortably beats its own control. Illinois corn is the closest, at 15.67 against 16.14. Corn ships on no customer surface. Iowa soybeans is the one cell that got slightly worse, with R^2 moving from 0.00 to -0.08, both at noise level. Some failing cells in Figure 6 sit below the RMSE gate and fail on R^2 , among them Illinois corn, Iowa soybeans and Nebraska soybeans. Every passing estimate carries an empirical 80 percent prediction interval, the 10th and 90th percentile of that cell's leave-one-year-out residuals, and we do not call it a forecast.

Management zones for variable-rate application

A management zone is a clustering of published terrain and vegetation indices and makes no prediction that can be scored against ground truth, so we print no accuracy figure for it. The same reasoning applies to our burn-severity and canopy-state screening products. The layers are the field register's current-season relative vegetation composite, the Topographic Wetness Index (Beven and Kirkby 1979²⁰; the digital raster formula of Moore, Grayson and Ladson 1991²¹) from a D8 flow-accumulation implementation²² over the Copernicus DEM, and a soil-moisture relative-wetness pattern from Sentinel-1 change detection, with a same-cost NDMI fallback²³ when that pattern cannot be read. Each layer is scaled 0 to 1 over the field, combined by a user-set weight (a missing layer's weight is dropped and the rest renormalised) and cut into quantile zones, then cleaned by a majority filter and a sieve.




Figure 7. The Seward pivot. Left: Topographic Wetness Index from the Copernicus DEM (colour bar 7 to 13). Right: three management zones of 17.8, 15.5 and 13.4 ha. Cells are 25.27 m wide and 33.17 m tall, and the scale bar is 200 m.

Figure 7 comes from a live run against the real Seward pivot boundary and is not a transcribed number: the run executed the shipped pipeline, with network reads to the Copernicus DEM, Sentinel-2 and, in this run, a successful Sentinel-1 soil-moisture read in place of the NDMI fallback. Zone 1, the lowest scoring (17.8 ha), has a relative vegetation of 0.97 and a TWI of 9.10. Zone 2 (15.5 ha) has

1.01 and 8.92. Zone 3, the highest (13.4 ha), has 1.03 and 9.57. The zones separate mainly on the vegetation layer, which carries half the combined weight by default, with a smaller terrain contribution, which is the behaviour that the weighted-combine self-tests check for.

Water use per look

This product is not new in this paper. It is the one reading we offer that already clears our validation bar: an independent test on 22 fresh sites and 242 observations gave a pooled bias of -6.6 percent, an RMSE of 1.17 mm/day and an r of 0.826, with all seven pre-registered gates passed. For scale, the published method it follows, SSEBop with the FANO cold boundary, reports a pooled bias of +3.0 percent, an RMSE of 1.36 and an r of 0.76 over 23 towers²⁴. We give that as literature context only. Irrigated cropland in arid climates is excluded from sale because the model reads about 40 percent low there.



Rolling hills with a dark ploughed field in the foreground, pale straw-coloured crop or stubble on the slope behind and a green strip on the far hill.

Limits and What Is Not Evaluated



Each product has a stated scope. This section lists what the numbers above do not cover.

2 of 8

state-crop yield cells that pass every gate

0.754

conditional precision on the gated Kansas box (gate 0.58)

0.5156

balanced accuracy over 6 crop classes (baseline 0.4403)

