ISSN: 2754-6659

Journal of Artificial Intelligence & Cloud Computing

SCIENTIFIC Research and Community

Review Article

Open Access

From Sensing to Safe Action: Digital Twin–Reinforcement Learning for Closed-Loop Oil Palm Plantation–Transport–Mill Optimization

Loso Judijanto

IPOSS Jakarta, Indonesia

ABSTRACT

Oil palm production is a coupled biological, logistical, and industrial system, yet most digital interventions optimize isolated functions. This qualitative literature review examines whether dynamic digital twins integrated with reinforcement learning (RL) can support coordinated, safety-constrained decisions across plantations, fresh fruit bunch (FFB) transport, and palm oil mills. It synthesizes recent peer-reviewed evidence on agricultural digital twins, interoperable architectures, sequential decision-making, safe and offline RL, oil palm sensing, harvest planning, mill simulation, and environmental performance. The synthesis distinguishes demonstrated component-level capabilities from an unvalidated system-level proposition. It develops a hierarchical Closed-Loop Palm Digital Twin architecture comprising palm/block, harvest, transport, mill, and environmental twins connected by shared events, state estimation, scenario simulation, and supervisory policy learning. A governance cycle—Sense → Synchronize → Simulate → Learn → Recommend → Authorize → Act → Verify—retains human authorization for consequential actions. The review identifies major design requirements: multiscale temporal synchronization, realistic fruit-quality decay and process constraints, uncertainty-aware counterfactuals, protected data sharing, and staged sim-to-real testing. It argues that the central value proposition is not autonomous control per se but demonstrably improved cross-stage coordination under agronomic, labor, safety, and environmental limits. Research priorities include representative benchmark environments, prospective shadow-mode trials, causal evaluation, and audited deployments. No empirical performance gains for an integrated oil palm digital twin–RL system are claimed.

***Corresponding author**

Loso Judijanto, IPOSS Jakarta, Indonesia

Received: October 30, 2026; **Accepted:** November 03, 2026; **Published:** November 13, 2026

Keywords: Digital Twin, Reinforcement Learning, Oil Palm, Plantation Management, Fresh Fruit Bunch, Harvest Scheduling, Agricultural Logistics, Palm Oil mill, Safe AI, Closed-Loop Optimization

JEL Classification: C61; L66; O33; Q16; Q56

Introduction

Background: A Coupled Biological–Industrial Production System

Oil palm production cannot be adequately understood as a collection of independent optimization problems. Agronomic decisions shape future bunch availability; ripeness and labor availability influence the precise timing of harvest; roads, vehicles, and mill reception capacity determine delivery times; and a mill must process a changing mix of incoming fruit within equipment, steam, quality, and environmental limits. A recommendation that maximizes field-level tonnage may therefore reduce overall economic value when it creates transport congestion or forces low-quality processing. Existing oil palm supply-chain optimization models recognize material and financial interdependence, but their static or episodically updated formulations do not by themselves constitute a continuously synchronized representation of operating assets (1,2).

The industry has simultaneously gained access to digital observations that were previously sparse or expensive. Agronomic machine-learning studies combine soil, climatic, plantation-age, and management variables to forecast yield, while visual and geospatial systems increasingly support assessment of fruit maturity and plantation conditions. Open image datasets facilitate FFB grading research, and hyperspectral approaches have demonstrated the potential to distinguish certain disease-associated signals. These contributions build components of digital observability, but a prediction, image classifier, or geographic dashboard does not establish how an action affects the next state of plantation and mill (3-6).

Urgency: Why Integration Matters Now

Digital twin research conceptualizes a persistent digital counterpart that tracks the evolving state of a physical entity and can support forecasts and interventions. In agriculture, digital twins are still heterogeneous in modeling depth, data connectivity, and implementation maturity; several reviews describe a gap between monitoring-oriented prototypes and closed-loop decision systems. Recent work has brought crop models, real-time observations, interoperable middleware, and RL decision agents together in agricultural pilot studies, offering a plausible architectural precedent rather than direct evidence from oil palm estates and mills (7-10).

The urgency is operational as well as technological. Field managers face uncertain rainfall, disease, production peaks, equipment breakdowns, road accessibility, labor attendance, and fluctuating processing constraints. Mill managers encounter a different clock speed, where steam pressure or sterilization settings may matter over minutes or hours rather than weeks or months.

Biogas utilization and process heat carry material environmental implications, suggesting that a coordinated objective must extend beyond tonnage and extraction throughput. Mill simulations and resource assessments show that integrated process and environmental modeling is feasible at component level, but also demonstrate the need for calibrated plant-specific models rather than general claims of universal performance (12-14).

Gap, Aims, and Scope

The central research gap is the absence of a sufficiently developed scholarly synthesis connecting dynamically synchronized crop and block representations, freshness-sensitive transport, mill process behavior, and safe policy learning in one oil palm decision environment. Digital twin studies often emphasize modeling and interoperability; oil palm studies commonly emphasize ripeness detection, yield prediction, harvesting optimization, or mill process simulation; RL literature develops techniques for adaptive actions but usually in other agricultural crops or industrial settings. These three bodies of scholarship have informative intersections, yet their combination should be treated as a proposed research program, not as a proven oil palm implementation (15-17).

This article asks how an interconnected digital twin–RL platform could be designed, validated, and governed to improve plantation–transport–mill coordination under biological uncertainty and binding operational constraints. Its objectives are to clarify terminology; derive a cross-stage architecture; specify state, action, reward, and constraint design; evaluate single- and multi-agent choices; outline interoperability and safety mechanisms; and propose a credible deployment and evidence agenda. The contribution is conceptual and critical: it reconstructs a design space through qualitative thematic synthesis without conducting a new field experiment, computing a pooled effect estimate, or asserting that the proposed system already works (18,19).

Three questions organize the review. First, which digital-twin components and interoperability mechanisms are necessary to represent biological, logistical, industrial, and environmental coupling (RQ1)? Second, under which observability, model-fidelity, safety, and organizational conditions could reinforcement learning improve cross-stage decisions relative to strong non-learning baselines (RQ2)? Third, what sequence of evidence, approval, and field evaluation is required before any increase in operational autonomy can be justified (RQ3)? These questions make the architectural proposal testable while explicitly separating technological possibility from demonstrated plantation–mill effectiveness.

Literature Review: Conceptual and Theoretical Foundations

Digital Twins Beyond Digital Models and Dashboards

Digital twins are frequently discussed as virtual replicas, but that shorthand obscures operational requirements. A model may describe a process without consuming current observations; a digital shadow may automatically receive data without returning executable recommendations; and a decision-capable twin must include valid state synchronization, model updating, accountable interventions, and feedback after action. Design patterns and model-based systems engineering help specify components, communication pathways, fidelity requirements, and appropriate degrees of autonomy. Reviews across manufacturing also show that a general label does not guarantee real-time connectivity, predictive validity, or validated control (7-10,18).

