



GroundTruth AI - Sanitized Evidence Summary

Purpose: provide enough technical confidence for early partnership discussions without exposing raw coordinates, training labels, model weights, source code, patent claims or proprietary extraction logic.

Positioning discipline
This document does not claim that GroundTruth AI has commercially validated leak detection, groundwater discovery, tunnel discovery or oil exploration. It summarizes the validation discipline behind the geospatial evidence engine and defines what a dedicated pilot should test.

1. What has actually been validated

- Core capability: a geospatial machine-learning workflow that turns satellite, terrain and temporal indicators into ranked locations for further review.
- Validation maturity: the project has moved from random-split performance toward leakage-controlled geographic validation and regional-model behavior analysis.
- Commercial translation: the strongest near-term claim is prioritization of field investigation, not replacement of field verification.
- First recommended wedge: Water & Infrastructure Intelligence - reduce where inspection begins, then verify with operational or field data.

The current work should be described as a decision-support evidence layer. It is strongest when used to narrow a search area, explain why a location is prioritized, and document uncertainty for a human or operational follow-up.

Disclosure boundary

Can be shared in screening materials	Should not be shared before NDA / formal pilot
High-level workflow; example map output; validation principles; aggregate internal results; pilot design.	Raw coordinates; training dataset; negative generation logic; model weights; code; exact feature list and weights; unpublished patent claims; sensitive local paths or account details.



2. Validation design and leakage controls

The value of GroundTruth AI is not simply the model score. It is the process used to prevent misleading geospatial performance claims.

Control	Why it matters	How it should be presented externally
Spatial holdout	Random train/test splits can overstate performance because nearby geography repeats similar signals.	Use separated geographic validation and report generalization honestly.
Coordinate discipline	Lat/lon or location proxies can create shortcut learning that does not transfer.	State that spatial leakage controls are part of the methodology.
Temporal discipline	Using future or repeated observations can contaminate evaluation.	Define time windows and avoid implying real-time proof where it was not tested.
Regional comparison	A global model may collapse in specific terrain, climate or data regimes.	Compare global and regional behavior; do not hide lower specificity.
Uncertainty reporting	Field teams need reliable priorities, not absolute claims.	Report confidence, limitations and failure modes with the map.

Recommended external wording

Safe claim
GroundTruth AI helps prioritize where investigation should begin by combining geospatial evidence layers with leakage-controlled validation and uncertainty reporting.

Avoid this claim until field validation exists
GroundTruth AI detects leaks, groundwater, tunnels or subsurface assets with confirmed operational accuracy.



3. Internal signal summary and limitations

The figures below are suitable only as sanitized aggregate indicators. They should not be presented as final commercial accuracy for a new use case.

Metric family	Current aggregate signal	Interpretation
Global model AUC	~0.653-0.700	Useful but not sufficient for strong cross-region deployment without additional controls.
Regional experts AUC	~0.773-0.803	Promising uplift in regional behavior under current internal experiments.
Regional experts recall	~0.864-0.907	High sensitivity; useful for triage where missing candidates is costly.
Specificity	Lower than desired	False positives remain a concern; ranking and field-verification workflow are essential.
Canonical status	Experimental uplift, not final replacement	Do not overstate the regional expert models as a locked commercial system.

Investor/partner interpretation

- The validation story is credible because it exposes degradation instead of hiding it.
- The commercial product should sell prioritization and investigation efficiency, not final truth.
- The next milestone is an externally bounded pilot with partner data or field feedback.
- A strong negative or limitation should be documented; it increases credibility with technical partners.



4. What the proposed pilots should test

A pilot should be designed to answer one operational question: does GroundTruth improve the quality and efficiency of where teams look first?

Pilot element	Recommended design
Geography	One non-sensitive test area with defined boundaries and stable access to imagery and/or operational records.
Inputs	Open EO/GIS layers plus one optional partner layer: utility records, thermal observations, imagery access or field reports.
Baseline	Partner current workflow, historical inspection records, random spatial baseline or expert-only prioritization.
Primary metric	Reduction in field-search area at comparable or improved hit-rate / review quality.
Secondary metrics	Explainability, review time, false-positive burden, confidence calibration and partner workflow fit.
Deliverable	Ranked map, evidence table, limitations page, concise executive report and next-step recommendation.

Suggested first ask to partners

Partner ask
Provide a small non-sensitive area, a sample data layer if available, and a 30-45 minute technical review. GroundTruth will return a ranked evidence output and a limitations note.

Sources for market context: HULO seed funding press release, 21 Oct 2025; SkyFi Series A press release, 14 Jan 2026; SatLeo Labs funding coverage, Apr 2026; Patronus AI Series B press release, 25 Jun 2026; AethexAI funding coverage, Jun 2026.
URLs: prnewswire.com/news-releases/hulo-raises-2-3m-seed-round-to-combat-global-water-loss-with-ai-powered-leak-detection-302589402.html | skyfi.com/en/press/skyfi-series-a-12-7-million-2026 | timesofindia.indiatimes.com/city/ahmedabad/ahmedabad-spacetechn-startup-raises-2-2mn-in-seed-round-to-accelerate-thermal-satellite-mission/articleshow/130152132.cms | prnewswire.com/news-releases/patronus-ai-raises-50-million-series-b-and-unveils-first-digital-world-models-for-ai-agent-training-and-simulation-302811248.html | thesaasnews.com/news/aethexai-raises-3m-pre-seed/