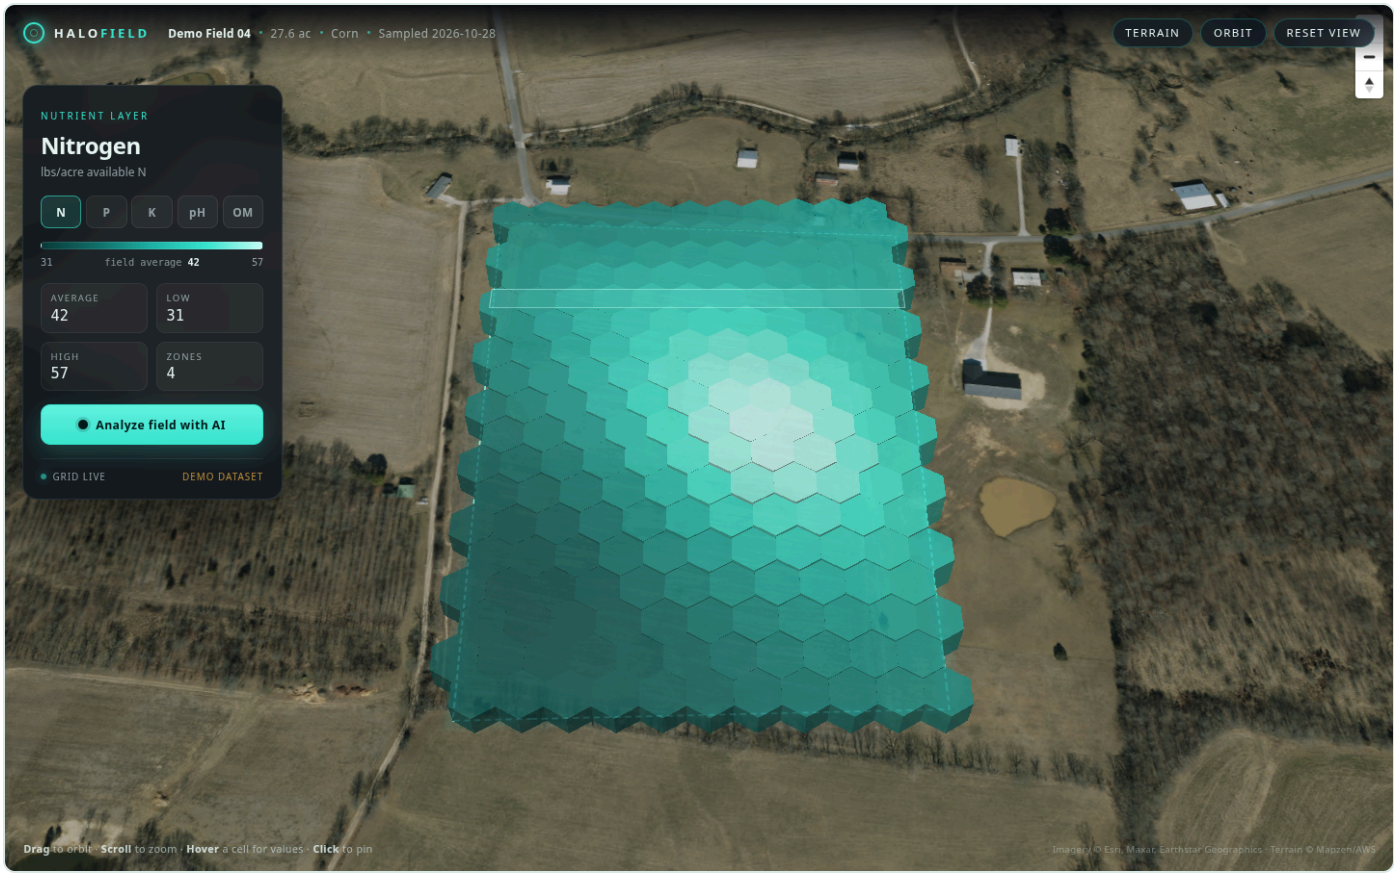
HALOFIELD

See your field from orbit.

Satellite soil intelligence: a real field on real imagery, a holographic nutrient grid, and a plain-language agronomy read. This guide shows how to run the demo and how to load your own samples.



Version 0.3 · November 2026 · Demo: http://100.80.207.66/halofield/

What you are looking at

HaloField puts a soil-sample grid back on the ground it came from. Instead of a spreadsheet, you see the parcel in satellite view with each cell coloured and raised by the nutrient you select. The AI panel turns the same numbers into a short, prioritized read.

