POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Thesis

Smart Plants: AI-Driven Orchestration for Scalable Plant Care

Laurea Magistrale in Computer Science and Engineering - Ingegneria Informatica

Author: Giulio Saulle

Advisor: Prof. William Fornaciari

Academic year: 2025-2026

1. Context and Motivation

Modern precision agriculture requires more than sensor connectivity: it needs an AI orchestrator able to fuse telemetry, forecasts, and operational rules into coherent control actions [5, 7]. In many real deployments, edge devices are constrained by limited power, unstable wireless links, and fragmented software stacks, so pushing advanced decision logic to each node is brittle. This thesis addresses these constraints through *Smart Plants*, where a Raspberry Pi gateway acts as an AI-driven orchestration core while ESP32 nodes remain lightweight and execution-focused.

The work builds on the previous mobile-centric Smart Plants stack [8], which exposed five critical limitations: (1) a *distributed intelligence barrier* preventing deployment of AI models on resource-constrained ESP32 nodes; (2) *data fragmentation* across ThingSpeak, Firebase, and local logs; (3) *configuration drift* from decentralised node management; (4) a *scalability ceiling* inherent to ad-hoc wiring of individual controllers; and (5) *opacity to operators*, offering limited insight into decision rationale. The proposed architecture shifts state management, scheduling, alerting, and analytics to a middleware layer, where the orchestrator evaluates AI outputs (stress classification, temporal forecasting, anomaly signals) and converts them into explainable, auditable actions across heterogeneous devices (plant sensors, room sensors, camera gateways, irrigation actuators, and external data plugins). The research follows a design science methodology with iterative prototyping and empirical evaluation across industrial (Agricola 2000) and academic (IMEM CNR, Parma) pilots.

2. Objectives and Research Questions

The thesis pursues three formal research questions:

RQ1 How can a centralized orchestration architecture intelligently coordinate heterogeneous plant sensors and actuators, deploy real-time AI models for decision-making, and scale seamlessly from hobbyist to institutional deployments without sacrificing configurability or transparency?

RQ2 Which machine learning techniques (classical, deep learning, and multimodal fusion) most effectively leverage distributed sensor data and visual information to predict plant stress, forecast soil moisture, and optimise watering schedules in uncontrolled greenhouse environments?

RQ3 What data governance, annotation, and experimentation workflows enable continuous improvement of deployed AI models through feedback loops, while maintaining reproducibility and enabling collaborative research across multiple sites?

Six objectives drive the work: (i) design a hub-and-spoke orchestration architecture; (ii) implement a middleware layer with multi-protocol support (serial, MQTT, Wi-Fi), unified persistence, and role-based access; (iii) architect horizontal scalability for dynamic device addition; (iv) integrate ML pipelines for real-time classification, forecasting, and multimodal fusion; (v) validate AI models against Bioristor ground-truth measurements; and (vi) validate the platform through industrial and academic case studies.

3. System Architecture and Platform Design

The implemented system adopts a hub-and-spoke topology replacing the previous mobile-centric stack where each ESP32 published directly to Firebase [8]. The Raspberry Pi gateway hosts a Flask middleware exposing web dashboards, REST APIs (13 endpoints with role-based access control), background schedulers (APScheduler), WebSocket live feeds, and automation services. ESP32 nodes publish telemetry and consume commands over serial or MQTT [1]. Data are persisted in a normalised SQLite schema (sensor readings, device registry, watering events, experiments, alert logs) with nightly Parquet exports for analytics.

Core architectural choices include:

- centralised configuration and scheduling policies;
- plugin-based extensibility enabling new data sources, notification channels, and analytics modules at runtime;
- dynamic device discovery that abstracts sensors and actuators into logical roles, allowing seamless addition of heterogeneous devices;
- separation between firmware responsibilities (acquisition, actuation, local safety) and orchestration responsibilities (logic, policy, observability);
- operational tooling for OTA workflows, logging, recovery, and incident mitigation.

This design improved onboarding speed and reduced operational friction during field trials. Total hardware cost for a complete two-plant installation is approximately €188, with marginal cost of ∼€35 per additional Plant Sensor node.