A plantation digital twin additionally differs from the machinery-oriented systems that originated in industrial engineering. A palm is a persistent biological asset with age-dependent growth, uncertain physiological responses, spatially variable soils, slow-changing productivity, and episodic harvest events. An industrial mill, by contrast, contains faster physical processes and specified engineering operating limits. A defensible architecture therefore requires multiple interconnected models with different temporal resolutions, not a single monolithic simulation forced onto one artificial update rate. Work on functional structural plant modeling and agricultural digital-twin reviews points toward modular model hierarchies and careful alignment between measurement and state representation (19,20).

Synchronization, Assimilation, and Interoperability

Synchronization is a scientific issue rather than a routine data-engineering step. Satellite observations may be intermittent and cloud-obstructed; drone surveys cover selected areas; weighing events arrive continuously; field harvest records are often entered retrospectively; and laboratory assays may appear after the affected fruit has already been processed. A twin must time-stamp events, distinguish observed from estimated variables, retain sensor uncertainty, and reconcile inconsistent identifiers for estate, block, palm, vehicle, batch, and mill lot. A model silently refreshed with mismatched timestamps can appear responsive while providing a misleading state estimate (8,9).

Interoperability should accordingly be defined at four levels: syntactic exchange of data, semantic agreement about operational entities, temporal consistency of events, and organizational permission to use data. The interoperable agricultural twin developed by Kallenberg and colleagues illustrates a modular approach linking crop and disease models to RL agents and an integration layer. Meanwhile, reviews of digital twins in food supply chains emphasize the importance of distributed information and provenance across organizations. For oil palm, these lessons imply canonical identity management, versioned agronomic and industrial ontologies, secure interfaces to mill control systems, and data lineage that follows fruit from harvesting block to processed batch (9,21-23).

Sequential Decision-Making and Reinforcement Learning

Reinforcement learning addresses sequential choices in which the consequences of an action unfold through later states. A policy maps an observed or inferred system state to a recommendation, and its quality depends on long-run returns rather than only immediate performance. This framing is attractive for the plantation–mill system because a vehicle dispatch choice affects a queue, fruit freshness, subsequent harvesting capacity, and downstream mill utilization. Model-based RL is especially relevant when the digital twin provides a simulated transition process, but the quality of learned behavior remains dependent on state observability, realism of the transition dynamics, and the incentives represented in the reward (7,24).

A single agent may be appropriate for bounded tasks, such as choosing the next dispatch batch under a fixed mill timetable. However, harvest supervisors, logistics coordinators, and mill operators possess partially distinct information and authority. Multi-agent RL can represent these roles but introduces non-stationarity: one agent's environment changes as the others learn. Semi-centralized agricultural irrigation research provides a useful precedent for a coordinator that makes global timing decisions while local agents select context-sensitive action magnitudes. Such findings motivate, but do not validate, hierarchical or cooperative policies for oil palm supply chains (25,26).

Safety, Uncertainty, Offline Learning, and Explanation

Unconstrained reward maximization is not equivalent to acceptable plantation management. A learned policy could increase simulated throughput by recommending unsafe pressures, excessive loads, implausible harvest rates, illegal environmental releases, or schedules inconsistent with rest periods. Safe RL research therefore distinguishes optimization objectives from constraints that must not be violated; constraints may be episode-level, state-wise, or enforced by an external safety mechanism. An operational design should privilege deterministic rule-based vetoes, audited approval pathways, and machine-controller interlocks over trust in a policy's estimated average reward (17,27).

Offline RL can exploit historical operational records before live intervention, yet those records usually reflect particular managerial policies, contingencies, and selection biases. Actions rarely observed in training data can receive unjustified value estimates, making deployment outside the behavior-data distribution risky. Sim-to-real studies also show that simulated transitions systematically omit real-world disturbances. A responsible architecture needs uncertainty estimation, domain randomization, external validation, and an explicit option to abstain when the twin diverges from reality (29-31).

Explanations are a governance requirement, not an ornamental user-interface element. Managers require a reason why a suggested harvest change improves expected value, which constraints bind, what inputs drive the result, and how sensitive the recommendation is to uncertain weather or mill availability. Explainable RL surveys distinguish feature attributions from causal or counterfactual accounts of sequential behavior, while interpretable policy design offers another route to accountability. For decisions that affect labor, chemical use, production safety, or environmental performance, transparent evidence and the ability to override a recommendation are integral to acceptance (32-35).

Oil Palm-Specific Technological Foundations

Research on FFB ripeness estimation offers a concrete starting point for sensing. Reviews document optical and non-destructive techniques, while supervised classifiers and datasets show how grading can be partly formalized across conditions. These tools do not eliminate variation in lighting, genotype, fruit position, or human ground truth; neither do they automatically solve the scheduling problem of when to inspect or harvest a block. The twin should therefore represent ripeness as an uncertain distribution and convert it into timing-sensitive decision inputs only after locally appropriate calibration (36-39).

Disease monitoring provides a different lesson: remotely sensed information can detect patterns consistent with pathological stress, but changes in spectral features are not themselves an infallible disease diagnosis. Hyperspectral and machine-learning studies of basal stem rot illustrate potential for early alerts and targeted inspection, while also underscoring the need for robust biological ground truth and site-specific transfer testing. For a plantation twin, the right action may be to trigger a field verification workflow rather than authorize chemical or silvicultural interventions automatically (6,40,41).

Newer research on adaptive site-specific crop protection distinguishes improved machine perception from truly feedback-based intervention policies. A 2026 review finds that coupling digital twins with reinforcement learning is a promising direction for intelligent spraying, while field-scale human–machine validation remains an important unresolved challenge. This distinction is particularly relevant to long-lived oil palms, where a false disease signal can precipitate costly or irreversible intervention, and it supports escalation from alerts to inspected evidence before any control decision (42).

Method: Qualitative Literature Review

Review Design and Rationale

This article adopts a qualitative literature review (QLR) approach. Its purpose is conceptual interpretation, critical comparison, and thematic integration of heterogeneous peer-reviewed evidence rather than exhaustive identification and statistical aggregation of every eligible publication. Unlike a systematic literature review, it does not present a preregistered protocol, PRISMA flowchart, pooled effect size, or claim of reproducible exhaustive coverage. Those methodological distinctions matter because digitally enabled agriculture brings together conceptual architectures, laboratory experiments, simulations, field trials, and reviews that cannot validly be ranked by a single numerical outcome measure (15,43).

Literature Identification, Appraisal, and Interpretation

The interpretive corpus emphasizes articles published from 2020 through October 2026 in established, peer-reviewed journals and covers five overlapping domains: agricultural and industrial digital twins; RL, safety, interpretability, and transfer; oil palm agronomic perception and harvest operations; supply-chain and mill process optimization; and sustainability-oriented production systems. Journal publisher pages, DOI records, research repositories, and scholarly research portals were used to verify article identities and publication metadata where accessible. Inclusion in the discussion depended on conceptual or technical relevance, not a claimed numerical screening threshold. The review therefore offers an argued thematic synthesis rather than the reproducibility guarantees associated with a systematic review (8,22).

