

ORIGINAL RESEARCH

Open access

Generative AI-Enabled Farm Advisory Systems: A Conceptual Architecture for Conversational Decision Support in Precision Agriculture

Nora Schneider^{1*}, Tobias Frank¹, Andreas Müller², Christoph Meier³

Abstract

Generative artificial intelligence is creating new possibilities for conversational access to agricultural information, farm data, and decision-support resources. However, fluent interaction does not establish factual accuracy, agronomic validity, farmer trust, or operational safety. This article proposes an original conceptual architecture for generative AI-enabled farm advisory systems in precision agriculture. The architecture treats conversational decision support as a governed socio-technical process rather than a standalone chatbot function. It connects farm advisory use-case definition, farmer query interpretation, heterogeneous agricultural data inputs, curated knowledge retrieval, generative model layers, response grounding, advisory-output controls, uncertainty communication, human oversight, expert escalation, farmer feedback, monitoring, and controlled system updating. Sensor observations, remote-sensing products, weather information, soil and crop data, farm-management records, and textual agricultural knowledge are positioned behind data-quality, provenance, documentation, and access-control requirements. Foundation models, large language models, vision-language models, multimodal models, prompting, domain adaptation, and retrieval-augmented generation are treated as optional architectural components whose appropriateness depends on the bounded advisory task. The architecture further distinguishes factual checking from agronomic validity review and separates farmer-facing explanations from advisor-facing evidence assessment. Safety and trust requirements include risk classification, hallucination monitoring, privacy, cybersecurity, auditability, farmer autonomy, explicit advisory boundaries, and escalation pathways. The principal contribution is a structured framework for designing and evaluating conversational farm advisory systems without implying autonomous decision authority, validated agronomic advice, guaranteed trust, operational safety, or deployment readiness.

Keywords Precision agriculture, Large language models, Generative AI, Human oversight, Farm advisory systems, Conversational decision support

*Correspondence:

Nora Schneider

nora.schneider@natfak.uni-freiburg.de

¹ Department of Generative AI for Farm Advisory, Faculty of Agriculture, University of Freiburg, Freiburg, Germany

² Department of Conversational Decision Support, Faculty of Agriculture, Heidelberg University, Heidelberg, Germany

³ Department of Precision Agriculture Advisory Systems, Faculty of Agriculture, University of Hohenheim, Stuttgart, Germany

Introduction

Precision agriculture increasingly depends on connected data environments in which field observations, sensor measurements, machinery records, remote-sensing products, weather information, farm-management data, and external knowledge are combined to support decisions. These environments are not merely technical pipelines: they involve data owners, service providers, advisors, researchers, platforms, and farmers whose roles influence how information is

collected, interpreted, exchanged, and acted upon. Consequently, a generative AI-enabled farm advisory system must be understood as part of a wider smart-farming data ecosystem rather than as an isolated conversational model [1].

The growing interest in conversational artificial intelligence also raises questions about who defines agricultural problems, whose knowledge is represented, and which users can meaningfully access or challenge system outputs. Responsible innovation in Agriculture 4.0 requires anticipation of possible consequences, inclusion of affected stakeholders, reflexive examination of design assumptions, and responsiveness to emerging concerns. These principles are particularly relevant when conversational systems communicate in authoritative language, because interface fluency can conceal uncertainty, data exclusions, institutional interests, or uneven distributions of benefits and risks [2].

Digital agriculture changes more than the speed or volume of information processing. It can reshape advisory relationships, professional roles, farmer skills, labour arrangements, knowledge ownership, and the balance of influence among agricultural actors. A conversational advisory system may therefore alter how farmers obtain information and how advisors contribute contextual interpretation. Its evaluation should consider institutional and social effects alongside technical performance, including farmer autonomy, local knowledge, digital inclusion, accountability, and the possibility that technology-mediated advice could reinforce existing inequalities or dependencies [3].

Agricultural decision-support systems generally connect data acquisition, analytical or predictive models, user interfaces, and decision processes. Generative AI adds a flexible natural-language layer to this established structure, but it does not remove the underlying requirements for reliable inputs, appropriate models, usable interaction, and contextual validation [4]. This article therefore proposes a conceptual architecture for generative AI-enabled farm advisory systems. The contribution links conversational support with data governance, knowledge retrieval, model control, advisory-output validation, human oversight, farmer feedback, safety boundaries, monitoring, and explicit limits on agronomic, autonomous-decision, trust, and deployment-readiness claims.

Generative AI in Farm Advisory Systems

Generative AI refers to computational models that produce new textual, visual, structured, or multimodal content in response to instructions and contextual inputs. Within crop production and broader farm advisory settings, large language models may help organize information, reformulate technical material, interpret questions, or support access to documented knowledge. Their value is nevertheless task dependent. A model that produces coherent agricultural language is not thereby qualified to determine whether a response is agronomically correct, locally appropriate, sufficiently current, or suitable for a consequential farm decision [5].

The possible use of large language models in food production must also be evaluated against risks of misinformation, inappropriate data collection, unequal access, changing labour roles, and unclear responsibility for generated outputs. A farmer-facing system may appear helpful because it responds rapidly and conversationally, yet its statements can remain incomplete, unsupported, or inconsistent with local conditions. Governance cannot therefore be reduced to a disclaimer placed beneath an interface. It requires defined use cases, restricted output categories, documented data practices, human review conditions, incident handling, and clear identification of accountable organizations and professionals [6].

