

Smart Agriculture for Hilly and Terraced Farming: A Literature Review and Decision-Support System Design

Nandini K N¹, M Varshitha², Shivam Kumar² & Suneha D Rati²

¹Assistant Professor, Er. Sri Sairam College of Engineering, Bengaluru, Karnataka

²Final Year, B.E (ISE), Er. Sri Sairam College of Engineering, Bengaluru, Karnataka

Department of Information Science and Engineering

DOI: doi.org/10.34293/iejcsa.v4i4.164

Abstract - Hilly agriculture presents conditions that are not adequately represented by uniform-field assumptions, including slope-dependent water movement, fragmented terraces, elevation variation, and difficult physical access. This paper reviews fifteen verified scholarly sources, including eight focal works on agricultural sensing, irrigation scheduling, machine learning, remote sensing, unmanned aerial systems, and conversational access. Seven supporting sources provide context on precision agriculture, data-intensive farming, decision-support adoption, digital-agriculture governance, and edge computing. The synthesis identifies design needs concerning input reliability, spatial interpretation, terrain-sensitive analysis, user access, and intermittent connectivity. These needs inform an Intelligent Smart Agriculture Platform for Hilly Cultivation that combines quality-aware ingestion, structured storage, Python analytics, FastAPI services, constrained LangChain/Gemini orchestration, and React-based chart and contour-map presentation. Numerical results originate from deterministic retrieval and analytics rather than the language model. The paper specifies a conceptual analytical framework and prospective evaluation protocol without claiming field validation, measured performance, or agronomic benefit.

Keywords: Smart agriculture, hilly-terrain cultivation, terrace farming, soil-erosion prediction, precision irrigation, conversational AI agents, GIS terrace mapping, crop-stress detection.

INTRODUCTION

Crop cultivation on hill and mountain slopes is constrained by irregular terrain, fragmented terraces, and elevation-sensitive microclimates. Manual inspection can be physically demanding and intermittent, while a condition observed at one location may not represent another elevation or orientation. The operational question is therefore not merely whether a farm is stressed, but which terrace requires attention and which observations support that priority.

Precision agriculture links sensors, information systems, and site-specific management to variability within agricultural production [1]. Data-intensive farming extends this pattern by integrating heterogeneous observations into analytical and decision processes [2]. IoT architectures enable distributed monitoring, but their practical value depends on sensor reliability, communication, storage, and interpretation rather than connectivity alone [3]. Agricultural decision-support research also shows that technical capability does not by itself ensure a useful tool; delivery, relevance, and fit with farmers'

decision practices affect adoption [4]. This matters in hilly regions, where a numerical estimate may be difficult to act on unless it is related to a mapped terrace, current data quality, and a feasible inspection route.

The reviewed literature provides mature component knowledge but limited integration around terrace-level hill agriculture. This paper contributes a structured synthesis, a corpus-scoped gap analysis, a literature-to-design mapping, the architecture of an Intelligent Smart Agriculture Platform, and a prospective evaluation framework. The platform combines Python analytics, FastAPI, LangChain with Gemini 1.5 Flash, SQLite, React.js, Recharts, and React-Leaflet. The contribution remains a design: no deployment, farmer trial, predictive score, or resource saving is reported.

A. Objectives of This Review

The review organizes sources by technical emphasis, compares their scope and limitations, identifies issues relevant to terrace monitoring, and converts those issues into requirements that can later be tested.

B. Technical Evolution of Smart Agriculture in Hilly Regions

Agricultural IoT research surveys sensing, communication, and application-layer concerns [5], while irrigation-scheduling literature distinguishes the strengths and limitations of plant- and environment-based indicators [6]. Soil-moisture measurement involves tradeoffs in scale, calibration, cost, and interpretation [7]. Machine-learning reviews show broad agricultural use but also establish that model suitability depends on data and task definition [8]. These strands motivate a pipeline in which sensing, validation, analysis, and communication remain separable.

C. Field Sensing and Terrain Visualization

Remote-sensing reviews describe the value of spatially resolved observations and the continuing importance of scale and timing [9]. Small unmanned aerial systems can provide high-resolution imagery for agricultural tasks, although acquisition, processing, regulation, and operational cost remain constraints [10]. For terraced land, these methods support mapping field condition; they do not establish that a model developed on flat fields will generalize to slopes.

D. The Farmer as Primary End User

Digital-agriculture research frames adoption as a socio-technical issue involving skills, institutions, trust, and work practices [11], while smart-farming digitization spans sensing, communication, computation, and applications [12]. Accordingly, a farmer-facing interface must expose evidence, not only predictions. Natural-language access is one possible route to direct questions, but charts, maps, timestamps, and conventional controls remain necessary.

LITERATURE REVIEW

A. Review Methodology

This paper conducts a structured review of fifteen peer-reviewed sources. Eight focal works are compared in Table I; seven supporting works strengthen the conceptual, adoption, spatial, and computing context. The review is curated and does not claim the database coverage of a systematic review.

B. Corpus and Scope

The corpus covers precision agriculture [1], smart-farm data [2], IoT sensing [3], [5], decision-support design [4], irrigation [6], soil-moisture measurement [7], machine learning [8], remote sensing and unmanned aircraft [9], [10], [13], digital-agriculture practice [11], [12], conversational farmer access [14], and edge computing [15]. Sources were selected for relevance to the proposed architecture, not because every study examined hilly terrain.

