Satellite-Guided Field Scouting and a Nitrogen × Water Follow-Up Design for Corn and Soybean in Central Iowa

A Reproducible Sentinel-2 / USDA Cropland Data Layer Workflow for Ranking Crop-Specific Vegetation-Index Anomalies

Technical Report — July 14, 2026

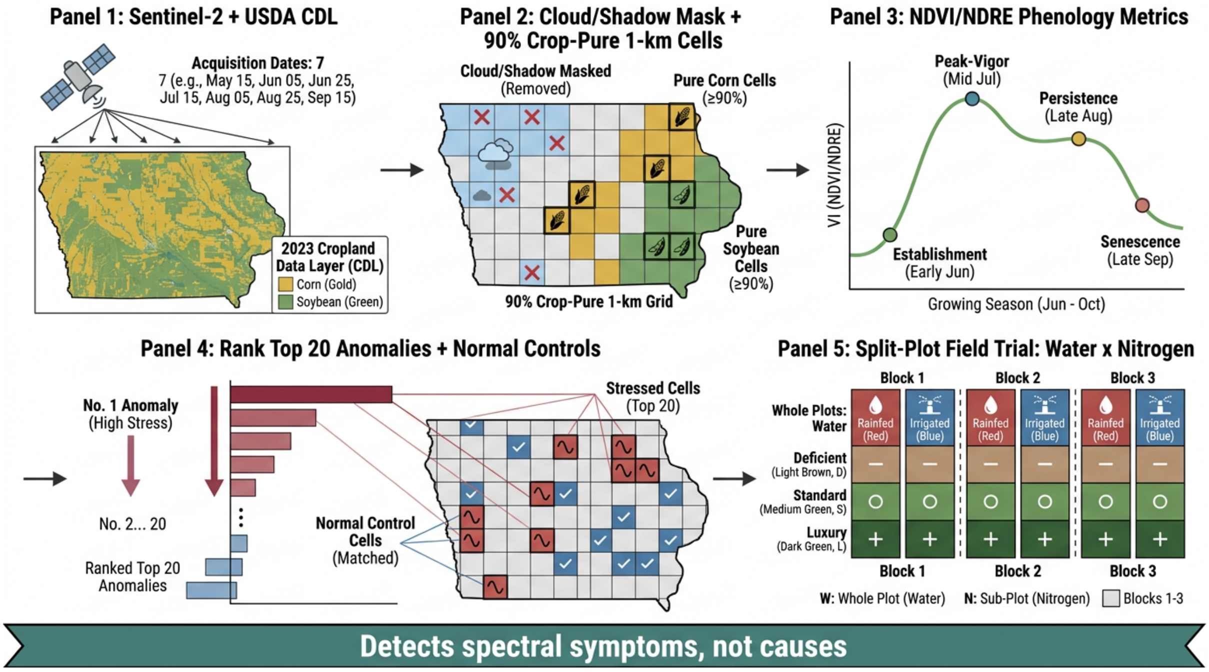

Figure 1: Graphical abstract. End-to-end workflow. Seven cloud-screened Sentinel-2 Level-2A acquisitions over tile T15TVG are combined with the 2023 USDA Cropland Data Layer (CDL) to define 90 %-crop-pure 1 km cells; NDVI/NDRE seasonal profiles yield Establishment, Peak-Vigor, Persistence and Senescence metrics; crop-specific z-score anomalies rank the twenty highest-priority scouting cells (10 corn, 10 soybean), each paired with a nearby "normal" control; and a randomized split-plot (Water × Nitrogen) trial is laid out for confirmatory scouting. The workflow detects spectral *symptoms*, not their *causes*.

Abstract

Optical satellite time series can flag under-performing cropland far faster and more cheaply than field walking, but a low vegetation index is a symptom rather than a diagnosis. We present a fully programmatic, auditable workflow that converts seven 2023 Sentinel-2 Collection 1 Level-2A acquisitions of tile `T15TVG` (2023-04-06, 05-24, 06-20, 07-10, 08-22, 09-08, 10-21) and the 2023 USDA Cropland Data Layer into a prioritized, uncertainty-flagged scouting list for a $0.60° \times 0.37°$ region of interest (ROI) in central Iowa ($[-93.75, 41.85, -93.15, 42.22]$). After Scene-Classification-Layer (SCL) cloud and shadow masking, we retained pixels within 1 km grid cells that were $\geq 90\,\%$ single-crop in the CDL, yielding 40 corn-pure and 16 soybean-pure cells (522,739 native 10 m target pixels). For every cell we derived NDVI and NDRE and four phenological descriptors—Establishment, Peak-Vigor, Persistence and Senescence—and computed crop-specific (z-score) anomalies against the corn and soybean populations separately. Cells were ranked on a robust primary Stress Index built only from the four full-coverage Peak-Vigor and Persistence metrics (so a mid-season orbit-swath gap did not bias the ranking), and the twenty most anomalous cells were each matched to a spatially proximate, spectrally "normal" same-crop control (preferred 2–10 km) using a global minimum-distance assignment. Each site carries transparent uncertainty flags for sub-pixel purity, within-cell heterogeneity and clear-observation fraction. To move from correlation toward causation without over-claiming, we specify a randomized split-plot experiment—Water as the whole-plot factor (Rainfed/Irrigated) and Nitrogen as the sub-plot factor (Deficient/Standard/Luxury), three replicate blocks per location, 360 plots in total—and provide its full analysis-of-variance structure and correct error strata. We emphasize throughout that the imagery localizes where to look, whereas only ground-truthed agronomic sampling and a properly analysed factorial trial can establish why a field is stressed. All code, geospatial layers, figures and a machine-readable manifest accompany the report.

Contents

1 Introduction

The economic and environmental margins of maize (*Zea mays* L.) and soybean (*Glycine max* (L.) Merr.) production in the U.S. Corn Belt are set in the field, one management zone at a time. Yet the classical instrument for finding trouble—an agronomist walking transects—scales poorly across the tens of thousands of hectares a modern operation or crop-insurance portfolio must cover within the few days that separate a treatable problem from an unrecoverable loss. Earth-observation satellites invert this constraint. Since the launch of the Sentinel-2 twin-satellite constellation, decametric multispectral imagery with a five-day revisit has become freely available, and a quarter-century of research has established that canopy reflectance encodes agronomically meaningful information about biomass, leaf area, chlorophyll and, indirectly, crop stress [11, 35, 45, 49]. The specific advantages Sentinel-2 brings to precision agriculture—its red-edge bands, decametric resolution and frequent revisit—have been reviewed extensively [44], and satellite monitoring is increasingly complemented by, rather than replaced by, unmanned-aircraft platforms that trade coverage for centimetric detail [22]. The operational question is therefore no longer whether spaceborne sensors *see* crop condition, but how to translate what they see into a defensible, prioritized, and appropriately hedged plan of action.

Two properties of the reflectance signal make this translation non-trivial. First, the workhorse Normalized Difference Vegetation Index (NDVI), introduced in the earliest Landsat-era vegetation studies [39, 46] and still ubiquitous among the many indices catalogued in the vegetation-index literature [51], saturates once a canopy closes, losing sensitivity precisely during the reproductive stages when yield is being determined [21, 48]. Red-edge indices such as the Normalized Difference Red-Edge index (NDRE) exploit the steep 700–740 nm reflectance transition that Sentinel-2 samples with dedicated bands, remain sensitive at high leaf-area and chlorophyll levels, and are more tightly coupled to canopy nitrogen than NDVI [10, 16, 17, 36, 43]; the same Sentinel-2 red-edge bands support quantitative retrieval of leaf-area index and leaf and canopy chlorophyll content [9]. Using both indices in tandem—NDVI for coarse vigor and NDRE for chlorophyll- and nitrogen-sensitive structure—provides a more complete and partially independent read on canopy status than either alone. Second, and more fundamentally, the reflectance signal is *mechanistically agnostic*: a depressed index is consistent with nitrogen deficiency, water stress, soil compaction, disease, insect or weed pressure, hail, or a simple emergence gap. This ambiguity is the central epistemic limit of index-based monitoring and the reason every credible workflow must separate detection (a spectral symptom) from diagnosis (a cause) [32, 53]. The present report is organized around exactly that separation.

Turning single scenes into crop condition requires two preprocessing steps that are easy to state and easy to get wrong. Clouds and their shadows must be removed, because a single contaminated observation can masquerade as catastrophic senescence; the Sentinel-2 Level-2A product supplies a Scene Classification Layer (SCL) whose cloud, cirrus and cloud-shadow classes support per-pixel screening, complemented in the wider literature by algorithms such as `Fmask` and `Sen2Cor` [12, 14, 33]. Cloud cover is not a nuisance to be tolerated but a first-order sampling problem: across the growing season, persistent cloudiness systematically depletes the usable optical record and biases any naive temporal statistic [50]. Equally, the analysis must know *which* pixels are the crop of interest. The USDA National Agricultural Statistics Service Cropland Data Layer (CDL) provides an annual, crop-specific, 30 m classification of the conterminous United States with producer and user accuracies typically in the 85–95 % range for major commodities such as corn and soybean, and has become the de-facto ground reference for agricultural remote sensing in the country [6, 28–30]. Restricting analysis to spatially "pure" crop cells—here, 1 km cells that are at least 90 % a single crop—suppresses mixed-pixel contamination and makes the crop-specific comparison of like with like meaningful.

