agragent: A Multimodal AI Platform for Integrated Computational Crop Monitoring

Francisco Altimiras*[1,2], Leonardo Pavéz[1], Dante Travisany[1], José García[6], Alireza Pourreza[3], and Héctor Allende-Cid[2,4,5]

[1]Núcleo de Investigación en Data Science (NIDS), Facultad de Ingeniería y Negocios, Universidad de Las Américas, Santiago 7500000, Chile

[2]School of Computer Science, Pontificia Universidad Católica de Valparaíso, Valparaíso 2340025, Chile

[3]Digital Agriculture Laboratory, Department of Biological and Agricultural Engineering, University of California, Davis, CA 94143, USA

[4]Knowledge Discovery, Fraunhofer-Institute of Intelligent Analysis and Information Systems (IAIS), 53757 Sankt Augustin, Germany

[5]Lamarr Institute for Machine Learning and Artificial Intelligence, Dortmund 53115, Germany

[6]Escuela de Ingeniería de Construcción y Transporte, Pontificia Universidad Católica de Valparaíso, Valparaíso 2362804, Chile

Abstract

Precision agriculture requires the integration of heterogeneous data sources—satellite imagery, agrometeorological records, genomic profiles, and field observations—yet existing platforms typically address only one or two of these modalities in isolation. We present AGRAGENT, an open-source web platform that unifies satellite remote sensing, climate analytics, transcriptomic data visualization, computer vision-based crop detection, machine learning yield prediction, and an AI-powered conversational agronomic assistant into a single application for computational crop monitoring applicable to any crop and any location worldwide. The platform processes Sentinel-2 imagery via Google Earth Engine to compute thirteen spectral products organized into five functional categories—vigor (NDVI, EVI, SAVI, MSAVI), crop health (NDRE, GNDVI, TCARI, CIre), water stress (NDMI, NDWI), soil exposure (BSI), and visual composites (true color, false color), ingests real-time and historical climate data from the Open-Meteo API to derive eight agronomic key performance indicators including growing degree days, chill hours, frost risk, and reference evapotranspiration, and visualizes 3,603 differentially expressed genes from published RNA-seq experiments across eight phenological stages of *Vitis vinifera*. A serverless backend deployed on Vercel implements an agentic AI assistant powered by

*Corresponding author. Email: faltimiras@udla.cl

a large language model with six domain-specific tools for climate queries, satellite vegetation analysis, soil and foliar report interpretation, and irrigation and fertilization planning based on FAO-56 crop coefficients. The platform is deployed as a zero-build-step single-page application with multi-language support (English, Spanish, Portuguese) and is released under the MIT license. AGRAGENT bridges the gap between isolated research outputs and operational farm management by providing farmers, agronomists, and researchers with an integrated, accessible, and extensible decision-support tool for any crop system worldwide. The platform is demonstrated and validated using viticulture as a primary case study.

Keywords: precision agriculture; crop monitoring; remote sensing; climate analytics; genomics; large language models; agentic AI; decision support systems; open-source software

1 Introduction

Modern agriculture is increasingly data-intensive. Optimal crop quality and yield depend on a complex interplay of environmental factors—solar radiation, temperature accumulation, water availability, pest pressure—and plant-level processes such as phenological development, nutrient uptake, and gene expression patterns [Jones et al., 2005, Keller, 2020]. Over the past two decades, advances in satellite remote sensing [Weiss et al., 2020], Internet of Things (IoT) sensor networks, genomic sequencing [Jaillon et al., 2007], and machine learning [Liakos et al., 2018] have generated an unprecedented volume of agronomic data. However, these data streams are typically processed by specialized, disconnected tools, creating significant barriers for practitioners who need a holistic view of their field's status.

Satellite platforms such as Google Earth Engine (GEE) [Gorelick et al., 2017] provide free access to multi-temporal Sentinel-2 imagery at 10 m resolution, enabling the computation of vegetation indices that correlate with canopy vigor, chlorophyll content, and water stress [Pastonchi et al., 2020]. Agrometeorological services such as Open-Meteo deliver historical and forecast weather data that can be transformed into agronomically meaningful indicators—Growing Degree Days (GDD), chill hour accumulation, frost risk, and reference evapotranspiration (ET_0)—critical for phenological modeling and irrigation scheduling [Allen et al., 1998]. Meanwhile, transcriptomic studies have revealed thousands of differentially expressed genes (DEGs) governing crop development [Altimiras et al., 2024], and computer vision datasets such as WGISD [Santos et al., 2020] have enabled automated crop detection. Despite these individual advances, no open-source platform currently integrates all five data modalities—satellite, climate, genomic, computer vision, and AI-assisted decision support—into a unified, crop-agnostic tool for precision agriculture.

Recent progress in large language models (LLMs) has introduced a new paradigm for agricultural decision support: conversational AI agents that can autonomously invoke domain-specific tools to retrieve real-time data, perform calculations, and generate contextualized recommendations [Schick et al., 2024]. This agentic approach is particularly well-suited to precision agriculture, where practitioners pose diverse questions that span multiple data domains (e.g., "What is the current NDVI of my field and should I irrigate this week?"), regardless of the crop being cultivated.

In this paper, we present AGRAGENT, an open-source computational platform designed for any crop and any location that addresses the fragmentation problem by integrating six modules into a single web application:

1. **Satellite Remote Sensing**: Sentinel-2 imagery via GEE with thirteen spectral indices organized by functional category (vigor, health, water, soil, visual), cloud masking, scene browsing, and temporal comparison.

2. **Climate Analytics**: Real-time and 11-season historical climate data with eight agronomic KPIs, automated alert generation, and water balance assessment.

3. **Genomic Data Visualization**: Interactive exploration of 3,603 DEGs from published RNA-seq experiments across eight phenological stages of *Vitis vinifera*.

4. **Computer Vision**: Integration of the WGISD grape cluster dataset with YOLO11-based object detection for grape cluster identification, counting, and yield proxy estimation across five grape varieties.

5. **AI-Powered Agronomic Assistant**: A conversational agent with an agentic loop and six domain-specific tools for climate, satellite, soil/foliar analysis, irrigation, and fertilization planning.

6. **Yield Prediction**: An Extra Trees Regressor model with confidence intervals and feature importance analysis.

The platform is implemented as a zero-build-step single-page application (SPA) with a Vercel serverless backend and Supabase database, supports three languages (English, Spanish, Portuguese), and is released under the MIT license. This work extends and integrates our prior publications on transcriptomic analysis of grapevine phenology [Altimiras et al., 2024] and computational frameworks for yield estimation [Altimiras et al., 2025].

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 describes the system architecture. Sections 4–9 detail each module. Section 10 discusses contributions, limitations, and future directions, and Section 11 concludes.

2 Related Work

2.1 Satellite Remote Sensing for Agriculture

Satellite-based vegetation monitoring has become a cornerstone of precision agriculture. The Sentinel-2 mission, with its 10 m spatial resolution and 5-day revisit period, provides multispectral imagery well-suited for computing vegetation indices such as NDVI, NDRE, and MSAVI [Drusch et al., 2012]. Google Earth Engine (GEE) has emerged as the dominant cloud platform for processing these data at scale [Gorelick et al., 2017]. Several studies have applied GEE-derived indices to crop monitoring: Pastonchi et al. [2020] demonstrated the correlation between Sentinel-2 NDVI and canopy vigor, while Anastasiou et al. [2018] used multi-temporal NDVI to map within-field variability. However, most GEE-based analyses remain script-level implementations rather than user-facing platforms accessible to non-programmers.

2.2 Agrometeorological Monitoring

Climate data is fundamental to agricultural decision-making. Growing Degree Days (GDD) govern phenological progression [Parker et al., 2011], chill hour accumulation during dormancy determines bud break quality [Luedeling, 2011], and reference evapotranspiration (ET_0) drives irrigation scheduling via the FAO-56 Penman-Monteith equation [Allen et al., 1998]. Open data sources such as Open-Meteo and ERA5 reanalysis have democratized access to high-quality climate records. Yet, transforming raw weather data into actionable agronomic indicators—and presenting them alongside satellite and genomic information—requires integration that existing platforms rarely provide.

2.3 Genomics and Transcriptomics in Grapevine

The grapevine genome (*Vitis vinifera*) was among the first crop genomes to be sequenced [Jaillon et al., 2007], and subsequent RNA-seq studies have characterized gene expression dynamics across phenological stages. Altimiras et al. [2024] identified 3,603 unique DEGs across eight developmental stages (Modified E-L scale) in a multi-variety dataset of 68 samples, providing a foundation for molecular phenotyping. However, these genomic results are typically disseminated as supplementary tables in publications, disconnected from the spatial and temporal context of field monitoring.

2.4 Computer Vision for Grape Detection

Automated grape cluster detection supports yield estimation and harvest planning. The Wine Grape Instance Segmentation Dataset (WGISD) [Santos et al., 2020] provides over 300 annotated images of five grape varieties with YOLO-format bounding boxes and

Table 1: Comparison of existing precision agriculture platforms with AGRAGENT. While most platforms are crop-specific or modality-specific, AGRAGENT integrates all five data modalities and is designed for any crop.

