USOD PUBLISHERS
United Scientific Open Discovery

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

Research Article

From Prediction to Intervention: Causal AI, Heterogeneous Treatment Effects, and Counterfactual Decision Intelligence for Sustainable Oil Palm Management

Loso Judijanto

IPOSS Jakarta, Indonesia

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

Citation: Loso Judijanto (2026) From Prediction to Intervention: Causal AI, Heterogeneous Treatment Effects, and Counterfactual Decision Intelligence for Sustainable Oil Palm Management. J of Adv Art Int Sys & Frontiers 1(2):1-14.

Received Date: October 07, 2026
Accepted Date: October 12, 2026
Published Date: November 11, 2026

Abstract

Artificial intelligence in agriculture has advanced rapidly in crop mapping, disease detection, stress recognition, yield forecasting, and resource monitoring. Yet most applications remain predictive rather than intervention-oriented. Accurate prediction can indicate which oil palm blocks are likely to experience nutrient deficiency, drought stress, disease, or yield decline, but it cannot by itself establish whether changing a management practice will improve the outcome. This qualitative literature review examines the transition from predictive artificial intelligence toward causal AI, heterogeneous treatment-effect estimation, and counterfactual decision intelligence for sustainable oil palm management. It synthesizes recent developments in causal inference, structural causal models, directed acyclic graphs, double/debiased machine learning, causal forests, meta-learners, causal representation learning, and counterfactual reasoning, and maps these approaches onto oil palm management questions involving fertilization, harvesting intervals, water management, Ganoderma control, and replanting. Particular attention is given to treatment-effect heterogeneity because agronomic responses differ across soils, palm ages, climatic conditions, nutrient status, prior management, and production systems. The review also examines methodological risks associated with using remote-sensing-derived covariates and predicted outcomes in causal analyses. A Predict–Explain–Intervene–Verify framework is proposed to integrate predictive intelligence, causal identification, intervention-effect estimation, human decision making, sustainability constraints, and post-intervention verification. The central argument is that the next frontier of plantation AI should move beyond asking what will happen toward determining which feasible actions are likely to cause better economic, agronomic, and environmental outcomes.

Keywords: causal artificial intelligence; causal machine learning; counterfactual decision intelligence; heterogeneous treatment effects; double machine learning; causal forests; precision agriculture; oil palm; sustainable intensification; remote sensing

Introduction

Artificial intelligence has become an increasingly important component of digital agriculture. Machine learning, computer vision, satellite remote sensing, unmanned aerial vehicles, proximal sensing, Internet of Things (IoT) devices, and data-intensive farm-management systems are being used to characterize crop conditions at spatial and temporal scales that would have been difficult to achieve using conventional field inspection alone. Within the oil palm sector, these technologies have been applied to plantation mapping, age estimation, nutrient diagnosis, disease identification, ripeness assessment, yield prediction, environmental monitoring, and plantation surveillance. High-resolution satellite products now permit mapping of plantation extent and approximate establishment year across major producing countries, while machine-learning models increasingly classify nutritional status, detect basal stem rot, and identify spatial patterns relevant to management (Danylo et al., 2021; Hashim et al., 2021; Tee et al., 2021; Tan et al., 2022; Chungcharoen et al., 2022). A recent sector-wide review further shows that plantation mapping, disease detection, and fruit grading are among the most prevalent Industry 4.0 applications in the oil palm industry, underscoring the current emphasis on observation and prediction rather than intervention (Bin Mohamad Zaki et al., 2025).

These developments provide substantial informational value, but the majority of agricultural AI still addresses predictive questions. A model may predict that a particular block will have low fresh fruit bunch (FFB) yield next season, classify a palm as nutrient-deficient, or estimate that a plantation has elevated disease risk. Such predictions are useful for prioritization, but they do not automatically determine what action should be taken. The management question is rarely only which palms are at risk; it is usually which feasible intervention will change that risk, by how much, for which palms, and at what cost. This distinction is fundamental because a variable can possess high predictive importance while having little or no manipulable causal effect on the outcome (Guo et al., 2020; Imbens, 2020; Kuang et al., 2020; Prosperi et al., 2020).

For example, a predictive yield model might find a strong association between fertilizer application and yield. Yet blocks receiving more fertilizer may also differ systematically in soil fertility, management quality, accessibility, previous yields, palm age, estate budget allocation, or agronomist attention. A conventional model optimized for out-of-sample prediction may exploit these relationships without separating the effect of fertilizer from the processes that determine where fertilizer is applied. Consequently, the conditional association ($P(Y|T,X)$) cannot generally be equated with the intervention distribution ($P(Y|do(T),X)$). The latter is the

quantity required when management asks what would happen if fertilizer treatment were deliberately changed (Imbens, 2020; Schölkopf et al., 2021; Feuerriegel et al., 2024).

The distinction is especially consequential for oil palm because plantation systems are highly heterogeneous. Yield response differs across soil types, climatic environments, stand ages, nutrient conditions, water availability, planting materials, topographies, historical management regimes, and socioeconomic settings. Recent Indonesian research has shown that smallholder yields can remain far below attainable levels and that nutrient management, harvesting interval, pruning, and weed control contribute materially to these gaps. Nutrient deficiencies are widespread, particularly for potassium, while balanced nutrient supply can significantly improve productivity. More recent paired-field evidence demonstrates that packages of better management practices can produce substantial yield and profit improvements, but the magnitude of response depends partly on initial conditions and the extent to which practices are adopted (Monzon et al., 2023; Sugianto et al., 2023; Lim et al., 2023; Thoumazeau et al., 2024; Sugianto et al., 2025). Recent field evidence reports roughly 40% higher annual FFB yield and 20% greater profit in paired better-management fields, while earlier national-scale work documented widespread nutrient deficiencies and large yield gaps.

These findings make oil palm a particularly relevant domain for heterogeneous treatment-effect analysis. A plantation-wide average fertilizer response may conceal blocks with large positive responses, blocks with marginal economic responses, and blocks in which additional fertilizer generates little benefit because another constraint—such as water deficit, rooting conditions, disease, or poor harvesting—is binding. Similar heterogeneity is likely for irrigation, drainage, harvest frequency, disease interventions, and replanting. The relevant decision quantity is therefore often not an average treatment effect but a conditional treatment effect that varies with covariates describing local circumstances (Nie and Wager, 2021; Brand, Zhou and Xie, 2023; Baiardi and Naghi, 2024; Feuerriegel et al., 2024).

Causal machine learning offers tools for addressing this problem. Double/debiased machine learning (DML) can flexibly control high-dimensional nuisance relationships while preserving valid estimation of target causal parameters. Causal forests and related methods can estimate heterogeneous treatment effects without requiring the analyst to pre-specify all effect modifiers. Meta-learners and orthogonal learners can combine flexible prediction with causal estimands, while structural causal models and directed acyclic graphs help clarify the assumptions required for identification. Counterfactual approaches extend the analysis from estimating treatment effects to comparing feasible alternative actions for an individual unit or management context (Nie and Wager, 2021; Chou et al., 2022; Brand, Zhou and Xie, 2023; Feuerriegel et al., 2024; Baiardi and Naghi, 2024; Richens, Lee and Johri, 2020).

Agriculture is beginning to adopt this intervention-oriented perspective. On-farm experimentation and causal forest approaches have been used to examine site-specific input responses, and recent work has combined causal and predictive machine learning to characterize fertilizer-effect heterogeneity. Digital agriculture research increasingly recognizes that explainability alone is insufficient when the objective is to recommend interventions, because feature attribution does not establish what would happen under a change in treatment. The emerging literature therefore supports an evolution from predictive analytics toward evidence-based intervention intelligence (Kakimoto et al., 2022; de Lara et al., 2023; Tanaka et al., 2023; Ryo, 2022; Tsoumas et al., 2025; Kouame, Heuvelink and Bindraban, 2026).

Remote sensing expands this opportunity by providing repeated, spatially continuous measurements of vegetation condition, plantation age, canopy characteristics, environmental stress, and land-use context. Nevertheless, remotely sensed outcomes create additional causal challenges. Satellite variables may be noisy proxies for latent agronomic states, and outcomes predicted from Earth observation (EO) imagery can carry model error into a second-stage causal analysis. If such predicted outcomes are treated as perfectly observed ground truth, intervention-effect estimates and uncertainty intervals can be distorted. Recent satellite-based impact-evaluation research shows that machine-learning prediction error can attenuate causal coefficients and understate uncertainty unless ground observations are used for diagnosis and correction (Pelletier et al., 2026). Consequently, causal applications of EO require explicit consideration of measurement error, calibration data, temporal alignment, and ground validation rather than simply inserting machine-learning predictions into causal models (Danylo et al., 2021; Wuepper, Oluoch and Hadi, 2025; Pelletier et al., 2026).

This review therefore addresses an intellectual and operational transition: from asking **what will happen?** toward asking **what will happen if management changes?** It synthesizes contemporary causal inference and causal machine-learning concepts and applies them to sustainable oil palm management. The review concentrates on fertilization, harvest management, water and drought management, Ganoderma intervention, and replanting because these represent important decisions for which predictive information alone is insufficient. It also develops a Predict–Explain–Intervene–Verify framework intended to connect prediction, causal identification, treatment-effect estimation, counterfactual reasoning, human management, sustainability objectives, and post-intervention learning.

The study addresses eight questions: how causal AI differs from predictive and explainable AI; which causal methods are appropriate for plantation management; how treatment-effect heterogeneity can support site-specific management; how remotely sensed information should enter causal analyses; which identification challenges are especially important in plantations; how causal estimates can be converted into economically and environmentally meaningful decisions; how recommended interventions should be verified; and which data and governance arrangements are necessary to operationalize causal decision intelligence. Together, these questions connect methodological validity with practical plantation decisions and define the scope of the thematic synthesis developed in the subsequent sections.

