IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

An Intelligent IoT and Generative AI-Based System for Plant Growth, Soil, and Ground Monitoring in Precision Agriculture

[[1]Vaijanath S. Khilari, [2]Dr. Suresh Halhalli]

[[1]M.Tech Student, [2]Professor, Department of E&TC Engineering, M. S. Bidve Engineering College, Latur, Maharashtra, India]

***Abstract*—**Precision agriculture is moving from isolated sensing and threshold-based automation toward integrated platforms that jointly interpret soil condition, irrigation-water quality, weather variability, crop response, and disease pressure. This paper presents a structured integrative review and an implementable framework for a Plant and Ground Monitoring System (PGMS) that combines calibrated field sensing, edge intelligence, cloud analytics, remote sensing, machine-learning inference, and a multilingual web dashboard. The review synthesizes evidence on soil nutrient monitoring, crop and fertilizer recommendation, evapotranspiration-informed irrigation, image-based plant disease diagnosis, and edge–cloud deployment. The resulting architecture uses an ESP32-S3 field node to acquire soil moisture, soil pH/electrical conductivity/nutrients, irrigation-water pH/electrical conductivity, atmospheric measurements, and leaf imagery; local rules provide safe operation during connectivity loss, while cloud services perform time-series analysis and decision fusion. Reported literature demonstrates approximately 97% crop-classification accuracy in one hybrid MLP–LSTM study, soil-parameter deviations below 13.26% against laboratory measurements, and water savings reported up to 70% in reviewed smart-irrigation deployments. These outcomes are not treated as results of the proposed prototype; instead, they define evidence-based performance targets and a staged validation protocol. The review identifies interoperability, field robustness, calibration drift, uncertainty communication, and secure deployment as the principal barriers to practical adoption. PGMS addresses these gaps through multimodal fusion, confidence-aware recommendations, offline control, and traceable human override.

***Index Terms*—**Climate-smart agriculture, edge computing, Internet of Things, irrigation scheduling, plant disease detection, precision agriculture, soil fertility.

I. INTRODUCTION

Agricultural production is increasingly exposed to interacting stresses rather than isolated constraints. Water scarcity changes nutrient mobility; irrigation-water salinity amplifies osmotic stress; high humidity modifies disease pressure; and a crop that is chemically suitable for a field may still be agronomically unsuitable when rainfall, temperature, or market timing is ignored. Digital agriculture therefore cannot be reduced to a sensor connected to a mobile application. It is a cyber-physical decision problem in which data quality, timing, model uncertainty, control safety, and farmer interpretation are inseparable.

The literature has established the value of data-intensive farming and the Internet of Things (IoT). Smart sensors and cloud platforms can support real-time operational decisions, but data ownership, interoperability, business models, and governance remain unresolved [1]. IoT-based smart agriculture has similarly been described as a convergence of sensing, communication, analytics, and actuation that can improve resource use while introducing challenges related to energy, connectivity, heterogeneity, and security [2]. Machine-learning reviews show successful applications in yield prediction, disease detection, weed identification, livestock management, and water management [3], while deep-learning studies demonstrate that image and time-series models can learn complex representations that are difficult to encode as fixed agronomic rules [4].

Despite this progress, most published prototypes remain functionally narrow. Soil-monitoring studies generally emphasize nutrients and crop recommendation [5], [6]; irrigation reviews focus on soil moisture, weather, and valve control [7]; computer-vision research concentrates on disease classification or localization [11]–[17]; and remote-sensing systems operate at a scale that can miss root-zone chemistry and local irrigation effects [18], [19]. Consequently, a farmer may receive several independent outputs—one for moisture, another for disease, and a third for fertilizer—without a mechanism for resolving conflicts among them.

This paper develops PGMS as an integrated edge–cloud framework rather than claiming a completed multi-season agronomic trial. Its contribution is fourfold. First, it provides a structured synthesis of the engineering evidence relevant to sensing, connectivity, irrigation, soil chemistry, machine learning, and crop imaging. Second, it identifies integration gaps that are obscured when these topics are reviewed separately. Third, it specifies an implementable architecture using an ESP32-S3 edge node, secure cloud services, a Next.js dashboard, and evidence-fusion logic. Fourth, it proposes a publication-grade validation protocol that separates sensor fidelity, network performance, model accuracy, agronomic benefit, and usability. This distinction is important because a system can be accurate in a laboratory yet unsafe, inaccessible, or economically ineffective in the field.

II. REVIEW METHOD AND SCOPE

A. Review Questions

A structured integrative review was used because the research problem spans embedded systems, agronomy, computer vision, remote sensing, cloud software, and human–computer interaction. The review was organized around five questions: (1) which field variables are necessary for actionable crop and fertilizer advice; (2) how should edge and cloud responsibilities be partitioned; (3) which physically grounded and data-driven methods support irrigation and disease decisions; (4) what performance has been reported by representative systems; and (5) which engineering controls are required for safe and trustworthy deployment?

The paper is not presented as a PRISMA meta-analysis. The included sources were selected for technical relevance, traceability, and coverage of the proposed architecture. Peer-reviewed journal articles were prioritized for scientific claims; standards and official documentation were used for hardware, communication, remote-sensing, and software specifications. Recent sources were combined with foundational work where the underlying method remains standard, such as FAO-56 reference evapotranspiration [9] and irrigation-water quality guidance [10].

B. Search and Eligibility Logic

Search concepts included precision agriculture, smart irrigation, soil nutrient monitoring, crop recommendation, fertilizer recommendation, plant disease detection, edge AI, ESP32, evapotranspiration, NDVI, and agricultural dashboards. Eligible studies had to describe an engineering method, dataset, architecture, measurable outcome, or deployment limitation directly relevant to PGMS. Sources were excluded when they repeated unverified marketing claims, lacked sufficient methodological detail, or could not be linked to a stable publisher, DOI, standards body, or official documentation page.

Evidence was synthesized narratively because the underlying studies use different crops, datasets, sensors, climates, definitions of accuracy, and field conditions. A classification accuracy obtained from a curated image dataset is therefore not compared directly with water savings from a farm trial. Instead, results are interpreted within four evidence domains: measurement fidelity, predictive performance, operational efficiency, and deployment readiness. This prevents a single headline metric from being mistaken for system-level validation.

C. Quality Controls

Three controls were applied during manuscript preparation. First, performance values were attributed to the study that reported them and were not presented as original PGMS results. Second, web references were restricted to verified official pages or publisher-hosted articles, and access dates were recorded for dynamic resources. Third, the proposed architecture was checked against device and platform documentation to ensure that stated capabilities—wireless connectivity, data synchronization, dashboard rendering, and model-service integration—are technically plausible [20], [22]–[25].

III. STATE OF THE ART

A. IoT and Edge–Cloud Precision Agriculture

The dominant architecture in current smart-agriculture systems contains a field layer, a communication layer, a processing layer, and a user-facing application. The field layer converts physical processes into data; the communication layer transports records through Wi-Fi, cellular, LoRaWAN, Zigbee, or wired interfaces; the processing layer stores and analyzes time series; and the application layer exposes decisions. The value of this architecture lies in continuous observation, but its weakness is that each additional sensor introduces calibration, synchronization, power, and failure modes.

Cloud-centric systems provide scalable storage and model execution, but they can create unacceptable dependence on network availability. Edge computing addresses this limitation by performing filtering, anomaly checks, compression, and urgent actuation close to the field. The ESP32-S3 is suitable for a low-cost edge node because it integrates Wi-Fi, Bluetooth Low Energy, a dual-core processor, ADC peripherals, and security features on a single device [20]. It cannot replace a server for large-model training, but it can execute deterministic safeguards, compact classifiers, and store-and-forward communication.