C. Inclusion and Comparison Criteria

Inclusion required a direct contribution to agricultural sensing, analysis, visualization, decision support, or user access, or a computing principle relevant to low-connectivity operation. Comparison considers data dependence, analytical role, spatial resolution, user interaction, and transfer to terraces. These axes serve the design and are not a ranking of source quality.

D. Extraction and Synthesis

For each source, the review recorded approach, reported scope, limitations, and architectural relevance. Claims were kept within the source domain. Evidence from lowland or general agricultural systems is not treated as validation for hilly terrain.

E. Evidence Boundaries

The synthesis separates three levels of evidence. Statements about sensing, irrigation, machine learning, remote sensing, and interaction are attributed to the relevant literature. Requirements for terrace identifiers, quality flags, contour presentation, and synchronization are design inferences drawn from the target problem. Claims about future accuracy, usability, or operational value remain hypotheses for evaluation. This separation prevents a capability reported in general precision agriculture from being presented as a demonstrated hill-farming outcome.

F. IoT Sensing and Data Foundations

Ayaz et al. survey IoT-based smart agriculture across sensing, networking, platforms, and applications [3]. Tzounis et al. similarly examine agricultural IoT technologies and implementation challenges [5]. Together, these works support continuous monitoring but show that an end-to-end system also depends on communication, interoperability, power, and data handling. Their scope is broader than terrace agriculture, leaving terrain-specific placement and runoff interpretation as design questions.

Dobriyal et al. review soil-moisture estimation methods and distinguish their spatial support, cost, and applicability [7]. A reading is meaningful only in relation to method, placement, calibration, and represented soil volume. The proposed platform therefore stores sensor metadata and quality state with each observation.

G. Irrigation and Predictive Analytics

Jones reviews plant-based irrigation scheduling methods and their practical strengths and limitations [6]. The design implication is that no indicator should be treated as context-free authority. For terraced fields, soil, plant, weather, and terrain evidence may need joint interpretation, with recommendations remaining subject to calibration and validation.

Liakos et al. survey machine-learning applications across crop, livestock, water, and soil management [8]. Their review establishes a broad analytical toolkit but does not justify a particular model for this platform. The present paper therefore defines outputs and evaluation criteria while leaving algorithms and parameters to future data-driven comparison.

H. Spatial Observation and Mapping

Mulla reviews advances and unresolved questions in remote sensing for precision agriculture, including spatial and temporal resolution [9]. Zhang and Kovacs examine small unmanned aerial systems for high-resolution agricultural observation [10]. Sishodia et al. further organize remote-sensing applications and platform choices [13]. These sources motivate terrace-linked visualization and possible imagery inputs, but none is treated as proof of erosion or yield performance on the target slopes.

I. Decision Support and Farmer Interaction

Rose et al. emphasize effective design and delivery of agricultural decision-support tools [4]. Klerkx et al. show that digital agriculture has social and organizational as well as technical dimensions [11]. Bacco et al. describe a smart-farming ecosystem spanning sensing, communication, computation, and applications [12]. These findings support modular architecture and explicit usability evaluation.

FarmChat provides direct evidence that conversational agricultural information access can be studied with intended users and that interaction modality matters [14]. It does not validate the proposed interface, language coverage, or hill-farming workflow; it supports treating conversation as an evaluable mechanism rather than an assumed benefit.

J. Comparative Analysis

Table 1 Summary of the Reviewed Studies

Ref.	Primary approach	Main design lesson and limitation
[3]	IoT-based smart agriculture survey	Integrates sensing, networks, and services; terrace deployment remains untested.
[5]	Agricultural IoT review	Identifies challenges across the IoT stack; does not supply a hill-terrain workflow.
[6]	Irrigation-scheduling review	Compares indicators and limitations; requires context and calibration.
[7]	Soil-moisture methods review	Clarifies measurement tradeoffs; readings depend on method, scale, and placement.
[8]	Machine learning in agriculture	Surveys predictive applications; model choice requires representative data.
[9]	Remote sensing review	Supports spatial management; resolution and timing constrain interpretation.
[10]	Small UAS review	Enables high-resolution observation; acquisition and processing add overhead.
[14]	Farmer conversational agent	Studies interaction modalities; results do not automatically transfer to this platform.

The focal works differ in emphasis: [3], [5] address infrastructure; [6], [7] examine decision inputs; [8] addresses predictive methods; [9], [10] provide spatial observation; and [14] studies interaction. Supporting literature adds two cautions. Decision-support value depends on delivery and user practice [4], [11], while data integration introduces governance and interoperability concerns [2], [12].

RESEARCH GAP

No focal source evaluates the complete combination of quality-aware terrace ingestion, terrain-linked analytics, contour visualization, conversational orchestration, and offline synchronization. This is a corpus-scoped integration gap, not a claim that no related system exists.

Five gaps guide the design. First, agricultural IoT and soil-moisture analysis depend on credible inputs [3], [5], [7], but a terrace workflow must explicitly handle duplicate, delayed, missing, stale, and displaced-sensor data. Second, predictive methods [8] and spatial observation [9], [10], [13] occupy separate literature strands; the use case requires them to meet at a stable terrace identifier. Third, weather, management, and terrain can change after model development, requiring versioning, drift checks, and recalibration. Fourth, spatial detail still entails acquisition, record association, and map maintenance after physical change. Fifth, FarmChat shows why conversational access should be evaluated [14], while the proposed setting also needs degraded operation. Edge computing can move selected services closer to data sources [15], but Gemini-based interpretation remains network-dependent.

