USOD PUBLISHERS
United Scientific Open Discovery

Journal of Advanced Artificial Intelligence Systems and Frontiers
Volume 1 | Issue 2

Research Article

Beyond Dashboards: Palm Agent-Human-Governed Agentic AI for Integrated Oil Palm Plantation-to-Mill Decision Intelligence

Loso Judijanto

IPOSS Jakarta, Indonesia

***Corresponding author:** Loso Judijanto IPOSS Jakarta, Indonesia.

Citation: Citation: Loso Judijanto (2026) Beyond Dashboards: Palm Agent-Human-Governed Agentic AI for Integrated Oil Palm Plantation-to-Mill Decision Intelligence. J of Adv Art Int Sys & Frontiers 1(2):1-12.

Received Date: October 08, 2026
Accepted Date: October 12, 2026
Published Date: November 11, 2026

Abstract

Digital transformation in plantation agriculture has progressed from electronic records and dashboards toward machine-learning prediction, digital twins, multimodal foundation models, and generative artificial intelligence. Nevertheless, most existing systems remain informational or task-specific: they visualize conditions, predict isolated outcomes, or provide recommendations that human users must reconcile manually. This qualitative literature review examines the emerging possibility of agentic artificial intelligence for vertically integrated oil palm plantation-to-mill operations. It distinguishes conversational large language models from genuinely agentic systems capable of retrieving evidence, decomposing objectives, invoking analytical tools, maintaining decision context, coordinating specialized agents, verifying outputs, and supporting controlled actions. The review synthesizes recent literature on LLM agents, multi-agent systems, retrieval-augmented generation, knowledge graphs, agricultural foundation models, digital twins, responsible AI, human–AI collaboration, and oil palm digital applications. Based on the synthesis, the article proposes PalmAgent, a human-governed multi-agent architecture comprising perception, data interoperability, knowledge grounding, reasoning and memory, analytical tools, specialist agents, authorization, and learning layers. Specialist Agronomy, Harvest, Logistics, Mill, Sustainability, and Financial Agents are coordinated by a supervisory orchestrator while consequential actions remain subject to risk-tiered human decision gates. The framework positions agentic AI not as a substitute for agronomists and managers but as an evidence-grounded orchestration layer connecting existing plantation intelligence. Research priorities include field validation, uncertainty calibration, cross-agent error containment, cybersecurity, multilingual interaction, organizational redesign, and measurement of operational, environmental, and economic outcomes.

Keywords: agentic AI; large language models; multi-agent systems; oil palm plantation; palm oil mill; decision intelligence; retrieval-augmented generation; knowledge graphs; human–AI collaboration; digital agriculture.

Introduction

From plantation digitalization to decision intelligence

Digital agriculture has evolved through several technological waves. Farm management information systems initially concentrated on collecting, storing, and presenting data for managerial use. Agriculture 4.0 subsequently expanded this model by integrating Internet of Things devices, cloud infrastructures, remote sensing, machine learning, robotics, geospatial information, and automated decision-support capabilities. The resulting systems can provide managers with increasingly rich representations of crop condition, weather, soil variability, machinery status, yield, resource use, and operational performance. Yet, the central operating logic of most systems has remained fundamentally human-centered in a specific sense: computers provide information or predictions, while people remain responsible for assembling multiple pieces of evidence into a coherent sequence of actions (Zhai et al., 2020; Giua, Materia and Camanzi, 2021; da Silveira, Lermen and Amaral, 2021; Yang et al., 2021; Shaikh, Rasool and Lone, 2022).

This limitation is particularly visible in operational dashboards and so-called control towers. Such systems can consolidate indicators that were previously distributed across disconnected databases and spreadsheets, improving visibility and situational awareness. However, visibility is not equivalent to reasoning. When a plantation dashboard simultaneously indicates a crop peak, labor shortage, deteriorating road condition, predicted rainfall, truck scarcity, and high mill utilization, a manager must still determine which indicators matter most, retrieve relevant standard operating procedures (SOPs), consult specialists, test alternatives, and authorize interventions. Conventional decision-support systems can assist some of these tasks, but they are commonly constructed around predefined workflows or specific models rather than flexible goal decomposition and cross-domain orchestration (Idoje et al., 2021; Runck et al., 2022; Purcell and Neubauer, 2023; Escribà-Gelonch et al., 2024).

The emergence of large language models (LLMs), multimodal foundation models, retrieval-augmented generation (RAG), tool-learning architectures, and multi-agent systems creates the possibility of another transition. Modern LLMs can act not only as generators of prose but also as interfaces to databases, application programming interfaces, search engines, geospatial systems, optimization models, and specialized software. More advanced architectures can decompose objectives, select tools, maintain memory, reason over retrieved evidence, delegate subtasks to other agents, assess intermediate outputs, and revise a plan when new evidence becomes available. These characteristics provide the basis for what is increasingly termed agentic AI (Wang et al., 2024; Li et al., 2024; Kumar, 2024; Acharya, Kuppan and Divya, 2025; Zhai et al., 2025; Xu et al., 2025; Chen et al., 2025; Abou Ali, Dornaika and Charafeddine, 2026; Chowa et al., 2026).

Why the oil palm plantation–mill system is an important test case

Oil palm represents an unusually demanding environment for agentic decision intelligence. Plantation performance emerges from interactions among perennial crop biology, soils, rainfall, fertilizer history, planting material, palm age, pests and diseases, harvesting rounds, labor availability, road conditions, transport

resources, fruit deterioration, mill reception, processing capacity, sustainability requirements, and financial constraints. Decisions also operate across markedly different temporal scales. Fruit evacuation can require decisions within hours; harvest planning occurs across days; nutrition, pest management, and infrastructure programs extend across months; while replanting, conservation, and capital planning involve multi-year horizons.

Existing oil palm digital research demonstrates substantial technical progress but also illustrates the fragmentation of current intelligence. Machine-learning models have been developed for yield prediction; remote sensing and Google Earth Engine have been applied to plantation mapping; deep learning has supported palm detection and counting; UAV imagery has enabled early identification of basal stem rot symptoms; image and geospatial platforms can classify bunch ripeness; IoT networks can monitor plantation soil conditions; optimization models can improve fresh fruit bunch (FFB) harvesting and evacuation; and blockchain or traceability architectures can strengthen environmental accounting and supply-chain transparency (Shaharum et al., 2020; Xu et al., 2021; Lim et al., 2021; Zeng et al., 2022; Putra and Wijayanto, 2023; Kent et al., 2023; El Hathat et al., 2023; Jamshidi et al., 2024; Puttinaovarat, Chai-Arayalert and Saetang, 2024; Khan et al., 2025; Sajak et al., 2025).

These applications are valuable, but most optimize or observe a bounded problem. A yield model does not automatically negotiate transport capacity with a mill schedule. A disease-detection model does not decide whether treatment should be prioritized relative to harvesting or water-management constraints. A routing model does not necessarily understand sustainability restrictions, estate SOPs, or a temporary maintenance shutdown at the mill. The principal opportunity for agentic AI is therefore not to replace these analytical systems, but to create an intelligent **orchestration layer** that can determine which tools and data are required, coordinate domain expertise, explain trade-offs, and package an evidence-grounded decision for authorized human approval.

Urgency, research gap and contribution

Agricultural foundation-model research is beginning to shift attention from narrow models toward generalized knowledge interaction, multimodal reasoning, agricultural question answering, knowledge graphs, and intelligent decision support. Recent literature specifically identifies domain adaptation, vector databases, multimodal data, and structured agricultural knowledge as important elements of future agricultural LLM ecosystems (Li et al., 2024; Yin et al., 2025; Li et al., 2025; Gong and Li, 2025; Shaikh, Rasool and Mir, 2025; Zhu et al., 2025).