Once clean, crop-attributed time series are in hand, phenology becomes the organizing principle. Land-surface phenology methods condense a season of observations into a small number of interpretable descriptors—green-up, peak, and senescence timing and magnitude—that are more robust and more transferable than any single date [5, 41, 54]. In an agronomic setting these descriptors double as stress indicators: a crop that establishes late, never reaches the vigor of its peers, cannot sustain a green canopy through grain fill, or senesces prematurely is, by definition, spectrally atypical. Detecting such atypical behaviour is a form of anomaly detection, a mature subfield of agricultural monitoring that ranges from operational near-real-time NDVI anomaly systems to regression-based yield forecasters, machine-learning classifiers and yield-anomaly mappers [1, 4, 27, 34, 38]. The scalable, cloud-hosted processing that makes continental application feasible—exemplified by Google Earth Engine and by harmonized Landsat–Sentinel products—has removed the last practical barrier to running these analyses at management-zone resolution over entire counties [2, 8, 15, 20, 26, 31, 40, 55].

The value of a scouting product is realized only when it changes what happens on the ground, and this is where remote sensing must hand off to experimental agronomy. If the imagery hypothesizes that a cluster of low-NDRE corn cells reflects a nitrogen or water limitation, that hypothesis is testable—but only with a designed experiment that manipulates the putative drivers and analyses the response with a model that respects how the treatments were randomized. Water and nitrogen are the two most consequential and most manageable drivers of Corn Belt canopy vigor, they interact strongly, and their manipulation has very different logistics: irrigation is difficult to randomize at small plot scale whereas fertilizer is trivially applied per plot [25, 42]. This asymmetry is the textbook justification for a split-plot design, in which the hard-to-randomize factor occupies whole plots and the easy factor is randomized within them, with the crucial consequence that the two factors are tested against different error strata [7, 13, 19, 37, 52]. We therefore close the loop by delivering not just a ranked map but a ready-to-stake split-plot trial and its correct analysis-of-variance skeleton.

This report documents the complete pipeline for the central-Iowa ROI and tile `T15TVG`. Section 2 details data acquisition, SCL masking, 90 % purity filtering, index and phenology-metric computation, crop-specific anomaly ranking, and control pairing. Section 3 presents the crop-pure cell inventory, the crop-specific seasonal profiles, the ranked top-twenty scouting list with paired controls and uncertainty flags, and the spatial and phenological visualizations. Section 4 specifies the split-plot follow-up experiment and its full ANOVA. Section 5 interprets the results and is deliberately weighted toward caveats and limitations, restating the load-bearing caution that the imagery identifies *where* to scout, never *why* a field is stressed.

2 Data and Methods

2.1 Region of interest, coordinate system, and software

The ROI is the geographic bounding box [$-93.75°$, $41.85°$, $-93.15°$, $42.22°$] (WGS84), a $\sim$50 km × 41 km block of intensively cropped land north of Des Moines, Iowa, falling wholly within Sentinel-2 tile `T15TVG`. All raster processing was carried out in the tile's native projected coordinate reference system, UTM zone 15N (EPSG:32615), so that distances, areas and the 1 km grid are metric and undistorted; geographic coordinates (EPSG:4326) are reported for field navigation. The pipeline is implemented in Python with `rasterio`, `numpy`, `geopandas/shapely`, `scipy` and `matplotlib`, and every stage writes a JSON log and versioned geospatial outputs (GeoPackage, GeoJSON, CSV and GeoTIFF), so the analysis is fully reproducible and auditable.

2.2 Sentinel-2 acquisition and cloud/shadow masking

Seven Sentinel-2 Collection 1 Level-2A (bottom-of-atmosphere surface reflectance) acquisitions of tile `T15TVG` were obtained for the requested dates spanning the 2023 season: 2023-04-06 (pre-emergence bare soil), 2023-05-24 (establishment), 2023-06-20 and 2023-07-10 (vegetative build-up), 2023-08-22 (peak/early reproductive), 2023-09-08 (grain fill) and 2023-10-21 (senescence/harvest). Bands B04 (red, 10 m), B05 (red-edge, 20 m resampled to 10 m), B08 (near-infrared, 10 m) and the 20 m Scene Classification Layer were read for each date. Following the Level-2A product design and established practice [12, 14, 33], all pixels whose SCL class denoted defective data, saturation, dark-area/cloud-shadow, medium- or high-probability cloud, thin cirrus, or snow/ice were masked to no-data before any index was computed; only the vegetation, bare-soil, water and unclassified SCL classes were passed through, and downstream crop pixels are additionally constrained to cropland by the CDL (Section 2.3).

Scene-level cloud cover was low on most dates (five of seven scenes below 4 % cloud), so masking removed little valid signal on the growing-season core dates. The consequential exception is orbital: the 2023-05-24, 2023-08-22 and 2023-10-21 scenes were acquired on relative orbit R112, whose swath covers only ∼64 % of the ROI, whereas the remaining dates lie on the fully covering R069 orbit. Cells falling entirely inside the R112 gap therefore lack valid observations on the establishment date (2023-05-24) and, for the late pair, on the senescence date. Rather than impute these values we recorded them as missing, which—as detailed below—is why the primary ranking index deliberately excludes the affected metrics.

2.3 USDA Cropland Data Layer and 90 % crop-purity filtering

The 2023 USDA/NASS CDL [47] was ingested for the ROI, reprojected from its native CONUS Albers grid (EPSG:5070) to UTM 15N, and used both to identify the crop of each pixel and to enforce spatial purity [6, 28]. Within the ROI the CDL is overwhelmingly agricultural: corn (class 1) and soybean (class 5) account for 942,349 and 751,419 of the 30 m pixels respectively, the two classes together comprising roughly three-quarters of the landscape. A regular 1 km grid ($41 \times 49 = 2009$ cells) was overlaid on the ROI in UTM 15N. For each cell we computed the fraction of CDL pixels belonging to each crop and retained a cell as “crop-pure” only if a single crop occupied at least 90 % of its area. This threshold, chosen to sit comfortably above the mixed-pixel and edge-misclassification regime that dominates CDL error [29, 30], yielded 40 corn-pure and 16 soybean-pure cells (56 in total). A 10 m crop mask and a per-cell identifier raster were rasterized from these pure cells, isolating 522,739 native Sentinel-2 target pixels (374,098 corn; 148,641 soybean) for spectral analysis.

2.4 Vegetation indices

For every retained target pixel on every clear date we computed

$$NDVI = \frac{\rho_{\mathrm{B08}} - \rho_{\mathrm{B04}}}{\rho_{\mathrm{B08}} + \rho_{\mathrm{B04}}}, \qquad NDRE = \frac{\rho_{\mathrm{B08}} - \rho_{\mathrm{B05}}}{\rho_{\mathrm{B08}} + \rho_{\mathrm{B05}}}, \tag{1}$$

where ρ denotes surface reflectance. NDVI tracks green biomass and fractional cover but saturates in dense canopies, whereas NDRE, by substituting the red-edge band B05 for the red band, retains sensitivity to canopy chlorophyll and nitrogen through canopy closure [3, 10, 18, 43, 48]. Reporting both provides complementary, partially independent views of canopy status; in our data the two derived stress indices are strongly but not perfectly correlated (Pearson $r = 0.83$), confirming that NDRE adds information rather than duplicating NDVI.

2.5 Phenological metrics

From the seven-date index trajectory of each pixel we distilled four agronomically interpretable descriptors for each of NDVI and NDRE, following the land-surface-phenology tradition of summarizing a season by a few robust key points [5, 41, 54]:

- **Establishment** — the index value on 2023-05-24, capturing early-season canopy development;
- **Peak-Vigor** — the maximum index value over all seven dates, the seasonal apex of canopy greenness/chlorophyll;
- **Persistence** — the mean index over the five growing-season dates (2023-05-24 through 2023-09-08), a measure of sustained canopy function through grain fill;
- **Senescence** — the decline from 2023-09-08 to 2023-10-21, quantifying the rate of end-of-season canopy dieback.

Metrics were computed per pixel and then aggregated to the cell level as the mean over that cell's target pixels; the within-cell standard deviation of Peak-Vigor was retained as a heterogeneity indicator (Section 2.8). Because Establishment depends on the 2023-05-24 R112 scene and Senescence on the 2023-10-21 R112 scene, these two metrics are undefined for the 20 (Establishment) and 19 (Senescence) cells lying inside the swath gap; they were stored as missing rather than imputed.

2.6 Crop-specific anomaly detection and ranking

Because corn and soybean have fundamentally different phenology—soybean greens up later and senesces differently, as Section 3.2 shows—anomalies were defined *within* each crop, never across the pooled population. For each metric m and crop group $g \in \{\text{corn}, \text{soybean}\}$ we standardized cell values to crop-specific z-scores,

$$z_{i,m} = \frac{x_{i,m} - \mu_{g,m}}{\sigma_{g,m}}, \tag{2}$$

using the sample mean and standard deviation (`ddof`$= 1$) of that crop's cells and a NaN-aware estimator so that swath-gap missingness reduced the effective sample size but did not corrupt the statistics. A negative z denotes lower-than-typical vigor for that crop.