PROBLEM STATEMENT

Three obstacles motivate the design. Repeated inspection is difficult across dispersed terraces; delayed observation may leave little time for preventive action; and a table or farm-level average does not identify the physical terrace to which a warning applies. Sparse observations also weaken temporal comparison. Integrated monitoring must therefore preserve time, provenance, quality, and spatial context rather than merely digitize a notebook. These statements define the target problem; they are not measured findings from a deployment.

The problem addressed in this paper is therefore how to design a terrace-level monitoring and decision-support platform for hilly cultivation that (i) accepts sensor, weather, and manual observations whose quality can be assessed; (ii) links every analytical output to a stable terrace identifier and to the evidence behind it; (iii) keeps numerical authority in deterministic analytics rather than in a language model; and (iv) remains usable, with explicit stale-data and unavailable-function indicators, when connectivity is intermittent. These requirements follow from the gaps identified above.

PROPOSED SYSTEM

The proposed Intelligent Smart Agriculture Platform for Hilly Cultivation combines quality-aware ingestion, structured terrace-level storage, Python analytics, FastAPI services, constrained LangChain/Gemini orchestration, and React-based chart and contour-map presentation. This section sets out the design rationale derived from the reviewed literature and the analytical modeling framework; the layered architecture, data model, and workflow follow in the next sections.

A. From Prediction to Operational Decision Support

The literature distinguishes data acquisition, analysis, and decision support. A moisture estimate or risk class becomes operational only when the user can identify the terrace, severity, evidence age, and reason for priority. This requirement combines decision-delivery concerns [4] with sensing and analytical capabilities [3], [7], [8].

This distinction leads to a requirement for evidence-linked outputs. Each alert should carry a terrace identifier, status category, timestamp, relevant input observations, and a concise reason. Priority should be computed in the analytics layer using explicit rules or models, not inferred solely from conversational phrasing. The interface can then present the same structured result as a ranked list, chart, map state, or short explanation. This design preserves a path from user-facing advice back to data.

The platform therefore produces evidence-linked outputs. An alert carries terrace identity, status, timestamp, supporting variables, and quality state. Python rules or models compute priority; the interface presents the same structured result as text, chart, or map. Automated actuation remains outside the claimed scope.

B. Spatial and Predictive Integration

A farm-level average can conceal local variation. Remote-sensing literature supports spatially explicit interpretation [9], [13], while UAS imagery can offer finer observation where acquisition is feasible [10]. In the proposed platform, the contour map is a spatial index for analytical output rather than decoration. Labels, timestamps, values, and non-color cues accompany status colors.

Contour-oriented representation is especially relevant because elevation and slope influence water movement and access. The proposed map binds each analytical record to a terrace or plot entity instead of asking the language model to infer a location from free text. Predicted condition, observed condition, and data-quality status should be visually distinguishable. This prevents a forecast from appearing to be a directly measured fact and prevents stale cached data from appearing current.

C. Data Reliability Before Prediction

Sensor readings can be missing, noisy, duplicated, delayed, irregularly sampled, or physically unrepresentative after displacement. Connectivity interruptions can produce a burst of late records whose device timestamps differ from upload time. Temporary failure may create a constant value that is syntactically valid but implausible over a long period. These conditions matter because cleaning after prediction cannot undo an alert generated from corrupted input.

The architecture consequently places validation before persistence into the curated analytical stream. Basic checks include valid terrace and sensor identifiers, plausible units and ranges, timestamp consistency, duplicate detection, and flags for gaps or abrupt changes. Cleaning decisions should not erase the source record silently. Raw observations and quality flags should be retained so that interpolation, exclusion, or correction can be audited. Where a value is estimated to bridge a short gap, the derived status must remain distinguishable from an observed reading.

This mechanism responds to measurement and IoT dependencies discussed in [3], [5], [7]. No cleaning routine can identify every physical fault; calibration and inspection events must therefore update sensor trust explicitly.

D. Terrain-Aware Modeling

Slope, elevation, terrace position, orientation, rainfall, soil moisture, crop information, and historical yield represent different parts of the decision context. Slope and terrace position can affect runoff and accumulation; orientation and elevation can be associated with microclimate differences; crop information changes the interpretation of moisture; and historical yield provides a temporal baseline. Their inclusion is a design hypothesis derived from the reviewed themes, not evidence that each variable will be predictive in every region.

Machine-learning and remote-sensing reviews [8], [9], [13] motivate prediction from environmental and spatial features but do not establish terrace performance. Models must be trained or calibrated on representative slopes, with evaluation separated by terrace, season, or site. Current observations and stale-data indicators are necessary when conditions differ from historical records.

E. Human-Centered Interaction

Conversational access may reduce navigation for direct questions, but FarmChat demonstrates that modality preferences and usability require empirical study [14]. The language interface calls constrained functions and communicates structured outputs; it must not create measurements or replace Python analytics.

Multiple forms of presentation serve different tasks. Text can summarize priority, a chart can show change, and a contour map can support route planning. Voice provides an alternative to touch under some conditions. Consistent with human-centered decision-support research [4], [11], usability and language coverage remain evaluation questions.