3 demo fields

Corn, soybeans and wheat parcels near Cape Girardeau, MO — real boundaries, synthetic sample values.

5 nutrient layers

Nitrogen, phosphorus, potassium, pH and organic matter, each with its own colour ramp and height scale.

Demo data: the demo fields use generated values so the app can be shown without a lab report. Every AI read is labelled accordingly. Your own uploads replace the demo numbers entirely.

Quick tour

1. **Open the demo** at `/halofield/`. The camera flies from orbit to the first field.
2. **Pick a field** from the selector in the top bar (corn, soybeans, wheat). The HUD updates acres, crop and sample date.
3. **Switch nutrient layers** with the N / P / K / pH / OM buttons. Colour and height both change; the legend shows the field average.
4. **Hover a cell** to read its exact values; **click** to pin the readout.
5. **Drag** to orbit, **scroll** to zoom, and use *Reset view* to return to the framed field.
6. **Hit "Analyze field with AI."** A side panel opens with the field read: variability, limiting factor, three actions and a zone plan with nitrogen rates.



The AI panel reads from the grid's summary statistics — it never sees your raw points.

Reading the AI report

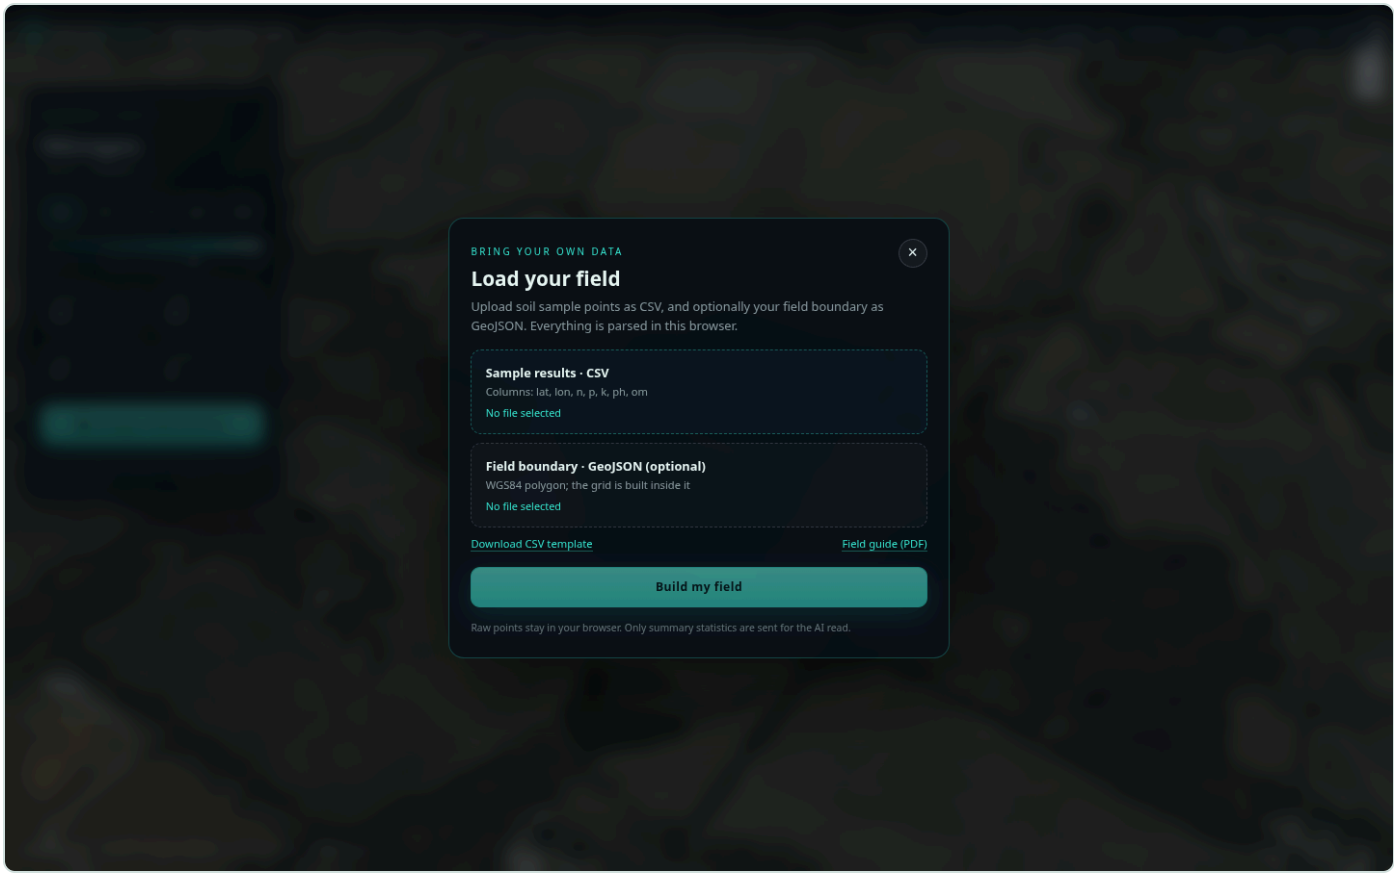
SECTION	WHAT IT MEANS
Headline	One-sentence verdict on uniformity and the main lever.
Variability	Which nutrient spreads most across the field, with the range that drives it.
Limiting factor	The nutrient or condition most likely to cap yield, and where.
Actions	Three ranked recommendations (high / medium / low) with specifics.
Zone plan	Acres and a nitrogen rate per fertility zone, ready to discuss with a spreader operator.
Confidence	How much the data supports the read. Low confidence is a signal to sample more, not a failure.