Evidence was interpreted along three cross-cutting axes. The first was whether an article demonstrated a measured result, proposed a method, or surveyed a field. The second was operational scope: palm or block, harvest, transport, mill, environmental resource, or multi-stage coordination. The third was deployment readiness: theoretical, simulation-only, pilot, field-tested, or operationally validated. Contradictions, especially the gap between laboratory classification and generalizable field performance, were retained rather than averaged away. These distinctions allow the review to infer design principles while avoiding the unsupported leap from a successful crop-control demonstration to proof of palm-sector economic benefit (9,16).

Synthesis Procedure and Limitations

The first synthesis passes clarified definitions of state, transition, intervention, feedback, and safety. A second pass grouped the literature into sensing, dynamic replication, sequential decision-making, vertical integration, environmental accounting, and governance. A third pass mapped each theme onto the proposed hierarchical twin and tested plausible cross-stage information flows against documented oil palm problems, such as batch maturity, truck constraints, mill bottlenecks, and steam demand. Finally, the review constructed evaluation criteria that could falsify an overly ambitious implementation claim. Because no new raw operational data or intervention experiments were analyzed, the resulting architecture is a structured hypothesis, and its cost-effectiveness and environmental impacts require prospective validation (18,45).

To reduce overinterpretation, evidence statements were classified as demonstrated field findings, simulation or laboratory findings, synthesis across published work, or design propositions developed by the present authors. These classes are not interchangeable: an algorithm that predicts a crop variable does not establish the causal value of acting on that prediction, and a simulated operational reward does not demonstrate an improvement in economic or environmental outcomes. Themes were retained when they directly informed a cross-stage decision requirement, a falsifiable safety condition, or a deployment gate, rather than simply sharing the vocabulary of digital twins or artificial intelligence.

Results: Thematic Findings from the Literature

Theme I: Observability is Improving Faster than Validated Control

The evidence suggests a striking asymmetry between improvements in observation and the maturity of integrated action. Oil palm yield predictors use broad agronomic datasets; FFB image datasets provide labels for model development; and outdoor measurements reveal realistic variations in fruit appearance. These contributions support better representation of crop condition and harvest readiness, but their outputs remain estimates with model error, sample bias, and shifting field conditions. An operational twin must translate each estimate into a time-stamped posterior state, complete with provenance and quality score, rather than treating model output as an unquestionable fact (3-5).

The same separation appears across the wider agricultural literature. Reviews identify monitoring, prediction, visualization, and limited closed-loop control as distinct maturity levels. A working greenhouse or crop-model twin can demonstrate synchronization in an environment where measurements and actuators are comparatively accessible, but open-field tree crops pose harder observational and biological challenges. Real-time communication between physical and digital components is necessary yet insufficient; the more demanding test is whether decision recommendations consistently improve outcomes when the underlying model encounters weather anomalies, disrupted roads, and production cycles outside its training period (20-22,45).

A recent review of agricultural machinery and robotics reinforces this implementation gap, identifying substantial variation in architecture, evidence quality, and the extent of rigorous numerical field validation. These findings should not be interpreted as palm-specific performance estimates; rather, they caution that apparent benefits from an agricultural twin are conditional on the physical assets, control objectives, and validation setting. For this review, the appropriate evidentiary threshold is successful prediction of held-out operating conditions followed by prospectively measured decision improvements (46).

Theme II: The Unit of Coordination should be the Fruit Flow, Not the Dashboard

A closed-loop palm system should treat the movement of FFB from block to loading point, vehicle, weighbridge, reception yard, and processing line as the organizing event chain. Blocks supply expected harvest volumes and ripeness distributions; harvest records create batches with origins and collection timestamps; vehicle dispatch changes delivery timing; weighing and grading update the mill's view of the arriving material. The combined state determines whether an action is desirable. Historical optimization models on supply-chain planning and harvest collection provide mathematical foundations but generally require an additional live-state and feedback layer to function as a persistent twin (1,2).

The relevant economic quantity is not simply the number of tons harvested or processed, but expected system value after considering oil yield, grade, losses, transport expense, energy, maintenance, environmental load, and delayed operations. The quality effects of ripeness, mechanical damage, and time between harvest and processing should therefore be treated as measurable, empirically calibrated process relationships. This is especially important because otherwise an RL policy may discover a reward artifact: for example, artificially increasing daily reception may look profitable in simulation even when it increases fruit deterioration or compromises the stability of steam and clarification operations. Food-supply-chain twin frameworks and food-processing digital twin methodology provide a helpful systems perspective on linking physical flows with traceable, multi-dimensional performance indicators (22,47,48).

Theme III: A Hierarchical Twin is Preferable to One Undifferentiated Model

The synthesis supports five coupled but separable components. Palm/Block Twins represent palm age, health, soil, moisture, forecast production, and intervention history; Harvest Twins represent ripeness, crew and machinery availability, and planned harvest rounds; Transport Twins represent locations, capacities, road conditions, dispatch, and queues; Mill Twins represent process units, energy, operating windows, and quality states; Environmental Twins estimate water, fuel, effluent, methane, and other boundary indicators. The approach mirrors architectural separation practiced in advanced digital-twin engineering and limits the computational and organizational costs of frequent whole-system updates (7-10,49).

The five twins should exchange standardized events through a common information contract while retaining task-specific model fidelity. A Palm/Block Twin might update growth forecasts weekly yet post harvest availability daily; the Transport Twin may require location updates within minutes; and selected mill process models may require control-system sampling at much higher frequencies. Cross-twin coupling should therefore be event-triggered or multirate, with explicit lag and uncertainty tracking. Local process controllers continue to manage fast safety-critical loops; the overarching twin optimizes schedules and bounded set points at timescales compatible with human authorization. Digital-twin design methodologies support this separation of levels and responsibilities (50,51). The connections among the shared data contract, five operating twins, model validation, safe policy learning, and authorized execution are summarized in Figure 1.

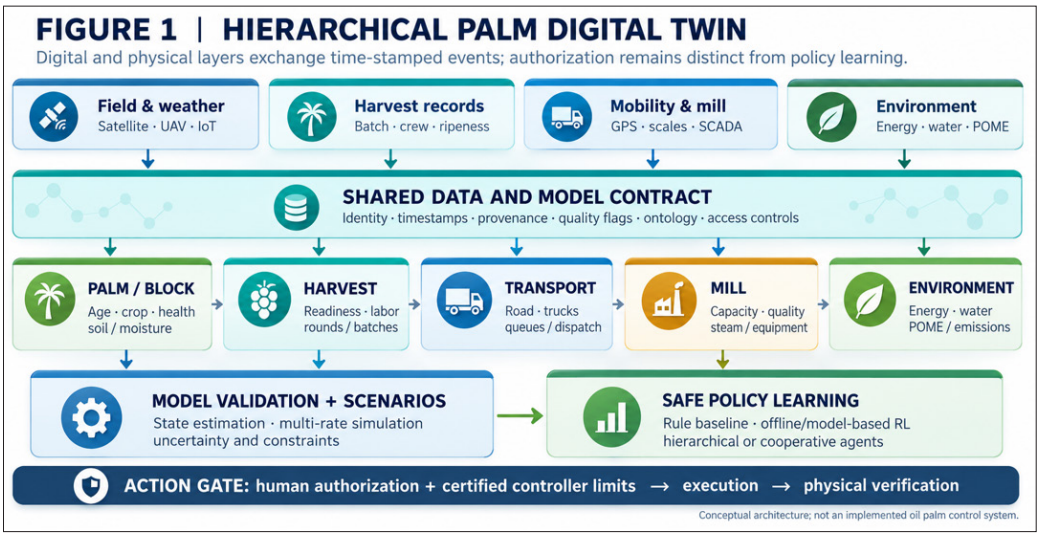


Figure 1: Proposed hierarchical Closed-Loop Palm Digital Twin architecture (conceptual)

Figure 1 illustrates the proposed architecture as a coordinated system of systems rather than a monolithic database. A shared identity and event layer connects the five twins, while uncertainty estimation and model validation continuously assess fidelity. The architecture deliberately separates an advisory RL layer from the supervisory authorization and operational execution layers. Its components are design propositions derived from transferable literature, not representations of an existing industrial installation (9,21).

