HarvestField: A Decision-Support Framework for Robot-Oriented Kiwifruit Canopy Management Using Spatial Harvestability Modelling

Changhao Wang
West Anhui University
Yu Hailu
West Anhui University
Wang Jicheng
West Anhui University
Li Rui
West Anhui University
Yu Zongliang
West Anhui University
Yang Lei
West Anhui University
Zhang Peng
04000141@wxc.edu.cn

West Anhui University

Research Article

Keywords: robotic harvesting, canopy management, harvestability field, functional–structural plant model, multi-objective optimisation, kiwifruit

Posted Date: June 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9528703/v1

Additional Declarations: No competing interests reported.

HarvestField: A Decision-Support Framework for Robot-Oriented Kiwifruit Canopy Management Using Spatial Harvestability Modelling

Abstract

Purpose Robotic harvesting systems are commonly evaluated in canopies treated as fixed operating environments, although kiwifruit canopy architecture is actively shaped by seasonal pruning, fruit thinning, tipping, and cane training. This study developed HarvestField, a decision-support framework for quantifying how robot-oriented canopy management affects harvesting feasibility before field operation.

Methods HarvestField defines a spatial harvestability field $\mathcal{H}(x) \in [0,1]$ as the product of reachability, visibility, clearance, and detachability sub-fields. A mass-weighted field average at fruit positions yields the Harvested Yield Index (HYI). The framework couples an L-system-based functional–structural plant model of Hayward kiwifruit with NSGA-III multi-objective optimisation over six management variables, jointly evaluating HYI, yield, and labour cost.

Results Calibration against a published 12,000-fruit robotic kiwifruit harvesting benchmark produced HYI=0.562, with an absolute deviation of 0.004 from the observed success rate of 0.558. HarvestField-optimised management increased mean HYI relative to standard management (0.642 versus 0.587). Under equal fruit counts, $\mathcal{H}$-guided thinning improved HYI by up to 0.301 compared with uniform thinning. Sobol analysis identified clearance neighbourhood radius as the dominant parameter ($S_T = 0.576$).

Conclusion HarvestField provides a spatially explicit and robot-specific framework for linking kiwifruit canopy management with harvesting feasibility. It supports management trade-off analysis while indicating the need for independent orchard-scale validation before deployment.

Keywords: robotic harvesting; canopy management; harvestability field; functional–structural plant model; multi-objective optimisation; kiwifruit

Introduction

The automation of selective harvesting has become a central engineering challenge in high-value horticulture because it requires the coordinated resolution of perception, motion planning, manipulation, and crop-handling problems within biologically variable production environments. Unlike cereal or forage harvesting, where mechanisation operates on bulk material streams, fresh-market fruit harvesting requires individual fruit to be located, approached, detached, and transferred without unacceptable damage to the product or canopy. This operational requirement is increasingly consequential as horticultural industries face persistent labour constraints, rising production costs, and seasonal workforce instability; consequently, agricultural robotics is no longer a peripheral technology but an increasingly necessary component of biosystems engineering for labour-sensitive crop systems (Zhou et al., 2022).

Despite rapid advances in machine vision, manipulator control, and end-effector design, robotic harvesting of high-value crops remains substantially less reliable than conventional manual picking. A widely cited review of harvesting robots showed that reported systems were constrained not only by sensing and actuation limits but also by the complexity of the crop environment, including occlusion, fruit clustering, collision risk, and heterogeneous accessibility (Zhou et al., 2022; Droukas et al., 2023). These constraints are particularly acute in trellised fruit crops, where the robot must operate inside a three-dimensional canopy whose local structure determines whether a fruit is visible to the sensor, reachable by the manipulator, accessible to the end-effector, and mechanically detachable without damaging neighbouring fruit or vegetation.

Kiwifruit provides a representative and practically important case for this problem. Commercial kiwifruit orchards are commonly trained on overhead pergola or T-bar systems, in which fruit are generally suspended below a vegetative canopy but may also occur within dense regions of shoots, leaves, peduncles, and trellis elements. Field evaluation of a multi-arm robotic kiwifruit harvester demonstrated the technical feasibility of autonomous picking, but also showed that harvesting success remained limited under realistic commercial conditions, with missed fruit and failed picking events arising from the combined effects of visibility, reachability, scheduling, and physical interaction constraints (Williams et al., 2019). This evidence indicates that, although improvements in robot hardware and algorithms are essential, the harvestable fraction of fruit is also governed by the spatial organisation of the crop environment in which the robot operates.

Most robotic-harvesting research has nevertheless been framed primarily as a robot-side adaptation problem. Within this paradigm, the canopy is treated as an exogenous and essentially fixed operating environment, while research effort is directed toward improving detection accuracy, three-dimensional localisation, collision-free trajectory generation, arm coordination, and compliant grasping. Such developments are indispensable; however, they do not fully address failure modes caused by inherently unfavourable fruit placement, excessive local canopy density, or insufficient approach clearance. More broadly, the technical literature on agricultural robotics has repeatedly emphasised that biological variability and unstructured crop environments are defining barriers to robust field operation (Yang et al., 2023; Liu et al., 2024). This observation implies that performance gains may also be obtained by reducing variability or improving structural compatibility on the crop side, rather than by requiring the robot to compensate for all unfavourable canopy conditions after they have already formed.

From the agronomic perspective, the premise that canopy structure is fixed is not appropriate. The architecture of a kiwifruit canopy at harvest is the cumulative outcome of management decisions made during the growing season, including winter pruning, retained cane number, bud retention density, fruit thinning intensity, summer tipping, and cane training angle. Functional-structural plant modelling has shown that kiwifruit vine architecture, carbon dynamics, organ growth, and management responses can be represented in an integrated three-dimensional simulation framework (Cieslak et al., 2011), and recent reviews further emphasise the role of FSPMs in linking plant architecture, phenotyping, and crop-management decisions (Soualiou et al., 2021).

Recent studies have begun to challenge the separation between crop design and robot design. In protected cultivation, modification of sweet pepper crop architecture has been shown to improve robotic harvesting by increasing visibility and reachability, thereby demonstrating that crop-side intervention can alter robot-relevant performance variables (van Herck et al., 2020). Similarly, orchard architecture and harvesting-robot kinematics have been co-considered in optimisation studies, establishing that plant training systems and robot design should not be treated as independent engineering decisions (Bloch et al., 2018). More recently, genome-edited flower morphology combined with autonomous robotic pollination has provided a high-profile demonstration of crop-robot co-design at the reproductive-architecture level (Xie et al., 2025). These studies collectively indicate a shift from robot-only adaptation toward crop-robot compatibility. Nevertheless, they do not provide a continuous, spatially explicit metric that quantifies harvestability throughout a canopy, nor do they translate standard seasonal management actions into a canopy-scale optimisation problem for robotic harvesting.

The missing element is therefore not simply another harvesting algorithm, but a quantitative representation of the canopy as a robot-operating environment whose spatial suitability can be evaluated before harvest. Habitat suitability and species-distribution modelling provide a relevant conceptual analogue, because spatial suitability models formalise how multiple environmental predictors can be documented, combined, and reported as decision-relevant suitability estimates (Zurell et al., 2020).

Transposed to robotic harvesting, the canopy can be regarded as a three-dimensional operating habitat for the harvester, and each spatial position can be assigned a degree of robot-oriented harvestability according to the local conditions required for successful picking. In this study, that suitability is represented by a harvestability field, $\mathcal{H}(x)\colon \mathbb{R}^3 \to [0,1]$, composed of four normalised sub-fields: reachability $R(x)$, visibility $V(x)$, clearance $C(x)$, and detachability $D(x)$. A mass-weighted average of $\mathcal{H}(x)$ evaluated at fruit positions defines the Harvested Yield Index (HYI), which is used as a prior indicator of robotic picking potential rather than as a direct substitute for observed field success.

On this basis, the present study addresses the following research question: for a specified robotic kiwifruit harvester, can routine canopy-management decisions be optimised to improve robot-oriented harvestability while maintaining agronomically acceptable yield and economically feasible labour demand? Three working hypotheses follow from this question. First, H1 states that a multiplicative spatial harvestability field can provide an interpretable prior indicator of robotic harvesting feasibility by integrating reachability, visibility, clearance, and detachability. Second, H2 states that fruit placement and canopy structure generated by standard management actions produce measurable variation in HYI, even when total fruit number is controlled. Third, H3 states that multi-objective optimisation can identify Pareto-structured management strategies that expose trade-offs among harvestability, yield, and labour cost, thereby providing more useful decision-support information than a single unconstrained optimum.

To test these hypotheses, this paper introduces HarvestField, a decision-support framework for robot-oriented kiwifruit canopy management. The framework contributes to biosystems engineering in three respects. First, it formalises robot-oriented harvestability as a spatially explicit field over canopy volume, thereby converting local physical constraints on harvesting into a computable management metric. Second, it couples this field with a functional-structural kiwifruit model to simulate how seasonal management variables reshape canopy geometry and fruit distribution. Third, it embeds the resulting HYI within a multi-objective optimisation framework that simultaneously considers harvestability, yield, and labour cost, consistent with the broader role of digital agriculture in supporting data-informed trade-off analysis across economic, biological, and operational objectives (Basso & Antle, 2020).

The remainder of the paper is organised as follows. Section 2 describes the materials and methods, including the HarvestField architecture, the management-variable encoding, the formulation of the harvestability field and HYI, the functional-structural canopy simulation, the optimisation procedure, and the calibration and validation boundary. Section 3 reports the results of benchmark calibration, baseline comparison, Pareto-front analysis, equal-load harvestability-guided thinning, sensitivity analysis, and cross-robot evaluation. Section 4 discusses the mechanistic interpretation of the results, practical implications for kiwifruit canopy management, the decision-support value of the Pareto formulation, limitations of the simulation-based validation, and requirements for future orchard-scale testing. Section 5 summarises the principal conclusions and outlines the conditions under which the proposed framework may be extended to other trellised fruit systems.

Materials and methods

This section describes the computational research design, data sources, plant and robot models, optimisation procedure, validation boundary, and statistical analysis used to evaluate HarvestField. The section is written as a reproducible simulation protocol rather than as a purely theoretical derivation: each mathematical quantity is connected to an implemented component of the pipeline, and each experimental comparison is defined in terms of its controlled variables, inputs, and outputs.

Overall research design

The study was designed as a simulation-based decision-support investigation for robot-oriented kiwifruit canopy management. The experimental unit was one representative New Zealand-style kiwifruit bay at harvest time, simulated under alternative seasonal management decisions and evaluated with a specified robotic harvesting configuration. The central methodological question was whether a canopy-management trajectory that modifies fruit position, local canopy density, and cane geometry can be assessed before harvest by a spatially explicit harvestability metric and then optimised jointly with agronomic yield and labour cost. The resulting design comprised four linked stages: functional–structural canopy simulation, spatial harvestability field computation, multi-objective management optimisation, and benchmark-calibrated evaluation.

The framework therefore differs from a robot-only simulation study. The robot was not treated as the sole design object; rather, its reachability, perception, clearance, and detachment constraints were used to evaluate the orchard environment that management practices produce. This design was selected because the research question concerns management-mediated changes in the harvesting environment, not post-hoc measurement of robotic picking success after a physical trial. The overall architecture is summarised in Fig. 1.