Generative AI includes more than a single model family. Foundation models may provide broad pretrained capabilities, while large language models emphasize linguistic processing and vision-language or multimodal models combine text with images or other data forms. Generative methods may support synthetic content, information synthesis, question answering, or interactive explanation across agri-food activities. However, each application presents different evidence requirements and failure modes. Architectural design should therefore begin with the advisory task and risk level rather than assuming that a larger or more multimodal model automatically provides more useful or trustworthy support [7].

Prompting, fine-tuning, domain adaptation, knowledge retrieval, and retrieval-augmented generation represent distinct methods for shaping model behaviour. Prompting can specify response structure and boundaries, whereas fine-tuning or

domain adaptation may align terminology and interaction patterns with a defined agricultural domain. Retrieval can supply external evidence at response time, allowing generated statements to be linked to curated sources. None of these methods independently establishes reliability. Agricultural language models continue to face limitations involving data representation, context specificity, hallucination, evaluation, infrastructure, and accessibility, requiring cautious integration within a controlled advisory workflow [8].

Conversational Decision Support in Precision Agriculture

Conversational decision support begins with interpretation of the user's question rather than immediate response generation. Farmer queries may contain local terminology, incomplete descriptions, uncertain timeframes, or implicit assumptions about a field, crop, disease symptom, management activity, or environmental condition. A retrieval component can transform the clarified question into a search across scientific or curated agricultural documents and return relevant passages for further review. Agricultural retrieval research shows that locating evidence for farmers' questions is itself a separate technical problem whose quality must be evaluated before retrieved material is used to ground a generated response [9].

A conversational agent or chatbot can make information easier to request, refine, and revisit through natural-language dialogue. In extension contexts, this may assist with navigation of educational resources, reformulation of technical concepts, or preparation of questions for a specialist. Such systems should complement rather than replace the relational and interpretive functions of extension professionals. The interface should disclose its informational role, identify circumstances requiring human assistance, avoid presenting generated language as professional judgment, and provide an escalation route when the question exceeds available evidence or enters a potentially consequential advisory domain [10].

Generative AI may also function as an interface between farmers and analytical tools that would otherwise require technical expertise. It can potentially translate model outputs, summarize observations, ask for missing context, or present alternative interpretations in accessible language. Its usefulness nevertheless depends on language support, connectivity, device access, literacy, local terminology, and representation of diverse production conditions. Farmer-facing explanations should therefore be designed around the intended users and decision context. Advisor-facing explanations may require more detailed evidence provenance, assumptions, data-quality indicators, and unresolved uncertainties than the simplified account presented to a farmer [11].

Conversational decision support must maintain a boundary between organizing information and issuing agronomic recommendations. Domain-specific agricultural question-answering systems can be compared using defined evaluation procedures, but performance within an evaluation setting does not establish local advisory validity or field readiness [12]. Scenario explanations should use conditional language and make assumptions visible. Uncertainty communication should identify missing data, conflicting sources, retrieval gaps, and questions requiring expert interpretation. The system must also preserve farmer autonomy by allowing users to question, correct, reject, or escalate outputs, and it must not authorize or execute farm-management actions autonomously.

Data Inputs, Model Layers, and Advisory Outputs

The proposed architecture accepts heterogeneous inputs only after determining their relevance and fitness for the defined advisory use case. Potential inputs include sensor measurements, weather observations, soil and crop records, satellite products, remote-sensing data, drone imagery where available, pest or disease observations, irrigation and nutrient records, machinery information, farm-management histories, and textual agricultural knowledge. Multimodal language models may combine textual and visual representations, but agricultural adaptation requires attention to modality alignment, dataset coverage, computational constraints, and task-specific evaluation. Multimodal capability should therefore be treated as an optional design choice rather than a universal requirement [13].

Vision-language models may support interaction with crop images, field observations, pest or disease symptoms, and other visual evidence. Nevertheless, image-conditioned responses remain dependent on acquisition quality, spatial and temporal relevance, resolution, labelling, and the model's ability to generalize beyond its development context [14]. Satellite and drone imagery should carry acquisition dates, spatial boundaries, processing information, and quality indicators. Soil, weather, crop, and sensor data should retain units, timestamps, collection methods, calibration status, and missingness. Visual interpretation should remain descriptive or evidential unless qualified review establishes that stronger conclusions are appropriate.

The knowledge and model layers should remain functionally distinct. A curated agricultural knowledge base provides controlled textual evidence; retrieval ranks passages relevant to a clarified query; grounding links response claims to selected evidence; and a generative model converts structured context into a provisional conversational output. Knowledge-guided agricultural language-model research indicates that retrieval and domain adaptation can improve knowledge assimilation while remaining vulnerable to corpus gaps, weak retrieval, omitted detail, and domain-specific data scarcity [15]. Accordingly, retrieval-augmented generation should include source versioning, relevance thresholds, conflict handling, abstention rules, and claim-level traceability rather than being treated as a general solution to hallucination.

The advisory-output layer should produce bounded informational support such as evidence summaries, clarification questions, conditional interpretations, uncertainty statements, source trails, or expert referrals. Model documentation should identify intended uses, exclusions, versions, adaptation history, and known limitations. Output traceability should preserve the relationship among the query, input data, retrieved evidence, model version, response controls, human review, and final release state. Generative AI must be embedded within wider digital agri-food governance and monitoring processes rather than evaluated as a self-contained content generator [16]. **Table 1** summarises the data inputs, model layers, conversational functions, advisory outputs, and validation needs for generative AI-enabled farm advisory systems.