F. Connectivity-Aware Operation

Explicit degraded modes follow from IoT communication requirements [3], [5] and the edge principle of placing selected services near data sources [15]. Terrace metadata, synchronized readings, map geometry, and recent alerts can be cached; new observations can be queued with device timestamps and synchronization state.

Local availability does not imply that all features are offline. Cached map and status viewing, queued entry, and simple deterministic filtering can operate without a network if their data and code are on the device. Gemini API calls require Internet access, so new language-model explanations are unavailable during an outage unless a separate local model is introduced in future work. The interface should identify stale data and unavailable functions rather than presenting a silent failure. This boundary makes the design testable and avoids claiming that a cloud-dependent conversational layer is fully offline.

G. Literature-to-Architecture Traceability

The final design follows a traceable sequence from evidence to requirement. The IoT reviews [3], [5] establish that field monitoring spans devices, communication, and applications; the corresponding requirement is an ingestion boundary that records device identity, source time, and communication status. The soil-moisture review [7] shows that measurement method and spatial support affect interpretation; the architecture therefore stores calibration and placement metadata instead of treating all readings as interchangeable.

Irrigation scheduling and machine-learning literature [6], [8] motivate analytical decision support, but neither permits the proposed language model to become the numerical authority. The resulting requirement is a reproducible Python function whose

inputs, transformation version, and output can be audited. A conversational request may select that function, but the numerical value must remain identical whether requested through voice, text, or a conventional dashboard.

Remote-sensing sources [9], [10], [13] motivate spatially resolved observation. Their limitations concerning scale, timing, acquisition, and processing lead to two design safeguards: every analytical state is linked to a stable terrace identifier, and every geometry record is versioned. This avoids treating a map as permanent ground truth and supports later replacement after erosion, landslides, or reterracing.

The decision-support and digital-agriculture literature [4], [11], [12] shifts attention from technical availability to delivery, governance, and work practice. Consequently, the interface exposes evidence age and quality, while cooperative extension is associated with authorization and audit requirements. FarmChat [14] supports investigating conversational access, not assuming its superiority. The proposed interface therefore preserves text, chart, map, and voice alternatives and subjects each to task-based evaluation.

Finally, edge-computing principles [15] inform the boundary between local and cloud functions. Caching, queued updates, quality checks, and selected deterministic retrieval can occur near the user or data source. Remote Gemini interpretation cannot. Stating this boundary turns “offline support” into a set of observable functions and failure behaviors rather than an undifferentiated architectural claim.

H. Analytical Modeling Framework

The analytics layer may support both regression and classification. Yield estimation is naturally framed as regression when the output is continuous, while moisture stress and erosion risk may be represented as ordered or categorical states. Irrigation priority and terrace inspection order are decision outputs that can combine model estimates with data freshness, severity, and operational rules. The appropriate algorithm cannot be selected solely from architecture; it requires representative data and validation.

At design level, let normalized indicators for terrace i at time t include moisture deficit $d_{i,t}$, recent rainfall intensity $r_{i,t}$, slope s_i , elevation or position e_i , crop sensitivity $c_{i,t}$, and erosion observation $u_{i,t}$. An illustrative risk index can be written as:

$$R(i,t) = w_d d(i,t) + w_r r(i,t) + w_s s(i) + w_e e(i) + w_c c(i,t) + w_u u(i,t) \quad (1)$$

Equation (1) is conceptual. The weights are not trained, specified, or validated, and an implemented relationship need not be linear. Its purpose is to place explicit normalized inputs in the analytics layer. A later study could compare a transparent score with machine-learning alternatives reviewed in [8].

Priority should also account for evidence quality. If $q_{i,t}$ denotes a reliability factor based on freshness, completeness, and sensor status, an illustrative inspection priority is:

$$P(i,t) = g(R(i,t), q(i,t), A(i,t), O(i,t)) \quad (2)$$

where A represents alert persistence and O represents operational context such as crop stage or an unresolved inspection. Function g is deliberately unspecified. A low reliability factor might raise a “verify sensor” task instead of allowing a severe but uncertain risk value to trigger an irrigation recommendation. This treatment converts data quality into decision context rather than hiding it during preprocessing.

Historical yield estimation could use soil, rainfall, crop, and terrain features consistent with data-intensive and machine-learning approaches [2], [8], but training and testing require real terrace outcomes. Thresholds and models should be versioned, and incoming feature ranges should be monitored for drift. None of these mechanisms implies present predictive validity.

SYSTEM ARCHITECTURE

A. Architectural Separation of Responsibilities

The platform follows a layered architecture in which authority over data and numbers remains outside the language model. Data sources include field sensors, weather observations, manual entries, and, where available, terrain or imagery-derived records. Ingestion validates and standardizes these inputs before structured storage. Python routines then retrieve records, compute indicators, estimate risks or yield, and generate alerts. FastAPI exposes these capabilities through typed service endpoints and domain functions.

LangChain coordinates the functions available to the conversational interface, while Gemini 1.5 Flash interprets user intent and formulates an explanation from tool results. The model does not serve as the numerical analytics engine and should not directly construct unrestricted database queries. React.js presents returned data; Recharts renders time-series or comparisons, and React-Leaflet associates terrace states with map features. This separation allows a numerical result to be reproduced independently of how it is phrased. The layers and data flow are summarized in Fig. 1.




Figure 1 Logical architecture and data flow of the proposed platform.