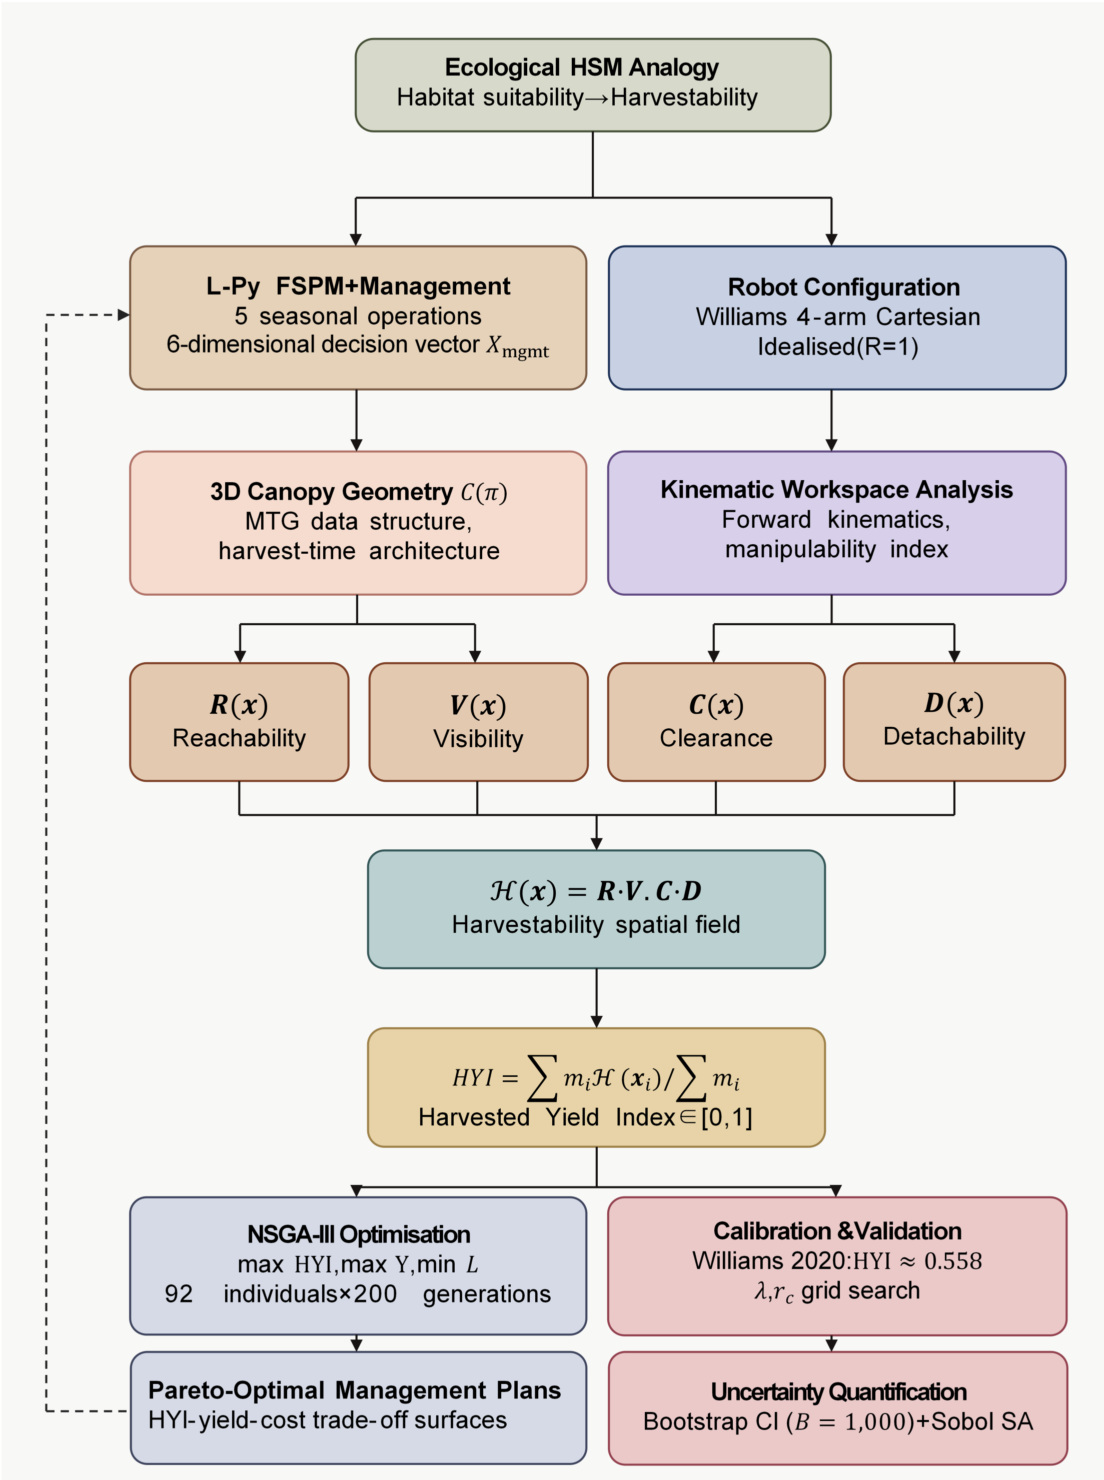

Fig. 1. Schematic overview of the HarvestField decision-support framework.

Management decision space and experimental unit

The management decision space was defined to represent actions that are routinely available to commercial kiwifruit growers and that can plausibly modify the three-dimensional distribution of fruit and vegetative structures before harvest. A management trajectory was encoded as a six-dimensional vector

$$\mathbf{u} = (s_c, n_c, \rho_b, n_f, n_t, \theta_c) \in \Omega, \qquad (1)$$

where s_c is cane spacing, n_c is the number of retained fruiting canes per bay side, ρ_b is bud retention density, n_f is the target fruit count per fruiting shoot, n_t is the maximum node count after summer tipping, θ_c is the cane training angle relative to the main leader, and Ω denotes the agronomically feasible decision space.

The variables were selected because they correspond to management interventions that directly affect the spatial distribution of fruit, leaves, and woody obstacles within the T-bar canopy.

Table 1. Management decision variables, commercial reference values, and optimisation bounds

Variable	Symbol	Unit	Reference value	Lower bound	Upper bound	Management operation
Cane spacing	s_c	m	0.30	0.20	0.60	Winter pruning
Retained fruiting canes per bay side	n_c	—	15	8	20	Winter pruning
Bud retention density	ρ_b	buds m^{-1}	20	10	65	Bud thinning
Target fruit count per fruiting shoot	n_f	—	6	3	10	Fruit thinning
Maximum node count after summer tipping	n_t	—	8	5	18	Summer tipping
Cane training angle	θ_c	°	90	45	135	Cane training

Each candidate vector $\mathbf{u}$ was translated into a phenologically ordered sequence of management actions and then simulated to harvest time. The simulation output for each candidate was a harvest-time canopy state

$$\mathcal{C}(\mathbf{u}) = \{\mathcal{G}, \mathcal{L}, \mathcal{B}, \mathcal{F}, \mathcal{P}\}, \qquad (2)$$

where $\mathcal{G}$ denotes the geometric support structure, $\mathcal{L}$ leaf elements, $\mathcal{B}$ branches and canes, $\mathcal{F}$ fruit objects, and $\mathcal{P}$ peduncle elements. Fruit i is represented by position $\mathbf{x}_i$ and mass m_i. The bay length was set to 5.0 m and the T-bar crossbar width to 1.5–1.8 m, corresponding to the geometric domain used for field computation.

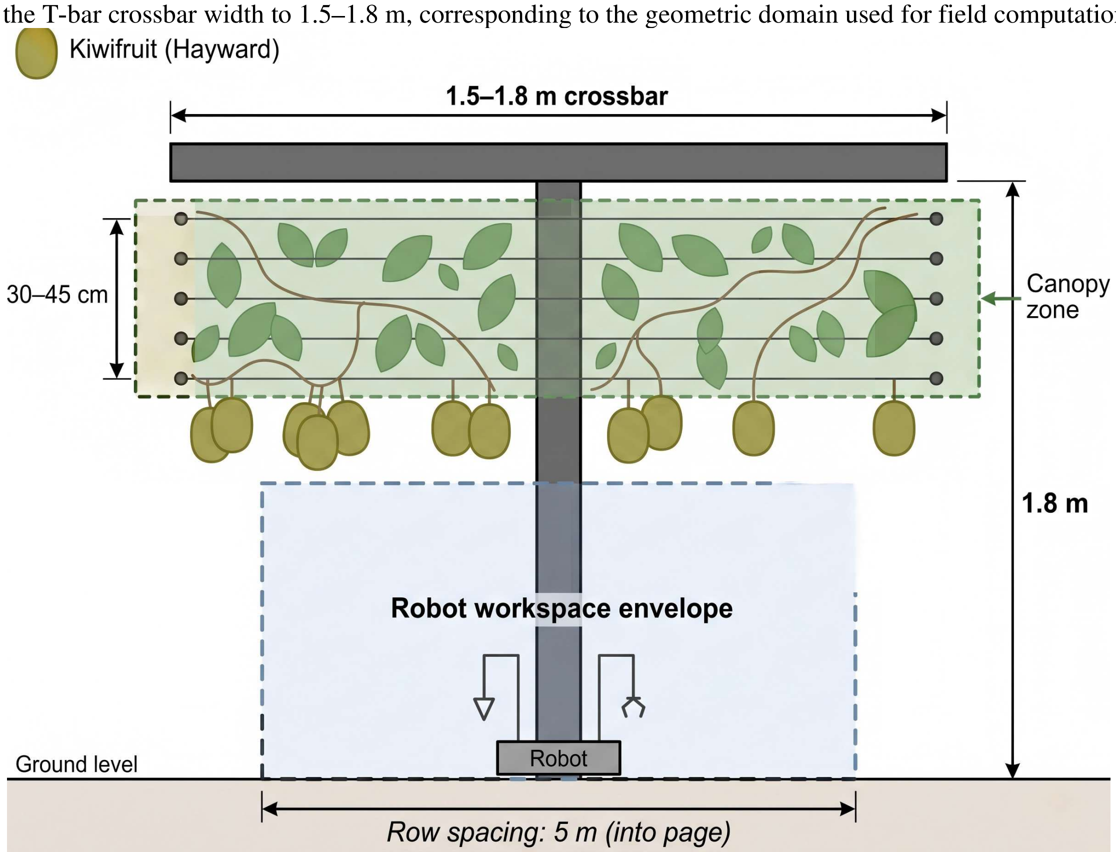

Fig. 2. Kiwifruit T-bar canopy geometry and field-computation domain.

Functional–structural plant simulation and data processing

Canopy development was simulated with an L-system-based functional–structural plant model (FSPM) implemented in L-Py, which couples procedural architectural development with Python-based data structures and plant-geometry processing (Boudon et al., 2012). The simulated species was *Actinidia deliciosa* cv. Hayward, and the model represented shoot emergence, internode elongation, budbreak, branching, fruit set, organ growth, and harvest-time geometry within a seasonal growth cycle. The canopy was stored as a multiscale tree graph (MTG), in which each structural unit carried geometric attributes, phenological state, and organ-type labels. This representation allowed management operations to be implemented as programmatic modifications to the plant graph rather than as purely statistical perturbations of fruit coordinates.

The seasonal management sequence comprised five operations. Winter pruning removed or retained cane substructures according to s_c and n_c. Bud thinning imposed the retained bud density ρ_b using a stochastic budbreak process modulated by position along the cane. Fruit thinning occurred after simulated flowering and fruit set and retained n_f fruits per fruiting shoot. In the HarvestField-guided condition, retention priority was determined by the local harvestability value $\mathcal{H}(\mathbf{x}_i)$, whereas uniform thinning retained fruits at approximately equal spacing along each shoot. Summer tipping truncated actively elongating shoots at n_t nodes, and cane training adjusted cane insertion angle to θ_c. These actions were implemented at their corresponding phenological stages in the Southern Hemisphere growing season, from winter dormancy to pre-harvest canopy maturity.