Table 1. Data Inputs, Model Layers, and Advisory Outputs for Generative AI-Enabled Farm Advisory Systems: Sensor Data, Remote Sensing, Weather, Soil, Crop, Pest and Disease Data, Farm Records, Agricultural Knowledge Bases, Retrieval Quality, Response Grounding, Output Traceability, Farmer-Facing Explanations, and Advisory Boundaries

Architecture component	Core system question	Required input or condition	Model or data function	Potential advisory output	Failure risk	Validation or governance need
Farmer and farm context	What farm, crop, location, production stage, and decision context are relevant?	Farmer-confirmed contextual information and defined advisory purpose	Constrains query interpretation and evidence selection	Restated question or request for missing context	Incorrect assumptions or inappropriate generalization	Farmer confirmation, data minimization, and context review
Sensor data	Are measurements current, calibrated, complete, and relevant?	Device identity, units, timestamps, calibration records, and missingness indicators	Ingests and normalizes field observations	Descriptive measurement summary or data-quality warning	Sensor drift, unit errors, missing observations, or device mismatch	Calibration review, plausibility checking, and provenance tracking
Remote-sensing and	Does the imagery	Geospatial boundaries,	Supplies spatial and	Image-derived	Cloud effects, outdated imagery,	Geospatial verification,

satellite data	represent the correct field, period, and spatial scale?	acquisition date, resolution, processing level, and quality metadata	temporal observations	contextual description or trend summary	spatial mismatch, or unsuitable resolution	quality screening, and task-specific validation
Drone imagery	Is high-resolution imagery sufficiently complete and properly processed?	Flight metadata, image quality, coverage, calibration, and authorization	Provides detailed visual evidence where supported	Visual-condition summary or referral for review	Sampling bias, artefacts, incomplete coverage, or false visual inference	Image-quality checks and local expert interpretation
Weather data	Are observed or forecast conditions sufficiently local and current?	Source identity, time horizon, geographic relevance, and uncertainty information	Contextualizes environmental conditions	Weather-context statement with limitations	Forecast uncertainty, location mismatch, or stale information	Source disclosure and uncertainty communication
Soil and crop data	Are soil measurements and crop observations representative of the relevant management area?	Sampling method, depth, date, crop identity, variety, growth stage, and units	Provides agronomic context for retrieval and interpretation	Soil or crop condition summary and identified information gaps	Non-representative sampling, incorrect stage assumptions, or unit confusion	Sampling documentation, farmer confirmation, and agronomic review
Pest and disease observations	Is available evidence sufficient for description, or is specialist diagnosis required?	Images, symptoms, timing, prevalence, crop context, and diagnostic evidence	Retrieves relevant knowledge and identifies uncertainty	Descriptive differential information or expert referral	Misidentification or inappropriate action based on incomplete evidence	High-risk classification, expert escalation, and no-treatment boundary
Irrigation, nutrient, and machinery records	Are operational records complete, interpretable, and authorized for use?	System type, application history, flow or machinery data, units, timestamps, and access rights	Supplies operational and historical context	Record summary, inconsistency notice, or missing-data request	Incorrect interpretation, proprietary-format errors, or security exposure	Access controls, record verification, and qualified review
Farm-management records	Are longitudinal records current, correctly	Field histories, activities, inputs, observations,	Provides temporal context across	Farm-history summary or contextual	Incomplete records, cross-farm mixing,	Consent, provenance, access control,

	linked, and used with consent?	outcomes, and retention permissions	seasons or operations	evidence packet	privacy breach, or outdated context	and retention governance
Textual agricultural knowledge	Which evidence sources are authoritative, current, and relevant to the bounded use case?	Curated peer-reviewed and validated agricultural materials with metadata	Populates the agricultural knowledge base	Source-linked evidence for retrieval	Outdated, conflicting, biased, or geographically unsuitable information	Editorial curation, domain review, version control, and conflict policy
Data quality, provenance, and documentation	Are inputs fit for purpose and fully traceable?	Completeness, accuracy, consistency, timeliness, source, transformations, and data dictionary	Scores input fitness and records lineage	Quality status, provenance record, abstention, or blocked processing	Invalid data transformed into persuasive language	Mandatory pre-generation quality gate and auditable documentation
Knowledge retrieval	Does the retrieval process return relevant, sufficient, and current evidence?	Clarified query, curated corpus, metadata, ranking method, and relevance criteria	Searches, ranks, filters, and assembles evidence	Ranked passages or evidence packet	Irrelevant retrieval, missing key evidence, or source conflict	Retrieval evaluation, thresholds, source display, and abstention
Retrieval-augmented grounding	Can material response claims be linked to retrieved evidence?	Selected passages, claim–source associations, response constraints, and source status	Conditions generation on evidence and preserves links	Grounded draft with source traceability	Unsupported synthesis despite apparently relevant retrieval	Claim-level checking, conflict handling, and human review
Generative model layer	Is the selected model suitable for the defined language, modality, and task?	Model documentation, version, intended use, exclusions, adaptation history, and evaluation evidence	Produces a provisional response from controlled context	Draft explanation, summary, clarification, or referral language	Hallucination, bias, context drift, or false specificity	Domain evaluation, version control, risk-based restrictions, and monitoring
Prompt and response control	Are scope, format, uncertainty, prohibited	Approved instructions, output schema, risk state, user	Constrains response form and screens	Structured response candidate	Prompt sensitivity, inconsistent compliance, or boundary bypass	Testing, change control, logging, and policy review