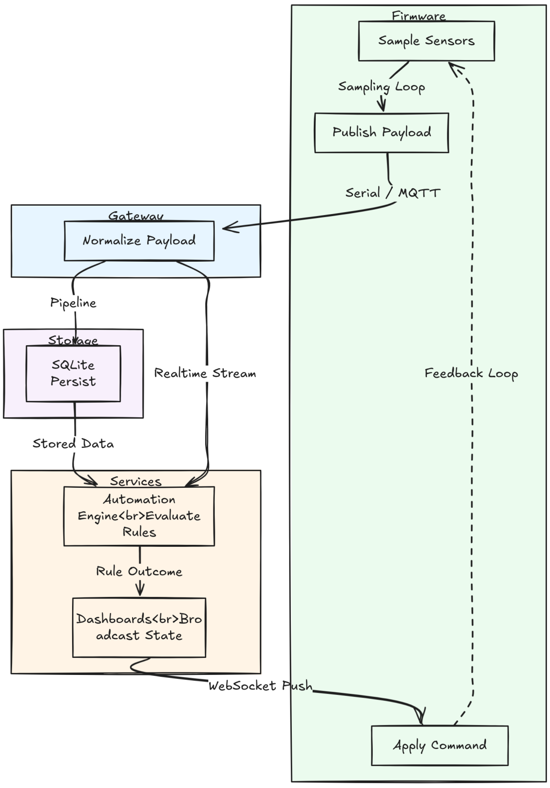

Figure 1: Smart Plants centralized hub-and-spoke architecture used in the thesis deployments.

4. AI Methodology and Ground-Truth Strategy

A central scientific contribution is the use of the Bioristor OECT sensor [4, 6] as physiological ground truth. While commodity probes measure environmental proxies (e.g., soil moisture, temperature, humidity, light), Bioristor measurements capture in-vivo ion flux within the plant's vascular system. The thesis defines binary stress labels from per-plant Bioristor R_{ds} thresholds (mean pre-stress resistance) and compares them against a Normalised Response (NR) variant that expresses fractional deviation from a per-plant baseline. Models are trained to infer stress from commodity inputs alone, so that Bioristor is needed only during a one-time calibration phase.

Model families evaluated span: (i) tree ensembles (Random Forest [3], XGBoost, Gradient

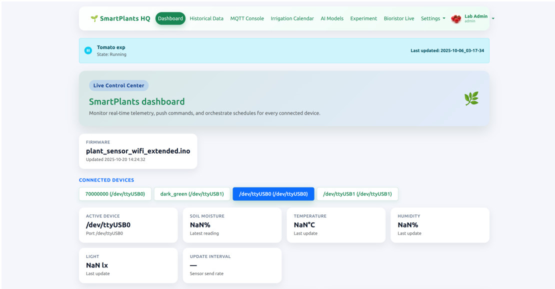

Figure 2: Web dashboard for real-time telemetry, historical charts, and watering schedule management.

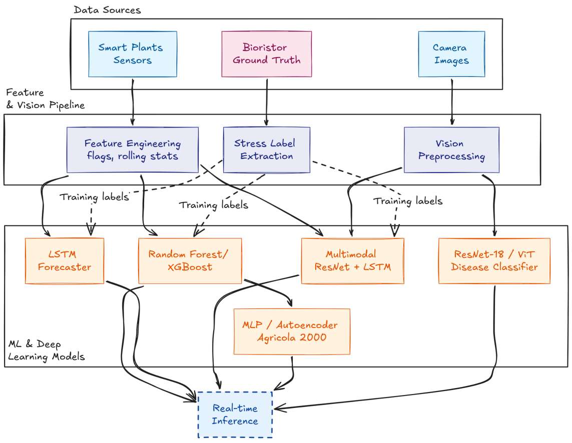

Figure 3: End-to-end Smart Plants pipeline from sensing and orchestration to AI analytics and control actions.

Boosting); (ii) an MLP neural baseline; (iii) a BiLSTM with self-attention for temporal sequences; (iv) a multimodal ResNet-18 + LSTM with CLIP-style contrastive learning; and (v) a dense autoencoder for semi-supervised anomaly detection. In addition, ResNet-18 and ViT-B/16 are compared on the PlantVillage tomato disease subset (10 classes, ∼18 000 images) to assess whether the camera infrastructure can also serve disease identification. Evaluation uses both random stratified splits and deployment-oriented temporal validation (TimeSeriesSplit CV, out-of-time holdout) to avoid optimistic estimates caused by temporal autocorrelation.