A robust partition places safety-critical and time-critical behavior at the edge. For example, a pump should stop when a maximum runtime is exceeded even if the cloud is unavailable. Conversely, historical trend analysis, satellite retrieval, retraining, and cross-farm learning belong in the cloud. PGMS therefore adopts a hybrid architecture in which local rules preserve minimum functionality and cloud services improve context and long-horizon decision quality.

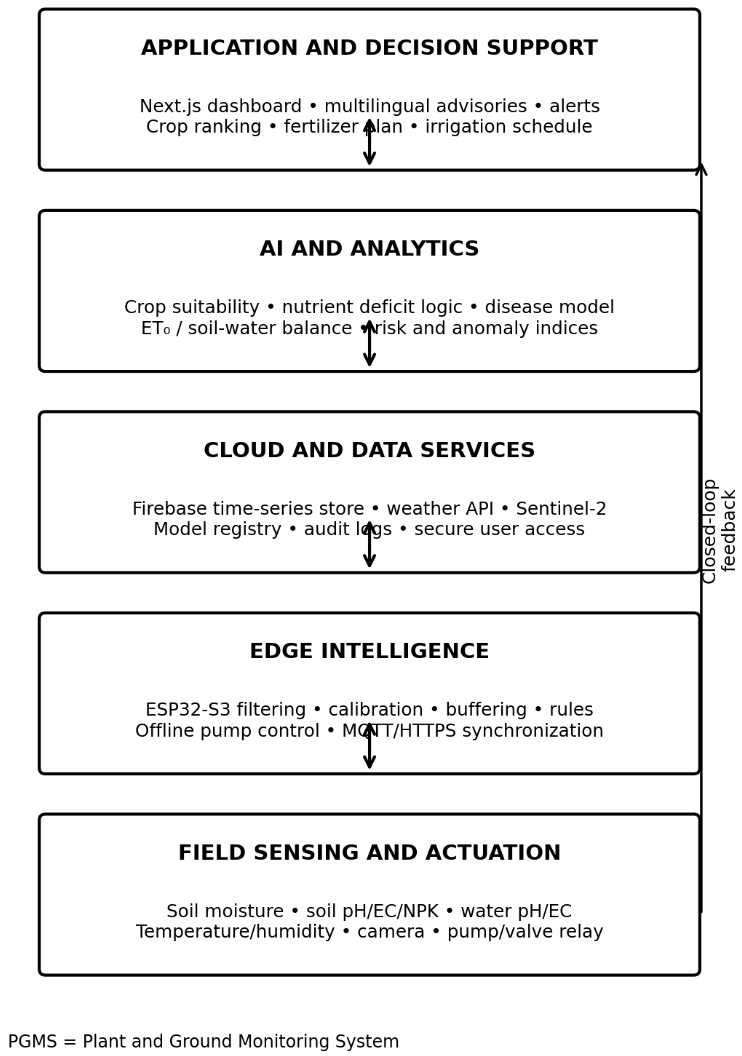

Fig. 1. Layered PGMS architecture and bidirectional decision flow.

B. Soil and Water Sensing

Soil moisture is the most common variable in low-cost irrigation prototypes, yet its interpretation is strongly dependent on soil texture, bulk density, salinity, sensor placement, and calibration. Low-cost sensors should be calibrated against gravimetric or otherwise trusted reference measurements for each representative soil class. Capacitive sensors are generally preferred for long-duration deployment because they avoid the exposed conductive electrodes used by simple resistive probes, but low cost does not guarantee accuracy. Sensor-specific transfer functions and temperature compensation remain necessary [21].

Soil pH and electrical conductivity (EC) provide chemical context. pH influences nutrient availability and microbial processes, whereas EC is an indirect measure of dissolved ionic content and salinity risk. NPK probes can support rapid screening but should not be treated as substitutes for accredited laboratory analysis without local validation. A practical system should report confidence and calibration status rather than presenting every digit returned by a sensor as equally reliable.

Irrigation-water quality is frequently omitted from IoT crop-recommendation systems even though salinity, sodicity, pH, and specific-ion toxicity can alter both crop suitability and fertilizer strategy. FAO guidance emphasizes that irrigation-water assessment must consider infiltration and salinity hazards in relation to soil and crop tolerance [10]. PGMS therefore treats water pH and EC as separate inputs rather than assuming that adequate soil moisture implies acceptable irrigation quality.

C. Crop and Fertilizer Recommendation

Recent systems combine N, P, K, pH, moisture, temperature, and humidity with classification models to recommend crops or identify fertility classes. The IoT system in [5] used soil and atmospheric sensors, MQTT transmission, machine-learning analysis, and a mobile application to generate crop and fertilizer guidance. A 2025 study extended this approach with a hardware testbed and ensemble/deep models, reporting approximately 92% accuracy for Random Forest soil-fertility prediction and 97% overall accuracy for a hybrid MLP–LSTM crop classifier [6]. The same study reported a maximum observed deviation of 13.26% when selected predictions were compared with laboratory measurements [6].

These results support the feasibility of rapid field screening, but they also illustrate the need to separate classification performance from sensor validity. A classifier can perform well on a dataset while the input sensor is biased, or a sensor can be accurate while a recommendation model is poorly calibrated for a local crop. PGMS therefore maintains independent quality indicators for measurement, model, and agronomic rule provenance.

Fertilizer recommendation requires more than mapping a crop to a generic NPK dose. The sequence of intervention matters. When pH is far outside the crop's effective range, additional nutrient application may have low efficiency and increase environmental loading. The proposed logic first identifies constraints such as extreme pH, salinity, or waterlogging; then estimates nutrient deficits relative to the crop and growth stage; and finally generates a staged plan with human-verifiable assumptions.

D. Climate-Aware Irrigation

Threshold irrigation turns a pump on when moisture falls below a fixed value. This is simple but ignores rainfall probability, evaporative demand, crop stage, and hysteresis. Predictive irrigation instead combines root-zone moisture with weather and crop coefficients. The FAO-56 Penman–Monteith method remains a widely accepted reference for computing reference evapotranspiration, ET_0, from net radiation, soil heat flux, air temperature, wind speed, and vapor-pressure deficit [9].

For a daily time step, the reference equation can be expressed as

$$ET_0 = [0.408\Delta(R_n - G) + \gamma(900/(T+273))u_2(e_s-e_a)] \,/\, [\Delta + \gamma(1+0.34u_2)] \qquad (1)$$

where ET_0 is in millimetres per day, Δ is the slope of the saturation vapor-pressure curve, R_n is net radiation, G is soil heat flux, γ is the psychrometric constant, T is mean air temperature, u_2 is wind speed at 2 m, and e_s-e_a is vapor-pressure deficit [9]. Crop evapotranspiration can then be estimated as $ETc = KcET_0$, with crop coefficient Kc selected by growth stage. In low-cost deployments where complete meteorological inputs are unavailable, the system should label substituted or API-derived variables and propagate the associated uncertainty.

A 2025 review of 56 IoT-enabled drip-irrigation studies reported water savings up to 70%, yield increases around 30% in cited cases, and predictive accuracies above 98% for several machine-learning configurations [7]. These numbers reflect heterogeneous studies and should not be interpreted as guaranteed outcomes. They do, however, justify the inclusion of predictive irrigation as a central PGMS function and establish water use, yield, uptime, and false-actuation rate as primary evaluation metrics.

E. Plant Disease and Growth Monitoring

Image-based plant disease detection has progressed from handcrafted color and texture features to convolutional neural networks, object detectors, attention mechanisms, and vision transformers. Foundational PlantVillage work demonstrated the value of large labeled repositories [13], and deep-learning classifiers achieved strong performance on controlled backgrounds [11], [12]. PlantDoc later introduced field-oriented images with complex backgrounds to address the gap between laboratory benchmarks and actual crop scenes [14].