Literature Review

From Prediction to Causal Questions

Predictive machine learning seeks functions that perform well when estimating unknown outcomes from observed features. Given predictors (X), a model approximates (E(Y|X)) or another statistical functional of the conditional outcome distribution. Performance is commonly assessed using predictive error, discrimination, calibration, or related metrics. Such objectives are appropriate when the central task is forecasting yield, classifying disease status, or detecting stress. They are not necessarily appropriate when the parameter of interest is the consequence of changing an action because intervention effects depend on a different data-generating question (Guo et al., 2020; Kuang et al., 2020; Baiardi and Naghi, 2024).

Causal inference instead specifies potential outcomes under alternative treatment states. Let $Y_i(1)$ denote the outcome that would occur for management unit i if treatment were applied and $Y_i(0)$ the outcome under a comparison regime. The individual treatment effect for unit i is then expressed as follows:

$$\tau_i = Y_i(1) - Y_i(0)$$

Only one potential outcome is normally observed for a given unit, producing the fundamental problem of causal inference. The analytical problem is therefore to recover meaningful contrasts using experimental design, identification assumptions, quasi-experimental variation, or appropriate adjustment for observed treatment-selection mechanisms (Imbens, 2020; Prosperi et al., 2020; Feuerriegel et al., 2024).

At the population level, the average treatment effect (ATE) summarizes the mean contrast between the two potential outcomes. It is defined as follows:

$$ATE = E[Y(1) - Y(0)]$$

Precision management often requires a more localized estimand, the conditional average treatment effect (CATE). It conditions the expected treatment contrast on covariates X and is written as follows:

$$CATE(x) = E[Y(1) - Y(0) \mid X = x]$$

The CATE is particularly important for precision agriculture because management frequently seeks to determine whether an intervention is advantageous under specific combinations of soil, weather, crop status, and historical management. A plantation manager may care less about whether an intervention works on average than about whether it works in Block A with acidic soil and mature palms under moderate water stress, or in Block B with high potassium deficiency but limited access during the wet season (Nie and Wager, 2021; Brand, Zhou and Xie, 2023; Shiba and Inoue, 2024).

Structural Causal Models and Directed Acyclic Graphs

Structural causal models provide a formal language for representing mechanisms connecting treatments, covariates, mediators, and outcomes. Directed acyclic graphs (DAGs) make such assumptions visible. In an oil palm fertilizer problem, soil characteristics may affect both fertilizer allocation and yield. Previous yield may influence the agronomist's decision to increase fertilizer while also predicting subsequent yield. Palm age affects yield directly and may determine fertilizer recommendations. Conditioning on the wrong variables can therefore introduce bias rather than remove it (Imbens, 2020; Schölkopf et al., 2021).

A DAG forces the analyst to distinguish pre-treatment confounders from post-treatment mediators. Suppose fertilizer improves leaf nutrient concentration, which subsequently increases bunch production. Leaf nutrient concentration measured after application is partly a mediator of fertilizer effect. Controlling mechanically for the post-treatment nutrient measurement could remove a portion of the effect that the analysis is trying to estimate. Likewise, selecting fields based on a variable influenced by both management quality and disease severity could create collider bias. These issues are not resolved merely by having more variables or a more powerful predictive algorithm (Guo et al., 2020; Jiao et al., 2024).

This observation is crucial for plantation AI. Digital platforms can collect hundreds of predictors, yet causal validity depends on **which** variables are adjusted for and **when** they are measured. Causal AI should therefore not be understood as predictive machine learning plus a causal label. It requires a clearly specified intervention, an estimand, temporal ordering, an identification strategy, and explicit assumptions regarding treatment assignment and potential outcomes (Prosperi et al., 2020; Feuerriegel et al., 2024).

Causal Machine Learning

Conventional regression can estimate causal parameters when the causal structure is sufficiently simple and model assumptions are appropriate. Modern plantation data, however, often contain nonlinear relationships, interactions, high-dimensional sensor variables, remotely sensed covariates, and complex treatment-assignment mechanisms. Causal machine learning seeks to use flexible algorithms while retaining an estimand tied to intervention effects rather than predictive accuracy alone (Brand, Zhou and Xie, 2023; Baiardi and Naghi, 2024).

Double/debiased machine learning is relevant because it uses machine learning to estimate nuisance functions—such as treatment propensity and expected outcome—while orthogonalization and sample splitting reduce bias in the target causal effect. The approach is attractive when numerous covariates potentially influence both treatment and outcome. In a plantation context, these may include topography, prior yield trajectories, weather histories, nutrient status, access, management quality proxies, vegetation indices, and estate characteristics (Baiardi and Naghi, 2024).

Causal forests are designed to identify variation in treatment response. Instead of optimizing splits primarily for predictive outcome accuracy, causal forests seek partitions that reveal differences in treatment effects. This can help characterize responders and non-responders under sufficient overlap and credible identification assumptions. Related R-learner methods separate the treatment-effect signal from nuisance outcome and treatment models, allowing flexible estimation of heterogeneous response functions (Nie and Wager, 2021).

Causal representation learning addresses an additional challenge: much agricultural data arrive as images, signals, or other high-dimensional representations. Satellite images do not naturally enter causal models as a small set of transparent covariates. Learning latent representations that preserve causally relevant information while improving generalization across environments is therefore an important research frontier. Such approaches remain methodologically demanding, however, and should not be treated as replacements for field knowledge regarding mechanisms and treatment assignment (Schölkopf et al., 2021; Jiao et al., 2024).

Explainability Is Not Causality

Explainable artificial intelligence (XAI) has become important because agronomists and plantation managers need to understand why algorithms generate particular predictions. Feature-attribution approaches can indicate that rainfall anomaly, canopy index, age, potassium status, or prior yield strongly influences a prediction. Such explanations improve transparency, debugging, and communication, but they describe the internal basis of a model prediction rather than necessarily identifying the effect of manipulating the feature (Ryo, 2022; Ding et al., 2022).

Counterfactual explanations likewise require careful interpretation. A conventional counterfactual explanation might state that a block would be classified as "low risk" if a vegetation index were higher. Yet vegetation index itself may not be directly actionable. Changing fertilizer, irrigation, drainage, disease treatment, or harvest timing might change vegetation condition, but the causal pathway has to be established separately. A plausible algorithmic counterfactual is therefore not automatically a feasible causal

intervention (Chou et al., 2022; Guidotti, 2022; Verma et al., 2024).

The distinction is operationally important. Causal counterfactual reasoning asks whether an actual management action would change the outcome under a structurally coherent model. Decision intelligence goes one step further by comparing multiple feasible treatments under economic, environmental, labour, and institutional constraints.

Oil Palm as a Heterogeneous Management System

Oil palm provides an unusually compelling context for treatment-effect heterogeneity. Plantation performance varies with weather, soil, age, management history, nutrient availability, planting material, drainage, harvesting practice, and organizational capacity. Indonesian analyses show substantial gaps between observed and attainable yields and identify management factors—including nutrients and harvesting interval—as important contributors (Monzon et al., 2021; Monzon et al., 2023).

Smallholder nutrition illustrates this heterogeneity. Sugianto et al. (2023) documented extensive nutrient deficiency across Indonesian smallholder fields, while Lim et al. (2023) showed that nutrient supply was frequently insufficient and imbalanced. Such results strongly justify improving fertilizer management, but they do not imply that every farmer or block should receive an identical nutrient rate. Response will depend upon baseline nutrient status, soil supply, achievable yield, water conditions, past fertilization, crop age, labour constraints, fertilizer price, and expected FFB value (Sugianto et al., 2023; Lim et al., 2023; Thoumazeau et al., 2024).

Field management is also economically heterogeneous. Mettauer et al. (2021) demonstrated links between smallholder practices and economic performance in Jambi, while later evidence from paired fields found improved agronomic management capable of increasing both yield and profit. Importantly, treatment response varied with baseline productivity and management adoption, reinforcing the argument that intervention policy should be conditional rather than universally uniform (Mettauer et al., 2021; Sugianto et al., 2025).

Remote Sensing as a Causal Data Infrastructure

Remote sensing has major advantages for causal agricultural research because it can provide repeated observations before and after interventions over large areas. Satellite products can characterize plantation age, vegetation status, rainfall exposure, flooding, land-use context, canopy condition, and neighbouring management. UAV imagery can provide finer-resolution disease or canopy information. These features can improve covariate balance, detect spatial spillovers, and lower the cost of longitudinal outcome monitoring (Danylo et al., 2021; Tee et al., 2021; Hashim et al., 2021).

However, remote-sensing information should be divided conceptually into at least three roles. First, remotely observed variables may be **pre-treatment covariates**, useful for adjusting confounding. Second, they may function as **treatment indicators**, for example detecting drainage changes or replanting. Third, they may be **predicted outcomes**, such as remotely inferred yield or disease severity. The third role is the most methodologically sensitive because errors in a predicted outcome can propagate into effect estimation.

Remote sensing therefore does not eliminate the need for field observation. Instead, causal plantation intelligence should combine strategically selected ground measurements with spatially extensive EO information. Such hybrid designs offer a realistic pathway to large-scale evaluation while avoiding the false assumption that remotely predicted variables are error-free measurements.

