

SmartFarming: a GIS-based agent framework for farm management

Ding Yang^{#a}, Yancheng Wang^{#a}, Keqiang Sun^{#a}, Yue Zhanga, Yanan Wanga, Yuru Yeb, Yifan Hongb, Yanchao Lib, Wei Zhoub, Jinhan Zhanga; Zhaode Yinc, Xiaohu Zhanga, b, d, Liu Jun Xiaoa, Bing Liua, Leilei Liua, Tao Chenga, Xia Yaoa, Chongya Jianga, Hengbiao Zhenga, Xiaojun Liua, Qiang Caoa, Yongcao Tiana, Yan Zhua, Weixing Caoa, Yihua Zhue*, Liang Tanga, b, c, d*

^a National Engineering and Technology Center for Information Agriculture, Engineering Research Center for Smart Agriculture, Ministry of Education, Key Laboratory for Crop System Analysis and Decision Making, Ministry of Agriculture and Rural Affairs, Jiangsu Key Laboratory for Information Agriculture, Jiangsu Collaborative Innovation Center for Modern Crop Production, Nanjing Agricultural University, Nanjing 210095, Jiangsu, PR China

^b Shennong Intelligent Agriculture Research Institute Nanjing Co., Ltd

^c Sanya Institute of Nanjing Agricultural University, Sanya, Hainan 572025, China

^d Nanjing Agricultural University (Legeous) AI Agriculture Research Institute

^e College of Information Science and Technology, Nanjing Agricultural University, Nanjing, Jiangsu Province China

These authors contributed equally: Ding Yang, Yancheng Wang and Keqiang Sun

Corresponding authors: Liang Tang (E-mail: tangl@njau.edu.cn) and Yihua Zhu (E-mail: zhyh@njau.edu.cn)

Abstract

Objective: Farm management requires continuous decisions constrained by field conditions, crops, weather, soil, agricultural inputs and business objectives, yet workflows built on predefined software modules still fail to transfer field boundaries, variable semantics, parameter units and intermediate outputs across platforms and models. This study develops SmartFarming, a GIS-based agent framework that translates natural-language farm management requests into constrained, executable and traceable multistep workflows.

Methods: The framework is organized around four components - Persona, Agent Harness, Capabilities and World Interface. The Agent Harness coordinates task orchestration, capability routing, context management, hierarchical memory, execution monitoring and governance. Skill contracts constrain capability composition, while artifacts and their manifests record spatial semantics, versions and data lineage. A GIS module connects field inputs, computation, result visualization and user review within a spatial feedback loop.

Results: Nine case studies instantiate the framework: satellite remote sensing monitoring, UAV remote sensing for rice growth, ground RGB image analysis for wheat-head and rice-panicle counting and disease identification, crop growth prediction, and model-remote-sensing coupling for yield prediction. The cases show that SmartFarming can organize, compose and check existing agronomic models, data services and spatial interactions in end-to-end workflows.

Conclusions: A GIS-based agent framework with explicit spatial, temporal and risk constraints provides a practical basis for reusing heterogeneous agronomic models and data services in farm management, and offers a transferable structure for building domain agent systems in agriculture.

Keywords: farm management; agent harness; skill contract; geographic information system; remote sensing

0 Introduction

Farm management is the continuous organization and regulation of fields, crops, fertilizers, pesticides, labor, machinery, and agricultural operations under given natural resources, production conditions, and operational objectives. Unlike one-off crop identification, disease diagnosis, or yield prediction, farm management spans planting planning, sowing, in-season growth, harvest, and post-season evaluation. Decisions on irrigation, fertilization, crop protection, and machinery operations require weather, soil, crop condition, input costs, and operational windows to be considered for fields with defined boundaries. Image recognition, remote sensing, geographic information systems (GIS), crop models, and Internet of Things (IoT) support field monitoring, precise input application, and production control [1]. These technologies also support a shift toward the integrated design and optimization of crop production systems [2]. The central challenge therefore extends beyond information acquisition at individual production stages to connecting observations, specialized models, and field operations in a continuous process of sensing, analysis, decision making, execution, and feedback.

To support these management processes, farm management information systems have evolved from early tools for farm-operation records and financial management into integrated software systems that combine field profiles, production planning, machinery data, input management, crop monitoring, and business analysis [3]. Commercial platforms emphasize different, often overlapping functions. Climate FieldView (Bayer) and John Deere Operations Center (John Deere) emphasize machinery connectivity, production-data analysis, and operational management, while Cropwise (Syngenta) and xarvio FIELD MANAGER (BASF) also support crop monitoring and agronomic decision-making using field, weather, and remote-sensing data and crop models [4,5,6,7]. These platforms collectively illustrate the breadth of digital support available for field-level management throughout the production cycle.

Although existing platforms integrate multiple business functions, this integration typically depends on predefined software modules, data interfaces, and business processes. Interoperability barriers persist among platforms in technical interfaces, variable semantics, data context, user management, and security and trust mechanisms [8]. Changes in data sources, crops, models, or management needs can require renewed requirements analysis, code changes, interface adaptation, testing, and deployment [9]. A task that combines field delineation, image retrieval, crop growth estimation, phenological assessment, crop simulation, and prescription generation may therefore require users to transfer boundaries, time ranges, parameters, units, and intermediate files manually between modules or platforms. These handoffs increase the workload and the risk of mismatched information. Dynamic coordination of heterogeneous data, models, and tools remains a challenge even when individual functions are already available.

Large language models (LLMs), which power applications such as ChatGPT and Gemini, support natural language understanding, question answering, text and code generation, and multimodal information processing [10,11]. Agricultural domain adaptation uses agricultural corpora, instruction tuning, and image-text data to improve how these models handle agricultural terminology, production questions, and visual features [12,13,14]. These approaches make agricultural knowledge and specialized information more accessible.

Agricultural LLMs can nevertheless generate factual hallucinations, omit key applicability conditions, or provide agronomic advice without verifiable evidence [15].

Tasks involving empirical equations, numerical calculations, and logical constraints, such as water and nutrient management, also expose errors in equation selection, inconsistent calculations, and unwarranted confidence in incorrect results [16]. Converting language-based advice into executable, checkable, and traceable farm management workflows therefore requires mechanisms for task planning, tool use, state maintenance, and execution feedback.

LLM-based agents provide a new approach to extending the practical execution capabilities of language models. An LLM primarily supports language understanding, reasoning, and action generation, whereas an agent additionally incorporates task objectives, external tools, environmental observations, state or memory management, and multistep execution loops. This enables the system to act continuously toward user objectives and adjust subsequent decisions in response to execution results [17,18]. Agent systems thus encompass planning, environmental feedback, state updates, and multirole collaboration in addition to individual tool calls [19,20,21].

Studies in agriculture and spatial information adapt these mechanisms to specialized tasks. Remote sensing and geospatial agents combine natural language task decomposition, specialized tools, and multiple modules for image or map analysis [22,23]. Cai et al. [24] combine image segmentation, LLM reasoning, and hierarchical analysis in an interpretable agent framework for few-shot plant phenotyping. These studies illustrate a transition from question answering and individual model inference to multistep execution with multiple tools. Farm management is not a one-off text-input-to-answer-output task. It requires continuous data queries, status analysis, management-plan generation, and result verification for specific fields. These processes may ultimately affect farming operations. Spatial objects, temporal states, and task semantics must remain consistent throughout this process. Existing studies, however, generally emphasize particular analytical tasks [25,26,27]; they do not by themselves establish a unified approach to field anchoring, state maintenance across growth stages, and artifact handoffs among remote sensing, crop models, and management tools. Mechanisms for map-based result review, governance of farming operations, and controlled skill extension have likewise not been systematically integrated. Real-world farm management therefore still lacks a field-centered system framework that can coordinate multisource agricultural data, remote sensing products, agronomic models, and operational tools while maintaining consistent spatial objects, variable semantics, data sources, and runtime states across tasks, thereby supporting task continuity, process traceability, and governable execution.

Accordingly, this study proposes SmartFarming, an artificial intelligence (AI) agent framework that translates natural language farm management requests into constrained, executable, and traceable multistep workflows. Four components, Persona, Agent Harness, Capabilities, and World Interface, define roles and behavior, maintain task state and coordinate capabilities, encapsulate agricultural algorithms and workflows, and connect GIS, data services, and execution environments, respectively. The framework incorporates agricultural objects, temporal states, and action risks into runtime coordination. The objectives of this study are threefold: (1) to develop a general GIS-based agent architecture that provides a unified structure for task orchestration, capability invocation, state maintenance, execution observability, and system governance; (2) to establish handoff mechanisms for spatial objects, variable semantics, and data lineage across skills, thereby supporting the composition and reuse of remote sensing monitoring, crop modeling, and management-analysis capabilities; and (3) to instantiate the SmartFarming agent and evaluate its end-to-end

closed-loop task capability using cases involving satellite and unmanned aerial vehicle (UAV) remote sensing, ground-image analysis, crop prediction, and model-remote sensing coupling.

1 SmartFarming agricultural agent architecture

Farm management tasks often begin with a textual or GIS-based request and develop into a sequence of spatial definition, data retrieval, crop-state assessment, and management planning. The resulting decisions may further trigger fertilization, irrigation, pesticide application, or machinery operations. Throughout this process, field boundaries and coordinate reference systems (CRSs) constrain where a task applies, while growth stages, weather windows, and management schedules constrain when it can be executed. Because errors in prescription or execution may lead to resource waste, yield loss, or operational risk, the agent must maintain consistent task semantics as information and actions move across these spatial and temporal contexts.

1.1 Architecture and conceptual boundaries under agricultural constraints

Farm management involves natural-language requests, geospatial objects, heterogeneous agricultural data, domain models, and management constraints. SmartFarming organizes these elements through four core components: Persona, Agent Harness, Capabilities, and World Interface (Fig. 1). Agricultural tasks are represented and coordinated with explicit spatial, temporal, and risk constraints, allowing the agent to maintain consistent task objects and conditions as information moves between data sources, analytical capabilities, and agricultural operations.




Fig. 1. SmartFarming architecture under spatial, temporal, and risk constraints. (a): Spatial, temporal, and risk constraints on agricultural tasks. (b): core components of the SmartFarming agricultural agent, including the foundation model, Persona, Agent Harness, Capabilities, and World Interface; and (c): information and action flows between the agent and the agricultural environment through Environment, Observation, and Actions.

Agricultural tasks remain subject to spatial, temporal, and risk constraints

throughout execution. Spatial constraints associate a task with a defined area of interest (AOI), spatial reference, and management unit. Temporal constraints determine whether a capability is applicable to a given crop growth stage, management window, or weather condition. Risk constraints govern actions that may alter agricultural production through authorization, execution gates, and human confirmation. SmartFarming organizes these constraints through four core components-Persona, Agent Harness, Capabilities, and World Interface-which define user roles, coordinate task execution, provide agricultural capabilities, and connect the agent with the agricultural environment, respectively.

Persona defines the intended user, task objectives, interaction strategy, and accessible capabilities. Farm managers may focus on production planning, input allocation, machinery scheduling, and harvesting; agronomists on crop condition assessment, nutrient and water management, pest and disease diagnosis, and agronomic recommendations; and researchers on remote sensing analysis, crop simulation, model evaluation, and scenario analysis. These Personas use the same Agent Harness while differing in their task priorities, accessible capabilities, and interaction strategies.

The Agent Harness coordinates execution across Persona, Capabilities, and World Interface. It coordinates skill selection and execution according to the current task context, while maintaining consistency in spatial objects, temporal conditions, and management objectives across task steps. The Harness also maintains task state and the relationships between upstream and downstream results. It translates task intent interpreted by the foundation model into constrained capability calls, linking language-based reasoning with agricultural computation and operations.

Capabilities comprise the tools, skills, models, and workflows available to agricultural tasks. A tool performs an individual operation, such as image clipping, coordinate transformation, or data retrieval, through an explicit interface. A skill organizes tools, models, workflows, or external services into a reusable agricultural capability, such as remote sensing monitoring, crop growth prediction, fertilization planning, or machinery planning. This organization allows different agricultural algorithms, models, and services to be incorporated without coupling the framework to a particular implementation.

The World Interface connects the agent with GIS, agricultural data services, computing environments, user confirmation interfaces, and external devices. It provides the interfaces through which the agent obtains information, presents results, and initiates permitted actions. Within the World Interface, GIS provides the spatial context for AOI definition, computation, result presentation, and user verification, keeping task inputs and outputs associated with the same field objects.

These components operate within an Environment-Observation-Action loop. The agricultural environment comprises crops, sensors, remote sensing platforms, and machinery. Observations include crop conditions, remote sensing imagery, field records, and yield records. Based on these observations, the agent interprets the task objective and coordinates the appropriate capabilities to produce management actions, such as fertilization, irrigation, harvesting, or machinery operations. Executed actions may alter field conditions and generate new observations, allowing subsequent tasks to proceed from the updated agricultural state.

1.2 Agent Harness runtime mechanisms and functions

The Agent Harness connects the LLM, Capabilities, and World Interface during

task execution. In addition to user objectives, conversation history, and tool results, it maintains agricultural task state, including the AOI, crop growth stage, weather window, and management objective. These state variables allow remote sensing, crop modeling, and management capabilities to operate within a consistent spatial and temporal context. For actions that may alter production states, the Harness also coordinates authorization checks and human confirmation.

SmartFarming builds on the reasoning-and-acting paradigm introduced by ReAct and related work on agent memory and environmental interaction [18,19]. It extends this interaction cycle with agricultural state and constraint management centered on the AOI, growth stage, and management objective (Fig. 2). Perceive captures user requests and external events; Retrieve and Compose assemble the information required for the current task; Decide, Act, and Observe select capabilities, execute calls, and process returned results; and Memorize updates task state and memory. Within this cycle, the Harness links user intent with spatial parsing, data acquisition, model analysis, GIS visualization, and, when required, human confirmation.

The Agent Harness comprises six functional modules: task orchestration, capability routing, context management, hierarchical memory, execution and observation, and governance control. These modules manage the information and control flow required for multistep agricultural tasks.