However, three gaps remain. First, much agricultural LLM literature is still oriented toward question answering, textual knowledge services, diagnosis, or generative interfaces rather than persistent agents capable of planning and tool use. Second, research rarely conceptualizes multi-agent coordination across a vertically integrated perennial production system that includes both biological production and industrial processing. Third, discussions of autonomy often remain technically driven, whereas plantation decisions can carry agronomic, labor, financial, environmental, and regulatory consequences. Agricultural AI therefore requires an explicit theory of **decision rights**, not merely higher model performance (Dara, Hazrati Fard and Kaur, 2022; Ryan, 2023; Alexander, Yarborough and Smith, 2024).

This article responds to those gaps by proposing PalmAgent: Human-Governed Multi-Agent Intelligence for Plantation-to-Mill Operations. PalmAgent is conceptualized as a layered architecture combining operational data, authoritative knowledge, LLM reasoning, conventional machine-learning and optimization tools, specialized functional agents, explicit verification, and risk-based human authorization. The contribution is therefore not another oil palm predictive model. It is a conceptual architecture for transforming numerous existing models and information systems into a coordinated decision process.

The review addresses eight questions: how agricultural decision intelligence is evolving toward agentic systems; what differentiates agentic AI from conversational LLMs; how plantation-to-mill specialist agents should be coordinated; how evidence grounding should operate; which decisions are suitable for different levels of autonomy; how explainability and accountability should be institutionalized; which technical and organizational constraints affect deployment; and how agentic plantation systems should be evaluated. Together, these questions organize the thematic synthesis and guide the development of the PalmAgent framework.

Literature Review

From farm information systems to digital twins

The intellectual lineage of PalmAgent begins with farm management information and decision-support systems. Digital farm systems have historically been designed to acquire observations, transform them into usable information, and support managerial decisions. Their increasing connectivity has produced ecosystems in which sensors, machinery, weather information, remote sensing, enterprise systems, and analytics can contribute to a common management environment. Although this substantially improves information availability, interoperability, usability, data ownership, connectivity, and user competence remain longstanding adoption barriers (Giua, Materia and Camanzi, 2021; Idoje et al., 2021; Amiri-Zarandi et al., 2022).

Agricultural digital twins extend the information-system paradigm by creating dynamic virtual representations of physical agricultural entities and processes. Rather than only displaying measurements, a digital twin may integrate state sensing, simulation, prediction, and what-if analysis. Reviews of agricultural digital twins describe potential applications in crop development, irrigation, fertilization, pest management, infrastructure, robotics, and supply chains, while also emphasizing challenges arising from data heterogeneity and the need to orchestrate many simultaneous operations (Purcell and Neubauer, 2023; Escribà-Gelonch et al., 2024; Tagarakis et al., 2024; Wang, 2024).

Digital twins are especially important conceptually because they demonstrate the transition from visualization toward dynamic operational representation. Guidani, Ronzoni and Accorsi (2024), for example, extend digital-twin thinking to agri-food supply chains in which agricultural, processing, storage, distribution, economic, logistical, and environmental indicators can be evaluated jointly. Such architectures provide an important technological substrate for PalmAgent because an agent requires a reasonably faithful picture of current state before it can reason about interventions. Nevertheless, a digital twin and an agent solve different problems: the twin represents and simulates the system, whereas an agent determines which information or simulation is required to accomplish an objective. Recent work coupling an LLM with the DSSAT crop model further illustrates this hybrid logic, because the language model can translate user intent and orchestrate model interaction while the established simulation engine retains responsibility for agronomic computation (Kpodo et al., 2026).

Foundation models and generative AI in agriculture

Conventional agricultural machine learning has achieved strong results in classification, prediction, detection, segmentation, and

optimization. However, these models are normally trained for bounded tasks and their functionality cannot easily be generalized to unrelated problems. Agricultural foundation models promise a different approach: a sufficiently broad model can potentially adapt to multiple downstream tasks, integrate different modalities, and act as a common interface to agricultural knowledge (Li et al., 2024; Yin et al., 2025).

The most consequential development for decision systems is the combination of language and multimodal models. General multimodal LLM research has demonstrated architectures that connect language reasoning with visual inputs and other modalities, while agricultural reviews increasingly envision systems that jointly interpret text, images, sensor readings, maps, and structured knowledge (Yin et al., 2024; Zhu et al., 2025).

Generative AI is also expanding across the broader agri-food value chain. Potential applications include knowledge retrieval, disease and crop analysis, decision support, supply-chain optimization, explanation of model outputs, and new interfaces between operators and complex agricultural information systems. Yet recent reviews consistently identify data quality, domain adaptation, hallucination, privacy, multimodal integration, and trust as unresolved barriers (Krupitzer, 2024; Shahriar et al., 2025; Shaikh, Rasool and Mir, 2025). Agriculture-specific evaluations likewise show that performance depends on prompting strategy, domain framing, retrieval quality, and structured knowledge constraints rather than on model fluency alone (Eldem and Eldem, 2026; Ravindran et al., 2026).

The emerging agricultural LLM literature therefore supplies critical components for PalmAgent but does not by itself establish agentic behavior. An LLM answering "Which blocks have the highest fertilizer priority?" may produce a useful response. An agentic system, in contrast, would determine that the question requires recent leaf analysis, soil information, yield history, fertilizer stock, rainfall forecast, estate nutrition policy, budget constraints, and possibly a response-curve model; retrieve or call each source; identify missing data; compare alternatives; estimate consequences; flag uncertainty; and then present a recommendation at the appropriate authorization level.

What makes AI genuinely agentic?

The term agentic AI should not be used simply because an application contains an LLM. LLM-based agent research generally identifies a combination of perception, memory, reasoning, planning, action, feedback, and interaction with external tools or agents as defining capabilities. Agentic systems operate within a loop rather than a single prompt-response exchange: they observe or retrieve state, formulate or revise a plan, perform actions, inspect intermediate results, and continue until a goal or stopping criterion is satisfied (Wang et al., 2024; Acharya, Kuppan and Divya, 2025; Abou Ali, Dornaika and Charafeddine, 2026; Chowa et al., 2026; Khapre et al., 2026).

Task planning is especially important. Complex operational goals rarely map directly to one analytical call. LLM-based task-planning research frames planning as identifying and sequencing actions needed to achieve a desired outcome under constraints (Zhai et al., 2025). Likewise, tool-learning research describes pipelines involving task interpretation, tool retrieval or selection, execution, and synthesis of outputs. Tools allow an LLM to extend beyond its internal statistical knowledge by calling deterministic calculators, databases, search engines, APIs, optimization models, geospatial engines, or specialized machine-learning systems (Xu et al., 2025; Chen et al., 2025).

Multi-agent architectures introduce a further level of decomposition. Instead of expecting one model context to embody all domain rules and operational goals, different agents can be assigned different profiles, knowledge resources, tools, and permissions. Recent LLM multi-agent literature identifies perception, self-action, mutual interaction, and evolution as important elements of such ecosystems (Li et al., 2024). In plantation settings, this modularity is attractive because agronomy, harvesting, transport, processing, sustainability, and finance have distinct vocabularies, data structures, performance objectives, and acceptable risk levels.

Grounding: RAG, knowledge graphs and authoritative evidence

Agentic autonomy increases the importance of factual grounding. LLMs can produce fluent statements that are unsupported, incorrect, internally inconsistent, or outdated. Hallucination research differentiates failures of factuality from failures of faithfulness to provided context and shows that retrieval alone does not eliminate the problem (Huang et al., 2025; Lin et al., 2024).