	claims, and escalation rules enforced?	role, and safety policies	prohibited content	with warnings or abstention		
Farmer-facing explanation	Is the information understandable without overstating certainty or authority?	User language, literacy needs, local terminology, evidence state, and uncertainty level	Translates technical evidence into accessible language	Plain-language explanation, missing-data notice, or next-step referral	Oversimplification, persuasive fluency, or inappropriate reliance	Farmer-centred usability and comprehension evaluation
Advisor-facing evidence view	Can a qualified reviewer inspect evidence, assumptions, and unresolved uncertainties?	Retrieved sources, input provenance, model status, risk classification, and draft output	Presents a detailed review packet	Approved, revised, rejected, or escalated output	Automation bias or superficial review	Qualified reviewer authority, workload support, and audit trail
Output traceability	Can the final output and its development state be reconstructed?	Query, input versions, evidence identifiers, model and prompt versions, review actions, and timestamps	Records response lineage	Auditable output record	Unreconstructable error or unclear accountability	Secure logging, access governance, retention policy, and independent audit
Advisory-validity boundary	Do technically valid inputs and grounded text establish agronomic correctness?	Explicit system limitation and independent domain-review requirement	Blocks unsupported inference from grounding to advisory validity	Caution, abstention, or expert referral	A well-sourced response treated as validated agronomic advice	Agronomic validity review and local validation planning

Human Oversight, Validation, and Farmer Feedback

Human oversight should be designed as an active architectural function rather than a symbolic assurance that a person is somewhere within the advisory process. Agricultural advisors contribute contextual sensemaking by interpreting digital information in relation to farm history, seasonal conditions, production constraints, farmer priorities, and locally established practices. A generative system may organize evidence or prepare a provisional explanation, but it cannot reproduce the relational knowledge, responsibility, and situated judgment through which advisors determine whether information is relevant to a particular farm. The architecture therefore assigns advisors a review and escalation role without assuming that every conversational exchange requires the same level of professional intervention [17].

Human-in-the-loop review should be triggered according to the risk, uncertainty, and intended use of an output. Evidence from combined digital and human advisory services indicates that digital information and interpersonal support can have different effects across farming populations and management contexts [18]. Consequently, human involvement should not be treated as a universal correction that automatically validates generated content. Reviewers require access to the

original query, data provenance, retrieved evidence, model and prompt versions, identified uncertainties, and proposed response. The architecture should allow a reviewer to approve, revise, reject, or escalate the output while recording the basis of that decision for subsequent audit and learning.

Farmer participation is equally important because the farmer may identify misunderstood terminology, incorrect field context, missing records, inappropriate assumptions, or inaccessible explanations. User-centred agricultural decision-support research emphasizes early stakeholder involvement, assessment of the real decision context, and iterative evaluation of system usability and relevance [19]. Farmer feedback should therefore be separated into categories such as contextual correction, usability feedback, source challenge, outcome observation, and complaint. A farmer's positive response cannot be treated as evidence of agronomic validity, while a correction should not flow directly into model training. Feedback must first be reviewed, documented, and routed to the relevant interface, data, retrieval, governance, or validation process.

Validation must distinguish factual accuracy, agronomic validity, local applicability, and operational reliability. Factual checking asks whether a statement is supported by evidence, whereas agronomic review asks whether the interpretation is meaningful under the relevant crop, soil, weather, management, and regional conditions. Agricultural knowledge and advice increasingly operate through networks involving farmers, advisors, researchers, firms, regulators, and digital intermediaries [20]. Validation should therefore include local evaluation planning, hallucination monitoring, advisory-output review, model-drift testing, knowledge-base maintenance, and controlled feedback updating. Neither a successful benchmark nor a carefully designed architecture is sufficient to establish field validity, operational safety, or deployment readiness.

Safety and Trust Requirements

Trust in a farm advisory system cannot be generated solely through a friendly interface, an explanation, or the display of source links. Farmers' trust in smart-farming services is influenced by how data are collected, who benefits from their use, whether institutional practices are transparent, and whether users can challenge or withdraw from the arrangement [21]. The architecture therefore treats trust as a relational and institutional outcome rather than a technical feature. It requires clear use-case disclosure, understandable data practices, meaningful consent, farmer control, visible limitations, complaint mechanisms, and evidence that responsible actors respond appropriately when errors or conflicts occur.

Hallucination represents a central safety concern because a generated response may be grammatically fluent while containing unsupported, contradictory, or fabricated content. Hallucination research in natural-language generation shows that factual inconsistency takes multiple forms and requires explicit definitions, evaluation methods, mitigation strategies, and continuing monitoring [22]. The proposed architecture combines retrieval checks, claim-level evidence links, factual review, uncertainty communication, abstention rules, and incident logging. Retrieval-augmented generation is not treated as a guarantee because an output may still rely on irrelevant passages, omit conflicting evidence, misinterpret a source, or introduce unsupported conclusions during generation.

Safety also requires a clear allocation of control between computational assistance and human judgment. Human-centred AI supports the combination of substantial automation with substantial human control when interfaces, responsibilities, and intervention mechanisms are deliberately designed [23]. For conversational farm decision support, this means preserving the farmer's authority to question, reject, or disregard outputs and providing qualified reviewers with real power to modify or block consequential responses. Risk classification should identify whether a query is informational, context-sensitive, potentially consequential, or beyond scope. High-risk or insufficiently grounded cases should trigger abstention or expert escalation rather than more confident generation.

Privacy and cybersecurity extend across the entire technical environment supporting conversational advice. Smart-farming systems may connect devices, gateways, communication networks, edge resources, cloud services, databases, analytical tools, and user applications, creating multiple points at which confidentiality, integrity, availability, or access

control can fail [24]. The architecture therefore requires data minimization, authentication, authorization, protected transmission and storage, secure logging, incident response, and review of third-party access. Accountability must identify who is responsible for data quality, knowledge curation, model operation, advisory review, security, and remediation. These safeguards may support trustworthiness, but they cannot guarantee farmer trust or safe deployment.