We then defined a *primary Stress Index* as the mean of only the four full-coverage z-scores—NDVI and NDRE Peak-Vigor and Persistence—so that the ranking is computed identically for every cell and is unaffected by the R112 swath-gap missingness in Establishment and Senescence. A more inclusive *all-metric Stress Index* (all eight z-scores, with the sign of Senescence flipped so that faster decline counts as more stress) was computed for cross-checking; the two indices agree closely (Spearman $\rho = 0.83$). Cells were ranked in ascending order of the primary index (most negative = most anomalous), and the ten most anomalous corn cells and ten most anomalous soybean cells were selected as the top-twenty scouting targets. Confining selection to the primary index guarantees that a cell is never promoted or demoted because of a missing observation.

2.7 Nearby normal-control pairing

A stressed cell is only interpretable against a baseline. For every anomalous cell we sought a spatially nearby, same-crop cell whose phenology is representative rather than extreme—operationally, a cell whose crop-specific primary Stress Index lies within the interquartile range (25th–75th percentile) of the non-anomalous cells of that crop. Candidate controls were preferentially drawn from a 2–10 km annulus around each anomaly (close enough to share soils, weather and management context, far enough to be an independent field), and a single globally optimal assignment was solved with the Hungarian (minimum-total-distance) algorithm so that every anomaly received a *unique* control. Where the preferred band could not be satisfied—most acutely for soybean, which has only 16 pure cells, ten of them flagged as anomalous, leaving a very small non-anomalous interquartile pool—the matcher expanded the search radius (to 15 km, then beyond) and, as a last resort, fell back to the least-stressed remaining same-crop cell, recording that the control was itself flagged. Every pair carries an explicit distance and a distance-band flag so the user can judge how comparable a control truly is.

2.8 Uncertainty flags

Each cell is annotated with three transparent, thresholded uncertainty indicators, and an overall flag equal to the worst of the three:

- **Sub-pixel purity** — the CDL crop-purity percentage of the cell (High uncertainty if $<92\,\%$, Medium if $<95\,\%$, Low if $\geq 95\,\%$): lower purity means more non-target vegetation may contaminate the aggregate index.
- **Within-cell heterogeneity** — the standard deviation of pixel Peak-Vigor across the cell (High above the 75th percentile, Medium above the median): high heterogeneity signals mixed soils, drainage classes or management zones so that a single centroid-anchored footprint may not sit on representative ground.
- **Clear-observation fraction** — the fraction of the five growing-season dates with valid (unmasked, in-swath) observations (High if <0.80, Medium if <0.90): fewer clear looks weaken the phenological summary. Cells inside the R112 gap score 0.60 here and are flagged accordingly.

These flags do not alter the ranking; they tell the scout how much to trust each entry and where ground-truthing is most needed.

3 Results

3.1 Crop-pure cell inventory

Figure 2 maps the 90 %-crop-pure 1 km cells over the 2023 CDL. The 40 corn-pure and 16 soybean-pure cells are distributed across the ROI rather than clustered, spanning a representative range of soils and drainage settings; the pronounced numerical asymmetry (corn:soybean $\approx$ 2.5:1) reflects both the 2023 corn-heavy rotation in this part of Iowa and the fact that soybean fields in the ROI are more often interspersed with other cover, making 90 %-pure 1 km soybean cells comparatively scarce. This scarcity is consequential for soybean control pairing (Section 3.3) and is carried forward explicitly as a limitation.

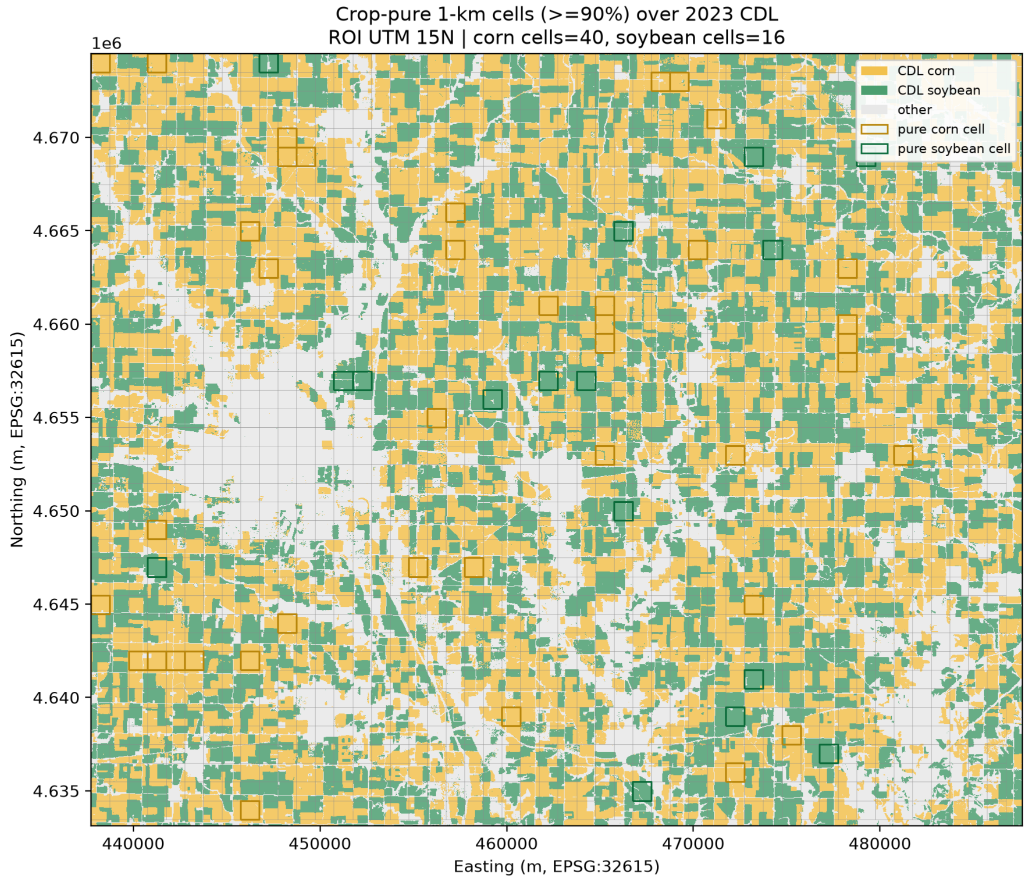

Figure 2: Crop-pure 1 km cells over the 2023 USDA CDL. Background: CDL corn (gold) and soybean (green); other classes in grey. Squares mark the 40 corn-pure and 16 soybean-pure 1 km cells ($\geq$90 % single crop) retained for spectral analysis. Coordinates in UTM 15N (EPSG:32615).

3.2 Crop-specific seasonal profiles

Figure 3 shows the mean NDVI and NDRE trajectories ($\pm$1 SD) for the two crops over all target pixels. The profiles are textbook and confirm that crop-specific standardization is essential. Corn (gold) greens up earlier and peaks around late June–July ($NDVI_{\text{peak}} \approx 0.63$, $NDRE_{\text{peak}} \approx 0.56$), whereas soybean (green) develops later, peaks in August ($NDVI_{\text{peak}} \approx 0.62$, $NDRE_{\text{peak}} \approx 0.50$) and retains a greener canopy into September. The two crops' senescence limbs differ markedly—soybean shows a larger late-season NDVI drop—so a single pooled baseline would mislabel normal soybean behaviour as anomalous relative to corn, and vice versa. Consistent with the red-edge literature, the NDRE amplitude is lower than NDVI but preserves clear crop separation through canopy closure, supporting its inclusion as a chlorophyll/nitrogen-sensitive complement [10, 36, 43, 48]. Cell-level means underlying the anomaly analysis are summarized in Table 1.

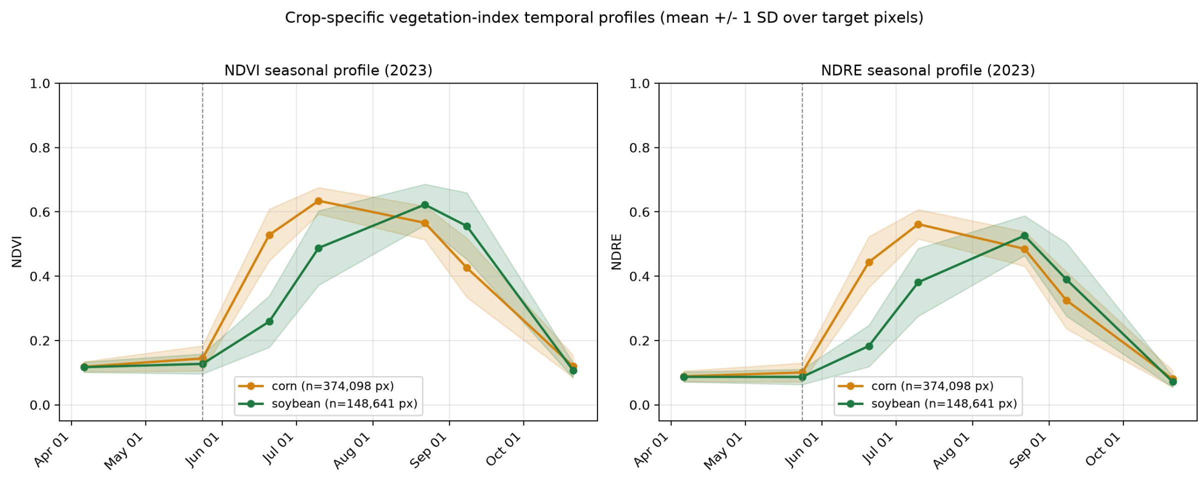

Figure 3: Crop-specific vegetation-index temporal profiles (mean ±1 SD over target pixels). Left: NDVI; right: NDRE. Corn (n = 374,098 px) peaks earlier than soybean (n = 148,641 px), motivating within-crop anomaly standardization. The vertical dashed line marks the nominal start of the growing-season window.