Theme IV: Rl Applications Differ Fundamentally by Timescale and Risk

At the agronomic level, RL could explore fertilizer allocation or conditional water-management strategies using disease, soil, rainfall, and nutrient information. Agricultural studies on irrigation and nitrogen decisions demonstrate that simulation-driven policies can be compared with conventional rules, but the evidence is context-dependent. For example, irrigation RL did not uniformly outperform optimized heuristics across rainfall settings, indicating that policy complexity should be justified by out-of-sample improvements and resource scarcity. Nitrogen-focused constrained RL is a more direct methodological example of balancing action constraints and nutrient efficiency, although transferring its crop assumptions to perennial palms requires independent calibration (10,16,52). A digital-twin-oriented irrigation study combining soil-water and crop-growth models with an RL agent offers a recent mechanistic example, but its sunflower and maize results cannot be generalized directly to perennial oil palm plantations (53).

Harvesting agents require different actions and guardrails. They could prioritize eligible blocks, recommend next-day crew assignments, or identify inspections when predicted maturity is uncertain. However, they cannot assume that images perfectly measure hidden fruit, labor is infinitely divisible, or all blocks are always accessible. An agent must also respect minimum harvest quality specifications and operationally feasible round lengths. Harvest status must be reconciled with field confirmations, while visual scores should be compared with established grading and inspection practices. Distinct sensor modalities and algorithm choices make local benchmarking indispensable (5). In a related but different crop system, a 2025 greenhouse-tomato study combined a dynamic twin and RL for coordinated harvesting decisions with physical testing, providing evidence that virtual-to-physical integration is possible without establishing transferable oil palm harvest gains (54).

A transport agent could recommend assignments of trucks to collection points, release times, sequencing through weighbridges, and diversion around blocked roads. This is a particularly promising early use case because transport actions are relatively frequent, observable, and reversible, and their effect on queues can be estimated from operational timestamps. Still, decisions must respect road safety, axle limits, driver hours, fuel constraints, contractual ownership of FFB, and mill reception windows. A useful baseline is a properly tuned dispatch heuristic or mixed-integer scheduling model, not an artificially weak first-in-first-out (1,2).

Mill optimization poses more demanding safety and process-engineering requirements. A high-level RL agent might choose bounded throughput targets, batch sequencing, heat and power planning, or maintenance-aware operating schedules, while programmable logic controllers and certified supervisory systems retain direct authority over critical actuators. Existing mill simulation studies show how steam, sterilization, throughput, and material balances can be analyzed; process-control RL reviews likewise warn that model mismatch and constraint handling complicate operational use. Consequently, an RL policy should not be allowed to directly bypass interlocks or command pressure, temperature, and equipment configurations outside validated envelopes (12,14,17,45).

Theme V: Environmental Outcomes Require Coupled Accounting

A narrow economic reward may shift damage from one process to another: congestion reductions could increase empty truck mileage, higher throughput may increase heat demand, and altered sterilization operations may change waste treatment conditions. Therefore, the Environmental Twin should not be relegated to a dashboard. It must calculate operational water and energy demand, effluent volumes, emissions-relevant conditions, and scenario-dependent resource balances with clearly defined system boundaries. Studies of mill biogas and integrated flowsheets show why environmental credits and costs must be recorded alongside production gains rather than assumed to follow automatically from digitalization (14).

The correct accounting unit should follow the decision horizon. A daily scheduling policy may track energy per ton of FFB processed and delayed fruit per batch, while seasonal agronomic planning

should include water or nutrient burden per hectare and longer-run productivity. Attribution requires attention: a decline in average emissions intensity may reflect higher utilization, changed fuel composition, or more favorable weather rather than improved decision quality. Food and agri-food digital twin scholarship supports multi-criteria visibility and scenario analysis, but the causal measurement of sustainability outcomes demands plant-specific inventories and reliable baselines Tzac (23,47).

Theme VI: Governance must be Part of the Control Loop

A technically plausible RL agent may still be institutionally unacceptable. Proposed actions affect agronomists, harvesting teams, transport contractors, mill managers, environmental officers, and sometimes third-party suppliers who hold distinct incentives and legal responsibilities. Effective governance requires role-specific authorization rights, auditable model versions, documented fallbacks, data ownership rules, and a process for reporting recommendation errors. Explanations should identify the material observations and binding constraints, not merely present a confidence score or an unexplained priority ranking (34,35).

A fully automatic feedback loop may be suitable for narrowly bounded, low-consequence actions following validation, but consequential decisions should retain an explicit human authorization gate. That gate has three functions: technical oversight of uncertainty and out-of-distribution conditions, institutional accountability for labor and environmental decisions, and organizational learning from accepted and rejected recommendations. The distinction between recommendation and actuation is therefore a substantive control-system design choice. Safe RL and explainability literature provide a rationale for such separations, even though no universally accepted palm-sector authorization threshold has yet been established (32).

Discussion and Analysis

Proposed Closed-Loop Palm Digital Twin

The principal synthesis is a proposed operating architecture: Sense → Synchronize → Simulate → Learn → Recommend → Authorize → Act → Verify. Sensing gathers field and mill observations; synchronization reconstructs a credible, time-aligned operational state; simulation evaluates feasible interventions and uncertainty; learning adjusts policies using historical and simulated experience; recommendation ranks bounded alternatives; authorization assigns explicit human accountability; action occurs through existing organizational and industrial systems; and verification observes both intended effects and unintended consequences. The sequence should be cyclical rather than a one-off analytics exercise, with verification failures triggering recalibration or temporary suspension (7,9,21).

Figure 2 depicts this sequence as a feedback-control cycle and highlights the human decision gate that separates learned recommendations from physical actuation.