Method

Review Design

This study adopts a **qualitative literature review**, not a systematic literature review. The objective is conceptual integration and theory-guided synthesis rather than exhaustive enumeration of every publication meeting a predetermined search protocol. Consequently, the review does not employ PRISMA screening counts or claim exhaustive coverage. Instead, literature is selected purposively for its relevance to causal inference, causal machine learning, counterfactual reasoning, digital agriculture, remote sensing, oil palm agronomy, and sustainable plantation management.

The qualitative design is appropriate because direct applications of advanced causal machine learning in oil palm remain limited, while relevant knowledge is distributed across statistics, econometrics, computer science, agricultural systems, precision agriculture, remote sensing, and crop-management literature. A narrowly bounded systematic search would risk excluding methodological advances developed in other domains that have clear conceptual relevance for oil palm management.

Literature Identification and Selection

Priority was given to peer-reviewed journal literature published from 2020 to 2026. Conceptual and methodological papers were identified around the terms causal inference, causal machine learning, heterogeneous treatment effects, double machine learning, causal forest, structural causal model, counterfactual explanation, causal representation learning, and digital agriculture. Agricultural evidence was then identified using combinations of precision agriculture, fertilizer response, on-farm experimentation, remote sensing, satellite outcomes, and agricultural decision making.

Oil palm literature was purposively identified around management domains central to the causal framework: yield gaps, fertilization, nutrient status, harvesting intervals, drought and water management, Ganoderma detection and management, remote sensing, plantation age, better management practices, sustainable intensification, and replanting. Literature was retained when it contributed to at least one of four functions: defining a causal concept, demonstrating an estimation approach, documenting an oil palm management problem, or providing empirical evidence regarding heterogeneity and intervention potential.

Thematic Analysis

The selected literature was coded conceptually according to treatment, outcome, potential confounders, heterogeneity variables, analytical method, identification logic, data type, spatial scale, and management implication. Rather than producing numerical meta-analytic effects, the review synthesizes recurring conceptual patterns and identifies where conventional predictive framing is inadequate.

The analysis organizes findings into interconnected themes: the dominance of prediction; causal identification; heterogeneous treatment effects; fertilizer management; harvest management; water intervention; Ganoderma intervention; replanting; remote sensing; sustainability-constrained decisions; and verification. The final synthesis uses these themes to construct the Predict–Explain–Intervene–Verify framework.

Results: Thematic Findings

Predictive Intelligence Dominates Oil Palm Digitalization

The oil palm digitalization literature demonstrates considerable technological sophistication. Satellite and radar data can distinguish plantation areas and estimate stand establishment periods. Hyperspectral and multispectral signals can assist nutrient diagnosis, while machine learning and remote sensing can identify symptoms associated with Ganoderma infection. These advances reduce the cost of observation and allow managers to prioritize field inspection over large production areas (Amirruddin et al., 2020; Danylo et al., 2021; Hashim et al., 2021; Tee et al., 2021; Chungcharoen et al., 2022; Tan et al., 2022). Recent synthesis of Industry 4.0 applications in the oil palm sector likewise finds that mapping and disease detection account for a substantial share of published digital applications, reinforcing the argument that current systems are predominantly observational and predictive (Bin Mohamad Zaki et al., 2025).

Yet an important asymmetry appears: digital systems are generally stronger at **state estimation** than **intervention estimation**. Detecting that a palm is stressed is different from determining whether fertilizer, drainage improvement, disease treatment, or another action will reverse the problem. Similarly, identifying an ageing plantation does not determine the economically optimal date for replanting. Prediction narrows the decision space but usually does not resolve it.

The conceptual gap can be organized into four related questions. Prediction asks what is likely to happen under the observed data-generating process. Explanation asks which observed variables or mechanisms plausibly structure that pattern. Intervention asks what would happen if management deliberately changed treatment T, while counterfactual decision intelligence compares feasible alternatives for the same management unit. The last two questions provide the greatest decision value but also require stronger causal assumptions than prediction or explanation alone (Imbens, 2020; Feuerriegel et al., 2024).

Figure 1 illustrates the conceptual transition from observational intelligence to actionable causal decision intelligence. The progression is not simply a sequence of increasingly complex algorithms; each stage changes the question being asked. Prediction estimates expected outcomes under the observed data-generating process, whereas intervention and counterfactual reasoning explicitly consider alternative management regimes.

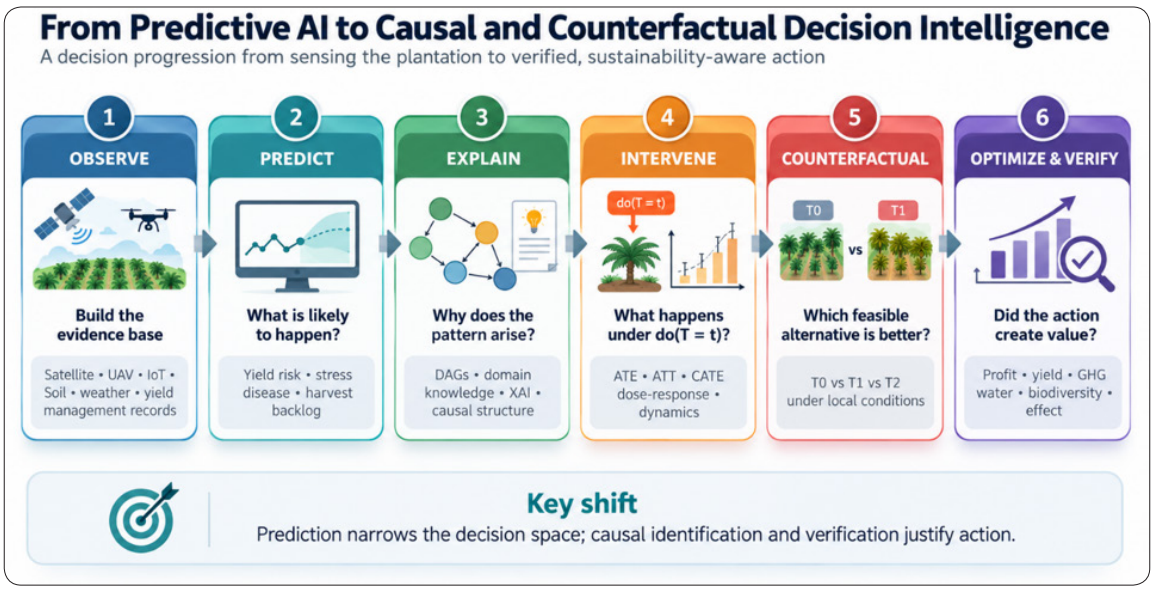


Figure 1: From Predictive AI to Causal and Counterfactual Decision Intelligence

The central message of Figure 1 is that feature importance should not be translated directly into agronomic prescription. A factor can be predictive yet non-manipulable; another variable can be manipulable but confounded; and an intervention may generate different outcomes in different locations. Causal AI therefore inserts an identification layer between prediction and recommendation.

Causal Identification Must Precede Algorithm Selection

The literature consistently indicates that powerful machine-learning estimators cannot compensate for poorly defined causal questions. The treatment must be explicit, the comparison condition meaningful, and treatment timing separated from outcome observation. Relevant confounders need to be identified based on domain knowledge, not only variable-selection algorithms (Imbens, 2020; Guo et al., 2020; Feuerriegel et al., 2024).

Consider a fertilizer example. Let (T) represent additional potassium and (Y) subsequent FFB yield. Soil potassium supply, palm age, historical yield, rainfall, management capability, and prior nutrient status may affect both treatment assignment and yield. A causal model needs to address these common causes. If some treated and untreated blocks have no overlap in relevant characteristics, positivity violations can make causal comparisons unstable even with very large datasets.

Plantation systems create additional challenges through temporal feedback. Prior yield influences current management, current management affects nutritional condition, and nutritional condition influences future yield and later treatment. Treatments are often continuous rather than binary, while multiple interventions are jointly applied. Causal models therefore need to evolve from simple one-time treatment contrasts toward dose–response relationships, time-varying treatments, and dynamic treatment regimes.

Heterogeneous Treatment Effects as the Core of Precision Intervention

Average treatment effects are useful for estate-level planning but can be misleading when treatment response varies widely.

The CATE formalizes this local treatment contrast and can be expressed as follows:
$\tau(x) = E[Y(1) - Y(0) \mid X = x]$

This estimand provides a conceptual bridge between causal inference and precision agriculture because it asks how treatment effects vary with local characteristics. It is especially useful when the same intervention is expected to produce different biological or economic responses across soils, palm ages, water regimes, or management histories.

Causal forests, R-learners, doubly robust learners, and related estimators can search flexibly for treatment-effect heterogeneity while reducing the need to specify all effect interactions in advance. Nevertheless, heterogeneity estimation is statistically demanding. Apparent subgroups can arise from small samples, inadequate overlap, noisy outcomes, or overfitting, requiring out-of-sample validation and agronomic interpretation (Nie and Wager, 2021; Inoue et al., 2024; Shiba and Inoue, 2024).

Agricultural studies increasingly demonstrate the value of this perspective. Causal forest approaches have been used in site-specific input management, while on-farm experimentation is generating datasets better suited to intervention estimation than ordinary farm records. Research on fertilizer responses shows that treatment effects can vary substantially with local conditions, reinforcing the inadequacy of uniform recommendations (Kakimoto et al., 2022; Tanaka et al., 2023; de Lara et al., 2023; Kouame, Heuvelink and Bindraban, 2026).

For oil palm, heterogeneity should be treated as a central design principle. A future causal recommendation system should identify at least four classes of management unit: high responders, economically marginal responders, near-zero responders, and potentially adverse responders. The objective would not be to maximize treatment intensity but to allocate resources where expected causal benefit exceeds financial and environmental cost.