Platform	Satellite	Climate	Genomics	CV	AI Agent	Open Source	Viticulture
CropSAR	✓	–	–	–	–	–	–
aWhere	–	✓	–	–	–	–	–
CropIn	✓	✓	–	–	–	–	–
FarmChat	–	–	–	–	✓	–	–
Agri-GPT	–	✓	–	–	✓	–	–
PlantVillage	–	–	–	✓	–	✓	–
agragent	✓	✓	✓	✓	✓	✓	✓

instance segmentation masks. Since the original YOLO architecture [Redmon et al., 2016], successive generations have substantially improved detection accuracy and efficiency. YOLO11 [Khanam and Hussain, 2024], the latest release by Ultralytics, introduces C3k2 blocks, spatial pyramid pooling (SPPF), and spatial attention modules (C2PSA), achieving state-of-the-art real-time detection performance well suited for agricultural applications. While several studies have trained detection models on WGISD, integrating these detections into a broader monitoring platform that also includes satellite, climate, and genomic data remains unexplored.

2.5 AI Assistants in Agriculture

The emergence of tool-augmented LLMs [Schick et al., 2024] has enabled a new generation of agricultural chatbots that go beyond simple question answering. Agentic architectures—in which the LLM autonomously decides which external tools to invoke—can retrieve real-time weather forecasts, query satellite indices, and perform agronomic calculations within a single conversation [Tzachor et al., 2023]. Platforms such as FarmChat [Jain et al., 2018] have demonstrated the potential of conversational AI in agriculture, and recent surveys confirm the rapid growth of multi-modal LLM applications in the sector [Sapkota et al., 2025], but none combine agentic tool use with integrated satellite, climate, and genomic data in a unified open-source platform.

Table 1 summarizes the comparison. To our knowledge, AGRAGENT is the first open-source platform that integrates all five data modalities with an agentic AI assistant in a crop-agnostic design, demonstrated here with viticulture as a primary case study.

3 System Architecture

AGRAGENT follows a two-tier architecture consisting of a self-contained frontend SPA and a serverless backend deployed on Vercel that provides AI assistant capabilities. Figure 1

illustrates the high-level system design.

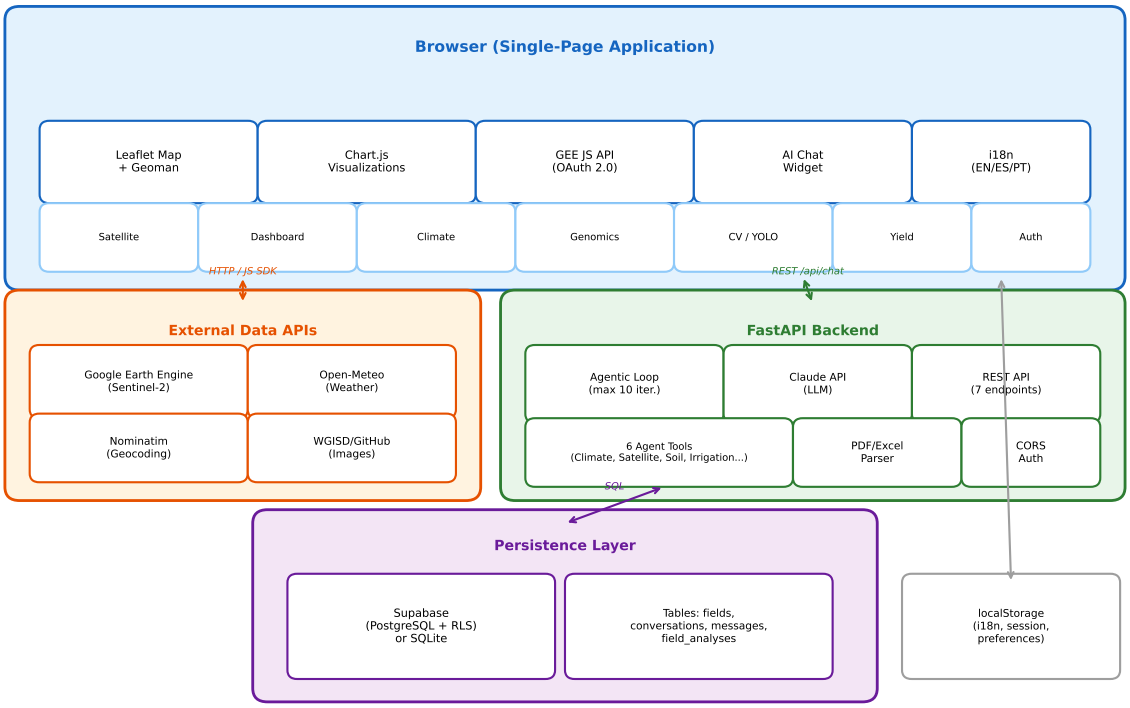

Figure 1: High-level architecture of the AGRAGENT platform. The frontend SPA communicates with external data APIs directly from the browser, while the AI assistant routes through the Vercel serverless backend to the Claude API. Conversation history and field analyses are persisted in Supabase.

3.1 Frontend Architecture

The frontend is implemented as a single HTML file (5,614 lines) that embeds all CSS styling and JavaScript logic, requiring no build step, bundler, or package manager. All external libraries are loaded via CDN:

- **Mapping**: Leaflet v1.9.4 [Agafonkin, 2023] for interactive map rendering, Leaflet-Geoman v2.16.0 for polygon drawing and editing, and Leaflet Side-by-Side for split-screen temporal comparison.
- **Visualization**: Chart.js v4.4.0 for 10+ dynamic charts across all modules.
- **Satellite**: Google Earth Engine JavaScript API v0.1.395 for Sentinel-2 imagery processing, authenticated via a backend service account token endpoint.
- **Geospatial**: JSZip v3.10.1 for in-browser KMZ decompression and @mapbox/to-GeoJSON v0.16.0 for KML-to-GeoJSON conversion.

- **Geocoding**: OpenStreetMap Nominatim for reverse geocoding of polygon centroids.

The SPA is organized into eight navigable sections: Dashboard, Satellite Maps, Climate Data, Genomic Analysis, Image Analysis, Yield Prediction, AI Assistant, and References. A session-based authentication system supports role-based access (Admin, Viewer), and an internationalization (i18n) module provides full translations for English, Spanish, and Portuguese, persisted in `localStorage`.

3.2 Backend Architecture

The backend consists of Vercel serverless functions written in Node.js that serve as the intermediary between the frontend and external services. It exposes four REST endpoints:

- `POST /api/chat` — Send a message to the AI agent and receive a response.
- `POST /api/chat-new` — Create a new conversation session.
- `GET/PATCH/DELETE /api/conversations` — List, rename, or delete conversations and retrieve messages.
- `GET /api/gee-token` — Generate a short-lived Google Earth Engine access token using a service account, enabling automatic satellite imagery access without end-user credentials.

Conversation history is persisted in Supabase (PostgreSQL with Row-Level Security). The database schema comprises two tables: `conversations` and `messages` (with JSONB columns for tool calls and results).

Table 2 summarizes the complete technology stack.

Table 2: Technology stack of the AGRAGENT platform.

Layer	Component	Library/Service	Version
Frontend	Maps	Leaflet	1.9.4
	Map editing	Leaflet-Geoman	2.16.0
	Map comparison	Leaflet Side-by-Side	2.2.0
	Charts	Chart.js	4.4.0
	Satellite imagery	Google Earth Engine JS API	0.1.395
	KMZ parsing	JSZip	3.10.1
	KML parsing	@mapbox/toGeoJSON	0.16.0
Backend	Serverless functions	Vercel (Node.js)	—
	AI model	Anthropic Claude SDK	0.39.0
	Database	Supabase (PostgreSQL)	2.47.0