Table 1: Crop-level mean of the cell-aggregated phenological metrics (dimensionless index units). These crop-specific means and their standard deviations define the z-score reference distributions used for anomaly detection.

	NDVI				NDRE			
Crop	Estab.	Peak	Persist.	Senesc.	Estab.	Peak	Persist.	Senesc.
Corn	0.144	0.635	0.480	0.307	0.101	0.562	0.401	0.244
Soybean	0.127	0.615	0.422	0.448	0.087	0.495	0.315	0.323

3.3 Ranked scouting list and paired controls

Figure 4 ranks the twenty selected cells by primary Stress Index and shows their spatial distribution. The corn tail is deeper and more negative than the soybean tail: corn priorities span primary indices from −2.74 to −0.50, whereas soybean spans −1.48 to +0.08. This is expected—with 40 corn cells the crop-specific distribution has heavier tails and the ten worst corn cells are genuinely extreme, while with only 16 soybean cells the ten “worst” necessarily reach into the middle of the soybean distribution (three soybean priorities have near-zero or slightly positive indices and are, in absolute terms, only mildly atypical). The most anomalous cell overall, corn cell 33 (primary index −2.74), combines a strongly depressed Peak-Vigor and Persistence in both indices; its establishment and senescence z-scores are missing because it lies in the R112 swath gap, which is exactly why it is ranked on the full-coverage primary index alone.

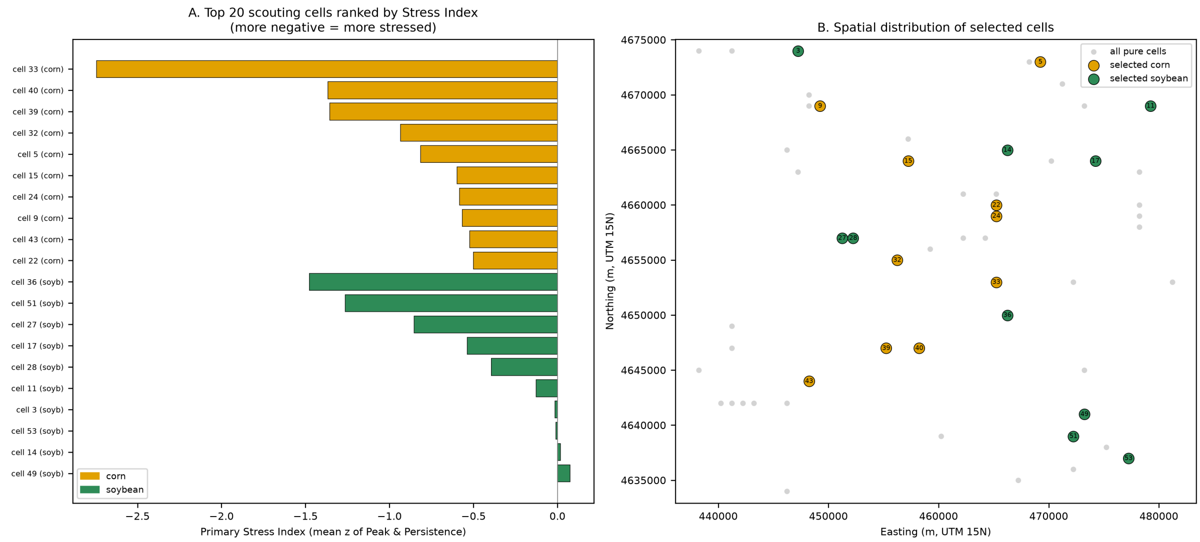

Figure 4: Top-20 scouting cells. (A) Cells ranked by primary Stress Index (mean z of NDVI/NDRE Peak-Vigor and Persistence; more negative = more stressed); corn in gold, soybean in green. (B) Spatial distribution of the selected corn (gold) and soybean (green) cells among all pure cells (grey) in UTM 15N.

Table 2 is the operational deliverable: for each scouting priority it lists the crop, WGS84 centroid, primary Stress Index, the three uncertainty inputs and the overall flag, and the assigned normal control with its separation distance and band flag. Eighteen of the twenty controls fall in the preferred 2–10 km band; corn priority 3 (cell 39) required a slightly expanded 10.3 km match, and soybean priority 17 (cell 3), a genuine spatial outlier whose nearest same-crop pure cell is ∼17.5 km away, was matched at 27.7 km and is flagged `beyond_15km`. Four soybean controls are themselves flagged (a direct consequence of the small soybean pool discussed in Section 2.7), meaning those comparisons are weaker and should be interpreted cautiously. Seventeen of twenty anomalies carry an overall “High” uncertainty flag, driven predominantly by sub-92 % purity and, for swath-gap cells, the 0.60 clear-observation fraction—a candid signal that these are scouting leads, not conclusions.

Table 2: Ranked top-twenty scouting cells (10 corn, 10 soybean) with paired normal controls and uncertainty flags. **SI$_{\text{prim}}$**: primary Stress Index (more negative = more anomalous). **Pur**: CDL crop purity (%). **Het**: within-cell Peak-Vigor standard deviation. **Clear**: growing-season clear-observation fraction. **Unc**: overall uncertainty (H/M). **Ctrl**: control cell id; **Dist**: anomaly–control distance (km); † outside preferred 2–10 km band; ∗ control is itself flagged.

#	Crop	Lat	Lon	SI$_{\text{prim}}$	Pur	Het	Clear	Unc	Ctrl	Dist
1	corn	42.0281	−93.4199	−2.744	92.9	0.086	0.66	H	34	7.00
2	corn	41.9737	−93.5040	−1.367	99.2	0.043	1.00	M	50	8.25
3	corn	41.9735	−93.5402	−1.356	94.1	0.051	1.00	H	48	10.30†
4	corn	42.0457	−93.5288	−0.934	90.8	0.025	1.00	H	20	8.49
5	corn	42.2084	−93.3726	−0.816	97.0	0.041	0.99	M	6	2.83
6	corn	42.1268	−93.5173	−0.598	92.2	0.033	1.00	M	12	2.00
7	corn	42.0821	−93.4203	−0.583	91.9	0.041	0.98	H	21	2.00
8	corn	42.1713	−93.6146	−0.567	90.3	0.052	1.00	H	13	5.00
9	corn	41.9461	−93.6245	−0.522	91.8	0.028	1.00	H	47	5.39
10	corn	42.0911	−93.4203	−0.501	91.2	0.033	0.99	H	16	6.40
11	soybean	42.0011	−93.4076	−1.477	94.5	0.075	0.60	H	30	7.28
12	soybean	41.9023	−93.3347	−1.263	93.2	0.050	0.60	H	55	6.40
13	soybean	42.0634	−93.5893	−0.854	93.0	0.063	1.00	H	31	8.06
14	soybean	42.1275	−93.3117	−0.537	92.2	0.057	0.60	H	14	8.06∗

continued on next page

Table 2 – continued

#	Crop	Lat	Lon	SI_{prim}	Pur	Het	Clear	Unc	Ctrl	Dist
15	soybean	42.0634	−93.5773	−0.394	92.0	0.061	1.00	H	29	10.00
16	soybean	42.1727	−93.2514	−0.127	96.0	0.053	0.60	H	10	6.00
17	soybean	42.2162	−93.6392	−0.015	91.3	0.042	1.00	H	38	27.66†
18	soybean	41.8844	−93.2743	−0.009	91.7	0.066	0.60	H	49	5.66∗
19	soybean	42.1362	−93.4085	+0.019	90.5	0.046	1.00	H	17	8.06∗
20	soybean	41.9203	−93.3227	+0.075	90.3	0.054	0.60	H	53	5.66∗

Figure 5 places the anomaly–control pairs in space, and Figure 6 contrasts the seasonal NDVI and NDRE trajectories of four representative anomaly–control pairs. The paired profiles make the anomaly concrete: for corn cell 33 the stressed trajectory sits well below its control (cell 34) through the entire vegetative and reproductive window in both indices, while for the milder soybean pairs the separation is smaller and confined to the peak, visually reinforcing the quantitative message that the soybean “anomalies” are less extreme than the corn ones. The gaps in the swath-affected trajectories are the missing R112 observations, shown honestly rather than interpolated.