Light interception was computed using the Caribu nested-radiosity model, which estimates the distribution of intercepted radiation over a triangulated plant canopy while accounting for leaf spatial heterogeneity and multiple scattering (Chelle & Andrieu, 1998). The radiation interception efficiency $RIE(\mathbf{u})$ was used as the principal yield-related canopy output. The normalised yield proxy $Y(\mathbf{u})$ combined radiation interception with the FSPM carbon-allocation module, whereas labour cost $L(\mathbf{u})$ was modelled as a weighted sum of operation-specific management intensities. Both $Y(\mathbf{u})$ and $L(\mathbf{u})$ were min–max normalised within each optimisation run to place them on a common scale with HYI.

Spatial harvestability field

The core metric of HarvestField is the spatial harvestability field, a continuous scalar field that assigns to each position in the canopy a normalised estimate of its suitability for robotic fruit retrieval. Given robot configuration θ_r and canopy geometry $\mathcal{C}(\mathbf{u})$, the field is defined as

$$\mathcal{H}(\mathbf{x} \mid \theta_r, \mathcal{C}(\mathbf{u})): \mathbb{R}^3 \to [0,1]. \qquad (3)$$

In this implementation, harvestability is represented as the multiplicative composition of four necessary spatial conditions:

$$\mathcal{H}(\mathbf{x}) = R(\mathbf{x})\, V(\mathbf{x})\, C(\mathbf{x})\, D(\mathbf{x}), \qquad (4)$$

where $R(\mathbf{x})$, $V(\mathbf{x})$, $C(\mathbf{x})$, and $D(\mathbf{x})$ denote reachability, visibility, clearance, and detachability, respectively. The multiplicative aggregation was used because failure of any one necessary condition can render a fruit effectively unharvestable, even if the remaining conditions are favourable. The field was computed over a 50 mm voxel grid spanning the canopy volume, with $100 \times 36 \times 42 = 151{,}200$ voxels in the standard bay geometry.

Reachability sub-field

The reachability sub-field quantifies whether a fruit position lies within the dexterous workspace of the harvesting robot. For a robot with joint configuration $\mathbf{q} \in \mathcal{Q}$, forward kinematics $FK: \mathcal{Q} \to \mathbb{R}^3$, and Jacobian $J(\mathbf{q})$, the manipulability score was computed as

$$w(\mathbf{q}) = \sqrt{\det(J(\mathbf{q})J(\mathbf{q})^T)}. \qquad (5)$$

This score follows the standard manipulability formulation for robotic mechanisms (Yoshikawa, 1985). Reachability was evaluated by Monte Carlo sampling $N_q = 50{,}000$ collision-free joint configurations and projecting the resulting end-effector positions onto the voxel grid:

$$R(\mathbf{x}) = \left\{ \max_{\mathbf{q}' \in \mathcal{Q}_{\text{valid}}} \frac{w(\mathbf{q})}{\max\limits_{\mathbf{q}' \in \mathcal{Q}_{\text{valid}}} w(\mathbf{q}')} \cdot I[\|FK(q) - x\| \le \Delta/2] \right\}, \qquad (6)$$

where $\mathcal{Q}_{\text{valid}}$ is the collision-free configuration set, $\Delta = 50$ mm is the voxel side length, and $\mathbb{I}(\cdot)$ is the indicator function. Voxels not reached by any sampled configuration were assigned $R(\mathbf{x}) = 0$.

Visibility sub-field

The visibility sub-field estimates whether a fruit at position $\mathbf{x}$ can be observed by the robot's perception system. Ray casting was conducted from $N_v = 8$ candidate camera viewpoints, consisting of four longitudinal platform positions and two lateral viewing angles at each position. The visibility score was defined as

$$V(\mathbf{x}) = \frac{1}{N_v} \sum_{v=1}^{N_v} \{1 - O_v(\mathbf{x})\}, \qquad (7)$$

where $O_v(\mathbf{x}) \in [0,1]$ is the occlusion fraction along the ray from camera position $\mathbf{c}_v$ to $\mathbf{x}$. Leaf and branch meshes were treated as occluding structures, whereas fruit geometry was excluded from the occluder set to avoid double-counting self-occlusion.

Clearance sub-field

The clearance sub-field quantifies obstruction around a fruit position, particularly the local density of branches, canes, and trellis elements that could impede end-effector approach. Local obstacle density was computed as

$$\rho_{\mathrm{obs}}(\mathbf{x}; r_c) = \sum_{j \in \mathcal{M}} A_j \, \mathbb{I}\{\|\mathbf{v}_j - \mathbf{x}\| \le r_c\}, \qquad (8)$$

where $\mathcal{M}$ is the set of triangulated obstacle mesh elements, A_j is the surface area of element j, $\mathbf{v}_j$ is its centroid, and r_c is the clearance neighbourhood radius. Clearance was then modelled as an exponential decay:

$$C(\mathbf{x}) = \exp[-\lambda \, \rho_{\mathrm{obs}}(\mathbf{x}; r_c)], \qquad (9)$$

where $\lambda > 0$ controls the severity with which local obstacles reduce passability. The two parameters λ and r_c were treated as calibration parameters because they encode the relationship between local geometric clutter and practical end-effector access.

Detachability sub-field

The detachability sub-field captures whether a fruit can be mechanically separated without excessive interference from neighbouring fruit or insufficient gripper force. It was computed as

$$D(\mathbf{x}) = \sigma\left\{\frac{F_g - F_d(\mathbf{x})}{\tau_d}\right\} \min\left\{1, \frac{d_{\mathrm{nn}}(\mathbf{x})}{d_{\min}}\right\}, \qquad (10)$$

where F_g is the gripper's maximum applicable force, $F_d(\mathbf{x})$ is the peduncle detachment force, $\tau_d = 1.0$ N is a temperature parameter, $d_{\mathrm{nn}}(\mathbf{x})$ is the Euclidean distance to the nearest neighbouring fruit, $d_{\min} = 30$ mm is the inter-fruit interference threshold, and $\sigma(\cdot)$ is the logistic sigmoid. The detachment force was parameterised by peduncle angle ϕ as

$$F_d(\mathbf{x}) = \beta_0 + \beta_1 \cos\phi(\mathbf{x}), \quad \beta_0 = 3.2 \text{ N}, \quad \beta_1 = 1.8 \text{ N}. \qquad (11)$$

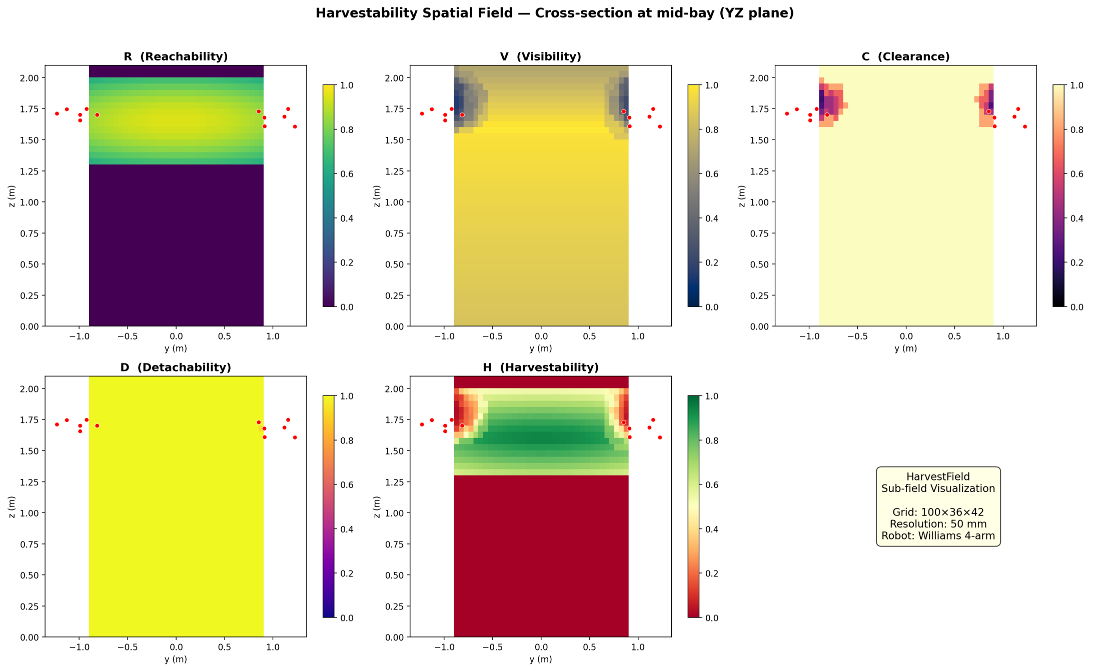

Fig. 3. Representative cross-sectional visualisation of the harvestability sub-fields and the composite field

Harvested Yield Index and optimisation objectives

The aggregate metric used to score a canopy configuration was the Harvested Yield Index (HYI). For N_f fruits with positions $\mathbf{x}_i$ and masses m_i, HYI was defined as the mass-weighted mean harvestability:

$$\mathrm{HYI}(\mathbf{u}; \theta_r) = \frac{\sum_{i=1}^{N_f} m_i \, \mathcal{H}\{\mathbf{x}_i \mid \theta_r, \mathcal{C}(\mathbf{u})\}}{\sum_{i=1}^{N_f} m_i}. \qquad (12)$$

HYI is a prior indicator: it is computed before a physical harvest and should be interpreted as a spatially informed estimate of harvesting potential rather than as a direct replacement for measured field success rate. This distinction is important because the four sub-fields are engineering descriptors of necessary harvesting conditions, not a complete stochastic model of robot control, fruit damage, or operator-level decision making.

The constrained management optimisation problem was formulated as

$$\max_{\mathbf{u}\in\Omega}\{\mathrm{HYI}(\mathbf{u};\theta_r), Y(\mathbf{u}), -L(\mathbf{u})\}, \qquad (13)$$

subject to

$$Y(\mathbf{u}) \geq Y_{\min},\ \ LAI(\mathbf{u}) \geq LAI_{\min}. \qquad (14)$$

Here $Y(\mathbf{u})$ is the normalised yield proxy, $L(\mathbf{u})$ is the normalised labour-cost proxy, $Y_{\min}$ enforces a minimum commercially meaningful yield, and $LAI_{\min}$ prevents solutions that improve robot access only by producing unrealistically sparse canopies. This formulation was designed to produce decision-support trade-offs rather than a single universal management prescription.

Multi-objective optimisation and implementation

The three-objective optimisation problem was solved with NSGA-III, a reference-direction evolutionary algorithm developed for maintaining well-distributed approximations to Pareto fronts in many-objective optimisation (Deb & Jain, 2014). Reference directions were generated using the Das–Dennis simplex-lattice construction with $H = 12$ divisions and $M = 3$ objectives, yielding 91 directions. The population size was set to 92, and each optimisation run was executed for 200 generations. Simulated binary crossover was applied with distribution index $\eta_c = 30$, and polynomial mutation was applied with distribution index $\eta_m = 20$. The implementation used pymoo v0.6.1, which provides reproducible evolutionary-optimisation operators, reference-direction generation, and Pareto-front post-processing (Blank & Deb, 2020).

Each candidate solution was evaluated by the same deterministic pipeline: (i) decode $\mathbf{u}$ into management actions; (ii) simulate canopy development to harvest time; (iii) compute radiation interception and yield proxy; (iv) compute the harvestability field at 50 mm resolution; (v) evaluate HYI at fruit positions; and (vi) compute labour cost. Approximate wall-clock times were 5.0 s for FSPM simulation, 0.3 s for light-interception computation, 2.0 s for field computation, and negligible time for analytical labour calculation. With 18,400 evaluations per run, parallel execution on 32 CPU cores required approximately 1.2 h per optimisation run.