B. Database and Data Management

The logical data model centers on terrace and plot entities. A plot can group one or more terraces, while each terrace stores stable identity and links to its current spatial representation. Sensors are registered to terraces with device type, installation or calibration information, and status. Sensor readings contain value, unit, observation time, receipt time, and quality state. Weather and soil observations are separate logical entities because their sources, coverage, and sampling intervals may differ.

Crop records connect a terrace and season with crop information. Yield history preserves harvest observations for temporal comparison. Irrigation events and resource-usage records capture when water or other tracked resources were applied, without assuming that every event is automatically controlled. Risk-state records preserve analytical output, model or rule version, source interval, and uncertainty or quality indicators. Alerts point to risk states and maintain lifecycle information such as open, acknowledged, or resolved.

Synchronization state is part of the data model rather than a hidden network detail. A queued record needs a local identifier, creation time, status, retry information, and a stable deduplication key. Map features should carry a representation version so that records remain related to the geometry current when they were created, consistent with the maintenance implications of remote and aerial mapping [9], [10].

Structured records enable terrace-level queries, time-series analysis, and historical comparison. They also let the conversational layer answer constrained questions without treating dialogue as the primary database. Audit fields connect an output to its inputs. At

cooperative scale, the organizational concerns identified in digital-agriculture research [11], [12] imply tenant and role boundaries, although the prototype does not claim production-grade multi-tenancy.

C. Operational Interaction and Degraded Mode

A typical online interaction begins when a user submits text or voice through the React client. FastAPI receives a structured request, and LangChain selects an approved retrieval or analytics function with Gemini-assisted intent interpretation. The function returns database rows or computed outputs. Gemini may summarize those results, after which the client displays the explanation and the same structured evidence in a chart or contour map.

During network loss, the client can show the last synchronized terrace state, mark stale data, accept selected inputs, and queue updates. Deterministic filtering over cached data can continue. Requests requiring remote state or Gemini are deferred or reported unavailable. Reconnection triggers synchronization before status is presented as current, applying the local-processing rationale of edge computing [15] without claiming fully offline language-model operation.

METHODOLOGY/WORKFLOW

A. Data Processing and Decision-Support Workflow

The workflow begins with field data. Sensor inputs may include soil-moisture measurements and device-health metadata. Weather inputs contain observations relevant to the analysis, while manual observations can record crop state, visible erosion, maintenance, or calibration. Each record should contain an observation timestamp and source identifier so that delayed upload is not mistaken for a new measurement.

During ingestion, the backend authenticates or identifies the source, checks the expected data type and unit, and assigns provenance. Validation then checks referential consistency, temporal ordering, plausible ranges, duplicate keys, sampling gaps, and abrupt changes. Records that fail a hard structural rule can be quarantined; records that are usable but uncertain can be stored with a quality flag. Cleaning may standardize units, align time windows, or apply a documented missing-data method. The raw record remains available for audit.

The storage stage writes normalized observations and their quality metadata to SQLite. Structured persistence supports later retrieval by terrace, sensor, crop, and time interval. Python analytics selects quality-eligible records, joins them to terrain and crop context, and applies deterministic thresholds, statistical routines, or trained models when such models become available. The resulting outputs may include moisture-stress class, erosion-risk class, yield estimate, irrigation priority, and a ranked list of terraces.

Alert generation is a rule-governed step after analysis. An alert should reference the analytical output, relevant terrace, severity, creation time, and acknowledgement state. This

avoids recomputing historical alerts from changing rules without a record of the original decision. FastAPI returns structured payloads containing values, units, quality status, and timestamps.

For conversational requests, LangChain maps a recognized intent to an approved domain function such as retrieving current terrace state, comparing a time series, or explaining a stored alert. Gemini receives the user request and structured tool output and can produce a concise natural-language response. React then renders text and, where appropriate, passes chart-ready data to Recharts or geo-referenced status to React-Leaflet. The final product is decision support, not autonomous agronomic authority: the farmer or officer remains responsible for interpreting the alert in local context.

B. Decision Provenance and Failure Handling

Every decision-support response should include enough metadata to reconstruct its origin: the requested terrace and interval, source observations, quality flags, analytical version, execution time, and synchronization state. The natural-language explanation is stored as presentation output rather than as the authoritative result. If a later review identifies a wrong alert, this provenance allows the error to be assigned to collection, cleaning, retrieval, analysis, orchestration, or rendering.

Failure handling should favor explicit uncertainty over silent substitution. Missing terrain context, stale weather, conflicting measurements, or an unavailable cloud service should produce a bounded response that identifies the missing requirement. The system may show the last synchronized state, but it must not label that state as current. Similarly, the conversational layer may request clarification when a terrace name or time interval is ambiguous; it must not choose an entity merely to complete a fluent response. These rules turn claim discipline into implementable behavior and provide test cases for the evaluation framework.

IMPLEMENTATION & EVALUATION

A. Software Requirements

The proposed client uses React.js for the interface, Recharts for quantitative plots, and React-Leaflet for the terrace map. Browser speech facilities may provide voice capture where supported. FastAPI exposes backend functions, while Python implements cleaning, retrieval, and analytics. LangChain constrains orchestration over approved functions, and Gemini 1.5 Flash provides remote intent interpretation and explanation. SQLite stores prototype data locally. These named technologies describe the proposed implementation stack rather than experimentally compared choices.