Proposed Conceptual Architecture

The proposed architecture begins with farm advisory use-case definition and stakeholder mapping. The intended users, decision context, included functions, prohibited outputs, data requirements, affected stakeholders, and responsible organizations must be specified before a model is selected. Explanation requirements should also be defined at this stage because explainable AI includes different purposes, audiences, methods, and limitations [25]. A farmer may need a concise explanation of evidence and uncertainty, whereas an advisor may require source passages, data provenance, model documentation, assumptions, and a record of response controls. Explainability is therefore positioned as a separate architectural function and not as evidence that the underlying output is correct.

The second stage covers farmer query intake, clarification, and conversational interaction. The system should identify ambiguity, request essential context, confirm the intended task, and determine whether the query remains within its approved use case. Explanations should be selective, contrastive, and responsive to the user's context rather than exhaustive technical descriptions [26]. Accessibility requirements include plain language, multilingual support where validated, compatibility with low-bandwidth channels, and alternatives for users with differing literacy or sensory needs. The interaction layer should disclose when it is generating information, communicate uncertainty, permit correction, and avoid interface patterns that pressure the farmer toward accepting a particular interpretation.

The third and fourth stages connect data assessment with knowledge retrieval and generative processing. Sensor, imagery, environmental, crop, operational, management, and textual inputs pass through quality, provenance, documentation, consent, privacy, and security checks. Curated knowledge is then searched and ranked, with retrieval quality evaluated independently from response generation. A generative model may produce a provisional synthesis using prompting, domain adaptation, fine-tuning, or multimodal processing where justified. Governance principles alone cannot make these processes responsible unless they are translated into assigned duties, professional practices, review procedures, auditing arrangements, and enforceable accountability [27]. The architecture therefore associates every technical layer with named oversight and change-control functions.

The fifth and sixth stages govern advisory-output generation, risk classification, validation, and human escalation. Draft responses are assessed for source grounding, factual consistency, agronomic-review requirements, uncertainty, prohibited claims, and potential consequences. Agricultural AI raises domain-specific concerns involving autonomy, labour, power, responsibility, data relations, environmental effects, and the treatment of different stakeholders [28]. These concerns cannot be reduced to generic model-performance measures. Farmer-facing outputs should remain informational and conditional, while advisor-facing views should expose assumptions and unresolved evidence gaps. Cases requiring diagnosis, prescriptions, consequential management choices, or unsupported local interpretation should be blocked or referred to appropriate human expertise.

The final stage establishes governed feedback, monitoring, and updating loops. Farmer corrections, advisor reviews, retrieval failures, security incidents, local validation findings, model-drift indicators, and knowledge-base changes are recorded and assessed before they alter any system component. Responsible smart-farming development requires socio-ethical considerations to be integrated throughout innovation rather than appended after technical construction [29]. Updating must therefore use version control, approval procedures, regression testing, and rollback capacity. The architecture retains explicit boundaries against agronomic-authority claims, autonomous farm decisions, guaranteed trust, certified safety, and deployment readiness. **Table 2** presents the proposed conceptual architecture for generative AI-enabled farm advisory systems in precision agriculture.

Table 2. Proposed Conceptual Architecture for Generative AI-Enabled Farm Advisory Systems in Precision Agriculture: Use-Case Definition, Farmer Query Intake, Data Quality, Knowledge Retrieval, Generative Model Layer, Response Grounding, Advisory Output Review, Human Oversight, Farmer Feedback, Safety Boundaries, Trust Requirements, Accountability, Monitoring, and Claim Boundaries

Architecture stage	Core purpose	Required input	Architectural function	Potential output	Validation or oversight need	Failure risk
Farm advisory use-case definition	Establish the bounded purpose of the system	Intended users, advisory context, included tasks, excluded tasks, and evidence requirements	Defines scope, output classes, and prohibited uses	Approved use-case statement	Farmer, advisor, domain, ethics, and governance review	Scope creep or use in unsupported decisions
Stakeholder and user mapping	Identify affected actors, needs, rights, and responsibilities	Farmers, advisors, developers, data stewards, service owners, researchers, and indirect stakeholders	Assigns roles, authority, access, and accountability	Stakeholder and responsibility map	Participatory review and inclusion assessment	Exclusion of marginal users or responsibility gaps
Farmer query intake	Capture the farmer's actual question	Natural-language query, preferred language, consent, and available context	Records and classifies the request	Query record and initial task classification	Privacy and usability review	Misinterpreted intent or unnecessary data collection
Query clarification and context confirmation	Identify missing, ambiguous, or inconsistent information	Crop, field, timing, management context, local terminology, and intended decision	Asks targeted clarification questions	Confirmed query or request for additional evidence	Farmer confirmation	False assumptions propagated into retrieval and generation
Data input assessment	Determine which data are relevant to the query	Sensor, weather, soil, crop, imagery, operational, farm-record, and textual inputs	Selects necessary inputs and rejects irrelevant data	Data inventory	Data steward and domain review	Excessive, irrelevant, or incomplete data use
Data quality and	Verify fitness, meaning, and	Source, timestamps,	Scores input quality and	Quality status, provenance	Automated checks and	Invalid or untraceable