Figure 2 depicts why average recommendations can be inferior to heterogeneous intervention rules. Even when the estate-wide average treatment effect is positive, individual blocks may respond very differently.

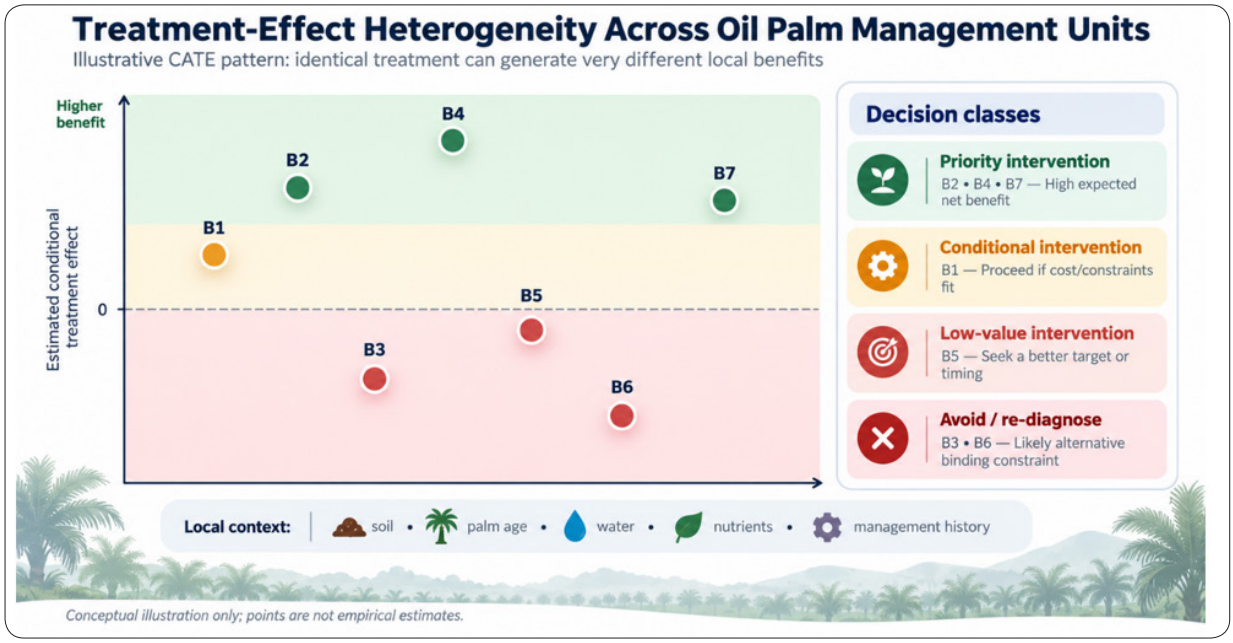


Figure 2: Treatment-Effect Heterogeneity Across Oil Palm Management Units

The managerial implication is significant. Precision agriculture should not merely use more spatially detailed predictions; it should use spatial information to estimate **spatially differentiated causal response**. This distinction converts site-specific agriculture from precision observation to precision intervention.

Fertilizer Management: From Correlation to Causal Dose–Response

Fertilizer is arguably the most immediate domain for causal AI in oil palm. Nutrient deficiency is widespread in many smallholder systems, and potassium represents a particularly important limitation. Fields with adequate nutrient status can substantially outperform nutrient-deficient fields, while poor balance among nitrogen, phosphorus, potassium, and magnesium contributes to productivity constraints (Sugianto et al., 2023; Lim et al., 2023). Indonesian field evidence has found widespread K deficiency and large yield differences between adequately and inadequately nourished fields, providing strong motivation for better nutrition management.

However, observational evidence relating nutrient application to yield is not automatically a causal fertilizer-response function. Farmers with better access to capital may apply more fertilizer while simultaneously practicing better harvesting and weed management. Estates may direct fertilizer toward historically productive blocks, creating treatment selection. Conversely, agronomists might target low-performing fields with additional fertilizer, potentially producing negative confounding.

The relevant causal quantity is frequently continuous rather than binary because fertilizer is applied at different rates. A conditional dose-response function can therefore be represented as follows:
$\mu(t, x) = E[Y(t) \mid X = x]$

Here, t denotes fertilizer dose and x represents local agronomic conditions. Management must then evaluate the marginal response of the expected outcome to a change in dose:
$\partial E[Y(t) \mid X = x] / \partial t$

This marginal response must be interpreted alongside fertilizer cost and FFB value before an economically meaningful recommendation can be made. The causal estimate therefore becomes an input to economic optimization rather than an endpoint in itself.

An economically rational treatment rule should combine the expected production response with private and environmental costs. A simplified net-benefit rule can be represented as follows:
$NB_i(t) = P_y[\hat{Y}_i(t) - \hat{Y}_i(t_0)] - C_i(t) - E_i(t)$

In this expression, $NB_i(t)$ is expected net benefit, P_Y is output price, $C_i(t)$ is private intervention cost, and $E_i(t)$ represents environmental cost or a sustainability penalty. The formulation makes explicit that treatment choice is an optimization problem rather than only a prediction problem.

This formulation changes the optimization target. The objective is no longer maximum predicted yield but maximum expected sustainable net benefit subject to agronomic constraints. It also makes clear why the optimal nutrient rate can differ among blocks even when all display deficiency.

Recent Indonesian evidence provides a strong empirical foundation for such research. Better management practices combining improvements in nutrient, soil, harvest, and weed management generated substantial yield and profit gains in paired smallholder fields, while the magnitude of the yield response varied with baseline conditions and adoption intensity (Sugianto et al., 2025). These findings demonstrate both the potential of intervention and the heterogeneity that causal ML should attempt to characterize.

Harvesting Interval as a Causal Intervention
Harvesting frequency directly influences fruit losses, ripeness, labour requirements, and delivery quality. Indonesian smallholder research has documented relatively long and variable harvest intervals and highlighted agronomic, socioeconomic, and institutional drivers of harvesting behaviour (de Vos et al., 2023). Yield-gap analysis likewise identifies harvesting practice as one contributor to lower observed productivity (Monzon et al., 2023).

Yet harvest interval is highly confounded. Low-yield fields may be harvested less frequently because the expected amount of fruit does not justify labour or transportation costs. Long harvest intervals could therefore be correlated with low yield partly because low expected yield causes infrequent harvesting, rather than only because infrequent harvesting causes yield loss. Distance to road, contractor availability, labour scarcity, farm size, FFB price, rainfall, and cooperative arrangements may simultaneously influence harvest timing and performance.

A causal study should therefore define harvest interval as treatment and distinguish direct causal effects from reverse causation and treatment selection. Useful designs could include randomized harvest scheduling, encouragement designs, staggered implementation, instrumental variables based on exogenous labour or logistical variation, difference-in-differences designs with multiple time periods, or carefully controlled panel analyses (Callaway and Sant'Anna, 2021).

Potential outcomes should also include more than FFB quantity. Treatment could affect loose-fruit loss, bunch ripeness distribution, free fatty acid (FFA) development, mill recovery, labour cost, transport efficiency, and net revenue. Causal decision intelligence should thus evaluate a multidimensional harvest-management outcome rather than a single predicted yield number.

Water Management and Drought Adaptation
Water deficit is a major determinant of oil palm performance, but the response to water intervention depends on climate, soil, drainage, rooting depth, palm age, and nutrient status. Irrigation research demonstrates that water management can alter palm physiological response during ENSO-related drought, while longer-term yield analyses show meaningful spatial and temporal relationships between climate and production (Brum et al., 2021; Monzon et al., 2022).

This domain illustrates treatment interaction. Fertilizer may have limited marginal benefit under severe water constraint, whereas improved nutrition may produce large benefits when moisture is adequate. Conversely, water intervention may have smaller effects where nutrient deficiency remains binding. Sustainable management therefore requires estimation of interactions among interventions rather than isolated one-factor predictions.

For water management, the relevant treatment effect can likewise be conditional on local agronomic context. A simple binary water-intervention contrast can be written as follows:
$\tau_W(x) = E[Y(W = 1) - Y(W = 0) \mid X = x]$

for water intervention and then allow (x) to include nutrient status, rainfall history, soil hydraulic properties, topographic indices, and palm age. A more advanced system could estimate combined treatment effects of water and fertilizer and identify whether they are complementary or substitutable.

Such analysis has substantial sustainability relevance. Water intervention should not be recommended only where biological yield response is positive; water availability, downstream consequences, energy use, drainage-related emissions, and ecological costs should constrain the decision. The causal estimand therefore needs to be integrated with resource accounting.

From Ganoderma Detection to Causal Disease Management
Ganoderma basal stem rot is among the most important disease threats in oil palm production, and substantial digital research has focused on detecting infection earlier and more accurately. Machine-learning reviews and remote-sensing studies show considerable potential for UAV, spectral, radar, and deep-learning methods to discriminate disease conditions and severity classes (Hashim et al., 2021; Tee et al., 2021).

Detection, however, is only the first component of decision support. A manager needs to know whether sanitation, removal, trenching, biological treatment, soil management, replanting, or a combination of actions will causally reduce subsequent disease progression or transmission. A highly accurate disease classifier cannot answer this question unless the data contain credible treatment comparisons.

Disease causality also introduces **interference**. Standard treatment-effect frameworks often assume that one unit's treatment does not affect another unit's outcome. That assumption is difficult to maintain for spatially transmissible plant diseases. Removing an infected palm may affect neighbouring palms by changing inoculum pressure; treating one zone can therefore alter outcomes in another zone.