Table 2. Principal computational settings used in the HarvestField evaluation pipeline

Setting	Symbol	Value
Canopy bay length	L_b	5.0 m
T-bar crossbar width	W_b	1.5–1.8 m
Voxel resolution	Δ	50 mm
Voxel grid size	—	100 × 36 × 42 = 151,200 voxels
Reachability samples	N_q	50,000 joint configurations
Visibility viewpoints	N_v	8 viewpoints
NSGA-III population size	N_{pop}	92
NSGA-III generations	G	200
Total evaluations per optimisation	N_{eval}	18,400
Primary implementation	—	Python 3.10; L-Py v3.14; OpenAlea/Caribu; pymoo v0.6.1

For robot-specific analysis, the field $\mathcal{H}(\mathbf{x} \mid \theta_r, \mathcal{C})$ was recomputed after changing the robotic configuration. Two robot conditions were used: a reconstructed Williams-type four-arm Cartesian harvester and an idealised upper-bound robot for which $R(\mathbf{x}) = 1$ throughout the canopy domain. This comparison tested whether management recommendations were primarily agronomic or dependent on robot kinematics.

Benchmark calibration and validation boundary

The two clearance-related parameters, λ and r_c, were calibrated against the large-scale kiwifruit robotic-harvesting benchmark reported by Williams et al. (2020), which evaluated an improved robotic kiwifruit harvester on more than 12,000 fruits and reported an overall picking success rate of 0.558. The calibration objective was

$$(\lambda^*, r_c^*) = arg \min_{\lambda\in[0.5,10.0],\, r_c\in[0.05,0.25]} \left|\mathrm{HYI}_{\lambda,r_c} - 0.558\right|. \qquad (15)$$

Calibration was considered acceptable if the absolute deviation from the benchmark success rate was less than 0.10. To represent uncertainty due to stochastic FSPM realisations, a non-parametric bootstrap with $B = 1{,}000$ replications was performed after parameter selection. Both percentile and bias-corrected and accelerated confidence intervals were computed.

This benchmark comparison defines the validation boundary of the study. HYI was calibrated to a published field benchmark and then used to evaluate relative management effects under controlled simulation conditions. The analysis should therefore be interpreted as benchmark-calibrated plausibility assessment and internal decision-support validation, rather than as prospective orchard-scale validation. Independent orchards, seasons, cultivars, and robotic platforms remain necessary for operational validation.

Baseline, ablation, and controlled thinning experiments

Six baseline strategies were defined to isolate the effect of robot-oriented spatial management from yield maximisation, random search, single-objective optimisation, and homogeneous-field assumptions. Each stochastic baseline was evaluated over $R = 10$ independent random seeds. The baseline definitions are provided in Table 3.

Table 3. Baseline strategies used to evaluate the contribution of the HarvestField framework

Code	Strategy	Definition and purpose
B1	Standard management	Commercial reference values in Table 1; status-quo canopy-management baseline.
B2	Yield-only optimisation	Optimises $Y(\mathrm{u})$ without HYI or labour-cost objectives.
B3	Random search	Uniform random sampling of Ω using the same 18,400-evaluation budget as NSGA-III.
B4	Single-objective HYI	Maximises HYI subject to yield and LAI constraints using a single-objective genetic algorithm.
B5	Uniform field assumption	Sets $\mathcal{H}(\mathrm{x}) = 1$ during optimisation and evaluates the resulting plan under the true heterogeneous field.
B6	HarvestField-optimised	Full NSGA-III optimisation of HYI, yield, and labour cost over the six-dimensional management space.

Four ablation experiments were then conducted. A1 removed the visibility sub-field by evaluating $\mathcal{H} = RCD$. A2 removed the clearance sub-field by evaluating $\mathcal{H} = RVD$. A3 replaced the Caribu radiation model with a Beer–Lambert extinction model to test whether spatially explicit radiosity computation affected the yield proxy. A4 tested the central management mechanism of HarvestField: harvestability-guided fruit thinning versus uniform fruit thinning under equal crop load. In A4, total fruit count, canopy geometry before thinning, robot configuration, random seed, FSPM parameters, light model, and field resolution were held constant. The only experimental difference was the spatial distribution of retained fruits. Three crop-load levels were examined: 4, 6, and 8 fruits per fruiting shoot.

Sensitivity analysis and statistical methods

Global sensitivity analysis was performed to identify which uncertain parameters most strongly influenced HYI. The analysis used a variance-based Sobol sensitivity workflow implemented through SALib-style structured model evaluations, following current recommendations for accessible and interpretable global sensitivity analysis (Saltelli, 2002; Iwanaga et al., 2022). Six parameters were varied: λ, r_c, F_g, fruit set rate, budbreak rate, and individual leaf area. With base sample size $N = 512$ and $p = 6$ parameters, the analysis required

$$N_{\mathrm{SA}} = N(2p + 2) = 512(2 \times 6 + 2) = 7{,}168 \qquad (16)$$

model evaluations. Parameters with total-order sensitivity $S_T > 0.10$ were treated as influential because they explained a non-negligible fraction of HYI variance after accounting for interaction effects.

Because baseline and ablation comparisons used $R = 10$ independent replications and did not assume normality, pairwise differences were analysed with the Wilcoxon signed-rank test (Wilcoxon, 1945). The family-wise error rate was controlled with Bonferroni correction across the eight pre-specified baseline comparisons:

$$\alpha^{*} = \frac{0.05}{8} = 0.00625. \qquad (17)$$

Raw p-values are reported, but statistical significance is interpreted against \(\alpha^{\ast}\). Effect magnitude was reported with Cliff's delta, an ordinal dominance statistic suitable for non-parametric comparisons (Cliff, 1993):

$$\frac{\delta = \#(x_i > y_j) - \#(x_i < y_j)\}}{n_x n_y}. \qquad (18)$$

Values of $|\delta| < 0.147$, $0.147 \leq |\delta| < 0.33$, $0.33 \leq |\delta| < 0.474$, and $|\delta| \geq 0.474$ were interpreted as negligible, small, medium, and large effects, respectively. This statistical protocol was selected to ensure that conclusions did not depend on distributional assumptions inappropriate for small stochastic-simulation samples.

Reproducibility and software environment

All stochastic components used explicit random seeds. Independent replications were generated by varying the seed while holding all model parameters, software versions, and evaluation settings fixed. The computational environment used Python 3.10, L-Py v3.14, OpenAlea/Caribu for plant geometry and light computation, pymoo v0.6.1 for multi-objective optimisation, and standard scientific Python packages for numerical analysis and tabulation. The primary outputs retained for reproducibility were candidate management vectors, simulated canopy geometries, fruit positions and masses, voxel-level field arrays, optimisation logs, bootstrap samples, sensitivity-index estimates, and statistical-test summaries. Processed data and configuration files are intended to be provided as online supplementary resources accompanying the manuscript.

Results

The results are reported in the same sequence as the evaluation design: calibration of the Harvested Yield Index (HYI), baseline comparisons, Pareto-front outcomes, equal-load fruit-thinning experiments, sub-field ablations, Sobol sensitivity analysis, and cross-robot comparison. This section reports numerical outcomes and statistical contrasts only; mechanistic explanations and practical implications are reserved for the Discussion.

Calibration and uncertainty of the Harvested Yield Index

The grid search over the two clearance-related parameters identified $\lambda = 1.0\ \mathrm{m}^{-1}$ and $r_c = 0.133$ m as the minimum-error calibration point. At this setting, the standard-management Hayward canopy evaluated with the Williams-type four-arm Cartesian robot produced HYI $= 0.562$. The absolute difference from the benchmark picking-success value of 0.558 was 0.004, corresponding to a relative deviation of 0.7% and remaining within the pre-specified acceptance threshold of 0.10. The calibration heatmap in Fig. 4a shows a narrower error gradient along the r_c dimension than along the λ dimension.

Bootstrap resampling with $B = 1{,}000$ stochastic FSPM realisations yielded a median HYI of 0.581 and a mean HYI of 0.582. The bootstrap standard deviation was 0.022. The percentile 95% confidence interval was [0.540,0.629], and the bias-corrected and accelerated (BCa) 95% confidence interval was [0.542,0.630]. Both intervals contained the benchmark reference value of 0.558, as shown in Fig. 4b.

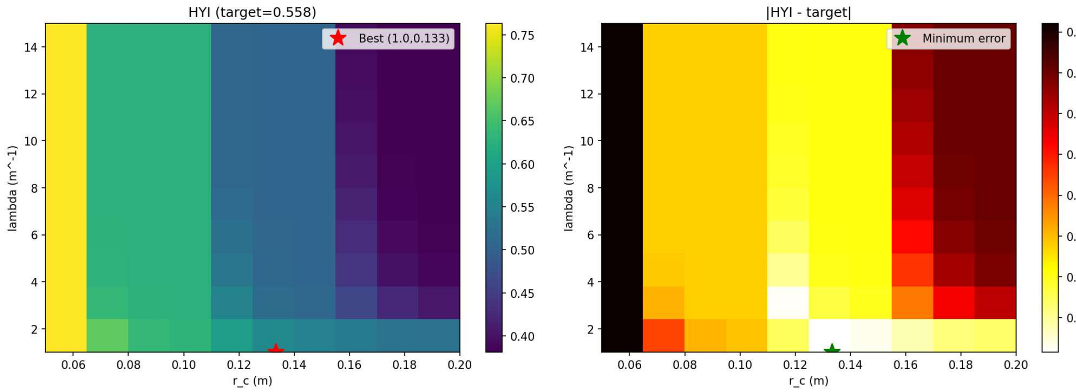

(a)

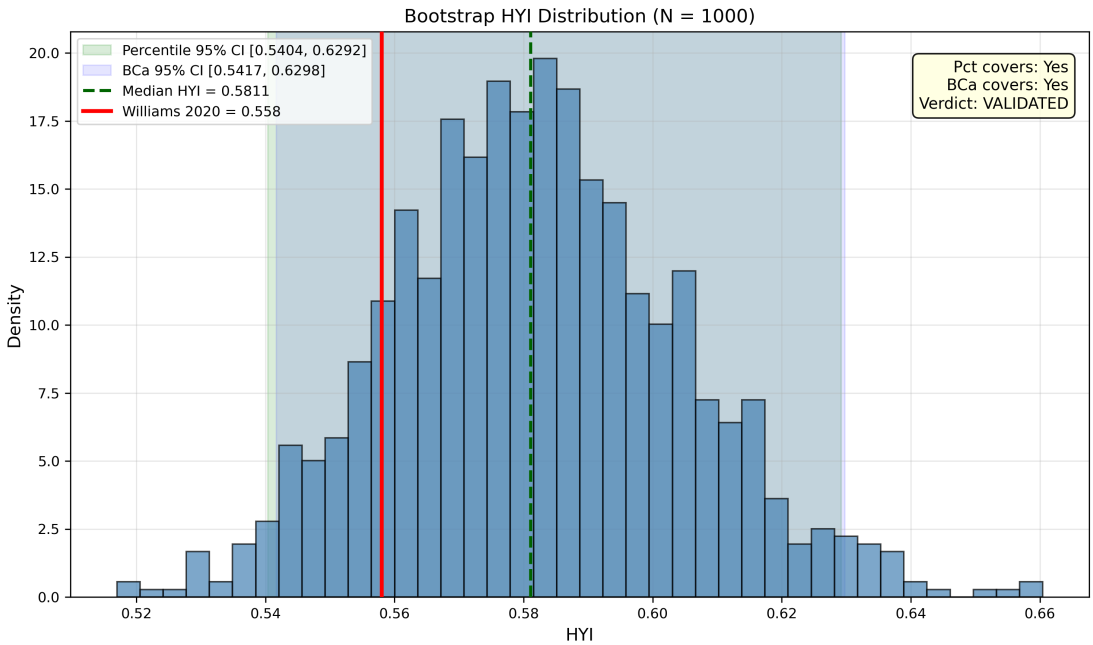

(b)

Fig. 4 Calibration and uncertainty analysis of the Harvested Yield Index. (a) Grid-search error surface for the calibration parameters; (b) Bootstrap distribution of HYI across stochastic FSPM realisations

Baseline comparison

The HYI values obtained under the six management strategies are presented in Fig. 5a, and the pre-specified pairwise statistical contrasts are reported in Table 4. HarvestField-optimised management (B6) produced the largest mean HYI among the tested strategies (0.642). Random search (B3) produced a mean HYI of 0.637, and standard commercial management (B1) produced a mean HYI of 0.587. The uniform-field assumption baseline (B5) was numerically indistinguishable from B1 at three decimal places, whereas the single-objective HYI baseline (B4) and the yield-only optimisation baseline (B2) yielded mean HYI values of 0.552 and 0.378, respectively.