provenance review	traceability	units, collection methods, calibration, transformations, and documentation	records lineage	record, or blocked input	qualified review	data entering the model
Privacy and cybersecurity review	Protect farm, personal, and operational data	Data-flow map, consent, threat model, access roles, storage and transmission controls	Applies minimization, authorization, security, and incident controls	Privacy and security status	Specialist review, testing, and monitoring	Leakage, unauthorized access, manipulation, or service disruption
Knowledge-base curation	Maintain authoritative and relevant agricultural evidence	Candidate literature, validated advisory resources, metadata, update schedules, and editorial criteria	Selects, classifies, versions, supersedes, and archives sources	Curated agricultural knowledge base	Agronomic and editorial governance	Outdated, biased, conflicting, or geographically unsuitable sources
Knowledge retrieval	Locate evidence directly relevant to the clarified query	Curated corpus, query representation, metadata filters, and ranking procedure	Searches, ranks, and filters documents and passages	Evidence packet	Relevance evaluation and coverage checks	Irrelevant, incomplete, or stale retrieval
Retrieval and grounding layer	Bind material claims to selected evidence	Evidence packet, claim–source links, source status, and conflict indicators	Supplies controlled context to generation and preserves traceability	Grounded response context	Claim-level review and conflict handling	Grounded-sounding synthesis unsupported by the sources
Generative model layer	Produce a provisional conversational synthesis	Model version, approved context, intended task, and documented limitations	Generates text or multimodal response candidates	Draft explanation, summary, clarification, or referral	Domain-specific evaluation and version control	Hallucination, bias, context drift, or false specificity
Prompt and response-control layer	Enforce task, format, risk, uncertainty, and claim boundaries	Approved prompts, output schema, user role, risk state,	Constrains generation and screens responses	Structured response candidate	Adversarial testing, logging, and change control	Prompt sensitivity, inconsistent

		and prohibited-content rules				compliance, or control bypass
Response grounding and factual checking	Determine whether claims match available evidence	Draft response, retrieved passages, farm inputs, and claim-support map	Detects unsupported or contradictory content	Verified, corrected, blocked, or uncertain response	Automated checks plus human review for consequential cases	Plausible but unsupported statements escape detection
Agronomic validity review	Assess whether interpretation is meaningful in context	Farm context, evidence, uncertainties, and qualified agricultural expertise	Reviews domain relevance and local applicability	Reviewed interpretation, revision, or referral	Agricultural advisor or domain-expert review	Correct facts combined into inappropriate farm guidance
Uncertainty and explanation layer	Communicate evidence, assumptions, unknowns, and limitations	Data-quality state, source agreement, retrieval coverage, model limitations, and user needs	Produces audience-specific explanation and uncertainty language	Farmer-facing explanation or advisor-facing evidence view	Comprehension and accuracy testing	False precision, vague warnings, or misleading explanation
Risk classification	Match safeguards to the consequence of error	Output type, intended use, uncertainty, reversibility, and affected parties	Assigns risk level and required controls	Risk status	Independent review of high-risk classification	Under-classification or inconsistent treatment
Advisory-output review	Confirm compliance before release	Draft output, evidence links, risk state, uncertainty, and policy checklist	Approves, revises, blocks, or escalates output	Release decision and review record	Qualified reviewer with meaningful authority	Rubber-stamping or excessive reliance on checklists
Human oversight and expert escalation	Preserve human intervention and responsibility	Reviewer roles, evidence packet, escalation criteria, and contact pathway	Transfers unresolved or consequential cases	Human-reviewed output or expert referral	Review-capacity and accountability monitoring	Unavailable expert, delayed referral, or nominal oversight
Farmer-facing conversational output	Present bounded, accessible decision-support information	Approved response, user preferences, uncertainty statement, and source information	Delivers explanation, options, questions, or referral	Informational conversational response	Fluent language interpreted as validated advice	Usability, accessibility, and comprehension evaluation

Farmer feedback and correction	Capture misunderstanding, contextual error, accessibility problems, and contestation	Voluntary feedback, correction controls, and consent	Routes feedback to the relevant governed process	Feedback or correction record	Satisfaction treated as validity or malicious feedback propagated	Feedback classification and human review
Validation and monitoring layer	Evaluate reliability before and during any future operational consideration	Test cases, local cases, incident logs, user research, drift indicators, and change history	Measures performance, detects failures, and restricts use	Validation, drift, or incident report	Independent technical, agronomic, and human-factors review	Narrow benchmarks conceal consequential failures
Knowledge-base updating	Maintain current, accurate, and context-relevant evidence	New studies, corrections, retractions, local evidence, and source challenges	Reviews and versions corpus changes	Updated knowledge-base version	Editorial and agronomic approval	Unreviewed update or retention of superseded evidence
Model or system updating	Introduce controlled technical or policy changes	Change request, validation evidence, risk analysis, and rollback plan	Tests, approves, versions, releases, or rejects changes	New controlled system version	Change-control board and post-change monitoring	Regression, undocumented behaviour, or lost traceability
Accountability and auditability	Identify responsible actors and reconstruct system events	Role assignments, data lineage, model versions, prompts, outputs, reviews, and incidents	Maintains audit records and supports investigation	Audit trail, corrective action, and responsibility record	Independent audit and protected access	Missing logs, excessive surveillance, or responsibility shifted to users
No agronomic-advice claim boundary	Prevent generated information from being represented as validated agronomic advice	Explicit use restrictions and expert-review criteria	Blocks unsupported prescriptions and diagnostic certainty	Caution, abstention, or expert referral	Informational content converted into a management prescription	Policy review and incident monitoring
No autonomous-decision claim boundary	Preserve farmer and authorized human control	Action permissions and separation from machinery or actuation systems	Prevents autonomous commitment or execution	Decision-support information only	Unintended action, automation bias, or loss of control	Authorization audits and technical separation

No trust-guarantee boundary	Prevent trust from being claimed as an engineered outcome	Trust evidence, communication rules, and complaint mechanisms	Restricts promotional or deterministic trust claims	Calibrated trustworthiness statement	Interface fluency or explanation used to manufacture trust	Independent farmer-centred evaluation
No deployment-readiness boundary	Separate conceptual design from operational maturity	Implementation evidence, local validation, security tests, governance practice, and field evaluation	Restricts readiness and safety claims	Conceptual-status declaration	Architecture completeness mistaken for operational proof	Independent staged readiness assessment

Figure 1 presents the proposed conceptual architecture for generative AI-enabled farm advisory systems, linking farm advisory use cases, data inputs, retrieval and model layers, conversational outputs, human oversight, safety controls, farmer feedback, and system updating.




