LeafGuard

A Decentralized Crop Disease Early-Warning Network on Bittensor

Bittensor Subnet Proposal | June 2026

1. Problem Statement

Smallholder farmers — roughly 500 million households worldwide, producing a third of the planet's food — lose an estimated 20 to 40 percent of crop yield to pests and disease every year. Early diagnosis is the single biggest lever against this loss: a leaf blight caught in week one is a local treatment, the same blight caught in week three is a regional outbreak.

Yet the farmers who need diagnostic tools most are the ones least served by the market. Centralized agri-tech companies build disease-detection apps for high-revenue commercial farms in wealthy regions, because that's where the paying customers are. A smallholder farmer in rural Bangladesh, Malawi, or Peru is typically left with a smartphone, a sick-looking plant, and no affordable way to know whether it's a minor fungal spot or the start of an outbreak that will hit every farm in the valley.

This is a market failure, not a technology failure: the models that could solve this already exist in research papers. What's missing is a sustainable incentive for anyone to operate, maintain, and continuously improve them for a low-revenue user base.

2. How the Subnet Works

LeafGuard is a Bittensor subnet where miners compete to diagnose crop disease from photos and forecast its spread, and validators score them against agronomist-verified ground truth and real-world outcomes.

Input

A farmer (or a field agent, NGO worker, or partner app) submits a photo of an affected leaf along with GPS location and crop type. The subnet's gateway layer pulls local weather data (temperature, humidity, recent rainfall) for that location automatically, since weather is a major driver of fungal and pest spread.

Miners

Miners run computer vision and forecasting models that:

- Classify the disease or pest from the leaf image (e.g. late blight, leaf rust, aphid damage)
- Estimate severity and likely progression stage
- Predict spread risk to nearby farms over the following 1–2 weeks, using the weather feed

Miners are free to use any model architecture — open-source vision transformers, fine-tuned CNNs, ensemble approaches — and are rewarded purely on accuracy, not on which technique they used. This is the same competitive, permissionless structure used in other Bittensor subnets: validators evaluate outputs and submit weight vectors on-chain, and Yuma Consensus converts those weights into emissions.

Validators

Validators score miners against two complementary signals:

- A held-out gold-standard set of leaf images independently labeled by agronomists, refreshed regularly so miners can't overfit to a static benchmark
- Real-world outcome tracking — when a miner predicts an outbreak will spread to a region, validators check weeks later (via field reports, partner NGOs, or follow-up photos) whether it actually did, and reward miners whose predictions held up

Validators submit a weight vector ranking miners on these two criteria. Per Bittensor's standard mechanism, validators who set weights close to the emerging consensus build stronger trust and earn higher dividends, while outlier or collusive scoring gets clipped and earns nothing — this keeps any single validator from unfairly favoring particular miners.

Output

The subnet aggregates the best-scoring miner predictions into a regional risk layer: a live map of disease pressure that farmers, agronomists, and partner organizations can query through a simple app or SMS interface, with no subscription cost to the farmer.