- + **Crop classification and land cover.** Both are agreement with a public classifier (the Cropland Data Layer, WorldCereal, WorldCover) on spatially held-out ground. Neither is accuracy against an independent field observation, and WorldCover's own 76.7 percent global accuracy is a ceiling on what training against its labels can show. Both products train fresh for each job and do not use a cached regional model, which is an open question for a future measurement.
- + **Field boundaries.** The adopted gain from overlapping tiles with test-time augmentation is real, but the only gated reference is a seventeen-year-old irrigated-only snapshot. Fields under 20 ha and over 60 ha remain the weakest size classes under every lever tried. No public CLU-derived reference covers the Iowa, Illinois or Nebraska counties we scored (a free 2008 CLU extract exists for parts of Iowa, Washington and Arkansas, and not for Benton County or the Nebraska and Illinois counties used here), so every number there is agreement with another classifier's output (NASS Crop Sequence Boundaries).
- + **Yield.** Passes occur in exactly two of eight state-crop cells (Illinois and Kansas soybeans). Corn ships nowhere a customer can see it, in any state. The 2026 forward record of 2,777 fields, locked by checksum, was built under the old model's gates and has not been re-scored under the pooled model.
- + **Management zones.** They read one season of vegetation and miss the multi-year pattern that the register's own multi-season zoning tool draws on, because an arbitrary area run through the service has no multi-year field history. The moisture layer falls back to NDMI on a field where Sentinel-1 cannot be read, and each run names its moisture source. The ISOXML (ISO 11783-10) export is written and round-trips through our own parser, but it has not been checked against the official schema or a physical ISOBUS controller, and it is not yet wired to a download route.
- + **Crop type.** WorldCereal exists for 2021 only. Every other season reads the Cropland Data Layer where the field is in the conterminous United States and returns nothing outside it. The agreement check covers four fields and is not a campaign.
- + **Geography.** Everything we offer stops at the US border until the weather source, presently a commercial plan, is replaced with one licensed for global commercial use. None of these products has been tried on a non-US field.

How the Record Is Used



One document per field, issued weekly and on demand, that a reader can check against its sources.

The customer-facing document is a single PDF for each field. It holds the register's current status, the 43-season history chart (the chart of Figure 1, without a figure number), the weather-derived growing-degree-day and rainfall lines, and the crop type with its source named. Where a product is validated or has a passing cell, the record also carries that product's own figure with its source scene list. A record issued under a partner's name says so wherever authorship appears, and the first page's footer still names Axine Laboratories as the provider. End to end on production, the free-tier version produced a 9-page record for the Seward pivot's own boundary in 16 minutes.

For an insurer, the record maps onto what we read in paragraph 62 of the handbook²: imagery within the insurance period with its source and capture time (the register's scene provenance, carried on every look), a map with the field boundary overlaid (the register's boundary layer), and a narrative tying the record to a date (the weekly status line). The remote inspection pilot is new in the 2026 edition. For a lender, the same discipline of dated, traceable scenes matches the shape of the records that Brazil's Resolution 5.267 asks institutions to keep¹.

What the record does not yet carry

Two things a claims-desk engagement would need us to build are absent. One is a narrative generator aimed at a final planting date. The other is a reconciliation of claimed acreage against the remotely determined acreage. Today's record states what stood on the field each week, and it does not say whether that satisfies the test of a specific claim.

How to read a figure in the record

A figure with a validated product behind it can be quoted with its scope. A figure from an agreement product, a candidate-for-review feature or a screening product is a record of what a public classifier or a published index shows, and it is labelled that way on the page. The question we ask of every reading is the agronomist's: does it justify the drive to the field? The numbers in this paper are meant to let a reader answer that for each product on its own terms, including the ones that did not pass.