5. Experimental Validation and Main Results

5.1. Agricola 2000 Use Case

The Agricola 2000 pilot [8] ran for four weeks in a controlled-environment chamber with two Plant Sensor devices monitoring matched groups of *Lactuca sativa* at 80% and 50% soil-moisture targets. Week 1 focused on calibration (threshold tuning against gravimetric measurements); weeks 2–3 enforced automated schedules; week 4 reintroduced manual interventions. A control group was irrigated by Agricola 2000 operators using traditional weight-based assessment.

Irrigation precision. Across 280 automated cycles and 64 manual overrides, automated control maintained the high-hydration cohort within ±6% of target and the low cohort within ±8%, with a median absolute error of 3.1% versus 5.6% for manual operation. Overshoot events (>15% target) dropped from 15 (manual) to 8 (automated), and mean time-to-target recovery fell from 42.3 to 18.5 min [8].

AI pipeline. The stress-classification pipeline processed 17,188 raw rows (60 s stress-level, 240 s sensor readings), producing 16,643 usable samples after forward-filling and anomaly filtering. Beyond the five raw sensor channels, 36 engineered features (rolling statistics, lag values, temporal encodings, moisture deltas) yielded a 41-feature representation. Three classifiers were compared:

- **Random Forest** (400 trees): macro-F1 0.967;
- **Gradient Boosting** (300 trees): macro-F1 0.975;
- **MLP** (256-128-64, dropout): macro-F1 0.960.

An LSTM forecaster outperformed the Prophet baseline [9] (MAE 3.2% vs. 4.8%), and a dense autoencoder (128-64-16-64-128) trained only on *Optimal* samples achieved AUROC 0.872 for unsupervised anomaly detection.

Temporal validation. Chronological evaluation drastically changed the picture: TimeSeriesCV macro-F1 fell to 0.2709 ± 0.1030 (RF) and 0.6114±0.3817 (GB), with the out-of-time holdout containing only 12 stressed test samples, too few for fair comparison. These results confirm that random-split scores are an optimistic upper bound and that temporally balanced stressed observations are essential for credible prospective claims.

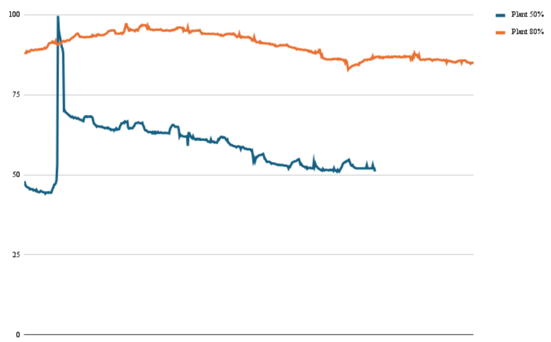

Figure 4: Agricola 2000 pilot: soil-moisture trajectories for 80% and 50% water-stress groups across the four-week protocol.

Figure 5: IMEM Parma experimental setup: Bioristor OECT electrodes implanted in tomato stems alongside Smart Plants sensor nodes.

5.2. IMEM Parma Bioristor Experiment

The IMEM Parma campaign [2, 6] ran for 42 days (October–November 2025) on four cherry-tomato plants (*Solanum lycopersicum*): two instrumented with Bioristor OECT electrodes and irrigated by Smart Plants, and two hand-watered controls. Bioristor electrodes were surgically implanted in the main stem at $\approx 5\,\text{cm}$ height, connected to a Keysight B2902A source-measure unit logging drain-source resistance (R_{ds}) hourly. After a baseline day at field capacity, controlled 8-day drought-irrigation cycles were imposed.

Labelling and dataset. Per-plant binary thresholds were computed as the mean R_{ds} before stress onset: $\theta_{\text{Bio1}} = 0.1497\,\Omega$ and $\theta_{\text{Bio2}} = 0.0611\,\Omega$. After merging, removing missing-Rds rows, and constructing 3 h and 6 h lag features, the dataset contained **7,368 labelled samples** (90.8% `not_stressed`, 9.2% `stressed`) with a 12-feature enhanced panel (4 base sensors + 8 temporal lags). The absolute-threshold scheme outperformed the Normalized Response (NR) variant: RF 5-fold CV accuracy 98.82% vs. 97.57%, macro-F1 0.9713 vs. 0.9597.