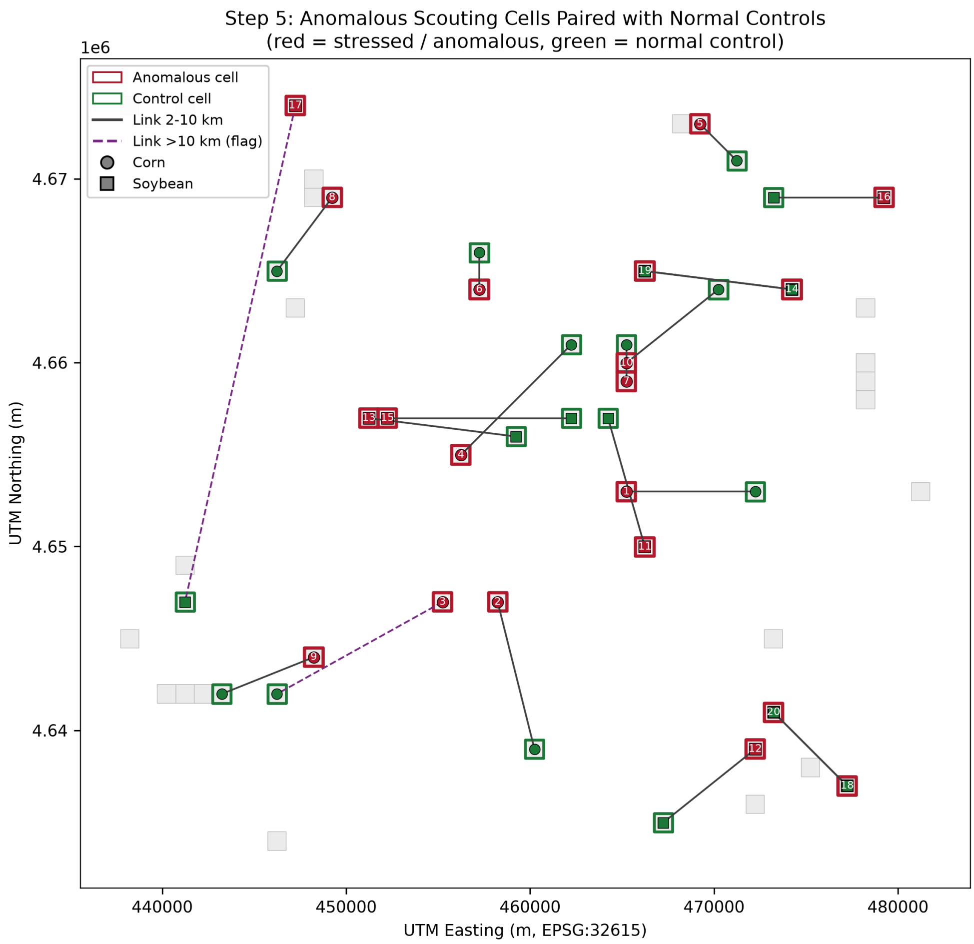

Figure 5: Anomalous cells paired with normal controls. Red squares: anomalous (stressed) cells; green squares: assigned normal controls. Solid links are within the preferred 2–10 km band; dashed links exceed 10 km. Circles denote corn, squares soybean. Note the long dashed link for the soybean spatial outlier (priority 17). UTM 15N.

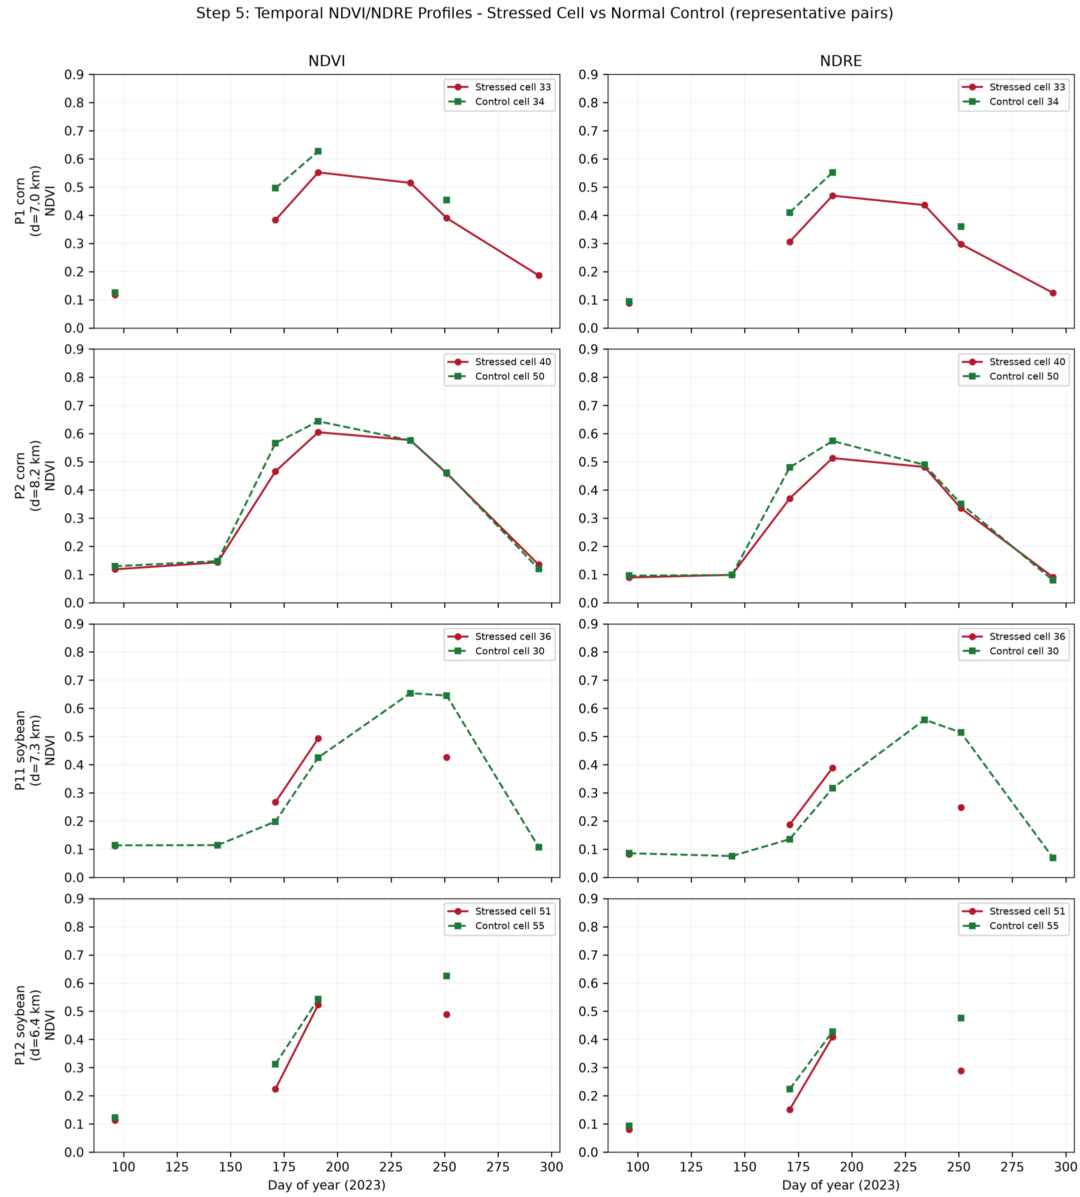

Figure 6: Stressed vs. normal-control seasonal profiles for four representative pairs (two corn, two soybean). Red: anomalous cell; green dashed: paired control. Left column NDVI, right column NDRE. Missing points on swath-gap dates are omitted rather than imputed.

4 Nitrogen × Water Follow-Up Experimental Design

4.1 Rationale and factor assignment

The scouting list localizes candidate stress but cannot attribute it. To test the two most plausible and most manageable agronomic drivers, we specify a two-factor field experiment crossing **Water management** (Rainfed vs. Irrigated) with **Nitrogen fertilization** (Deficient, Standard, Luxury) at each of the twenty flagged locations. The two factors are not logistically symmetric: supplemental irrigation is difficult to confine and randomize at small plot scale, whereas nitrogen rate is trivially applied per plot. This asymmetry is the classical trigger for a *split-plot* design [13, 19, 52]: Water is assigned to whole plots (main plots) and Nitrogen is randomized to sub-plots nested within each whole plot. The design choice is not cosmetic—it dictates that the two factors are estimated with different precision and, critically, tested against different error terms [7, 37].

4.2 Layout and randomization

At each location the experiment comprises three replicate blocks; within each block the two Water levels are randomly assigned to two main-plot strips, and within each strip the three Nitrogen levels are independently randomized across three sub-plots. This gives $2 \times 3 = 6$ treatment combinations per block, 18 plots per location, and $18 \times 20 = 360$ plots in total. Sub-plots measure 12×24 m (288 m^2), with 1.5 m sub-plot alleys, 3 m main-plot alleys and 6 m block alleys, producing a compact $\sim$129 m $\times$ 51 m footprint that is centred on each 1 km cell centroid and reprojected from UTM 15N to WGS84 for staking. Randomization used a fixed seed (`numpy` seed 42) so the layout is exactly reproducible; validation confirmed a fully balanced design (360/360 plots, unique plot identifiers, exactly one of each of the six treatment combinations per block, and—by design of the independent per-block randomization—no systematic placement, with 33 of 60 blocks placing Rainfed in the south row). Figure 7 shows the conceptual split-plot hierarchy alongside four actual randomized location layouts, which differ from one another precisely because randomization is independent per site.