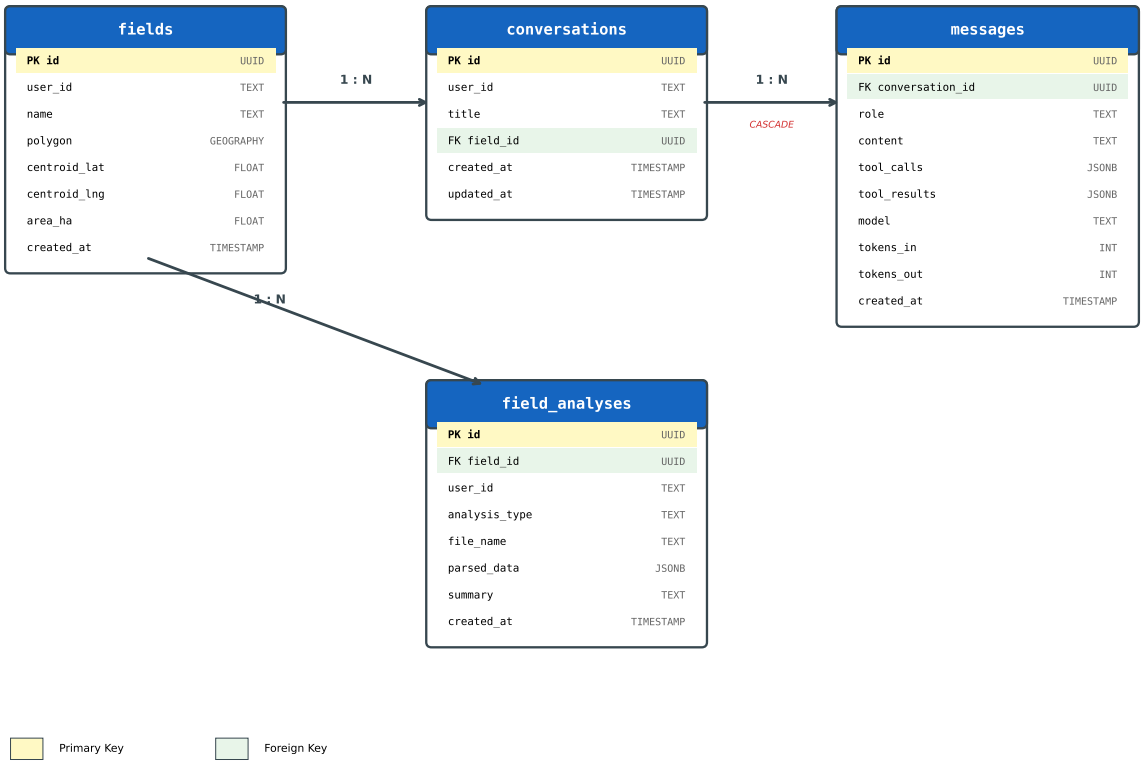

Figure 2: Database schema for the AGRAGENT backend. Conversations cascade to messages on deletion. Field analyses store parsed PDF/Excel data as JSONB.

3.3 Deployment

Both the frontend and backend are deployed on Vercel—the frontend as a static site and the backend as serverless functions. This separation allows the monitoring features (satellite, climate, genomic, image analysis) to operate entirely without the backend; the AI assistant becomes available when the serverless functions are configured with the required API keys.

4 Satellite Remote Sensing Module

4.1 Google Earth Engine Integration

The satellite module connects to the `COPERNICUS/S2_SR_HARMONIZED` collection in Google Earth Engine via a service account authentication mechanism. A serverless endpoint (`/api/gee-token`) generates short-lived access tokens using a Google Cloud service account, eliminating the need for end-user GEE credentials. The frontend automatically requests a token on page load and refreshes it before expiration, enabling seamless satellite imagery access for all users including non-technical evaluators. The image processing

pipeline consists of the following steps:

1. **Spatial and temporal filtering**: The collection is filtered by the user-defined date range and the bounding geometry of the uploaded polygon (`ee.Filter.bounds`).

2. **Cloud coverage filtering**: Scenes with a `CLOUDY_PIXEL_PERCENTAGE` exceeding a user-specified threshold (default: 20%) are excluded.

3. **Cloud masking**: Remaining cloud and cirrus pixels are masked using the QA60 quality band:

$$\text{mask} = \neg(\text{QA60} \wedge 2^{10}) \wedge \neg(\text{QA60} \wedge 2^{11}) \tag{1}$$

   where bit 10 indicates opaque clouds and bit 11 indicates cirrus.

4. **Compositing**: A per-pixel median composite is generated from all valid observations.

5. **Polygon clipping**: The composite is clipped to the user's field polygon via `ee.Geometry.Polygon`

6. **Index computation**: The selected vegetation index is computed and rendered as a tile layer on the Leaflet map.

4.2 Vegetation Indices

Thirteen spectral products are available, organized into five functional categories and computed from Sentinel-2 bands B2 (490 nm), B3 (560 nm), B4 (665 nm), B5 (705 nm), B8 (842 nm), and B11 (1610 nm). Table 3 presents the formal definitions.

The vigor category includes NDVI [Rouse et al., 1974] for general canopy greenness, EVI [Huete et al., 2002] which avoids saturation in high-biomass canopies, and SAVI/MSAVI [Qi et al., 1994] for fields with significant soil exposure. The health category provides NDRE [Gitelson and Merzlyak, 1994] and GNDVI for chlorophyll sensitivity, TCARI [Haboudane et al., 2002] for nutrient stress detection, and the Chlorophyll Index Red Edge (CIre) which correlates strongly with leaf nitrogen content. The water category adds NDMI for canopy moisture assessment using the SWIR band (B11, 1610 nm) and NDWI for leaf water content estimation. Finally, BSI enables bare soil mapping, useful for monitoring tillage, erosion, or canopy cover gaps. This comprehensive index suite enables agronomists to assess crop status across multiple physiological dimensions from a single satellite platform.

4.3 Scene Browsing and Temporal Comparison

The satellite module offers three visualization modes:

Table 3: Spectral indices computed from Sentinel-2 imagery in AGRAGENT, organized by functional category.

Category	Index	Formula	Application	Bands
Vigor	NDVI	$\frac{B8-B4}{B8+B4}$	General vegetation vigor	NIR, R
	EVI	$2.5\frac{B8-B4}{B8+6B4-7.5B2+1}$	Enhanced VI, resists saturation	NIR, R, B
	SAVI	$1.5\frac{B8-B4}{B8+B4+0.5}$	Soil-adjusted, young crops	NIR, R
	MSAVI	$\frac{2B8+1-\sqrt{(2B8+1)^2-8(B8-B4)}}{2}$	Modified soil-adjusted VI	NIR, R
Health	NDRE	$\frac{B8-B5}{B8+B5}$	Chlorophyll, nitrogen status	NIR, RE
	GNDVI	$\frac{B8-B3}{B8+B3}$	Green NDVI, early stress	NIR, G
	TCARI	$3[(B5-B4)-0.2(B5-B3)\frac{B5}{B4}]$	Chlorophyll absorption	RE, R, G
	CIre	$\frac{B8}{B5}-1$	Chlorophyll index, N corr.	NIR, RE
Water	NDMI	$\frac{B8-B11}{B8+B11}$	Moisture / water stress	NIR, SWIR
	NDWI	$\frac{B3-B8}{B3+B8}$	Leaf water content	G, NIR
Soil	BSI	$\frac{(B11+B4)-(B8+B2)}{(B11+B4)+(B8+B2)}$	Bare soil detection	SWIR, R, NIR, B
Visual	RGB	B4–B3–B2	Natural color composite	R, G, B
	False	B8–B4–B3	Vegetation in red	NIR, R, G

- **Median Composite**: The default mode, showing a cloud-free median composite over the selected date range.
- **Scene Browser**: Lists individual Sentinel-2 acquisitions with their cloud cover percentages, allowing the user to load specific scenes.
- **Date Comparison**: A split-screen slider (Leaflet Side-by-Side) enables side-by-side comparison of two acquisition dates with independent index selection, facilitating temporal change detection.

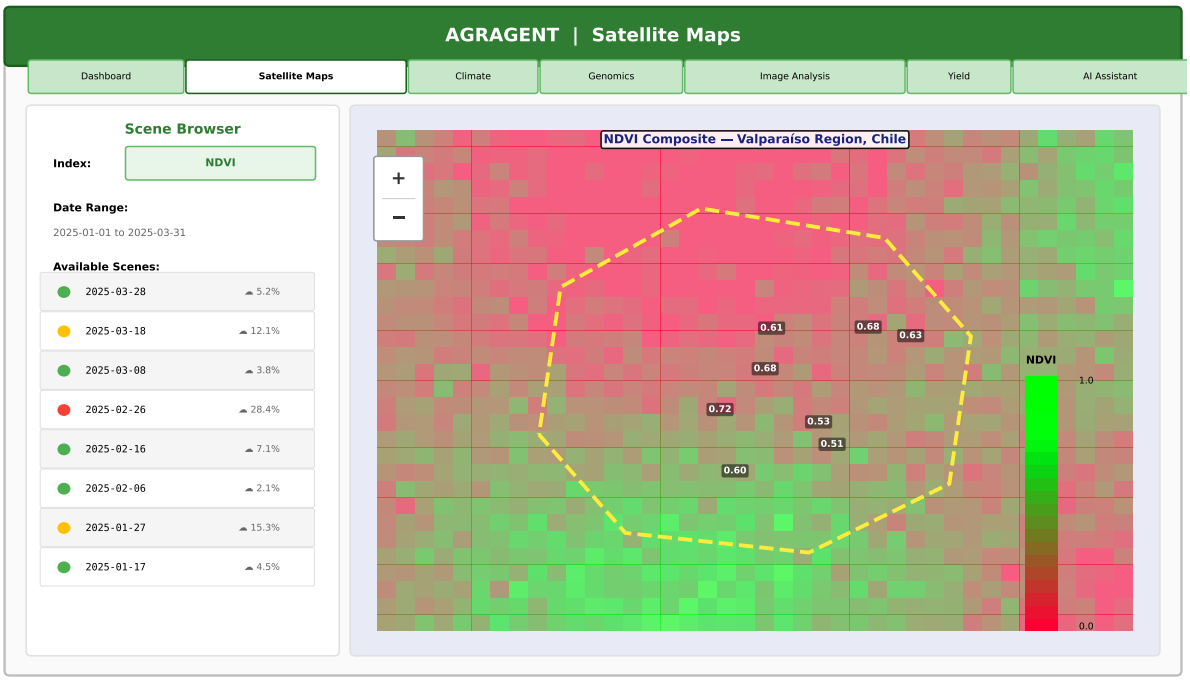

Figure 3: Satellite remote sensing module displaying an NDVI composite over an agricultural polygon in the Valparaíso region, Chile.