Task orchestration organizes a user request into executable task steps and maintains their dependencies. It determines the sequence in which skills are invoked, tracks intermediate results, and updates the task state as execution proceeds. For multistep agricultural tasks, orchestration also preserves the relevant AOI, temporal conditions, and management objective across successive operations.

Capability routing maintains an index of callable tools and skills together with their descriptions, metadata, and authorization states. The Harness retrieves candidate capabilities according to user intent and filters them using input requirements, crop applicability, spatial scope, and Persona-defined visibility. Capability discovery, visibility, and authorization are treated separately: a registered and visible capability cannot be executed unless the requesting user or service identity also satisfies its authorization requirements.

Context management controls the information supplied to the LLM at each reasoning step. The context contains structured representations of the Persona configuration, current task state, agricultural environment state, relevant memory, artifact references, and applicable constraints. Large data objects are not inserted directly into the model context; instead, the Harness retains references to the corresponding artifacts and resolves their contents when required.

Hierarchical memory connects the current task with previous management activities and longer-term agricultural knowledge. Working memory stores information required for the current execution, episodic memory records completed management episodes, and long-term memory is designed to retain field history, cropping systems, and accumulated management experience. Together with execution logs and versioned artifacts, these records provide traceable links between previous operations and subsequent decisions.

The execution-and-observation module separates LLM reasoning from operations that modify data or external state. The LLM generates structured capability requests, while the execution environment performs data queries, file processing, model computation, map updates, and other deterministic operations. Short operations may return results synchronously, whereas longer downloads, simulations, or image-inference jobs can execute independently and report progress, warnings, errors,

and completion states. Structured outputs and execution states are returned to the Harness for subsequent reasoning and task-state updates, while results requiring review are presented through the corresponding user interface. The detailed skill execution workflow is illustrated in Fig. S1.

Governance control applies layered checks before and during capability execution. Persona determines capability visibility, authorization determines whether the requesting identity may invoke a capability, and capability interfaces validate required inputs and execution conditions. The World Interface further enforces spatial scope, data-access permissions, and resource constraints during execution. The Harness coordinates these controls but does not replace the validation or decision logic implemented by the responsible component.

SmartFarming assigns different control requirements to information display, advice generation, prescription export, dispatch to management systems, and machinery execution. Actions that may directly change external production states, such as prescription dispatch or device execution, are paused by default until confirmation is provided by an authorized user. Execution records retain the AOI, input data, execution parameters, skill and model versions, quality warnings, and approval status. Iteration limits, failure stop conditions, and version tracking further constrain repeated execution, unauthorized invocation, and outputs without traceable provenance.

1.3 Agricultural data foundation and service integration

The agricultural data foundation organizes heterogeneous environmental, crop, and management information by location, time, crop, and production scenario. Standardized access allows the agent to retrieve relevant data and provenance. Storage, standardization, and query services remain separate from agent reasoning; the Harness requests data according to task objectives and current state. Details of the data sources, standardization procedures, and API-based integration are provided in Text S1.

Environmental data include weather, soil, and terrain information for environmental assessment, crop modeling, and spatial analysis. Crop data include phenology, cultivar parameters, and growth states. Management data include cropping systems, agricultural operations, and administrative boundaries that describe production contexts and management constraints. These data jointly support the representation of the agricultural environment and task state (Table S3).

Data integration requires spatial, temporal, and semantic standardization. Spatial data use a consistent CRS and indexing framework, temporal fields use a common format, and units and category codes follow defined conversions. Vertically layered data use a common layer index. Standardization retains provenance, spatial resolution, temporal validity, processing methods, and versions to support subsequent tracing and reproduction.

Databases and service interfaces expose standardized data to the agent. Spatial and temporal indexes support queries by coordinates, time range, crop type, or task parameters. Standard interfaces connect external data resources and model services while keeping them separate from agent reasoning.

During execution, the Harness queries information for the AOI, time, crop, and task conditions. It retains source, version, spatial extent, and query-time metadata so that agricultural capabilities can share a consistent data context and trace their inputs and outputs.

1.4 GIS spatial feedback and artifact handoffs

Growth monitoring, yield prediction, fertilization prescriptions, and irrigation advice must refer to explicit fields, spatial extents, CRSs, and time windows. GIS supports the management, computation, and presentation of these spatial objects [28]. SmartFarming uses GIS to connect user intent, input data, and management outputs to a common AOI.

The spatial loop connects user input, AOI normalization, spatial computation, result return, and confirmation. Users define the target through GIS interaction or vector uploads, such as Shapefile or GeoJSON, a geographic format based on JavaScript Object Notation (JSON). The system parses the geometry, checks coordinate consistency, and determines task extent to create a standardized spatial object. This object anchors subsequent data retrieval, model computation, and result presentation.

The Harness then organizes data queries, model analysis, and management planning around this AOI. Outputs retain their spatial extent, temporal information, and provenance, allowing remote sensing results, predictions, and management plans to be linked and displayed for the same spatial object.

GIS returns georeferenced outputs and their semantics, including GeoJSON and georeferenced Tagged Image File Format (GeoTIFF) rasters. Users inspect boundaries, anomalous areas, and management results on the map, then confirm, revise, or terminate the task. Their decisions update task state and memory. Changes to spatial conditions trigger AOI normalization again. GIS therefore serves both as a display interface and as a feedback channel that maintains the link from spatial intent through AOI and geospatial artifacts to user confirmation (Fig. 2).




Fig. 2. SmartFarming task execution from user intent to result confirmation.

1.5 Domain skill hierarchy and representative workflows

SmartFarming organizes domain capabilities into three layers, including shared foundations, sensing and modeling, and decisions and operations. The first layer prepares reusable spatial objects and agricultural data, the second derives crop states and predictions from observations, and the third uses these results to support management planning and production operations. Each agricultural skill specifies its applicable crops, spatial scope, temporal conditions, input requirements, and outputs.

The shared-foundation layer provides reusable AOIs, field boundaries, agricultural data query results, standardized remote sensing imagery, and raster or vector layers. It acquires, organizes, and preprocesses data without estimating crop

states or generating agronomic decisions. Quality-controlled AOIs, imagery, and spatial data can be reused across satellite monitoring, UAV analysis, crop simulation, and management planning. For example, remote sensing imagery after quality control and spatial clipping can support phenological-stage identification, crop-state estimation, and yield prediction. Reusing these common data products avoids repeated preprocessing of the same spatial objects and data across different agricultural tasks.

The sensing-and-modeling layer converts observations into crop-state estimates and predictions. Remote sensing skills calculate vegetation indices, identify phenological stages, and estimate leaf area index (LAI), aboveground biomass (AGB), and crop nitrogen status. Ground-image skills identify plants and plant traits, weeds, and disease symptoms. Crop-model skills combine weather, soil, cultivar, management, and remote sensing observations to simulate crop development, biomass accumulation, and yield. The resulting artifacts provide standardized inputs for downstream skills.

The decision-and-operation layer combines spatial objects, agricultural data, and estimated or simulated crop states to formulate management plans and, where appropriate, connect them to production systems. Representative capabilities include fertilization planning, irrigation planning, harvest scheduling, machinery configuration, and route planning. Agronomic parameters are calculated by domain models, rules, or optimization algorithms, while the agent is responsible for task interpretation, capability selection, and execution coordination. For example, a variable-rate fertilization skill can combine field information and crop-state estimates to generate a spatial prescription that is subsequently presented in GIS for review. In this way, SmartFarming retains agricultural models and algorithms as the computational basis of domain operations while organizing them with data and computing resources as reusable skills.

Standardized artifacts provide the interfaces between these layers. The shared-foundation layer prepares spatial objects and data, the sensing-and-modeling layer produces crop-state estimates and predictions, and the decision-and-operation layer uses these results to formulate management plans. The associated task state retains the AOI, temporal context, and provenance as information moves across the layers. The growth-monitoring cases presented in this study mainly involve the first two layers together with GIS-based result review, whereas workflows that generate fertilization, irrigation, or other management plans extend into the decision-and-operation layer.

2 Controlled evolution of SmartFarming capabilities

2.1 Dual-plane design separating capability use from evolution

SmartFarming separates production use from capability development through a runtime plane and a development plane. The runtime plane allows end users to invoke released skills, whereas the development plane supports the creation, testing, and maintenance of candidate skills. Unverified capabilities cannot enter the production environment, as they may compromise task reproducibility and traceability or introduce risks to agricultural operations.

The runtime plane serves farm managers, agronomists, and agricultural researchers. Authorized users provide task parameters, inspect results, and confirm actions when required, but cannot directly modify released capability definitions,

execution logic, or interface constraints.

The development plane serves skill developers, system maintainers, and controlled development assistants. New algorithms, models, data services, and workflows are packaged as candidate skills according to SmartFarming specifications and then submitted for validation, registration, and release. Candidate skills remain separate from production task execution until they satisfy the required release conditions.

A controlled release process connects the two planes. Candidate skills become available in the runtime plane only after interface validation, compatibility checks, runtime testing, and risk review. This separation allows capabilities to evolve without permitting an active agent or end user to modify released skills during production execution.

Capability evolution occurs by adding, composing, or upgrading standardized skills rather than modifying the agent core. A skill may encapsulate algorithms, model services, external tools, or workflows and is associated with its structural specification, interface constraints, version, and release status.

2.2 Skill-Contract-Artifact object model and composition

SmartFarming uses a hierarchical skill system to organize agricultural tasks and coordinates multiple skills to complete typical workflows. For these capabilities to be retrieved, validated, and composed by the agent, a skill cannot exist only as a natural language description, prompt template, or loosely associated toolset; it requires a consistent structural representation. SmartFarming therefore represents each skill as an agricultural capability object and uses a Skill-Contract-Artifact model to make composition explicit, machine-readable, verifiable, and traceable.

A skill is the basic unit of agricultural capability. Inspired by modular software design and the principle of explicit interface contracts [29], SmartFarming represents each runtime skill as a five-tuple:

$$S = \langle \text{Spec}, \text{Impl}, \text{Contract}, \text{Metadata}, \text{Policy} \rangle \quad (1)$$

Spec defines the objective, steps, and expected behavior; Impl denotes the algorithm, model, tool, or external service; contract defines inputs, outputs, semantic constraints, and execution conditions; Metadata records applicable crops, spatial scale, time range, tags, and version; and Policy declares risk level, invocation prerequisites, and governance requirements. The Harness uses these fields for retrieval, applicability checks, and execution coordination. Runtime authorization makes the final admission decision from the current identity, role, and task state.

SmartFarming transfers data and task state through standardized artifacts instead of relying on direct skill-to-skill calls. An artifact exchanges information and maintains task continuity between upstream and downstream skills, carrying data, files, model states, analytical results, or management outputs produced, consumed, or transferred during execution. The basic relationship is:

$$\text{Skill}_i \rightarrow \text{Artifact}_j \rightarrow \text{Skill}_k \quad (2)$$

The upstream skill produces the artifact, and the downstream skill consumes it. The Harness schedules execution from skill dependencies and uses artifacts to preserve state, spatial semantics, and results along the workflow.

A contract specifies the artifact types a skill accepts and produces, together with execution constraints. The Harness connects skills only when their input and output types and semantics are compatible. Otherwise, it pauses the task and requests new data or missing information.

Each artifact has an associated manifest that records provenance, version, spatial extent, generation time, execution context, and dependencies. In this design, the manifest is a trace record attached to an artifact; it does not itself determine composition or scheduling. Together, the artifact and manifest support recovery, accountability, and reproduction of execution paths.

The Skill-Contract-Artifact model separates capability definition, interface constraints, and runtime data handoff. The Harness uses these structures to identify applicable skills, verify input–output compatibility, coordinate their execution, and maintain data lineage across multistep agricultural workflows.

2.3 Skill structure and domain-specific type contracts

Before admission, every skill declares a consistent structural description, interface constraints, runtime states, and governance requirements. These declarations allow capabilities from different sources to be parsed, validated, invoked, and composed consistently. The declaration specifies the agricultural problem, applicable crops, spatial scale, time range, input conditions, and expected outputs. A rice-monitoring skill, for example, states supported growth stages, remote sensing data, resolution, and output indicators. A yield-prediction skill states its weather, soil, cultivar, and management requirements. Metadata therefore prevents invocation when crop, time, or data conditions are unsuitable.

Contracts specify interface structures, semantics, and applicability. The OpenAPI Specification can describe parameter structures, data types, and service calls, while International Organization for Standardization/International Electrotechnical Commission (ISO/IEC) 25010 provides a reference for quality characteristics such as reliability, compatibility, and maintainability [30,31]. SmartFarming adds agricultural spatial, temporal, variable, and data-quality constraints. Spatial constraints cover CRS, AOI, resolution, and scale. Temporal constraints cover observation time, simulation period, growth stage, and validity. Variable constraints cover names, units, ranges, and agronomic meaning. Data-quality constraints cover completeness, validity flags, uncertainty, and warnings.

Skills also follow a common runtime protocol for long or externally dependent tasks. In addition to final outputs, they report waiting, running, completed, failed, paused, or cancelled states with progress, warnings, and errors. Start conditions, exception handling, and termination conditions are explicit. Image retrieval reports when no image satisfies the requested time and quality criteria, and crop simulation pauses rather than inferring critical missing parameters. Skill outputs conform to their output contracts so that downstream skills can determine whether invocation conditions are met. Artifact objects manage the associated trace information.

Policy records visibility, resource needs, risk level, and confirmation requirements. Runtime authorization combines these declarations with the current identity and task state. Monitoring and information display can run automatically when authorized. Consequential skills that dispatch fertilizer prescriptions, control irrigation, or operate machinery require parameter bounds, confirmation by a responsible person, and execution records. These structural and type requirements support consistent registration, invocation, composition, and execution, and provide a basis for versioning, lifecycle control, and managed evolution.