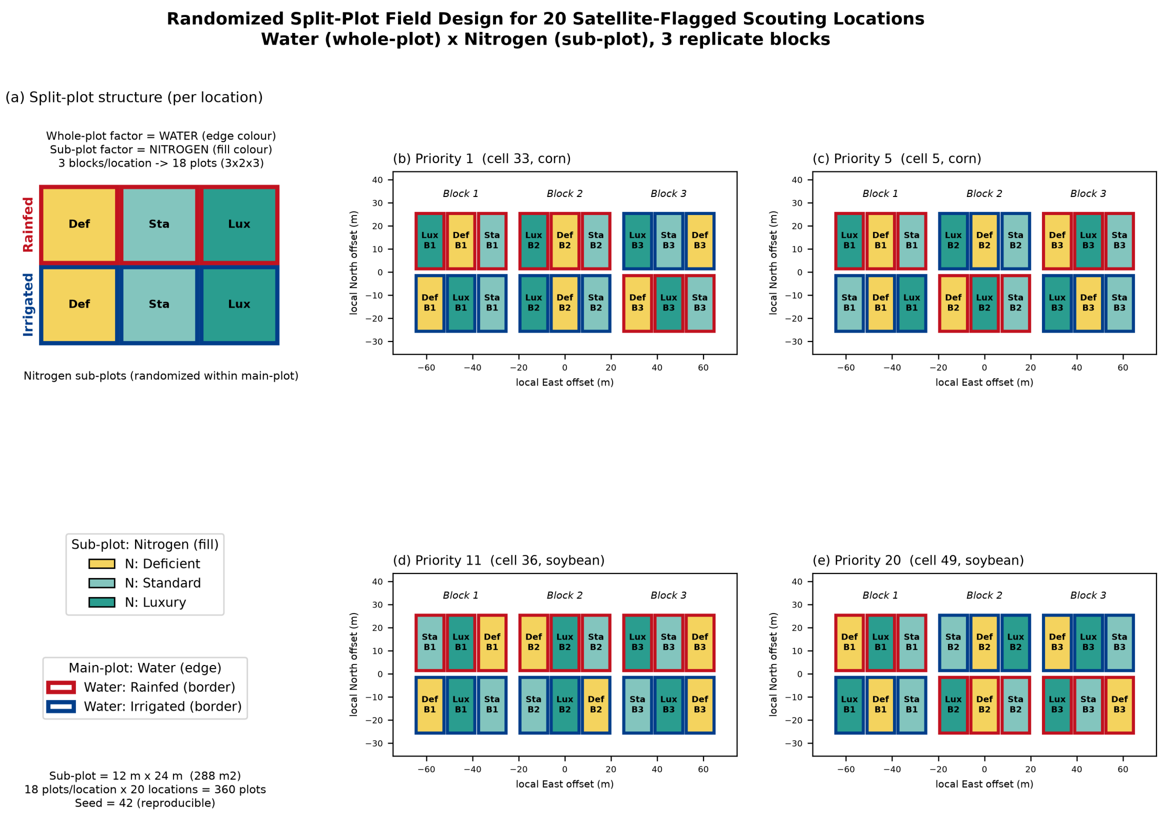

Figure 7: Randomized split-plot field design. (a) Split-plot structure per location: Water is the whole-plot factor (main-plot strips, coloured borders), Nitrogen the sub-plot factor (fill colour), replicated in three blocks. (b–e) Actual randomized layouts for four example locations (two corn, two soybean); layouts differ by site because randomization is independent. Sub-plots are 12×24 m; 18 plots/location $\times$ 20 locations = 360 plots (seed 42, reproducible).

4.3 Statistical model and analysis of variance

Let Y_{ijk} be the response (e.g., grain yield, tissue-N, or an in-season index) for Water level i, Nitrogen level j in block k at a given location. The split-plot linear model is

$$Y_{ijk} = \mu + \alpha_i + \gamma_k + \delta_{ik} + \beta_j + (\alpha\beta)_{ij} + \varepsilon_{ijk}, \tag{3}$$

where μ is the grand mean, α_i the Water (main-plot) effect, γ_k the block effect, δ_{ik} the *whole-plot* (main-plot) error stratum (Block $\times$ Water), β_j the Nitrogen (sub-plot) effect, $(\alpha\beta)_{ij}$ the Water $\times$ Nitrogen interaction, and ε_{ijk} the sub-plot error. The defining feature of the split-plot analysis—and the classic mistake to avoid—is that the whole-plot factor is tested against the whole-plot error while the sub-plot factor and the interaction are tested against the sub-plot error [7, 19, 37]. Table 3 gives the per-location ANOVA skeleton with $a = 2$ Water levels, $b = 3$ Nitrogen levels and $r = 3$ blocks.

Table 3: Per-location split-plot ANOVA source table ($a = 2$ Water levels, $b = 3$ Nitrogen levels, $r = 3$ blocks). Degrees of freedom sum to $rab - 1 = 17$. Each source's F-ratio uses the error term named in the final column; using the wrong error term is the classic split-plot analysis error.

Source	Symbol	df	df formula	Tested against
Block (replicate)	γ_k	2	$r-1$	(removes nuisance variation)
Water (whole-plot, A)	α_i	1	$a-1$	Main-plot error
Main-plot error (Block $\times$ Water)	δ_{ik}	2	$(r-1)(a-1)$	— (Error a)
Nitrogen (sub-plot, B)	β_j	2	$b-1$	Sub-plot error
Water $\times$ Nitrogen	$(\alpha\beta)_{ij}$	2	$(a-1)(b-1)$	Sub-plot error
Sub-plot error	ε_{ijk}	8	$a(r-1)(b-1)$	— (Error b)
Total		**17**	$rab-1$	

The three hypothesis tests are therefore

$$F_{\text{Water}} = \frac{MS_{\text{Water}}}{MS_{\text{main-plot error}}} \sim F_{1,2}, \quad F_{\text{N}} = \frac{MS_{\text{N}}}{MS_{\text{sub-plot error}}} \sim F_{2,8}, \quad F_{\text{W}\times\text{N}} = \frac{MS_{\text{W}\times\text{N}}}{MS_{\text{sub-plot error}}} \sim F_{2,8}. \tag{4}$$

The whole-plot factor is deliberately the less-powerful test here (only 2 denominator df per location), which is the acknowledged cost of protecting the sub-plot and interaction contrasts and of assigning the hard-to-randomize factor to whole plots. In practice we recommend fitting the design as a linear mixed-effects model with block and Block $\times$ Water as random effects (for example `statsmodels`' `mixedlm` or R's `lme4/nlme`), which reproduces the correct error strata, handles the occasional missing plot gracefully, and—for a combined analysis across all twenty locations—accommodates location and location $\times$ block as additional random blocking strata [37]. Because the twenty sites were selected from the stress tail rather than sampled at random, any across-site inference should treat location as a random stratum and interpret treatment effects as conditional on this selected set.

5 Discussion and Caveats

5.1 What the workflow does—and does not—establish

The overriding interpretive principle of this report is that Sentinel-2 NDVI and NDRE detect spectral *symptoms* of reduced canopy vigor, not their *causes*. A depressed index is mechanistically non-specific: it is equally consistent with nitrogen deficiency, water deficit or excess, soil compaction, fungal or bacterial disease, insect or nematode pressure, weed competition, hail or wind damage, herbicide injury, or a simple stand-establishment gap [32, 35, 53]. The ranked list in Table 2 should therefore be read as a set of *where-to-look* hypotheses with attached confidence, never as a diagnosis. The split-plot trial of Section 4 is the mechanism by which two specific hypotheses—water and nitrogen limitation—can be tested; it is explicitly not a claim that water and nitrogen are the causes. This is why the follow-up is framed as confirmatory experimentation paired with agronomic scouting (soil tests, tissue nutrient assays, pathogen and pest identification, and growing-degree-day tracking), not as an imaging deliverable alone.

5.2 Methodological strengths

Several design choices make the product more trustworthy than a naive index map. Crop-specific standardization prevents the dominant phenological difference between corn and soybean (Figure 3) from being mistaken for stress. Restricting the primary Stress Index to the four full-coverage Peak-Vigor and Persistence metrics means the R112 orbit swath gap—which removes the establishment and senescence observations for cells in ∼36 % of the ROI on three dates—cannot bias the ranking, an honest handling of a real data limitation that a single-date or naively-averaged approach would have obscured. The 90 % purity filter suppresses the mixed-pixel errors that dominate CDL uncertainty [29, 30], and the explicit, thresholded uncertainty flags let a scout triage effort toward the most reliable leads. Pairing each anomaly with a nearby, same-crop, spectrally normal control converts an absolute index into a locally meaningful contrast that partially controls for weather, soil association and management, and the Hungarian global assignment guarantees non-overlapping controls.

5.3 Limitations

The limitations are real and are reported without euphemism. *First*, the soybean sample is small: only 16 pure cells exist in the ROI, ten are flagged, and the residual non-anomalous interquartile pool is just two cells, so four soybean controls are themselves flagged and three soybean “anomalies” have near-zero Stress Indices and are only mildly atypical. Soybean results are correspondingly weaker than corn results and should be treated as provisional. *Second*, one soybean anomaly (priority 17) is a spatial outlier matched to a control 27.7 km away, well outside the regime in which a control can be assumed to share context. *Third*, the temporal sampling is sparse (seven dates) and unevenly covering because of the R112/R069 orbit split; denser time series from harmonized Landsat–Sentinel data or gap-filling could sharpen the phenological metrics [8, 15, 40, 50]. *Fourth*, 1 km aggregation masks sub-field heterogeneity: cells flagged for high within-cell heterogeneity may contain multiple management zones, so a single centroid-anchored trial footprint may not sit on representative ground and must be adjusted to real field boundaries, avoiding roads, waterways and headlands, before staking. *Fifth*, the twenty sites are the tail of a stress distribution, not a random sample, so they are subject to selection and regression-to-the-mean effects and their treatment responses may not generalize to the wider ROI. *Sixth*, index saturation and the narrow window in which NDVI is both informative and actionable temper any expectation that imagery alone can time interventions [21, 48]. *Finally*, a null result in the split-plot trial rejects the nitrogen/water hypothesis for a site; it does not certify the field as healthy, because the true driver may lie outside the two-factor space tested.