Recent reviews confirm that many models exceed 97% accuracy on curated datasets, while field performance varies more widely because of illumination, occlusion, symptom similarity, cultivar variation, and background clutter [15]–[17]. This domain shift is critical: a system that reports 99% on detached leaves may not be dependable when disease occupies only a small region of a field image. For this reason, PGMS combines image evidence with environmental context and explicitly supports an ‘uncertain—request another image or expert review’ outcome.

OpenCV can support preprocessing, color-space transformations, segmentation, edge detection, and geometric measurements [25]. A lightweight CNN or detector can then estimate disease class and lesion severity. Growth monitoring can use periodic images to derive canopy area, plant height proxies, color indices, or organ counts; however, these measurements require a fixed camera geometry or scale reference. The dashboard screenshots supplied with the prototype demonstrate the intended interaction pattern, but the visual indicators must be experimentally calibrated before they can be treated as quantitative agronomic measurements.

Fig. 2. Grayscale composite of the PGMS prototype dashboard: image analysis, live soil/weather fusion, and advisory output.

F. Remote Sensing and NDVI

Ground sensors provide high temporal resolution at a small number of points; satellite imagery provides spatial coverage at lower temporal resolution and is affected by cloud cover. These modalities are complementary. NDVI is computed from near-infrared and red reflectance and is widely used as a vegetation-greenness indicator [18].

$$\mathrm{NDVI} = (\rho\mathrm{NIR} - \rho\mathrm{Red}) / (\rho\mathrm{NIR} + \rho\mathrm{Red}) \qquad (2)$$

The Sentinel-2 constellation provides 13 spectral bands, including 10 m bands suitable for vegetation analysis, with a nominal five-day revisit under the mission configuration [19]. PGMS uses NDVI as a field-scale corroboration signal, not as a direct diagnosis. A declining NDVI can indicate stress but does not uniquely distinguish drought, nutrient deficiency, disease, harvest, or cloud contamination. The value of multimodal fusion is therefore explanatory: a spatial decline aligned with low moisture supports a water-stress hypothesis, whereas a decline with adequate moisture and image-detected lesions supports a disease hypothesis.

TABLE I
REPRESENTATIVE EVIDENCE USED IN THE PGMS SYNTHESIS

Ref.	Focus	Method/Data	Relevant evidence
[1]	Smart-farming data	Structured review	Data governance, interoperability and platform design remain central.
[5]	Soil nutrients and crop advice	IoT sensors + MQTT + ML	End-to-end monitoring and recommendation pipeline demonstrated.
[6]	Soil fertility and crop prediction	Sensors + RF/Extra Trees + MLP–LSTM	~92% RF fertility accuracy; 97% hybrid crop accuracy; max deviation 13.26%.
[7]	Smart drip irrigation	Review of 56 studies	Reported water savings up to 70%; scale and standardization remain barriers.
[11], [12]	Disease classification	CNN transfer learning	High controlled-dataset accuracy; generalization depends on image conditions.
[14]	Field disease dataset	2,598 field images	Plant Doc targets complex backgrounds and detection use cases.
[15]–[17]	Disease review	Classification, detection, segmentation	Field robustness, lightweight models and multimodal fusion remain priorities.
[19]	Remote sensing	13-band Sentinel-2 imagery	10 m bands and five-day revisit support field-scale vegetation monitoring.
[21]	Low-cost soil sensing	Sensor modeling and calibration	Calibration and environmental dependence are decisive for usable measurements.

IV. PROPOSED PGMS FRAMEWORK

A. System Objectives and Design Principles

PGMS is designed around six principles. First, every recommendation must be traceable to measurements, rules, or models. Second, edge operation must remain safe during network loss. Third, measurement confidence and model confidence must be represented separately. Fourth, recommendations should be crop- and growth-stage-specific. Fifth, the system must support human override and preserve an audit trail. Sixth, the dashboard must communicate uncertainty in language that a farmer or agronomist can act upon.

The framework is intentionally modular. A minimal deployment may include moisture, temperature/humidity, a relay, and a cloud dashboard. A research deployment can add calibrated pH/EC/NPK probes, irrigation-water quality, camera imagery, weather APIs, and Sentinel-2 data. Modularity reduces the risk that a single expensive or unreliable sensor makes the whole system unusable.

B. Hardware Architecture

The reference field node uses an ESP32-S3 with protected power input, watchdog reset, local nonvolatile buffering, and isolated relay or MOSFET control for pumps and valves. Analog sensors are sampled only after stabilization and are separated in time when cross-interference is plausible. The pH interface requires a high-impedance front end and two- or three-point calibration using standard buffers; direct connection of an arbitrary probe output to an ADC is not assumed to be accurate. EC and NPK interfaces must follow the electrical requirements of the selected probe, often through RS-485/Modbus rather than an unconditioned analog pin.

A DHT-class sensor may be used in early prototypes, but field validation should consider a more stable temperature/humidity sensor and a radiation shield. Camera input may be obtained from an ESP32 camera module, a fixed field camera, or a smartphone upload. The camera source is stored with metadata so that model performance can be stratified by acquisition device and environment.

Power architecture depends on deployment duration. For short tests, regulated mains or a power bank is sufficient. For field deployment, a solar panel, charge controller, battery health monitoring, and duty-cycled operation are required. The firmware should persist the last safe actuator state, maximum daily runtime, and manual lockout so that reboot or communication failure cannot produce uncontrolled irrigation.

C. Firmware and Communication

Firmware is divided into acquisition, validation, decision, communication, and diagnostics tasks. Acquisition reads sensors according to a deterministic schedule. Validation rejects impossible ranges, detects stuck values, and applies calibration coefficients. The decision task executes local hysteresis and safety rules. Communication publishes compact timestamped records and receives signed configuration updates. Diagnostics records reboot causes, signal strength, queue depth, and sensor health.

MQTT is appropriate for telemetry because it supports publish–subscribe communication and quality-of-service levels [26]. HTTPS/REST can be retained for configuration, images, and services that do not require persistent messaging. The field message schema should include device ID, farm/plot ID, UTC timestamp, sequence number, firmware version, raw reading, calibrated value, unit, quality flag, and calibration identifier. Sequence numbers allow the server to detect missing or duplicated records.

Offline resilience is implemented through a bounded local queue. When connectivity returns, queued records are uploaded in chronological order while real-time alarms are prioritized. Edge rules are versioned and stored with a checksum. Cloud updates are applied only after validation and can be rolled back. These controls are more important than nominal network latency because a farm system will eventually encounter weak coverage, power cycling, and backend maintenance.