2.4 Capability admission and lifecycle management

SmartFarming manages capability admission through Discovery, Engineering, Governance, and Operation (Fig. 3). Machine learning systems require continued management of models and dependencies across development, validation, deployment, and maintenance [32,33,34]. The four-stage lifecycle applies this principle to changing agricultural algorithms, models, data sources, and knowledge without altering active production tasks.




Fig. 3. Controlled evolution of SmartFarming capabilities.

Discovery identifies unmet capability needs from runtime logs, failures, performance changes, user feedback, new cultivars, new remote sensing sources, new crop models, or new management needs. The system combines runtime signals with identified capability gaps to specify the target capability, applicability scope, and constraints, thereby producing a requirement definition.

Engineering designs and implements candidate skills. Developers use the capability requirements to design algorithms, data requirements, interfaces, and dependencies and encapsulate new agricultural algorithms, models, data services, or workflows as skills conforming to SmartFarming specifications. Supporting a new remote sensing satellite source, crop model, or pest and disease recognition algorithm, for example, requires input-output interface adaptation, definition of agricultural semantics, and compatibility design with the existing skill system so that the capability can reliably join existing task chains.

Governance validates, registers, and releases candidate skills. Checks include functional validation, contract consistency, regression tests, and end-to-end workflow tests. Registration records metadata, versions, dependencies, and runtime configuration. Authorization, risk assessment, and governance review precede release. Consequential skills for fertilizer prescriptions, irrigation control, or machinery operations require agronomic rule checks, expert review, and human confirmation beyond the checks applied to lower-risk monitoring capabilities.

Operation invokes released skills and monitors states, outputs, artifact quality, resources, and user feedback. Performance degradation, changing data sources, new cultivars, or new user needs return to Discovery. Consistent with lifecycle risk management [35], changes to algorithms, interfaces, or dependencies must pass through Engineering and Governance again. An active agent cannot directly modify or deploy a released skill.

This lifecycle provides a common route from need discovery through engineering, validation, release, operation, and feedback. Its controls are intended to preserve continuity, reproducibility, traceability, and safety as capabilities change; these properties require empirical evaluation in sustained deployments.

2.5 Capability extension mechanisms

Capabilities can originate from manual development, external resources, foundation-model assistance, or runtime feedback. Regardless of origin, each candidate follows the same packaging, validation, registration, and release process before entering production.

The Model Context Protocol (MCP) can connect external databases, remote sensing platforms, crop models, computing tools, and knowledge services [36]. MCP provides connectivity and message exchange, but does not define agricultural semantics, spatial constraints, data quality, authorization, or risk levels. Every external capability integrated into SmartFarming must therefore still be encapsulated through the skill object model, satisfy Contract, Policy, and type-contract requirements, and undergo validation, registration, and release before becoming an agricultural capability in the production environment.

Foundation models can assist with code generation, tool packaging, and workflow construction [37]. In SmartFarming, they support development by proposing candidate implementations, skill templates, interface descriptions, and upgrade suggestions. They cannot modify active skills directly. AI-assisted candidates still require lifecycle validation, registration, human review, governance checks, and release.

Runtime feedback can also initiate extension. Repeated failures, declining artifact quality, data-source changes, or new agricultural needs become upgrade proposals that return to Discovery. An upgraded skill passes through Engineering and Governance rather than directly replacing a production version.

These controls allow external connections, model-assisted development, and feedback-driven upgrades to follow a common release process while keeping the production environment stable and traceable.

3 SmartFarming applications in farm management

The application chain links crop-state sensing, management decisions, growth prediction, and yield updating. Satellite and UAV remote sensing characterize field variation; existing nutrient and water models are wrapped as callable skills; RiceGrow and WheatGrow support full-season simulation and scenario analysis; and model-remote sensing coupling uses an ensemble Kalman filter (EnKF) to update simulated crop states from remote sensing observations. These components illustrate composable, traceable, and reviewable workflows for precision management and yield prediction. The models and algorithms underlying the case-study workflows are summarized in Table S1.

3.1 Remote sensing monitoring skills

Remote sensing forms the sensing entry point. Satellite images provide broad spatial coverage and time series, while UAV images provide finer spatial resolution and flexible acquisition. Two types of data are encapsulated as skills and output growth indicators, field statistics, and visualization products. Their main functions and input-output specifications are summarized in Table S2. Their outputs can enter fertilization, irrigation, or yield workflows only when the AOI, CRS, date, algorithm version, and units are retained.

For satellite monitoring, users draw an AOI on GIS or upload a vector boundary. The workflow uses Microsoft Planetary Computer (MPC) as a cloud geospatial catalog and accesses its SpatioTemporal Asset Catalog (STAC) interface to search by AOI, observation window, and cloud threshold. Available products include Sentinel-2 Level-2A, Landsat Collection 2 Level-2, and Moderate Resolution Imaging Spectroradiometer (MODIS) data. The workflow extracts surface reflectance and vegetation indices and exposes modules for field-boundary detection, LAI estimation, topdressing prescriptions, and irrigation-demand estimation. When revisit frequency and cloud-free coverage permit, the system can update regional observations and produce spatial management maps.

The UAV skill accepts uploaded multispectral images, reads geolocation and red, red-edge, near-infrared, and other bands, and calculates vegetation indices such as the normalized difference vegetation index (NDVI). Registered algorithms estimate AGB, LAI, leaf nitrogen accumulation (LNA), leaf nitrogen concentration (LNC), grain protein content (GPC), and yield, with continuous and classified maps as outputs (Table S2). The monitored crop-status results can likewise be used to generate topdressing prescription maps for the target field and guide optimized production management.

3.2 Ground red-green-blue image recognition skill

The ground red-green-blue (RGB) image skill complements remote sensing with object-level information from phones, cameras, or low-altitude UAVs. It detects wheat spikes, rice panicles, weeds, and disease regions. After a user uploads an image and describes the task, the agent checks the image channels and task type, then matches a registered skill from its description, input schema, and supported targets. This routing separates RGB tasks from remote sensing tasks that require red-edge or near-infrared bands.

The agent packages the image path, task category, and parameters as a structured request and loads registered You Only Look Once (YOLO) model weights in a sandbox. YOLO predicts classes, bounding boxes, and confidence scores in one forward pass [38]. The skill counts objects and returns annotated images, label files, and JSON or comma-separated values (CSV) outputs. The artifact records classes, counts, confidence scores, and file locations for user review. Occlusion, illumination, motion blur, and scale can affect detection; use is therefore limited to crops, target classes, and imaging conditions covered by the registered model [39,40].

3.3 Water and nutrient management decision skills

Water and nutrient management connects sensed crop states to field decisions. This module does not design new fertilization or irrigation algorithms. It wraps existing agronomic models as domain skills with explicit input and output contracts, and the agent selects and composes them from user objectives, data state, and constraints.

Three skills address basal fertilizer recommendations, topdressing, and irrigation demand. Basal fertilizer recommendations use crop type, target yield, soil properties, and management history. Remote sensing topdressing combines observations with growth stage and agronomic rules. Irrigation demand uses weather, soil moisture, and crop state. The skills share AOIs, agricultural data, and model services through contracts and artifacts.

The agent interprets requests, builds the workflow, checks invocation conditions, and organizes feedback; validated agricultural models calculate fertilizer or irrigation quantities. Outputs record inputs, model versions, applicability, and quality state. For a request to prepare a rice water and nutrient plan, the agent identifies the field and crop state, retrieves remote sensing, soil, and weather information, and selects the relevant skill. Missing critical data, failed applicability checks, or conflicting outputs cause the workflow to request information or human review rather than generate a prescription.

3.4 Process-based crop growth prediction skills

Crop growth prediction converts a natural language scenario into parameters, model execution, and results. The computational core is the CropGrow developed by the National Engineering and Technology Center for Information Agriculture (NETCIA) at Nanjing Agricultural University [41]. It includes RiceGrow and WheatGrow [42,43], which have been evaluated and adapted through Agricultural Model Intercomparison and Improvement Project (AgMIP) studies in multiple regions [44,45,46,47].

4.4.1 Rice and wheat prediction with crop growth models

SmartFarming exposes RiceGrow and WheatGrow as standardized skills with crop-specific deterministic kernels. Both represent genotype $\times$ environment $\times$ management (G $\times$ E $\times$ M) interactions to simulate development, growth, and yield formation. Before simulation, the agent converts the request into crop-specific structured parameters, retrieves weather and soil profiles for the selected location, and seeks an exact cultivar match in the parameter database. If no exact match exists, it returns similar candidates for user confirmation. The system standardizes dates, variable names, and units, and checks sowing and transplanting order, weather coverage, cultivar parameter completeness, and irrigation and fertilizer events. Missing or ambiguous critical parameters prompt a question or a predefined fallback; the agent does not invent values. The scheduler then invokes RiceGrow or WheatGrow from sowing to maturity.

At each time step, the model updates phenology and growth to produce a traceable trajectory. Standard outputs include predicted dates of key stages, LAI, AGB, and final grain yield. Users can compare sowing or transplanting dates, irrigation regimes, and fertilization plans under common weather, soil, and cultivar conditions. These scenarios support model-based assessment of growth, biomass, and yield responses.

4.4.2 Outputs and presentation

The skill produces daily simulation data, a natural language summary, and interactive visualizations. Visualizations show AGB and grain yield, biomass partitioning among roots, leaves, stems, and panicles or spikes, and key stages from sowing to maturity. The daily table records LAI, organ biomass, AGB, and yield. The summary reports predicted yield, development, management responses, weather source, stress settings, assumptions, and applicability limits.

4.4.3 Technical workflow

The workflow has four stages: scenario parsing, parameter matching and data assembly, model execution, and result validation and presentation. Scenario parsing extracts location, cultivar, dates, and management events from natural language. A conservative deterministic policy requests missing or ambiguous fields and accepts

multiple Chinese date and fertilizer expressions. Users can alternatively select a location on the map to define the spatial query.

Parameter matching and data assembly retrieves station or gridded weather for the coordinates and period, historical observations or forecasts as appropriate, soil profiles through spatial indexes, and genetic parameters from the cultivar database. Exact cultivar matches have priority; otherwise, candidates are marked with their match status and uncertainty. Weather, soil, cultivar, and management data are standardized by variable name, unit, and time scale in a traceable package that records values, sources, and fallback status.

Before execution, the system checks schema, units, chronology, and agronomic constraints. Validated parameters are serialized as JSON and sent to a model service through a representational state transfer (REST) application programming interface (API). Depending on complexity, duration, scenarios, and load, a full-season simulation may take seconds to minutes. Range, completeness, and internal-consistency checks precede conversion to summaries, time-series plots, and scenario comparisons. The interface displays inputs, outputs, sources, assumptions, and warnings. The result interpreter describes validated outputs without adding unsupported agronomic conclusions.

4.4.4 Deployment in the agent system

The models run as separate microservices connected through MCP [36]. RiceGrow and WheatGrow maintain independent environments containing dependencies, cultivar parameters, and weather and soil data, together with layered weather fallback strategies.

If local weather data are absent or fail checks for period, variables, or completeness, the system applies a time-scale-dependent fallback. Spatial retrieval uses map-selected longitude and latitude in the World Geodetic System 1984 (WGS 84). A field polygon can be represented by its centroid for point queries or by area-weighted values from intersecting grids. Historical simulations first rank online stations within a predefined radius by geodesic distance, elevation difference, and completeness. If no station meets the quality criteria, the system queries the NASA Prediction Of Worldwide Energy Resources (POWER) Daily Point API or ERA5-Land reanalysis. Point queries use bilinear interpolation from four surrounding cells; field queries use intersection-area weighting. This architecture integrates process models with natural language task coordination.

3.5 Model-remote sensing coupling for yield prediction

Crop models provide continuous daily predictions from weather and cultivar parameters, but regional parameter applicability can limit them. Remote sensing, including Sentinel-2, observes field-scale crop states intermittently because of revisit intervals and cloud. Model-remote sensing coupling uses data assimilation to update a dynamic crop model with these observations. This can bring simulated states closer to assimilated observations, although improved independent yield accuracy must be demonstrated with held-out measurements.

The SmartFarming workflow contains six steps (Fig. 4): AOI and agronomic parameter collection; daily RiceGrow simulation; Sentinel-2 STAC retrieval within model-derived phenological windows; LAI estimation from Sentinel-2 red-edge data; daily ensemble Kalman filter (EnKF) assimilation; and yield-raster export and visualization. EnKF links steps two and five by returning updated ensemble states to

the model for the next forecast and analysis cycle.




Fig. 4. Model-remote sensing coupling workflow

4 Results and analysis

4.1 Satellite-based growth monitoring and management

To evaluate the end-to-end execution chain of the satellite remote sensing skill-from field delineation and image retrieval to crop-status diagnosis and prescription generation. We selected a rice field in Chongming District, Shanghai, China, as a demonstration case. Based on the field boundary drawn by the user, SmartFarming standardized the boundary vector and calculated a field area of 6.45 ha. Sentinel-2 imagery was then retrieved for the period from 1 to 31 August 2026 under a cloud-cover constraint of $\leq 20\%$. The system returned three candidate scenes, from which a Sentinel-2B image acquired on 2 August 2026, with 0.7% cloud cover, was selected for analysis. This acquisition occurred 79 days after sowing on 15 May 2026, corresponding to the late-tillering to stem-elongation stage of rice. Within the same workflow, the system generated NDVI, NDWI, NDRE, LAI, and variable-rate topdressing prescription products (Fig. 5).




Fig. 5. Demonstration of SmartFarming satellite remote sensing monitoring skills. (a) NDVI; (b) NDWI; (c) NDRE; (d) recommended amount of topdressing; (e) the corresponding prompt and response.

The vegetation indices revealed spatial heterogeneity in crop-canopy status within the study area. Within the ranges shown in Fig. 5, NDVI varied from 0.2 to 0.6, NDRE from 0.3 to 0.5, and NDWI from -0.6 to -0.3. Based on the predefined crop-status classification and topdressing-mapping rules, the system subsequently converted pixel-level crop-status classes into a variable-rate topdressing prescription. The prescribed rate at the example point was 232.5 kg ha⁻¹.