The panel also shows the model and response time; repeat analyses of the same field are cached for an hour.

> **Decision support, not a lab certification.** HaloField summarizes the numbers you give it and applies general agronomic reasoning. It does not replace a certified lab report, a crop adviser, or your own knowledge of the field.

Load your own field

Click **Upload data** in the top bar. Two files are accepted; only the CSV is required.



1 · Sample results (CSV, required)

One row per sample point, with a header row. Column names are matched loosely — `lat` / `latitude`, `lon` / `lng` / `longitude`, and nutrient names such as `n`, `nitrogen`, `p`, `k`, `ph`, `om` all work.

LAT	LON	N	P	K	PH	OM
37.39581	-89.70400	52	30	190	6.6	2.4
37.39569	-89.70069	44	26	170	6.4	2.2
37.39225	-89.70075	36	22	150	6.0	1.9

A template is linked inside the upload dialog. At least three valid points are required; missing nutrient columns fall back to regional defaults and are flagged in the interface.

2 · Field boundary (GeoJSON, optional)

If you have the parcel polygon, upload it and the grid is clipped exactly inside it. Without a boundary, HaloField builds a hull around your sample points, padded by 10 m. Boundaries must be WGS84 (EPSG:4326).

3 · Build

Press **Build my field**. The browser interpolates your points onto a hex grid (inverse-distance weighting, eight nearest samples), computes acreage from the polygon, and drops you into the field view. From there the layers, hover readout and AI panel work exactly as in the demo.

> **Privacy:** the CSV and boundary are parsed in your browser. Raw points are never uploaded. Only the summary statistics (per-nutrient min/max/average and zone aggregates) are sent for the AI read.

A two-minute demo script

1. **Start wide.** Open the demo without parameters and let the globe settle. “This is your field, from orbit.”
2. **Land.** The camera flies to the corn parcel. Point out the boundary hugging the actual field edge.
3. **Make it physical.** Switch N → P → K. “Height is the value; colour is the value. You can see the low corner without reading a table.”
4. **Read one cell.** Hover the low corner. “Thirty-one pounds of N here, fifty-seven on the high side.”
5. **Ask for the plan.** Run the AI read. “It says nitrogen is the lever, and here's a rate per zone.”
6. **Hand it over.** Open Upload data. “Send me your lab CSV and your boundary; this becomes your field in about a minute.”

What to ask a grower

- Do you have sample points with coordinates, or only a lab summary?
- Do you have a field boundary (shapefile, GeoJSON, or a farm-management export)?
- Which crop and yield goal should the nitrogen rates target?
- Who writes the variable-rate prescription today, and in what file format?

Current limits

- Demo values are synthetic; uploads are the real path.
- Interpolation is IDW, not kriging — fine for a demo, worth upgrading for production prescriptions.
- No VRT export yet (ISO-XML / shapefile); the zone plan is a discussion aid, not a machine file.
- Boundaries are straight-edged unless you upload a real polygon.
- AI output is decision support and should be reviewed by an agronomist.