For a plantation agent, internal model knowledge should consequently be treated as a reasoning resource rather than the sole source of agronomic truth. Retrieval-augmented generation can supply current SOPs, fertilizer policies, approved chemical lists, maintenance manuals, sustainability requirements, field records, or research literature at inference time. Knowledge graphs provide a complementary mechanism by encoding explicit entities and relationships—for example, estate-division-block-soil palm-age-treatment-disease-harvest relationships—that are difficult to represent reliably through unstructured text alone. Literature on knowledge graphs and LLMs suggests that structured knowledge can strengthen factual control, interpretability, and relational reasoning, although knowledge integration itself introduces unresolved alignment and updating challenges (Lavrinovics et al., 2025; Wang et al., 2025; Gong and Li, 2025). Recent agricultural RAG research similarly indicates that domain rules and structured validation layers can reduce factual and numerical errors, reinforcing the need for domain-governed retrieval rather than unconstrained generation (Ravindran et al., 2026).

Agricultural implementations are already moving toward multimodal knowledge structures. Lv et al. (2024), for example, demonstrate agricultural multimodal knowledge-graph construction integrating visual and textual information. Such research is relevant because PalmAgent must reason not only over documents but also over maps, images, structured observations, sensor streams, and model outputs. Grounding should therefore be conceived as an evidence fabric, not merely a vector database.

Human–AI collaboration and decision rights

A particularly important warning comes from human–AI decision research: combining humans and AI does not automatically produce superior decisions. Vaccaro, Almaatouq and Malone (2024) found that human–AI combinations in the reviewed experimental literature did not, on average, outperform the better of humans or AI, with decision tasks showing particular challenges. The implication is that PalmAgent should not be justified through the vague assumption that "human plus AI is better." The architecture must deliberately allocate subtasks according to comparative strengths.

Puranam (2021) frames this issue as an organization-design problem: collaboration requires an explicit division of labor between humans and algorithms. More recent interaction research similarly observes that many AI-assisted decision systems provide recommendations without creating genuinely collaborative interaction patterns (Gomez et al., 2025). Explainability can

improve performance when it helps domain experts interrogate predictions, but explanations alone do not resolve accountability, poor task allocation, or weak model governance (Ryo, 2022; Senoner et al., 2024).

For this reason, PalmAgent adopts **human governance by design**. Human involvement is not an emergency fallback after automation; it is an architectural layer specifying which agents may observe, analyze, recommend, simulate, request approval, or execute actions.

Method

Qualitative literature review design

This study applies a qualitative literature review (QLR) rather than a systematic literature review. The purpose is interpretive synthesis and conceptual framework development. Accordingly, the study does not claim exhaustive literature capture, calculate pooled effects, employ PRISMA reporting, or interpret publication counts as evidence of relative importance. Instead, relevant contemporary literature is brought together across several adjacent knowledge domains to explain how technological concepts developed elsewhere can be translated into an integrated oil palm decision architecture.

The review prioritizes peer-reviewed journal literature published between 2020 and 2026. Earlier theoretical traditions may inform interpretation, but the evidentiary corpus used in the manuscript emphasizes contemporary developments because foundation models, generative AI, LLM agents, and modern digital-twin architectures have changed rapidly during this period. Literature was considered from the areas of agentic AI, LLM agents, multi-agent systems, task planning, tool use, hallucination and trustworthiness, multimodal models, knowledge graphs, digital agriculture, agricultural foundation models, digital twins, responsible AI, human–AI collaboration, and oil palm digital technologies.

Interpretive synthesis

The analysis used thematic interpretation around seven questions: what an agent perceives; what it knows; how it reasons; which tools it can invoke; how agents coordinate; how actions are authorized; and how the system learns from outcomes. Literature from computer science provided agent architecture concepts, while agricultural and oil palm studies supplied domain constraints and examples of available analytical components.

The review then used cross-domain translation. For example, generic "tool use" was translated into calls to GIS systems, harvest forecasts, optimization solvers, weather services, mill databases, or financial models. "Memory" was translated into persistent operational context concerning blocks, previous recommendations, realized outcomes, and unresolved tasks. "Multi-agent collaboration" was translated into coordinated agronomy, harvest, logistics, mill, sustainability, and financial reasoning.

The final stage was conceptual integration into PalmAgent. This method is appropriate because the principal research gap is architectural: many necessary technologies already exist independently, but a coherent theory of how they should interact within a human-governed plantation-to-mill decision system remains underdeveloped.

Results: Thematic Findings

Theme 1: the transition from dashboards to cognitive orchestration

The first thematic finding is that agricultural digital intelligence can be represented as a maturity progression rather than as a binary distinction between "traditional" and "AI-enabled" management. Dashboards answer questions about what has happened or what is currently happening. Predictive systems estimate what is likely to happen. Prescriptive systems evaluate what should be done under a model's assumptions. Conversational copilots allow users to query information using natural language. Tool-using agents can initiate analytical operations, while multi-agent systems can decompose complex goals among specialized actors. Agentic intelligence therefore shifts the analytical unit from a prediction toward a **decision process**.

Figure 1 summarizes the conceptual transition at the core of the article. The progression should not be interpreted as meaning that newer layers eliminate earlier technologies. Instead, each successive layer incorporates and orchestrates capabilities developed at previous levels.



Figure 1: From dashboards to agentic plantation decision intelligence

Figure 1 shows why an agentic plantation architecture should not be reduced to a chatbot. Its most consequential feature is not natural-language generation but the ability to organize a sequence of analytical and organizational activities. The system can recognize that an answer depends on information it does not currently possess, retrieve that information, call the appropriate analytical tool, request input from another specialized agent, and revise the recommendation. This distinction is consistent with contemporary definitions of LLM agents, planning, and tool use (Wang et al., 2024; Zhai et al., 2025; Xu et al., 2025; Chowa et al., 2026).

Theme 2: oil palm AI already provides many of the required tools

The second finding is that PalmAgent need not begin from an empty analytical landscape. Numerous components already exist in the oil palm literature. Yield models can analyze agronomic and climatic variables and produce block or regional production forecasts. Jamshidi et al. (2024) demonstrate the value of comparing multiple machine-learning and deep-learning approaches using detailed agronomic information, while Khan et al. (2025) extend oil palm forecasting using long historical and agrometeorological time series. These models can become callable forecasting tools rather than standalone end-user applications.

Remote sensing constitutes another mature tool category. Shaharum et al. (2020) use Google Earth Engine and machine-learning classifiers for large-scale oil palm mapping, while Xu et al. (2021) and Zeng et al. (2022) illustrate the continuing development of satellite-based plantation identification using multi-source data. Putra and Wijayanto (2023) and Kipli et al. (2023) demonstrate deep-learning approaches to palm detection and counting. Such technologies could supply PalmAgent with spatial observations that are difficult to obtain from enterprise databases alone.

Computer vision can also contribute operationally. Kent et al. (2023) demonstrate early symptom detection of basal stem rot from UAV images, and Puttinaovarat, Chai-Arayalert and Saetang (2024) combine ripeness classification with geospatial plantation verification. Sensor infrastructures similarly provide live field observations; Sajak et al. (2025), for example, describe LoRa-based real-time soil monitoring for oil palm plantations. These technologies are important because agent quality is bounded by the freshness and reliability of perception.

Operational optimization is another available capability. Lim et al. (2021) formulate harvesting and evacuation-route optimization for FFB, demonstrating how operational planning can be represented mathematically. El Hathat et al. (2023) show how predictive modeling and blockchain can support emissions tracking across palm oil supply chains. In a PalmAgent architecture, such models would remain specialized tools. The LLM-based agent would decide when to call them and how to integrate their outputs with qualitative constraints, SOPs, managerial priorities, and other agents' findings.

Theme 3: specialist agents are preferable to one universal plantation agent

A vertically integrated plantation enterprise contains multiple decision domains that are related but not interchangeable. This suggests a multi-agent rather than monolithic architecture. The Agronomy Agent would focus on crop health, nutrition, soils, water, pests, diseases, and field treatments. It would use agronomic records, laboratory data, spatial observations, weather, field inspection, and scientific or internal technical knowledge.