The skill also provided recommendations for water management and pest and disease monitoring. Alternate wetting and drying irrigation was recommended: the field should be maintained under moist conditions, allowed to drain naturally, and then re-irrigated as appropriate to sustain root activity and reduce lodging risk. Because the rice crop exhibited vigorous growth and relatively high canopy closure, particular attention should be paid to the risks posed by rice planthoppers, sheath blight, and rice blast. More intensive field scouting was therefore recommended from late August to early September.

4.2 UAV remote sensing task

To evaluate the applicability of the SmartFarming unmanned aerial vehicle (UAV) remote sensing skill, a case study in a rice cultivar trial area in Chongming, Shanghai was conducted. The rice was sown on May 15, 2024, and the UAV images are acquired on July 25, 2024, corresponding to 71 days after sowing and the booting stage. A DJI Phantom 4 Multispectral captures images in five bands: blue, green, red, red edge, and near-infrared. The five-band orthomosaic and a Shapefile containing the experimental plot boundaries serve as inputs. The case evaluates the capabilities of system for automated image mosaicking and recognition, spatial registration, multi-trait estimation, plot-level statistical analysis, and visualization of UAV remote sensing data.

During execution, SmartFarming first checks the consistency of the image bands, coordinate reference system, spatial extent, and invalid values. It then aligns the images with the plot boundaries in the World Geodetic System 1984 / Universal Transverse Mercator zone 51N coordinate reference system (WGS 84 / UTM zone 51N; EPSG:32651). The system calculates the normalized difference vegetation index (NDVI) from multispectral reflectance and invokes the integrated estimation algorithms to generate aboveground biomass (AGB), leaf area index (LAI), leaf age (LA), leaf nitrogen accumulation (LNA), leaf nitrogen concentration (LNC), grain protein content (GPC), and yield (YIELD), together with a UAV-based topdressing prescription (Fig. 6). The georeferenced Tagged Image File Format (GeoTIFF) rasters for these indicators provide the basis for quantitative analysis. The system subsequently uses each experimental plot as a zonal statistical unit, calculates the mean values of the eight indicator rasters, writes the results to the Shapefile attribute table, and standardizes the field names and measurement units.

Spatially, LAI, LA, LNA, GPC, AGB, and estimated yield are generally within the medium-to-high ranges across the main planted strips. Lower values occur primarily along field margins, near roads crossing the fields, and in localized areas in the south. Several estimated indicators exhibit similar spatial patterns, showing that the system consistently maps common variation in canopy reflectance to products representing crop growth, nitrogen status, grain quality, and yield. Higher prescribed topdressing rates occur mainly in the southern area and in a small number of isolated patches, whereas low-to-moderate rates predominate elsewhere. The resulting map illustrates how the workflow converts a uniform fertilization strategy into a spatially variable prescription for localized management.




Fig. 6. SmartFarming unmanned aerial vehicle remote sensing case for rice growth monitoring. (a) LAI; (b) LA; (c) LNA; (d) LNC; (e) GPC; (f) AGB; (g) yield; (h) amount of topdressing; and (i) the corresponding prompt and response.

4.3 Image recognition tasks

5.3.1 Test modules and tasks

The Crop Image Analysis Assistant integrates ground RGB detection. Users upload phone, camera, or low-altitude UAV images and describe spike or panicle counting, weed detection, or disease-region recognition. The agent selects the skill from image type and task semantics, without requiring users to specify model weights or scripts.

5.3.2 Task interpretation and skill invocation

The workflow proceeds through request input, task parsing, skill matching, sandbox execution, and result organization. It checks channels and task type, generates model parameters from the registered skill schema, loads YOLO weights in a sandbox, and returns annotated images, labels, and JSON or CSV outputs. The response includes detected classes, counts, confidence scores, and file locations.

Test images cover wheat-spike and rice-panicle counting and the detection of disease regions, pests, and weeds. The system parses the requests, invokes the matched models, and returns annotated images and structured outputs (Fig. 7).




Fig. 7. Wheat-head and rice-panicle counts and disease identification. (a) Rice image recognition; (b) wheat image recognition; (c) neck blast; (d) bacterial leaf blight; (e) rice blast; (f) rice narrow brown leaf spot; and (g) corresponding prompt and response.

4.4 Growth and yield prediction task

We define a rice production scenario to exercise full-season simulation and result presentation. The location is Nanjing, Jiangsu, and the cultivar is Nangeng 46. Sowing and transplanting dates are 16 May and 10 June 2025. The scenario specifies conventional water management, straw incorporation, 300 kg ha⁻¹ of 15-15-15 compound fertilizer as basal fertilizer, and 150 kg ha⁻¹ of urea at both tillering and panicle initiation. The model predicts a 162-day season. AGB increases rapidly after approximately 60 days and approaches a plateau after 120 days, while grain biomass accumulates during filling(Fig. 8). The organ trajectories reflect modeled transfer from vegetative organs to grain.




Fig. 8. Rice growth prediction output. (a) Structured scenario information summarizing location, cultivar, season, water management, fertilization management, and predicted yield; (b) dynamics of AGB and grain yield; (c) biomass allocation among root, leaf, stem, and panicle; (d) phenological and management timeline; and (e) the corresponding prompt and response.

10.12343/Agrixiv.202610.00002V2

4.5 Application of model-remote sensing coupling

A synthetic single-season rice simulation experiment for the Jianghuai region was constructed for the testing of model-remote sensing coupling for yield prediction from 1 May to 17 September. Simulated drivers range from 20 to 35° C and from 12 to 25 MJ m⁻² d⁻¹ for daily radiation. Scenario A runs RiceGrow with one parameter set and no remote sensing update. Scenario B uses an EnKF with 80 ensemble members and assimilates Sentinel-2-derived LAI at tillering, jointing, booting, heading, and grain filling with an observation-error standard deviation of 0.5 (Fig. 9). Ensemble spread increases around booting and heading because maximum LAI and growth rate are perturbed. The posterior LAI moves closer to the assimilated LAI observations, confirming that the analysis step transfers those observations into the modeled state.




Fig. 9. Model-remote sensing coupling results. (a) Prior (model-only simulation), posterior (EnKF assimilation), and reference LAI trajectories; (b) LAI trajectories of ensemble members; (c) dynamic yield formation over the full growing season; (d) model-remote sensing coupling results; and (e) the corresponding prompt and response.

The case instantiates a reusable workflow from AOI collection and Sentinel-2 STAC retrieval through LAI estimation and daily EnKF assimilation. It demonstrates model-state updating and workflow composition. Because the assimilated observations are also the target to which posterior LAI is compared, this agreement is not an independent accuracy test. Independent LAI and yield measurements are required before claiming more reliable yield estimates.

5 Discussion

5.1 Characteristics of the SmartFarming Agent Architecture

Existing LLM-based agents have developed relatively mature general-purpose operating mechanisms [48]. SmartFarming uses the ReAct approach to cycle through perception, reasoning, action, and environmental observation, allowing the LLM to invoke external tools and continue task execution based on execution feedback. Subsequent studies have further integrated task planning, memory management, capability invocation, and environmental interaction into agent architectures, enabling them to support more complex multistep task execution [18,19,49]. On this basis, SmartFarming is further tailored to the spatial, temporal, and action constraints of farm management. It incorporates agricultural task states, including field AOIs,

growth stages, observation times, management windows, and action conditions, into the continuous maintenance process of the Agent Harness. Compared with general-purpose LLM-agent frameworks, which typically operate through iterative reasoning, planning, and tool use and adapt their actions using environmental or tool-generated feedback [50], SmartFarming treats farm-management states as explicit and persistent operational context. Recent domain-specific agent systems have likewise demonstrated the value of integrating specialized tools with memory or shared workflow context for multistep scientific tasks [51,52]. This design helps different farm-management capabilities remain continuously connected around clearly defined spatial entities, temporal conditions, and management objectives. It also reduces the risk of drift in entities, states, and task conditions during multistep tasks.

SmartFarming and existing farm management platforms differ in their functional emphases. Conventional farm management software can already support field records, production records, crop monitoring, and farming operations. However, these functions typically depend on predefined modules, fixed interfaces, and predefined business processes. When a task simultaneously involves remote sensing data, crop models, and management tools, data conversion and manual coordination may still be required [3,9]. SmartFarming places greater emphasis on dynamically organizing heterogeneous data, models, and tools according to specific management objectives. Through contracts and artifacts, it explicitly transfers information on field boundaries, coordinate reference systems, variable semantics, measurement units, temporal conditions, and data provenance. This enables different capabilities to form continuous task chains. Compared with conventional software that relies on predefined business processes, this approach provides greater flexibility for composing cross-module tasks and helps reduce spatial, semantic, and parameter mismatches that may arise during manual handoffs.

With respect to the division of responsibilities between specialized computation and language-model reasoning, SmartFarming adopts an operating mode that separates agent orchestration from specialized-module computation. Existing research on domain-specific agents indicates that combining LLMs with specialized tools can leverage the task-understanding and workflow-organization capabilities of language models while retaining the computational capabilities of specialized tools for deterministic tasks [53]. This division of responsibilities is particularly important for agricultural management. Although agricultural large language models (AgriLLMs) have advantages in domain knowledge understanding and natural-language interaction, they may still generate factual hallucinations, omit applicability conditions, and produce unstable results when handling formulas, parameters, and numerical calculations [15,16]. Therefore, SmartFarming uses language models primarily for task understanding, skill selection, parameter organization, and workflow coordination. Tasks with explicit formulas, parameters, and applicability conditions, including fertilization, irrigation, remote-sensing inversion, and crop model simulation, are executed by validated algorithms, models, and rule-based modules. Compared with directly generating agronomic judgments or deterministic numerical results using language models, this separation of responsibilities helps reduce the effects of formula misuse, unsupported parameter completion, and model hallucinations on results. It also preserves the established computational logic and applicability boundaries of specialized models.

Building on these mechanisms, SmartFarming further uses contracts to constrain the inputs, outputs, and execution conditions of different skills. It also records the data and parameters used during task execution, skill and model versions, intermediate

products, exception information, and the human-confirmation outcomes. For tasks that may alter actual agricultural production states, such as issuance of fertilization prescriptions, irrigation control, and agricultural machinery operation, this study treats permission checks, execution gating, and human confirmation as control requirements for high-risk actions. Compared with general-purpose agents that primarily focus on task planning, tool invocation, and task completion [54], SmartFarming places greater emphasis on the consistency of spatial entities, variable semantics, and operational states across the multistep execution of agricultural tasks. It also emphasizes the inspectability, reproducibility, and traceability of the process through which task results are produced. Accordingly, the principal architectural additions in SmartFarming are the continuous maintenance of agricultural spatiotemporal states, the dynamic organization of heterogeneous agricultural capabilities, the controlled execution of deterministic computation, and the verifiability of multistep tasks.

5.2 Agricultural skill Orchestration and Controlled Capability Extension

SmartFarming builds on a general-purpose agent harness and external capability-invocation mechanisms. It can access external tools, data services, and model capabilities through unified interfaces or the Model Context Protocol (MCP). Similar approaches to capability organization have been applied in multiple domain-specific agents [52,53,55]. These studies indicate that connecting specialist tools and external capabilities through an agent harness or standardized interfaces has become an important technical approach for domain-specific agents to extend task capabilities. SmartFarming further encapsulates agricultural capabilities, including spatial processing, remote-sensing monitoring, image analysis, fertilization and irrigation calculations, and crop models, as independent skills. Unlike general tools and APIs, which primarily specify functions and parameters, upstream and downstream capabilities in agricultural tasks must also maintain consistency in field extent, observation time, variable semantics, measurement units, and data provenance [56,57]. To this end, SmartFarming uses contracts to specify the inputs, outputs, and execution conditions of skills. It uses artifacts and their manifests to retain actual outputs together with their spatial, temporal, provenance, and version metadata. When an upstream result does not meet the requirements of a downstream skill, the system can halt the task or request supplementary information rather than allowing the language model to independently infer missing parameters [15,16]. This design helps reduce handoff errors caused by inconsistent spatial extent, variable definitions, or parameter units in complex agricultural tasks.

Skill encapsulation enables agricultural algorithms and models to be reused and recombined without modifying the Agent Harness; recent work likewise emphasizes standardized, permissioned interfaces for integrating external tools into agent systems [58]. Because agricultural systems continually incorporate new data sources, models, and regional algorithms, SmartFarming separates operational use from skill development. Users invoke validated and released skills, whereas new algorithms, models, and data services must undergo contract checks, testing, review, and versioned release before entering the runtime environment. Constrained tool use and auditable intermediate steps are important safeguards for maintaining traceability in multistep agent workflows [59]. Human oversight also remains necessary when agent outputs may inform complex management decisions [60]. Its long-term maintenance efficiency, compatibility, and stability still require quantification.

More broadly, SmartFarming separates a stable Agent Harness from replaceable agricultural skills and their external interfaces. This modular separation is consistent with the perception-memory-action decomposition commonly used to describe LLM-based agents [54]. Recent systems show that specialist tools can be combined with memory to support multistep workflows [51], while dedicated execution roles can coordinate tool-based actions [61]. In SmartFarming, the World Interface connects skills to GIS, agricultural data services, computing environments, and external devices. The same design could be adapted to other specialized domains by redefining Skill, Contract, Artifact, and Policy semantics for their own entities, temporal requirements, data quality, and risk constraints.

5.3 Validation of the SmartFarming Framework Through Multisource Agricultural Cases

Existing studies have begun to explore applications of LLMs and agents to specialized agricultural tasks, mainly targeting individual tasks such as knowledge synthesis [62], irrigation decision-making [26], and plant phenotyping [24,52]. These agricultural agents can already combine domain knowledge, sensor data, and specialized algorithms to support different types of agricultural tasks, including knowledge synthesis, management decision-making, and phenotypic analysis. However, SmartFarming organizes such capabilities through unified skill invocation, input-condition checks, and structured task artifacts. For tasks involving spatial entities, the AOI further serves as a spatial anchor for data retrieval, model computation, and result presentation. Thus, rather than focusing on a single remote-sensing, vision, or crop model task, as many existing agents do, SmartFarming focuses on organizing agricultural capabilities from different sources and with different computational approaches into continuously connected task workflows. When data sources, analytical models, or task requirements change, organizing capabilities around skills and their input-output contracts also provides a unified interface basis for capability replacement and recombination.

