

FIELDWORK / AGRICULTURAL INTELLIGENCE

From machine records to a defensible field trial.

Two worked projects: machinery-data alignment and quality control, followed by replication planning for a small yield difference.

TECHNICAL DEMONSTRATION REPORT / 04 OCTOBER 2026

52,827	2,724	33,011
Raw yield observations	Unique QA-flagged records	QA-retained, spatially matched records

01 / What errors did we find, and how did we handle them?

Using frozen public machinery files, we aligned the 2018 harvest and application coordinate systems, removed duplicate records from the analysis view, screened implausible measurements, and withheld ambiguous spatial matches. Every original row remains in an audit ledger.

The two harvests have different matching methods. Their combined counts summarize processing only; they are not pooled experimental evidence. Quality thresholds can flag genuine extremes, so flagged observations are not all proven instrument errors.

02 / How many replicate plots are needed?

For two treatments randomized within independent paired blocks, detecting **2 bu/ac** at 80% power requires **52 complete paired blocks / 104 plots** if the standard deviation of the block differences is **5 bu/ac**. At 90% power, the same assumptions require 68 blocks / 136 plots.

A conditional calculation, not a field prescription.

The required count depends on the effect worth detecting, the variance of plot-level treatment differences, the trial design and independence between blocks. Neither GPS point count nor historical 20 m analysis cells establishes experimental replication.

EVIDENCE STATUS

Machinery findings are measured from OFPEDATA, pinned to commit ae70c5d8ad4fe154fd4b746a7ec43f3031a6e0a9. Replication results are prospective scenarios. GPS drift, machine offsets, sensor delay and yield calibration were not established or corrected. [1]

Align the coordinate frame. Preserve the measurements.


The 2018 harvest uses longitude and latitude in WGS84 (EPSG:4326); application points use WGS84 / UTM zone 12N (EPSG:32612). Comparing those numeric coordinates directly would mix degrees and meters. We transformed the harvest positions with an explicit longitude-first axis order. [1, 2]

Layer / action	Coordinate system	Unit
2018 harvest: source	EPSG:4326 / WGS84	Degrees
2018 harvest: working view	EPSG:32612 / UTM 12N	Meters
2018 application: unchanged	EPSG:32612 / UTM 12N	Meters
2016 harvest + application	Source UTM 12N / GRS80, unspecified datum	Meters

ONE REAL TRANSFORMED RECORD / 2018 SOURCE ROW 1

Longitude -110.5533800102; latitude 47.8816623215
Easting 533391.9291 m; northing 5303244.1565 m

The record still carries yield 50.1587 bu/ac and moisture 11.8%. Reprojection changes the coordinate representation; it does not change the observation or prove the physical location is accurate. The reverse-transform numerical error is below 3×10^{-14} degrees, a consistency check rather than a GNSS accuracy estimate.



Same local map frame and yield scale. Left: original yield records, with QA flags in orange. Right: records passing QA and matching an application point within 12 m. The loss of coverage is visible; no missing measurements are invented.

2016 datum boundary

Both 2016 layers have the same source projection definition, but its datum is unspecified. We preserve their coordinates for the within-year polygon join. We do not assign a guessed EPSG code or export a supposedly verified WGS84 layer. A rigorous cross-year geodetic comparison remains unresolved.

FIELDWORK METHODS / 01B

What was found. What was changed or withheld.

The balanced screening profile is a documented analytical choice, not a universal calibration rule. Reasons overlap; count unique flagged records rather than summing the rows below.

Diagnostic	2018	2016	Handling in analysis view
Duplicate x / y / yield	5	197	Keep first occurrence; preserve all source rows.
Non-positive yield or crop flow	0	1,146	Flag and exclude; do not impute.
Recorded speed below 1.5 mph	342	n/a	Flag 2018 low-speed observations. Speed was not supplied for 2016.
Moisture outside 8-22%	247	82	Flag outside stated plausibility range.
Non-positive recorded swath	0	833	Flag and exclude. Source units retained.
Yield distribution tails	626	1,501	Flag outside central 98% of positive yield values.
Invalid required value	0	0	Check required finite values; no replacements needed.
Unique records flagged	1,143	1,581	One exclusion per row, even with multiple flags.

The yield-tail limits were 6.90145-80.889875 bu/ac in 2018 and 2.77095-54.51302 bu/ac in 2016. These are dataset-specific distribution cutoffs. They can exclude legitimate low- or high-yield areas, so the raw view and flag ledger are essential for review.

