



UMAY ANA

RESPONSIBLE AI POLICY

Human-Centered Agricultural Intelligence & Institutional Decision Support

Version 1.0 | September 2026

Enterprise AI & Agricultural Intelligence Services

Banks • Insurers • Reinsurers • Public Institutions • Agribusiness

Important drafting note. This policy is written for Umay Ana’s actual positioning: AI-supported agricultural intelligence and institutional risk support. It avoids presenting advisory signals as guaranteed diagnoses, certified MRV outcomes, or autonomous credit/insurance decisions.

1. Purpose

Umay Ana uses artificial intelligence to help farmers and institutions interpret agricultural observations, weather and climate context, farm practices, images, structured records, and derived signals. This Responsible AI Policy establishes principles for designing, deploying, monitoring, and communicating those capabilities in a manner that supports human judgment, safety, fairness, privacy, and traceability.

2. Core Principles

Principle	Commitment
Human accountability	AI supports people; responsibility for material decisions remains with authorized human decision-makers.
Purpose limitation	Models and data are used for defined agricultural and institutional purposes, not unrelated profiling.
Transparency	Clearly distinguish AI-generated guidance, derived signals, self-recorded evidence, and verified external facts where practical.
Uncertainty	Avoid implying certainty where model confidence, evidence quality, data coverage, or agronomic conditions are uncertain.
Fairness	Evaluate high-impact use cases for unjustified differential outcomes and avoid prohibited or irrelevant sensitive attributes.
Privacy & security	Apply data minimization, access controls, and appropriate contractual/privacy safeguards.
Traceability	Preserve sufficient metadata and provenance for enterprise review where technically and legally appropriate.

3. Decision Boundaries

Unless a separate validated and legally compliant solution expressly provides otherwise, Umay Ana outputs are decision-support signals and must not be used as the sole basis for:

- Approving or denying credit or changing legally significant credit terms.
- Automatically paying, rejecting, or materially determining an insurance claim.
- Making a final underwriting, pricing, cancellation, or coverage decision.
- Certifying carbon credits, sustainability compliance, food safety, or regulatory eligibility.
- Replacing professional agronomic, veterinary, medical, legal, financial, or engineering judgment where such professional judgment is required.

4. Agricultural Advice and Safety

Agricultural conditions vary materially by location, soil, cultivar, weather, farm practice, local regulation, and data quality. Umay Ana should present recommendations as context-aware guidance rather than guaranteed outcomes. Higher-risk recommendations involving pesticides, veterinary treatment, food safety, legal limits, or regulated inputs should be constrained, clearly qualified, and referred to authoritative local instructions or qualified professionals where appropriate.

5. Image and Disease Analysis

Image-based outputs may identify patterns consistent with plant conditions or risks but are not guaranteed diagnoses. Product interfaces and institutional integrations should avoid converting uncertain image observations into definitive claim, credit, or regulatory outcomes without corroborating evidence. Confidence, evidence status, and review prompts should be used where useful and technically feasible.

6. Institutional Risk Signals

Derived signals - such as early field-risk indicators, claim-triage context, portfolio monitoring signals, or sustainability cohorts - are intended to prioritize attention and support analysis. They do not, by themselves, prove causation, default, fraud, loss, eligibility, or compliance. Enterprise customers are responsible for establishing lawful decision policies, thresholds, validation, and human review appropriate to their regulated use case.

7. Data Quality and Provenance

- Track whether information is user-provided, externally sourced, system-generated, AI-derived, or aggregated where practical.
- Communicate material data gaps and avoid precision that exceeds underlying evidence.
- Use temporal and geographic context when agricultural relevance depends on place and time.
- Prefer validated structured fields for institutional workflows over free-form model text when feasible.

8. Model and Provider Governance

Umay Ana may use one or more third-party or internally configured AI models. Model selection may change over time. Provider will consider capability, safety, privacy, security, latency, cost, and reliability when selecting or replacing models. Material customer-specific restrictions may be documented in an Order Form or DPA.

9. Evaluation and Monitoring

- Test representative scenarios before materially expanding higher-impact workflows.
- Monitor operational failures, user feedback, safety issues, and unexpected output patterns.
- Review significant model or prompt changes in proportion to risk.
- Maintain rollback, feature-disable, or rule-based mitigation options where practical for critical functionality.
- Use outcome validation when Customer lawfully provides sufficient ground-truth data; avoid claiming predictive performance that has not been validated.

10. Fairness and Prohibited Use

Umay Ana should not be intentionally used to discriminate unlawfully or infer sensitive traits that are irrelevant to the agricultural purpose. Enterprise customers must not use the Services to circumvent applicable fair-lending, insurance, employment, consumer-protection, privacy, or anti-discrimination obligations. Provider may restrict or suspend use that creates a material legal, safety, or platform-integrity risk.

11. Explainability and Review

For institutional use cases, Umay Ana aims to provide decision context that can be reviewed by people - for example source class, time/location context, evidence status, contributing agricultural factors, or supporting observations - rather than presenting a single unexplained score as conclusive. The appropriate level of explanation will depend on the use case, model, data, and contractual scope.

12. Governance and Continuous Improvement

Responsible AI is an operating practice rather than a one-time document. Provider may update this policy as the product, laws, model ecosystem, and customer use cases evolve. Material contractual commitments remain governed by the MSA, DPA, SLA, Order Form, and other signed terms.

Acceptance and Relationship to the MSA

This document forms part of the Umay Ana enterprise contracting framework when incorporated by reference into a Master Services Agreement, Order Form, Statement of Work, or other written agreement executed by the parties. If there is a conflict, the order of precedence stated in the applicable agreement or Order Form will apply.

This template is intended to support commercial contracting and enterprise readiness. It should be reviewed by qualified counsel for the governing jurisdiction, the customer's regulatory status, data flows, and negotiated risk allocation before signature or public reliance.