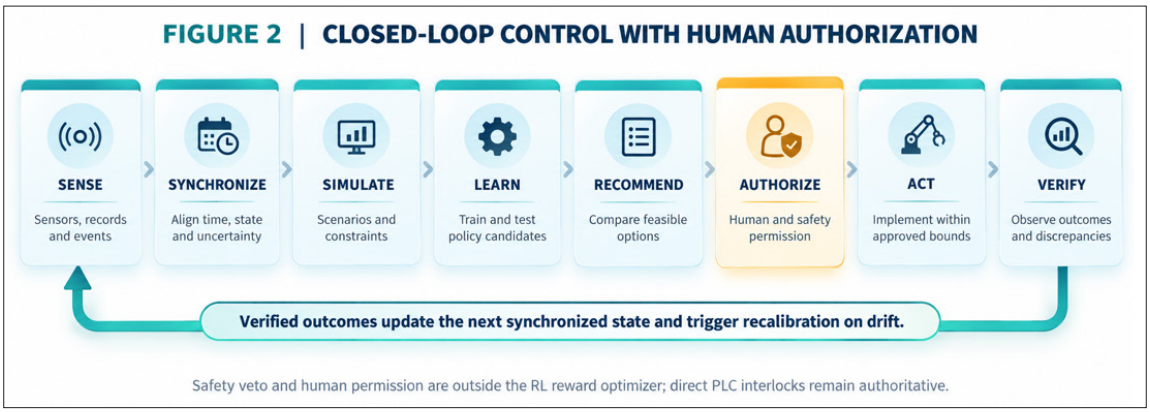


Figure 2: Sense–Synchronize–Simulate–Learn–Recommend–Authorize–Act–Verify feedback cycle (conceptual)

Figure 2 makes two conditions visible. First, recommendations pass through an independent safety and human-authorization layer before operational action; training does not confer permission to execute. Second, actual physical outcomes return to the twin as evidence for state correction and performance evaluation. If a mill receives fewer tons than forecast or the dispatch policy creates unexpected queues, the digital twin must update its state and the organization must evaluate the disagreement. Without the final Verify step, the platform risks becoming a sophisticated but unfalsifiable simulator (18).

State, Action, Reward, and Constraint Specification

The system should be defined as a partially observable, constrained, multiscale decision process. A shared state might include block age and expected FFB availability, maturity distributions, soil-water indicators, disease inspection flags, harvest capacity, vehicle positions and loads, road availability, mill queues, tank storage, steam balance, line availability, and environmental indicators. Not every variable will be directly observed; estimates must carry confidence or uncertainty intervals. Actions should be explicitly scoped: selecting inspection targets, proposing harvest-round changes, assigning transport capacity, recommending mill intake scheduling, or changing bounded supervisory targets. Operational constraints can then be checked before policy execution (24,29).

Reward design must avoid collapsing heterogeneous outcomes into an opaque metric. One transparent objective can combine expected recoverable oil value, quality compliance, labor and transport cost, energy and water use, and a penalty for waiting or spoilage. However, safety-critical limits should not be traded away merely because the reward gives them a sufficiently small penalty. Hard constraints should represent maximum machine capacities, certified process limits, legal operating conditions, environmental discharge rules, mandated worker protections, and physically feasible agronomic intervention bounds. Sensitivity analysis should examine whether the preferred policy changes under alternative weights, market prices, and plausible process uncertainty (10). Comparative crop-rotation research also shows that interpretable tabular RL can remain competitive with deep policies under reward perturbations; thus, interpretability and

policy simplicity should be evaluated as design choices rather than sacrificed by default for model complexity (55).

Time-scale separation is essential to prevent unstable interactions. A seasonal fertilizer planner should not alter a daily transport policy's reward structure with each sensor ping. Conversely, a dispatch policy must react to road closures and mill queues more often than a growth model can be recalibrated. A supervisory coordinator can exchange forecasts and capacity envelopes while lower-level agents act inside those limits. These principles reflect the logic of multi-agent coordination, with the important qualification that non-stationarity can render jointly trained policies brittle when production cycles or equipment availability change (25,26).

Digital Twin Calibration, Uncertainty, And Sim-To-Real Transfer
High-fidelity simulation is not inherently accurate simulation. Calibration should use independent data segments representing different seasons, planting ages, vehicle fleets, collection routes, and mill operating regimes. Investigators need to distinguish errors due to sensors, incorrect initial states, mistaken parameters, and missing mechanisms. Physiological model uncertainty should be represented alongside stochastic arrivals and equipment failures. Training on a single clean simulator creates false confidence, so a robust protocol should perturb rainfall, road speeds, fruit arrivals, processing yields, energy demand, and sensor latency within empirically defensible ranges (19,20,30,31). Beyond crop systems, twin-assisted agricultural tractor control further illustrates that real-time state updating and virtual policy training can be coordinated for a bounded mechanical action, yet the process dynamics and certified safety protections of a palm oil mill still require independent engineering validation (27).

Historical logs are valuable for offline initialization but provide incomplete coverage of counterfactual actions. A mill log may contain few examples of intentional throughput reduction before an upstream quality spike; a harvest log rarely records what would have happened had the team chosen a different block. Offline policy evaluation must therefore identify support limits and avoid claiming causal improvement from purely observational correlations. Counterfactual simulation can extend the comparison set, but only after the transition model has passed predictive validation against held-out physical outcomes. This circular dependency is a core research difficulty and explains why promising offline reward estimates cannot substitute for supervised pilot deployment (24,29).

A realistic evaluation ladder should start with shadow recommendations, then controlled decision support, then bounded human-approved trials, and finally limited automation for well-understood reversible decisions. Shadow mode records what the policy would recommend without allowing it to affect physical operations. Controlled trials can alternate recommended and existing dispatch practices across suitable routes, shifts, or comparable periods while protecting against spillovers. Measurement should include not only mean economic performance but also tail risks, rejection rates, incidents, model drift, operator burden, and downstream quality differences. Such staged deployment translates lessons from safe RL and industrial control without borrowing their success claims wholesale (17,45).

Interoperability, Traceability, Cybersecurity, and Organizational Design
A fruit batch requires a persistent identity linking block, harvest date, collection point, transport movements, weighbridge transaction, grading result, processing interval, and downstream output estimate. Such linkage may be imperfect when loads are mixed, suppliers share vehicles, or batch-level physical segregation is absent; the data model should express mixtures and uncertainty rather than invent exact provenance. Semantic interoperability enables systems to agree on the definition of a harvest round, delivered weight, rejected fruit, downtime, or oil extraction metric. Food supply-chain digital twins offer transferable lessons about cross-organization visibility, while agricultural interoperability reviews show why semantic standards remain a practical bottleneck (9,22,56).

Cybersecurity needs to be addressed from the outset because digital twin connectivity increases the exposure of operational systems. The architecture should segregate advisory analytics from safety-critical control networks, authenticate sensors and users, log state and model changes, and fail safely when cloud connectivity or identity reconciliation breaks down. Contractors and third-party fruit suppliers may reasonably resist sharing commercially sensitive data; interfaces should request only the information needed for a specific decision and enforce role-based access. These governance choices are not incidental implementation details: invalid or manipulated weighbridge, inventory, or sensor data can redirect learned policies and undermine both economic and environmental claims (7,23,49).

A Comparative Performance and Evaluation Framework
An appropriate study compares the integrated architecture against several progressively stronger baselines: manual planning, established deterministic or mixed-integer optimization, prediction-plus-rules, and isolated RL agents. A twin–RL platform should only be judged superior if it produces a statistically and operationally material benefit after accounting for data collection, maintenance, computation, operator workload, and model governance. This requirement guards against comparing an advanced method with deliberately weak controls. Comparable lessons are visible in agricultural RL studies where optimized rules remain competitive and in mathematical oil palm scheduling studies that already extract significant value without learning policies (2,16).