The stack should run in a modern browser and on a workstation capable of hosting the local service and database. Remote language-model access and cloud synchronization require Internet connectivity. Local caching should use durable browser or device storage rather than volatile interface state so that queued records survive a restart.

B. Field and Hardware Considerations

A field implementation would require sensors appropriate to the selected variables, reliable timestamps, power provision, protected enclosures, and a store-and-forward communication mechanism. Hardware requirements depend on sampling frequency, radio technology, terrain, and the analytics assigned to the edge. Consequently, this design does not prescribe fixed processor, memory, storage, or sensor counts as universally sufficient.

Sensor installation and calibration procedures are as important as nominal hardware capability. Each device should be registered to a terrace and installation position, and maintenance events should be recorded. Drone or handheld imagery is a potential data source, but imagery-based crop or erosion analysis is not claimed as a validated current capability.

C. Prototype Output Screen

Fig. 2 shows an output screen of the monitoring application, in which the soil moisture, humidity, and temperature readings are presented to the user as labelled gauges with units. The screen is a single snapshot of displayed readings and serves only to illustrate the presentation of field data; it is not a performance or validation result.



Figure 2 Output screen of the monitoring application showing soil moisture, humidity, and temperature readings.

D. Evaluation and Validation Framework

The platform should be evaluated in stages because correct software operation, predictive validity, field usefulness, and usability are different claims. A reproducible evaluation should describe the data source, terrace and season coverage, split strategy, missing-data treatment, model version, and comparator. No numerical results are reported in this paper.

E. Predictive Performance

Regression outputs such as yield or continuous moisture estimates can be assessed with mean absolute error (MAE) and root mean square error (RMSE). MAE describes average absolute deviation in the output's unit, while RMSE places greater weight on large errors. Both should be reported with the evaluation distribution and, where possible, by terrace or terrain subgroup so that aggregate performance does not conceal failure on steep or high-elevation locations.

Categorical moisture-stress or erosion-risk outputs can be assessed through a confusion matrix, precision, recall, and F1-score. Recall is important where missed high-risk terraces are costly, whereas precision indicates how often generated warnings correspond to the target class. F1 summarizes their balance but should not replace class-wise reporting. Evaluation should use field labels or independently established reference observations and should prevent records from the same temporal episode from leaking across training and test sets.

F. Sensor and Data-Quality Robustness

A robustness study can introduce controlled missing readings, noise, outliers, duplicated records, delayed uploads, irregular intervals, and communication gaps into a clean test sequence. Sensor displacement requires either a field protocol or a realistic proxy because it may produce plausible but unrepresentative values rather than obvious corruption. The study should measure fault-detection rate, false alarms, the proportion of outputs suppressed or downgraded, and the change in predictive error after each disturbance.

The cleaning pipeline should also be tested for traceability. Evaluators should verify that raw values remain recoverable, quality flags propagate to risk states, and imputed observations are distinguishable from measurements. Scenarios with simultaneous faults can examine whether the system fails safely by requesting inspection instead of providing unjustifiably precise advice.

G. System and Database Performance

System testing can measure database query latency, analytics execution time, FastAPI response time, synchronization delay, dashboard rendering time, and local computational resource use. Workloads should vary the number of terraces, sensors, historical records, and concurrent users. Separate timing for deterministic analytics and remote language-model response is necessary because network and provider delay could otherwise be misattributed to the local platform.

Correctness under concurrency should be examined alongside latency. Tests can verify stable deduplication of queued observations, consistent alert states, and map refresh behavior when updates arrive from multiple devices. These measurements can also identify when SQLite ceases to be appropriate for the intended workload.

H. Connectivity Robustness

Connectivity evaluation should compare normal online operation, temporary loss, cached operation, and synchronization after reconnection. Test cases can inspect whether the client clearly marks stale content, preserves queued entries across restart, orders observations by source time, avoids duplicate insertion, and refreshes derived states only after synchronization. Cloud-dependent language interaction should be shown as unavailable or deferred rather than simulated locally without evidence.

I. Human–Computer Interaction

A future user study can compare the proposed interface with conventional dashboard navigation for representative tasks: locating an at-risk terrace, identifying the observations supporting an alert, comparing two periods, and recording a field observation. Measures can include task completion time, task success, navigation errors, request reformulation, and the proportion of natural-language queries mapped to the intended function. Voice evaluation should consider recognition success, correction effort, environmental noise, language, accent, and dialect.

Such a study should avoid assuming that voice is universally preferable. Participants should be able to use map, chart, text, or voice modes, and the analysis should examine which mode supports each task. Qualitative interviews may identify trust concerns, confusing terminology, or cases where concise explanations omit information that farmers consider essential.

J. Operational Usefulness

Field validation can assess whether ranked alerts improve irrigation or inspection prioritization relative to an existing practice. Potential outcomes include agreement with expert prioritization, time required to identify a terrace, alert acknowledgement and resolution patterns, and the number of avoidable inspection trips. Any later claim about resource saving, yield, or erosion outcomes would require a defined baseline, adequate observation period, and controls for rainfall, crop, and management differences. A full monsoon-cycle study would be particularly relevant, but the present paper reports no such experiment.

K. Baselines and Component Attribution

Evaluation should include baselines that match each claim. A persistence or recent-average baseline can contextualize a continuous estimate, while a documented threshold rule can contextualize a learned classifier. Map-based terrace identification can be compared with identifier-only or tabular presentation for localization tasks. Conversational routing can be compared with direct menu selection using the same underlying function, preventing interface fluency from being confused with analytical improvement.