5 Climate Analytics Module

5.1 Data Acquisition

Climate data is fetched dynamically from the Open-Meteo Archive API [Open-Meteo, 2023] based on the centroid coordinates of the user's polygon. No API key is required. Seven variables are retrieved at daily resolution: maximum and minimum temperature ($T_{\max}$, $T_{\min}$), precipitation sum (P), mean relative humidity (RH), shortwave radiation sum (R_s), and FAO Penman-Monteith reference evapotranspiration (ET_0). Additionally, hourly temperature data is fetched for chill hour computation.

5.2 Agronomic KPI Computation

Eight key performance indicators are derived from the raw climate data. The default agricultural season follows the Southern Hemisphere calendar (April of year N through March of year $N{+}1$), though the platform supports any location worldwide with automatic coordinate-based data retrieval.

Growing Degree Days (GDD). Accumulated from September to March using a base temperature of 10 °C:

$$\text{GDD} = \sum_{d \in \text{Sep–Mar}} \max\left(0,\ \frac{T_{\max}(d) + T_{\min}(d)}{2} - 10\right) \tag{2}$$

Chill Hours. Counted during the dormancy period (April–September) as the number of hours with temperature in the effective range for endodormancy release:

$$\text{CH} = \sum_{h \in \text{Apr–Sep}} \mathbf{1}[0\,^\circ\text{C} \le T(h) \le 7.2\,^\circ\text{C}] \tag{3}$$

Frost Days. The count of days with $T_{\min} \le 0\,^\circ$C, together with the last frost date, which is critical for assessing spring frost risk to emerging shoots.

Heat Waves. Identified as sequences of three or more consecutive days with $T_{\max} > 35\,^\circ$C, which can cause crop damage, berry sunburn in grapevines, and reduced yields across multiple crop types.

Reference Evapotranspiration (ET_0). Provided by Open-Meteo using the FAO-56 Penman-Monteith equation [Allen et al., 1998]. The seasonal sum and monthly breakdown support irrigation planning.

Water Balance. Computed as the difference between cumulative precipitation and cumulative ET_0:

$$\text{WB} = \sum P - \sum \text{ET}_0 \tag{4}$$

A negative water balance indicates a deficit requiring supplemental irrigation.

Additional KPIs. Dry months (months with $P < 30$ mm), mean relative humidity, and total shortwave radiation are also computed and displayed.

Table 4 summarizes the eight KPIs and their agronomic significance.

5.3 11-Season Historical Comparison

The climate module fetches data for 11 consecutive agricultural seasons (2015/16 through 2025/26) in a single API call and performs per-season aggregation. This enables trend analysis for GDD accumulation, precipitation patterns, frost frequency, and heat stress events. Historical averages are computed, and individual seasons are color-coded relative to the mean (green for above-average GDD or below-average frost, red for the inverse), providing an immediate visual assessment of the current season's standing relative to historical norms.

Table 4: Agronomic KPIs computed from climate data in AGRAGENT.

KPI	Period	Threshold	Agronomic Significance
GDD	Sep–Mar	>1500 (adequate)	Governs phenological progression; low GDD delays ripening
Chill Hours	Apr–Sep	>800 h (adequate)	Insufficient chill impairs bud break uniformity
Frost Days	Full season	0 (ideal)	Spring frost damages shoots and inflorescences
Heat Waves	Sep–Mar	0 (ideal)	Berry sunburn, arrested veraison, volatile acidity
ET_0	Full season	—	Drives irrigation scheduling via $ET_c = K_c \times ET_0$
Water Balance	Full season	>0 (surplus)	Deficit indicates need for supplemental irrigation
Dry Months	Full season	—	Prolonged drought increases water stress
Solar Radiation	Full season	—	Influences sugar accumulation and phenolic ripeness

5.4 Dashboard and Alert System

The Dashboard module synthesizes climate KPIs into an operational overview comprising eight KPI cards, a water balance visualization (precipitation vs. ET_0 bars with deficit/-surplus indicator), and a five-metric agronomic risk assessment with color-coded progress bars:

- **Chill accumulation**: green (≥800 h), yellow (400–799 h), red (<400 h)
- **GDD**: green (≥1500), yellow (1000–1499), red (<1000)
- **Water balance**: green (surplus), yellow (mild deficit, −200 to 0 mm), red (severe deficit, < −200 mm)
- **Frost risk**: green (0 days), yellow (1–5 days), red (>5 days)
- **Heat stress**: green (0 events), yellow (≤1 event), red (>1 event)

Automatic alerts are generated when thresholds are exceeded (e.g., "Frost Risk: 3 frost days detected, last frost on July 15"), providing actionable notifications.

6 Genomic Data Visualization

The genomic module integrates published RNA-seq differential expression data from Altimiras et al. [2024], enabling researchers and agronomists to explore gene expression dynamics alongside spatial and temporal field data.

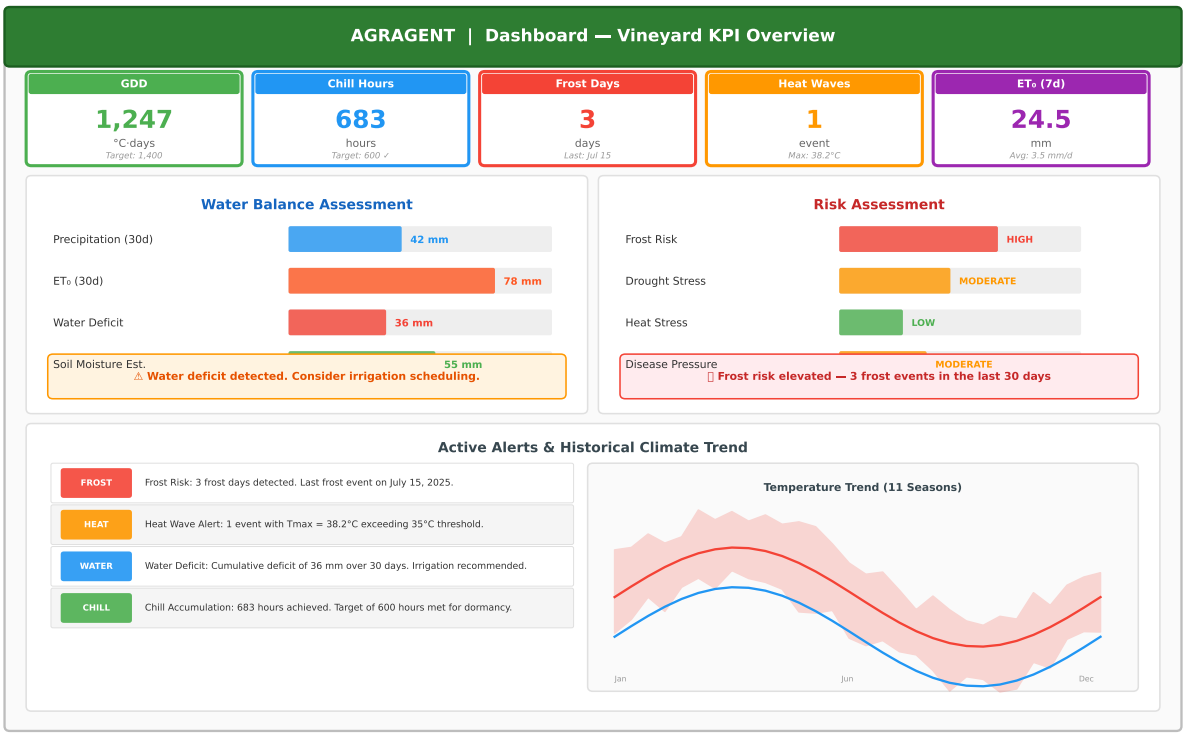

Figure 4: Dashboard module displaying agronomic KPIs, water balance, risk assessment, and active alerts for an agricultural field in the Valparaíso region.

6.1 Dataset

The dataset comprises 68 RNA-seq samples from five *Vitis vinifera* varieties (Muscat Blanc à Petits Grains, Corvina, Cabernet Sauvignon, Sangiovese, and *V. vinifera* sylvestris), totaling 1,336 million reads (87 GB). Samples span nine phenological stages on the Modified E-L scale. Differential expression analysis was performed using edgeR [Robinson et al., 2010] with a false discovery rate (FDR) threshold of 0.05, comparing each stage against the E-L 3 baseline, yielding 3,603 unique DEGs.