Figure 1. Conceptual Architecture for Generative AI-Enabled Farm Advisory Systems in Precision Agriculture

Alt-text description: A conceptual architecture showing farmer queries and farm data entering a generative AI advisory system with retrieval, grounding, model layers, advisory outputs, validation gates, human oversight, safety boundaries, farmer feedback, monitoring, and updating loops.

Practical Implications and Future Research

For farmers, advisors, extension services, cooperatives, and data stewards, the architecture provides a vocabulary for distinguishing conversational access from validated advice. Practical evaluation should ask whether data are current and interpretable, whether knowledge sources are curated, whether retrieval is relevant, whether uncertainty is understandable, and whether escalation pathways are reachable. Knowledge-guided model development may improve agricultural information processing, but its practical value remains dependent on corpus quality, retrieval performance, domain evaluation, and local context [15]. These implications describe areas for assessment and research; they are not instructions for selecting a platform, procuring a system, or making a farm-management decision.

Technology developers and platform providers should evaluate the complete advisory workflow rather than optimizing linguistic quality alone. Research is needed on multilingual farmer-query interpretation, low-bandwidth interaction, multimodal data quality, retrieval benchmarking, evidence-conflict handling, hallucination detection, advisor workload, farmer comprehension, and the effects of interface design on appropriate reliance. Advisory organizations should also investigate how professional roles, service relationships, and local knowledge change when conversational systems mediate access to agricultural information [17]. Such research should involve farmers and advisors before system requirements are fixed and should report failures, exclusions, and contextual limitations as carefully as technical performance.

Long-term research should connect local validation, security, governance, accessibility, and socio-ethical evaluation. Privacy and cybersecurity studies should examine threats across devices, communications, knowledge services, user accounts, audit records, and update mechanisms rather than treating the conversational interface as the only exposure point [24]. Institutional research should clarify accountability among developers, data custodians, advisory organizations, reviewers, and service operators. Responsible innovation should remain iterative, allowing use cases, safeguards, and system boundaries to be reconsidered as evidence and stakeholder concerns evolve [29]. These directions are conceptual research priorities, not legal, policy, procurement, engineering, or deployment prescriptions.

Conclusion

This article proposes a conceptual architecture for generative AI-enabled farm advisory systems that connects conversational decision support with farm and environmental data, curated agricultural knowledge, retrieval and grounding, generative model layers, advisory-output controls, uncertainty communication, validation, human oversight, expert escalation, farmer feedback, safety requirements, trust requirements, monitoring, and governed updating. Its central premise is that conversational fluency, model scale, multimodal capability, or source retrieval cannot independently establish factual reliability, agronomic validity, farmer trust, operational safety, or readiness for deployment. Responsible development requires bounded use cases, traceable evidence, contextual interpretation, farmer-centred interaction, meaningful human authority, privacy and cybersecurity protection, accountability, and explicit decision boundaries. The architecture is intended to structure future design, evaluation, and interdisciplinary research while preserving farmer autonomy and the indispensable role of agricultural expertise.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 14 Oct 2025

Revised: 31 Oct 2025

Accepted: 08 Dec 2025

Published online: 25 May 2026

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

Wolfert S, Ge L, Verdouw C, Bogaardt MJ. Big data in smart farming—A review. *Agric Syst.* 2017;153:69-80.

<https://doi.org/10.1016/j.agry.2017.01.023>.

Rose DC, Chilvers J. Agriculture 4.0: Broadening responsible innovation in an era of smart farming. *Front Sustain Food Syst.* 2018;2:87.

<https://doi.org/10.3389/fsufs.2018.00087>.

Klerkx L, Jakku E, Labarthe P. A review of social science on digital agriculture, smart farming and agriculture 4.0: New contributions and a future research agenda. *NJAS Wagening J Life Sci.* 2019;90-91:100315.

<https://doi.org/10.1016/j.njas.2019.100315>.

Zhai Z, Martínez JF, Beltran V, Martínez NL. Decision support systems for agriculture 4.0: Survey and challenges. *Comput Electron Agric.* 2020;170:105256.

<https://doi.org/10.1016/j.compag.2020.105256>.

Kuska MT, Wahabzada M, Paulus S. AI for crop production—Where can large language models (LLMs) provide substantial value?. *Comput Electron Agric.* 2024;221:108924.

<https://doi.org/10.1016/j.compag.2024.108924>.

De Clercq D, Nehring E, Mayne H, Mahdi A. Large language models can help boost food production, but be mindful of their risks. *Front Artif Intell.* 2024;7:1326153.

<https://doi.org/10.3389/frai.2024.1326153>.

Krupitzer C. Generative artificial intelligence in the agri-food value chain—Overview, potential, and research challenges. *Front Food Sci Technol.* 2024;4:1473357.