The **Harvest Agent** would interpret crop forecasts, ripeness observations, harvesting-round status, labor availability, productivity, and block access. It should be capable of asking whether a predicted harvest peak is biologically plausible, which blocks require priority, and what labor reallocation might be feasible. A **Logistics Agent** would use road conditions, truck availability, payload, distance, queue times, and route optimization to evaluate FFB evacuation alternatives.

The **Mill Agent** would focus on crop reception, inventory, throughput, processing availability, maintenance status, extraction-related quality risks, and operating constraints. It would not necessarily optimize mill operations autonomously; instead, it would expose mill consequences to cross-functional planning. A **Sustainability Agent** would examine emissions, water, biodiversity, traceability, certification commitments, environmental restrictions, and policy obligations. Finally, a **Financial Agent** would translate operational alternatives into cost, budget, margin, working-capital, and investment consequences.

The agents require a **Supervising Orchestrator** because individually rational recommendations may conflict. Increasing harvest labor may make agronomic sense but exceed budget. Rapid evacuation may reduce fruit deterioration but be constrained by mill downtime. Opening a temporary access route may improve logistics while violating environmental safeguards. The orchestrator's function is therefore not to average recommendations mechanically but to identify conflicts, request additional analysis, enforce non-negotiable constraints, and create a decision package suitable for human approval.

Theme 4: PalmAgent requires a layered rather than model-centric architecture

A central result of the synthesis is that agentic intelligence should be treated as system architecture rather than as a single "AI model." Foundation-model research emphasizes the flexibility of generalized models, but operational reliability depends equally on data quality, interfaces, external tools, knowledge management, governance, and system evaluation (Li et al., 2024; Yin et al., 2025; Li et al., 2025).

Figure 2 illustrates the principal conceptual contribution of the review. It organizes PalmAgent into layers that separate observation, authoritative knowledge, probabilistic reasoning, specialist analytical tools, agent roles, and authorization.

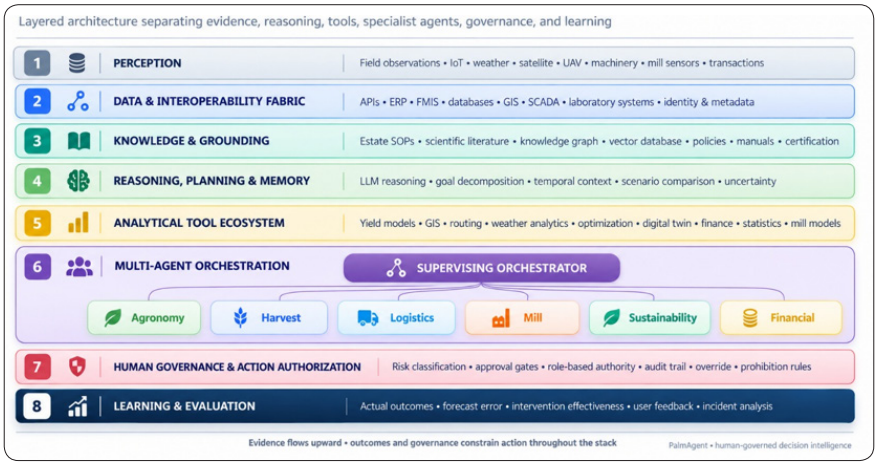


Figure 2: PalmAgent: Human-Governed Multi-Agent Intelligence for Plantation-to-Mill Operations

The layered design protects against a common misconception that the LLM itself should "know everything." PalmAgent instead treats the model as a reasoning and orchestration mechanism operating over external authoritative sources. Knowledge graphs can represent relationships and constraints; RAG can retrieve current documents; analytical tools can perform calculations; and human governance determines whether a resulting action is permissible. This architecture is particularly important because LLM hallucination and model staleness cannot be eliminated merely by increasing model scale (Huang et al., 2025; Lavrinovics et al., 2025; Huang et al., 2024).

Theme 5: cross-agent reasoning creates the principal system-level value

The main value of multi-agent orchestration becomes visible in scenarios that cross departmental boundaries. Consider a predicted FFB peak following several rainy days. The Harvest Agent could identify blocks likely to produce above-plan volumes and estimate labor requirements. The Logistics Agent could identify weak road segments, vehicle shortages, or excessive travel time. The Mill Agent could detect that a planned maintenance period reduces reception or sterilization capacity. The Financial Agent could estimate the costs of overtime, rented transport, delayed harvest, or lost quality, while the Sustainability Agent could prohibit routes or practices that conflict with environmental constraints.

A dashboard would present many of these facts. PalmAgent's potential advantage is that the supervising agent could formulate a decision sequence: verify the harvest forecast; determine priority blocks; request route optimization under road constraints; obtain mill capacity scenarios; quantify financial trade-offs; verify sustainability restrictions; and present two or three feasible plans with evidence, assumptions, uncertainties, and approval requirements.

Figure 3 emphasizes that coordination, rather than independent agent intelligence, is the defining feature of the proposed plantation-to-mill system. It therefore visualizes how domain-specific agents contribute to a single governed decision sequence rather than operating as isolated digital silos.

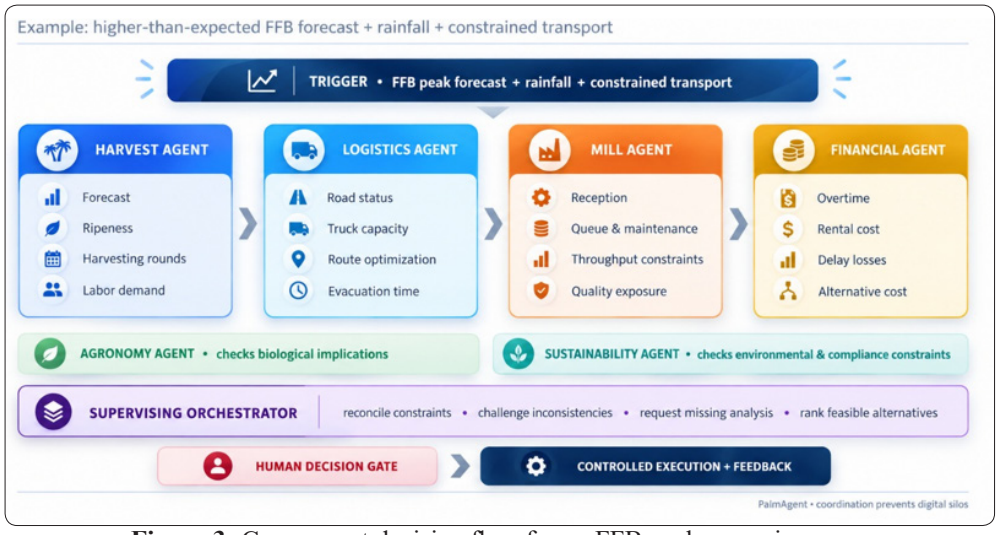


Figure 3: Cross-agent decision flow for an FFB peak scenario

The figure illustrates why multi-agent specialization should be paired with shared evidence and supervisory reconciliation. Without coordination, multiple agents could reproduce the organizational silo problem digitally. Multi-agent architecture is valuable only when it improves division of analytical labor and preserves a shared system objective (Li et al., 2024; Puranam, 2021).

Theme 6: human governance must be risk-tiered

Responsible agricultural AI literature emphasizes fairness, transparency, accountability, privacy, robustness, and sustainability. These requirements become even more important when systems progress from predictions toward tool execution. A recommendation error can be inconvenient; an automatically executed error can alter field operations or physical processes. Agricultural data infrastructures also involve privacy, data ownership, security, and access-control concerns that are magnified when agents can autonomously query or act across multiple systems (Dara, Hazrati Fard and Kaur, 2022; Amiri-Zarandi et al., 2022; Yazdinejad et al., 2021; Žalik and Žalik, 2023; Rezazadeh Baee et al., 2026).