Recommended outcome measures include harvested and delivered FFB by grade; time from harvest to mill processing; truck utilization and queue duration; oil extraction and quality indicators; mill throughput, unplanned downtime, and energy intensity; water and effluent indicators; labor compliance; financial contribution after operating costs; and safety events. Each measure should be linked to a causal pathway and a data-source definition. Leading indicators such as forecast accuracy and simulator calibration error are necessary but insufficient: a more accurate fruit-arrival model is valuable only if its use improves decisions or reduces risk. Environmental measurement must declare allocation rules and the boundary of emissions accounting (13,47).

Economic Feasibility, Inclusion, and Adoption Conditions
The architecture's economic case should be incremental and context dependent. A large integrated estate with multiple mills may have extensive telemetry and enough operational scale to justify expensive digital infrastructure. A smaller estate or smallholder collection network may obtain greater benefit from simple batch timestamps, route planning, and reliable grading rather than a high-fidelity per-palm simulator. Digitalization can thus widen performance gaps unless interoperable low-cost participation mechanisms accompany advanced analytics. Modular deployment allows organizations to prioritize interventions with clear measurable value while preserving future integration options (22,57).

The first commercial candidate is likely to be harvest–transport–mill coordination rather than full autonomous agronomy. Its decision cycle is fast enough for feedback, event data may be available, decisions are comparatively reversible, and the intended benefit can be estimated using dispatch and quality metrics. More ambitious nutrition, disease, and water-management RL should follow only once biological models can generalize across soils, ages, weather, genetics, and management histories. The contrast between crop-model demonstrations and oil palm-specific prediction studies underlines the need to stage benefits instead of treating the enterprise-wide twin as a single purchase or deployment event (3-5,10,21).

Research Agenda and Roadmap for Evidence Generation

Future research should develop an openly specified benchmark simulator with realistic palm maturation, variable harvest readiness, batch mixing, road disruptions, mill constraints, and environmental accounting. Such a benchmark would allow comparison of heuristics, mathematical programming, model-predictive control, offline RL, model-based RL, and cooperative agents using the same information and action opportunities. The benchmark should include intentionally difficult seasons and failures, since performance only under average conditions provides weak evidence of operational robustness. Researchers should report not only aggregate returns but also constraint violations, calibration diagnostics, sensitivity to reward weights, and generalization gaps (24,25).

Figure 3 operationalizes this evidence agenda as a four-phase deployment sequence with defined advancement gates and retained rollback procedures.

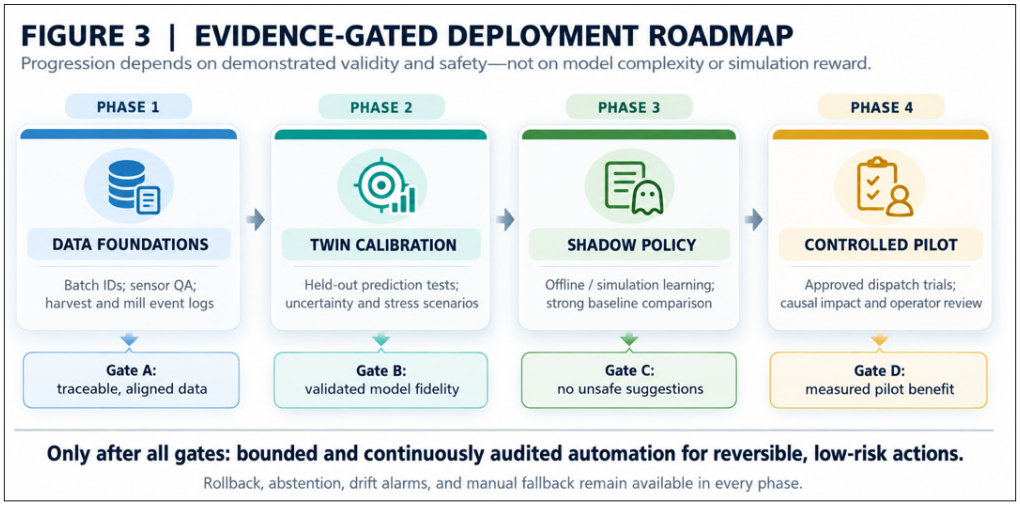


Figure 3: Evidence-gated validation and deployment roadmap (proposed).

Figure 3 translates the synthesis into an evidence-gated implementation sequence. Each phase specifies tests that must be passed before escalating operational authority: data reliability before twin calibration; out-of-sample model fit before policy comparison; shadow-mode robustness before human-approved trials; and credible measured benefits before constrained automation. The resulting program makes the proposed architecture falsifiable and protects operations from the mistaken assumption that virtual optimization is equivalent to validated real-world improvement (29,30).

A complementary agenda concerns the causal value of information. Does better block-level ripeness detection materially improve a mill's quality-adjusted economic return? Does a more frequently synchronized queue twin improve decisions beyond inexpensive radio or mobile coordination? Does adding an Environmental Twin change preferred production schedules when methane, steam, water, and effluent constraints are applied? Studies answering these questions would distinguish core value-creating features from attractive but redundant digital infrastructure. They would also help determine how much sensor precision or model detail is economically justified in different (12,14,58).

Answering the Research Questions Through an Integrative Synthesis

The literature supports a conditional answer to the architectural question: yes, an integrated twin–RL architecture is conceptually feasible because its component technologies and selected agricultural examples exist. However, feasibility is not equivalent to demonstrated performance in a coupled palm plantation and mill. Evidence for cross-stage optimization is strongest as an engineering rationale and weakest as prospective causal validation. In practical terms, the most defensible architecture is hierarchical, interoperable, event-driven, multirate, and supervised. It must preserve provenance across mixed fruit flows and accept uncertain states as a normal operating condition) (10,21,22,59).

The policy-design question also has a conditional answer. RL is most defensible where decisions recur, consequences can be modeled, action spaces are bounded, feedback is observable, and standard rules leave measurable improvement potential. Multi-agent RL is justified only when independently controlled subsystems must coordinate and simpler optimization is insufficient. Safety requires a separate veto mechanism and human authorization of consequential decisions. Adoption is warranted only after a rigorous comparison with competitive baselines, validation through shifted seasons and disruptions, and monitoring of economic, environmental, labor, and operational risks. These conditions convert technological ambition into testable managerial and scientific propositions (16,60-64).

Conclusion and Recommendations

The convergence of digital twins and reinforcement learning offers a plausible way to move the palm oil industry beyond isolated sensing and prediction toward coordinated, feedback-informed management. Its distinctive potential comes from

treating plantation biology, harvest decisions, transport flows, mill operations, and environmental constraints as interacting components of one decision system. The proposed Closed-Loop Palm Digital Twin provides a hierarchical structure and a cycle of sensing, synchronization, simulation, learning, recommendation, authorization, action, and verification. Crucially, the contribution is a conceptual synthesis and research framework, not a claim that an integrated oil palm RL controller has already been field-proven.