The planned contrast between B6 and B1 showed a positive difference of $\Delta\mathrm{HYI} = +0.055$ with a large ordinal effect size (Cliff's $\delta = 0.94$). The contrast between B6 and B3 was not statistically significant ($\Delta\mathrm{HYI} = +0.005$, $p = 0.461$). The contrast between B1 and B2 showed $\Delta\mathrm{HYI} = +0.209$, and the contrast between B1 and B4 showed $\Delta\mathrm{HYI} = +0.035$. The contrast between B1 and B5 yielded $\Delta\mathrm{HYI} < 0.001$. All p-values in Table 4 are raw Wilcoxon signed-rank p-values, evaluated against the Bonferroni-adjusted threshold $\alpha^* = 0.00625$ for eight planned comparisons.

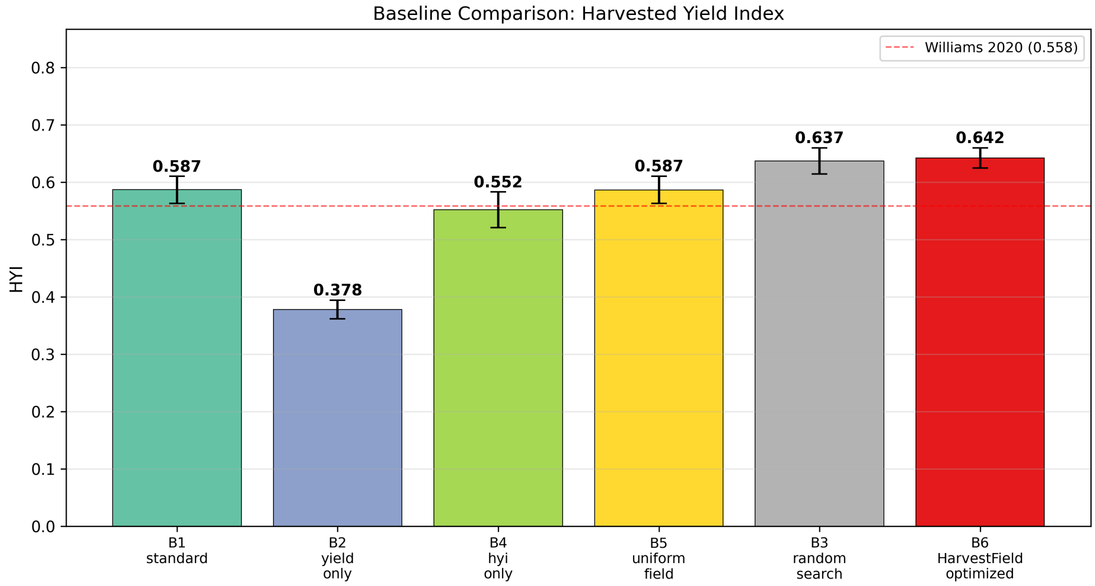

(a)

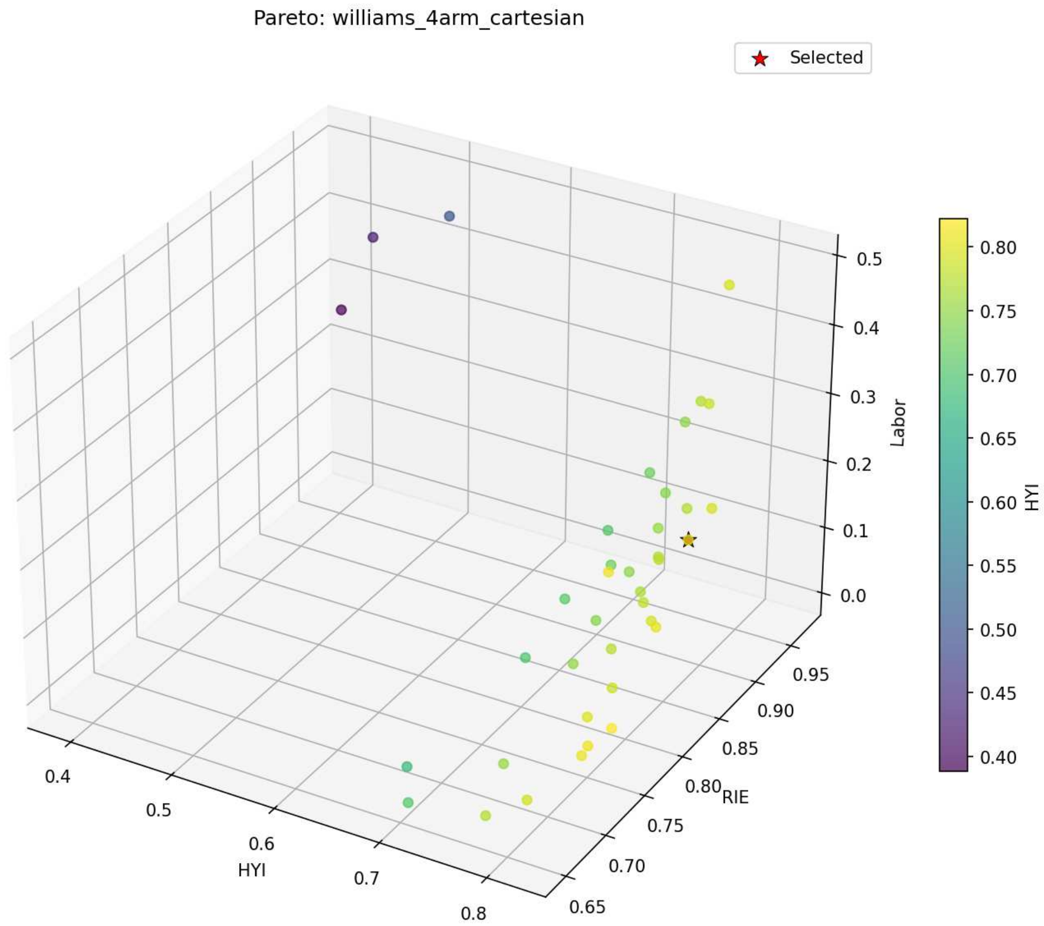

(b)

Fig. 5 Baseline comparison and Pareto-front outcomes. (a) HYI comparison across the six baseline management strategies; (b) Three-objective Pareto front under the Williams-type robot configuration

Table 4. Pre-specified statistical contrasts for baseline HYI comparisons

Contrast	Mean HYI A	Mean HYI B	ΔHYI	Cliff's δ	Raw p	Significance
B6 vs B1	0.642	0.587	+0.055	0.94	0.00195	Yes
B6 vs B3	0.642	0.637	+0.005	0.22	0.461	No
B1 vs B2	0.587	0.378	+0.209	1.00	0.00195	Yes
B1 vs B4	0.587	0.552	+0.035	0.62	0.00195	Yes
B1 vs B5	0.587	0.587	<0.001	0.00	1.000	No
B3 vs B1	0.637	0.587	+0.050	0.88	0.00195	Yes
B3 vs B4	0.637	0.552	+0.085	1.00	0.00195	Yes
B3 vs B5	0.637	0.587	+0.050	0.90	0.00195	Yes

Note. Mean HYI A and Mean HYI B denote the means of the first and second conditions in each contrast. ΔHYI is computed as A minus B. Raw p-values are from the Wilcoxon signed-rank test. Significance is assessed against $\alpha^* = 0.05/8 = 0.00625$. Cliff's δ is reported as the non-parametric effect size.

Pareto-front outcomes under the Williams-type robot configuration

The NSGA-III run for the Williams-type four-arm Cartesian robot produced a three-objective Pareto front spanning HYI, radiation interception efficiency (RIE), and normalised labour cost (Fig. 5b). The final hypervolume indicator was $\mathrm{HV} = 3.521$. The hypervolume trajectory reached a near-plateau by generation 120, after which the remaining increase was less than 2% of the final value.

The selected knee-point solution on this Pareto front yielded $\mathrm{HYI} = 0.786$, $\mathrm{RIE} = 0.900$, and $L_{\mathrm{norm}} = 0.149$. The corresponding management vector was $s_c = 0.499\ \mathrm{m}$, $n_c = 15$ retained fruiting canes per bay side, $\rho_b = 56.6\ \mathrm{buds\ m^{-1}}$, $n_f = 8$ fruits per fruiting shoot, $n_t = 16$ nodes after summer tipping, and $\theta_c = 89.6^{\circ}$.

Equal-load harvestability-guided fruit thinning

The equal-quantity-different-placement thinning experiment compared $\mathcal{H}$-guided fruit retention with spatially uniform fruit retention at matched crop loads. Results are shown in Fig. 6a. At 4 fruits per cane, $\mathcal{H}$-guided thinning produced a mean HYI of 0.877, whereas uniform thinning produced a mean HYI of 0.576. The mean paired difference was $\Delta\text{HYI} = +0.301$, with a 95% confidence interval of [0.286,0.314].

At 6 fruits per cane, $\mathcal{H}$-guided thinning produced a mean HYI of 0.851, and uniform thinning produced a mean HYI of 0.605. The mean paired difference was$\Delta\text{HYI} = +0.246$, with a 95% confidence interval of [0.236,0.258]. At 8 fruits per cane, the corresponding means were 0.716 and 0.639, with $\Delta\text{HYI} = +0.077$ and a 95% confidence interval of [0.073,0.082]. All three contrasts had Wilcoxon $p = 0.00195$ and Cliff's $\delta = 1.00$, evaluated against the Bonferroni-adjusted threshold $\alpha^* = 0.0167$ for the three thinning levels.

For the 4- and 6-fruits-per-cane treatments, paired guided and uniform simulations retained exactly 60 and 90 fruits, respectively, in every replication. For the 8-fruits-per-cane treatment, retained fruit counts ranged from 117 to 120 because of stochastic fruit set, but the guided and uniform treatments retained identical fruit counts within each paired seed.