PalmAgent should therefore use graded autonomy. Low-consequence tasks may be completed automatically. For instance, an agent could summarize yesterday's rainfall, identify blocks with missing inspections, or prepare a morning exception report. Medium-consequence tasks could generate recommendations but require confirmation before operational communication. High-consequence decisions—such as major chemical-use changes, substantial budget commitments, deviation from approved SOPs, or interventions with environmental or worker-safety implications—should require designated managerial or technical approval.

Some actions should remain outside autonomous authority irrespective of technical capability. This principle is important because "can the system execute?" is not the same question as "should the system be permitted to execute?". Human oversight thus becomes an enforceable decision-rights framework.

Figure 4 illustrates a governance model linking decision consequence to permissible autonomy. The model makes explicit that human accountability should increase as the operational, financial, environmental, regulatory, or safety consequences of an action become more material.



Figure 4: Risk-tiered human authorization model for PalmAgent

Risk should therefore determine autonomy rather than the apparent confidence of the LLM. This model also reduces the temptation to equate technological sophistication with managerial authority. Human–AI research indicates that effective collaboration requires appropriate allocation of tasks, while responsible agricultural AI research stresses accountability and clear governance throughout deployment (Vaccaro, Almaatouq and Malone, 2024; Alexander, Yarborough and Smith, 2024; Gomez et al., 2025).

Discussion and Analysis

Answering RQ1: from information visibility to orchestration

The literature supports a clear distinction among information systems, predictive models, and agentic decision architectures. Agricultural digitalization has progressively improved data capture and analytics, while digital twins bring dynamic representations and what-if analysis closer to physical operations. Agentic AI extends this progression by adding goal decomposition, tool selection, iterative planning, and coordination across analytical resources.

PalmAgent should therefore be understood as an additional layer rather than a wholesale replacement. Existing dashboards remain valuable for visual supervision. Predictive models remain necessary for forecasting. Optimization algorithms remain preferable to general-purpose LLMs for rigorously defined mathematical problems. GIS engines remain superior for spatial computation, and process-control systems should continue to handle deterministic control loops. Agentic AI adds value when a managerial objective requires **selecting and sequencing multiple forms of intelligence**.

This distinction is important for avoiding technological overreach. The strongest architecture is hybrid: deterministic tools perform tasks where deterministic computation is appropriate, specialized machine-learning models perform bounded predictions, digital twins represent system dynamics, knowledge systems provide evidence, and LLM agents orchestrate interactions among these resources. Tool-learning literature strongly supports such a decomposition because external tool use is intended to overcome limitations in the model's embedded knowledge and native calculation abilities, while emerging agricultural systems show the practical value of coupling LLM interfaces with validated crop-simulation engines rather than substituting for them (Xu et al., 2025; Chen et al., 2025; Kpodo et al., 2026).

Answering RQ2: conversational capability is necessary but insufficient

A conversational interface improves accessibility but does not create agency. A system becomes meaningfully agentic when it can maintain an objective across multiple steps, recognize information gaps, create a plan, retrieve current evidence, invoke appropriate external tools, inspect intermediate outcomes, revise its plan, coordinate with other agents, and stop or escalate when confidence or permission is insufficient.

This distinction has practical consequences for evaluation. A conventional agricultural chatbot can be benchmarked partly through answer accuracy and relevance, as recent agriculture-specific LLM studies have done across question-answering tasks (Eldem and Eldem, 2026). PalmAgent requires additional metrics, including correct task decomposition, tool selection, API-call accuracy, evidence fidelity, constraint satisfaction, plan feasibility, recognition of missing data, error recovery, cross-agent consistency, and compliance with human authorization rules. Recent agentic-AI taxonomies likewise emphasize evaluation at architectural, cognitive, security, and system-performance levels rather than relying on a single language-model score (Khapre et al., 2026).

The article therefore rejects a superficial architecture in which six differently named chatbots simulate specialist agents while sharing the same undifferentiated knowledge and permissions. Genuine specialization should involve different retrieval collections, tool access, operating instructions, constraints, performance indicators, and escalation pathways.

Answering RQ3: why specialization matters in vertically integrated plantations

Oil-palm enterprises create an especially strong rationale for agent specialization because plantation and mill activities operate under different physical and managerial logics. Agronomy is dominated by biological response, spatial variability, and lagged effects. Harvest management is shaped by ripeness, harvesting intervals, labor productivity, and crop availability. Logistics is a network and resource-allocation problem. Milling is an industrial processing problem with throughput, maintenance, quality, energy, and inventory constraints. Sustainability introduces boundary conditions that may override local economic optimization. Finance converts operational choices into organizational trade-offs.

Specialist agents create modularity. New models or policies can be added to one agent without redesigning the entire system. Permissions can differ: the Sustainability Agent may have read access to sensitive compliance documentation but no authority to modify operational records; the Financial Agent may access cost models but not worker personal data. Domain experts can evaluate the agent most closely associated with their responsibilities, improving validation and accountability.

However, specialization also creates coordination risk. The system could generate internally incompatible recommendations, agents could rely on different versions of the same dataset, or one incorrect output could cascade through subsequent agents. PalmAgent therefore requires a shared ontology, common timestamps, provenance metadata, confidence representation, and an orchestrator that distinguishes hard constraints from negotiable objectives. These requirements are consistent with recent agentic-AI work that treats coordination, security, explainability, and system-level evaluation as interdependent architectural concerns (Wang et al., 2024; Khapre et al., 2026).

Answering RQ4: evidence grounding should be multi-layered
RAG is necessary but insufficient for plantation decision grounding. A useful PalmAgent evidence architecture should contain at least five forms of authority.

First, **normative internal knowledge** includes SOPs, work instructions, delegated authorities, agronomic standards, approved product lists, and technical circulars. Second, **scientific knowledge** consists of peer-reviewed research and curated domain references. Third, **operational truth** comes from current enterprise systems: block records, harvest data, stock, maintenance, laboratory analyses, mill status, and workforce availability. Fourth, **environmental and spatial evidence** comes from weather, remote sensing, GIS, UAVs, and IoT systems. Fifth, **computed evidence** comprises outputs from validated predictive, optimization, financial, and simulation tools.

The system should preserve provenance for each category. When PalmAgent recommends accelerating harvest in a particular division, users should be able to see which crop forecast, field observations, road data, mill constraints, SOP clauses, and model outputs supported the conclusion. If the recommendation depends on a source that is outdated or missing, the agent should state the deficiency rather than silently infer a value.

Knowledge graphs are valuable because plantation information is relational. A block belongs to a division, contains palms of a certain planting year, has soil and topographic characteristics, receives treatments, produces harvest events, connects to roads and collection points, and contributes FFB to particular mills. Representing these relationships explicitly can improve retrieval and constraint checking compared with treating every source as unstructured text (Lavrinovics et al., 2025; Wang et al., 2025).

Answering RQ5: defining appropriate autonomy
The most suitable early applications of PalmAgent are likely to be **decision preparation** rather than physical autonomy. Examples include daily exception synthesis, investigation of production variances, harvest-plan scenario generation, fertilizer-program evidence compilation, FFB evacuation planning, mill–estate coordination, sustainability documentation checks, and financial scenario comparison.

These applications are attractive because they involve high cognitive coordination but can preserve human authority over execution. They can also create measurable benefits through reduced search time, faster cross-functional coordination, more systematic evidence use, and better traceability.

Greater autonomy may eventually be appropriate for bounded, reversible, low-risk tasks such as generating routine reports or creating draft schedules under validated rules. The threshold for autonomous execution should rise sharply when an intervention becomes expensive, environmentally consequential, safety-relevant, difficult to reverse, or dependent on uncertain biological assumptions.

This conservative architecture is supported by the broader trustworthiness literature. LLM systems remain vulnerable to hallucination, input manipulation, outdated evidence, unexpected tool behavior, and cascading errors. Verification and validation therefore have to operate throughout the system lifecycle rather than being confined to model testing before deployment (Huang et al., 2024; Lin et al., 2024).