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. Conselho Monetário Nacional. Resolução CMN nº 5.267, 28 November 2025, amending Section 7 (Monitoramento e Fiscalização), Chapter 2 of the Manual de Crédito Rural. https://www.normasbrasil.com.br/norma/resolucao-5267-2025_487044.html
2. Federal Crop Insurance Corporation. *2026 Prevented Planting Standards Handbook*, FCIC-25370, paragraph 62 (TALV Remote Inspections, Pilot Process). USDA Risk Management Agency, 2025. <https://www.rma.usda.gov/sites/default/files/2025-12/2026-25370-Prevented-Planting-Standards-Handbook.pdf>
3. USDA National Agricultural Statistics Service. *Acreage* report, 30 June 2026. <https://www.nass.usda.gov/Newsroom/2026/06-30-2026.php>
4. Crop Protection Network. Corn disease loss estimates from the United States and Ontario, Canada, 2024. <https://cropprotectionnetwork.org/publications/corn-disease-loss-estimates-from-the-united-states-and-ontario-canada-2024>
5. Crop Protection Network. Corn disease loss estimates from the United States and Ontario, Canada, 2025, published 6 March 2026. <https://cropprotectionnetwork.org/publications/corn-disease-loss-estimates-from-the-united-states-and-ontario-canada-2025>
6. USDA Risk Management Agency. *Summary of Business*, national totals as of 5 October 2026. https://pubfs-rma.fpac.usda.gov/pub/sob/current_week/sobrpt2023-2026.pdf
7. USDA National Agricultural Statistics Service. *2022 Census of Agriculture*, release of 13 February 2024. <https://www.nass.usda.gov/Newsroom/2024/02-13-2024.php>
8. USDA National Agricultural Statistics Service. *Cropland Data Layer*, documentation and frequently asked questions. https://www.nass.usda.gov/Research_and_Science/Cropland/SARS1a.php
9. Van Tricht, K., et al. (2023). WorldCereal: a dynamic open-source system for global-scale, seasonal, and reproducible crop and irrigation mapping. *Earth System Science Data* 15:5491-5512. doi:10.5194/essd-15-5491-2023
10. ESA WorldCover 10 m 2021 v200 Product Validation Report. https://pure.iiasa.ac.at/id/eprint/18981/1/WorldCover_PVR_V2.0.pdf
11. Kansas Irrigated Cropland c.2007. Kansas Applied Remote Sensing (KARS), Kansas Biological Survey, University of Kansas.
12. Kerner, H., et al. (2025). Fields of The World: A Machine Learning Benchmark Dataset for Global Agricultural Field Boundary Segmentation. *Proceedings of the AAAI Conference on Artificial Intelligence* 39(27):28151-28159. arXiv:2409.16252
13. US Geological Survey. Landsat 5. <https://www.usgs.gov/landsat-missions/landsat-5>
14. Kussul, N., Lavreniuk, M., Skakun, S. and Shelestov, A. (2017). Deep learning classification of land cover and crop types using remote sensing data. *IEEE Geoscience and Remote Sensing Letters* 14(5). doi:10.1109/LGRS.2017.2681128
15. Lavreniuk, M., et al. (2025). DelineateAnything. arXiv:2504.02534
16. Zhang, T.Y. and Suen, C.Y. (1984). A fast parallel algorithm for thinning digital patterns. *Communications of the ACM* 27(3):236-239.

- 17.** Lobell, D.B. and Burke, M.B. (2010). On the use of statistical models to predict crop yield responses to climate change. *Agricultural and Forest Meteorology* 150:1443-1452.
- 18.** Johnson, D.M. (2014). An assessment of pre- and within-season remotely sensed variables for forecasting corn and soybean yields in the United States. *Remote Sensing of Environment* 141:116-128.
- 19.** Johnson, D.M., Rosales, A., Mueller, R., Reynolds, C., Frantz, R., Anyamba, A., Pak, E. and Tucker, C. (2021). USA crop yield estimation with MODIS NDVI: are remotely sensed models better than simple trend analyses? *Remote Sensing* 13(21):4227. doi:10.3390/rs13214227
- 20.** Beven, K.J. and Kirkby, M.J. (1979). A physically based, variable contributing area model of basin hydrology. *Hydrological Sciences Bulletin* 24:43-69.
- 21.** Moore, I.D., Grayson, R.B. and Ladson, A.R. (1991). Digital terrain modelling: a review of hydrological, geomorphological, and biological applications. *Hydrological Processes* 5:3-30.
- 22.** O'Callaghan, J.F. and Mark, D.M. (1984). The extraction of drainage networks from digital elevation data. *Computer Vision, Graphics, and Image Processing* 28:323-344.
- 23.** Gao, B.-C. (1996). NDWI: a normalized difference water index for remote sensing of vegetation liquid water from space. *Remote Sensing of Environment* 58(3):257-266.
- 24.** Senay, G.B., Parrish, G.E.L., Schauer, M., Friedrichs, M., Khand, K., Boiko, O., Kagone, S., Dittmeier, R., Arab, S. and Ji, L. (2023). Improving the Operational Simplified Surface Energy Balance evapotranspiration model using the Forcing and Normalizing Operation. *Remote Sensing* 15:260.



THE PER-FIELD AGRICULTURAL RECORD