Model results. Five architectures were evaluated:

- **Random Forest** (1000 trees, 5-fold CV): accuracy 0.9882 ± 0.0030, macro-F1 0.9713 ± 0.0081, $\kappa = 0.9687$; held-out 80/20: 97.63%.
- **XGBoost**: statistically indistinguishable from RF ($p = 0.8523$, paired t-test).
- **BiLSTM + Attention** (6-step $\times$ 12-feature sequences, 80 epochs): 97.1% accuracy, $\kappa = 0.8382$, 95% stressed-class recall.
- **Multimodal** (ResNet-18 + sensor LSTM, contrastive + CE loss, 3,252 image–sensor pairs): 92.0% accuracy, macro-F1 0.91.
- **MLP** (256-128-64, batch-norm, dropout): 90.8% accuracy baseline.

Explainability. SHAP analysis on the 12-feature panel shows `soil_moisture_lag_6h`, `soil_moisture`, and `soil_moisture_lag_3h` as the top-3 predictors in both RF and XGBoost, followed by light and its lags (4th–6th), while temperature and humidity rank lowest. On the 4-feature base panel, importance is more evenly distributed ($\approx$30% each for temperature and soil moisture, 28% light, 12% humidity). This finding motivates retaining the light sensor even in greenhouse deployments.

Temporal validation gap. Random-split accuracy of 98.8% dropped to $79.2\% \pm 28.3\%$ under temporal splitting ($\Delta \approx 19.6\,\text{pp}$), confirming that deployment-aware evaluation is mandatory. Cross-domain Rds analysis showed the training plants cover a broader stress range (median $0.084\,\Omega$) than the hand-watered controls (median $0.110\,\Omega$), supporting model robustness once temporal leakage is addressed.

6. Contributions

The thesis advances precision agriculture through interconnected contributions:

1. **Orchestration architecture.** A hub-and-spoke gateway with modular services,

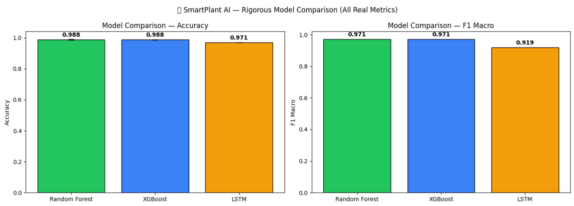

Figure 6: Model accuracy comparison in the IMEM Parma campaign (classical ML, deep learning, and multimodal approaches).

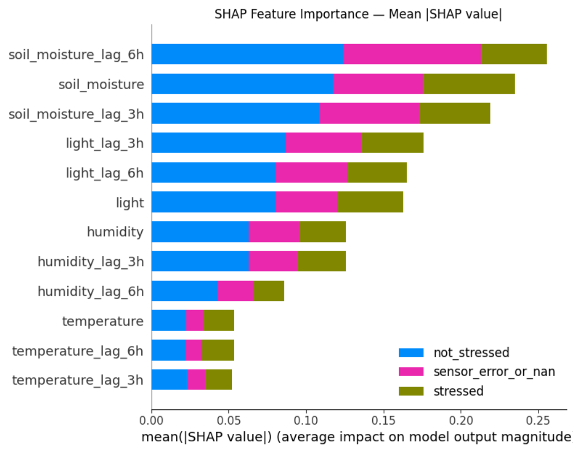

Figure 7: SHAP mean |SHAP| feature importance for the Random Forest model on the IMEM Parma 12-feature panel.

unified persistence, and role-based access, reusable as an IoT plant-monitoring reference.

2. **Bioristor-validated AI framework.** Commodity-sensor models trained with Bioristor-derived labels achieve accurate stress detection without Bioristor at inference, cutting instrumentation cost by an order of magnitude.
3. **AI integration and firmware patterns.** ML embedding methodologies for resource-constrained gateways, plus ESP32 firmware for sensing, actuation, and multi-transport connectivity.
4. **Data governance and operational tooling.** Reproducible notebooks, annotated datasets, deployment playbooks, and security-hardening guides validated across both pilot sites.
5. **Validated evidence.** Empirical results from Agricola 2000 and IMEM confirming that centralised orchestration improves watering precision, enables early stress detection, and scales across mixed environments.