Answering RQ6: human supervision must become operational governance
"Human in the loop" is too vague for a high-consequence plantation system. PalmAgent requires defined roles. The requesting user specifies objectives, while a domain-accountable specialist retains ownership of the technical decision. An approver may hold delegated financial or operational authority. System administrators govern digital permissions but should not automatically possess agronomic decision authority. Audit personnel may inspect decision provenance without participating in the original decision.

Each agent recommendation should therefore include an authorization label. For example: "advisory only," "supervisor approval required," "estate manager approval required," or "technical authority approval required." The system must not infer its own authority from linguistic confidence.

A complete audit record should ideally preserve the request, retrieved sources, tool calls, agent messages, major intermediate conclusions, data versions, human modifications, final approval, execution status, and observed outcome. This transforms explainability from a narrative feature into operational accountability.

Answering RQ7: technical and organizational barriers
Data quality and interoperability
PalmAgent depends on consistent identifiers and reliable timestamps across estate, mill, GIS, sensor, and financial systems. In many agricultural environments, digital transformation has proceeded through separate applications acquired at different times. Duplicate block names, different coding schemes, missing metadata, manual spreadsheets, and inconsistent spatial boundaries can undermine agent reasoning even if the LLM itself is highly capable.

Interoperability is therefore a precondition rather than a secondary engineering problem. A sophisticated agent cannot compensate safely for unclear master data. Agricultural information-system and platform literature repeatedly identifies interoperability, standards, and data governance as core challenges (Zhai et al., 2020; Runck et al., 2022; Rezazadeh Baee et al., 2026).

Cybersecurity and tool permissions
Agentic systems enlarge the attack surface because they can connect to many tools and repositories. Prompt injection or compromised retrieved documents could attempt to manipulate an agent into revealing sensitive information or invoking unauthorized tools. Tool permissions should therefore follow least-

privilege principles. Read access should be separated from write or execution access, and each high-risk call should be authenticated, logged, and checked against policy.

Smart-farming cybersecurity literature already documents threats associated with connected devices, networks, and agricultural platforms. Agentic architectures add a new layer because an attacker may seek to manipulate not only data but the reasoning sequence linking data to action (Yang et al., 2021; Yazdinejad et al., 2021).

Privacy and organizational power

Plantation data may contain operationally sensitive information, employee-related information, location data, commercial costs, supplier information, and potentially proprietary agronomic practices. Agricultural data-governance literature demonstrates that control, ownership, privacy, and trust can significantly influence digital-technology adoption. Federated or privacy-preserving approaches may therefore become relevant where multi-estate learning requires information sharing without centralizing raw records (Amiri-Zarandi et al., 2022; Žalik and Žalik, 2023).

Multilingual and multi-level communication

Plantation enterprises also have a communication problem that conventional AI architecture discussions often overlook. Senior managers, scientists, field assistants, operators, contractors, and workers may require different levels of technical detail and may use different languages. Multilingual LLM capability creates a possibility for translating the same evidence into managerial, technical, and operational formats while preserving the underlying decision logic.

Yet simplification should not alter substantive meaning. A field instruction derived from an approved recommendation must maintain quantities, timing, safety conditions, and restrictions precisely. Multilingual performance should thus be validated as a safety-critical interface issue rather than treated merely as user convenience.

Answering RQ8: how PalmAgent should be evaluated

Evaluation of PalmAgent should extend far beyond generic LLM benchmarks. At least six levels are required. Model-level evaluation examines factuality, instruction adherence, calibration, and hallucination. Retrieval-level evaluation examines whether the system retrieves the correct and most current SOPs, records, and evidence. Tool-level evaluation measures tool selection, parameter accuracy, execution success, and correct interpretation of outputs. Agent-level evaluation measures task decomposition, memory, error recovery, and domain-constraint compliance. System-level evaluation examines coordination, latency, resilience, security, escalation behavior, and prevention of unauthorized actions. Finally, business and sustainability evaluation determines whether the system improves actual plantation outcomes; this layered approach is aligned with recent agentic-AI evaluation taxonomies that explicitly combine performance, explainability, security, privacy, and safety dimensions (Khapre et al., 2026).

The last level is decisive. An agent that produces impressive textual reasoning but fails to improve decision quality or erodes user trust is not successful. Relevant outcome indicators could include FFB evacuation time, overdue crop, fruit quality, forecast error, harvesting productivity, mill waiting time, extraction-related outcomes, fertilizer efficiency, environmental incidents, cost variance, managerial decision time, the number of unresolved exceptions, and auditability. Field experiments should compare at least three conditions where feasible: existing decision processes, AI-assisted processes without orchestration, and PalmAgent-assisted processes. Such designs would help establish whether the multi-agent architecture provides additional value rather than simply introducing technological novelty.

Research agenda

The literature synthesis suggests eight priority research streams. First, **domain ontology and plantation knowledge-graph research** should establish standardized representations of blocks, palms, operations, agronomic events, assets, and processing relationships. Second, **agent-grounding research** should compare different combinations of RAG, knowledge graphs, and structured SQL/API access.

Third, **cross-agent reliability research** should examine how errors propagate and how agents can challenge one another constructively. Fourth, **uncertainty research** should distinguish model confidence, data quality, source disagreement, and operational uncertainty rather than compressing all uncertainty into a single score. Fifth, **human–agent interaction research** should test how recommendations, explanations, alternatives, and approval gates affect decision quality.

Sixth, **organizational research** should study how agentic AI changes roles of agronomists, estate managers, mill managers, controllers, sustainability teams, and IT functions. Seventh, **economics research** should determine when the additional infrastructure and governance costs of agentic AI are justified. Finally, **longitudinal field validation** should examine whether system performance remains reliable across crop cycles, seasonal transitions, unusual weather, model updates, policy changes, and organizational turnover.

Conclusion and Recommendations

The transition from dashboards to agentic AI represents a potentially important shift in agricultural decision-system design. The defining change is not that managers can communicate with software using natural language. It is that an AI system can participate in the structure of decision work: decomposing objectives, retrieving evidence, invoking validated analytical tools, coordinating specialized perspectives, comparing alternatives, verifying constraints, escalating uncertainty, and preparing controlled actions.

Oil palm plantation-to-mill operations provide a compelling setting for this transition because biological, spatial, logistical, industrial, sustainability, and financial decisions are continuously interdependent. Contemporary oil palm research already provides many of the building blocks required for an agentic system, including yield prediction, remote sensing, UAV analysis, IoT monitoring, routing optimization, geospatial platforms, and supply-chain traceability. The major unresolved challenge is how those components should be integrated into a coherent cognitive and organizational architecture.

This review proposes **PalmAgent** as that architecture. PalmAgent combines a perception layer, interoperable operational data, authoritative knowledge, LLM-based reasoning and memory, external analytical tools, specialized Agronomy, Harvest, Logistics, Mill, Sustainability, and Financial Agents, a Supervising Orchestrator, and a risk-tiered human authorization layer. The framework deliberately places governance above autonomy. High-quality evidence, transparent provenance, clear decision rights, role-based permissions, and human accountability are fundamental requirements rather than optional controls.

The framework also rejects technological substitution as its central objective. Agronomists, estate managers, mill specialists, sustainability professionals, and financial managers possess contextual knowledge, authority, and accountability that an

LLM cannot simply inherit. PalmAgent should instead reduce the transaction cost of finding evidence, coordinating expertise, testing alternatives, and documenting decisions.

Future research should therefore move beyond demonstrations of agricultural question answering toward measurable operational trials. The scientific question is no longer only whether LLMs can discuss agriculture convincingly. It is whether a governed multi-agent system can improve the quality, speed, traceability, resilience, and sustainability of real plantation decisions without creating unacceptable new risks. Establishing that evidence would mark the transition from generative agricultural interfaces to genuinely trustworthy agentic decision intelligence.