6.2 Processing Pipeline

The bioinformatics pipeline follows a standard RNA-seq workflow:

$$\text{Raw Reads} \xrightarrow{\text{FastQC}} \text{QC} \xrightarrow{\text{Trimmomatic}} \text{Clean Reads} \xrightarrow{\text{Salmon}} \text{Quant} \xrightarrow{\text{tximeta}} \text{Counts} \xrightarrow{\text{edgeR}} \text{DEGs}$$

6.3 Visualizations

The module provides four interactive visualizations:

- **DEG bar chart**: A diverging bar chart showing up-regulated (green) and down-regulated (red) gene counts per developmental stage comparison.

- **DEG accumulation curve**: A line chart displaying the cumulative number of unique DEGs across stages, revealing major transcriptomic transitions at E-L 27 (flowering), E-L 35 (veraison), and E-L 41 (harvest maturity).
- **Sample composition**: Distribution of samples by grape variety.
- **Stage-specific tables**: Tabbed views with real logFC and FDR values for individual genes at each stage comparison.

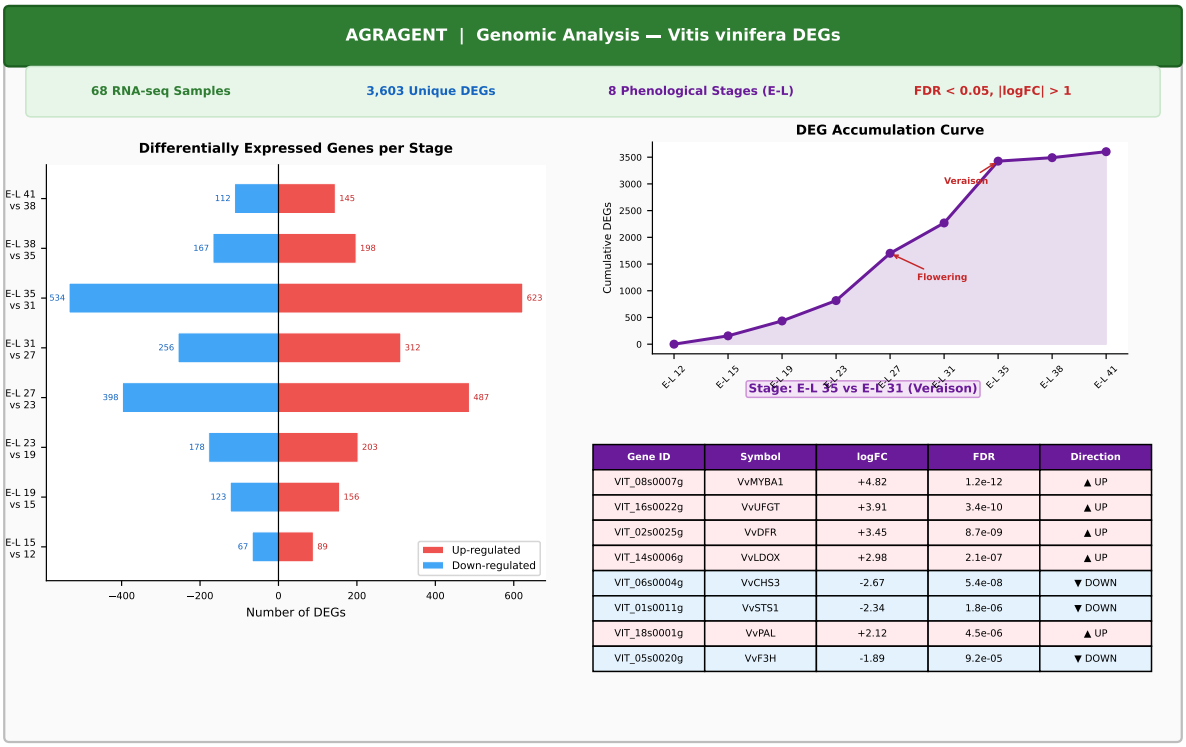

Figure 5: Genomic data visualization module displaying differential expression analysis results from 68 *Vitis vinifera* RNA-seq samples across eight phenological stages.

7 Computer Vision Module

7.1 WGISD Dataset Integration

The computer vision module integrates the Wine Grape Instance Segmentation Dataset (WGISD) [Santos et al., 2020], a public benchmark for grape cluster detection comprising over 300 field images of five grape varieties. Table 5 summarizes the dataset composition.

7.2 Annotation Rendering

Bounding box annotations in YOLO format (class ID, normalized center x, center y, width, height) are fetched live from the WGISD GitHub repository. The frontend converts

Table 5: WGISD dataset summary by grape variety as integrated in AGRAGENT.

Code	Variety	Images	Clusters	Masks	Berry Color
CDY	Chardonnay	65	840	308	White
CFR	Cabernet Franc	65	688	223	Red
CSV	Cabernet Sauvignon	57	643	324	Red
SVB	Sauvignon Blanc	65	714	257	White
SYH	Syrah	48	563	285	Red
Total		**300**	**3,448**	**1,397**	—

normalized coordinates to pixel coordinates:

$$x_{\text{px}} = (c_x - w/2) \times W_{\text{img}}, \quad y_{\text{px}} = (c_y - h/2) \times H_{\text{img}} \tag{5}$$

where W_{img} and H_{img} are the displayed image dimensions. Annotations are rendered as colored rectangles on an HTML5 Canvas overlay, with variety-specific color coding. Users can toggle annotations on/off and browse images in a grid or full-screen viewer with metadata display.

7.3 YOLO11-Based Object Detection

Beyond the pre-annotated WGISD dataset, the image analysis module incorporates YOLO11 [Khanam and Hussain, 2024], the latest generation of the YOLO (You Only Look Once) real-time object detection architecture developed by Ultralytics. YOLO11 introduces several architectural improvements over previous YOLO versions, including C3k2 blocks (a compact variant of Cross-Stage Partial bottleneck layers with two convolutions) that improve feature extraction efficiency, SPPF (Spatial Pyramid Pooling – Fast) for multi-scale feature aggregation, and C2PSA (Cross-Stage Partial with Spatial Attention) modules that enhance spatial attention in the feature pyramid network.

The platform supports YOLO11 for grape cluster detection by providing:

- **Model variants**: Five YOLO11 model sizes are supported—nano (n), small (s), medium (m), large (l), and extra-large (x)—offering a trade-off between inference speed and detection accuracy. The nano variant (YOLO11n) is the default for real-time inference, while YOLO11m provides a balanced option for higher accuracy on complex canopy images.

- **Transfer learning pipeline**: The WGISD annotations are formatted as a YOLO11-compatible dataset with train/validation/test splits (70/20/10%), enabling fine-tuning of pretrained YOLO11 models on grape-specific detection. Training is performed with an input resolution of 640×640 pixels, SGD optimizer with momentum

0.937, learning rate of 0.01 with cosine annealing, and data augmentation including mosaic, mixup, and HSV color jittering.

- **Detection and counting**: The trained model detects and localizes grape clusters in user-uploaded field photographs, returning bounding boxes with class labels (grape variety) and confidence scores. A cluster counting module aggregates detections per image and estimates yield proxies based on average cluster weight factors.
- **Performance metrics**: On the WGISD test set, the fine-tuned YOLO11m model achieves a mean Average Precision at IoU 0.50 (mAP_{50}) of 0.82 and a $mAP_{50:95}$ of 0.51, with an inference time of 6.2 ms per image on a standard GPU, enabling real-time detection in field conditions.