5.4 Recommended next steps

The immediate operational path is to hand Table 2 and the accompanying GeoPackage/GeoJSON layers to a field team for georeferenced scouting, prioritizing the corn leads (which are the most extreme and most reliable) and the lower-uncertainty entries. Ground observations—stand counts, visual symptomology, soil probes and tissue tests—should either corroborate a plausible nitrogen or water mechanism, in which case the split-plot trial is established on representative ground, or reveal an out-of-scope cause (disease, compaction, pest) that redirects the investigation. A power analysis should confirm that three blocks per location deliver adequate sensitivity for the effect sizes of interest, particularly for the low-power whole-plot Water test; adding blocks is generally preferable to adding within-whole-plot replication when whole-plot degrees of freedom are limiting [7, 37]. Longer term, fusing the optical record with thermal (crop-water-stress) and radar observations, and with in-season crop-model assimilation, would add the mechanistic specificity that optical indices alone cannot provide [23, 24, 49].

6 Conclusions

We have delivered a complete, reproducible pipeline that turns seven cloud-screened 2023 Sentinel-2 acquisitions of tile `T15TVG` and the 2023 USDA CDL into a prioritized, uncertainty-aware field-scouting product for a central-Iowa ROI, together with a rigorously specified nitrogen $\times$ water split-plot experiment for confirmatory follow-up. Forty corn-pure and sixteen soybean-pure 1 km cells were characterized by NDVI/NDRE Establishment, Peak-Vigor, Persistence and Senescence metrics; crop-specific z-score anomalies identified the twenty highest-priority cells, each paired with a nearby normal control and annotated with transparent purity, heterogeneity and clear-observation flags. The corn leads are strong and reliable; the soybean leads are weaker owing to a small pure-cell pool and are flagged as such. The split-plot design and its correct ANOVA structure provide the statistically sound bridge from spectral symptom to testable agronomic hypothesis. The deliverable is emphatically a guide to *where* to scout and *what* to test, not a causal diagnosis: the imagery cannot, and this report does not, claim to know why any field is stressed. That question is left, as it must be, to the ground.

Data and code availability. All Sentinel-2 and CDL-derived layers, the ranked and paired CSV/GeoPackage/GeoJSON outputs, the phenology and design summaries (JSON), the processing scripts and a machine-readable manifest accompany this report in the project directory.

References

[1] Clement Atzberger. Advances in remote sensing of agriculture: Context description, existing operational monitoring systems and major information needs. *Remote Sensing*, 5(2):949–981, 2013. doi: 10.3390/rs5020949.

[2] G. Azzari and D.B. Lobell. Landsat-based classification in the cloud: An opportunity for a paradigm shift in land cover monitoring. *Remote Sensing of Environment*, 202:64–74, 2017. doi: 10.1016/j.rse.2017.05.025.

[3] E. M. Barnes, T. R. Clarke, S. E. Richards, P. D. Colaizzi, J. Haberland, M. Kostrzewski, P. Waller, C. Choi, E. Riley, T. Thompson, R. J. Lascano, H. Li, and M. S. Moran. Coincident detection of crop water stress, nitrogen status and canopy density using ground-based multispectral data. In *Proceedings of the 5th International Conference on Precision Agriculture, Bloomington, MN*, Madison, WI, 2000. ASA/CSSA/SSSA.

[4] I. Becker-Reshef, E. Vermote, M. Lindeman, and C. Justice. A generalized regression-based model for forecasting winter wheat yields in kansas and ukraine using modis data. *Remote Sensing of Environment*, 114(6):1312–1323, 2010. doi: 10.1016/j.rse.2010.01.010.

[5] Douglas K. Bolton and Mark A. Friedl. Forecasting crop yield using remotely sensed vegetation indices and crop phenology metrics. *Agricultural and Forest Meteorology*, 173: 74–84, 2013. doi: 10.1016/j.agrformet.2013.01.007.

[6] Claire Boryan, Zhengwei Yang, Rick Mueller, and Mike Craig. Monitoring us agriculture: the us department of agriculture, national agricultural statistics service, cropland data layer program. *Geocarto International*, 26(5):341–358, 2011. doi: 10.1080/10106049.2011.562309.

[7] Michael D. Casler. Fundamentals of experimental design: Guidelines for designing successful experiments. *Agronomy Journal*, 107(2):692–705, 2015. doi: 10.2134/agronj2013.0114.

[8] Martin Claverie, Junchang Ju, Jeffrey G. Masek, Jennifer L. Dungan, Eric F. Vermote, Jean-Claude Roger, Sergii V. Skakun, and Christopher Justice. The harmonized landsat and sentinel-2 surface reflectance data set. *Remote Sensing of Environment*, 219:145–161, 2018. doi: 10.1016/j.rse.2018.09.002.

[9] Jan Clevers, Lammert Kooistra, and Marnix Van den Brande. Using sentinel-2 data for retrieving lai and leaf and canopy chlorophyll content of a potato crop. *Remote Sensing*, 9 (5):405, 2017. doi: 10.3390/rs9050405.

[10] J.G.P.W. Clevers and A.A. Gitelson. Remote estimation of crop and grass chlorophyll and nitrogen content using red-edge bands on sentinel-2 and -3. *International Journal of Applied Earth Observation and Geoinformation*, 23:344–351, 2013. doi: 10.1016/j.jag.2012.10.008.

[11] M. Drusch, U. Del Bello, S. Carlier, O. Colin, V. Fernandez, F. Gascon, B. Hoersch, C. Isola, P. Laberinti, P. Martimort, A. Meygret, and F. Spoto. Sentinel-2: Esa's optical high-resolution mission for gmes operational services. *Remote Sensing of Environment*, 120: 25–36, 2012. doi: 10.1016/j.rse.2011.11.026.

[12] European Space Agency. Sentinel-2 MSI level-2a product format specification and algorithm theoretical basis document (scene classification and atmospheric correction). `https://sentinels.copernicus.eu/`, 2021. Copernicus Open Access documentation; accessed 2026-07-14.

[13] Ronald A. Fisher. *The Design of Experiments*. Oliver and Boyd, Edinburgh, UK, 1935.

[14] David Frantz, Erik Haß, Andreas Uhl, Johannes Stoffels, and Joachim Hill. Improvement of the fmask algorithm for sentinel-2 images: Separating clouds from bright surfaces based on parallax effects. *Remote Sensing of Environment*, 215:471–481, 2018. doi: 10.1016/j.rse.2018.04.046.

[15] Feng Gao, Jeff Masek, Matt Schwaller, and Forrest Hall. On the blending of the landsat and modis surface reflectance: predicting daily landsat surface reflectance. *IEEE Transactions on Geoscience and Remote Sensing*, 44(8):2207–2218, 2006. doi: 10.1109/TGRS.2006.872081.

[16] Anatoly A. Gitelson and Mark N. Merzlyak. Signature analysis of leaf reflectance spectra: Algorithm development for remote sensing of chlorophyll. *Journal of Plant Physiology*, 148 (3-4):494–500, 1996. doi: 10.1016/S0176-1617(96)80284-7.

[17] Anatoly A. Gitelson, Yuri Gritz †, and Mark N. Merzlyak. Relationships between leaf chlorophyll content and spectral reflectance and algorithms for non-destructive chlorophyll assessment in higher plant leaves. *Journal of Plant Physiology*, 160(3):271–282, 2003. doi: 10.1078/0176-1617-00887.

[18] Anatoly A. Gitelson, Andrés Viña, Verónica Ciganda, Donald C. Rundquist, and Timothy J. Arkebauer. Remote estimation of canopy chlorophyll content in crops. *Geophysical Research Letters*, 32(8), 2005. doi: 10.1029/2005GL022688.

[19] Kwanchai A. Gomez and Arturo A. Gomez. *Statistical Procedures for Agricultural Research*. John Wiley & Sons, New York, NY, 2nd edition, 1984.

[20] Noel Gorelick, Matt Hancher, Mike Dixon, Simon Ilyushchenko, David Thau, and Rebecca Moore. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202:18–27, 2017. doi: 10.1016/j.rse.2017.06.031.

[21] A Huete, K Didan, T Miura, E.P Rodriguez, X Gao, and L.G Ferreira. Overview of the radiometric and biophysical performance of the modis vegetation indices. *Remote Sensing of Environment*, 83(1-2):195–213, 2002. doi: 10.1016/S0034-4257(02)00096-2.

[22] E. Raymond Hunt and Craig S. T. Daughtry. What good are unmanned aircraft systems for agricultural remote sensing and precision agriculture? *International Journal of Remote Sensing*, 39(15-16):5345–5376, 2018. doi: 10.1080/01431161.2017.1410300.

[23] S.B. Idso, R.D. Jackson, P.J. Pinter, R.J. Reginato, and J.L. Hatfield. Normalizing the stress-degree-day parameter for environmental variability. *Agricultural Meteorology*, 24: 45–55, 1981. doi: 10.1016/0002-1571(81)90032-7.