Future Ganoderma causal AI should incorporate spatial exposure mapping, infection networks, distance-dependent spillovers, and longitudinal disease states. Remote sensing could help observe these spillovers, but causal models need to define treatment and exposure at spatial neighbourhood scales rather than treating palms

as independent observations.

Replanting as a Dynamic Counterfactual Decision

Replanting illustrates why counterfactual decision intelligence requires a longer horizon than one-season prediction. Ageing palms can remain productive but become more difficult to harvest, may experience higher disease risk, and eventually exhibit declining economic performance. Replanting creates an immediate production interruption and substantial investment cost, followed by future gains from a new stand.

Indonesian research has shown that strategic replanting and yield improvement are important for maintaining production without continued area expansion (Zhao et al., 2023). At the same time, replanting can alter biodiversity and ecosystem conditions, indicating that the intervention should be considered beyond production alone (Pashkevich et al., 2021). The national importance of yield intensification is reinforced by climate-smart analyses showing that narrowing yield gaps on existing land can reduce pressure for expansion into fragile ecosystems (Monzon et al., 2021).

Replanting is inherently a dynamic treatment decision because current action changes a sequence of future outcomes rather than only the next-period yield. A simplified discounted value representation is:

$$V_i(a_t) = \Sigma_{k=0}^{H} \delta^k E[B_{i,t+k}(a_t, a_{t+1}, \ldots)]$$

In this expression, $V_i(a_t)$ is the expected present value of an intervention sequence, a_t is the action at time t, H is the decision horizon, and δ is the discount factor. The representation highlights why replanting analysis must compare trajectories of future benefits and costs rather than a single contemporaneous outcome.

The relevant question is not simply whether old palms have lower yield. It is whether replanting now produces a superior discounted trajectory relative to continued management, delayed replanting, rehabilitation, or partial replanting under uncertainty regarding prices, labour, climate, disease, planting material, and financing.

Causal AI could support this decision by combining treatment-effect models with dynamic programming or policy-learning approaches. For smallholders, the model must additionally incorporate liquidity constraints and temporary income loss. A biologically optimal replanting decision can be economically impossible for a farmer without bridging finance, illustrating why feasible counterfactuals must include socioeconomic constraints.

Remote Sensing: Powerful but Not Causally Neutral

Oil palm remote sensing has created detailed datasets describing plantation extent, age, disease symptoms, vegetation condition, and other properties. Danylo et al. (2021), for example, demonstrated large-scale mapping of plantation extent and detection year in Indonesia, Malaysia, and Thailand. Such data can strengthen causal analysis because plantation age and landscape conditions become observable across large spatial domains.

Nevertheless, causality requires careful temporal alignment. If a vegetation index is measured after treatment, it may be a mediator rather than a confounder. If a predicted disease score is generated partly from features influenced by intervention, using it as a control can introduce post-treatment bias. If remotely predicted yield contains systematic error correlated with treatment or geography, causal estimates can inherit that error. Evidence from satellite-based agricultural impact evaluation shows that prediction error can attenuate treatment-effect estimates and understate standard errors, making validation samples and explicit correction procedures important for credible inference (Pelletier et al., 2026).

Remote sensing should therefore be integrated through a causal data architecture. Pre-treatment imagery can help characterize baseline states and treatment propensity. Post-treatment imagery can monitor mechanisms and outcomes. Ground-truth subsets should be preserved for calibration and causal validation. The system must record which variables are observed, predicted, or constructed, because these categories are not statistically interchangeable.

Sustainability-Constrained Treatment Effects

Sustainable intensification requires more than proving that an intervention increases yield. An action can have positive production effects while increasing fertilizer loss, emissions, water use, soil degradation, or biodiversity pressure. Conversely, a management strategy can reduce environmental impact while lowering production sufficiently to encourage land expansion elsewhere.

Oil palm research demonstrates the importance of these trade-offs. Climate-smart intensification analyses show that raising yields on existing land can potentially reduce pressure to expand cultivation into forests and peatlands, while conservation set-asides and biodiversity-sensitive management demonstrate that within-plantation strategies also influence ecological performance (Fleiss et al., 2020; Monzon et al., 2021; Hanafiah et al., 2022; Popkin et al., 2022). These trade-offs include greenhouse-gas (GHG) emissions as well as water, soil, and biodiversity outcomes, so causal decision rules should not optimize production in isolation.

The future causal question should therefore be multi-outcome rather than yield-only. One useful representation is a vector of conditional effects across economic and environmental endpoints:

$$\tau(x) = [\tau_{yield}(x), \tau_{profit}(x), \tau_{GHG}(x), \tau_{water}(x), \tau_{biodiversity}(x)]^T$$

A treatment may be desirable only when expected yield or profit gains are sufficiently large relative to environmental costs. Decision rules can use explicit constraints, shadow prices, or Pareto frontiers rather than collapsing all outcomes into a single opaque model objective.

Discussion and Analysis

The Fundamental Shift: From Accuracy to Decision Validity

The most important implication of this review is that plantation AI requires a change in evaluation criteria. For predictive models, accuracy is central. For intervention systems, **decision validity** becomes central. A model with excellent predictive performance may still produce poor recommendations if treatment effects are confounded or if the decision rule is applied outside the population where effects can be identified.

Causal intelligence therefore does not replace prediction. Prediction remains essential for risk stratification, early warning, baseline characterization, nuisance-function estimation, and outcome forecasting. Instead, causal AI changes the relationship between prediction and action. Prediction becomes one layer of the decision system rather than the final basis for recommendation.

This distinction parallels broader causal ML research. Modern causal methods use flexible prediction algorithms internally but place them within an estimand-aware architecture. Double machine learning, for example, may use highly accurate models for treatment and outcome nuisance components while preserving the intervention parameter through orthogonalization. Causal forests use recursive methods but optimize for treatment-effect heterogeneity. The analytical objective determines whether machine learning is merely predictive or part of a causal system (Brand, Zhou and Xie, 2023; Baiardi and Naghi, 2024; Feuerriegel et al., 2024).

PEIV: Predict–Explain–Intervene–Verify

To operationalize this transition, this review proposes a Predict–Explain–Intervene–Verify (PEIV) framework. The framework is intended to separate state estimation from causal identification while retaining an explicit human decision and post-intervention learning loop.

The **Predict** stage characterizes the plantation system. Satellite, UAV, weather, soil, sensor, disease, yield, economic, and management data are combined to estimate current state and future risk. Outputs might include yield forecasts, disease probabilities, nutrient-deficiency risk, drought vulnerability, harvest backlog, or replanting urgency.

The **Explain** stage develops the causal problem. Agronomists and analysts specify the intervention, outcome, time horizon, causal graph, pre-treatment confounders, possible mediators, effect modifiers, and assumptions required for identification. XAI can support diagnosis but does not substitute for causal structure.

The **Intervene** stage estimates potential outcomes under feasible actions. Treatment-effect models generate ATEs, CATEs, continuous dose–response curves, or dynamic effects. Counterfactual alternatives are filtered by feasibility, profitability, labour, sustainability, and institutional constraints.

The **Verify** stage closes the learning loop. Recommended actions are implemented through randomized trials, paired management units, staggered adoption, quasi-experiments, or monitored operational deployment. Observed outcomes are compared with expected effects, and deviations are used to revise the model and intervention policy.

Figure 3 summarizes the PEIV architecture. Its defining feature is the verification loop: the system is not considered successful simply because it produces recommendations. Recommendations generate new evidence that should improve causal knowledge over time.

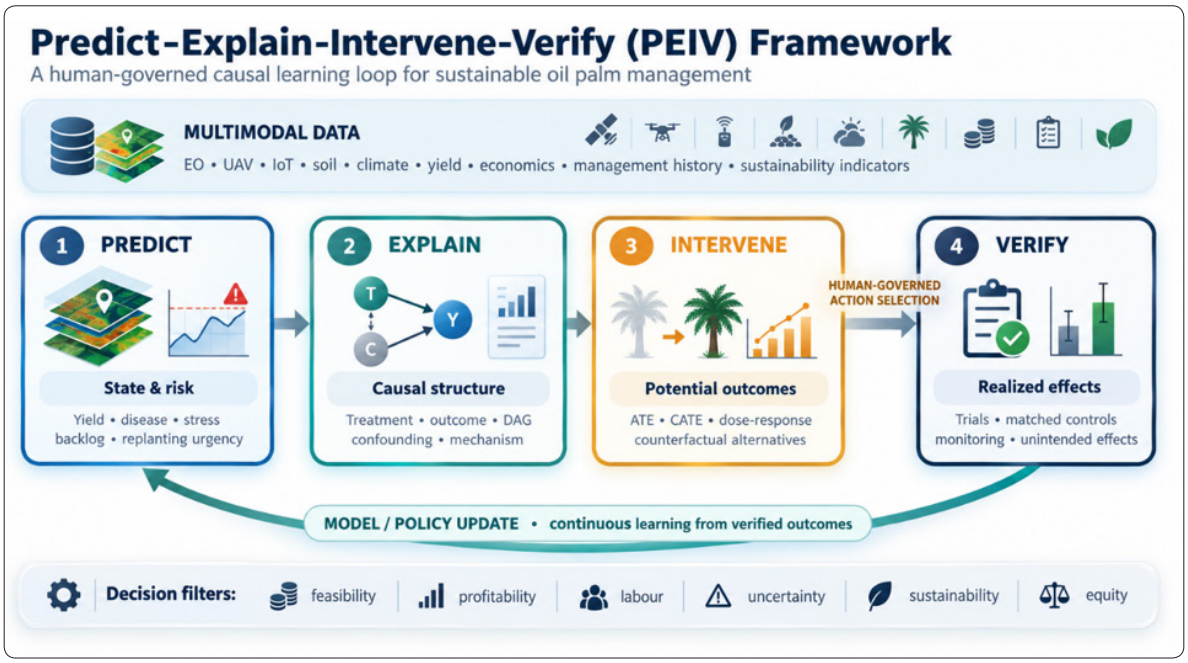


Figure 3: Predict–Explain–Intervene–Verify Framework for Sustainable Oil Palm Management

PEIV differs from a conventional AI pipeline because human decision makers remain explicitly positioned between causal estimation and implementation. Agronomists possess domain knowledge about intervention feasibility and biological plausibility that may not be encoded in the historical dataset. Human governance is particularly necessary when estimated effects are uncertain or when interventions have environmental and social consequences.