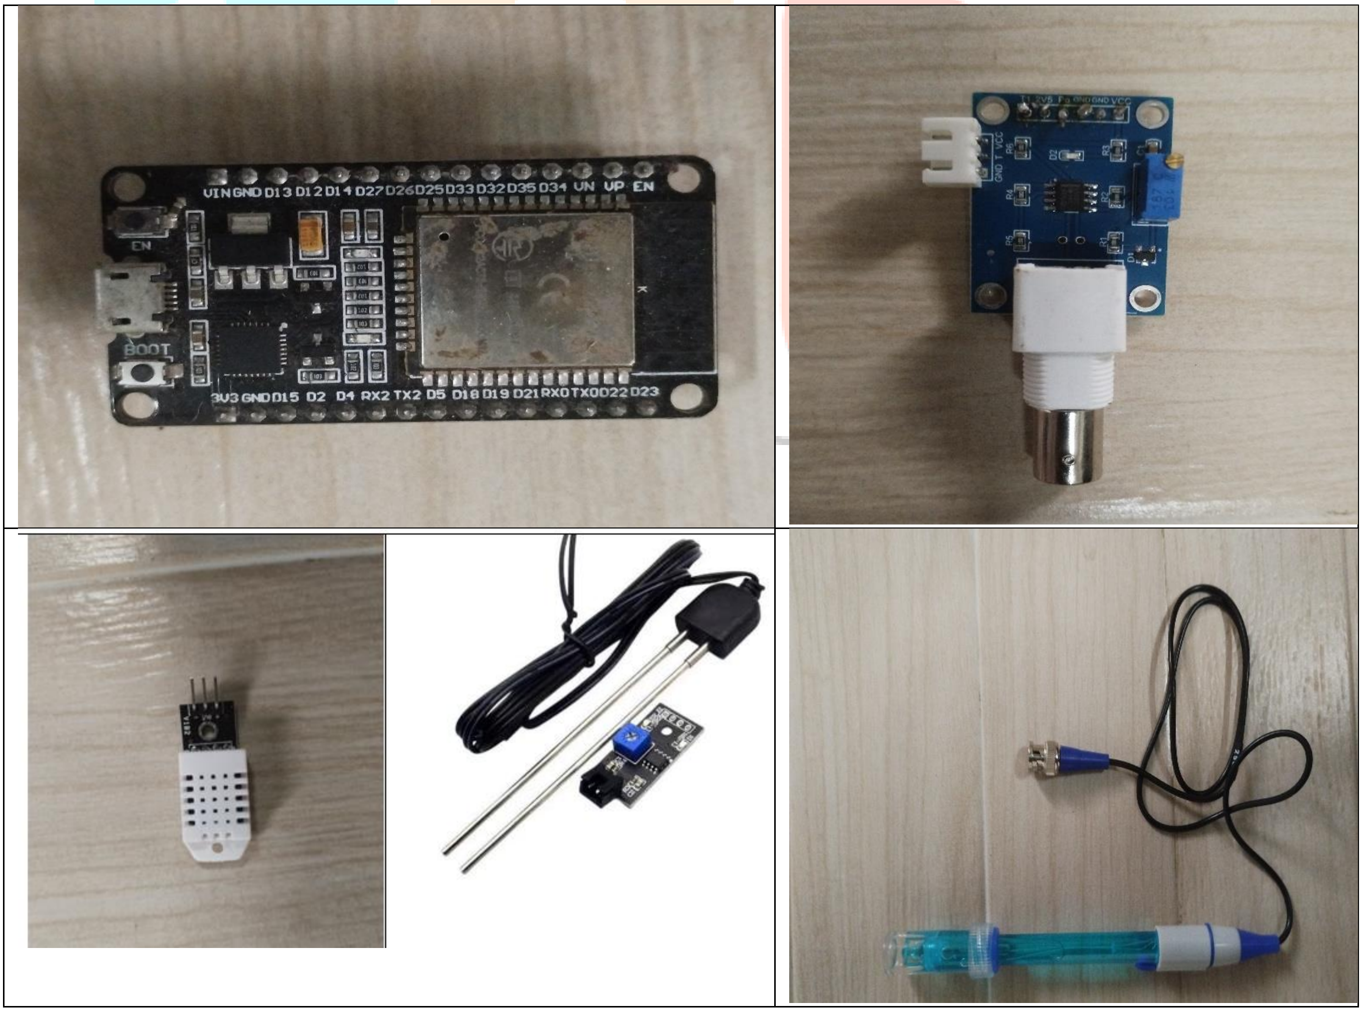

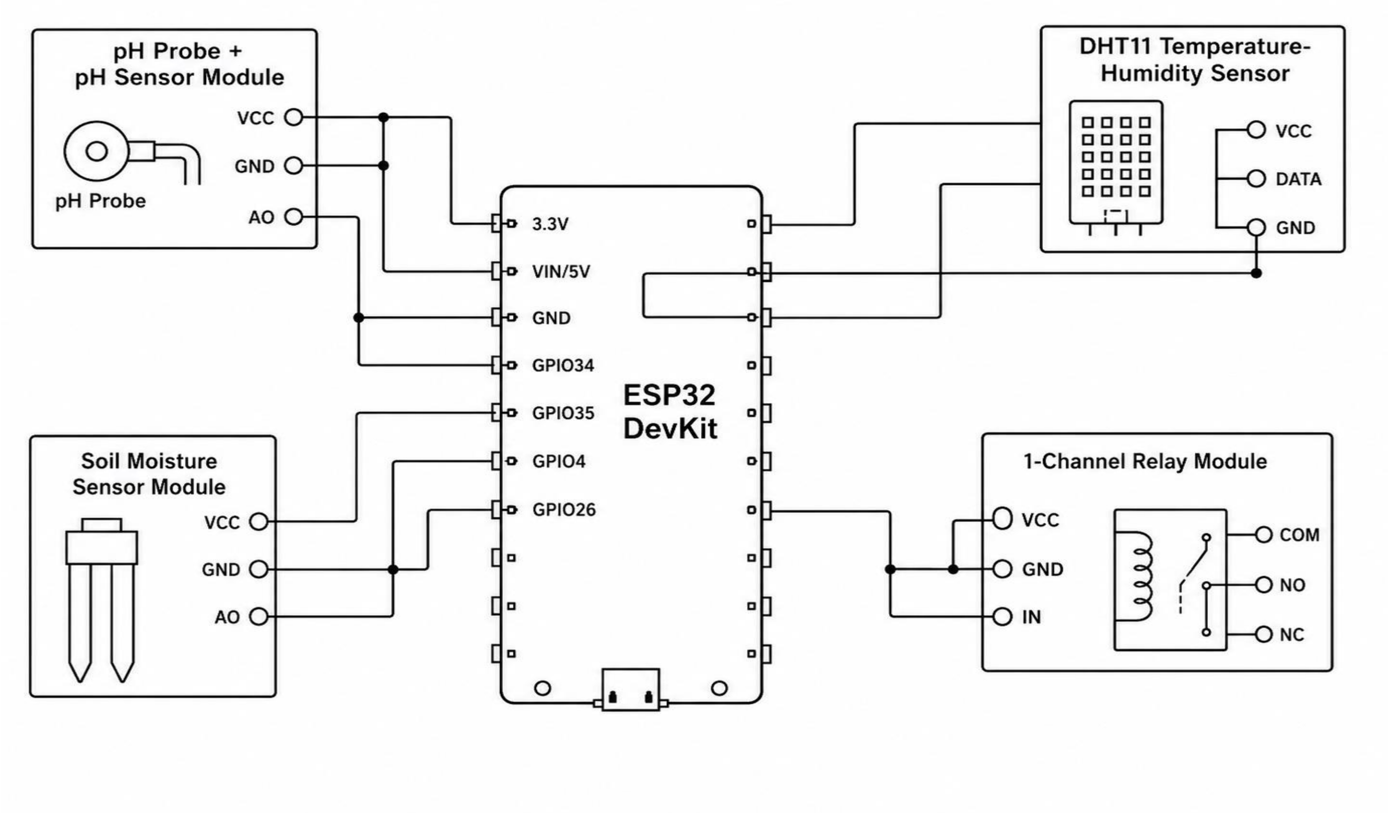

D. Cloud, Dashboard, and AI Services

Firebase Realtime Database or Firestore can provide synchronized cloud storage, authentication integration, and security rules for a prototype [22]. A Next.js application can deliver server-rendered and client-interactive dashboard views [23]. The data model separates raw observations from derived features, recommendations, actuator commands, and human acknowledgments. This prevents a later model update from overwriting the evidence used for an earlier decision.