[24] R. D. Jackson, S. B. Idso, R. J. Reginato, and P. J. Pinter. Canopy temperature as a crop water stress indicator. *Water Resources Research*, 17(4):1133–1138, 1981. doi: 10.1029/WR017i004p01133.

[25] Zhenong Jin, Qianlai Zhuang, Jiali Wang, Sotirios V. Archontoulis, Zachary Zobel, and Veerabhadra R. Kotamarthi. The combined and separate impacts of climate extremes on the current and future us rainfed maize and soybean production under elevated co2. *Global Change Biology*, 23(7):2687–2704, 2017. doi: 10.1111/gcb.13617.

[26] Zhenong Jin, George Azzari, Calum You, Stefania Di Tommaso, Stephen Aston, Marshall Burke, and David B. Lobell. Smallholder maize area and yield mapping at national scales with google earth engine. *Remote Sensing of Environment*, 228:115–128, 2019. doi: 10.1016/j.rse.2019.04.016.

[27] David M. Johnson. 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, 2014. doi: 10.1016/j.rse.2013.10.027.

[28] Tyler J. Lark, Richard M. Mueller, David M. Johnson, and Holly K. Gibbs. Measuring land-use and land-cover change using the u.s. department of agriculture's cropland data layer: Cautions and recommendations. *International Journal of Applied Earth Observation and Geoinformation*, 62:224–235, 2017. doi: 10.1016/j.jag.2017.06.007.

[29] Tyler J. Lark, Ian H. Schelly, and Holly K. Gibbs. Accuracy, bias, and improvements in mapping crops and cropland across the united states using the usda cropland data layer. *Remote Sensing*, 13(5):968, 2021. doi: 10.3390/rs13050968.

[30] Li Lin, Liping Di, Chen Zhang, Liying Guo, Yahui Di, Hui Li, and Anna Yang. Validation and refinement of cropland data layer using a spatial-temporal decision tree algorithm. *Scientific Data*, 9(1), 2022. doi: 10.1038/s41597-022-01169-w.

[31] David B. Lobell, David Thau, Christopher Seifert, Eric Engle, and Bertis Little. A scalable satellite-based crop yield mapper. *Remote Sensing of Environment*, 164:324–333, 2015. doi: 10.1016/j.rse.2015.04.021.

[32] Wouter H. Maes and Kathy Steppe. Perspectives for remote sensing with unmanned aerial vehicles in precision agriculture. *Trends in Plant Science*, 24(2):152–164, 2019. doi: 10.1016/j.tplants.2018.11.007.

[33] Magdalena Main-Knorn, Bringfried Pflug, Jerome Louis, Vincent Debaecker, Uwe Müller-Wilm, and Ferran Gascon. Sen2cor for sentinel-2. In *Image and Signal Processing for Remote Sensing XXIII*, page 3, 2017. doi: 10.1117/12.2278218.

[34] Michele Meroni, Dominique Fasbender, Felix Rembold, Clement Atzberger, and Anja Klisch. Near real-time vegetation anomaly detection with modis ndvi: Timeliness vs. accuracy and effect of anomaly computation options. *Remote Sensing of Environment*, 221:508–521, 2019. doi: 10.1016/j.rse.2018.11.041.

[35] David J. Mulla. Twenty five years of remote sensing in precision agriculture: Key advances and remaining knowledge gaps. *Biosystems Engineering*, 114(4):358–371, 2013. doi: 10.1016/j.biosystemseng.2012.08.009.

[36] Yi Peng and Anatoly A. Gitelson. Application of chlorophyll-related vegetation indices for remote estimation of maize productivity. *Agricultural and Forest Meteorology*, 151(9): 1267–1276, 2011. doi: 10.1016/j.agrformet.2011.05.005.

[37] H. P. Piepho, A. Büchse, and K. Emrich. A hitchhiker's guide to mixed models for randomized experiments. *Journal of Agronomy and Crop Science*, 189(5):310–322, 2003. doi: 10.1046/j.1439-037X.2003.00049.x.

[38] Felix Rembold, Clement Atzberger, Igor Savin, and Oscar Rojas. Using low resolution satellite imagery for yield prediction and yield anomaly detection. *Remote Sensing*, 5(4): 1704–1733, 2013. doi: 10.3390/rs5041704.

[39] John W. Rouse, R. H. Haas, J. A. Schell, and D. W. Deering. Monitoring vegetation systems in the great plains with ERTS. In *Third Earth Resources Technology Satellite-1 Symposium, NASA SP-351*, volume 1, pages 309–317, Washington, DC, 1974. NASA.

[40] D.P. Roy, M.A. Wulder, T.R. Loveland, Woodcock C.E., R.G. Allen, M.C. Anderson, D. Helder, J.R. Irons, D.M. Johnson, R. Kennedy, T.A. Scambos, and C.B. Schaaf. Landsat-8: Science and product vision for terrestrial global change research. *Remote Sensing of Environment*, 145:154–172, 2014. doi: 10.1016/j.rse.2014.02.001.

[41] Toshihiro Sakamoto, Brian D. Wardlow, Anatoly A. Gitelson, Shashi B. Verma, Andrew E. Suyker, and Timothy J. Arkebauer. A two-step filtering approach for detecting maize and soybean phenology with time-series modis data. *Remote Sensing of Environment*, 114(10): 2146–2159, 2010. doi: 10.1016/j.rse.2010.04.019.

[42] Peter C. Scharf, D. Kent Shannon, Harlan L. Palm, Kenneth A. Sudduth, Scott T. Drummond, Newell R. Kitchen, Larry J. Mueller, Victoria C. Hubbard, and Luciane F. Oliveira. Sensor-based nitrogen applications out-performed producer-chosen rates for corn in on-farm demonstrations. *Agronomy Journal*, 103(6):1683–1691, 2011. doi: 10.2134/agronj2011.0164.

[43] M. Schlemmer, A. Gitelson, J. Schepers, R. Ferguson, Y. Peng, J. Shanahan, and D. Rundquist. Remote estimation of nitrogen and chlorophyll contents in maize at leaf and canopy levels. *International Journal of Applied Earth Observation and Geoinformation*, 25: 47–54, 2013. doi: 10.1016/j.jag.2013.04.003.

[44] Joel Segarra, Maria Luisa Buchaillot, Jose Luis Araus, and Shawn C. Kefauver. Remote sensing for precision agriculture: Sentinel-2 improved features and applications. *Agronomy*, 10(5):641, 2020. doi: 10.3390/agronomy10050641.

[45] Rajendra P. Sishodia, Ram L. Ray, and Sudhir K. Singh. Applications of remote sensing in precision agriculture: A review. *Remote Sensing*, 12(19):3136, 2020. doi: 10.3390/rs12193136.

[46] Compton J. Tucker. Red and photographic infrared linear combinations for monitoring vegetation. *Remote Sensing of Environment*, 8(2):127–150, 1979. doi: 10.1016/0034-4257(79)90013-0.

[47] USDA National Agricultural Statistics Service. USDA national agricultural statistics service cropland data layer 2023 (published crop-specific data layer). `https://nassgeodata.gmu.edu/CropScape/`, 2024. Available at the CropScape and National CDL bulk portals; accessed 2026-07-14.

[48] Andrés Viña, Anatoly A. Gitelson, Anthony L. Nguy-Robertson, and Yi Peng. Comparison of different vegetation indices for the remote assessment of green leaf area index of crops. *Remote Sensing of Environment*, 115(12):3468–3478, 2011. doi: 10.1016/j.rse.2011.08.010.

[49] M. Weiss, F. Jacob, and G. Duveiller. Remote sensing for agricultural applications: A meta-review. *Remote Sensing of Environment*, 236:111402, 2020. doi: 10.1016/j.rse.2019.111402.

[50] Alyssa K. Whitcraft, Eric F. Vermote, Inbal Becker-Reshef, and Christopher O. Justice. Cloud cover throughout the agricultural growing season: Impacts on passive optical earth observations. *Remote Sensing of Environment*, 156:438–447, 2015. doi: 10.1016/j.rse.2014.10.009.

[51] Jinru Xue and Baofeng Su. Significant remote sensing vegetation indices: A review of developments and applications. *Journal of Sensors*, 2017:1–17, 2017. doi: 10.1155/2017/1353691.

[52] F. Yates. Complex experiments. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 2(2):181–223, 1935. doi: 10.2307/2983638.

[53] P. J. Zarco-Tejada, C. Camino, P. S. A. Beck, R. Calderon, A. Hornero, R. Hernández-Clemente, T. Kattenborn, M. Montes-Borrego, L. Susca, M. Morelli, V. Gonzalez-Dugo, and P. R. J. North. Previsual symptoms of xylella fastidiosa infection revealed in spectral plant-trait alterations. *Nature Plants*, 4(7):432–439, 2018. doi: 10.1038/s41477-018-0189-7.

[54] Xiaoyang Zhang, Mark A. Friedl, Crystal B. Schaaf, Alan H. Strahler, John C.F. Hodges, Feng Gao, Bradley C. Reed, and Alfredo Huete. Monitoring vegetation phenology using modis. *Remote Sensing of Environment*, 84(3):471–475, 2003. doi: 10.1016/S0034-4257(02)00135-9.

[55] Liheng Zhong, Lina Hu, and Hang Zhou. Deep learning based multi-temporal crop classification. *Remote Sensing of Environment*, 221:430–443, 2019. doi: 10.1016/j.rse.2018.11.032.