Processing stage	2018 point join	2016 polygon join
Original yield observations	31,276	21,551
Pass quality screening	30,133	19,970
Pass QA + spatial matching	15,302	17,709
Retained, but unmatched	14,831	2,261

Corrected, screened, unresolved

Corrected representation: the 2018 coordinate-system mismatch. **Screened:** duplicates and the explicit QA flags above. **Unresolved:** instrument calibration, timestamp-based synchronization, machine offsets and the 2016 datum. No raw measurement values were overwritten.

A spatial join needs a coverage and ambiguity audit.

2018 / nearest application point

The nearest point is assigned only when its distance is within the selected radius. This is a proximity rule; it does not establish synchronized machine activity or randomized treatment assignment.

Radius	Matched after QA	Coverage of retained yield	Still unmatched
8 m	14,151	47.0%	15,982
12 m	15,302	50.8%	14,831
20 m	15,464	51.3%	14,669

Moving from 12 m to 20 m adds only 162 records. It does not repair the broader coverage gap. Radius selection should reflect geometry, implement footprint and positional uncertainty, then be checked for analytical sensitivity.

Numerical coincidence is not positioning accuracy.

13,136 QA-retained harvest points lie within 1 mm of an application coordinate. The median matched separation is approximately 0.000122 m. These nearly coincident source positions are not independent survey control, so they cannot substantiate sub-millimeter GNSS accuracy.

2016 / point inside application polygon

When a yield point intersects multiple application polygons, we use the median recorded rate only if the rate spread is at most 5 source lb/ac. Wider disagreement is withheld as ambiguous. That threshold is an explicit demo rule, not a validated agronomic tolerance.

Overlap diagnostic	Records
Conflicting overlap, all raw yield rows	592
Conflicting overlap among QA-retained rows	546
Other retained rows without an accepted polygon match	1,715

The stored distance of zero for a polygon match means "inside a polygon". It is not a measured zero GPS error. Source application-rate units remain lb/ac; unresolved product-versus-nitrogen mass basis is not silently converted.

Inspect the full chain

The deliverable includes one audit CSV per harvest, two filtered analysis CSVs, a 2018 WGS84 GeoJSON, all raw/screened/matched maps, source hash checks and the reproducible scripts. Each row preserves source identity, source coordinates, working coordinates, quality reasons and match status.

FIELDWORK METHODS / 02A

Define the replicate before counting the plots.

For this calculation, each independent block contains one plot of treatment A and one plot of treatment B. Randomize A/B allocation within each block. The plot receiving a treatment is the experimental unit; the block produces one paired difference for analysis. [3]

Quantity	What counts
r independent paired blocks	r plots per treatment; 2r treatment plots in total.
One block difference	D = plot-average yield A minus plot-average yield B.
GPS records inside a plot	Subsamples used to estimate that plot's yield, not independent treatment replicates.
20 m cells or bootstrap blocks	Historical analytical summaries, not evidence of randomized experimental units.

The decisions that determine r

- 1. Minimum relevant effect:** specify the absolute difference worth detecting before looking at trial results. Here the worked example is 2 bu/ac, about 3.61% of the 2018 screened median of 55.4157 bu/ac; that median is a reference, not a forecast.
- 2. Error rate and power:** this report uses a two-sided significance level of 0.05, and either 80% or 90% power. Power is the probability of rejecting the null under the specified effect and variance assumptions.
- 3. Plot-scale noise:** estimate the standard deviation of within-block treatment differences from a suitable pilot or compatible replicated trials. Raw point variation and yield-map residuals are not interchangeable with this variance.
- 4. Independence and operations:** choose block placement, plot dimensions, machinery width, buffers and harvest aggregation so the intended design is implementable and its error structure is credible. If residual block differences are spatially correlated, use a design-specific model and simulation.


The variance input must match the planned analysis.

If both treatment plots have standard deviation σ_P and within-block correlation ρ , then $SD(D) = \sigma_P \times \sqrt{2(1-\rho)}$. Correlation is a planning input to estimate, not a benefit to assume without evidence.

The paired-t calculation assumes independent, normally distributed block differences with a common variance. Multi-arm trials, clustered treatments, repeated seasons and multi-farm inference need a calculation for that actual design.

FIELDWORK METHODS / 02B

A small difference can require a large experiment.



All curves: difference 2 bu/ac, two-sided alpha 0.05, paired t test. Dots show the smallest integer count reaching at least 80% power. Plot count is twice the block count.

SD of block differences	Difference 1 bu/ac	Difference 2 bu/ac	Difference 3 bu/ac
2.5 bu/ac	52 / 68	15 / 19	8 / 10
5 bu/ac	199 / 265	52 / 68	24 / 32
10 bu/ac	787 / 1,053	199 / 265	90 / 119

Each cell gives paired blocks at 80% / 90% power. Multiply by two for total treatment plots. These are complete-block requirements under the stated model, before losses.