Subnet Architecture

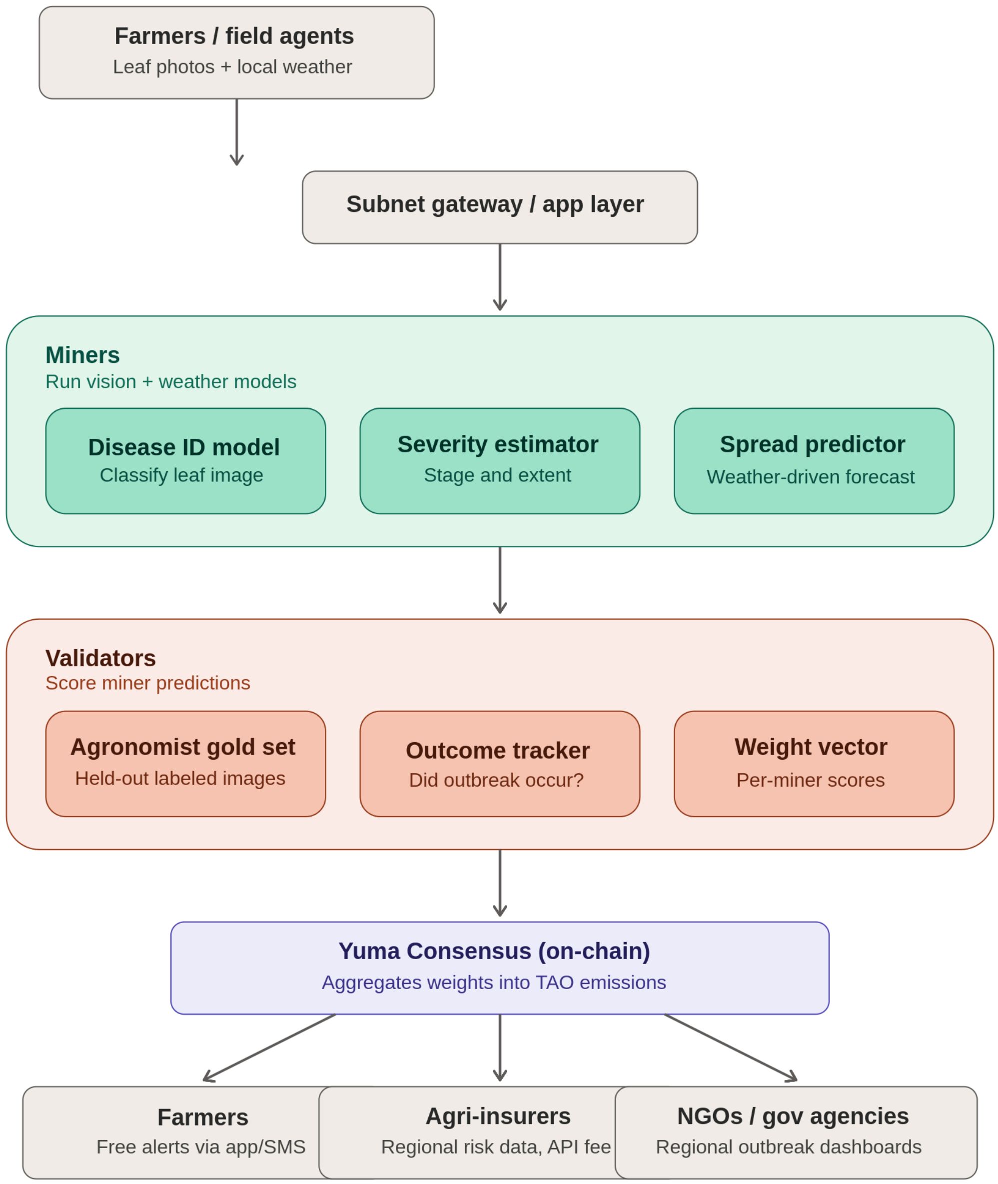

Figure 1. Data flow from farmer submission through miner diagnosis, validator scoring, and Yuma Consensus to end users.

3. Why This Needs to Be Decentralized

Three properties of this problem make it a poor fit for a single centralized company and a good fit for Bittensor's incentive design:

- No company has a commercial reason to serve this population well. Smallholder farmers can't pay enterprise prices, so a centralized provider has no incentive to keep improving the model once it's "good enough" for a grant report. A subnet's miners are paid continuously in TAO for staying state-of-the-art, regardless of whether the end users can pay.

- The ground truth is inherently distributed. No single lab has leaves from every crop variety, climate zone, and disease strain on Earth. A global, permissionless network of contributors — farmers submitting photos, local agronomists labeling them, field agents reporting outcomes — can build a training and evaluation set that no centralized research team could assemble alone.
- Competitive, continuous re-scoring beats a static model. Plant pathogens evolve and regional conditions shift. A subnet's validators keep refreshing the gold-standard set and outcome data, so miners are constantly re-evaluated against current reality instead of a model that was "accurate enough" at launch and never revisited.

4. Who Benefits

Stakeholder	Benefit
Smallholder farmers	Free or near-free early disease diagnosis and regional outbreak alerts, in their own language and over SMS where smartphones are scarce.
Agricultural NGOs and extension services	A live regional risk dashboard to prioritize where field agents and treatment supplies go, instead of reacting after an outbreak is already visible.
Agri-insurers and commodity buyers	A paid API for regional crop-risk data to price insurance and forecast supply, becoming a sustainable revenue source that funds subnet emissions beyond pure TAO speculation.
Miners	TAO emissions for building and continuously improving an open, high-value vision and forecasting model, with full freedom over architecture.
Validators	Dividends for sourcing and maintaining credible ground truth (agronomist labels and outcome data) — itself a valuable, defensible role.

5. Miner–Validator Interaction Summary

- Miners register on the subnet and run inference on a rolling stream of farmer-submitted images plus weather context.
- Validators periodically query miners with both gold-standard test images and live farmer submissions, then compare miner outputs against agronomist labels and tracked outcomes.
- Each validator computes a weight vector ranking miner performance and submits it on-chain.
- Yuma Consensus aggregates these weight vectors across all validators, weighting by stake, to compute final TAO emissions — rewarding miners for accuracy and validators for evaluations that hold up against the broader validator community's consensus.
- Subnet emissions follow Bittensor's standard split: miners and validators each receive a share of emissions based on performance, with a smaller share routed to the subnet creator for continued development.

6. Closing Note

LeafGuard isn't a hypothetical use case bolted onto Bittensor for novelty's sake — it's a problem that's underserved precisely because it's unprofitable to solve centrally, and well-suited to decentralization because the data and ground truth needed to solve it are inherently scattered across millions of small farms worldwide. That's exactly the gap Bittensor's incentive design is built to fill.