In spatially related cases involving satellite remote sensing, unmanned aerial vehicle remote sensing, weather-soil data matching, and crop models, GIS and AOI not only support result presentation but also participate in defining spatial inputs and constraining task scope. After a user specifies a target field, the AOI can serve as a common spatial anchor for constraining image retrieval, spatial-data processing, environmental-data matching, and model execution. This enables different data and computational results to correspond to the same management object. It should be noted that GIS in agricultural applications primarily supports spatial-data storage, query, overlay analysis, and visualization. Even when satellite, UAV, weather, and soil data are clipped to the same AOI, they cannot necessarily enter the same analytical or modeling workflow directly. Different data sources may still differ in field-boundary versions, coordinate reference systems, spatial resolution, observation times and temporal windows, variable definitions, measurement units, and processing-history records. Conventional GIS workflows do not usually treat these conditions as explicit constraints that are continuously transferred and checked across tool and model invocations. Consequently, when data move from spatial processing to remote-sensing inversion, environmental matching, or crop model simulation, problems may still arise, including inconsistent input semantics, poor result traceability, or unclear model applicability [56,63,64]. Against this background, SmartFarming does not replace the spatial-analysis capabilities of GIS. Instead, it extends GIS and AOI from tools for

result presentation or data filtering into spatial constraints that can be carried with a task. Contracts specify the inputs, outputs, and execution conditions of skills, whereas artifacts and manifests retain data provenance, spatial extent, temporal conditions, and processing versions. Satellite, UAV, and ground-based images provide crop information at different scales and in different modalities. SmartFarming primarily organizes these results for subsequent analysis, simulation, or management steps. Numerical calculations for fertilization and irrigation, as well as crop model simulations, remain the responsibility of existing agronomic algorithms or process-based crop models, while the agent primarily performs input organization, condition checking, capability invocation, and upstream-downstream task handoff.

5.4 Limitations and Future Directions

Although this study has designed and evaluated SmartFarming through its agent architecture, agricultural skill organization, and multiple agricultural task cases, the current work still has several limitations. First, the existing cases primarily validate the operational feasibility of the framework and task chains. The fields, crops, and management scenarios involved remain limited, and long-term stability and generalizability have not yet been evaluated through multi-region, multi-crop, and multi-season experiments. Second, this study focuses on how SmartFarming organizes existing agricultural data, algorithms, and models. It has not yet conducted systematic performance comparisons with existing farm-management platforms or other agricultural agents under a common dataset, tool pool, and set of evaluation metrics [65]. Therefore, the current results primarily demonstrate its technical characteristics in task organization, spatial anchoring, cross-skill handoff, and process tracing. They are not sufficient to determine whether its overall performance is superior to that of existing systems. The performance of SmartFarming is also constrained by the quality of underlying data and agricultural models. Errors in remote-sensing inversion, image recognition, crop-model parameters, and data completeness may propagate through task chains [17,66]. Production-relevant tasks, such as fertilization, irrigation, crop protection, and agricultural machinery operation, therefore require model-applicability checks, exception-triggered termination, result review, and human confirmation.

Future research should focus on three areas. First, unified benchmarks and repeated evaluations across multiple scenarios should incorporate ground-truth data to quantify closed-loop execution efficiency, spatial consistency, and the extent of human intervention required under fixed protocols [67]. Second, long-term memory modules should be developed, deployed, and evaluated using state-consolidation and retrieval strategies informed by historical field and crop conditions [54,68]. Third, the engineering pipeline for controlled capability evolution should be refined while preserving review procedures and safety safeguards [32,35]. Structured runtime feedback could be used to generate proposals for verifiable, traceable, and reversible skill upgrades, thereby supporting the controlled extension of agent capabilities.

6 Conclusions

This study proposes SmartFarming, an AI GIS-based agent framework for farm management. The framework uses an Agent Harness to coordinate task orchestration, capability routing, hierarchical memory, execution observability, and governance

mechanisms. It composes capabilities through skills and contracts and transfers task products through artifacts, while manifests record data lineage. Together with a GIS-based spatial closed loop and a controlled skill-upgrade mechanism, it unifies field inputs, operational computations, result visualization, and user review within a single task control plane. The main contribution is an agricultural agent runtime architecture characterized by spatial anchoring, data lineage, and contract-based handoff. This architecture provides a concrete engineering pathway for systematically integrating agronomic models, remote-sensing data, and management tools. Through representative cases spanning satellite remote-sensing monitoring, UAV multispectral inversion, ground-based RGB image recognition, crop model simulation, and model-remote sensing coupling, this study demonstrates that SmartFarming can be instantiated and that its end-to-end task chains are operational. The results show that the system can translate natural-language-driven management requirements into multistep task orchestration constrained by spatial, temporal, and agronomic conditions. In the test cases, it preserves the structured transfer of field boundaries, variable semantics, measurement units, and algorithm versions, forming a continuous closed loop from data sensing and model analysis to review of management outputs. This study provides an organizational foundation for the coordinated application of remote-sensing monitoring, crop models, and farm-management tools in practical farm-management settings. It also provides technical support for the development of traceable and extensible agricultural agents.

Acknowledgement

This study was supported by the National major agricultural science and technology projects (0305), National Natural Science Foundation of China (32572202), the Jiangsu Collaborative Innovation Center for Modern Crop Production (JCICMCP), the 111 Project and the Program of China Scholarship Council (Grant No. 202406850105).

CRediT author contributions

Ding Yang: Conceptualization; Data curation; Formal analysis; Investigation; Software; Writing – original draft; Writing – review & editing.

Yancheng Wang: Data curation; Formal analysis; Investigation; Methodology; Writing – original draft; Writing – review & editing.

Keqiang Sun: Data curation; Investigation; Methodology; Software; Writing – original draft; Writing – review & editing.

Yue Zhang: Data curation; Investigation; Software; Visualization.

Yanan Wang: Data curation; Investigation; Software; Validation.

Yuru Ye: Software; Data curation.

Yifan Hong: Software; Data curation.

Yanchao Li: Software; Investigation.

Wei Zhou: Software; Resources.

Jinhan Zhang: Data curation; Investigation.

Zhaode Yin: Investigation; Data curation.

Xiaohu Zhang: Formal analysis; Writing – review & editing.

Liujun Xiao: Formal analysis; Writing – review & editing.

Bing Liu: Formal analysis; Writing – review & editing.

Leilei Liu: Formal analysis; Writing – review & editing.

Tao Cheng: Formal analysis; Writing – review & editing.

Xia Yao: Formal analysis; Writing – review & editing.

Chongya Jiang: Formal analysis; Writing – review & editing.

Hengbiao Zheng: Formal analysis; Writing – review & editing.

Xiaojun Liu: Formal analysis; Writing – review & editing.

Qiang Cao: Formal analysis; Writing – review & editing.

Yongcao Tian: Formal analysis; Writing – review & editing.

Yan Zhu: Formal analysis; Writing – review & editing.

Weixing Cao: Formal analysis; Writing – review & editing.

Yihua Zhu: Conceptualization; Project administration; Supervision; Writing – review & editing.

Liang Tang: Conceptualization; Formal analysis; Funding acquisition; Investigation; Project administration; Supervision; Writing – review & editing.

References

- [1] WALTER A, FINGER R, HUBER R, et al. Smart farming is key to developing sustainable agriculture[C]. Proceedings of the National Academy of Sciences of the United States of America, 2017, 114(24): 6148-6150. <https://doi.org/10.1073/pnas.1707462114>.
- [2] BASSO B, ANTLE J. Digital agriculture to design sustainable agricultural systems[J]. Nature Sustainability, 2020, 3: 254-256. <https://doi.org/10.1038/s41893-020-0510-0>.
- [3] FOUNTAS S, CARLI G, SØRENSEN C G, et al. Farm management information systems: Current situation and future perspectives[J]. Computers and Electronics in Agriculture, 2015, 115: 40-50. <https://doi.org/10.1016/j.compag.2015.05.011>.
- [4] BAYER. Crop Science Sustainability Progress Report 2024[EB/OL]. 2024. <https://www.bayer.com/sites/default/files/bayer-cs-sustainability-progress-report-2024-11-07>.
- [5] JOHN DEERE. John Deere Operations Center[EB/OL]. <https://www.deere.com/en-us/products-solutions/technology-solutions/precision-ag-technology/operations-center/features>
- [6] SYNGENTA. Cropwise: Digital Farming Connected[EB/OL]. <https://www.cropwise.com/>
- [7] BASF. BASF showcases data-driven farming with xarvio FIELD MANAGER at 2024 Science Behind event[EB/OL]. <https://www.basf.com/us/en/media/news-releases/2024/03/basf-showcases-data-driven-farming-with-xarvio-field-manager-at->. 2024
- [8] URDU D, BERRE A J, SUNDMAEKER H, et al. Aligning interoperability architectures for digital agri-food platforms[J]. Computers and Electronics in Agriculture, 2024, 224: 109194. <https://doi.org/10.1016/j.compag.2024.109194>.
- [9] KRUIZE J W, WOLFERT J, SCHOLTEN H, et al. A reference architecture for Farm Software Ecosystems[EB/OL]. Computers and Electronics in Agriculture 125, 12–28. 2016. <https://doi.org/10.1016/j.compag.2016.04.011>.
- [10] OPENAI. GPT-4 Technical Report[EB/OL]. arXiv. 2023. <https://doi.org/10.48550/arXiv.2303.08774>.
- [11] GEMINI TEAM, ANIL R, BORGEAUD S, et al. Gemini: A Family of Highly Capable Multimodal Models[EB/OL]. arXiv. 2023. <https://doi.org/10.48550/arXiv.2312.11805>.
- [12] DIDWANIA K, SETH P, KASLIWAL A, et al. AgriLLM: Harnessing Transformers for Framer Queries[C]. In Proceedings of the Third Workshop on NLP for Positive Impact, 179-187. Association for Computational Linguistics. 2024. <https://doi.org/10.18653/v1/2024.nlp4pi-1.16>.

- [13] AWAIS M, ALHARTHI A H S A, KUMAR A, et al. AgroGPT: Efficient Agricultural Vision-Language Model with Expert Tuning[C]. In 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 5687-5696. 2025.
<https://doi.org/10.1109/WACV61041.2025.00555>.
- [14] ZHU H, QIN S, SU M, et al. Harnessing large vision and language models in agriculture: A review[EB/OL]. arXiv. 2024. <https://doi.org/10.48550/arXiv.2407.19679>.
- [15] KUSKA M T, WAHABZADA M, PAULUS S. AI for crop production - Where can large language models (LLMs) provide substantial value?[EB/OL]. Computers and Electronics in Agriculture, 221, 108924. 2024. <https://doi.org/10.1016/j.compag.2024.108924>.
- [16] KPODO J, NEJADHASHEMI A P, EESWARAN R. Evaluating the logical and mathematical reasoning capabilities of language models in agricultural water management[J]. Computers and Electronics in Agriculture, 2026, 245: 111533.
<https://doi.org/10.1016/j.compag.2026.111533>.
- [17] WANG B, JÄGERMEYER J, O'LEARY G J, et al. Pathways to identify and reduce uncertainties in agricultural climate impact assessments[J]. Nature Food, 2024, 5: 550-556.
<https://doi.org/10.1038/s43016-024-01014-w>.
- [18] SUMERS T R, YAO S, NARASIMHAN K, et al. Cognitive architectures for language agents[EB/OL]. Transactions on Machine Learning Research. 2024.
<https://doi.org/10.48550/arXiv.2309.02427>.
- [19] YAO S, ZHAO J, YU D, et al. ReAct: Synergizing reasoning and acting in language models[C]. International Conference on Learning Representations. 2023.
<https://doi.org/10.48550/arXiv.2210.03629>.
- [20] XI Z, CHEN W, GUO X, et al. The rise and potential of large language model based agents: a survey[EB/OL]. Science China Information Sciences 68, 121101. 2025.
<https://doi.org/10.1007/s11432-024-4222-0>.
- [21] BAJORIA S, RANJAN S, MA, et al. From Language Models to Agentic AI: A Survey of Autonomous, Action-Enabled, and Collaborative LLM Agents[J]. Cognitive Computation, 2026, 18: 103. <https://doi.org/10.1007/s12559-026-10619-1>.
- [22] XU W, YU Z, MU B, et al. RS-Agent: Automating remote sensing tasks through intelligent agent[EB/OL]. arXiv. 2024. <https://doi.org/10.48550/arXiv.2406.07089>.
- [23] CHEN Y, WANG W, LOBRY S, et al. An LLM agent for automatic geospatial data analysis[EB/OL]. arXiv. 2024. <https://doi.org/10.48550/arXiv.2410.18792>.
- [24] CAI W, QI F, ZHA L, et al. Few-shot and interpretable agentic framework based on large language models for data-efficient plant phenotyping[J]. Computers and Electronics in Agriculture, 2026, 243: 111382. <https://doi.org/10.1016/j.compag.2025.111382>.
- [25] JIANG J, YAN L, LIU H, et al. Knowledge assimilation: Implementing knowledge-guided agricultural large language model[J]. Knowledge-Based Systems, 2025, 314: 113197.
<https://doi.org/10.1016/j.knosys.2025.113197>.
- [26] KANDAMALI D F, PORTER W M, PORTER E, et al. CottonBot: An AI-driven cotton farming assistant and irrigation advisor using LLM-RAG and agentic AI tools[J]. Smart Agricultural Technology, 2025, 12: 101640. <https://doi.org/10.1016/j.atech.2025.101640>.
- [27] CHAMARA N, PAN Y, TAGHVAEIAN S, et al. Can generative AI make farming decisions?[J]. Current status and future pathways-A case study in row crop production with ChatGPT. Artificial Intelligence in Agriculture, 2026, 16: 837-853.
<https://doi.org/10.1016/j.aiia.2026.03.008>.
- [28] JANOWICZ K, GAO S, MCKENZIE G, et al. GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond[EB/OL]. Int. J. Geogr. Inf. Sci. 34, 625–636. 2020. <https://doi.org/10.1080/13658816.2019.1684500>.
- [29] MEYER B. Applying “Design by contract”[J]. Computer, 1992, 25(10): 40-51.
<https://doi.org/10.1109/2.161279>.
- [30] ISO/IEC. ISO/IEC 25010:2023 - Systems and Software Engineering - SQuaRE - Product Quality Model, ISO/IEC International Standard[S]. 2023
- [31] OPENAPI INITIATIVE. OpenAPI Specification, Version 3.1.0[EB/OL].
<https://spec.openapis.org/oas/v3.1.0.html>. 2021
- [32] AMERSHI S, BEGEL A, BIRD C, et al. Software Engineering for Machine Learning: A Case Study[C]. In ICSE-SEIP 2019. 2019. <https://doi.org/10.1109/ICSE-SEIP.2019.00042>.