References

1. Abou Ali M, Dornaika F, Charafeddine J (2026) Agentic AI: a comprehensive survey of architectures, applications, and future directions. Artificial Intelligence Review 59: 11. https://doi.org/10.1007/s10462-025-11422-4
2. Acharya DB, Kuppan K, Divya B (2025) Agentic AI: autonomous intelligence for complex goals—a comprehensive survey. IEEE Access 13: 18912–18936. https://doi.org/10.1109/ACCESS.2025.3532853
3. Alexander CS, Yarborough M, Smith A (2024) Who is responsible for "responsible AI"?: navigating challenges to build trust in AI agriculture and food system technology. Precision Agriculture 25: 146–185. https://doi.org/10.1007/s11119-023-10063-3
4. Amiri-Zarandi M, Dara RA, Duncan E, Fraser EDG (2022) Big data privacy in smart farming: a review. Sustainability 14(15): 9120. https://doi.org/10.3390/su14159120
5. Chen J, Wu H, Pang J, et al. (2025) Tool learning with language models: a comprehensive survey of methods, pipelines, and benchmarks. Vicinagearth 2: 16. https://doi.org/10.1007/s44336-025-00024-x
6. Chowa SS, Alvi R, Rahman SS, et al. (2026) From language to action: a review of large language models as autonomous agents and tool users. Artificial Intelligence Review 59: 71. https://doi.org/10.1007/s10462-025-11471-9
7. da Silveira F, Lermen FH, Amaral FG (2021) An overview of agriculture 4.0 development: systematic review of descriptions, technologies, barriers, advantages, and disadvantages. Computers and Electronics in Agriculture 189: 106405. https://doi.org/10.1016/j.compag.2021.106405
8. Dara R, Hazrati Fard SM, Kaur J (2022) Recommendations for ethical and responsible use of artificial intelligence in digital agriculture. Frontiers in Artificial Intelligence 5: 884192. https://doi.org/10.3389/frai.2022.884192
9. El Hathat Z, et al. (2023) Analyzing the greenhouse gas emissions in the palm oil supply chain in the VUCA world: a blockchain initiative. Business Strategy and the Environment. https://doi.org/10.1002/bse.3436
10. Eldem A, Eldem H (2026) The development and evaluation of agricultural question-answering systems based on large language models. Scientific Reports 16: 5357. https://doi.org/10.1038/s41598-026-35003-9
11. Escribà-Gelonch M, et al. (2024) Digital twins in agriculture: orchestration and applications. Journal of Agricultural and Food Chemistry 72(19): 10737–10752. https://doi.org/10.1021/acs.jafc.4c01934
12. Giua C, Materia VC, Camanzi L (2021) Management information system adoption at the farm level: evidence from the literature. British Food Journal 123(3): 884–909. https://doi.org/10.1108/BFJ-05-2020-0420
13. Gomez C, Cho SM, Ke S, Huang C-M, Unberath M (2025) Human-AI collaboration is not very collaborative yet: a taxonomy of interaction patterns in AI-assisted decision making from a systematic review. Frontiers in Computer Science 6: 1521066. https://doi.org/10.3389/fcomp.2024.1521066
14. Gong R, Li X (2025) The application progress and research trends of knowledge graphs and large language models in agriculture. Computers and Electronics in Agriculture 235: 110396. https://doi.org/10.1016/j.compag.2025.110396
15. Guidani B, Ronzoni M, Accorsi R (2024) Virtual agri-food supply chains: a holistic digital twin for sustainable food ecosystem design, control and transparency. Sustainable Production and Consumption 46: 161–179. https://doi.org/10.1016/j.spc.2024.01.016
16. Huang L, Yu W, Ma W, et al. (2025) A survey on hallucination in large language models: principles, taxonomy, challenges, and open questions. ACM Transactions on Information Systems 43(2): Article 42, 1–55. https://doi.org/10.1145/3703155
17. Huang X, Ruan W, Huang W, et al. (2024) A survey of safety and trustworthiness of large language models through the lens of verification and validation. Artificial Intelligence Review 57: 175. https://doi.org/10.1007/s10462-024-10824-0
18. Idoje G, et al. (2021) Survey for smart farming technologies: challenges and issues. Computers and Electrical Engineering 92: 107104. https://doi.org/10.1016/j.compeleceng.2021.107104
19. Jamshidi EJ, Yusup Y, Hooy CW, et al. (2024) Predicting oil palm yield using a comprehensive agronomy dataset and 17 machine learning and deep learning models. Ecological Informatics 81: 102595. https://doi.org/10.1016/j.ecoinf.2024.102595
20. Kent OW, Chun TW, Choo TL, Kin LW (2023) Early symptom detection of basal stem rot disease in oil palm trees using a deep learning approach on UAV images. Computers and Electronics in Agriculture 213: 108192. https://doi.org/10.1016/j.compag.2023.108192
21. Khan N, Kamaruddin MA, Sheikh UU, Rahman AAH, Bakht MP (2025) Prediction of oil palm yield using machine learning: comparison of linear and non-linear algorithms with multivariate time series data. Journal of Oil Palm Research. https://doi.org/10.21894/jopr.2025.0044
22. Khapre S, Mersha MA, Olalekan Z, Shakil H, Iskander S, Moustafa A, Kalita J (2026) Agentic AI systems: a systematic survey of multi-agent architectures, cognitive foundations, interaction, explainability, security, and performance evaluation. Neurocomputing 696: 134049. https://doi.org/10.1016/j.neucom.2026.134049
23. Kipli K, et al. (2023) Deep learning applications for oil palm tree detection and counting. Smart Agricultural Technology 5: 100241. https://doi.org/10.1016/j.atech.2023.100241
24. Kpodo J, Nejadhashemi AP, Kordjamshidi P, Eeswaran R, Dong Y (2026) Closing the gap between crop simulation and practical decision-making. Computers and Electronics in Agriculture 247: 111655. https://doi.org/10.1016/j.compag.2026.111655
25. Krupitzer C (2024) Generative artificial intelligence in the agri-food value chain: overview, potential, and research challenges. Frontiers in Food Science and Technology 4: 1473357. https://doi.org/10.3389/frfst.2024.1473357
26. Kumar P (2024) Large language models (LLMs): survey, technical frameworks, and future challenges. Artificial Intelligence Review 57: 260. https://doi.org/10.1007/s10462-024-10888-y
27. Lavrinovics E, Biswas R, Bjerva J, Hose K (2025) Knowledge graphs, large language models, and hallucinations: an NLP perspective. Journal of Web Semantics 85: 100844. https://doi.org/10.1016/j.websem.2024.100844