Component-level tests should determine where an error originates. The same records can be processed with and without selected quality checks to examine whether

anomaly handling changes downstream stability. Cached and live retrieval can be compared after controlled disconnection to verify stale-state labeling and synchronization. For language responses, evaluators should separately score intent selection, parameter extraction, tool execution, and faithfulness to the returned data. This decomposition is important because a correct numerical result can still be communicated incorrectly, while a fluent explanation can conceal a failed retrieval.

A future field protocol should predefine exclusion criteria, missing-data handling, and the unit of analysis. Repeated readings from one terrace are not independent evidence of transfer to another terrace. Temporal splits can test later conditions; terrace- or site-held-out splits can test spatial transfer. Reporting both aggregate and subgroup results would reveal whether performance differs by slope range, sensor availability, crop, or season. These procedures are proposed evaluation controls and do not imply that a dataset or result already exists.

CONCLUSION

A. Summary of Contributions

The reviewed literature establishes strong foundations in sensing, irrigation scheduling, machine learning, spatial observation, decision support, and conversational access. For the hilly-terrain use case, input quality, prediction-to-location linkage, map maintenance, user interaction, and intermittent connectivity remain integration requirements rather than validated outcomes.

The proposed Intelligent Smart Agriculture Platform responds with quality-aware ingestion, structured terrace-level storage, terrain-sensitive Python analytics, constrained service functions, conversational orchestration, and chart and contour-map presentation. Numerical values originate in the database and analytics layer; LangChain and Gemini are limited to function orchestration and communication. Cached status, queued updates, synchronization state, and explicit stale-data indicators define a partial offline mode without implying that cloud language services operate without connectivity.

B. Field Data and Seasonal Validation

The immediate priority is validation with real terrace data. A future dataset should preserve observations across terraces, elevations, crops, rainfall events, irrigation actions, and verified outcomes. Collection over a complete monsoon cycle can expose abrupt weather, communication loss, sensor movement, and changing soil conditions that a short demonstration may miss. Longitudinal studies across multiple seasons are needed to examine drift and whether historical yield and risk relationships remain stable.

C. Analytical and Interaction Extensions

Analytical research can compare transparent rules, statistical models, and machine-learning methods under the same splits and metrics. Adaptive or online learning may help when weather and management change, but any update mechanism must protect against

learning from faulty sensors. Edge computing could perform quality checks, caching, and selected lightweight estimation near the field. Improved offline operation might later include local language processing, but the current architecture relies on remote Gemini access for model-based interpretation.

Local-language support should be developed with intended users rather than through translation alone. Regional vocabulary and dialect variation can be studied through task-oriented query sets, while voice systems should be tested under field noise. Long-term usability research can examine whether alerts are understood, whether explanations support trust calibration, and whether interaction adds or removes work from daily practice.

Computer vision using handheld, drone, or satellite imagery is another future direction [9], [10], [13]. Such data could assist crop-stress observation, terrace-boundary maintenance, or erosion inspection, subject to ground truth and image-quality constraints. Imagery should complement rather than automatically replace soil observations. Automated terrain-change detection may reduce remapping work, but its accuracy and cost require evaluation.

D. Staged Deployment and Reproducibility

Development should proceed through stages whose evidence claims remain separate. A first stage can test ingestion, validation, storage, analytics, and synchronization with controlled records and deliberately injected faults. This stage can verify software correctness but cannot establish agronomic validity. A second stage can operate on a limited number of mapped terraces while recommendations remain advisory and are compared with independent field observations. Only after sensor behavior, mapping accuracy, and alert interpretation are understood should a later study examine multiple holdings or cooperative use. This sequence limits the risk that an apparently complete interface is mistaken for a validated decision system.

The evaluation materials should be reproducible. Future reports should document sensor types, calibration events, units, sampling intervals, terrace identifiers, map versions, missing-data rules, and the exact analytical version responsible for each result. Data partitions should reflect the intended claim: a temporal holdout can test performance on a later period, whereas a terrace- or site-level holdout can test transfer beyond locations used for model development. Randomly dividing adjacent records from the same sensor between training and testing may inflate apparent performance because nearly identical observations occur in both sets.

Reproducibility also applies to conversational output. A test set of representative farmer questions can be paired with the intended domain function, required entities, and acceptable evidence. The evaluation can then distinguish successful intent routing from fluent wording. Stored tool calls and structured outputs should make it possible to determine whether a wrong answer arose from data retrieval, numerical analysis,

orchestration, or explanation. This separation follows the architecture's central principle that generated language must not become the hidden source of a numerical recommendation.

A field protocol should record user overrides and disagreements rather than treating them as errors by default. A farmer may know that a sensor was moved, that irrigation already occurred, or that a mapped path is temporarily inaccessible. Capturing the reason for an override can improve later data-quality rules and reveal operational variables absent from the initial model. However, these records should be reviewed before they are used as training labels; otherwise, inconsistent interface actions could be learned as if they were verified agronomic outcomes.

Future work should compare the cost and information value of alternative sensing arrangements. The dependencies discussed in [3], [5], [7] do not imply that every terrace needs every sensor. Studies could compare dense sensing, sentinel terraces, shared weather observations, and mixed automated/manual input under realistic maintenance, power, and communication constraints.