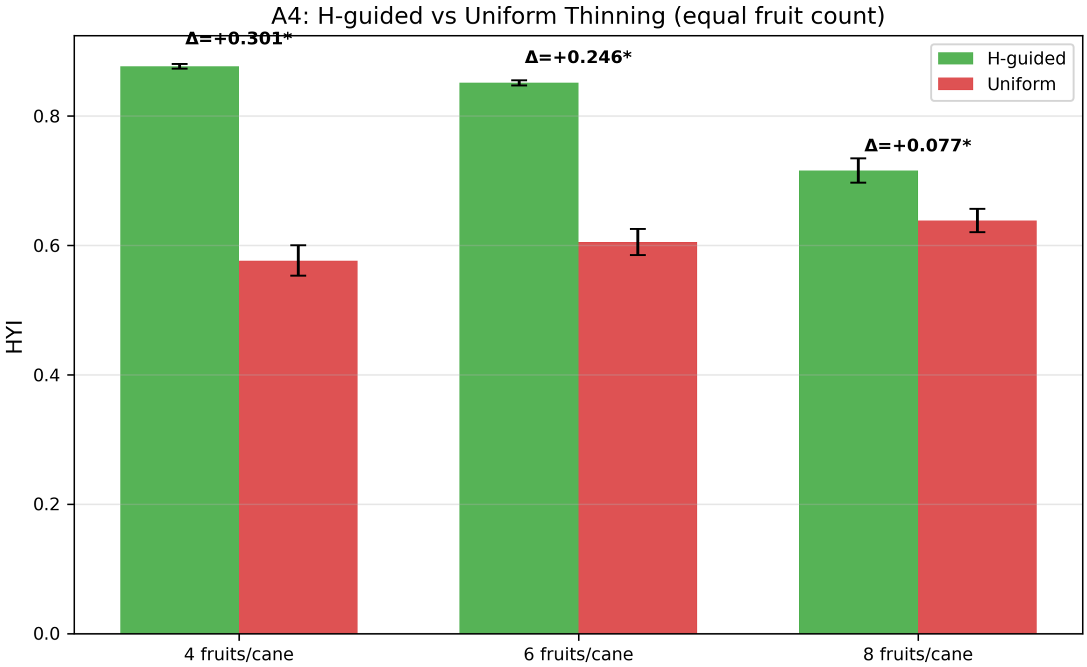

(a)

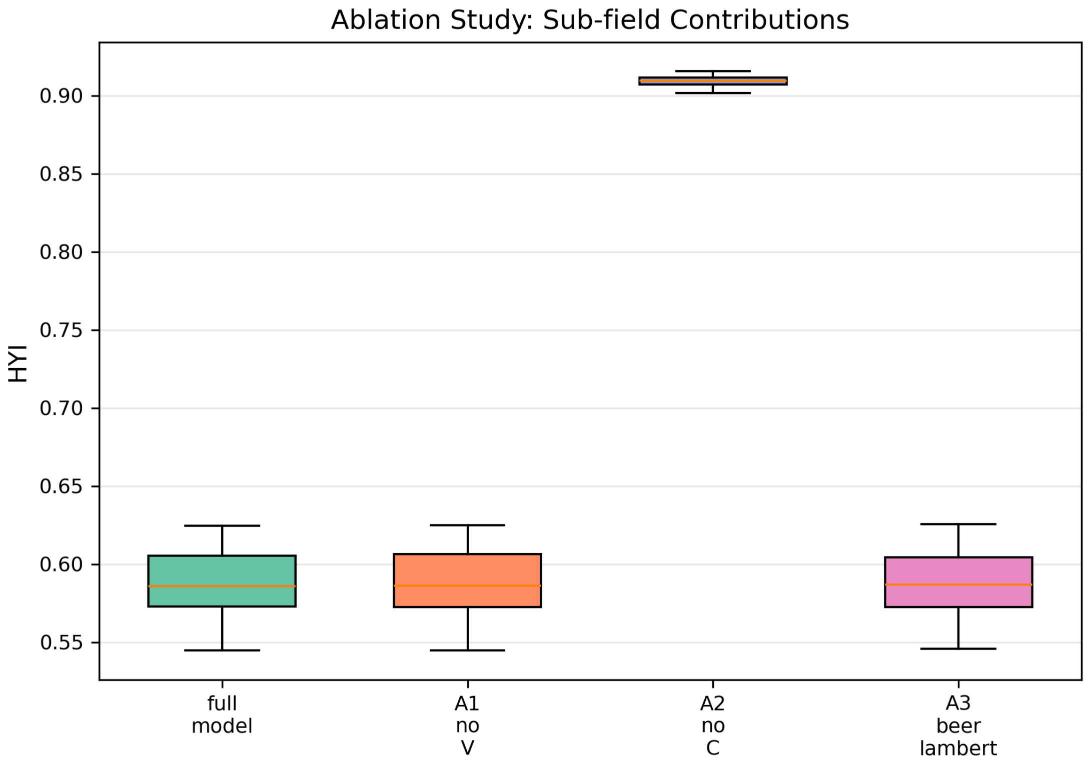

(b)

Fig. 6 Equal-load thinning and sub-field ablation results. (a) Comparison between $\mathcal{H}$-guided thinning and uniform thinning at equal fruit loads; (b) Effects of sub-field ablation on HYI

Sub-field ablation results

The sub-field ablation experiment is summarised in Fig. 6b. The full field model yielded a mean HYI of 0.587 with a standard deviation of 0.023 over ten replications. Removing the visibility sub-field (A1; $\mathcal{H} = R \cdot C \cdot D$) yielded a mean HYI of 0.587 with a standard deviation of 0.024. Replacing the Caribu radiation model with the Beer–Lambert approximation (A3) yielded a mean HYI of 0.587 with a standard deviation of 0.024.

Removing the clearance sub-field (A2; $\mathcal{H} = R \cdot V \cdot D$) yielded a mean HYI of 0.909 with a standard deviation of 0.004. The increase relative to the full model was $\Delta\text{HYI} = +0.323$. The distribution of A2 values ranged from 0.901 to 0.916, whereas the full-model distribution ranged from 0.545 to 0.625.

Sobol sensitivity analysis

The Sobol global sensitivity analysis comprised $N = 512$ base samples, six uncertain parameters, and 7,168 total model evaluations. The resulting first-order and total-order indices for HYI are shown in Fig. 7a and summarised in Table 5. The clearance neighbourhood radius r_c had the largest first-order index ($S_1 = 0.541 \pm 0.083$) and total-order index ($S_T = 0.576 \pm 0.071$). Gripper force F_g ranked second, with $S_1 = 0.268 \pm 0.077$ and $S_T = 0.291 \pm 0.047$. Fruit set rate ranked third, with $S_1 = 0.121 \pm 0.053$ and $S_T = 0.170 \pm 0.027$.

The clearance decay coefficient λ had $S_1 = 0.008 \pm 0.014$ and $S_T = 0.013 \pm 0.005$. Budbreak rate and individual leaf area each contributed less than 0.001 to the total-order index. The largest second-order interaction was observed between r_c and fruit set rate, with $S_2 = 0.034$; all other second-order terms were below 0.035 in absolute value.

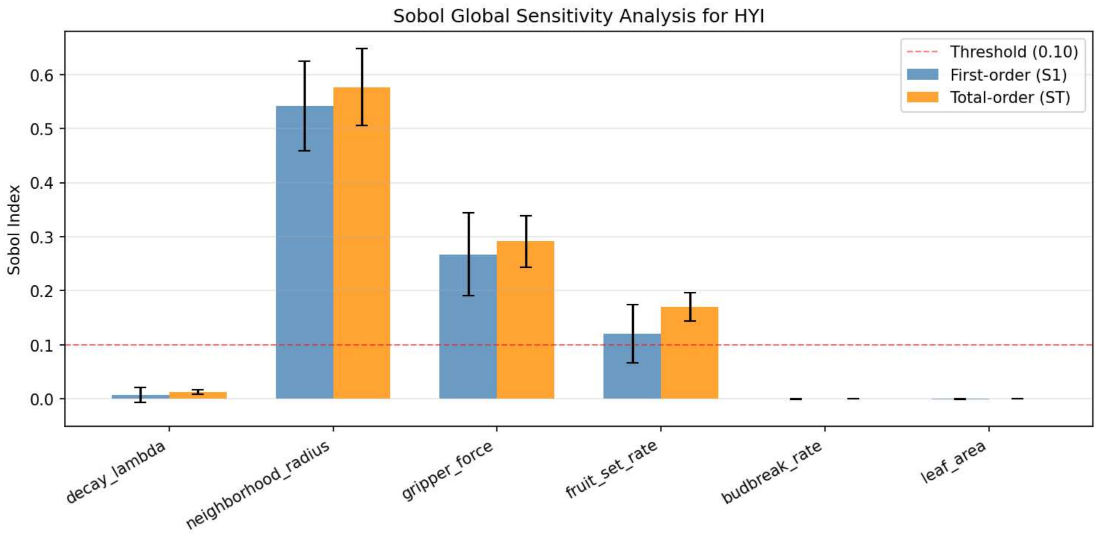

(a)

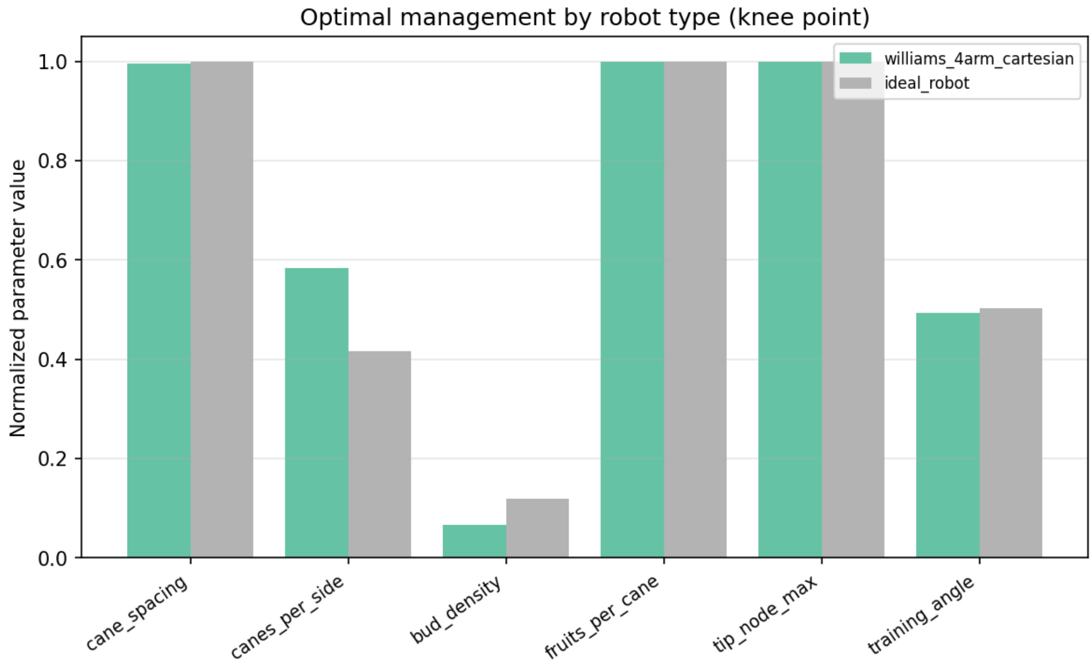

(b)

Fig. 7 Sensitivity analysis and cross-robot comparison. (a) Sobol first-order and total-order sensitivity indices for HYI; (b) Normalised comparison of Pareto-knee management parameters between robot configurations

Cross-robot comparison

Parallel optimisation under an idealised robot configuration produced a Pareto front with HV = 3.677, compared with HV = 3.521 for the Williams-type four-arm Cartesian robot. This difference corresponds to a 4.4% increase in hypervolume. The knee-point outcomes and management parameters for the two robot configurations are listed in Table 6 and shown in Fig. 7b.