28. Li H, Wu H, Li Q, Zhao C (2025) A review on enhancing agricultural intelligence with large language models. Artificial Intelligence in Agriculture 15(4): 671–685. https://doi.org/10.1016/j.aiia.2025.05.006
29. Li J, Xu M, Xiang L, et al. (2024) Foundation models in smart agriculture: basics, opportunities, and challenges. Computers and Electronics in Agriculture 222: 109032. https://doi.org/10.1016/j.compag.2024.109032
30. Li X, Wang S, Zeng S, Wu Y, Yang Y (2024) A survey on LLM-based multi-agent systems: workflow, infrastructure, and challenges. Vicinagearth 1: 9. https://doi.org/10.1007/s44336-024-00009-2
31. Lim CH, Cheah ZH, Lee XH, et al. (2021) Harvesting and evacuation route optimisation model for fresh fruit bunch in the oil palm plantation site. Journal of Cleaner Production 307: 127238. https://doi.org/10.1016/j.jclepro.2021.127238
32. Lin Z, Guan S, Zhang W, et al. (2024) Towards trustworthy LLMs: a review on debiasing and dehallucinating in large language models. Artificial Intelligence Review 57: 243. https://doi.org/10.1007/s10462-024-10896-y
33. Lv B, Wu H, Chen W, Chen C, Miao Y, Zhao C (2024) VEG-MMKG: multimodal knowledge graph construction for vegetables based on pre-trained model extraction. Computers and Electronics in Agriculture 226: 109398. https://doi.org/10.1016/j.compag.2024.109398
34. Puranam P (2021) Human–AI collaborative decision-making as an organization design problem. Journal of Organization Design 10: 75–80. https://doi.org/10.1007/s41469-021-00095-2
35. Purcell W, Neubauer T (2023) Digital twins in agriculture: a state-of-the-art review. Smart Agricultural Technology 3: 100094. https://doi.org/10.1016/j.atech.2022.100094
36. Putra YC, Wijayanto AW (2023) Automatic detection and counting of oil palm trees using remote sensing and object-based deep learning. Remote Sensing Applications: Society and Environment 29: 100914. https://doi.org/10.1016/j.rsase.2022.100914
37. Puttinaovarat S, Chai-Arayalert S, Saetang W (2024) Oil palm bunch ripeness classification and plantation verification platform: leveraging deep learning and geospatial analysis and visualization. ISPRS International Journal of Geo-Information 13(5): 158. https://doi.org/10.3390/ijgi13050158
38. Ravindran DJS, Skarga-Bandurova I, Sivakumar V, Awais M, Mithra S (2026) AgroLLM: connecting farmers and agricultural practices through large language models for enhanced knowledge transfer and practical application. AgriEngineering 8(1): 38. https://doi.org/10.3390/agriengineering8010038
39. Rezazadeh Baee MA, Simpson L, Kane L, et al. (2026) Toward data trustworthiness in digital agriculture: taxonomy and building blocks. EURASIP Journal on Wireless Communications and Networking 2026: 19. https://doi.org/10.1186/s13638-025-02553-x
40. Runck BC, Joglekar A, Silverstein KAT, Chan-Kang C, Pardey PG, Wilgenbusch JC (2022) Digital agriculture platforms: driving data-enabled agricultural innovation in a world fraught with privacy and security concerns. Agronomy Journal 114(5): 2635–2643. https://doi.org/10.1002/agj2.20873
41. Ryan M (2023) The social and ethical impacts of artificial intelligence in agriculture: mapping the agricultural AI literature. AI & Society 38: 2473–2485. https://doi.org/10.1007/s00146-021-01377-9
42. Ryo M (2022) Explainable artificial intelligence and interpretable machine learning for agricultural data analysis. Artificial Intelligence in Agriculture: 257–265. https://doi.org/10.1016/j.aiia.2022.11.003
43. Sajak AAB, Shukor NFM, Kassim MS, Rahman SR, Aziz AA (2025) Comprehensive LoRa based IoT real-time soil monitoring for oil palm plantation. International Journal of Technology 16(5): 1453–1466. https://doi.org/10.14716/ijtech.v16i5.5139
44. Senoner J, Schallmoser S, Kratzwald B, et al. (2024) Explainable AI improves task performance in human–AI collaboration. Scientific Reports 14: 31150. https://doi.org/10.1038/s41598-024-82501-9
45. Shaharum NSN, Shafri HZM, Ghani WAWAK, et al. (2020) Oil palm mapping over Peninsular Malaysia using Google Earth Engine and machine learning algorithms. Remote Sensing Applications: Society and Environment 17: 100287. https://doi.org/10.1016/j.rsase.2020.100287
46. Shahriar S, Corradini MG, Sharif S, Moussa M, Dara R (2025) The role of generative artificial intelligence in digital agri-food. Journal of Agriculture and Food Research 20: 101787. https://doi.org/10.1016/j.jafr.2025.101787
47. Shaikh TA, Rasool T, Lone FR (2022) Towards leveraging the role of machine learning and artificial intelligence in precision agriculture and smart farming. Computers and Electronics in Agriculture 198: 107119. https://doi.org/10.1016/j.compag.2022.107119
48. Shaikh TA, Rasool T, Mir WA (2025) Fields of the future: digital transformation in smart agriculture with large language models and generative AI. Computer Standards & Interfaces 94: 104005. https://doi.org/10.1016/j.csi.2025.104005
49. Tagarakis AC, Benos L, Kyriakarakos G, Pearson S, Sørensen CG, Bochtis D (2024) Digital twins in agriculture and forestry: a review. Sensors 24(10): 3117. https://doi.org/10.3390/s24103117
50. Vaccaro M, Almaatouq A, Malone T (2024) When combinations of humans and AI are useful: a systematic review and meta-analysis. Nature Human Behaviour 8: 2293–2303. https://doi.org/10.1038/s41562-024-02024-1
51. Wang L (2024) Digital twins in agriculture: a review of recent progress and open issues. Electronics 13(11): 2209. https://doi.org/10.3390/electronics13112209
52. Wang L, Ma C, Feng X, et al. (2024) A survey on large language model based autonomous agents. Frontiers of Computer Science 18: 186345. https://doi.org/10.1007/s11704-024-40231-1
53. Wang X, Chen Z, Wang H, et al. (2025) Large language model enhanced knowledge representation learning: a survey. Data Science and Engineering 10: 315–338. https://doi.org/10.1007/s41019-025-00285-y
54. Xu K, Qian J, Hu Z, et al. (2021) A new machine learning approach in detecting the oil palm plantations using remote sensing data. Remote Sensing 13(2): 236. https://doi.org/10.3390/rs13020236
55. Xu W, Huang C, Gao S, et al. (2025) LLM-based agents for tool learning: a survey. Data Science and Engineering 10: 533–563. https://doi.org/10.1007/s41019-025-00296-9
56. Yang X, Shu L, Chen J, Ferrag MA, Wu J, Nurellari E, Huang K (2021) A survey on smart agriculture: development modes, technologies, and security and privacy challenges. IEEE/CAA Journal of Automatica Sinica 8(2): 273–302. https://doi.org/10.1109/JAS.2020.1003536
57. Yazdinejad A, Zolfaghari B, Azmoodeh A, et al. (2021) A review on security of smart farming and precision agriculture: security aspects, attacks, threats and countermeasures. Applied Sciences 11(16): 7518. https://doi.org/10.3390/app11167518
58. Yin S, Fu C, et al. (2024) A survey on multimodal large language models. National Science Review 11(12): nwae403. https://doi.org/10.1093/nsr/nwae403
59. Yin S, Xi Y, Zhang X, Sun C, Mao Q (2025) Foundation

models in agriculture: a comprehensive review. Agriculture 15(8): 847. https://doi.org/10.3390/agriculture15080847
60. Žalik KR, Žalik M (2023) A review of federated learning in agriculture. Sensors 23(23): 9566. https://doi.org/10.3390/s23239566
61. Zeng J, Tan ML, Tew YL, et al. (2022) Optimization of open-access optical and radar satellite data in Google Earth Engine for oil palm mapping in the Muda River Basin, Malaysia. Agriculture 12(9): 1435. https://doi.org/10.3390/agriculture12091435
62. Zhai W, Liao J, Chen Z, Su B, Zhao X (2025) A survey of task planning with large language models. Intelligent Computing 4: 0124. https://doi.org/10.34133/icomputing.0124
63. Zhai Z, Martínez JF, Beltran V, Martínez NL (2020) Decision support systems for agriculture 4.0: survey and challenges. Computers and Electronics in Agriculture 170: 105256. https://doi.org/10.1016/j.compag.2020.105256.
64. Zhu H, Qin S, Su M, Lin C, Li A, Gao J (2025) Harnessing large vision and language models in agriculture: a review. Frontiers in Plant Science 16: 1579355. https://doi.org/10.3389/fpls.2025.1579355.

Copyright: © 2026 Loso Judijanto. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.