Substantively, the review recommends starting with auditable data and a reliable event backbone rather than with autonomous agents. Palm and mill models should be calibrated for specific assets, uncertainty should accompany predictions, and safety-critical processes must remain under certified controls. The most promising near-term experiment is a human-approved harvest–transport–mill scheduling pilot evaluated against credible existing methods. Subsequent phases may investigate agronomic interventions and coordinated multi-agent policies, provided that realistic biological constraints, operational worker protections, and environmental limits are enforced throughout.

For researchers, the priority is to create shared test environments, report negative results, establish transportability across seasons and estates, and measure realized business and sustainability outcomes rather than simulation rewards alone. For plantation companies and mills, the priority is staged investment, explicit model ownership, cybersecure connectivity, workforce involvement, and a policy that permits operators to reject and explain recommendations. For regulators and sector organizations, interoperable data vocabularies and transparent validation guidance could lower adoption barriers while improving accountability. Closed-loop optimization becomes scientifically and managerially credible only when every proposed action can be explained, authorized, measured, and corrected.

References

1. Peña González D, Cortés Borda D, Mele FD, Barrios Sarmiento A, Domínguez Santiago M (2021) An optimization approach for the design and planning of the oil palm supply chain in Colombia, Computers & Chemical Engineering 146: 107208.
2. Gonzalez-Negrete T, García-Cáceres RG, Escobar-Velásquez JW (2024) Optimal planning of oil palm fruit harvest, International Journal of Services and Operations Management 49: 311–340.
3. Jamshidi EJ, Yusup Y, Hooy CW, Kamaruddin MA, Mat Hassan H, et al, (2024) Predicting oil palm yield using a comprehensive agronomy dataset and 17 machine learning and deep learning models, Ecological Informatics 81: 102595.
4. Suharjito, Junior FA, Koeswandy YP, Debi, Nurhayati PW, Asrol M, Marimin (2023) Annotated Datasets of Oil Palm Fruit Bunch Piles for Ripeness Grading Using Deep Learning, Scientific Data 10: 72.
5. Puttinaovarat S, Chai-Arayalert S, Saetang W (2024) Oil Palm Bunch Ripeness Classification and Plantation Verification Platform: Leveraging Deep Learning and Geospatial Analysis and Visualization, ISPRS International Journal of Geo-Information 13: 158.
6. Lee CC, Koo VC, Lim TS, Lee YP, Abidin H (2022) A multi-layer perceptron-based approach for early detection of BSR disease in oil palm trees using hyperspectral images, Heliyon 8: e09252.
7. Ariesen-Verschuur N, Verdouw C, Tekinerdogan B (2022) Digital Twins in greenhouse horticulture: A review, Computers and Electronics in Agriculture 199: 107183.
8. Verdouw C, Tekinerdogan B, Beulens A, Wolfert S (2021) Digital twins in smart farming, Agricultural Systems 189: 103046.
9. Wang L (2024) Digital Twins in Agriculture: A Review of Recent Progress and Open Issues, Electronics 13: 2209.
10. Baja H, Kallenberg MGJ, Berghuijs HNC, Athanasiadis IN (2025) Adaptive fertilizer management for optimizing nitrogen use efficiency with constrained reinforcement learning, Computers and Electronics in Agriculture 237: 110554.
11. Pereira PHM, Piazza GC, Usman K, da Silveira LSC, Ceriotti VC, Pazin YS, de Freitas EP (2026) Digital twins in agriculture: A systematic literature review on modeling, semantics, and interoperability, Smart Agricultural Technology 14: 102283.
12. Mohd Nadzim UKH, Yunus R, Omar R, Ismail MHS (2023) Simulation of Energy Consumption in Sterilisation Process Using Aspen Plus and Response Surface Methodology, Journal of Oil Palm Research 35: 281–295.
13. Nasrin AB, Abdul Raman AA, Sukiran MA, Bukhari NA, Buthiyappan A, Subramaniam V, Abdul Aziz A, Loh SK (2024) Renewable Energy and Greenhouse Gases Emission Reduction Potential of Biogas from Palm Oil Mill Effluent, Journal of Oil Palm Research 36: 456–468.
14. Sulin SN, Mokhtar MN, Baharuddin AS, Mohammed MAP (2025) Simulation and techno-economic evaluation of integrated palm oil mill processes for advancing a circular economy, Cleaner Environmental Systems 19: 100323.
15. Tagarakis AC, Benos L, Kyriakarakos G, Pearson S, Sørensen CG, Bochtis D (2024) Digital Twins in Agriculture and Forestry: A Review, Sensors 24: 3117.
16. Kelly TD, Foster T, Schultz DM (2024) Assessing the value of deep reinforcement learning for irrigation scheduling, Smart Agricultural Technology 7: 100403.
17. Faria RdR, Capron BDO, Secchi AR, de Souza MB (2022) Where Reinforcement Learning Meets Process Control: Review and Guidelines, Processes 10: 2311.
18. Sharma A, Kosasih E, Zhang J, Brintrup A, Calinescu A (2022) Digital Twins: State of the art theory and practice, challenges, and open research questions, Journal of Industrial Information Integration 30: 100383.
19. Gu S, Yang L, Du Y, Chen G, Walter F, et al, (2024) A Review of Safe Reinforcement Learning: Methods, Theories, and Applications, IEEE Transactions on Pattern Analysis and Machine Intelligence 46: 11216–11235.
20. Mitsanis C, Hurst W, Tekinerdogan B (2024) A 3D functional plant modelling framework for agricultural digital twins, Computers and Electronics in Agriculture 218: 108733.
21. Kallenberg MGJ, Baja HA, Ilić M, Tomčić A, Tošić M (2025) Interoperable agricultural digital twins with reinforcement learning intelligence, Smart Agricultural Technology 12: 101412.
22. Huang Y, Ghadge A, Yates N (2024) Implementation of digital twins in the food supply chain: a review and conceptual framework, International Journal of Production Research 62: 6400–6426.
23. Melesse TY (2025) Digital twin-based applications in crop monitoring, Heliyon 11: e42137.
24. Moerland TM, Broekens J, Plaat A, Jonker CM (2023) Model-based Reinforcement Learning: A Survey, Foundations and Trends in Machine Learning 16: 1–118.
25. Gronauer S, Diepold K (2022) Multi-agent deep reinforcement learning: a survey, Artificial Intelligence Review 55: 895–943.
26. Agyeman BT, Decard-Nelson B, Liu J, Shah SL (2024) A semi-centralized multi-agent RL framework for efficient irrigation scheduling, Control Engineering Practice 153: 106183.

