



Agricultural ML & Climate Modeller

Remote (Canada) | \$60,000–\$85,000/year

ABOUT US

Vintality Tech Inc. is a precision agriculture consultancy building practical decision-support tools for growers across BC and into California. We combine IoT sensor networks, geospatial data, and agronomic models to give vineyards and fruit farms real-time, actionable intelligence — crop stress signals, irrigation triggers, disease risk, phenology forecasts. We're a small, fast-moving team, and the models you build here will be used by real farmers making real decisions every growing season.

THE ROLE

We're looking for someone who sits at the intersection of machine learning, climate science, and applied agriculture. Your core job is to build, validate, and push forward the models that drive our platform's agronomic intelligence — ET-based water balance, heat and water stress indices, plant stress signals derived from soil and canopy data, disease risk, phenology, and climate-informed variety and site analysis.

This is primarily a research and programming role. You'll work closely with our full-stack developer to get models into production as Lambda pipelines, and with agronomic advisors and growers to understand what the data should be telling us — but you won't be the primary client-facing voice. What we need from you is the ability to find innovative ways to combine sensor streams, weather data, and crop science into better signals, and to ship those improvements into a live system used by real farms.

WHAT YOU'LL DO

Climate & Crop Stress Modelling (Core Focus)

- Build, validate, and improve models that quantify plant stress from sensor data: ET-based water balance (FAO-56 Penman-Monteith), heat stress accumulation, VPD stress, and soil moisture deficit models that go beyond simple threshold alerts.
- Develop and validate GDD-based phenology models for *Vitis vinifera*, calibrated to Okanagan, Pacific Northwest, and California growing conditions — and push them further using ML to improve prediction accuracy and lead time.
- Research and implement novel approaches to detecting plant stress signals from the data we already collect — canopy temperature, soil moisture by depth, microclimate variation — finding combinations that produce earlier and more reliable agronomic signals.
- Build climate-informed models for variety suitability, heat accumulation trends, and long-range risk indicators as the platform evolves.
- Implement new disease risk models (Botrytis, downy mildew) and nutrient stress indices grounded in peer-reviewed viticulture and crop science literature.



Model Validation & Production

- Design rigorous validation approaches: backtesting against historical sensor data, cross-site generalization, uncertainty quantification, and failure mode analysis.
- Build repeatable evaluation pipelines — metrics, drift monitoring, and documentation of what changed between model versions and why.
- Identify when a model produces implausible outputs and diagnose whether the cause is a parameter issue, data quality problem, or a model assumption that doesn't hold at a specific site.
- Produce season-end validation summaries comparing model outputs against observed field outcomes — the feedback loop that drives model improvement.

Data Pipelines & Production Delivery

- Write and maintain Python agronomic calculation pipelines running as AWS Lambda functions — reading from PostgreSQL, computing model outputs, writing results back.
- Collaborate with the full-stack developer on database schema design for model tables — specifying inputs, outputs, and update cadence in enough detail for the infrastructure to be built around it.
- Define how models should behave when sensor inputs are incomplete, missing, or suspect — graceful degradation matters in production.
- Work with geospatial data as a modelling input: soil variability zones, NDVI, spatial field observations from Mergin Maps — not GIS administration, but using spatial context to improve model outputs.

Research & Innovation

- Stay current with the viticulture, precision agriculture, and climate modelling literature and bring relevant findings back to the team — we read and we apply.
- Propose and scope new model features: what data would we need, what's the agronomic basis, how would we validate it, how does it fit into the platform?
- Contribute to Vintality's longer-term research direction as the company expands into new crops and geographies.

WHAT WE'RE LOOKING FOR

Must-have:

- Graduate degree (M.Sc. or equivalent research experience) in Agricultural Engineering, Bioresource Engineering, Environmental Science, Applied ML, or a closely related field with a strong agricultural or climate component.
- Solid Python for scientific computing and modelling — numpy, pandas, scikit-learn; comfortable implementing algorithms from literature, not just tuning existing libraries.
- Strong ML fundamentals: train/validation discipline, time-series awareness, overfitting, uncertainty quantification, and the ability to explain what a model does and where it fails.
- Working understanding of plant-water relations, ET-based irrigation science, or crop stress physiology — enough to ground modelling choices in agronomic reality, not just optimize a loss function.



- Ability to work independently in a remote setting and drive a research question from literature through implementation to validated production output.

Nice-to-have:

- Viticulture-specific background — *Vitis vinifera* physiology, GDD phenology, Okanagan or Pacific Northwest growing conditions.
- Experience with climate modelling, downscaling, or long-range agroclimate forecasting.
- Familiarity with geospatial data and tools (GeoPandas, PostGIS, QGIS) as modelling inputs rather than cartographic outputs.
- Experience with time-series anomaly detection, sensor data quality, or IoT data pipelines.
- Exposure to AWS (Lambda, RDS) or similar cloud platforms — enough to deploy and debug a calculation pipeline in production.
- Published or conference-presented work in agricultural modelling, precision agriculture, or climate-crop systems.

PORTFOLIO OR WORK SAMPLES (ENCOURAGED)

Show us how you think about a modelling problem. Strong examples include:

- A model or algorithm you implemented from the literature — including the agronomic or scientific basis, your validation approach, and what you'd do differently now.
- A research project or thesis chapter involving agricultural, climate, or environmental data — even a methods section or results summary works.
- A GitHub repo, notebook, or script implementing a crop or climate calculation from raw sensor or weather data.
- A write-up of a modelling decision you made: why you chose a particular approach, what you tested against, what you learned.

WHAT SUCCESS LOOKS LIKE (FIRST 90 DAYS)

By 30 Days

- You understand every model currently running in production — inputs, outputs, assumptions, limitations — and have a documented view on what you'd improve and why.
- You can run the data pipeline end-to-end: from raw sensor readings through Lambda calculation to the output that appears on the dashboard.

By 60 Days

- You've shipped at least one model improvement into production — a recalibrated parameter, a better stress signal, an improved validation metric — with before/after documentation.
- You've proposed at least one new modelling approach based on the data we have and the literature, with a clear validation plan.



By 90 Days

- You own the modelling and calculation layer of the platform — you can explain every output, defend every threshold with a reference, and catch when something is wrong before it reaches a grower.
- You've produced a validation summary for the most recent season and identified at least two concrete improvements for the next one.
- You've specced or started building at least one new model feature — defining inputs, outputs, and the scientific basis — that extends what the platform can currently do.

WHAT WE OFFER

- Salary of \$60,000–\$75,000/year, based on experience
- Annual Health Spending Account
- Extended health & dental benefits
- Fully remote — work from anywhere in Canada
- Flexible schedule in a collaborative, small-team environment
- Real data, real farms, real impact — your models inform actual farming decisions every season
- Professional development and room to grow as the company expands

HOW TO APPLY

Send your resume and a short cover letter to:

hire@vintality.com

Please include portfolio links or work samples if available, and a brief note on what draws you to agricultural or climate modelling — and what you'd want to build that we haven't thought of yet.