Worked planning case / 2 bu/ac, SD(D) = 5 bu/ac

52 complete blocks (104 plots) attain 80.78% power; 51 blocks attain 79.99%. For 90% power, 68 complete blocks (136 plots) attain 90.16%. The normal approximation gives 50 blocks at 80%; the exact small-sample t calculation requires 52.

Why the old yield map cannot settle the count

The 2018 historical 20 m cell model has residual SD 10.5541 bu/ac. If we assume it equals planned plot SD, and additionally assume within-block correlation 0.5, SD(D) equals 10.5541. That illustrative case needs 221 / 295 blocks at 80% / 90% power for 2 bu/ac. Neither assumption was verified, so this is sensitivity analysis, not the recommended design.

FIELDWORK METHODS / 02C

Validate the calculation. Then qualify the field design.

Exact noncentral-t power [4] was checked against 20,000 simulated sets of normal block differences per scenario, with fixed seed 20261004. All differences were within four Monte Carlo standard errors. This checks calculation behavior under the model; it does not validate the model for a particular field.

Blocks	Effect	Exact rejection probability	Simulated probability
4	2 bu/ac	8.92%	8.75%
12	2 bu/ac	24.46%	24.61%
52	2 bu/ac	80.78%	80.87%
68	2 bu/ac	90.16%	90.35%
52	0 bu/ac	5.00%	4.88%

Blocking sensitivity / same assumed plot SD of 5 bu/ac

Assumed within-block correlation	SD(D), bu/ac	Blocks for 2 bu/ac / 80%
0	7.07	101
0.25	6.12	76
0.5	5.00	52
0.75	3.54	27

A strong pairing structure can reduce noise, but neighboring blocks are not automatically independent. Confirm that treatment assignment, spatial layout and the eventual analysis support the assumed variance reduction.

Allow for loss; plan the protocol before planting

If 52 complete pairs are required and 10% pair loss is assumed, schedule $\text{ceil}(52 / 0.90) = \mathbf{58 \text{ paired blocks / 116 plots}}$. This provides expected retention, not a guarantee of 52 complete pairs. Losing either treatment plot breaks a complete pair; handling missing plots should be pre-specified.

Before selecting a final count, collect: the minimum useful yield effect; pilot plot means and treatment/block labels; plot size and implement width; usable field area and buffers; harvest calibration and sensor-delay information; expected losses; and the population of farms/seasons to which the result should apply.

FIELDWORK METHODS / APPENDIX

Methods, provenance and reproducibility.

Data and execution

OPPEDATA, Montana State University (2020), MIT license. Frozen revision ae70c5d8ad4fe154fd4b746a7ec43f3031a6e0a9. Ten original source-file hashes were verified against the pinned manifest. The machinery audit independently reloads the raw shapes, recomputes coordinate transforms, flags and matches, and asserts the counts agree with the original demo pipeline. [1]

The project package contains the frozen original source archive, the two analysis scripts, generated findings, audit and analysis CSVs, GeoJSON, charts, a browser demonstration and this PDF. README.md documents offline reproduction. Original source files are preserved byte-for-byte. Neither raw yield values nor source positions were replaced.

Exact power computation

For r blocks, degrees of freedom = $r - 1$; noncentrality = $\sqrt{r} \times \text{delta} / \text{SD}(D)$. Let c be the central- t 97.5th percentile. Power is $P(T < -c) + P(T > c)$, for the corresponding noncentral- t distribution. Search integer r from 2 upward until the target is reached. The CSV includes achieved power and the result for one fewer block. [4]

References

[1] OPPEDATA / frozen machinery source

<https://github.com/paulhegedus/OPPEDATA/tree/ae70c5d8ad4fe154fd4b746a7ec43f3031a6e0a9>

[2] pyproj / Transformer API and always_xy axis order

<https://pyproj4.github.io/pyproj/stable/api/transformer.html>

[3] University of Minnesota Extension / How to do research on your farm

<https://extension.umn.edu/courses-and-events/for-professionals/on-farm-research/how-to-do-research-on-your-farm>

[4] SciPy / Noncentral Student t distribution

<https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.nct.html>

Interpretation limits

The quality rules and spatial joins demonstrate auditable processing; they do not establish causal treatment effects, accurate timestamps, calibrated nitrogen doses or commercial yield gains. Historical match coverage is not experimental randomization. Final replicate counts require the intended treatment design and justified plot-scale variance. No client trial or independently verified physical GNSS correction is claimed.

FIELDWORK / BRAND & SERVICE CONCEPT

Live demonstration: fieldwork.geointeladvantage.com. The public-data work is demonstrated capability; team qualifications and commercial contact arrangements remain to be established.