7. Limitations and Future Work

Current limitations include species coverage restricted to tomato and lettuce, class imbalance degrading prospective accuracy, a small Bioristor-instrumented plant count, no supervised cross-domain evaluation, and constrained multimodal dataset size (3 252 pairs). Future work targets:

- cross-season, cross-crop, and cross-domain validation with sensors on control plants;
- online learning and domain adaptation to bridge the temporal generalisation gap;
- automated ADC-to-moisture calibration using substrate-specific curves;
- battery/solar variants, security hardening (OpenVPN, HTTPS, 2FA), and CI/CD pipelines;
- RL-based irrigation policies and SHAP-based explainability dashboards.

8. Conclusions

The thesis demonstrates that centralised AI-driven orchestration is a practical and scalable strategy for precision plant care. The three research questions are answered as follows:

RQ1 Dynamic device registries, USB/MQTT abstraction layers, and automated provisioning workflows unify heterogeneous sensors under a single orchestration plane. The hub maintained >99% uptime over multi-week trials and onboarded new devices without firmware reflashing.

RQ2 Random Forest and XGBoost achieve 98.8% cross-validated accuracy (macro-F1 0.97) on the IMEM binary labelling scheme; BiLSTM adds temporal context (97.1%); multimodal fusion reaches 92.0% by combining RGB imagery with sensor data. Temporal splitting reveals a ∼19.6 pp accuracy gap, identifying deployment-aware evaluation as the primary challenge.

RQ3 Reproducible data pipelines, annotated notebooks, and deployment scripts link field telemetry to deployable intelligence, establishing a practical experimentation workflow validated across Agricola 2000 and

IMEM sites.
Moving orchestration to the Raspberry Pi gateway increased modularity and deployment control, while Bioristor-grounded labelling enabled high-performing classifiers from low-cost inputs alone. Smart Plants is presented as an extensible blueprint for controlled-environment agriculture with validated field utility and clear directions for continued maturation.

9. Acknowledgements

I would like to thank Prof. William Fornaciari for supervision and guidance, Agricola 2000 for industrial collaboration and greenhouse access, and the IMEM CNR laboratory for support in the Bioristor experimental campaign and for providing the physiological ground-truth measurements that grounded the AI models.

References

[1] Manuele Bettelli, Filippo Vurro, Riccardo Pecori, Michela Janni, Nicola Coppedè, Andrea Zappettini, and Daniele Tessera. Classification and forecasting of water stress in tomato plants using bioristor data. *IEEE Access*, 11:34795–34806, 2023. Reference study for water stress classification using bioristor ground truth.

[2] Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

[3] Nicola Coppedè, Michela Janni, Manuele Bettelli, Calogero Leandro Maida, Francesco Gentile, Marco Villani, Roberta Ruotolo, Salvatore Iannotta, Nelson Marmiroli, Marta Marmiroli, and Andrea Zappettini. An in vivo biosensing, biomimetic electrochemical transistor with applications in plant science and precision farming. *Scientific Reports*, 7(1):16217, 2017.

[4] Andreas Kamilaris and Francesc X. Prenafeta-Boldú. Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147:70–90, 2018.

[5] Enrico Masi et al. Bioristor: A plant sensor for continuous and real-time sap monitoring. *Scientific Reports*, 10(1):4002, 2020.

[6] Mqtt version 3.1.1, 2014. Lightweight publish/subscribe messaging transport.

[7] A. G. S. Prem Rajak et al. Internet of things and smart sensors in agriculture: Scopes and challenges. *Journal of Agriculture and Food Research*, 2023.

[8] Gabriele Romano. Smart plants: An easier way to handle your plants. Master's thesis, Politecnico di Milano, Milan, Italy, 2025. Master's thesis collaboratively developed with the Smart Plants team.

[9] Sean J. Taylor and Benjamin Letham. Forecasting at scale. `https://facebook.github.io/prophet/`, 2017. Prophet open-source forecast library.