The AI layer contains deterministic agronomic constraints, statistical or machine-learning models, and an optional language-generation service. The Gemini API can transform structured recommendations into multilingual explanations [24], but it must not be the source of unbounded fertilizer or pesticide decisions. The model receives a constrained JSON object and is instructed to restate approved actions, highlight uncertainty, and advise expert consultation when thresholds are exceeded. Deterministic services retain authority over quantities and actuator control.

The dashboard exposes live status, trend charts, device health, calibration age, weather context, crop ranking, fertilizer rationale, disease images, and irrigation events. A recommendation is displayed with contributing variables and a confidence category. Users can accept, postpone, or reject actions and provide reasons; these responses create data for later usability analysis and model improvement.

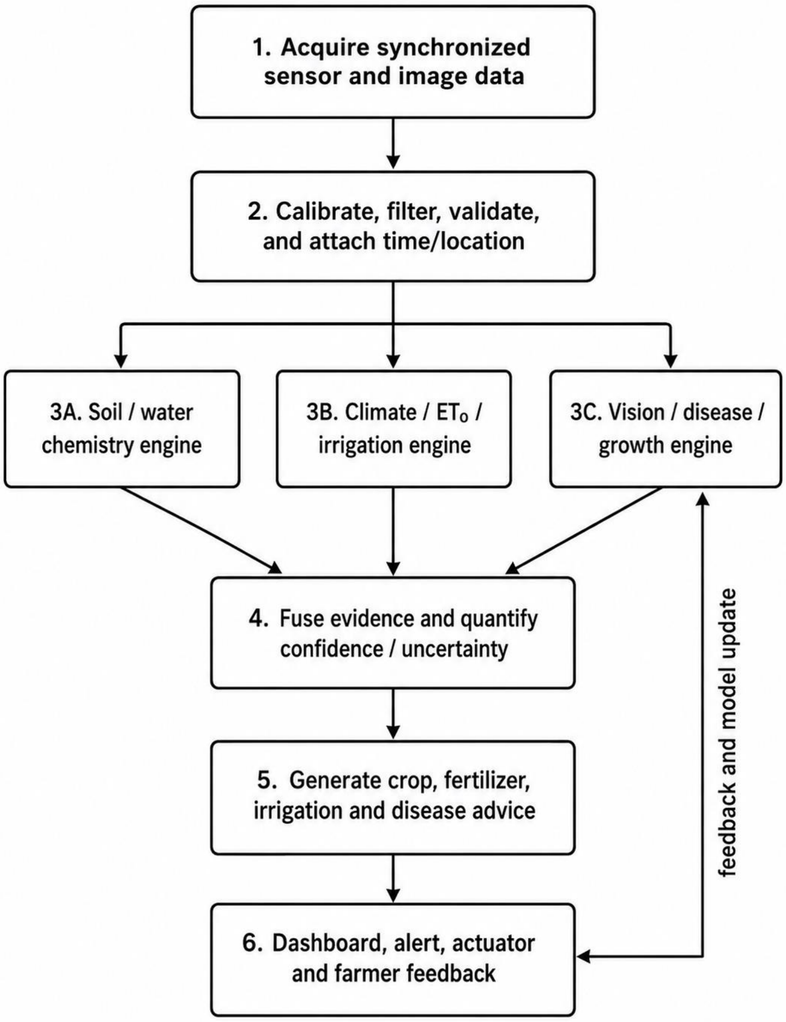

Fig. 3. PGMS data processing and evidence-fusion workflow.

E. Derived Indices and Decision Fusion

PGMS uses interpretable indices to support—but not replace—raw measurements. A normalized water-stress index can combine root-zone moisture deficit with atmospheric demand:

$$WSI = w_1 \cdot clip((\theta target - \theta)/R\theta, 0, 1) + w_2 \cdot clip(ETc/RET, 0, 1) \quad (3)$$

where θ is measured volumetric water content or a calibrated proxy, θtarget is the crop-stage target, $R\theta$ is an acceptable moisture range, ETc is crop evapotranspiration, RET is a normalization constant, and $w_1+w_2=1$. The weights are calibrated from field data rather than fixed universally.

A crop compatibility score ranks candidate crops by normalized distance from preferred pH, moisture, temperature, rainfall, salinity, and nutrient conditions:

$$CScrop = 100 \cdot [1 - \Sigma j\ \alpha j\ dj(xj, \Omega crop,j)] \qquad (4)$$

where dj is a bounded distance between observed variable xj and the crop-specific acceptable interval Ωcrop,j, and αj are importance weights summing to one. Hard constraints such as severe salinity or season incompatibility can override the score. The output is therefore a transparent ranking, not a claim that the highest score guarantees yield.

A disease-risk index combines temperature, humidity, leaf-wetness proxy, recent rainfall, crop susceptibility, and image-model output. When image and environment disagree, the system lowers confidence rather than forcing agreement. A nutrient-lockout flag is raised when pH or EC lies outside crop-specific limits despite apparently sufficient nutrient readings. These mechanisms make conflicts visible and reduce overconfident recommendations.

F. Crop Profiles and Advisory Logic

TABLE II
ILLUSTRATIVE CROP PROFILE FIELDS USED BY THE RULE ENGINE

Crop	Indicative soil pH	Moisture (%)	Nutrient priority
Wheat	6.0–7.5	40–60	High N; moderate P/K
Rice	5.5–6.5	70–90	High N and K
Cotton	6.0–8.0	35–55	Moderate N; high K
Sugarcane	6.0–7.5	55–75	High N and K
Soybean	6.0–7.0	45–65	Low N; high P
Tomato	6.0–6.8	50–70	Moderate N; high K
Onion	6.0–7.0	30–50	Balanced NPK
Chili	6.0–7.0	40–60	Moderate N; high K
Maize	5.8–7.0	50–70	High N; moderate P/K
Groundnut	6.0–7.0	35–55	Low N; high P/Ca

Table II contains indicative profile fields for the initial rule base. These values are not universal fertilizer prescriptions. A production system must localize profiles by cultivar, soil type, season, irrigation method, and extension-service recommendations. The engine stores provenance and effective dates so that an agronomist can replace a generic range with an approved regional profile without modifying firmware.

The fertilizer workflow computes nutrient deficits only after excluding invalid or stale sensor data. It then checks pH and EC constraints, identifies crop stage, and produces a split-application plan. For high-risk conditions—extreme pH, high salinity, uncertain NPK, or suspected disease—the output is a diagnostic action such as laboratory soil testing or expert inspection rather than a precise dose.

G. Security, Privacy, and Safety

Agricultural IoT devices can create physical consequences, so cybersecurity is part of functional safety. NISTIR 8259A identifies core device capabilities including identification, configuration, data protection, logical access control, software update, and cybersecurity-state awareness [27]. PGMS maps these capabilities to unique device credentials, least-privilege database rules, encrypted transport, signed updates, credential rotation, event logging, and revocation of lost devices.

The system stores farm location, crop status, images, and potentially economically sensitive data. Role-based access separates farmer, agronomist, technician, and administrator functions. Public dashboard links are disabled by default, and images are not used for model retraining without explicit consent. Actuator commands require server authorization, sequence checks, and a local safety envelope. Manual emergency stop always overrides cloud commands.

V. EVIDENCE SYNTHESIS AND RESULTS

A. Measurement and Crop-Model Performance

The strongest evidence for integrated soil intelligence comes from systems that compare field sensing with laboratory or structured dataset references. The 2023 system in [5] demonstrates the feasibility of combining moisture, humidity, temperature, and NPK sensing with MQTT and crop/fertilizer recommendations. The 2025 SISFMA study in [6] reported approximately 92% accuracy for Random Forest fertility prediction and 97% overall accuracy for a hybrid MLP–LSTM crop classifier. Its hardware-to-laboratory comparison reported a maximum percentage deviation of 13.26% among selected sample evaluations, with parameter-level deviations of 4.7% for N, 8.6% for P, 6.9% for K, and 5.2% for pH in the reported aggregate table [6].