The Williams-type robot knee point had HYI = 0.786, RIE = 0.900, and $L_{\text{norm}} = 0.149$. The idealised-robot knee point had HYI = 0.854, RIE = 0.891, and $L_{\text{norm}} = 0.114$. Both configurations selected $n_f = 8$, $n_t = 16$, and a training angle close to 90°. The Williams-type configuration selected $n_c = 15$ and $\rho_b = 56.6$ buds m^{-1}, whereas the idealised configuration selected $n_c = 13$ and $\rho_b = 61.9$ buds m^{-1}.

Table 5. Sobol sensitivity indices for HYI

Parameter	S_1	S_1 95% CI half-width	S_T	S_T 95% CI half-width
Clearance neighbourhood radius, r_c	0.541	±0.083	0.576	±0.071
Gripper force, F_g	0.268	±0.077	0.291	±0.047

Fruit set rate	0.121	±0.053	0.170	±0.027
Clearance decay coefficient, λ	0.008	±0.014	0.013	±0.005
Budbreak rate	≈0	±0.001	≈0	≈0
Individual leaf area	≈0	±0.001	≈0	≈0

Note. S_1 denotes the first-order index and S_T denotes the total-order index. Confidence intervals are bootstrap-derived half-widths from the sensitivity analysis outputs.

Table 6. Pareto-knee outcomes and management parameters for the cross-robot comparison

Outcome or parameter	Williams-type four-arm robot	Idealised robot
Hypervolume,HV	3.521	3.677
Knee-point HYI	0.786	0.854
Knee-point RIE	0.900	0.891
Normalised labour cost, L_{norm}	0.149	0.114
Cane spacing, s_c (m)	0.499	0.500
Retained canes per bay side, n_c	15	13
Bud retention density, ρ_b (buds m^{-1})	56.6	61.9
Fruits per fruiting shoot, n_f	8	8
Tipping node count, n_t	16	16
Training angle, θ_c (°)	89.6	90.2

Note. The Williams-type configuration represents the calibrated four-arm Cartesian harvesting platform. The idealised robot is evaluated with reachability set to unity throughout the computational domain.

Discussion

The results indicate that robot-oriented canopy management can be formulated as a spatial decision problem, rather than being treated only as a post-harvest assessment of robot performance. The calibrated Harvested Yield Index (HYI) reproduced the benchmark picking-success value with a small absolute deviation, whereas the baseline, ablation, sensitivity, and cross-robot analyses showed that harvestability within a kiwifruit canopy is spatially heterogeneous and operationally consequential. The following discussion interprets these findings in relation to the physical structure of the harvestability field, prior work on agricultural robotics and digital decision support, and the practical conditions under which HarvestField could be evaluated in commercial orchards.

Mechanistic interpretation of spatial harvestability

The main mechanistic implication of the results is that the composite field $\mathcal{H}(\mathbf{x}) = R(\mathbf{x})V(\mathbf{x})C(\mathbf{x})D(\mathbf{x})$ provides a compact representation of the serial constraints that determine whether an individual fruit can be retrieved by a robotic harvester. In a physical picking sequence, failure in reachability, visibility, approach clearance, or detachability is sufficient to prevent successful retrieval, even if the remaining conditions are favourable. The multiplicative formulation therefore does not merely aggregate four descriptors; it encodes the conjunctive structure of the harvesting task. This interpretation is consistent with recent field evaluations of multi-arm robotic fruit harvesters, which indicate that coordinated perception, planning, and manipulation remain constrained by dense orchard structure and access limitations (Lammers et al., 2024).

The ablation and sensitivity results further specify which spatial constraint dominated under the simulated T-bar kiwifruit geometry. Removing the clearance term produced an inflated HYI, and the Sobol analysis identified the clearance neighbourhood radius r_c as the strongest contributor to HYI variance. These results indicate that the limiting condition was not only whether fruit positions were inside the nominal manipulator workspace, but whether the local geometry around each fruit allowed an unobstructed approach by the end-effector. The limited effect of removing visibility should be interpreted within the below-canopy viewing geometry of kiwifruit: many suspended fruits remained observable from at least one upward-facing viewpoint, whereas branch, trellis, and fruit-cluster obstruction remained spatially variable. This finding extends reports from greenhouse harvesting systems in which improved sensing and gripper design did not remove all performance constraints imposed by plant architecture (Arad et al., 2020).

The equal-quantity fruit-thinning experiment provides direct evidence that harvestability is a management-relevant spatial property. Because total fruit count, canopy realisation, robot configuration, random seed, light model, and field resolution were controlled within each paired comparison, the observed HYI differences between $\mathcal{H}$-guided and uniform thinning can be attributed to fruit placement rather than to crop-load variation. The decreasing advantage of $\mathcal{H}$-guided thinning as fruit load increased is consistent with a saturation effect: low crop loads allow retained fruit to be concentrated in high-harvestability regions, whereas high crop loads force retention of a larger proportion of candidate fruit positions, including positions with lower $\mathcal{H}$. Thus, the experiment converts the harvestability field from a diagnostic descriptor into an agronomically interpretable principle: for robotic harvesting, fruit position is an optimisation variable, not merely a consequence of yield management.

Relationship to existing literature and theoretical contribution

Most agricultural-robotics research has sought to improve harvesting through robot-side capabilities, including perception, guidance, manipulation, and system integration. Reviews of field agricultural robots have emphasised that biological environments are unstructured, variable, and difficult to standardise (Yang et al., 2023; Liu et al., 2024). HarvestField addresses the same problem from the complementary direction: it treats the canopy as a managed biological structure whose spatial configuration can be deliberately altered before harvest. This perspective does not reduce the importance of machine vision or robot control, both of which remain central technological bottlenecks in harvesting robotics (Zhou et al., 2022; Droukas et al., 2023). Rather, it shows that the return on robot-side improvements can be constrained by canopy states that are physically unfavourable for access, approach, or detachment.

The baseline results support this interpretation. Yield-only optimisation reduced HYI relative to standard management, whereas HarvestField-optimised management increased HYI without requiring a change in robot hardware. This divergence indicates that agronomic productivity and robotic accessibility are not automatically aligned objectives. The finding therefore challenges a sequential design logic in which canopy management is first optimised for yield and robots are subsequently expected to adapt to the resulting structure. Instead, it supports a joint decision logic in which yield, labour cost, and robot-oriented harvestability are evaluated within a common optimisation space.

The theoretical contribution of the study is threefold. First, it reframes harvestability as a prior, continuous spatial property of the crop environment, expressed as $\mathcal{H}: \mathbb{R}^3 \to [0,1]$, rather than as a posterior statistic observed after a harvesting trial. Second, the HYI translates position-level harvestability into a canopy-scale indicator that can be used before harvest to compare management trajectories. Third, the framework establishes canopy management as an intermediate design layer between plant genotype and robot hardware. This contribution differs from purely robot-centred optimisation because it preserves existing cultivars and robotic platforms while enabling co-adaptation through pruning, thinning, tipping, and cane training decisions. The resulting decision-support logic is consistent with broader digital-agriculture perspectives in which data-driven models are used to balance economic, biological, and operational objectives rather than optimise a single technical metric in isolation (Basso & Antle, 2020).

Practical implications for orchard management and robotic-harvesting services

For orchard management, the immediate implication is that robot readiness should be assessed at the scale of spatial canopy structure rather than at the level of average crop load alone. The sensitivity results suggest that practices preserving approach corridors around fruit clusters are likely to be particularly relevant for below-canopy kiwifruit harvesting. In practical terms, cane spacing, retained cane number, bud density, fruit-load targets, tipping intensity, and cane angle should be interpreted not only as determinants of yield and labour demand, but also as variables that shape local reachability, visibility, clearance, and detachability. HarvestField therefore adds a robot-oriented criterion to management operations that are already familiar to growers, rather than proposing a disruptive change in cultivar or trellis system.

For robotic-harvesting service providers, the framework offers a pre-harvest diagnostic workflow. A service provider could reconstruct or simulate a canopy block, compute $\mathcal{H}$, identify regions with low clearance or reachability, and communicate management recommendations before the harvesting window. This would shift robotic harvesting from a deployment-only service to an integrated advisory process in which orchard preparation and robotic operation are coordinated. Such a workflow is consistent with the broader transition from isolated agricultural machines to digital farming systems that combine sensing, simulation, and operational decision support (Basso & Antle, 2020; Liu et al., 2024).

For equipment designers, the cross-robot comparison shows that canopy-management recommendations are conditional on robot configuration. The Williams-type Cartesian robot and the idealised robot converged on some parameters but differed in retained cane number and bud density. This result indicates that a canopy considered robot-friendly for one platform may not be optimal for another platform with a different workspace envelope, camera position, end-effector size, or approach direction. Consequently, universal management rules for robotic harvesting should be treated with caution unless the robot configuration is explicitly specified.

Limitations and future work

Several limitations define the interpretation and scope of the present results. First, the study remains simulation-based. Although HYI was calibrated against a large published robotic-harvesting benchmark and evaluated through baseline, ablation, sensitivity, and cross-robot analyses, the reported improvements are model-supported estimates rather than prospective orchard-scale performance guarantees. Field validation should compare conventional and HarvestField-guided management under matched crop-load treatments, replicated orchard bays, and a fixed robotic platform. Outcomes should include HYI, observed picking success, failure-mode decomposition, cycle time, fruit quality, labour requirement, and yield.

Second, each sub-field simplifies physical processes that are more complex under field conditions. Reachability is approximated through sampled kinematics; visibility is evaluated from a finite set of candidate viewpoints; clearance is represented by local obstacle density rather than full end-effector trajectory simulation; and detachability is estimated from simplified force and neighbour-distance terms. These assumptions are appropriate for a management-level decision-support model, but they do not replace dynamic manipulation simulation or physical picking trials. Future studies should quantify how each sub-field correlates with observed failure modes, including missed fruit, approach collision, fruit drop, peduncle breakage, and excessive cycle time.

Third, the plant and environmental representations are incomplete. Wind-induced canopy motion, shoot stiffness, fruit maturity heterogeneity, leaf water status, cultivar-dependent peduncle mechanics, and seasonal management variability were not explicitly represented. These factors can alter visibility, clearance, and detachability during actual harvesting. The present HYI should therefore be interpreted as a static prior indicator at the start of harvest, rather than as a full temporal prediction of harvesting performance. A dynamic extension, $\mathcal{H}^{(t)}(\mathbf{x})$, would allow fruit removal, changing occlusion, and sequence-dependent access to be incorporated into a joint model of canopy management and harvesting policy.

Fourth, the current six-dimensional management vector captures major commercial interventions but not all local decisions made by growers. Branch-level adjustment, differential pruning within vigour zones, localised leaf removal, and spatially variable fruit-load targets would require higher-resolution sensing and more detailed prescription maps. Deep-learning-based orchard fruit detection and yield estimation provide a practical route for replacing simulated fruit positions with orchard-specific digital reconstructions (Maheswari et al., 2021; He et al., 2022). Integrating such sensing with HarvestField would enable block-specific diagnosis and could clarify whether the framework generalises across cultivars, trellis systems, and robot architectures. Comparative studies across kiwifruit, apple, sweet pepper, and grape systems would further determine whether the dominant sub-field remains clearance-dependent or shifts toward visibility, reachability, or detachability under different crop geometries.

Conclusions

This study addressed a central limitation in current robotic fruit-harvesting research: the canopy is usually treated as a fixed operational constraint, although its structure is actively shaped by grower decisions throughout the growing season. To convert this observation into a computable engineering framework, HarvestField integrates three components: a spatial harvestability field $\mathcal{H}(x)$, an L-system-based functional–structural plant model of Hayward kiwifruit, and NSGA-III multi-objective optimisation over six canopy management variables. The resulting framework evaluates how pruning, bud retention, fruit thinning, summer tipping, and cane training jointly affect HYI, yield, and labour cost.

