

GroundTruth AI

Decision-grade geospatial evidence for infrastructure, water and environmental teams

Partnership screening deck

Prepared for potential data, analytics and validation partners. No raw coordinates, code, model weights, patent claims or training data are included.



The gap: imagery is abundant, decisions are still hard

Customers do not buy imagery. They buy confidence on where to inspect, where to invest, and where not to waste field time.

Raw Earth observation


Satellite scenes, SAR products, DEM layers and weather data are powerful but difficult for non-specialists to interpret consistently.

Unexplained model scores

A single heatmap or probability score rarely tells a field team why a location is worth investigating or how much uncertainty remains.

Field teams need priorities

Water, infrastructure and environmental teams need ranked areas, evidence traces, and reviewable assumptions.



Market signal: SkyFI Series A press release described demand for satellite imagery plus AI-powered analytics across government and commercial sectors; HULO announced seed funding for AI-powered water leak detection.

What GroundTruth produces



Output is designed as a decision artifact, not a black-box prediction.

1. Ranked areas

A prioritized short-list of locations for field investigation or review.

2. Evidence trace

Which data layers and time windows contributed to each finding.

3. Confidence + limits

Calibrated score, uncertainty bands and known failure modes.

4. Partner-ready report

A concise PDF / API object for a municipal, utility or engineering workflow.

Primary first use case: reduce the area that must be inspected manually for water and infrastructure risk triage.

Evidence sources: modular and partner-friendly

The platform can consume standard open EO/GIS layers first, then add proprietary partner data when available.

Optical time-series

Vegetation, surface change, seasonal patterns

SAR backscatter

Moisture-sensitive and structure-sensitive signals

Terrain / DEM

Slope, local relief, flow accumulation and positional context

Thermal layer


Optional partner layer for heat, stress or water-related anomalies

Operational data

Optional utility logs, pressure, flow, inspections or complaints

Human review

Field validation, expert override, failure labeling



Validation is the differentiator

Many geospatial AI demos fail when moved to a new region. GroundTruth treats reliability controls as a product feature.

Leakage control

Spatial holdout and metadata discipline reduce shortcuts from repeated geography, coordinate leakage or temporal contamination.

Regional behavior

Global and regional models are compared to identify performance collapse or region-specific signal strength.

Uncertainty reporting

Results are expressed as prioritized evidence rather than definitive claims when field validation is still pending.

Claim discipline

Separate validated model behavior from proposed commercial use cases.

Current internal signal summary



Internal results are included only to show validation maturity. They do not yet prove commercial water-leak detection without a dedicated pilot.

Partner integration models

Each prospective partner gets a narrow, low-risk entry point. The goal is not to expose core IP; it is to prove additional decision value.

WaterTech partner

Combine network intelligence with external geospatial screening for areas with sparse sensors or limited history.

Target: HULO-style collaboration.

EO platform partner

Offer GroundTruth as a specialized analytics layer that turns imagery into ranked investigation outputs.

Target: SkyFi-style integration.

Thermal data partner

Evaluate thermal observations as an incremental evidence layer for water, heat and infrastructure stress.

Target: SatLeo-style pilot.

AI reliability partner

Define a geospatial AI evaluation benchmark for leakage, unsupported claims and regional collapse.

Target: Patronus-style discussion.

Voice AI partner


Explore Arabic field reporting for non-GIS users.

Target: AethexAI-style concept.

Non-disclosure boundary: no code, raw coordinates, model weights, patent claims, or training labels are included in screening materials.

Proposed limited pilot

A useful pilot should test whether GroundTruth reduces field-search effort, not whether it replaces field verification.



Pilot scope

One non-sensitive geography, one partner dataset if available, a fixed review period, and pre-agreed validation criteria.

Success metric

Reduction in field investigation area, quality of ranked priorities, and explainability of the evidence trail.

Deliverables

Ranked map, concise PDF report, evidence table, limitations note, and next-step commercialization recommendation.

Suggested first commercial wedge: Water & Infrastructure Intelligence - prioritizing where inspection should begin.

Contact: Ali Faraj | GroundTruth AI