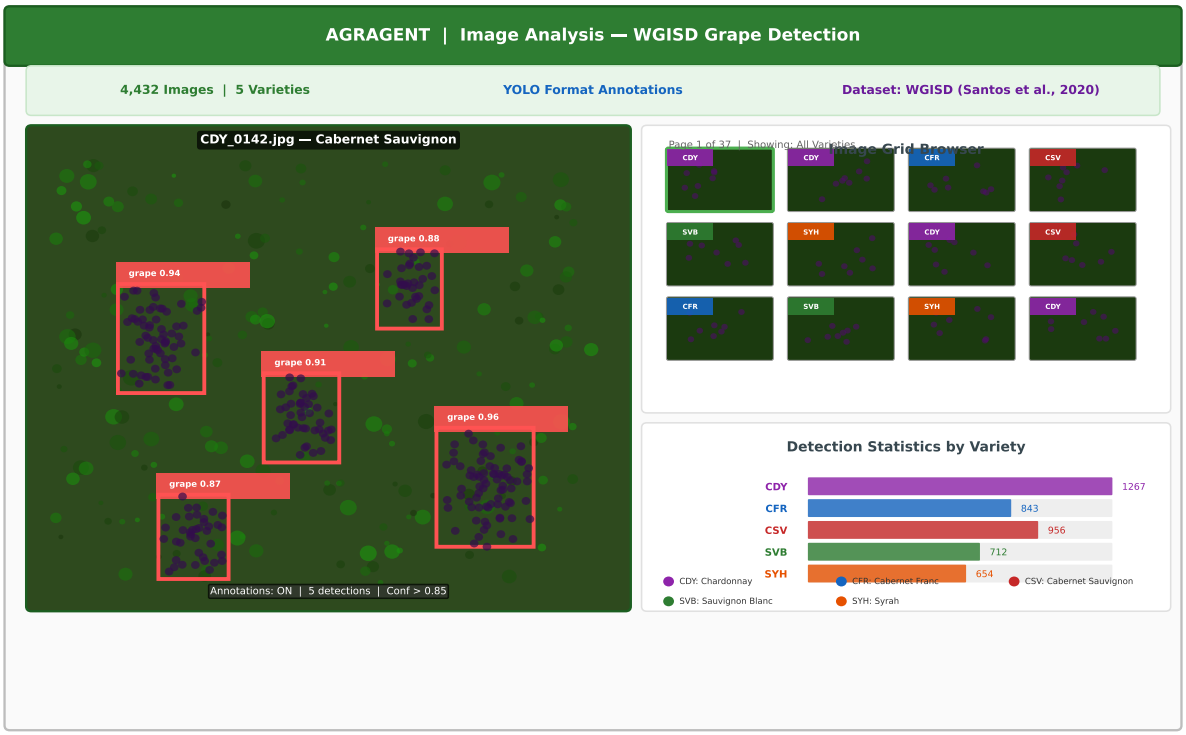

Figure 6: Computer vision module displaying WGISD grape cluster images with YOLO11 bounding box detections for Cabernet Sauvignon.

8 AI-Powered Agronomic Assistant

8.1 Agentic Architecture

The AI assistant (AgrAgent) is implemented as an agentic loop on the Vercel serverless backend. Upon receiving a user message, the system loads the last 20 messages from the conversation history, appends the new message with its application context, and enters

an iterative loop (maximum 10 iterations) in which the LLM can autonomously decide to invoke domain-specific tools. Algorithm 1 describes the procedure.

Algorithm 1 Agentic loop for the AI agronomic assistant.

Require: User message m, conversation ID c, maximum iterations $K = 10$
Ensure: Agent response r, tools used T

$H \leftarrow \textsc{LoadHistory}(c, \text{limit} = 20)$
$M \leftarrow H \cup \{(\textsc{user}, m)\}$
$T \leftarrow \emptyset$
$\textsc{SaveMessage}(\text{c}, \textsc{user}, \text{m})$
for $k = 1$ **to** K **do**
  response $\leftarrow \textsc{Claude.Create}(M, \text{tools} = \textsc{Tools}, \text{system} = \textsc{SystemPrompt})$
  if response.stop_reason = `end_turn` **and** no `tool_use` blocks **then**
    $r \leftarrow$ response.text
    $\textsc{SaveMessage}(\text{c}, \textsc{assistant}, \text{r})$
    return (r, T)
  end if
  for all tool call t in response **do**
    result $\leftarrow \textsc{ExecuteTool}(t.\text{name}, t.\text{input})$
    $T \leftarrow T \cup \{t.\text{name}\}$
    Append (t, result) to M
  end for
end for
return (response.text, T)

8.2 Agent Tools

The agent has access to six domain-specific tools, each defined with a JSON schema that the LLM uses to generate structured function calls. Table 6 describes each tool.

8.3 Agronomic Knowledge Base

The irrigation planning tool uses FAO-56 crop coefficients (K_c) [Allen et al., 1998] to convert reference evapotranspiration into crop-specific water demand:

$$\text{ET}_c = K_c \times \text{ET}_0 \tag{6}$$

Table 7 presents the K_c values for selected crops implemented in the system, organized by growth stage.

The irrigation tool also incorporates soil available water capacity (AWC) by texture class (sandy: 80 mm/m, loam: 160 mm/m, clay: 200 mm/m) and irrigation system efficiency (drip: 90%, sprinkler: 78%, furrow: 55%, flood: 45%) to compute gross irrigation

Table 6: Agent tools available to the AI agronomic assistant.

Tool	Description	Data Source
`get_climate_data`	Fetches 7-day weather forecast and current conditions for given coordinates	Open-Meteo API
`get_ndvi_data`	Retrieves Sentinel-2 NDVI statistics (mean, min, max, std) for a location and date range	Google Earth Engine
`analyze_soil_report`	Interprets soil analysis parameters (pH, EC, OM, N, P, K, Ca, Mg, micronutrients) with crop-specific recommendations	Domain knowledge
`analyze_foliar_report`	Interprets leaf tissue analysis against sufficiency ranges, detects deficiencies and toxicities	Domain knowledge
`calculate_irrigation_plan`	Computes net and gross irrigation depth, application frequency, and weekly schedule using $\mathrm{ET}_c = K_c \times \mathrm{ET}_0$	FAO-56 coefficients
`calculate_fertilization_plan`	Generates NPK requirements based on yield target, splits into base, side-dressing, and foliar applications	Crop uptake tables

depth:

$$I_{\text{gross}} = \frac{\mathrm{ET}_c}{\eta} \tag{7}$$

where η is the system efficiency.

The fertilization tool uses crop-specific NPK uptake rates per ton of yield to calculate total nutrient requirements, distributed across base application (40–60%), side-dressing (30–40%), and foliar supplementation (10–20%).

8.4 Context Injection

Every user message is automatically enriched with the current application state before being sent to the LLM. The context block includes:

- The active section the user is viewing (e.g., Dashboard, Climate, Maps).
- Field information: location name, area (ha), polygon coordinates, centroid.
- Climate KPIs: GDD, chill hours, frost days, heat waves, humidity, ET_0, solar radiation.
- Water balance: precipitation vs. ET_0 with deficit/surplus status.

Table 7: FAO-56 crop coefficients (K_c) by growth stage implemented in AGRAGENT.

Crop	Initial	Development	Mid-season	Late season
Tomato	0.60	0.75	1.15	0.80
Potato	0.50	0.75	1.15	0.75
Corn	0.30	0.70	1.20	0.50
Wheat	0.30	0.70	1.15	0.40
Soybean	0.40	0.70	1.15	0.50
Onion	0.70	0.85	1.05	0.75
Garlic	0.70	0.85	1.00	0.70
Pepper	0.60	0.75	1.05	0.90
Lettuce	0.70	0.80	1.00	0.95
Grapevine	0.30	0.60	0.85	0.45

- Active alerts: frost risk, drought, heat stress, low GDD, low chill, high ET_0.
- Satellite status: GEE connection state, active vegetation index, available scenes.
- Yield prediction: predicted value if available.

This context injection enables the assistant to provide recommendations grounded in the user's actual field conditions without requiring them to re-state information already visible in the interface.

8.5 Conversation Management

The assistant is accessible through two synchronized interfaces: a floating widget (bottom-right button on any screen) and a dedicated full-page section. Conversations are persisted in the database, supporting creation, renaming, deletion, and history browsing. Messages are rendered with Markdown formatting and tool-use badges indicating which tools were invoked during the response generation.

9 Yield Prediction Module

The yield prediction module implements a machine learning pipeline that combines satellite-derived vegetation indices, climate indicators, and historical yield records to produce per-parcel yield forecasts with uncertainty quantification [Altimiras et al., 2025].

9.1 Feature Engineering

The feature set integrates three data modalities into a unified tabular representation for each parcel–season pair:

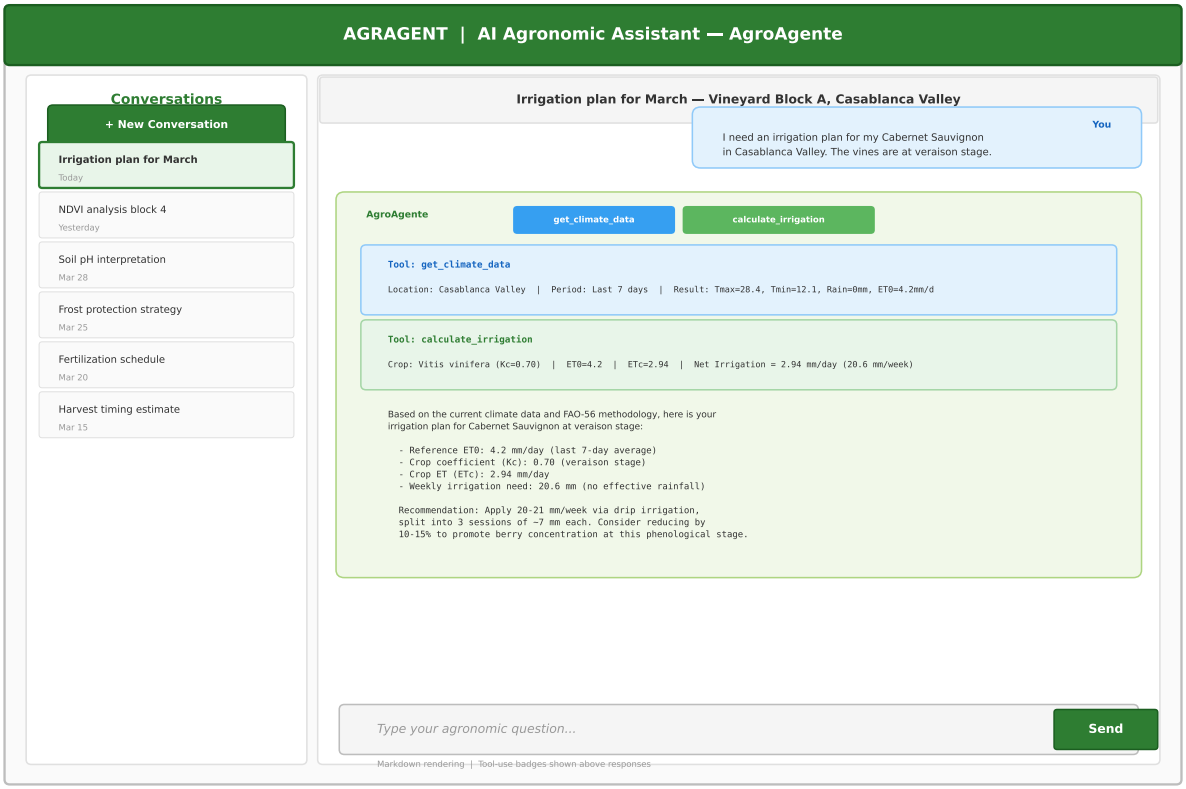

Figure 7: AI agronomic assistant demonstrating a multi-tool interaction: the agent autonomously retrieves climate data and computes an irrigation plan based on the user's field conditions.

- **Satellite features**: Temporal statistics (mean, median, standard deviation, maximum, and minimum) of four vegetation indices—NDVI, NDRE, MSAVI, and TCARI—extracted from cloud-free Sentinel-2 composites during the growing season (September–March). These capture canopy vigor dynamics, chlorophyll content evolution, and stress indicators throughout the crop cycle.
- **Climate features**: Seasonal aggregates derived from the climate module, including cumulative GDD, total precipitation, mean and maximum temperature, chill hours, number of frost days, heat wave events, cumulative ET_0, water balance (precipitation minus ET_0), and total solar radiation.
- **Temporal features**: Season identifier and month-of-prediction to capture inter-annual trends and enable intra-season forecasting at different lead times.

This multi-modal feature engineering yields a total of 32 input variables per observation, enabling the model to capture interactions between remote sensing phenology signals and agrometeorological conditions.

9.2 Model Selection and Training

Five regression algorithms were evaluated: Extra Trees Regressor, CatBoost, Random Forest, XGBoost, and Support Vector Regression (SVR). All models were trained using 5-fold stratified cross-validation with temporal blocking to prevent data leakage between seasons. Hyperparameter optimization was performed via randomized search with 100 iterations per model.

The Extra Trees Regressor was configured with 500 estimators, no maximum depth constraint (fully grown trees), a minimum of 2 samples per leaf, and bootstrap sampling disabled (the default for Extra Trees, which instead uses random thresholds on all samples). Unlike Random Forest, Extra Trees introduces additional randomness by selecting split thresholds at random rather than searching for the optimal split, which reduces variance and improves generalization on high-dimensional multi-modal inputs.

Table 8 compares the performance of the five candidate models.

Table 8: Yield prediction model comparison using 5-fold cross-validation. The Extra Trees Regressor achieves the best performance across all metrics.

Model	R^2	**MAPE (%)**	**RMSE (t/ha)**	**Train Time (s)**
Extra Trees Regressor	**0.972**	**3.8**	**0.41**	2.3
CatBoost	0.969	4.1	0.44	18.7
Random Forest	0.964	4.5	0.47	3.1
XGBoost	0.958	5.0	0.52	6.4
SVR	0.931	6.8	0.66	0.9

The Extra Trees Regressor achieved an R^2 of 0.972, a mean absolute percentage error (MAPE) of 3.8%, and a root mean square error (RMSE) of 0.41 t/ha, outperforming all other candidates. CatBoost and Random Forest performed comparably ($R^2 > 0.96$), while SVR exhibited the weakest performance, likely due to its sensitivity to feature scaling in the multi-modal input space.

9.3 Confidence Intervals and Feature Importance

Prediction uncertainty is quantified by leveraging the ensemble structure of the Extra Trees model: the individual predictions from each of the 500 decision trees are collected, and the 5th and 95th percentiles define the 90% prediction interval. This approach provides non-parametric, instance-specific confidence bounds (e.g., 9.8 ± 0.8 t/ha) without requiring distributional assumptions.

Feature importance is computed using the mean decrease in impurity (MDI) criterion, aggregated across all trees. The top contributing features are precipitation (importance: 0.92), GDD (0.87), maximum temperature (0.81), NDVI (0.74), NDRE (0.65), MSAVI (0.58), and TCARI (0.45). The dominance of climate variables (precipitation, GDD,

temperature) confirms that agrometeorological conditions are the primary yield drivers, while satellite-derived indices provide complementary information about canopy status that further refines predictions.

9.4 Visualization and Interpretation

The module displays three interactive components: (1) a per-parcel yield prediction card with the point estimate, confidence interval, and a qualitative assessment relative to the historical mean; (2) a historical yield chart comparing actual versus predicted values from 2020 to 2025, enabling visual assessment of model accuracy over multiple seasons; and (3) a horizontal bar chart of feature importances ranked by MDI score, helping users understand which environmental factors most influence yield in their specific location.

10 Discussion

10.1 Integration as Contribution

The primary contribution of AGRAGENT is not any individual module but rather their integration into a single, coherent platform. Existing tools require practitioners to switch between GEE code editors for satellite analysis, spreadsheet software for climate records, bioinformatics portals for genomic data, and standalone chatbots for advisory support. AGRAGENT eliminates this fragmentation by enabling data to flow across modules: a polygon uploaded in the satellite module automatically triggers climate data retrieval, dashboard population, and contextualizes subsequent AI assistant interactions. This cross-module data flow—from spatial definition to climate characterization to AI-assisted recommendation—represents a qualitative improvement in the user experience for precision agriculture, applicable to any crop system worldwide.

10.2 Practical Impact

Several design decisions enhance the platform's accessibility:

- **Zero build step**: The frontend is a single HTML file loaded directly by the browser, eliminating Node.js, webpack, or other build toolchain dependencies.
- **Free data sources**: Open-Meteo (no API key) and GEE (free for noncommercial use) ensure that users incur no data costs.
- **Multi-language support**: Full i18n coverage in English, Spanish, and Portuguese addresses the linguistic needs of Latin American, Iberian, and global agricultural communities.

- **Graceful degradation**: The monitoring features (satellite, climate, genomic, image analysis) function independently of the AI backend, allowing standalone deployment on any static file server.
- **Open-source licensing**: The MIT license encourages adoption, modification, and integration into existing agricultural workflows.

10.3 Limitations

Several limitations should be acknowledged:

- The yield prediction model is currently demonstrated with illustrative data; deploying a production model requires site-specific training data and validation.
- The YOLO11-based detection model has been evaluated on the WGISD benchmark; generalization to other grape varieties, training systems, and crops requires additional annotated datasets and domain-specific fine-tuning.
- GEE access is provided via a shared service account; high-traffic deployments may require dedicated quota management.
- The single-page architecture, while eliminating build complexity, may face scalability challenges for very large field management operations.
- The AI assistant requires an Anthropic API key and incurs per-token costs, though the platform functions fully without it.

10.4 Future Work

Planned extensions include: (1) on-device YOLO11 inference using ONNX or TensorFlow.js exported models for edge deployment without server-side GPU requirements; (2) a site-specific yield model training pipeline with user-uploaded historical data; (3) IoT soil sensor integration for real-time soil moisture monitoring; (4) multi-tenant authentication via Supabase Auth with team-based field sharing; (5) a mobile-native application for in-field data collection; and (6) real-time push notifications for critical alerts (frost, heat waves, irrigation deficit).

11 Conclusions

We have presented AGRAGENT, an open-source computational platform for crop monitoring that integrates satellite remote sensing (thirteen Sentinel-2 spectral indices across five functional categories via Google Earth Engine with automatic service account authentication), climate analytics (eight agronomic KPIs with 11-season historical comparison),

genomic data visualization (3,603 DEGs from published RNA-seq data), YOLO11-based computer vision for grape cluster detection ($mAP_{50} = 0.82$ on the WGISD benchmark), machine learning yield prediction (Extra Trees Regressor, $R^2 = 0.972$), and an AI-powered agronomic assistant with six domain-specific tools and context-aware recommendations.

The platform is deployed as a zero-build-step single-page application with multi-language support and is freely available under the MIT license at https://agragent.com.