Causal DAGs as Operational Management Tools

DAGs are often treated as statistical diagrams, but in plantation management they can become interdisciplinary decision tools. A fertilizer DAG can bring agronomists, data scientists, soil specialists, estate managers, and economists into the same causal discussion.

Figure 4 presents a simplified DAG for fertilizer decision making. It distinguishes common causes of fertilizer treatment and yield from mediating pathways through nutritional status.

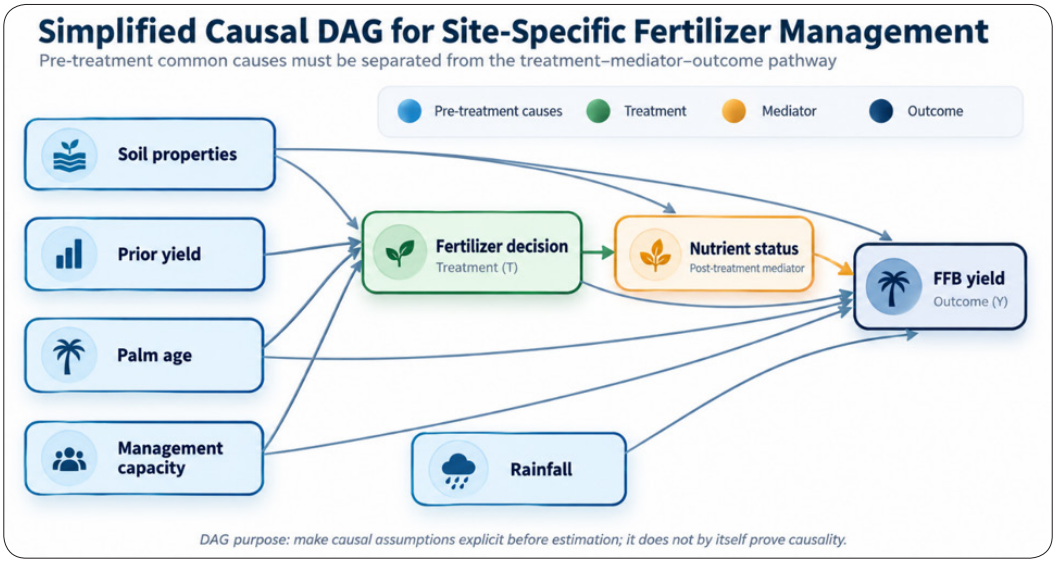


Figure 4: Simplified Causal DAG for Site-Specific Fertilizer Management

The DAG highlights several reasons naïve regression may fail. Management capacity can increase fertilizer use while independently improving harvest and weed management; prior yield can influence fertilizer decisions and subsequent yield; and nutrient status after treatment is a mediator rather than necessarily an appropriate adjustment variable. The purpose of the DAG is not to prove the causal model but to make assumptions visible and contestable before estimation.

Why Heterogeneity Changes Plantation Economics

Treatment-effect heterogeneity changes the economics of precision management. Consider a fixed fertilizer budget. A uniform allocation may spread fertilizer across all blocks regardless of expected response. A causal allocation can instead rank blocks according to expected incremental net benefit.

For a management unit i, incremental net benefit can be used to translate a heterogeneous treatment effect into an economic decision quantity. A simplified expression is:

$$INB_i = P_{FFB}\,\hat{\tau}_i - C_i - \lambda E_i$$

Here, $\hat{\tau}_i$ is the estimated treatment effect, C_i is application cost, E_i represents environmental burden, and λ converts environmental impact into a constraint or decision weight. This formulation allows agronomic response, private cost, and sustainability considerations to be evaluated within the same decision rule.

If treatment effects vary strongly, targeting can increase return to the same resource budget. This principle is especially valuable for smallholders with limited capital. Rather than advising a general increase in fertilizer use, an intervention system could identify which nutrient, rate, timing, and field combination has the highest probability of generating a profitable response.

However, such targeting must avoid creating a false impression of individual-level certainty. CATEs are expectations conditional on measured characteristics, not deterministic predictions of biological response. Recommended systems should therefore communicate uncertainty intervals, treatment overlap, and evidence strength.

From Single Treatment to Intervention Portfolios

Oil palm management rarely involves isolated treatments. Fertilizer, harvest interval, pruning, weed management, drainage, soil cover, disease sanitation, and labour allocation interact. Evidence that packages of better management practices can improve smallholder yield and profit reinforces the importance of intervention portfolios rather than single-factor optimization (Sugianto et al., 2025).

Causal analysis of treatment bundles is more difficult because the number of possible combinations increases rapidly. One solution is factorial experimentation where practical. Another is sequential learning that evaluates a limited set of biologically plausible treatment regimes. Multi-treatment causal forests and policy-learning methods may also contribute, but direct empirical validation is essential.

The concept of binding constraints provides a useful agronomic interpretation. A fertilizer treatment may fail because water deficit is limiting; irrigation may fail because palms are severely nutrient deficient; replanting may be ineffective if poor planting material or management persists. Causal decision intelligence should therefore diagnose which constraint is most likely to be relaxed by a particular intervention.

Spatial and Temporal Interference

Traditional causal estimators often rely on versions of the stable unit treatment value assumption. Plantation systems frequently violate simple forms of this assumption. Drainage in one area can affect neighbouring hydrology. Disease treatment can change inoculum pressure. Pest management can alter population dynamics beyond the treated block. Replanting can influence labour availability, transportation, local microclimate, or ecological connectivity.

Consequently, treatment needs to include both direct intervention and exposure to neighbouring intervention. Spatial causal models could define treatment intensity within distance bands or hydrological units, while network-based methods could represent disease transmission and operational connections.

Temporal interference is equally important. Fertilizer applied now can influence crop response across multiple production cycles; water deficit can affect later flowering and bunch development; replanting creates years of delayed production. Dynamic causal models and longitudinal g-methods therefore deserve more attention in future oil palm research.

External Validity and Transportability
A treatment effect estimated in one estate cannot automatically be transferred to another estate. Soil, climate, labour systems, genetics, baseline management, disease pressure, and institutional capacity may differ. Climate-smart oil palm research has already demonstrated the importance of spatial domains when extrapolating productivity interventions, while more recent spatial intensification frameworks similarly emphasize the conditions under which field evidence can be transferred (Monzon et al., 2021; Agus et al., 2024).

Causal transportability should therefore become an explicit design problem. Rather than asking whether a model trained in Estate A predicts Estate B accurately, causal AI should ask whether the **treatment-response mechanism** remains invariant enough for effect estimates to travel. Causal representation learning and invariant approaches may eventually contribute to this problem, but agronomic domain definition remains essential (Schölkopf et al., 2021).

Smallholder Inclusion and Algorithmic Feasibility
The smallholder context introduces constraints that purely technical models can overlook. Farmers may face limited access to fertilizer, finance, harvesting labour, transport, or extension. Certification studies and field research show that institutional and resource constraints strongly influence management practice and sustainability outcomes (Watts et al., 2021; Mettauer et al., 2021; Ogahara et al., 2022).

A counterfactual recommendation such as "increase potassium by X kg per palm" is not meaningful if the farmer cannot access the fertilizer or finance the purchase. Similarly, "shorten harvest interval" may be infeasible where contractors are unavailable. Feasible counterfactuals must therefore include socioeconomic variables and potentially interventions on constraints themselves—credit access, farmer groups, extension, logistics, or collective purchasing.

The same principle applies to equity. A policy-learning algorithm optimized only for total output may allocate scarce support toward farms expected to produce the greatest absolute gains, potentially excluding resource-poor farmers. Sustainable causal decision intelligence should therefore incorporate distributional objectives and minimum-access constraints rather than treating equity as an external consideration.

Human Governance, Accountability, and Verification
High-stakes agronomic interventions should not be executed merely because a model assigns them a high expected treatment effect. Causal estimates depend on assumptions about confounding, measurement, temporal order, and model transportability. These assumptions need to be visible to agronomists and managers.

Human governance should involve at least four checks. First, **biological plausibility**: is the estimated response consistent with crop physiology? Second, **identification credibility**: was treatment sufficiently comparable across units? Third, **operational feasibility**: can the treatment actually be implemented at the required timing and quality? Fourth, **sustainability acceptability**: are likely environmental effects within agreed constraints? Verification is equally important. Plantation companies possess a particular advantage because management interventions are repeated over many blocks and years. This creates opportunities for embedded experimentation. Small randomized strip trials, staggered implementation, paired blocks, or adaptive experiments can convert routine operations into cumulative causal evidence.

Such an architecture changes the role of AI from static model deployment to learning intervention governance. Its value depends on repeated verification because observed outcomes provide new evidence about where the intervention model succeeds, where it fails, and when the treatment policy should be revised.