These results are encouraging but do not eliminate calibration risk. A deviation of 13.26% may be acceptable for screening in one context and unacceptable for a narrow agronomic threshold in another.

PGMS therefore defines acceptance criteria by use case: sensor output may be sufficient for trend detection while a fertilizer quantity still requires laboratory confirmation. The key result of the synthesis is that measurement uncertainty must be connected to decision consequence.

B. Irrigation Outcomes

The irrigation literature reports the most direct resource-efficiency outcomes. The review in [7] found studies reporting water savings from approximately 37% to 70%, and cited cases of yield improvement around 30%. It also summarized machine-learning irrigation classifiers with accuracy up to 99.8%. These values are not directly comparable because datasets, crops, baselines, and accuracy definitions differ. Some results arise from small controlled experiments, while others refer to broader initiatives.

For PGMS, a credible irrigation evaluation must therefore measure applied water volume, soil-moisture response, crop yield or biomass, energy consumption, missed irrigation, unnecessary irrigation, and operator overrides. Model accuracy alone is insufficient; a classifier can be 99% accurate in an imbalanced dataset while still making the few false-positive pump activations that matter most.

C. Disease-Model Performance

Plant-disease reviews document many benchmark accuracies above 97% and several values near 99% [15]. However, the same literature contains substantially lower field results for difficult crops and imaging conditions. The difference is explained by domain shift: plain backgrounds, balanced classes, and well-developed symptoms simplify the benchmark, whereas field images include shadows, overlapping leaves, multiple disorders, and early symptoms.

The synthesis supports three PGMS decisions. First, the disease module should be evaluated on locally captured field images, not only PlantVillage. Second, confidence calibration and rejection performance should be reported alongside accuracy and F1-score. Third, environmental variables should modify the alert priority but must not be used to ‘confirm’ an incorrect image prediction. Multimodal fusion is valuable when it exposes consistency and inconsistency, not when it hides disagreement.

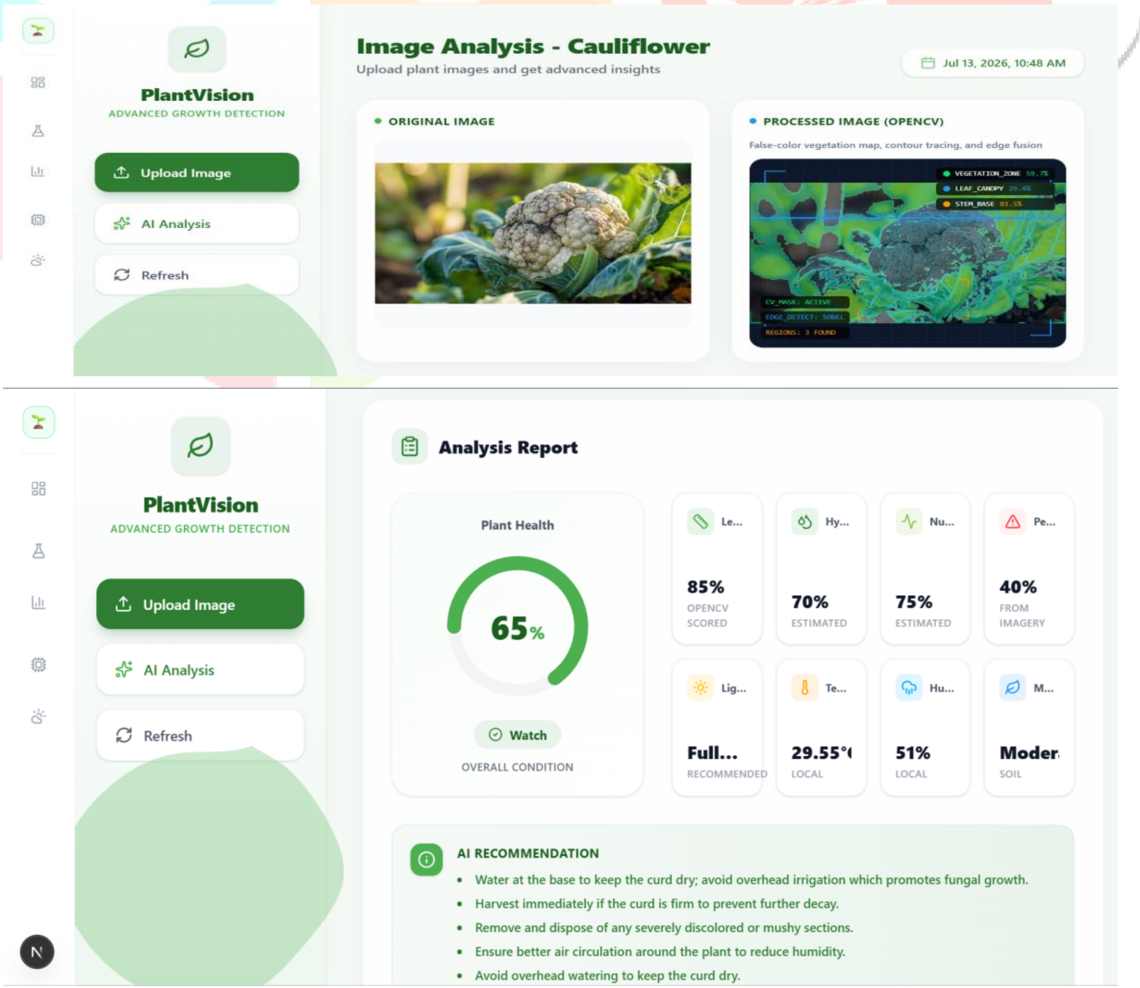

D. System-Level Findings

The review identifies five system-level findings. (1) Soil and water chemistry are rarely integrated into predictive irrigation, even though salinity and pH affect the consequences of irrigation. (2) Disease models and environmental sensing are typically developed separately. (3) cloud dashboards often display sensor values without communicating calibration age or uncertainty. (4) reported machine-learning accuracy is frequently disconnected from agronomic or economic outcomes. (5) offline behavior, cybersecurity, and manual override receive less attention than sensing and classification.

PGMS addresses these findings through a common data model, explicit provenance, confidence-aware fusion, staged intervention, offline-safe control, and user feedback. Its novelty is therefore not a single new classifier. The proposed contribution is a coherent engineering integration in which each subsystem is evaluated on its own terms and then tested as part of a closed-loop decision process.

TABLE III
FUNCTIONAL COMPARISON OF REPRESENTATIVE SYSTEM CLASSES

System	Chemistry	Climate	Disease	Advisory	Interface
Moisture-threshold prototype	Yes	No	No	No	Basic
Soil nutrient + crop ML [5]	Partial	No	No	Yes	Mobile
SISFMA [6]	Partial	Climate vars	No	Yes	Mobile
Smart irrigation review [7]	Moisture	Yes	No	No	Varied
Disease vision systems [11]–[17]	No	No	Yes	No	Research
Proposed PGMS	Soil + water	ET_0 + forecast	Image + context	Crop + fertilizer	Multilingual

VI. VALIDATION PROTOCOL

A. Stage 1: Bench Calibration

Each sensor is evaluated over its operating range against a traceable reference. pH calibration uses at least pH 4 and pH 7 buffers, with pH 10 added when alkaline measurements are expected. Moisture calibration uses oven-dry/gravimetric references or a validated volumetric method across representative soil textures. EC is tested with standard conductivity solutions. Repeatability, hysteresis, temperature dependence, warm-up time, drift, and between-unit variation are recorded.