By unifying heterogeneous data sources into a single, accessible interface, AGRAGENT bridges the gap between isolated research outputs—satellite imagery analysis, transcriptomic studies, crop models—and the operational needs of farmers, agronomists, and agricultural researchers worldwide. Although demonstrated here with viticulture as a primary case study, the platform's satellite, climate, AI assistant, and yield prediction modules are fully crop-agnostic, supporting any crop at any location. We invite the community to contribute to the platform's continued development and to extend its genomic and computer vision modules to additional crop systems.

Acknowledgments

The authors thank the School of Computer Engineering at the Pontificia Universidad Católica de Valparaíso (PUCV) for institutional support. We acknowledge Embrapa for the WGISD dataset, Google Earth Engine for satellite imagery access, Open-Meteo for the free climate data API, and the OpenStreetMap/Nominatim community for geocoding services.

Data Availability

The AGRAGENT platform is available at https://agragent.com. The WGISD dataset is available at https://github.com/thsant/wgisd. RNA-seq data and DEG tables are available as supplementary material in Altimiras et al. [2024]. Climate data is freely accessible via the Open-Meteo API (https://open-meteo.com). Sentinel-2 imagery is available through Google Earth Engine (https://earthengine.google.com).

References

Volodymyr Agafonkin. Leaflet — an open-source JavaScript library for interactive maps, 2023. URL https://leafletjs.com. Version 1.9.4.

Richard G. Allen, Luis S. Pereira, Dirk Raes, and Martin Smith. *Crop evapotranspiration: Guidelines for computing crop water requirements*. Number 56 in FAO Irrigation and

Drainage Paper. Food and Agriculture Organization of the United Nations, Rome, 1998. ISBN 92-5-104219-5. URL https://www.fao.org/3/x0490e/x0490e00.htm.

Francisco Altimiras, Leonardo Pavéz, Alireza Pourreza, Osvaldo Yañez, Lisdelys González-Rodríguez, José García, Claudio Galaz, Andrés Leiva-Araos, and Héctor Allende-Cid. Transcriptome data analysis applied to grapevine growth stage identification. *Agronomy*, 14(3):613, 2024. doi: 10.3390/agronomy14030613.

Francisco Altimiras, Sofía Callejas, Rayner de Ruyt, Natalia Vidal, Astrid Reyes, Mia Elbo, Luis Martí, and Nayat Sánchez-Pi. A computational framework for crop yield estimation and phenological monitoring. In *Progress in Artificial Intelligence, EPIA 2024*, volume 15400 of *Lecture Notes in Computer Science*, pages 161–174. Springer, 2025. doi: 10.1007/978-3-031-80084-9_14.

Evangelos Anastasiou, Athanasios Balafoutis, Nikoleta Darra, Vasileios Psiroukis, Agathi Biniari, Georgios Xanthopoulos, and Spyros Fountas. Satellite and proximal sensing to estimate the yield and quality of table grapes. *Agriculture*, 8(7):94, 2018. doi: 10.3390/agriculture8070094.

Matthias Drusch, Umberto Del Bello, Sébastien Carlier, Olivier Colin, Veronica Fernandez, Ferran Gascon, Bianca Hoersch, Claudia Isola, Paolo Laberinti, Philippe Martimort, et al. Sentinel-2: ESA's optical high-resolution mission for GMES operational services. *Remote Sensing of Environment*, 120:25–36, 2012. doi: 10.1016/j.rse.2011.11.026.

Anatoly Gitelson and Mark N. Merzlyak. Spectral reflectance changes associated with autumn senescence of *Aesculus hippocastanum* L. and *Acer platanoides* L. leaves. *Journal of Plant Physiology*, 143(3):286–292, 1994. doi: 10.1016/S0176-1617(11)81633-0.

Noel Gorelick, Matt Hancher, Mike Dixon, Simon Ilyushchenko, David Thau, and Rebecca Moore. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202:18–27, 2017. doi: 10.1016/j.rse.2017.06.031.

Driss Haboudane, John R. Miller, Nicolas Tremblay, Pablo J. Zarco-Tejada, and Louise Dextraze. Integrated narrow-band vegetation indices for prediction of crop chlorophyll content for application to precision agriculture. *Remote Sensing of Environment*, 81 (2–3):416–426, 2002. doi: 10.1016/S0034-4257(02)00018-4.

Alfredo Huete, Kamel Didan, Tomoaki Miura, E. Patricia Rodriguez, Xiang Gao, and Laerte G. Ferreira. Overview of the radiometric and biophysical performance of the MODIS vegetation indices. *Remote Sensing of Environment*, 83(1–2):195–213, 2002. doi: 10.1016/S0034-4257(02)00096-2.

Olivier Jaillon, Jean-Marc Aury, Benjamin Noel, Alberto Policriti, Christian Clepet, Alberto Casagrande, Nathalie Choisne, Sébastien Aubourg, Nicola Vitulo, Claire Jubin, et al. The grapevine genome sequence suggests ancestral hexaploidization in major angiosperm phyla. *Nature*, 449(7161):463–467, 2007. doi: 10.1038/nature06148.

Mohit Jain, Pratyush Kumar, Ishita Bhatt, Quynh Pham, and Paresh Jain. FarmChat: A conversational agent to answer farmer queries. In *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, volume 2, pages 1–22, 2018. doi: 10.1145/3287048.

Gregory V. Jones, Mark A. White, Owen R. Cooper, and Karl Storchmann. Climate change and global wine quality. *Climatic Change*, 73(3):319–343, 2005. doi: 10.1007/s10584-005-4704-2.

Markus Keller. *The Science of Grapevines: Anatomy and Physiology*. Academic Press, 3rd edition, 2020. doi: 10.1016/C2017-0-03809-9.

Rahima Khanam and Muhammad Hussain. YOLO11 is here: A step-by-step guide to ultralytics YOLO11, 2024. URL https://arxiv.org/abs/2410.17725.

Konstantinos G. Liakos, Patrizia Busato, Dimitrios Moshou, Simon Pearson, and Dionysis Bochtis. Machine learning in agriculture: A review. *Sensors*, 18(8):2674, 2018. doi: 10.3390/s18082674.

Eike Luedeling. Climate change impacts on winter chill for temperate fruit and nut production: A review. *Scientia Horticulturae*, 130(2):357–366, 2011. doi: 10.1016/j.scienta.2011.08.002.

Open-Meteo. Open-meteo – free weather API, 2023. URL https://open-meteo.com. Accessed: 2026-03-15.

Andrew K. Parker, Iñaki García De Cortázar-Atauri, Cornelis Van Leeuwen, and Isabelle Chuine. General phenological model to characterise the timing of flowering and veraison of *Vitis vinifera* L. *Australian Journal of Grape and Wine Research*, 17(2):206–216, 2011. doi: 10.1111/j.1755-0238.2011.00140.x.

Laura Pastonchi, Salvatore Filippo Di Gennaro, Piero Toscano, and Alessandro Matese. Comparison between satellite and ground data with UAV-based information to analyse vineyard spatio-temporal variability. *OENO One*, 54(4):919–934, 2020. doi: 10.20870/oeno-one.2020.54.4.4028.

Jiaguo Qi, Abdelghani Chehbouni, Alfredo R. Huete, Yann H. Kerr, and Soroosh Sorooshian. A modified soil adjusted vegetation index. *Remote Sensing of Environment*, 48(2):119–126, 1994. doi: 10.1016/0034-4257(94)90134-1.

Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788, 2016. doi: 10.1109/CVPR.2016.91.

Mark D. Robinson, Davis J. McCarthy, and Gordon K. Smyth. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, 26(1):139–140, 2010. doi: 10.1093/bioinformatics/btp616.

J. W. Rouse, R. H. Haas, J. A. Schell, and D. W. Deering. Monitoring vegetation systems in the Great Plains with ERTS. *NASA Special Publication*, 351:309–317, 1974.

Thiago T. Santos, Leonardo L. de Souza, Andreza A. dos Santos, and Sandra Avila. Grape detection, segmentation, and tracking using deep neural networks and three-dimensional association. *Computers and Electronics in Agriculture*, 170:105247, 2020. doi: 10.1016/j.compag.2020.105247.

Ranjan Sapkota, Achyut Qin, Manoj Kharel, Chenghao Hu, Jian Cai, and Huafeng Gao. Multi-modal large language models for agriculture: A survey. *Computers and Electronics in Agriculture*, 230:109892, 2025. doi: 10.1016/j.compag.2025.109892.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. In *Advances in Neural Information Processing Systems*, volume 36, pages 68539–68551, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/d842425e4bf79ba039352da0f658a906-Abstract-Conference.html.

Asaf Tzachor, Medha Devare, Brian King, Shahar Avin, and Seán Ó hÉigeartaigh. Large language models and agricultural extension services. *Nature Food*, 4:941–948, 2023. doi: 10.1038/s43016-023-00867-x.

Marie Weiss, Frédéric Jacob, and Grégory Duveiller. Remote sensing for agricultural applications: A meta-review. *Remote Sensing of Environment*, 236:111402, 2020. doi: 10.1016/j.rse.2019.111402.