Research Agenda
The first priority is to create intervention-ready datasets. Historical operational records are useful but often lack explicit reasons for treatment assignment. Future plantation data systems should record treatment timing, intended treatment, realized treatment, dose, implementation quality, decision rationale, baseline conditions, and multiple downstream outcomes.

Second, oil palm research needs more on-farm randomized and quasi-randomized experiments. On-farm experimentation in other crops demonstrates how operational variation can be deliberately structured to learn site-specific treatment effects (Kakimoto et al., 2022; Tanaka et al., 2023; de Lara et al., 2023). Oil palm estates offer long-lived spatial structures that can support longitudinal experimentation if operational constraints are addressed.

Third, multimodal causal datasets should integrate ground agronomy, remote sensing, weather, economic records, and sustainability indicators. Remote sensing should support causal inference without replacing field measurement, particularly because prediction error in satellite-derived outcomes can bias downstream impact estimates if it is ignored (Pelletier et al., 2026).

Fourth, benchmarking is needed. Future studies should compare conventional regression, predictive ML, DML, causal forests, and experimental estimates on the same plantation problems. Evaluation should include bias, calibration of treatment effects, policy value, robustness, transportability, and economic outcomes rather than only predictive accuracy.

Fifth, causal digital twins represent a longer-term opportunity. A digital twin that merely simulates observed correlations is insufficient for intervention planning; a credible causal digital twin would combine crop physiology, data-driven causal estimation, uncertainty, and repeated intervention validation. Recent agricultural digital-twin reviews emphasize the value of linking real-time sensing with simulation and decision support, while also identifying interoperability, model validity, and field validation as major unresolved constraints (Escribà-Gelonch et al., 2024; Melesse, 2025; Pereira et al., 2026).

Sixth, sustainability should be incorporated directly into causal estimands. Research should estimate intervention effects on yield and profitability alongside GHG emissions, nutrient losses, soil health, water use, disease spread, and biodiversity outcomes.

Finally, deployment research should study how agronomists and farmers interact with causal recommendations. Human-AI decision making depends not only on model quality but on how uncertainty and causal assumptions are communicated (Wang and Yin, 2022). A management system that correctly estimates a treatment effect but induces overconfidence or inappropriate automation can still fail.

Figure 5 translates these research priorities into an integrated oil palm intervention map. It shows how fertilization, water and drainage management, harvesting, disease management, and replanting can be evaluated through a common sustainability-verification layer before policy is updated for the next management cycle.

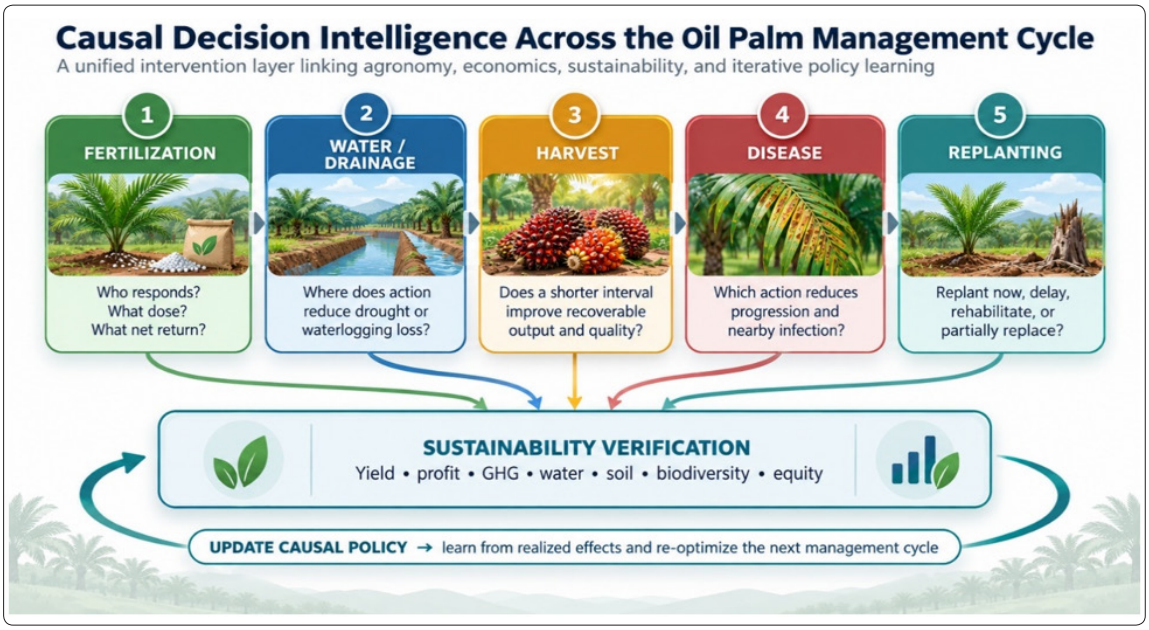


Figure 5: Causal Decision Intelligence Across the Oil Palm Management Cycle

The figure emphasizes that causal intelligence is not a separate AI application but a decision layer that can link multiple management domains. The common principle is explicit comparison of alternative actions rather than passive forecasting of outcomes.

Conclusion

Oil palm digitalization has made substantial progress in observing plantations and predicting future states. Remote sensing, machine learning, UAVs, and sensor systems can increasingly indicate where palms are stressed, where disease may occur, how plantation age varies, and where future yield is likely to fall. Yet management ultimately requires more than accurate prediction. It requires defensible answers about which feasible actions will change outcomes.

Causal AI provides the conceptual and methodological foundation for this transition. Its value lies not in replacing agronomy or experimentation, but in integrating them with data-intensive estimation. Structural causal models can clarify assumptions; causal machine learning can flexibly estimate intervention effects; heterogeneous treatment-effect methods can identify where responses differ; and counterfactual reasoning can compare feasible alternatives under local conditions.

Treatment-effect heterogeneity should become a foundational principle of precision oil palm management. Uniform recommendations are unlikely to be economically or environmentally optimal across heterogeneous plantations. Fertilizer, water management, harvesting, disease control, and replanting should increasingly be framed as conditional treatment decisions whose value depends on soils, crop age, weather, baseline management, disease status, economics, and operational constraints.

Remote sensing can greatly scale causal intelligence, but its role must be methodologically disciplined. Predicted variables are not automatically equivalent to observed outcomes, and temporal ordering matters when remotely sensed features are influenced by treatment. Field measurements and deliberately designed interventions therefore remain indispensable.

The proposed Predict–Explain–Intervene–Verify framework positions prediction as the beginning rather than the end of decision intelligence. Prediction identifies risk; causal explanation structures the decision problem; intervention models estimate effects; human managers select feasible and responsible actions; and verification determines whether expected benefits actually occur. Repeated verification converts plantation operations into a continuously learning causal system.

The strategic implication is clear: the next generation of oil palm AI should be judged not only by how accurately it predicts yield, stress, or disease, but by whether it can support evidence-based actions that causally improve productivity, profitability, resilience, environmental performance, and smallholder welfare. Moving from prediction to intervention is therefore not simply a methodological refinement. It represents a fundamental shift from observing agricultural systems more intelligently toward changing them more responsibly.

References

1. Agus F, Tenorio FA, Saleh S, Purwantomo DKG, Yustika RD, Marwanto S, Suratman, Sidhu MS, Cock J, Kam SP, Fairhurst T, Rattalino Edreira JI, Donough CR, Grassini P (2024) Guiding oil palm intensification through a spatial extrapolation domain framework. Agricultural Systems 213: 103778.
2. Amirruddin AD, Muharam FM, Ismail MH, Tan NP, Ismail MF (2020) Hyperspectral spectroscopy and imbalance data approaches for classification of oil palm macronutrients observed from frond 9 and 17. Computers and Electronics in Agriculture 178: 105768.
3. Baiardi A, Naghi AA (2024) The value added of machine learning to causal inference: Evidence from revisited studies. The Econometrics Journal 27(2): 213–234.
4. Bin Mohamad Zaki MA, Ooi J, Ng WPQ, How BS, Lam HL, Foo DCY, Lim CH (2025) Impact of Industry 4.0 technologies on the oil palm industry: A literature review. Smart Agricultural Technology 10: 100685.
5. Brand JE, Zhou X, Xie Y (2023) Recent developments in causal inference and machine learning. Annual Review of Sociology 49: 81–110.
6. Brum M, et al. (2021) Effects of irrigation on oil palm