Acceptance metrics include mean absolute error (MAE), root-mean-square error (RMSE), bias, coefficient of variation, and calibration residuals. A sensor receives a validity interval and recalibration date. The firmware stores calibration coefficients with a version identifier, and every uploaded observation references that version.

B. Stage 2: Controlled Pot Experiment

A controlled pot trial isolates water and nutrient treatments. At minimum, the design should include adequate-water control, water-deficit, over-irrigation, nutrient-deficit, and pH-stress groups with replication. Soil and plant measurements are collected independently from PGMS to prevent circular validation. The irrigation engine is compared with a fixed schedule and a simple threshold baseline.

Disease-model testing uses healthy plants and naturally or safely acquired disease images from the target crops. Images are split by plant or plot rather than by random image to prevent near-duplicate leakage. Augmentation is applied only to training data. Performance is reported by class using precision, recall, F1-score, confusion matrix, calibration error, and rejection rate.

C. Stage 3: Field Trial

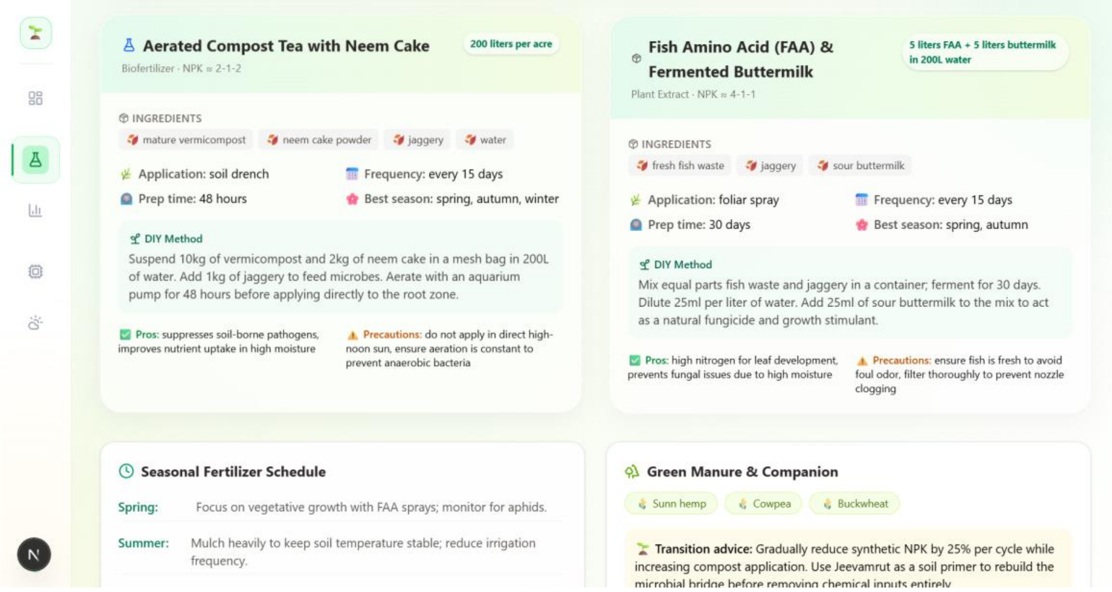

The field phase uses multiple plots and at least one complete crop cycle. Treatments compare conventional practice, sensor-threshold control, and PGMS decision support. Primary outcomes are water applied per unit area, yield, fertilizer input, energy consumption, disease-detection lead time, and labor events. Secondary outcomes include uptime, packet loss, battery performance, false alarms, manual overrides, and user satisfaction.

Statistical analysis should specify the experimental unit, randomization, missing-data handling, and confidence intervals. Agronomic outcomes are assessed with analysis appropriate to the design, while time-series sensor errors use repeated-measures or mixed models where necessary. The paper should report negative or null outcomes and not only best-case days.

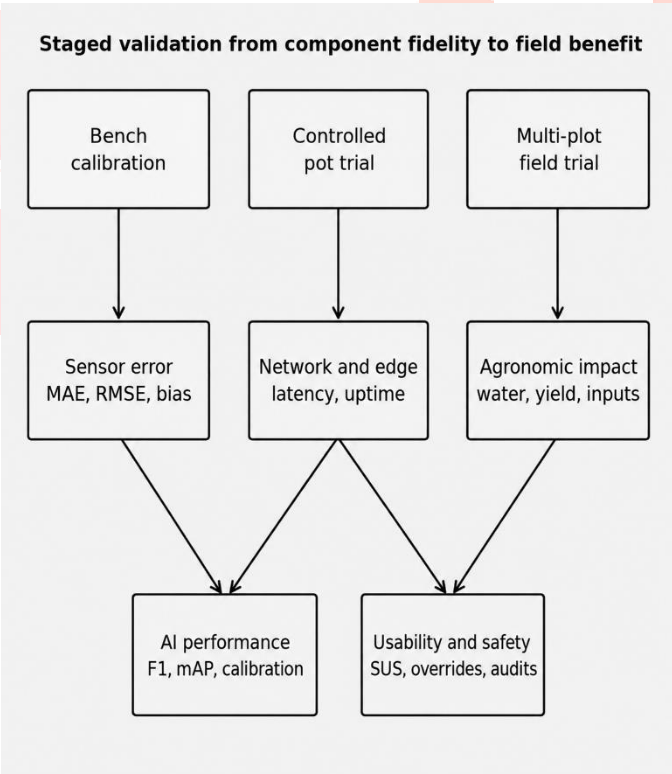

Fig. 4. Staged validation protocol for measurement, computation, agronomic benefit, and usability.

D. Target Metrics

TABLE IV
MINIMUM VALIDATION METRICS FOR A FIELD-READY PGMS

Domain	Metrics	Reference condition
Sensor fidelity	MAE, RMSE, bias, drift	Against laboratory/traceable reference
Communication	Delivery ratio, latency, queue recovery	Online and forced-offline tests
Irrigation	Water used, false activation, soil response	Against threshold and schedule baselines
Domain	**Metrics**	**Reference condition**
Crop model	Macro-F1, top-k accuracy, calibration	External/local dataset
Disease model	Macro-F1, mAP, ECE, rejection	Field images split by plant/plot
System	Uptime, energy, recovery time	Full crop-cycle operation
Usability	SUS score, task completion, overrides	Farmer/agronomist study

Table IV prevents selective reporting. A manuscript claiming a 50 ms dashboard latency, for example, should state whether the value is device-to-cloud, cloud-to-browser, or local rendering time and should report the distribution rather than a single minimum. Similarly, an accuracy target of 95% is meaningful only when class balance, split strategy, confidence interval, and error consequences are specified.

VII. DISCUSSION

A. Why Integration Is Difficult

Integration creates dependencies that are absent in isolated prototypes. Sensor clocks drift, weather APIs use different locations and time zones, images arrive at irregular intervals, and satellite data may be unavailable because of cloud cover. Models are trained at different temporal and spatial scales. A single ‘farm health score’ can therefore conceal incompatible evidence. PGMS avoids premature aggregation by retaining modality-specific features and generating a fused recommendation only after quality checks.

Another difficulty is causal ambiguity. Low canopy vigor may be caused by insufficient water, excess water, salinity, nutrient deficiency, root disease, or a combination. Multimodal data narrows the hypothesis space but does not automatically establish causality. The dashboard must explain which evidence supports an action and which alternative explanations remain.

B. Edge AI and Model Lifecycle

On-device inference can reduce latency, bandwidth, and dependence on the cloud. Lightweight agricultural models have been proposed for resource-constrained devices [28]. Yet model compression can reduce accuracy, and an embedded model can become stale as crop, season, or camera conditions change. PGMS therefore uses edge AI selectively: deterministic safety rules and compact anomaly detectors run locally, while heavier image analysis and retraining remain cloud-based until field evidence justifies migration.