27. Zhang R, Zhu H, Chang Q, Mao Q (2025) A Comprehensive Review of Digital Twins Technology in Agriculture, Agriculture 15: 903.
28. Gu S, Yang L, Du Y, Chen G, Walter F, et al, (2024) A Review of Safe Reinforcement Learning: Methods, Theories, and Applications, IEEE Transactions on Pattern Analysis and Machine Intelligence 46: 11216–11235.
29. Prudencio RF, Maximo MROA, Colombini EL (2024) A Survey on Offline Reinforcement Learning: Taxonomy, Review, and Open Problems, IEEE Transactions on Neural Networks and Learning Systems 35: 10237–10257.
30. Ju H, Juan R, Gomez R, Nakamura K, Li G (2022) Transferring policy of deep reinforcement learning from simulation to reality for robotics, Nature Machine Intelligence 4: 1077–1087.
31. Zhu W, Guo X, Owaki D, Kutsuzawa K, Hayashibe M (2023) A Survey of Sim-to-Real Transfer Techniques Applied to Reinforcement Learning for Bioinspired Robots, IEEE Transactions on Neural Networks and Learning Systems 34: 3444–3459.
32. Bekkemoen Y (2024) Explainable reinforcement learning (XRL): a systematic literature review and taxonomy, Machine Learning 113: 355–441.
33. Milani S, Topin N, Veloso M, Fang F (2024) Explainable Reinforcement Learning: A Survey and Comparative Review, ACM Computing Surveys 56.
34. Glanois C, Weng P, Zimmer M, Li D, Yang T, et al, (2024) A Survey on Interpretable Reinforcement Learning, Machine Learning 113: 5847–5890.
35. Dazeley R, Vamplew P, Cruz F (2023) Explainable reinforcement learning for broad-XAI: a conceptual framework and survey, Neural Computing and Applications 35: 16893–16916.
36. Lai JW, Ramli HR, Ismail LI, Wan Hasan WZ (2023) Oil Palm Fresh Fruit Bunch Ripeness Detection Methods: A Systematic Review, Agriculture 13: 156.
37. Goh JY, Md Yunos Y, Mohamed Ali MS (2025) Fresh Fruit Bunch Ripeness Classification Methods: A Review, Food and Bioprocess Technology 18: 183–206.
38. Mansour MYMA, Dambul KD, Yeep CK (2023) A Review of Non-Destructive Ripeness Classification Techniques for Oil Palm Fresh Fruit Bunches, Journal of Oil Palm Research 35: 543–554.
39. Soetrisno YAA, Handoyo E, Sumardi, Sinuraya EW (2024) Oil Palm Level of Ripeness Classification Using EfficientDet-Lite CNN Architecture, Journal of Oil Palm Research 36: 618–629.
40. Khairunniza-Bejo S, Shahibullah MS, Noor Azmi AN, Jahari M (2021) Non-Destructive Detection of Asymptomatic Ganoderma boninense Infection of Oil Palm Seedlings Using NIR-Hyperspectral Data and Support Vector Machine, Applied Sciences 11: 10878.
41. Izzuddin MA, Hamzah A, Nisfariza MN, Idris AS (2025) Object-Based Image Analysis (OBIA) on Hyperspectral Imagery from Drone for Ganoderma Basal Stem Rot Disease Detection in Oil Palm, Journal of Oil Palm Research 37: 288–299.
42. Rai N, Choi D, Boyd NS, Schumann AW (2026) Advancing site-specific disease and pest management in precision agriculture: From reasoning-driven foundation models to adaptive, feedback-based learning, Computers and Electronics in Agriculture 250: 111859.
43. Tzachor A, Richards CE, Jeen S (2022) Transforming agrifood production systems and supply chains with digital twins, npj Science of Food 6: 47.
44. Rajasekhar N, Radhakrishnan TK, Samsudeen N (2025) Exploring reinforcement learning in process control: a comprehensive survey, International Journal of Systems Science 56: 3528–3557.
45. Bua C, Borgianni L, Adami D, Giordano S (2025) Reinforcement Learning-Driven Digital Twin for Zero-Delay Communication in Smart Greenhouse Robotics, Agriculture 15(12): 1290.
46. Lokavarapu NT, Sun X (2026) Integrated digital twins for agricultural machinery and robotics: A systematic review of architectures, quantified impacts, and deployment gaps, Journal of Agriculture and Food Research 29: 103051.
47. Guidani B, Ronzoni M, Accorsi R (2024) Virtual agri-food supply chains: A holistic digital twin for sustainable food ecosystem design, control and transparency, Sustainable Production and Consumption 46: 161–179.
48. Purlis E (2024) Digital Twin Methodology in Food Processing: Basic Concepts and Applications, Current Nutrition Reports 13: 914–920.
49. Tekinerdogan B, Verdouw C (2020) Systems Architecture Design Pattern Catalog for Developing Digital Twins, Sensors 20: 5103.
50. He B, Bai K-J (2021) Digital twin-based sustainable intelligent manufacturing: a review, Advances in Manufacturing 9: 1–21.
51. Psarommatis F, May G (2023) A literature review and design methodology for digital twins in the era of zero-defect manufacturing, International Journal of Production Research 61: 5723–5743.
52. Alibabaei K, Gaspar PD, Assunção ET, Alirezazadeh S, Lima TM (2022) Irrigation optimization with a deep reinforcement learning model: Case study on a site in Portugal, Agricultural Water Management 263: 107480.
53. Xu M, Wang X, Kang M, Chang F, Hua J (2026) IrriTwin: Integrating reinforcement learning with soil water and crop models toward digital-twin-oriented irrigation scheduling, Journal of Hydrology 678: 136034.
54. Lang Y, Zhang Y, Sun T, Chai X, Zhang N (2025) Digital twin-driven system for efficient tomato harvesting in greenhouses, Computers and Electronics in Agriculture 236: 110451.
55. Goldenits G, Neubauer T, Raubitzek S, Mallinger K, Weippl E (2025) Tabular reinforcement learning for reward robust, explainable crop rotation policies matching deep reinforcement learning performance, Computers and Electronics in Agriculture 237: 110571.
56. Pereira PHM, Piazza GC, Usman K, da Silveira LSC, Ceriotti VC, Pazin YS, de Freitas EP (2026) Digital twins in agriculture: A systematic literature review on modeling, semantics, and interoperability, Smart Agricultural Technology 14: 102283.
57. Awais M, Wang X, Hussain S, Aziz F, Mahmood MQ (2025) Advancing Precision Agriculture Through Digital Twins and Smart Farming Technologies: A Review, AgriEngineering 7: 137.
58. Omar Z, Abdul Majeed APP, Rosbi M, Ghazalli SA, Selamat H (2024) Outdoor oil palm fruit ripeness dataset, Data in Brief 55: 110667.
59. Liu M, Fang S, Dong H, Xu C (2021) Review of digital twin about concepts, technologies, and industrial applications, Journal of Manufacturing Systems 58: 346–361.
60. Melesse TY, Franciosi C, Di Pasquale V, Riemma S (2023) Analyzing the Implementation of Digital Twins in the Agri-Food Supply Chain, Logistics 7: 33.
61. Monteiro J, Barata J (2025) Digital twin-enabled regional food supply chain: A review and research agenda, Journal of

Industrial Information Integration 45: 100851.
62. Onaji I, Tiwari D, Soulatiantork P, Song B, Tiwari A (2022) Digital twin in manufacturing: conceptual framework and case studies, International Journal of Computer Integrated Manufacturing 35: 831–858.
63. Tao F, Xiao B, Qi Q, Cheng J, Ji P (2022) Digital twin modeling, Journal of Manufacturing Systems 64: 372–389.
64. Tekinerdogan B (2023) On the Notion of Digital Twins: A Modeling Perspective, Systems 11: 15.

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