transpiration during ENSO-induced drought in the Brazilian Eastern Amazon. Agricultural Water Management 245: 106569.
7. Callaway B, Sant'Anna PHC (2021) Difference-in-differences with multiple time periods. Journal of Econometrics 225(2): 200–230.
8. Chou YL, Moreira C, Bruza P, Ouyang C, Jorge J (2022) Counterfactuals and causability in explainable artificial intelligence: Theory, algorithms, and applications. Information Fusion 81: 59–83.
9. Chungcharoen T, et al. (2022) Machine learning-based prediction of nutritional status in oil palm leaves using proximal multispectral images. Computers and Electronics in Agriculture 198: 107019.
10. Danylo O, Pirker J, Lemoine G, Ceccherini G, See L, McCallum I, Hadi, Kraxner F, Achard F, Fritz S (2021) A map of the extent and year of detection of oil palm plantations in Indonesia, Malaysia and Thailand. Scientific Data 8: 96.
11. de Lara A, et al. (2023) Predicting site-specific economic optimal nitrogen rate using machine learning methods and on-farm precision experimentation. Precision Agriculture 24: 1792–1812.
12. de Vos RE, et al. (2023) Shortening harvest interval, reaping benefits? A study on harvest practices in oil palm smallholder farming systems in Indonesia. Agricultural Systems 211: 103753.
13. Ding W, Abdel-Basset M, Hawash H, Ali AM (2022) Explainability of artificial intelligence methods, applications and challenges: A comprehensive survey. Information Sciences 615: 238–292.
14. Escribà-Gelonch M, Liang S, van Schalkwyk P, Fisk I, Long NVD, Hessel V (2024) Digital twins in agriculture: Orchestration and applications. Journal of Agricultural and Food Chemistry 72(19): 10737–10752.
15. Feuerriegel S, Frauen D, Melnychuk V, Schweisthal J, Hess K, Curth A, Bauer S, Kilbertus N, Kohane IS, van der Schaar M (2024) Causal machine learning for predicting treatment outcomes. Nature Medicine 30: 958–968.
16. Fleiss S, et al. (2020) Conservation set-asides improve carbon storage and support associated plant diversity in certified sustainable oil palm plantations. Biological Conservation 248: 108631.
17. Guidotti R (2022) Counterfactual explanations and how to find them: Literature review and benchmarking. Data Mining and Knowledge Discovery.
18. Guo R, Cheng L, Li J, Hahn PR, Liu H (2020) A survey of learning causality with data: Problems and methods. ACM Computing Surveys 53(4): Article 75.
19. Hanafiah MM, et al. (2022) Impact of Malaysian palm oil on sustainable development goals: Co-benefits and trade-offs across mitigation strategies. Sustainability Science 17: 1639–1661.
20. Hashim IC, Mohamed Shariff AR, Bejo SK, Muharam FM, Ahmad K (2021) Machine-learning approach using SAR data for the classification of oil palm trees that are non-infected and infected with basal stem rot disease. Agronomy 11(3): 532.
21. Imbens GW (2020) Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. Journal of Economic Literature 58(4): 1129–1179.
22. Inoue K, et al. (2024) Machine learning approaches to evaluate heterogeneous treatment effects in randomized controlled trials: A scoping review. Journal of Clinical Epidemiology 176: 111538.
23. Jiao K, et al. (2024) Causal inference meets deep learning: A comprehensive survey. Research 7: 0467.
24. Kakimoto S, et al. (2022) Causal forest approach for site-specific input management via on-farm precision experimentation. Computers and Electronics in Agriculture 199: 107164.
25. Kouame KF, Heuvelink GBM, Bindraban PS (2026) Modeling and explaining fertilizer effect heterogeneity on maize yield in Ghana using causal and predictive machine learning. Field Crops Research 337: 110287.
26. Kuang K, Li L, Geng Z, Xu L, Zhang K, Liao B, Huang H, Ding P, Miao W, Jiang Z (2020) Causal inference. Engineering 6(3): 253–263.
27. Lim YL, Tenorio FA, Monzon JP, Sugianto H, Donough CR, Slingerland MA, Agus F, Darlan NH, Dwiyahreni AA, Farrasati R, Mahmudah N, Muhamad T, Nurdwiansyah D, Palupi S, Pradiko I, Rahutomo S, Saleh S, Syarovy M, Wiratmoko D, Grassini P (2023) Improving yield in oil palm smallholder fields requires higher and balanced nutrient use. Agricultural Systems 210: 103729.
28. Melesse TY (2025) Digital twin-based applications in crop monitoring. Heliyon 11(2): e42137.
29. Mettauer R, Baron V, Demitria P, Smit H, Alamsyah Z, Penot E, Bessou C, Chambon B, Ollivier J, Thoumazeau A (2021) Investigating the links between management practices and economic performances of smallholders' oil palm plots: A case study in Jambi province, Indonesia. Agricultural Systems 194: 103274.
30. Monzon JP, et al. (2022) Influence of weather and endogenous cycles on spatiotemporal yield variation in oil palm. Agricultural and Forest Meteorology 314: 108789.
31. Monzon JP, Lim YL, Tenorio FA, Farrasati R, Pradiko I, Sugianto H, Donough CR, Rattalino Edreira JI, Rahutomo S, Agus F, Slingerland MA, Zijlstra M, Saleh S, Nashr F, Nurdwiansyah D, Ulfaria N, Winarni NL, Zulhakim N, Grassini P (2023) Agronomy explains large yield gaps in smallholder oil palm fields. Agricultural Systems 210: 103689.
32. Monzon JP, Slingerland MA, Rahutomo S, Agus F, Oberthür T, Andrade JF, Couëdel A, Rattalino Edreira JI, Hekman W, van den Beuken R, Hidayat F, Pradiko I, Purwantomo DKG, Donough CR, Sugianto H, Lim YL, Farrell T, Grassini P (2021) Fostering a climate-smart intensification for oil palm. Nature Sustainability 4: 595–601.
33. Nie X, Wager S (2021) Quasi-oracle estimation of heterogeneous treatment effects. Biometrika 108(2): 299–319.
34. Ogahara Z, Jespersen K, Theilade I, Nielsen MR (2022) Review of smallholder palm oil sustainability reveals limited positive impacts and identifies key implementation and knowledge gaps. Land Use Policy 120: 106258.
35. Pashkevich MD, et al. (2021) Assessing the effects of oil palm replanting on arthropod biodiversity. Journal of Applied Ecology 58(1): 27–43.
36. Pelletier J, Korb M, Alemu S, Yonis MB, Lybbert TJ, Stigler M (2026) Causal inference with predicted outcomes: Correcting prediction error bias in satellite-based impact evaluation. Journal of Development Economics 179: 103655.
37. 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.
38. Popkin M, Reiss-Woolever VJ, Turner EC, Luke SH (2022) A systematic map of within-plantation oil palm management practices reveals a rapidly growing but patchy evidence base. PLOS Sustainability and Transformation 1(7): e0000023.
39. Prosperi M, Guo Y, Sperrin M, Koopman JS, Min JS, He X, Rich S, Wang M, Buchan IE, Bian J (2020) Causal inference and counterfactual prediction in machine learning for actionable healthcare. Nature Machine Intelligence 2:

369–375.
40. Richens JG, Lee CM, Johri S (2020) Improving the accuracy of medical diagnosis with causal machine learning. Nature Communications 11: 3923.
41. Ryo M (2022) Explainable artificial intelligence and interpretable machine learning for agricultural data analysis. Artificial Intelligence in Agriculture 6: 257–265.
42. Schölkopf B, Locatello F, Bauer S, Ke NR, Kalchbrenner N, Goyal A, Bengio Y (2021) Toward causal representation learning. Proceedings of the IEEE 109(5): 612–634.
43. Shiba K, Inoue K (2024) Harnessing causal forests for epidemiologic research: Key considerations. American Journal of Epidemiology 193(6): 813–818.
44. Sugianto H, Donough CR, Monzon JP, Sunawan, Pradiko I, Lim YL, Tenorio FA, Rizzo G, Rahutomo S, Agus F, Marwanto S, Slingerland M, Cock J, Grassini P (2025) Improving yield and profit in smallholder oil palm fields through better agronomy. Agricultural Systems 224: 104269.
45. Sugianto H, Monzon JP, Pradiko I, Tenorio FA, Lim YL, Donough CR, Sunawan, Rahutomo S, Agus F, Cock J, Amsar J, Farrasati R, Iskandar R, Rattalino Edreira JI, Saleh S, Santoso H, Tito AP, Ulfaria N, Slingerland MA, Grassini P (2023) First things first: Widespread nutrient deficiencies limit yields in smallholder oil palm fields. Agricultural Systems 210: 103709.
46. Tan YP, Cheor WL, Yeo SF, Leow SM (2022) Expert systems in oil palm precision agriculture: A decade systematic review. Journal of King Saud University – Computer and Information Sciences 34: 1569–1594.
47. Tanaka TST, et al. (2023) Toward an effective approach for on-farm experimentation: Lessons learned from a case study of fertilizer application optimization in Japan. Precision Agriculture 24: 2044–2060.
48. Tee CAT, Teoh YX, Por LY, Lai KW (2021) Discovering the Ganoderma boninense detection methods using machine learning: A review of manual, laboratory, and remote approaches. IEEE Access 9: 105780–105807.
49. Thoumazeau A, Mettauer R, Turinah, Junedi H, Baron V, Chéron-Bessou C, Ollivier J (2024) Effects of fertilization practices and understory on soil health and oil palm performances in smallholdings: An Indonesian case study. Agricultural Systems 213: 103802.
50. Tsoumas I, et al. (2025) Leveraging causality and explainability in digital agriculture. Environmental Data Science 4: e23.
51. Verma S, et al. (2024) Counterfactual explanations and algorithmic recourses for machine learning: A review. ACM Computing Surveys 56(12): 1–42.
52. Wang X, Yin M (2022) Effects of explanations in AI-assisted decision making: Principles and comparisons. ACM Transactions on Interactive Intelligent Systems 12(4): Article 27.
53. Watts JD, Pasaribu K, Irawan S, Tacconi L, Martanila H, Wisnu CG, Wiratama, Musthofa FK, Sugiarto BS, Manvi UP (2021) Challenges faced by smallholders in achieving sustainable palm oil certification in Indonesia. World Development 146: 105565.
54. Wuepper D, Oluoch WA, Hadi H (2025) Satellite data in agricultural and environmental economics: Theory and practice. Agricultural Economics 56: 493–511.
55. Zhao J, Elmore AJ, Lee JSH, Numata I, Zhang X, Cochrane MA (2023) Replanting and yield increase strategies for alleviating the potential decline in palm oil production in Indonesia. Agricultural Systems 210: 103714.