Model lifecycle management includes dataset versioning, training configuration, validation results, deployment approval, rollback, and post-deployment monitoring. A new model is first run in shadow mode, where predictions are logged but do not affect recommendations. Only after agreement and error analysis does it become active. This process is more conservative than directly updating a farm model from an API endpoint, but it is appropriate for a system that influences physical inputs.

C. Explainability and Farmer-Centered Design

Agricultural advice must be concise enough for action and detailed enough for verification. A farmer may need 'Do not irrigate today; rain probability is high and root-zone moisture is adequate,' while an agronomist needs the underlying moisture trend, ET estimate, forecast source, and confidence. PGMS provides layered explanations rather than one universal view.

Multilingual generation improves accessibility but introduces the risk of changing technical meaning. Approved agronomic terms, units, crop names, and warnings should therefore be stored in a controlled glossary. Generated text is checked against structured action fields, and quantities are never created solely by a language model. User studies should evaluate comprehension and decision quality, not merely whether the interface looks attractive.

D. Limitations

This paper is a review and architecture study. It does not claim that the PGMS prototype has achieved the literature-reported accuracies, water savings, or yield improvements. Crop ranges in Table II are illustrative and require regional agronomic validation. Low-cost NPK and pH sensors may be unsuitable for quantitative fertilizer dosing without laboratory calibration. Weather APIs may not represent farm microclimates, and satellite imagery may be temporally incomplete.

The literature synthesis is integrative rather than exhaustive and does not calculate pooled effect sizes. Some recent review articles summarize heterogeneous primary studies, which can propagate inconsistent definitions. Future publication should supplement this framework with registered experimental protocols, raw datasets, calibration files, source code, and independent field replication.

E. Future Research

Priority research directions include: local calibration libraries for Indian soil classes; multimodal disease datasets pairing images with weather and soil records; uncertainty-aware crop ranking; causal evaluation of fertilizer interventions; low-power wide-area communication; federated learning that protects farm data; digital twins for testing irrigation policies; and economic analysis that includes maintenance and calibration labor.

A second priority is interoperability. The long-term value of PGMS depends on documented schemas and exportable data, not a single cloud vendor. Open interfaces allow a farmer or institution to migrate storage, compare models, and preserve years of measurements. The architecture should therefore treat Firebase, Next.js, and Gemini as replaceable implementations of storage, interface, and language services rather than immutable scientific components.

VIII. CONCLUSION

This paper synthesized current work on smart-farm sensing, soil fertility and crop recommendation, predictive irrigation, plant disease detection, remote sensing, and edge–cloud deployment, and translated the evidence into a coherent PGMS architecture. The literature shows that strong component-level performance is possible: representative studies report high crop-classification accuracy, useful agreement between sensor systems and laboratory measurements, and substantial irrigation savings. At the same time, these outcomes are heterogeneous and do not by themselves validate an integrated farm platform.

The proposed PGMS contribution is a traceable multimodal decision pipeline. Soil and irrigation-water chemistry constrain crop and fertilizer advice; weather and FAO-56 evapotranspiration inform irrigation; images and environmental context support disease and growth interpretation; edge rules preserve safe function during connectivity loss; and the dashboard communicates evidence, confidence, and override options. The staged validation protocol separates sensor fidelity, network behavior, machine-learning performance, agronomic benefit, usability, and security. A future field paper can therefore report genuine PGMS results without confusing prototype screenshots or literature benchmarks with experimental evidence.

Publication readiness ultimately depends on disciplined validation. The next step is a replicated crop-cycle trial with calibrated sensors, pre-specified baselines, local field imagery, recorded water and fertilizer inputs, and transparent reporting of failures. Under those conditions, PGMS can be evaluated as a practical climate-smart agricultural decision-support system rather than only as a collection of promising technologies.

References

[1] K. Singh, K. A. Momin, M. Nishal, C. Sultania, and M. Rao, “Agri-Guard: IoT-Based Network for Agricultural Health Monitoring with Fault Detection,” in Proc. 10th Int. Conf. Internet of Things, Big Data and Security (IoTBDS), 2025, pp. 223–230, doi: 10.5220/0013209500003944.

[2] M. Ayaz, M. Ammad-Uddin, Z. Sharif, A. Mansour, and E.-H. M. Aggoune, “Internet-of-Things (IoT)-Based Smart Agriculture: Toward Making the Fields Talk,” IEEE Access, vol. 7, pp. 129551–129583, 2019, doi: 10.1109/ACCESS.2019.2932609.

[3] Y. Kalyani and R. Collier, “A Systematic Survey on the Role of Cloud, Fog, and Edge Computing Combination in Smart Agriculture,” Sensors, vol. 21, no. 17, Art. no. 5922, 2021, doi: 10.3390/s21175922.

[4] S. Adla, N. K. Rai, S. H. Karumanchi, S. Tripathi, M. Disse, and S. Pande, “Laboratory Calibration and Performance Evaluation of Low-Cost Capacitive and Very Low-Cost Resistive Soil Moisture Sensors,” Sensors, vol. 20, no. 2, Art. no. 363, 2020, doi: 10.3390/s20020363.

[5] P. Placidi, L. Gasperini, A. Grassi, M. Cecconi, and A. Scorzoni, “Characterization of Low-Cost Capacitive Soil Moisture Sensors for IoT Networks,” Sensors, vol. 20, no. 12, Art. no. 3585, 2020, doi: 10.3390/s20123585.

[6] Espressif Systems, ESP32-S3 Series Datasheet, ver. 2.2, Mar. 2026.

[7] R. G. Allen, L. S. Pereira, D. Raes, and M. Smith, Crop Evapotranspiration—Guidelines for Computing Crop Water Requirements, FAO Irrigation and Drainage Paper 56. Rome, Italy: Food and Agriculture Organization of the United Nations, 1998. [Online].

[8] W. Jia, Y. Zhang, Z. Wei, Z. Zheng, and P. Xie, “Daily Reference Evapotranspiration Prediction for Irrigation Scheduling Decisions Based on the Hybrid PSO-LSTM Model,” PLOS ONE, vol. 18, no. 4, Art. no. e0281478, 2023, doi: 10.1371/journal.pone.0281478.

[9] M. Drusch et al., “Sentinel-2: ESA’s Optical High-Resolution Mission for GMES Operational Services,” Remote Sens. Environ., vol. 120, pp. 25–36, 2012, doi: 10.1016/j.rse.2011.11.026.

[10] R. Hassanpour, A. Majnooni-Heris, A. Fakheri Fard, and J. Verrelst, “Monitoring Biophysical Variables and Cropland Dynamics at Field Scale Using Sentinel-2 Time Series,” Remote Sens., vol. 16, no. 13, Art. no. 2284, 2024, doi: 10.3390/rs16132284.

[11] K. G. Liakos, P. Busato, D. Moshou, S. Pearson, and D. Bochtis, “Machine Learning in Agriculture: A Review,” Sensors, vol. 18, no. 8, Art. no. 2674, 2018, doi: 10.3390/s18082674.

[12] S. P. Mohanty, D. P. Hughes, and M. Salathé, “Using Deep Learning for Image-Based Plant Disease Detection,” Front. Plant Sci., vol. 7, Art. no. 1419, 2016, doi: 10.3389/fpls.2016.01419.

[13] W. Zhang and A. Yadav, “Generative AI as an Interface Between Farmers and AI Algorithms: Opportunities and Challenges,” J. Indian Inst. Sci., vol. 105, pp. 305–321, 2025, doi: 10.1007/s41745-025-00490-8.

[14] Google AI for Developers, “Gemini API Documentation,” 2026.

[15] Google Firebase, “Firebase Realtime Database Documentation,” 2026.