<https://doi.org/10.3389/frfst.2024.1473357>.

Shaikh TA, Rasool T, Venington K, Yaseen SM. The role of large language models in agriculture: Harvesting the future with LLM intelligence. *Prog Artif Intell.* 2025;14(2):117-64.

<https://doi.org/10.1007/s13748-024-00359-4>.

Koopman B, Mourad A, Li H, van der Vegt A, Zhuang S, Gibson S, et al. AgAsk: An agent to help answer farmer's questions from scientific documents. *Int J Digit Libr.* 2024;25(4):569-84.
<https://doi.org/10.1007/s00799-023-00369-y>.

Prestegaard-Wilson J, Vitale J. Generative artificial intelligence in extension: A new era of support for livestock producers. *Anim Front.* 2024;14(6):57-9.
<https://doi.org/10.1093/af/vfae024>.

Zhang W, Yadav A. Generative AI as an interface between farmers and AI algorithms: Opportunities and challenges. *J Indian Inst Sci.* 2025;105(2):305-21.
<https://doi.org/10.1007/s41745-025-00490-8>.

Eldem A, Eldem H. The development and evaluation of agricultural question-answering systems based on large language models. *Sci Rep.* 2026;16(1):5357.
<https://doi.org/10.1038/s41598-026-35003-9>.

Haghighat M, Saleh A, Rahimi Azghadi M. Multimodal language models in agriculture: A tutorial and survey. *Inf Fusion.* 2026;129:104042.
<https://doi.org/10.1016/j.inffus.2025.104042>.

Zhu H, Qin S, Su M, Lin C, Li A, Gao J. Harnessing large vision and language models in agriculture: A review. *Front Plant Sci.* 2025;16:1579355.
<https://doi.org/10.3389/fpls.2025.1579355>.

Jiang J, Yan L, Liu H, Xia Z, Wang H, Yang Y, et al. Knowledge assimilation: Implementing knowledge-guided agricultural large language model. *Knowl Based Syst.* 2025;314:113197.
<https://doi.org/10.1016/j.knosys.2025.113197>.

Shahriar S, Corradini MG, Sharif S, Moussa M, Dara R. The role of generative artificial intelligence in digital agri-food. *J Agric Food Res.* 2025;20:101787.
<https://doi.org/10.1016/j.jafr.2025.101787>.

Eastwood C, Ayre M, Nettle R, Dela Rue B. Making sense in the cloud: Farm advisory services in a smart farming future. *NJAS Wagening J Life Sci.* 2019;90-91:100298.
<https://doi.org/10.1016/j.njas.2019.04.004>.

Ding J, Jia X, Zhang W, Klerkx L. The effects of combined digital and human advisory services on reducing nitrogen fertilizer use: Lessons from China's national research programs on low-carbon agriculture. *Int J Agric Sustain.* 2022;20(6):1136-49.
<https://doi.org/10.1080/14735903.2022.2057643>.

Rose DC, Parker C, Fodey J, Park C, Sutherland WJ, Dicks LV. Involving stakeholders in agricultural decision support systems: Improving user-centred design. *Int J Agric Manag.* 2018;6(3-4):80-9.
<https://doi.org/10.5836/ijam/2017-06-80>.

Fielke S, Taylor B, Jakku E. Digitalisation of agricultural knowledge and advice networks: A state-of-the-art review. *Agric Syst.* 2020;180:102763.
<https://doi.org/10.1016/j.agry.2019.102763>.

Jakku E, Taylor B, Fleming A, Mason C, Fielke S, Sounness C, et al. "If they don't tell us what they do with it, why would we trust them?" Trust, transparency and benefit-sharing in smart farming. *NJAS Wagening J Life Sci.* 2019;90-91:100285.
<https://doi.org/10.1016/j.njas.2018.11.002>.

Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, et al. Survey of hallucination in natural language generation. *ACM Comput Surv*. 2023;55(12):248:1-38.
<https://doi.org/10.1145/3571730>.

Shneiderman B. Human-centered artificial intelligence: Reliable, safe and trustworthy. *Int J Hum Comput Interact*. 2020;36(6):495-504.
<https://doi.org/10.1080/10447318.2020.1741118>.

Gupta M, Abdelsalam M, Khorsandroo S, Mittal S. Security and privacy in smart farming: Challenges and opportunities. *IEEE Access*. 2020;8:34564-84.
<https://doi.org/10.1109/ACCESS.2020.2975142>.

Barredo Arrieta A, Díaz-Rodríguez N, Del Ser J, Bennetot A, Tabik S, Barbado A, et al. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf Fusion*. 2020;58:82-115.
<https://doi.org/10.1016/j.inffus.2019.12.012>.

Miller T. Explanation in artificial intelligence: Insights from the social sciences. *Artif Intell*. 2019;267:1-38.
<https://doi.org/10.1016/j.artint.2018.07.007>.

Mittelstadt B. Principles alone cannot guarantee ethical AI. *Nat Mach Intell*. 2019;1(11):501-7.
<https://doi.org/10.1038/s42256-019-0114-4>.

Ryan M. The social and ethical impacts of artificial intelligence in agriculture: Mapping the agricultural AI literature. *AI Soc*. 2023;38(6):2473-85.
<https://doi.org/10.1007/s00146-021-01377-9>.

Eastwood C, Klerkx L, Ayre M, Dela Rue B. Managing socio-ethical challenges in the development of smart farming: From a fragmented to a comprehensive approach for responsible research and innovation. *J Agric Environ Ethics*. 2019;32(5-6):741-68.
<https://doi.org/10.1007/s10806-017-9704-5>.