- [33] BRECK E, CAI S, NIELSEN E, et al. The ML Test Score: A Rubric for ML Production Readiness and Technical Debt Reduction[C]. In 2017 IEEE International Conference on Big Data. 2017. <https://doi.org/10.1109/BigData.2017.8258038>.
- [34] SCULLEY D, HOLT G, GOLOVIN D, et al. Hidden Technical Debt in Machine Learning Systems[J]. In Advances in Neural Information Processing Systems 28. 2015
- [35] TABASSI E. Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1[R]. 2023. <https://doi.org/10.6028/NIST.AI.100-1>.
- [36] ANTHROPIC. Model Context Protocol Specification (2024-11-05) Model Context Protocol[EB/OL]. <https://modelcontextprotocol.io/specification/2024-11-05>. 2024
- [37] LU C, LU C, LANGE R T, et al. The AI Scientist: Towards fully automated open-ended scientific discovery[EB/OL]. arXiv. 2024. <https://doi.org/10.48550/arXiv.2408.06292>.
- [38] REDMON J, DIVVALA S, GIRSHICK R, et al. You only look once: Unified, real-time object detection[C]. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR): 779-788. 2016. <https://doi.org/10.1109/CVPR.2016.91>.
- [39] DAVID E, MADEC S, SADEGHI-TEHRAN P, et al. Global Wheat Head Detection (GWHD) dataset: A large and diverse dataset of high-resolution RGB-labelled images to develop and benchmark wheat head detection methods[EB/OL]. Plant Phenomics, 3521852. 2020. <https://doi.org/10.34133/2020/3521852>.
- [40] WANG X, YANG W, LV Q, et al. Field rice panicle detection and counting based on deep learning[J]. Frontiers in Plant Science, 2022, 13: 966495. <https://doi.org/10.3389/fpls.2022.966495>.
- [41] ZHU Y, LIU L, LIU B, et al. The CropGrow model[M]. In G. Hoogenboom (Ed.), Current crop models: State-of-the-art and future developments (pp. 689-733). Burleigh Dodds Science Publishing. 2025. <https://doi.org/10.19103/AS.2025.0155.24>.
- [42] TANG L, ZHU Y, HANNAWAY D, et al. RiceGrow: A Rice Growth and Productivity Model[J]. NJAS: Wageningen Journal of Life Sciences, 2009, 57(1): 83-92. <https://doi.org/10.1016/j.njas.2009.12.003>.
- [43] YAN M, CAO W, LUO W, et al. A Mechanistic Model of Phasic and Phenological Development of Wheat[J]. I. Assumption and Description of the Model. Chinese Journal of Applied Ecology, 2000, 11(3): 355-359
- [44] ROSENZWEIG C, JONES J W, HATFIELD J L, et al. The Agricultural Model Intercomparison and Improvement Project (AgMIP): Protocols and pilot studies[J]. Agricultural and Forest Meteorology, 2013, 170: 166-182. <https://doi.org/10.1016/j.agrformet.2012.09.011>.
- [45] LI T, HASEGAWA T, YIN X, et al. Uncertainties in predicting rice yield by current crop models under a wide range of climatic conditions[J]. Global Change Biology, 2015, 21(3): 1328-1341. <https://doi.org/10.1111/gcb.12758>.
- [46] ASSENG S, EWERT F, MARTRE P, et al. Rising temperatures reduce global wheat production[J]. Nature Climate Change, 2015, 5: 143-147. <https://doi.org/10.1038/nclimate2470>.
- [47] WANG E, MARTRE P, ZHAO Z, et al. The uncertainty of crop yield projections is reduced by improved temperature response functions[J]. Nature Plants, 2017, 3: 17102. <https://doi.org/10.1038/nplants.2017.102>.
- [48] CHOWA S S, ALVI R, RAHMAN S S, et al. From language to action: A review of large language models as autonomous agents and tool users[J]. Artificial Intelligence Review, 2026, 59: 71. <https://doi.org/10.1007/s10462-025-11471-9>.
- [49] RAMOS M C, COLLISON C J, WHITE A D. A review of large language models and autonomous agents in chemistry[J]. Chemical Science, 2025, 16(6): 2514-2572. <https://doi.org/10.1039/D4SC03921A>.
- [50] KIM Y, GU K, PARK C, et al. Capable language models can outgrow the benefits of collaboration[J]. Nature Machine Intelligence, 2026, 8: 1157-1172. <https://doi.org/10.1038/s42256-026-01268-y>.
- [51] BU D, SUN J, LI K, et al. Empowering AI data scientists using a multi-agent LLM framework with self-evolving capabilities for autonomous, tool-aware biomedical data analyses[EB/OL]. Nature Biomedical Engineering. Advance online publication. 2026. <https://doi.org/10.1038/s41551-026-01634-6>.

- [52] CHEN F, STOGIANNIDIS I, WOOD A, et al. A conversational multi-agent AI system for automated plant phenotyping[J]. *Nature Communications*, 2026, 17: 6391.
<https://doi.org/10.1038/s41467-026-71090-y>.
- [53] BRAN A M, COX S, SCHILTER O, et al. Augmenting large language models with chemistry tools[J]. *Nature Machine Intelligence*, 2024, 6: 525-535.
<https://doi.org/10.1038/s42256-024-00832-8>.
- [54] ZHENG J, SHI C, CAI X, et al. Lifelong learning of large language model based agents: A roadmap[J]. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026, 48(5): 5552-5571. <https://doi.org/10.1109/TPAMI.2025.3650546>.
- [55] WANG Z, CHEN Z, YANG Z, et al. Empowering biomedical evidence exploration and synthesis with deep knowledge graph research[J]. *Nature Machine Intelligence*, 2026, 8: 1142-1156. <https://doi.org/10.1038/s42256-026-01266-0>.
- [56] WOLFERT S, GE L, VERDOUW C N, et al. Big Data in Smart Farming - A review[J]. *Agricultural Systems*, 2017, 153: 69-80. <https://doi.org/10.1016/j.agsy.2017.01.023>.
- [57] CUNO S, LÄMMEL P. Use of Semantic Artefacts in Agricultural Data-Driven Service Development[EB/OL]. *DATA* 2024. 2024. <https://doi.org/10.5220/0012760400003756>.
- [58] DONG H, YANG H, MIAO Y, et al. Operating smart grids by customizing large model agents[J]. *Communications Engineering*, 2026, 5: 114.
<https://doi.org/10.1038/s44172-026-00709-1>.
- [59] KLANG E, OMAR M, RAUT G, et al. Orchestrated multi agents sustain accuracy under clinical-scale workloads compared to a single agent[J]. *npj Health Systems*, 2026, 3: 23.
<https://doi.org/10.1038/s44401-026-00077-0>.
- [60] MA C, RAN M, SONG J, et al. Multi-agent AI for end-to-end air quality analysis and decision support[J]. *npj Clean Air*, 2026, 2: 59. <https://doi.org/10.1038/s44407-026-00102-4>.
- [61] SHI T, LI Y, LI Z, et al. A collaborative agent with two lightweight synergistic models for autonomous crystal materials research[J]. *Nature Machine Intelligence*, 2026, 8: 1441-1456.
<https://doi.org/10.1038/s42256-026-01298-6>.
- [62] MOURTZINIS S, SILVA T S, LO J C M, et al. A human in the loop approach to applying large language models for farm management insight[J]. *Scientific Reports*, 2026, 16: 1273.
<https://doi.org/10.1038/s41598-025-30991-6>.
- [63] PINCHEIRA M, ANTONELLI F, VECCHIO M. AGRICLIMA: Towards a Federated Platform for Spatiotemporal Risk Analysis in Agriculture[J]. *Agriculture*, 2025, 15(23): 2450.
<https://doi.org/10.3390/agriculture15232450>.
- [64] JIANG J, ZHANG Q, JOHANSEN K, et al. UAV-satellite synergies for crop monitoring: current landscape, challenges, and future directions[J]. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2026, 239: 908-927. <https://doi.org/10.1016/j.isprsjprs.2026.06.038>.
- [65] LIU X, YU H, ZHANG H, et al. AgentBench: Evaluating LLMs as Agents[C]. In *International Conference on Learning Representations (ICLR 2024)*. 2024
- [66] CHAPAGAIN R, REMENYI T A, HARRIS R M B, et al. Decomposing crop model uncertainty: A systematic review[J]. *Field Crops Research*, 2022, 279: 108448.
<https://doi.org/10.1016/j.fcr.2022.108448>.
- [67] YAO S, SHINN N, RAZAVI P, et al. τ -bench: A Benchmark for Tool-Agent-User Interaction in Real-World Domains[C]. In *International Conference on Learning Representations (ICLR 2025)*. 2025
- [68] ZHONG W, GUO L, GAO Q, et al. MemoryBank: Enhancing Large Language Models with Long-Term Memory[C]. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024, 38(17): 19724-19731. <https://doi.org/10.1609/aaai.v38i17.29946>.
- [69] CHEN Z, PETSCH A N, ISRAELSKI A J, et al. An agentic artificially intelligent X-ray scientist[J]. *Nature Machine Intelligence*, 2026, 8: 1075-1086.
<https://doi.org/10.1038/s42256-026-01261-5>.
- [70] WANG L, MA C, FENG X, et al. A Survey on Large Language Model Based Autonomous Agents[J]. *Frontiers of Computer Science*, 2024, 18(6): 186345.
<https://doi.org/10.1007/s11704-024-40231-1>.

Supplementary material

SmartFarming: a GIS-based agent framework for farm management

This file contains the supplementary material for the manuscript above. The following items are provided:

Text S1. Supplementary text.

Table S1. Supplementary table 1.

Table S2. Supplementary table 2.

Table S3. Supplementary table 3.

Note: The main text (33 pages, 9 figures) precedes this supplementary section. In-text references such as "Text S1" and "Table S1" refer to the items listed above.

SmartFarming: a GIS-based agent framework for farm management

Names of authors: Ding Yang^{#a}, Yancheng Wang^{#a}, Keqiang Sun^{#a}, Yue Zhang^a, Yanan Wang^a, Yuru Ye^b, Yifan Hong^b, Yanchao Li^b, Wei Zhou^b, Jinhan Zhang^a; Zhaode Yin^c, Xiaohu Zhang^{a, b, d}, LiuJun Xiao^a, Bing Liu^a, Leilei Liu^a, Tao Cheng^a, Xia Yao^a, Chongya Jiang^a, Hengbiao Zheng^a, Xiaojun Liu^a, Qiang Cao^a, Yongcao Tian^a, Yan Zhu^a, Weixing Cao^a, Yihua Zhu^{e*}, Liang Tang^{a, b, c, d*}

^a National Engineering and Technology Center for Information Agriculture, Engineering Research Center for Smart Agriculture, Ministry of Education, Key Laboratory for Crop System Analysis and Decision Making, Ministry of Agriculture and Rural Affairs, Jiangsu Key Laboratory for Information Agriculture, Jiangsu Collaborative Innovation Center for Modern Crop Production, Nanjing Agricultural University, Nanjing 210095, Jiangsu, PR China

^b Shennong Intelligent Agriculture Research Institute Nanjing Co., Ltd

^c Sanya Institute of Nanjing Agricultural University, Sanya, Hainan 572025, China

^d Nanjing Agricultural University (Legeous) AI Agriculture Research Institute

^e College of Information Science and Technology, Nanjing Agricultural University, Nanjing, Jiangsu Province China

These authors contributed equally: Ding Yang, Yancheng Wang and Keqiang Sun

Corresponding authors: Liang Tang (E-mail: tangl@njau.edu.cn) and Yihua Zhu (E-mail: zhyh@njau.edu.cn)

Table S1. Models and algorithms used in the SmartFarming case studies

Skill	Model or algorithm	Computation	Function	Inputs	Outputs	References
Remote sensing monitoring	Normalized difference vegetation index (NDVI)	Two-band spectral index	Characterizes canopy greenness and relative growth and supports growth classification	Sentinel-2 near-infrared (NIR) and red reflectance	NDVI raster	Rouse et al., 1974
	Enhanced vegetation index (EVI)	Three-band spectral index	Accounts for soil background and aerosol effects in greenness estimates	NIR, red, and blue reflectance	EVI raster	Huete et al., 2002
	Normalized difference red-edge index (NDRE)	Red-edge spectral index	Characterizes canopy chlorophyll and nitrogen status and reduces saturation relative to red-based NDVI	NIR and red-edge (RE) reflectance	NDRE raster	Barnes et al., 2000; Delegido et al., 2011; Clevers and Gitelson, 2013
	Normalized difference water index (NDWI)	Two-band spectral index	Produces an open-water-sensitive layer using the displayed McFeeters formulation	Green and NIR reflectance	NDWI raster	McFeeters, 1996
Variable-rate nitrogen application	Rice critical nitrogen dilution and nitrogen-deficit topdressing	OSAVI-based biomass and NDRE-based plant N accumulation; nitrogen balance	Generates rice jointing and panicle topdressing prescriptions from estimated nitrogen deficit	NIR, RE, and red reflectance; region, cultivar, fertilizer nitrogen fraction, and application frequency	Fertilizer product rate (kg ha ⁻¹)	Fu et al., 2024