The main methodological contribution is the formulation of $\mathcal{H}(x) = R(x)V(x)C(x)D(x)$ as a spatially resolved measure of robot-oriented canopy suitability. By aggregating this field at fruit positions, the Harvested Yield Index (HYI) provides a prior indicator of harvesting feasibility that can be computed before physical picking. This shifts the assessment of robotic harvestability from post-harvest performance measurement toward pre-harvest management evaluation, while retaining a direct connection to the physical conditions required for robotic retrieval: reachability, visibility, clearance, and detachability.

The simulation results support four principal findings. First, calibration against the Williams robotic-harvesting benchmark produced HYI=0.562, an absolute difference of 0.004 from the observed success rate of 0.558, and bootstrap uncertainty intervals contained the reference value. Second, the HarvestField-optimised strategy increased mean HYI from 0.587 under standard management to 0.642, while the Pareto front displayed explicit trade-offs among harvestability, yield proxy, and labour cost. Third, under equal total fruit counts, $\mathcal{H}$-guided thinning outperformed uniform thinning at all tested crop-load levels, with ΔHYI reaching +0.301 at the lowest fruit density. This result supports the hypothesis that the spatial arrangement of retained fruit, and not only total crop load, affects robot-oriented harvestability. Fourth, Sobol analysis identified clearance neighbourhood radius as the dominant source of HYI variance, indicating that local obstacle conditions constitute a critical component of harvestability under the T-bar kiwifruit system.

The practical value of the framework lies in the fact that its decision variables correspond to standard horticultural operations rather than to genetic modification, specialised sensing infrastructure, or new robot hardware. In this form, HarvestField can serve as a computational decision-support layer for comparing management plans, identifying Pareto-efficient compromises, and examining whether a proposed canopy structure is compatible with a given harvesting platform. The cross-robot comparison further indicates that such recommendations should be interpreted as robot-specific rather than universal, because the same canopy geometry can generate different harvestability profiles under different kinematic constraints.

The study also defines the boundaries of the present evidence. The framework was calibrated against a published benchmark rather than independently validated through new orchard trials; therefore, HYI should be interpreted as a prior harvestability indicator, not as a direct substitute for measured picking success. The FSPM simplifies some sources of biological and environmental variability, including wind-induced canopy motion, maturity heterogeneity, and operator-level variation during pruning or thinning. The current implementation also evaluates only two robot configurations and does not yet model dynamic changes in the canopy during sequential harvesting.

Future work should therefore prioritise orchard-scale validation using independent multi-season data, extension to additional harvesting platforms, and integration of dynamic picking-sequence optimisation with seasonal management optimisation. More broadly, the framework suggests a pathway for moving from robot adaptation alone toward bidirectional crop–robot system design. Within that broader trajectory, HarvestField provides a quantitative language for making crop environments themselves part of the engineering design space for robotic agriculture.

Declarations

Funding This study was supported by the National Natural Science Foundation of China (NSFC) (12104343), the Crosswise Project of West Anhui University in 2025 "Improvement and Research on High-Efficiency Integrated Load-bearing Wheel Rolling Forming Technology"(0045026049)，and the National-level College Students' Innovation and Entrepreneurship Training Program Project(202510376040S).

Ethics Approval Not applicable. This study did not involve human participants, human data, animals, or field collection of plant materials requiring ethical approval or collection permits.

References

Arad, B., Balendonck, J., Barth, R., Ben-Shahar, O., Edan, Y., Hellström, T., Hemming, J., Kurtser, P., Ringdahl, O., Tielen, T., & van Tuijl, B. (2020). Development of a sweet pepper harvesting robot. Journal of Field Robotics, 37(6), 1027–1039. https://doi.org/10.1002/rob.21937

Basso, B., & Antle, J. (2020). Digital agriculture to design sustainable agricultural systems. Nature Sustainability, 3, 254–256. https://doi.org/10.1038/s41893-020-0510-0

Blank, J., & Deb, K. (2020). pymoo: Multi-objective optimization in Python. IEEE Access, 8, 89497–89509. https://doi.org/10.1109/ACCESS.2020.2990567

Bloch, V., Degani, A., & Bechar, A. (2018). A methodology of orchard architecture design for an optimal harvesting robot. Biosystems Engineering, 166, 126–137. https://doi.org/10.1016/j.biosystemseng.2017.11.006

Boudon, F., Pradal, C., Cokelaer, T., Prusinkiewicz, P., & Godin, C. (2012). L-Py: An L-system simulation framework for modelling plant architecture development based on a dynamic language. Frontiers in Plant Science, 3, 76. https://doi.org/10.3389/fpls.2012.00076

Chelle, M., & Andrieu, B. (1998). The nested radiosity model for the distribution of light within plant canopies. Ecological Modelling, 111(1), 75–91. https://doi.org/10.1016/S0304-3800(98)00100-8

Cieslak, M., Seleznyova, A. N., & Hanan, J. (2011). A functional–structural kiwifruit vine model integrating architecture, carbon dynamics and effects of the environment. Annals of Botany, 107(5), 747–764. https://doi.org/10.1093/aob/mcq180

Cliff, N. (1993). Dominance statistics: Ordinal analyses to answer ordinal questions. Psychological Bulletin, 114(3), 494–509. https://doi.org/10.1037/0033-2909.114.3.494

Deb, K., & Jain, H. (2014). An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, Part I: Solving problems with box constraints. IEEE Transactions on Evolutionary Computation, 18(4), 577–601. https://doi.org/10.1109/TEVC.2013.2281535

Droukas, L., Doulgeri, Z., Tsakiridis, N. L., Triantafyllou, D., Kleitsiotis, I., Mariolis, I., Giakoumis, D.,

Tzovaras, D., Kateris, D., & Bochtis, D. (2023). A survey of robotic harvesting systems and enabling technologies. Journal of Intelligent & Robotic Systems, 107, 21. https://doi.org/10.1007/s10846-022-01793-z
Elith, J., & Leathwick, J. R. (2009). Species distribution models: Ecological explanation and prediction across space and time. Annual Review of Ecology, Evolution, and Systematics, 40, 677–697. https://doi.org/10.1146/annurev.ecolsys.110308.120159
He, L., Fang, W., Zhao, G., Wu, Z., Fu, L., Li, R., Majeed, Y., & Dhupia, J. (2022). Fruit yield prediction and estimation in orchards: A state-of-the-art comprehensive review for both direct and indirect methods. Computers and Electronics in Agriculture, 195, 106812. https://doi.org/10.1016/j.compag.2022.106812
Iwanaga, T., Usher, W., & Herman, J. (2022). Toward SALib 2.0: Advancing the accessibility and interpretability of global sensitivity analyses. Socio-Environmental Systems Modelling, 4, 18155. https://doi.org/10.18174/sesmo.18155
Kootstra, G., Wang, X., Blok, P. M., Hemming, J., & van Henten, E. (2021). Selective harvesting robotics: Current research, trends, and future directions. Current Robotics Reports, 2, 95–104. https://doi.org/10.1007/s43154-020-00034-1
Lammers, K., Zhang, K., Zhu, K., Chu, P., Li, Z., & Lu, R. (2024). Development and evaluation of a dual-arm robotic apple harvesting system. Computers and Electronics in Agriculture, 227, 109586. https://doi.org/10.1016/j.compag.2024.109586
Liu, L., Yang, F., Liu, X., Du, Y., Li, X., Li, G., Chen, D., Zhu, Z., & Song, Z. (2024). A review of the current status and common key technologies for agricultural field robots. Computers and Electronics in Agriculture, 227, 109630. https://doi.org/10.1016/j.compag.2024.109630
Maheswari, P., Raja, P., Apolo-Apolo, O. E., & Pérez-Ruiz, M. (2021). Intelligent fruit yield estimation for orchards using deep learning based semantic segmentation techniques—A review. Frontiers in Plant Science, 12, 684328. https://doi.org/10.3389/fpls.2021.684328
Saltelli, A. (2002). Making best use of model evaluations to compute sensitivity indices. Computer Physics Communications, 145(2), 280–297. https://doi.org/10.1016/S0010-4655(02)00280-1
Soualiou, S., Wang, Z., Sun, W., de Reffye, P., Collins, B., Louarn, G., & Song, Y. (2021). Functional–structural plant models mission in advancing crop science: Opportunities and prospects. Frontiers in Plant Science, 12, 747142. https://doi.org/10.3389/fpls.2021.747142
van Herck, L., Kurtser, P., Wittemans, L., & Edan, Y. (2020). Crop design for improved robotic harvesting: A case study of sweet pepper harvesting. Biosystems Engineering, 192, 294–308. https://doi.org/10.1016/j.biosystemseng.2020.01.021
Wilcoxon, F. (1945). Individual comparisons by ranking methods. Biometrics Bulletin, 1(6), 80–83. https://doi.org/10.2307/3001968
Williams, H. A. M., Jones, M. H., Nejati, M., Seabright, M. J., Bell, J., Penhall, N. D., Barnett, J. J., Duke, M. D., Scarfe, A. J., Ahn, H. S., Lim, J. Y., & MacDonald, B. A. (2019). Robotic kiwifruit harvesting using machine vision, convolutional neural networks, and robotic arms. Biosystems Engineering, 181, 140–156. https://doi.org/10.1016/j.biosystemseng.2019.03.007
Williams, H. A. M., Ting, C., Nejati, M., Jones, M. H., Penhall, N. D., Lim, J. Y., Seabright, M. J., Bell, J., Ahn, H. S., Scarfe, A. J., Duke, M. D., & MacDonald, B. A. (2020). Improvements to and large-scale evaluation of a robotic kiwifruit harvester. Journal of Field Robotics, 37(2), 187–201. https://doi.org/10.1002/rob.21890
Xie, Y., Zhang, T., Yang, M., Lyu, H., Zou, Y., Sun, Y., Xiao, J., Lian, W., Tao, J., Han, H., & Xu, C. (2025). Engineering crop flower morphology facilitates robotization of cross-pollination and speed breeding. Cell, 188(21), 5809–5830.e27. https://doi.org/10.1016/j.cell.2025.07.028
Yang, Q., Du, X., Wang, Z., Meng, Z., Ma, Z., & Zhang, Q. (2023). A review of core agricultural robot technologies for crop productions. Computers and Electronics in Agriculture, 206, 107701. https://doi.org/10.1016/j.compag.2023.107701
Yoshikawa, T. (1985). Manipulability of robotic mechanisms. The International Journal of Robotics Research, 4(2), 3–9. https://doi.org/10.1177/027836498500400201
Zhou, H., Wang, X., Au, W., Kang, H., & Chen, C. (2022). Intelligent robots for fruit harvesting: Recent developments and future challenges. Precision Agriculture, 23, 1856–1907. https://doi.org/10.1007/s11119-022-09913-3
Zurell, D., Franklin, J., König, C., Bouchet, P. J., Dormann, C. F., Elith, J., Fandos, G., Feng, X., Guillera-Arroita, G., Guisan, A., Lahoz-Monfort, J. J., Leitão, P. J., Park, D. S., Peterson, A. T., Rapacciuolo, G., Schmatz, D. R., Schröder, B., Serra-Diaz, J. M., Thuiller, W., Yates, K. L., Zimmermann, N. E., & Merow, C. (2020). A standard protocol for reporting species distribution models. Ecography, 43(9), 1261–1277. https://doi.org/10.1111/ecog.04960

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM2ExperimentalProtocolandParameterSettings.docx
- ESM3SupplementaryFigures.docx
- ESM4ReproducibilityPackage.zip