E. Cooperative Governance and Long-Term Maintenance

Larger cooperative deployments require a stronger database service, tenant isolation, role-based access, audit logging, backup, and synchronization. Digital-agriculture governance research [11], [12] also motivates attention to sensor authenticity, credentials, least-privilege access, retention, and the risks of sending farm information to cloud services.

Long-term operation requires ownership of maintenance decisions. A cooperative should define who can register a sensor, approve a changed boundary, alter a threshold, acknowledge an alert, and authorize aggregated access. Such rules should be tested with actual practice rather than inferred only from software roles.

Map and model maintenance should be coordinated. After erosion, a landslide, or reterracing, remotely derived geometry [9], [10], [13] may no longer represent the field. The platform should preserve prior geometry, date replacements, and identify outputs that used an older version. Updated models should likewise create new versions.

Future comparisons should define the operational baseline precisely. "Manual practice" may mean scheduled inspection, inspection after visible symptoms, notebook review, or judgement informed by local weather reports; these alternatives require different effort and provide different information. Recording the baseline permits later researchers to distinguish the contribution of map-based localization, predictive ranking, conversational access, and faster data availability. Before wider use, a candidate release should also preserve raw observations, identify stale inputs, reproduce deterministic outputs, enforce tenant access, survive interrupted synchronization, and expose the evidence behind each alert. Acceptance thresholds require agreement with domain experts and intended users and are not prescribed without data.

Finally, future studies should establish calibration schedules and compare system-supported prioritization with current field practice. Maintenance effort, training needs, connectivity costs, and the time required to resolve false or uncertain alerts should be recorded alongside technical performance. These are future possibilities, not functions demonstrated by the present paper. Their purpose is to move from an internally coherent prototype toward evidence of predictive, technical, and operational validity.

F. Concluding Remarks

The paper contributes a design and evaluation framework, not evidence of field performance. Predictive accuracy, robustness, latency, usability, and operational usefulness must be established through real terrace data and future studies. This distinction keeps the proposed integration testable while avoiding unsupported claims of deployment, savings, or agronomic benefit.

REFERENCES

- [1] R. Gebbers and V. I. Adamchuk, "Precision agriculture and food security," *Science*, vol. 327, no. 5967, pp. 828–831, 2010, doi: 10.1126/science.1183899.
- [2] S. Wolfert, L. Ge, C. Verdouw, and M.-J. Bogaardt, "Big data in smart farming—A review," *Agricultural Systems*, vol. 153, pp. 69–80, 2017, doi: 10.1016/j.agsy.2017.01.023.
- [3] 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.
- [4] D. C. Rose et al., "Decision support tools for agriculture: Towards effective design and delivery," *Agricultural Systems*, vol. 149, pp. 165–174, 2016, doi: 10.1016/j.agsy.2016.09.009.
- [5] A. Tzounis, N. Katsoulas, T. Bartzanas, and C. Kittas, "Internet of Things in agriculture, recent advances and future challenges," *Biosystems Engineering*, vol. 164, pp. 31–48, 2017, doi: 10.1016/j.biosystemseng.2017.09.007.
- [6] H. G. Jones, "Irrigation scheduling: Advantages and pitfalls of plant-based methods," *Journal of Experimental Botany*, vol. 55, no. 407, pp. 2427–2436, 2004, doi: 10.1093/jxb/erh213.
- [7] P. Dobriyal, A. Qureshi, R. Badola, and S. A. Hussain, "A review of the methods available for estimating soil moisture and its implications for water resource management," *Journal of Hydrology*, vols. 458–459, pp. 110–117, 2012, doi: 10.1016/j.jhydrol.2012.06.021.
- [8] 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.

- [9] D. J. Mulla, "Twenty five years of remote sensing in precision agriculture: Key advances and remaining knowledge gaps," *Biosystems Engineering*, vol. 114, no. 4, pp. 358–371, 2013, doi: 10.1016/j.biosystemseng.2012.08.009.
- [10] C. Zhang and J. M. Kovacs, "The application of small unmanned aerial systems for precision agriculture: A review," *Precision Agriculture*, vol. 13, pp. 693–712, 2012, doi: 10.1007/s11119-012-9274-5.
- [11] L. Klerkx, E. Jakku, and P. Labarthe, "A review of social science on digital agriculture, smart farming and agriculture 4.0," *NJAS—Wageningen Journal of Life Sciences*, vols. 90–91, Art. no. 100315, 2019, doi: 10.1016/j.njas.2019.100315.
- [12] M. Bacco, P. Barsocchi, E. Ferro, A. Gotta, and M. Ruggeri, "The digitisation of agriculture: A survey of research activities on smart farming," *Array*, vols. 3–4, Art. no. 100009, 2019, doi: 10.1016/j.array.2019.100009.
- [13] R. P. Sishodia, R. L. Ray, and S. K. Singh, "Applications of remote sensing in precision agriculture: A review," *Remote Sensing*, vol. 12, no. 19, Art. no. 3136, 2020, doi: 10.3390/rs12193136.
- [14] M. Jain, P. Kumar, I. Bhansali, Q. V. Liao, K. Truong, and S. Patel, "FarmChat: A conversational agent to answer farmer queries," in *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 2, no. 4, pp. 1–22, 2018, doi: 10.1145/3287048.
- [15] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet Things Journal*, vol. 3, no. 5, pp. 637–646, 2016, doi: 10.1109/JIOT.2016.2579198.