Foliar topdressing	Wheat critical nitrogen dilution and nitrogen-deficit topdressing	CIre-based biomass and EVI-based LAI; LAI-normalized nitrogen status; nitrogen balance	Generates wheat jointing and ear topdressing prescriptions from estimated nitrogen deficit	NIR, RE, red, and blue reflectance; region, cultivar, and fertilizer nitrogen fraction	Fertilizer product rate (kg ha ⁻¹)	Fu et al., 2024
	P5-P95 indicator mapping	Quantile-based linear interpolation	Maps crop nitrogen or canopy indicators to foliar urea rates	Rice NDRE	Urea rate (kg ha ⁻¹)	Zhang et al., 2024
	Wheat P5-P95 mapping of leaf area index (LAI)	EVI-derived LAI and quantile-based linear interpolation	Maps wheat LAI to foliar urea rates	Wheat EVI	Urea rate (kg ha ⁻¹)	Li et al., 2023a
UAV monitoring	UAV NDVI	Two-band spectral index	Maps canopy greenness and may provide a predictor for trait estimation	UAV NIR and red reflectance	NDVI raster	Rouse et al., 1974
UAV trait estimation	Aboveground biomass (AGB) model	Machine learning mapping from spectral indices, textural indices and plant height quantiles	Estimates the spatial distribution of AGB	UAV RGB, multispectral imagery, crop, and growth stage	AGB (kg ha ⁻¹)	Lu et al., 2019b; Zheng et al., 2019
	Leaf area index (LAI) model	Linear mapping from NDVI quantiles	Estimates plot-scale canopy leaf area	UAV NDVI, crop, and growth stage	LAI (m ² m ⁻²)	Yu et al., 2024
	Leaf age (LA) model	Linear mapping from NDVI	Generates a leaf-age layer	UAV NDVI, crop, and growth stage	Leaf-age index	Wang et al., 2024

	Leaf nitrogen accumulation (LNA) model	Linear mapping from NDVI quantiles	Maps leaf nitrogen accumulation for nitrogen-status assessment	UAV NDVI, crop, and growth stage	LNA (kg ha ⁻¹)	Zheng et al., 2018a
	Leaf nitrogen concentration (LNC) model	Linear mapping from NDVI quantiles	Estimates LNC for nitrogen-status assessment	Multispectral or texture features	LNC (%)	Lu et al., 2019a; Zheng et al., 2018b; Zheng et al., 2020
	Grain protein content (GPC) model	Linear mapping from NDVI quantiles	Estimates spatial variation in grain protein content	Multispectral features	GPC (%)	Li et al., 2023b
	Grain yield model	Linear mapping from NDVI quantiles	Produces spatial yield estimates before harvest	Multispectral features	Grain yield (kg ha ⁻¹)	Lu et al., 2025; Zheng et al., 2022
	You Only Look Once (YOLO) object detectors	Single-stage convolutional object detection	Detects and counts rice panicles and wheat spikes and locates rice disease regions	Red, green, and blue (RGB) images	Detections, classes, confidence scores, counts, and annotated images	Redmon et al., 2016; Zang et al., 2022; Wang et al., 2022
Rice growth prediction	RiceGrow	Deterministic process-based rice model	Simulates daily phenology, LAI, organ biomass, AGB, and grain yield under genotype × environment × management conditions	Daily weather, soil, cultivar parameters, sowing or transplanting dates, irrigation, and fertilization	Daily rice states, phenology, and final yield	Tang et al., 2009; Zhu et al., 2020

Wheat growth prediction	WheatGrow	Deterministic process-based wheat model	Simulates daily phenology, LAI, organ biomass, AGB, and grain yield under genotype $\times$ environment $\times$ management conditions	Weather, soil, cultivar, and management data	Daily wheat states, phenology, and yield	Yan et al., 2000; Zhu et al., 2020
Remote sensing regression	Modified triangular vegetation index 2 (MTVI2) quadratic regression and random forest	Remote sensing regression and machine learning	Converts Sentinel-2 observations into LAI inputs for RiceGrow assimilation	Sentinel-2 surface-reflectance bands, band mapping, and image and sowing dates	LAI ($\text{m}^2 \text{m}^{-2}$)	Reyes J.S., et al. 2023
	S2MREP red-edge-position regression as labeled in the source	Remote sensing regression	Converts Sentinel-2 red-edge observations into LAI inputs for WheatGrow assimilation	Sentinel-2 B04, B05, B06, and B07 surface reflectance	LAI ($\text{m}^2 \text{m}^{-2}$)	Li et al., 2023a
Biomass estimation	Two-leaf light-use efficiency (TL-LUE) and carbon-to-biomass conversion	Daily gross primary production integration	Estimates cumulative AGB from light use and environmental stresses	LAI, Sentinel-2, daily weather, carbon dioxide, crop, and period	AGB (g m^{-2})	Wu et al., 2026
Sequential assimilation	Ensemble Kalman filter (EnKF)	Sequential Bayesian state estimation	Updates RiceGrow ensemble states with remotely sensed LAI and propagates the analysis into subsequent forecasts and yield formation	Prior ensemble states and LAI observations with an error model	Posterior ensemble states and yield	Evensen, 1994, 2003; Huang et al., 2013

Yield mapping	Posterior RiceGrow yield projection and raster export	Process-model output mapping	Converts posterior state trajectories into pixel- or plot-level yield estimates and georeferenced layers	Posterior RiceGrow states and an area of interest (AOI) grid	Yield trajectories and GeoTIFF files or maps	Guo et al., 2019
---------------	---	------------------------------	--	--	--	------------------

Table S2. Remote-sensing skills in SmartFarming.

Platform	Skill category	Main function	Main inputs	Main outputs
Satellite remote sensing	Image retrieval and clipping	Retrieves Sentinel-2 and other satellite images for the AOI and exports clipped images	AOI, date range, and cloud-cover threshold	Multiband GeoTIFF files and satellite image manifest
	Crop growth estimation	Derives crop growth indicators, such as LAI, from multispectral satellite images	Satellite image manifest, band mapping, and field boundaries	Growth-indicator rasters, field statistics, and growth manifest
UAV remote sensing	Multispectral image parsing	Reads UAV multispectral GeoTIFF files and identifies key bands, including red, red edge, and NIR	UAV multispectral imagery and band_map or band metadata	Standardized UAV image inputs and band mapping
	Crop growth estimation	Estimates growth indicators using UAV NDVI and growth-stage models	UAV multispectral imagery, growth stage, and field boundaries	NDVI, AGB, LA, LAI, LNA, LNC, GPC, and YIELD layers
	Growth classification and field statistics	Classifies estimated indicators and summarizes them by field	Indicator GeoTIFF files and field boundaries	Class maps, comma-separated values (CSV) files of field statistics, and interactive map layers

Table S3. Datasets in the SmartFarming data foundation.

Category	Dataset	Spatial resolution	Temporal coverage	Format	Key attributes and framework use
Meteorological data	Station observations	Station locations	Daily; annual files	CSV	Temperature and precipitation drive crop models and provide weather context for risks at a given place and time.
Soil data (1 km)	BD, CL, SA, SI, SOM, pH, CEC, TN, TP, TK, AN, AP, AK	1 km	Static	NetCDF	Soil physical properties, fertility, and total and available N, P, and K support fertilization, suitability assessment, and growth models.
Soil data (250 m)	Ksat, WCpF2, WCsat	250 m	Static	GeoTIFF	Hydraulic conductivity, field capacity, and saturated water content at five depths (0-5, 5-15, 15-30, 30-60, and 60-100 cm)support water-status diagnosis and irrigation estimation.
Terrain data	MERIT DEM, slope, and aspect	1 km working layers	Static (2017)	GeoTIFF	Terrain constraints support drainage assessment, crop suitability analysis, and machinery planning.
Crop management	Nitrogen rates for wheat and rice	10 km working layers	1961-2020	HDF5	Historical nutrient inputs provide regional context and do not replace field measurements.
	Cultivar parameters	Subregion	Static	CSV	Cultivar-specific parameters drive crop growth models after matching by location and cultivar.
Phenological data	Wheat, maize, and rice phenology	1 km	2000-2015	GeoTIFF	Crop-specific seasonal onset, heading, and maturity define crop calendars, image time windows, and farm operation stages.
Administrative boundaries	County boundaries	County	Static	Shapefile	Boundaries support spatial aggregation and policy-related queries and link the AOI to administrative units.
Cropping systems	Crop distributions and cropping intensity	1 km	Decadal working layers	CSV	Cropping intensity and types provide background on cropping systems and rotations for interpreting regional tasks.




Fig. S1. Skill execution workflow.

Text S1. SmartFarming data foundation: data sources, standardization, and integration with the agent system.

1 Data sources and categories

1.1 Meteorological data

Daily precipitation and maximum and minimum air temperatures for 1961-2025 are obtained from the global station network of the National Oceanic and Atmospheric Administration (NOAA). These records provide meteorological forcing for crop model skills and contextual information for agent reasoning (Smith et al., 2011). Missing-value entries are removed, temperatures are converted from degrees Fahrenheit to degrees Celsius, and precipitation is converted from inches to millimeters. Data from the nearest available station are used for locations without observations.

1.2 Soil data

The soil data describe physical and chemical properties relevant to soil fertility, water retention, and nutrient availability. The 1 km products from the China Soil

Science Data Center (<https://soil.geodata.cn/>) include bulk density (BD), clay content (CL), sand content (SA), silt content (SI), soil organic matter (SOM), soil pH, cation exchange capacity (CEC), total nitrogen (TN), total phosphorus (TP), total potassium (TK), available nitrogen (AN), available phosphorus (AP), and available potassium (AK). These products are stored in Network Common Data Form (NetCDF) files with eight depth bands, each representing a standard soil depth interval.

Saturated hydraulic conductivity (Ksat), water content at pF 2, used here as field capacity (WCpF2), and saturated water content (WCsat) are obtained from HiHydroSoil v2.0 (Simons et al., 2020). This global dataset provides soil hydraulic properties at 250 m resolution, derived from SoilGrids 2.0 (Poggio et al., 2021) using the pedotransfer functions of Tóth et al. (2015). The selected layers are stored as multiband GeoTIFF files and ordered from 0-5 to 60-100 cm depth to support simulations of soil water dynamics by depth. For framework-level regional queries, the soil layers are resampled to a common $0.1^{\circ} \times 0.1^{\circ}$ working grid using nearest-neighbor resampling..

1.3 Terrain data

The Multi-Error-Removed Improved-Terrain digital elevation model (MERIT DEM; Yamazaki et al., 2017) provides elevation data for deriving slope and aspect. These terrain variables influence solar radiation interception, drainage patterns, and conditions for mechanized operations. They therefore provide context for agent reasoning about crop suitability and field operation planning.

1.4 Crop management data and cultivar parameters

Crop management inputs draw on several datasets. Historical nitrogen fertilizer inputs for wheat and rice are represented in the working database at 10 km resolution, with Lu and Tian (2017) cited as the source of the historical fertilizer data. The published source covers 1961-2013 at $0.5^{\circ} \times 0.5^{\circ}$ resolution. The framework carries the 2013 values forward to fill missing records for 2014-2020. Irrigation status is assigned using the Global Map of Irrigation Areas from the Food and Agriculture

Organization of the United Nations (FAO; <https://data.apps.fao.org/catalog/dataset/f79213a0-88fd-11da-a88f-000d939bc5d8>).

Areas with an irrigation fraction below 10% are classified as rainfed under the classification rule of the framework. Rice and wheat cultivar parameters are used in RiceGrow (Tang et al., 2009) and WheatGrow (Yan et al., 2000; Zhu et al., 2020; Cao et al., 2024). Calibration and evaluation of these parameter sets draw on the cited studies (Wang et al., 2025; Cao et al., 2024).

Rice parameters include intrinsic earliness (IE), photoperiod sensitivity (PS), temperature sensitivity (TS), basic filling factor (BFF), optimal temperature (To), harvest index (HI), and maximum CO₂ assimilation rate (AMX), whereas wheat parameters include photoperiod sensitivity (PS), physiological vernalization time (PVT), temperature sensitivity (TS), intrinsic earliness (IE), basic filling factor (FDF), specific leaf area (SLA), harvest index (HI), and maximum CO₂ assimilation rate (AMX), as defined in Tables S1 and S2. The parameter sets are cataloged by subregion to support cultivar recommendations for different cropping systems.

1.5 Phenological data

Gridded phenological data for rice and wheat have a spatial resolution of 1 km and cover 2000-2015 (Luo et al., 2020). They describe crop-specific developmental stages, including the start of the growing season, heading, and maturity. Crop calendars are assigned to grid cells using the nearest available location and constrain farm recommendations that depend on crop timing.

1.6 Administrative boundaries and cropping systems

County boundaries support spatial aggregation and queries involving administrative and policy contexts. The framework also uses 1 km cropping-system layers at decadal intervals to represent the spatial distribution of cropping intensity and rotation patterns. Together with grain-crop land information (Liu et al., 2024), these layers provide regional land-use context for agent reasoning.

2 Data processing workflow

The data workflow comprises acquisition, standardization, and deployment as data services.

Source data formats include NetCDF (.nc), GeoTIFF (.tif), Hierarchical Data Format version 5 (HDF5; .h5), CSV, and Shapefile (.shp). The source files are retained in their original formats before database ingestion to preserve metadata and avoid information loss during conversion.

All datasets are transformed to the World Geodetic System 1984 (WGS 84) geographic coordinate reference system, identified as EPSG:4326 in the European Petroleum Survey Group (EPSG) geodetic parameter dataset. Raster datasets are aligned to a common grid to support joint queries across variables. Soil properties with depth information use a shared layer index covering 0-5, 5-15, 15-30, 30-60, and 60-100 cm. Temporal attributes follow the International Organization for Standardization (ISO) 8601 date and time format, and categorical variables use standardized codes. Station observations are georeferenced by longitude and latitude.

Standardized datasets are loaded into a MySQL relational database. The schema organizes records by spatial extent and temporal validity, with spatial indexing of geographic locations to accelerate location-based retrieval. A representational state transfer (REST) API provides a parameterized interface to the database and accepts longitude, latitude, and an optional time range. The API returns structured JSON responses containing the attributes available for the requested location and time window. SmartFarming skills for satellite image analysis, UAV data fusion, and yield prediction invoke these APIs through the Model Context Protocol (MCP) to support decisions with retrieved data.

3 Integration with the agent system

The API-mediated architecture separates data services from agent reasoning. SmartFarming can therefore retrieve agricultural data for a specific location without embedding large geospatial datasets directly in the model context. For a request to predict wheat growth under current soil and weather conditions, the relevant skill

sends an API request with the location to the data foundation. It retrieves soil properties, recent weather records, and historical phenological information and adds them to the reasoning context of the large language model (LLM). This design is intended to ground recommendations in location-specific data, support reproducibility, and remain compatible with a modular architecture that uses external tools.

References

- Barnes, E. M., Clarke, T. R., Richards, S. E., Colaizzi, P. D., Haberland, J., Kostrzewski, M., Waller, P., Choi, C., Riley, E., Thompson, T., Lascano, R. J., Li, H., & Moran, M. S. (2000). Coincident detection of crop water stress, nitrogen status and canopy density using ground-based multispectral data. *Proceedings of the Fifth International Conference on Precision Agriculture*.
- Cao, Y., Yin, T., Zhang, Y., Yang, X., Liu, B., Zhu, Y., Tang, L., Xiao, L., Cao, W., & Liu, L. (2024). Quantitative assessment of the effects of rising temperature on the grain protein of winter wheat in China and its adaptive strategies. *Computers and Electronics in Agriculture* 226, 109474. <https://doi.org/10.1016/j.compag.2024.109474>
- Clevers, J. G. P. W., & Gitelson, A. A. (2013). Remote estimation of crop and grass chlorophyll and nitrogen content using red-edge bands on Sentinel-2 and -3. *International Journal of Applied Earth Observation and Geoinformation* 23, 344-351. <https://doi.org/10.1016/j.jag.2012.10.008>
- Delegido, J., Verrelst, J., Alonso, L., & Moreno, J. (2011). Evaluation of Sentinel-2 red-edge bands for empirical estimation of green LAI and chlorophyll content. *Sensors* 11, 7063-7081. <https://doi.org/10.3390/s110707063>
- Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research: Oceans* 99, 10143-10162. <https://doi.org/10.1029/94JC00572>
- Evensen, G. (2003). The ensemble Kalman filter: theoretical formulation and practical implementation. *Ocean Dynamics* 53, 343-367. <https://doi.org/10.1007/s10236-003-0036-9>
- Fu, Z., Zhang, J., Jiang, J., Zhang, Z., Cao, Q., Tian, Y., Zhu, Y., Cao, W., & Liu, X. (2024). Using the time series nitrogen diagnosis curve for precise nitrogen management in wheat and rice. *Field Crops Research*, 307: 109259.
- Guo, C., Tang, Y., Lu, J., Zhu, Y., Cao, W., Cheng, T., Zhang, L., & Tian, Y. (2019). Predicting wheat productivity: Integrating time series of vegetation indices into crop modeling via sequential assimilation. *Agricultural and Forest Meteorology*, 272-273, 69-80.
- Huang, Y., Zhu, Y., Li, W., Cao, W., & Tian, Y. (2013). Assimilating remotely sensed information with the WheatGrow model based on the ensemble square root filter for improving regional wheat yield forecasts. *Plant Production Science* 16, 352-364. <https://doi.org/10.1626/pps.16.352>
- Huete, A., Didan, K., Miura, T., Rodriguez, E. P., Gao, X., & Ferreira, L. G. (2002). Overview of the radiometric and biophysical performance of the MODIS vegetation indices. *Remote Sensing of Environment* 83, 195-213. [https://doi.org/10.1016/S0034-4257\(02\)00096-2](https://doi.org/10.1016/S0034-4257(02)00096-2)

- Li, W., Li, D., Liu, S., Baret, F., Ma, Z., He, C., Warner, T. A., Guo, C., Cheng, T., Zhu, Y., Cao, W., & Yao, X. (2023a). RSARE: A physically-based vegetation index for estimating wheat green LAI to mitigate the impact of leaf chlorophyll content and residue-soil background. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2003, 138-152
- Li, W., Wu, W., Yu, M., Tao, H., Yao, X., Cheng, T., Zhu, Y., Cao, W., & Tian, Y. (2023b). Monitoring rice grain protein accumulation dynamics based on UAV multispectral data. *Field Crops Research*, 294, 108858
- Liu, S., Wang, L., & Zhang, J. (2024). The dataset of main grain land changes in China over 1985-2020. *Scientific Data* 11, 1430. <https://doi.org/10.1038/s41597-024-04292-y>
- Lu, C., & Tian, H. (2017). Global nitrogen and phosphorus fertilizer use for agriculture production in the past half century: shifted hot spots and nutrient imbalance. *Earth System Science Data* 9, 181-192. <https://doi.org/10.5194/essd-9-181-2017>
- Lu, J., Li, J., Fu, H. K., Zou, W. L., Kang, J. R., Yu, H. W., & Lin, X. L. (2025). Estimation of rice yield using multi-source remote sensing data combined with crop growth model and deep learning algorithm. *Agricultural and Forest Meteorology*, 370
- Lu, N., Wang, W., Zhang, Q., Li, D., Yao, X., Tian, Y., Zhu, Y., Cao, W., Baret, F., Liu, S., & Cheng, T. (2019a). Estimation of nitrogen nutrition status in winter wheat from unmanned aerial vehicle based multi-angular multispectral imagery. *Frontiers in Plant Science*, 10, 1601
- Lu, N., Zhou, J., Han, Z., Li, D., Cao, Q., Yao, X., Tian, Y., Zhu, Y., Cao, W., & Cheng, T. (2019b). Improved estimation of aboveground biomass in wheat from RGB imagery and point cloud data acquired with a low-cost unmanned aerial vehicle system. *Plant Methods*, 15, 17
- Luo, Y., Zhang, Z., Chen, Y., Li, Z., & Tao, F. (2020). ChinaCropPhen1km: a high-resolution crop phenological dataset for three staple crops in China during 2000-2015 based on leaf area index (LAI) products. *Earth System Science Data* 12, 197-214. <https://doi.org/10.5194/essd-12-197-2020>
- McFeeters, S. K. (1996). The use of the Normalized Difference Water Index (NDWI) in the delineation of open water features. *International Journal of Remote Sensing* 17, 1425-1432. <https://doi.org/10.1080/01431169608948714>
- Poggio, L., de Sousa, L. M., Batjes, N. H., Heuvelink, G. B. M., Kempen, B., Ribeiro, E., & Rossiter, D. (2021). SoilGrids 2.0: Producing soil information for the globe with quantified spatial uncertainty. *SOIL* 7, 217-240. <https://doi.org/10.5194/soil-7-217-2021>
- Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. *IEEE CVPR*, 779-788. <https://doi.org/10.1109/CVPR.2016.91>
- Rouse, J. W. Jr, Haas, R. H., Schell, J. A., & Deering, D. W. (1974). Monitoring vegetation systems in the Great Plains with ERTS. Third ERTS-1 Symposium, NASA SP-351, 309-317. <https://ntrs.nasa.gov/citations/19740022614>
- Serrano Reyes, J., Jiménez, J. U., Quirós-McIntire, E. I., Sanchez-Galan, J. E., & Fábrega, J. R. (2023). Comparing two methods of leaf area index estimation for rice (*Oryza sativa* L.) using in-field spectroradiometric measurements and multispectral satellite images. *AgriEngineering* 5, 965-981. <https://doi.org/10.3390/agriengineering5020060>

- Simons, G. W. H., Koster, R., & Droogers, P. (2020). HiHydroSoil v2.0: A high resolution soil map of global hydraulic properties. *FutureWater Report* 213. <https://www.futurewater.nl/projects/hihydrosoil-v2-0-global-maps-of-soil-hydraulic-properties-at-250m-resolution/>
- Smith, A., Lott, N., & Vose, R. (2011). The Integrated Surface Database: Recent developments and partnerships. *Bulletin of the American Meteorological Society*, 92, 704-708. <https://doi.org/10.1175/2011BAMS3015.1>
- Tang, L., Zhu, Y., Hannaway, D., Meng, Y., Liu, L., Chen, L., & Cao, W. (2009). RiceGrow: A rice growth and productivity model. *NJAS - Wageningen Journal of Life Sciences* 57, 83-92. <https://doi.org/10.1016/j.njas.2009.12.003>
- Tóth, B., Weynants, M., Nemes, A., Makó, A., Bilas, G., & Tóth, G. (2015). New generation of hydraulic pedotransfer functions for Europe. *European Journal of Soil Science*, 66, 226-238. <https://doi.org/10.1111/ejss.12192>
- Wang, S., Liu, Y., Asseng, S., Harrison, M. T., Tang, L., Liu, B., Liu, K., Luo, Z., Wang, E., Chang, J., Qiu, X., Liu, L., Zhang, X., Cao, W., Zhu, Y., & Xiao, L. (2025). Rice yield stability and its determinants across different rice-cropping systems in China. *Agricultural and Forest Meteorology*, 364, 110452. <https://doi.org/10.1016/j.agrformet.2025.110452>
- Wang, X., Yang, W., Lv, Q., Huang, C., Liang, X., Chen, G., Xiong, L., Duan, L. (2022). Field rice panicle detection and counting based on deep learning. *Frontiers in Plant Science* 13, 966495 <https://doi.org/10.3389/fpls.2022.966495>
- Wang, Y., Gu, Y., Tang, J., Guo, B., Warner, T. A., Guo, C., Zheng, H., Hosoi, F., Cheng, T., Zhu, Y., Cao, W., & Yao, X. (2024). Quantify wheat canopy leaf angle distribution using terrestrial laser scanning data. *IEEE Transactions on Geoscience and Remote Sensing*, 62
- Wu, Y., Xiong, Y., Yu, W., Gu, Y., Zheng, H., Jiang, C., Yao, X., Zhu, Y., Cao, W., & Cheng, T. (2026). In-season estimation of aboveground biomass and yield in winter wheat with a UAV-based LUE model and machine learning. *Plant Phenomics*, 8: 100162.
- Yamazaki, D., Ikeshima, D., Tawatari, R., Yamaguchi, T., O'Loughlin, F., Neal, J. C., Sampson, C. C., Kanae, S., & Bates, P. D. (2017). A high-accuracy map of global terrain elevations. *Geophysical Research Letters*, 44, 5844-5853. <https://doi.org/10.1002/2017GL072874>
- Yan, M., Cao, W., Luo, W., & Jiang, H. (2000). A mechanistic model of phasic and phenological development of wheat. I. Assumption and description of the model. *Chinese Journal of Applied Ecology* 11, 355-359.
- Yu, M., He, J., Li, W., Zheng, H., Wang, X., Yao, X., Cheng, T., Zhang, X., Zhu, Y., Cao, W., & Tian, Y. (2024). Estimation of rice leaf area index utilizing a Kalman filter fusion methodology based on multi-spectral data obtained from Unmanned Aerial Vehicles (UAVs). *Remote Sensing*, 16
- Zang, Y., He, L., Yuan, M., & Li, P. (2022). Detection method of wheat spike improved YOLOv5s based on the attention mechanism. *Frontiers in Plant Science*, 13, 993244.
- Zhang, J., Pan, Y., Wang, W., Shi, Z., Zhang, Z., Fu, Z., Cao, Q., Tian, Y., Zhu, Y., Liu, X., & Cao, W. (2024). Potential benefits of variable rate nitrogen topdressing strategy coupled with zoning technique: a case study in a town-scale rice production system. *European Journal of Agronomy*, 155, 127132.
- Zheng, H., Cheng, T., Li, D., Zhou, X., Yao, X., Tian, Y., Cao, W., & Zhu, Y. (2018a). Evaluation of

RGB, color-infrared and multispectral images acquired from unmanned aerial systems for the estimation of nitrogen accumulation in rice. *Remote Sensing*, 10, 824

Zheng, H., Cheng, T., Zhou, M., Li, D., Yao, X., Tian, Y., Cao, W., & Zhu, Y. (2019). Improved estimation of rice aboveground biomass combining textural and spectral analysis of UAV imagery. *Precision Agriculture*, 20, 611-629

Zheng, H., Ji, W., Wang, W., Lu, J., Li, D., Guo, C. , Yao, X., Tian, Y. , Cao, W. , Zhu, Y., & Cheng, T. (2022). Transferability of models for predicting rice grain yield from unmanned aerial vehicle (UAV) multispectral imagery across years, cultivars and sensors. *Drones*, 6, 423

Zheng, H., Li, W., Jiang, J., Liu, Y., Cheng, T., Tian, Y., Zhu, Y., Cao, W., Zhang, Y., & Yao, X. (2018b). A comparative assessment of different modeling algorithms for estimating leaf nitrogen content in winter wheat using multispectral images from an unmanned aerial vehicle. *Remote Sensing*, 10, 2026

Zheng, H., Ma, J., Zhou, M., Li, D., Yao, X., Cao, W., Zhu, Y., & Cheng, T. (2020). Enhancing the nitrogen signals of rice canopies across critical growth stages through the integration of textural and spectral information from Unmanned Aerial Vehicle (UAV) multispectral imagery. *Remote Sensing*, 12, 957

Zhu, Y., Tang, L., Liu, L. L., Liu, B., Zhang, X. H., Qiu, X. L., Tian, Y. C., & Cao, W. X. (2020). Research progress on the crop growth model CropGrow. *Scientia Agricultura Sinica* 53, 3235-3256. <https://doi.org/10.3864/j.issn.0578-1752.2020.16.004>